

Energetically Consistent Inelasticity for Optimization Time Integration

XUAN LI, University of California, Los Angeles, USA

MINCHEN LI, University of California, Los Angeles & TimeStep Inc., USA

CHENFANFU JIANG, University of California, Los Angeles & TimeStep Inc., USA




Fig. 1. Our energetically consistent inelasticity model can not only be applied to the Material Point Method (top row), but also easily extend to the Finite Element Method (bottom row with decreasing hardening coefficients from left to right). The stability under large time steps is guaranteed by the optimization time integration.

In this paper, we propose Energetically Consistent Inelasticity (ECI), a new formulation for modeling and discretizing finite strain elastoplasticity/viscoelasticity in a way that is compatible with optimization-based time integrators. We provide an in-depth analysis for allowing plasticity to be implicitly integrated through an augmented strain energy density function. We develop ECI on the associative von-Mises J2 plasticity, the non-associative Drucker-Prager plasticity, and the finite strain viscoelasticity. We demonstrate the resulting scheme on both the Finite Element Method (FEM) and the Material Point Method (MPM). Combined with a custom Newton-type optimization integration scheme, our method enables simulating stiff and large-deformation inelastic dynamics of metal, sand, snow, and foam with larger time steps, improved stability, higher efficiency, and better accuracy than existing approaches.

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

Additional Key Words and Phrases: MPM, Elastoplasticity, Inelasticity, Implicit Integration, Optimization Time Integration

Authors' addresses: Xuan Li, University of California, Los Angeles, USA, yunuoch@gmail.com; Minchen Li, University of California, Los Angeles & TimeStep Inc., USA, minchernl@gmail.com; Chenfanfu Jiang, University of California, Los Angeles & TimeStep Inc., USA, 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 the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2022/7-ART52 \$15.00

<https://doi.org/10.1145/3528223.3530072>

ACM Reference Format:

Xuan Li, Minchen Li, and Chenfanfu Jiang. 2022. Energetically Consistent Inelasticity for Optimization Time Integration. *ACM Trans. Graph.* 41, 4, Article 52 (July 2022), 16 pages. <https://doi.org/10.1145/3528223.3530072>

1 INTRODUCTION

Since the pioneering work of Terzopoulos and Fleischer [1988], the computer graphics community has observed increasing interests in modeling inelastic deformations governed by elastoplasticity, viscoelasticity, and viscoplasticity. These inelastic mechanical properties govern the behaviors of a wide range of everyday objects. Drawing inspirations from continuum mechanics, computer graphics researchers have successfully modeled and simulated many inelastic materials, ranging from metal, sand, snow and mud to foam, paint and organic tissues.

Inelasticity (mainly elastoplasticity and viscoplasticity) has been widely explored using mesh-based Finite Elements. During inelastic deformation, extreme element distortion and fracture commonly co-exist. Thus, remeshing [O'Brien et al. 2002] and virtual node [Hegemann et al. 2013; Molino et al. 2004] techniques are often applied. More recently, the Material Point Method (MPM) has emerged as a popular alternative for inelastic materials [Jiang et al. 2016] due to its natural support of topologically changing continuum materials.

Despite a large amount of work in modeling inelasticity, a loss of accuracy occurs in almost all existing work. In particular, when implicit time integration schemes are performed, the plastic strain

is often treated as a constant, and the real plastic deformation is imagined to happen instantaneously at the beginning or the end of a time step. Such a semi-implicit lagged treatment of inelasticity results in unnoticeable visual artifacts for certain material models such as the heuristic snow plasticity in Stomakhin et al. [2013] but significant errors such as excessive artificial cohesion for others [Gao et al. 2018; Tampubolon et al. 2017].

The choice of semi-implicit is largely due to the prominent challenge in modeling implicit inelasticity. Klár et al. [2016] was the first to explore differentiating the plastic flow for Drucker-Prager soil plasticity and incorporating it into the implicit momentum balance. The authors proposed an implicit force formulation that resembles a similar format to semi-implicit formulations [Stomakhin et al. 2013]. Unfortunately, their formulation cannot be expressed as the negative gradient of analytical energy. Resultingly, the stiffness matrix is asymmetric, and GMRES became necessary for the associated nonlinear root-finding problem – a problem that by itself has no stability or convergence guarantees when solved with Newton’s method. Fang et al. [2019] used alternating direction method of multipliers (ADMM) to shift the asymmetry to local small linear systems, however without an energy, they could not perform global convergence techniques such as line search.

This paper tackles the challenge by revisiting the derivation of implicit plasticity. Our objective is to construct an analytical, augmented potential energy function whose derivative exactly reproduces the implicit force. Related work in classic engineering literature [Ortiz and Stainier 1999; Radovitzky and Ortiz 1999] formulated variational constitutive model updates based on the principle of maximum plastic dissipation and minimizing over the so-called dual inelastic potential. Taking a different path, we derive our method based on constructing a smooth energy that is consistent with existing return mapping-based plasticity treatments [Simo and Hughes 1998] in explicitly integrated inelasticity simulation systems. As a result, our implicit inelasticity formulation can be directly incorporated into recently advanced optimization time integrators [Gast et al. 2015; Li et al. 2020; Wang et al. 2020] to enable large time step integration with guaranteed stability, theoretical consistency with return mapping, and a symmetric energy Hessian. Our contributions include:

- An implicit internal force formulation for fully implicit finite strain elastoplasticity;
- A strain energy augmentation method that yields analytically integrable elastoplastic forces and symmetric force derivatives for von Mises J2 plasticity;
- An extension of our model to support strain hardening, pressure-dependent soil plasticity, and rate-dependent viscoelasticity;
- Algorithms for incorporating our model in optimization-based time integrators with the Material Point Method and the Finite Element Method.

We demonstrate our results by simulating a wide range of inelastic materials, including metal, sand, snow, and foam. Our method allows the simulations of inelasticity to enjoy the advantages of guaranteed stability, global convergence, and large time step sizes brought by optimization-based time integrators without suffering from inaccuracy and numerical artifacts from prior work.

2 RELATED WORK

Inelasticity with FEM. Elastoplastic simulation with FEM has been extensively explored by the computer graphics community. O’Brien et al. [2002] used the additive decomposition of strain to separate elastic deformations and plastic deformations and used the von-Mises yield criterion. However, as Irving et al. [2004] pointed out, this decomposition does not support incompressibility for finite strain. Instead, Irving et al. [2004] used the multiplicative decomposition of deformation gradient with the volume-preserving return mapping algorithm. Our model is based on this decomposition as well. Under this framework, large plastic deformations may make the dynamic system ill-conditioned. To solve this problem, Molino et al. [2004] proposed the virtual node algorithm to allow topology changes when the simulated mesh is severely distorted, and Bargteil et al. [2007] used remeshing technique to maintain a high-quality mesh throughout the simulation. For high-performance simulation, Wojtan and Turk [2008] used frequently remeshed high-resolution surfaces combined with low-resolution interior tetrahedral mesh to resolve thin features near the boundaries. Wojtan et al. [2009] further improved the framework to allow topology changes in inelasticity simulations. These methods introduced extra computational costs or complexities. Instead, we use optimization time integrators to maintain long-time stability and global convergence. Furthermore, Bargteil et al. [2007] proposed a volume-preserving plasticity model incorporating creep and work hardening/softening, which is also followed by Wojtan and Turk [2008]. These are important requirements for obtaining physical accuracy, which are all supported by our model as well. Jones et al. [2016b] proposed an examples-based approach for the mesh-based discretization, which search rest shapes on a predefined example manifold. This method is efficient for animation purposes but are less physically accurate.

Inelasticity with MPM. Extending the work of Harlow [1964] and Brackbill and Ruppel [1986] on PIC/FLIP, MPM was proposed as a hybrid Lagrangian/Eulerian method for solid mechanics by Sulsky et al. [1994]. Since its appearance in the graphics community [Hegemann et al. 2013; Stomakhin et al. 2013], it has attracted a lot of attentions. The most prominent advantage of MPM on modeling inelastic materials is its flexibility in handling extreme deformation and topological changes, which pose significant challenges to Lagrangian mesh-based approaches. Snow plasticity was first simulated by Stomakhin et al. [2013] in a semi-implicit fashion, enforcing thresholds on principal stretches with post-projections. Yue et al. [2015] used the Herschel-Bulkley model of non-Newtonian viscoplastic flow to approximate foam behaviors. Fei et al. [2019] derived an analytic plastic flow approach for Herschel-Bulkley fluid to simulate compressible, shear-dependent liquids. Daviet and Bertails-Descoubes [2016] modeled the granular materials as compressible viscoplastic fluids combined with the Drucker-Prager yield criterion. Their method suits the granular material simulations well, but follows a different perspective from ours. From the perspective of large strain solid mechanics, Klár et al. [2016] simulated granular continuum using the return mapping algorithm for the Drucker-Prager plasticity. Following Klár et al. [2016], Yue et al. [2018] proposed a hybrid method combining both discrete and continuum treatments to achieve a high level of details with less computational costs. Fang



Fig. 2. **Snow Ball.** A free-falling snow ball hits on a static dragon and smashes into pieces.

et al. [2019] applied the return mapping approach to handle elastoplasticity and viscoelasticity in an ADMM framework. Except for Klár et al. [2016], these methods all temporally discretize inelasticity in an explicit or semi-implicit way, where the plastic correction was performed as an extra step at the end of each time step, fully decoupled from elasticity. Decoupled treatment in an explicit integration can be justified via operator splitting; however, it will cause artifacts for a (semi-)implicit integration. We use the return mapping framework as well for our fully implicit elastoplasticity and viscoelasticity, and we will show that ours is more temporally consistent compared to Klár et al. [2016].

Inelasticity with Other Discretizations. Inelasticity simulations are also explored with other types of spatial discretizations, e.g., Smoothed Particle Hydrodynamics (SPH), Position Based Dynamics (PBD), peridynamics, etc.

SPH is a mesh-free Lagrangian method originally invented for fluid simulations. Inspired by SPH, Jones et al. [2014]; Müller et al. [2004] applied the plasticity model in O'Brien et al. [2002] to moving least square particles for elastoplastic objects. Clavet et al. [2005] used springs between particles to mimic elasticity and achieved plasticity by modifying rest lengths during the simulation. These two plasticity models are not derived from the finite strain framework. Alduán and Otaduy [2011] simulated granular materials using an incompressible SPH framework combined with the Drucker-Prager yield criterion. Their plastic correction was performed in a Jacobi-like manner until convergence, while ours is performed with fixed-point iterations. Yang et al. [2017] proposed an elastoplastic model based on the Drucker-Prager yield criterion as well within an SPH framework. Gerszewski et al. [2009] introduced deformation gradients to the SPH framework so that plasticity models based on the multiplicative decomposition of deformation gradient can be applied. They used explicit time integrators combined with the plasticity model in [Irving et al. 2004]. Gissler et al. [2020] used an implicit compressible SPH solver to simulate the compression of snow. The plasticity is handled by an extra correction step on the deformation gradient following Stomakhin et al. [2013], which is still a semi-implicit method.

PBD was proposed by [Müller et al. 2007] for real-time simulations. This method replaced internal forces in force-based methods with constraints on positions. Plastic deformations can be introduced by the shape matching framework [Bender et al. 2017; Falkenstein et al. 2017; Jones et al. 2016a; Müller et al. 2005]. However,

this simulation framework sacrifices physical accuracy for better efficiency.

The peridynamic theory is an emerging field in simulations, which was proposed by Silling [2000] to handle discontinuities caused by deformations, such as cracks. It defines pairwise force functions between particles and uses the integration over the interactions from neighboring particles to describe dynamics. He et al. [2017] used the peridynamics framework to simulate elastoplastic materials in a projective dynamics way. They adopted the Drucker-Prager criterion for plasticity. Their solver can also be extended to simulate viscoelasticity. Chen et al. [2018] derived a form of force functions based on the isotropic linear elasticity model to simulate elastoplastic materials. They used explicit time integrators and an additive plasticity model.

Optimization Time Integration. Large-scale implicit simulation methods usually require solving large systems of nonlinear equations. To solve these systems, the Newton method for root-finding problems is usually adopted, which needs careful tuning of the time step size to ensure convergence. In fact, many of these implicit equations can be integrated to get variational forms, where the equivalent minimization problem can be solved by applying robust optimization techniques. The optimization time integrators have advantages in terms of long-time stability even when simulating severe deformation with large time step sizes.

Bouaziz et al. [2014] proposed Projective Dynamics (PD), which reformulated the backward Euler time integration for a specific type of material into a local-global alternating solver. Both the local and global steps have simple variational forms that can be solved in a robust and efficient way. This framework was later extended to simulate hyperelastic materials [Liu et al. 2017], support Laplacian damping [Li et al. 2018], and utilize other time integration schemes [Dinev et al. 2018]. Narain et al. [2016] then extended PD to a more general form within the ADMM framework. Brown and Narain [2021] improved the ADMM framework to resolve large rotations. Gast et al. [2015] recast the backward Euler time integration with hyperelastic materials, Rayleigh dampings, and collision penalties as a minimization problem. Li et al. [2019] and Wang et al. [2020] explored domain decomposed and hierarchical preconditioning strategies respectively within a quasi-Newton optimization framework for robust and efficient time integration. Wang and Yang [2016] proposed a gradient descent solver for GPUs to accelerate optimization time integrations. Li et al. [2020] proposed Incremental Potential

Contact (IPC), a variational form for frictional contacts. Their friction bases are iterated in a similar manner to our iterative yield stresses. IPC is later proven effective for simulating codimensional objects [Li et al. 2021b], rigid bodies [Ferguson et al. 2021], reduced elastic solids [Lan et al. 2021], and FEM-MPM coupled domains [Li et al. 2021a], all within the optimization time integration framework. In this paper, we follow Gast et al. [2015] and Li et al. [2020] for the optimization time integration of MPM and FEM respectively. The hessian matrices are enforced to be positive definite following the per-stencil projection technique in Teran et al. [2005].

In addition to hyperelastic solids, Batty et al. [2007] reformulated the classical pressure projection step in solid-fluid coupling as a kinetic energy minimization. Narain et al. [2010] used a hybrid method for sand simulation, where the pressure and friction are solved on the Eulerian grid with a staggered projection method, alternating between two coupled quadratic programs. Narain et al. [2012] posed the strain limiting in cloth simulation as a nonlinear optimization problem. Karamouzas et al. [2017] proposed an energy-based crowd model for crowd simulation. Inglis et al. [2017] formulated fluid simulation under a primal-dual optimization framework. Brown et al. [2018] proposed an energy for dissipative forces.

3 FOUNDATIONS

In this section we start with reviewing finite strain elastoplasticity (Section 3.1), MPM spatial discretization (Section 3.2), optimization-based time integration (Section 3.3), and discretized plastic flow rule (Section 3.4). Our review is by no means complete, and they are provided as necessary background knowledge for our new model.

In Section 3.5, we present a new implicit force formulation that is consistent with the variational weak form. It has a remarkable advantage – integrability, and thus lays important theoretical foundations for our method.

3.1 Finite Strain Elastoplasticity

Our variational inelasticity model is derived under the finite strain elastoplasticity framework. Here we review some basic concepts and refer to [Simo 1992; Simo and Hughes 1998] for more details.

Let $\Omega^0 \subset \mathbb{R}^3$ be the reference configuration of the continuum body and denote $\mathbf{x} := \Phi(\mathbf{X}, t)$ the deformation map from Ω^0 (with coordinate $\mathbf{X}$) to the world space Ω^t (with coordinate $\mathbf{x}$). The deformation gradient $\mathbf{F} = \frac{\partial \Phi}{\partial \mathbf{X}}(\mathbf{X}, t)$ measures the local deformation of the infinitesimal region around $\mathbf{X}$. With finite strain elastoplasticity, $\mathbf{F}$ is multiplicatively decomposed into $\mathbf{F} = \mathbf{F}^E \mathbf{F}^P$, where $\mathbf{F}^P$ denotes the permanent plastic deformation, and $\mathbf{F}^E$ denotes the elastic deformation which results in elastic forces (Figure 3). Plasticity requires that the Kirchhoff stress $\boldsymbol{\tau}$ associated with $\mathbf{F}^E$ is inside the admissible area defined by a yield condition $y(\boldsymbol{\tau}) \leq 0$. The surface characterized by $y(\boldsymbol{\tau}) = 0$ is often referred to as the yield surface. When $\mathbf{F}$ changes, $\mathbf{F}^E$ will follow some plastic flow to evolve so that it lies within the yield surface. In this paper, we follow the volume preserving plastic flow from [Klár et al. 2016].

From the Lagrangian view point, the state of dynamics of an elastoplastic continuum can be described by a Lagrangian density field $R(\mathbf{X}, t)$ on Ω^0 and a Lagrangian velocity field $\mathbf{V}(\mathbf{X}, t) := \frac{\partial \Phi(\mathbf{X}, t)}{\partial t}$ on Ω^0 . The two fields are governed by the conservation of




Fig. 3. The elastoplastic decomposition of the deformation gradient.

mass

$$R(\mathbf{X}, t)J(\mathbf{X}, t) = R(\mathbf{X}, 0), \quad (1)$$

where $J = \det \mathbf{F}$, and the conservation of momentum,

$$R(\mathbf{X}, 0) \frac{\partial \mathbf{V}}{\partial t}(\mathbf{X}, t) = \nabla^{\mathbf{X}} \cdot \mathbf{P} + R(\mathbf{X}, 0)\mathbf{g}, \quad (2)$$

where $\mathbf{P}$ is the first Piola-Kirchhoff stress and $\mathbf{g}$ is the gravity.

3.2 MPM Discretization

The Material Point Method (MPM) discretizes a continuum by a set of disconnected Lagrangian material particles. The continuous time variable t is discretized by consecutive time steps $t^0, t^1, \dots, t^n$. Without loss of generality, we assume a fixed time step size Δt . The advection is carried on material particles, so the conservation of mass across time steps is trivially satisfied. Assuming backward difference on $\frac{\partial \mathbf{V}}{\partial t}$, the weak form of the momentum equation is [Jiang et al. 2016]

$$\frac{1}{\Delta t} \int_{\Omega^0} R(\mathbf{X}, 0)(\mathbf{V}^{n+1} - \mathbf{V}^n) Q_\alpha d\mathbf{X} = - \int_{\Omega^0} \mathbf{P} \nabla^{\mathbf{X}} Q_\alpha d\mathbf{X} \quad (3)$$

for an arbitrary test function $Q_\alpha(\mathbf{X}, t^n)$.

Here for simplicity, we drop the gravity term and admit the free surface assumption.

MPM uses the previous time step t^n as the reference configuration, so the integration on Ω^0 is pushed forward onto Ω^n :

$$\frac{1}{\Delta t} \int_{\Omega^n} \rho(\mathbf{x}, t^n)(\hat{\mathbf{v}}^{n+1} - \mathbf{v}^n) q_\alpha d\mathbf{x} = - \int_{\Omega^n} \frac{1}{J^n} \mathbf{P} \mathbf{F}^{nT} \nabla^{\mathbf{x}} q_\alpha d\mathbf{x}, \quad (4)$$

where $\rho, \mathbf{v}^n, \hat{\mathbf{v}}^{n+1}$, and q_α are Eulerian counterpart of $R, \mathbf{V}^n, \mathbf{V}^{n+1}$, and Q_α , obtained by pushing forward from Ω^0 onto Ω^n .

In MPM, B-Spline-based interpolations are often applied to define fields on Ω^n , and material particles serve as quadratures to approximate the volume integration. Let $\mathbf{x}_p^n, \mathbf{X}_p$ be the coordinate of particle p in Ω^n and Ω^0 respectively, and w_{ip}^n and ∇w_{ip}^n be the weight and weight gradient between particle p and grid i . With the mass lumping technique, the force equilibrium of grid i can be discretized as

$$\frac{1}{\Delta t} m_i^n (\hat{\mathbf{v}}_i^{n+1} - \mathbf{v}_i^n) = - \sum_p \mathbf{P}_p \mathbf{F}_p^{nT} \nabla w_{ip}^n V_p^0, \quad (5)$$

where V_p^0 is the initial volume of particle p , $m_i^n = \sum_p m_p w_{ip}^n$ is the lumped mass on grid i and m_p is approximated by $R(\mathbf{X}_p, 0)V_p^0$. The right hand side of Equation 5 is the internal elastic force on grid i .

At each time step, the velocity field $\mathbf{v}^n$ is transferred from material particles to grid nodes, and the new velocity field $\hat{\mathbf{v}}^{n+1}$ is solved and transferred back to material particles for advection. In this paper, we



Fig. 4. **Car collision.** The red car hits the yellow car against the wall by 56 mph, creating a big dent on the right side of the yellow car.

use the quadratic MLS kernel [Hu et al. 2018] as the weight function, and APIC [Jiang et al. 2015] as the particle-grid transfer scheme.

3.3 Optimization Time Integration

Assuming implicit integration with BDF1 (backward Euler), the first Piola-Kirchhoff stress $\mathbf{P}$ in Equation 5 is associated with the deformation gradient $\mathbf{F}^{n+1}$ at time step t^{n+1} . The deformation gradients $\mathbf{F}^n$ and $\mathbf{F}^{n+1}$ are related by

$$\mathbf{F}^{n+1} = (\mathbf{I} + \Delta t \nabla \hat{\mathbf{v}}_p^{n+1}) \mathbf{F}^n, \quad (6)$$

where $\nabla \hat{\mathbf{v}}_p^{n+1} = \sum_i \hat{\mathbf{v}}_i^{n+1} \nabla w_{ip}^n$.

Existing optimization-based time integrators in computer graphics often assume hyperelastic materials. Without plasticity, the first Piola-Kirchhoff stress is simply the derivative of the corresponding elastic strain energy density function: $\mathbf{P}(\mathbf{F}) = \frac{\partial \Psi}{\partial \mathbf{F}}$. Equation 5 is then equivalent to the following optimization problem:

$$\Delta \hat{\mathbf{v}} = \argmin_{\Delta \mathbf{v}} E(\Delta \mathbf{v}) = \sum_i m_i \|\Delta \mathbf{v}_i\|^2 + \sum_p \Psi(\mathbf{F}_p^{tr}(\mathbf{v}^n + \Delta \mathbf{v})) \mathbf{V}_p^0, \quad (7)$$

$$\hat{\mathbf{v}}^{n+1} = \mathbf{v}^n + \Delta \hat{\mathbf{v}},$$

where $\mathbf{F}_p^{tr}(\mathbf{v}) = (\mathbf{I} + \Delta t \nabla \mathbf{v}_p) \mathbf{F}_p^n$ is the elastic predictor. The optimization problem can be robustly solved by projected Newton's method with backtracking line search [Wang et al. 2020].

Gravity. For the effect of gravity, we add the term $m_i \mathbf{g}$ on the right-hand side in Equation 5, which corresponds to the extra term $-\sum_i m_i \mathbf{g}^\top \mathbf{v}$ in Equation 7.

3.4 Discretization of Plastic Flow

In the discrete setting, plasticity is most commonly achieved by the return mapping algorithm [Klár et al. 2016; Simo and Hughes 1998], which is equivalent to solving for a strain that satisfies the plastic flow rule. Geometrically, the return mapping defines how elastic predictors outside the yield surface should be corrected so that the effective stresses lie inside the yield surface. We follow the notations in Klár et al. [2016] to describe discrete plastic flows in this paper.

For elasticity, we adopt the St. Venant-Kirchhoff (StVK) model with Hencky strains. The elastoplasticity of isotropic materials can be characterized in the principal stretch space using the singular value decomposition (SVD) [Stomakhin et al. 2012]. Let $\mathbf{F}^{tr} =$




Fig. 5. Return mapping for a discrete plastic flow.

$\mathbf{U} \Sigma^{tr} \mathbf{V}^\top$ be the SVD of an elastic predictor $\mathbf{F}^{tr}$. The Hencky strain is defined as $\epsilon = \log \Sigma^{tr}$, and the Kirchhoff stress for the StVK model is $\tau = 2\mu\epsilon + \lambda \text{tr}(\epsilon) \mathbf{I}$, where μ, λ are Lamé parameters. For a discrete plastic flow, if the stress associated with an elastic predictor is outside the yield surface, then the stress is projected back onto the yield surface. The projection procedure in the principal stress space is illustrated in Figure 5. Note that along the perpendicular direction to the diagonal, we would have $\det \mathbf{F}^P = 1$, which corresponds to a volume-preserving plastic deformation. We denote the endpoint of the return mapping (also known as the corrector or the effective stress) as $\mathbf{F}^E = \mathbf{Z}(\mathbf{F}^{tr})$ where $\mathbf{Z}(\cdot)$ is the return mapping.

3.5 Force Balance with Implicit Plasticity

For elastoplastic materials, the first Piola-Kirchhoff stress $\mathbf{P}$ in Equation 5 should be rewritten as [Bonet and Wood 1997]

$$\mathbf{P} = \frac{\partial \Psi^E}{\partial \mathbf{F}^E} \mathbf{F}^{P-\top}. \quad (8)$$

Here Ψ^E is the elastic strain energy density function. We add a superscript to emphasize the elastic energy is only associated with the elastic deformation gradient $\mathbf{F}^E$.

Through a weak form derivation of the updated Lagrangian dynamics (see the supplemental document for details), we show that

the implicit internal force on grid node i is:

$$f_i^{n+1} = - \sum_p V_p^0 \frac{\partial \Psi^E}{\partial F^E} (F_p^{E,n+1}) F_p^{P,n+1-\top} F_p^{n\top} \nabla w_{ip}^n, \quad (9)$$

In practice one does not need to track the plastic deformation gradients F^P on the particles. The nodal force can be expressed in terms of $F^{E,tr}$ (defined as $F^{E,tr}(\mathbf{v}) = (\mathbf{I} + \Delta t \nabla \mathbf{v}_p) F_p^{E,n}$):

$$f_i^{n+1} = - \sum_p V_p^0 \frac{\partial \Psi^E}{\partial F^E} (Z(F^{E,tr})) Z(F^{E,tr})^\top F^{E,tr-\top} F^{E,n\top} \nabla w_{ip}^n. \quad (10)$$

Note that when doing explicit time integration, we can directly replace $F_p^{E,n+1}$, $F_p^{E,tr}$ both with $F_p^{E,n}$, which gives the common force expression for explicit MPM:

$$f_i^n = - \sum_p V_p^0 \frac{\partial \Psi^E}{\partial F^E} (F_p^{E,n}) F_p^{E,n\top} \nabla w_{ip}^n. \quad (11)$$

For implicit plasticity, Klár et al. [2016] directly replace $\frac{\partial \Psi^E}{\partial F^E} (F_p^{E,n})$ in Equation 11 with $\frac{\partial \Psi^E}{\partial F^E} (Z(F^{E,tr}))$ and define the resulting expression as the implicit force. We can clearly observe that

$$f_i^{n+1} \neq - \sum_p V_p^0 \frac{\partial \Psi^E}{\partial F^E} (Z(F_p^{E,tr})) F_p^{E,n\top} \nabla w_{ip}^n, \quad (12)$$

i.e., the implicit force in [Klár et al. 2016] is not equivalent to our formulation (Equation 10). As we discuss in the supplemental document, the choice of Klár et al. [2016] is only semi-implicit. Furthermore, their formulation is not integrable because their force derivative is not symmetric. Therefore in [Klár et al. 2016] it is necessary to adopt Newton-Raphson root finding with GMRES for the asymmetric linear system solve. In the next section, we elaborate on our new model which enables the existence of an analytical energy.

4 ENERGY CONSISTENT INELASTICITY (ECI)

4.1 One-Dimensional Investigation

To motivate ECI, let's start with applying a discrete plasticity model to a one-dimensional spring with a constant yield stress.

Consider a one-dimensional elastoplastic spring with rest length $V_0 = 1$. We fix its one end at $x = 0$, and place the other end at 1 initially. With the initial state being the reference configuration, we model the spring with finite strain elastoplasticity, where the deformation gradient can be conveniently calculated as $F(x) = x$ with x being the coordinate of its free end.

Discretizing time into steps $t^0, t^1, \dots, t^n$ with equal time step size Δt , for time step n , the elastic predictor $F^{E,tr}$ by assuming a purely elastic deformation is given by

$$F^{E,tr}(x) = \frac{1}{F^{P,n}} x. \quad (13)$$

We assign the following strain energy density function:

$$\Psi^E(F^E) = \frac{k}{2} (\log F^E)^2, \quad (14)$$

where k is the stiffness. Viewing the spring as a single-element FEM discretization, since $V_0 = 1$, Ψ^E equals the total elastic potential.




Fig. 6. The strain-energy (left) and the strain-stress (right) plot of an elastoplastic spring. Here $\epsilon = \log(F)$ and $\tau = \frac{\partial \Psi}{\partial F} F$.

Assume $k = 1$ for brevity, the Kirchhoff stress is then given by

$$\tau(F^E) := \frac{\partial \Psi^E}{\partial F^E} F^E = \log(F^E). \quad (15)$$

Let the constant yield stress be $\tau_Y = \log(F_Y)$, where $F_Y \in [1, \infty)$ is a critical strain, and define the yield function to be $\tau - \tau_Y \leq 0$, we can then follow standard plasticity treatment [Simo and Hughes 1998] to define a simple return mapping procedure with the form

$$F^{E,n+1} = Z(F^{E,tr}) = \begin{cases} F^{E,tr} & F^{E,tr} \leq F_Y \\ F_Y & \text{otherwise} \end{cases}. \quad (16)$$

In terms of the Hencky strain $\epsilon^E = \log(F^E)$, the yield condition is equivalent to

$$\epsilon^{E,tr} - \epsilon^{E,n+1} = \delta\gamma > 0. \quad (17)$$

Geometrically, the quantity $\delta\gamma$ measures how far away the elastic strain predictor is from the yield surface in the principal strain space. This quantity plays an important role in our variational modeling of plasticity. Specifically, we have the following theorem for the elastoplastic springs:

THEOREM 4.1 (AUGMENTED ENERGY DENSITY FOR SPRINGS). *In the problem setting described above ($V^0 = 1, k = 1$), using the following energy density function*

$$\Psi(x) = \begin{cases} \Psi^E(Z(F^{E,tr}(x))) + \tau_Y \delta\gamma(F^{E,tr}(x)) & F^{E,tr} > F_Y \\ \Psi^E(F^{E,tr}(x)) & \text{otherwise} \end{cases} \quad (18)$$

reveals a force that is equivalent to what one would get if one performed the force-based implicit plasticity.

We include in the supplemental document the proof for this theorem as well as details showing that $\Psi(x)$ is piecewise C^∞ and everywhere C^1 . A comparison between the augmented energy and the pure-elastic energy is shown in Figure 6.

Taking the inertia into consideration, we test the model on a small dynamic mass-spring system. At the end of each time step, F_p is updated from the following relation:

$$F^{n+1} = x^{n+1} = Z(F^{E,tr}) F^{P,n+1} = F^{E,tr} F^{P,n}. \quad (19)$$

The ECI simulation results quantitatively match the results using explicit integration (see Figure 8).




Fig. 7. **Squeeze armadillo.** Rolling a plastic armadillo through a gear left the gear teeth permanently distorts the armadillo body.




Fig. 8. **Spring simulation with ECI.** We simulate under the same initial velocity but different critical strains. The results from our large-time-step ECI all match with the results from explicit small-time-step time integration.

4.2 Extending to Von-Mises Plasticity

A natural analogy of elastoplastic spring with constant yield stress for the plasticity of isotropic hyperelastic materials is the von-Mises plasticity model, which also associates all stress predictors with a constant yield stress τ_Y (the norm of the deviatoric Kirchhoff stress on the yield surface). We study von-Mises plasticity under the St. Venant-Kirchhoff constitutive model with Hencky strains. Following the notations from Section 3.4, the yield surface is defined as

$$y(\tau) = \|\hat{\tau}\|_F - \tau_Y = 0, \quad (20)$$

where $\hat{\tau} = \tau - \frac{1}{d} \text{tr}(\tau) \mathbf{I}$ is the deviatoric part of the Kirchhoff stress.

The equivalent yield condition is

$$\delta\gamma = \|\hat{\epsilon}\| - \frac{\tau_Y}{2\mu} > 0, \quad (21)$$




Fig. 9. Yield surface of the **von-Mises plasticity** model.

where $\epsilon = \log(\Sigma^{tr})$ is the trial Hencky strain and $\hat{\epsilon} = \epsilon - \frac{1}{d} \text{tr}(\epsilon) \mathbf{I}$ is the deviatoric part of the Hencky strain. The corresponding return mapping (Figure 9) is

$$Z(F^{E,tr}) = \begin{cases} F^{E,tr} & \delta\gamma \leq 0 \\ U \exp(\epsilon - \delta\gamma \frac{\hat{\epsilon}}{\|\hat{\epsilon}\|}) V^T & \text{otherwise} \end{cases} \quad (22)$$

We have the following key lemma for the von-Mises plasticity:

LEMMA 4.2. *Define the **augmented elastoplastic energy density function** as:*


$$\Psi(F) = \begin{cases} \Psi^E(F), & \delta\gamma(F) \leq 0 \\ \Psi^E(Z(F)) + \tau_Y \delta\gamma(F), & \text{otherwise} \end{cases} \quad (23)$$

This energy density function satisfies the following identity for any F :

$$\frac{\partial \Psi(F)}{\partial F} \equiv \frac{\partial \Psi^E}{\partial F^E}(Z(F)) Z(F)^T F^{-T}. \quad (24)$$

The proof of this lemma is provided in the supplementary document. With this lemma, it is easy to prove the following theorem:

THEOREM 4.3 (AUGMENTED ENERGY THEOREM FOR VON-MISES PLASTICITY). *The augmented elastoplastic energy density function (Equation 23) viewed as a hyperelastic strain energy density function*

Fig. 10. Yield surface of the **Drucker-Prager plasticity** model.

reveals a force that is equivalent to what one would get if one performed the force-based implicit plasticity, i.e.

$$\begin{aligned} f_i &= - \sum_p V_p^0 \left[\frac{\partial \Psi}{\partial \mathbf{F}} (\mathbf{F}_p^{E,tr}) \right] \mathbf{F}_p^{E,n \top} \nabla \mathbf{w}_{ip}^n \\ &= - \sum_p V_p^0 \left[\frac{\partial \Psi^E}{\partial \mathbf{F}_p^E} (\mathbf{Z}(\mathbf{F}_p^{E,tr})) \mathbf{Z}(\mathbf{F}_p^{E,tr})^\top \mathbf{F}_p^{E,tr - \top} \right] \mathbf{F}_p^{E,n \top} \nabla \mathbf{w}_{ip}^n. \end{aligned} \quad (25)$$

When performing optimization time integration, we can simply view the elastoplastic free energy density as a new strain energy density function. At the end of each time step, we update $\mathbf{F}^E$ with $\mathbf{Z}(\mathbf{F}^{E,tr})$. The detailed pipeline is elaborated in Section 5.

4.3 Extending to Pressure Dependent Soil Plasticity

Drucker-Prager plasticity is widely applicable to the simulations of granular materials such as sand. The yield surface under the St. Venant-Kirchhoff constitutive model with Hencky strains is defined as

$$y(\boldsymbol{\tau}) = \|\hat{\boldsymbol{\tau}}\|_F + \alpha \text{tr}(\boldsymbol{\tau}) = 0, \quad (26)$$

where $\alpha = \sqrt{\frac{2}{3} \frac{2 \sin \phi_f}{3 - \sin \phi_f}}$ and ϕ_f is the friction angle.

The equivalent yield condition is then

$$\text{tr}(\boldsymbol{\epsilon}) > 0, \quad \text{or} \quad \delta\gamma = \|\hat{\boldsymbol{\epsilon}}\|_F + \alpha \frac{(d\lambda + 2\mu) \text{tr}(\boldsymbol{\epsilon})}{2\mu} > 0. \quad (27)$$


The corresponding return mapping (Figure 10) is

$$\mathbf{Z}(\mathbf{F}^{E,tr}) = \begin{cases} \mathbf{U} \mathbf{V}^\top & \text{tr}(\boldsymbol{\epsilon}) > 0 \\ \mathbf{F}^{E,tr} & \delta\gamma \leq 0, \text{tr}(\boldsymbol{\epsilon}) \leq 0 \\ \mathbf{U} \exp(\boldsymbol{\epsilon} - \delta\gamma \frac{\hat{\boldsymbol{\epsilon}}}{\|\hat{\boldsymbol{\epsilon}}\|}) \mathbf{V}^\top & \text{otherwise} \end{cases} \quad (28)$$

The augmented elastoplastic energy introduced above for our 1D spring and von-Mises model essentially comes from the integrability of the following vector field over $\mathbb{R}^{d \times d}$:

$$\frac{\partial \Psi^E}{\partial \mathbf{F}^E} (\mathbf{Z}(\mathbf{F})) \mathbf{Z}(\mathbf{F})^\top \mathbf{F}^{-\top}. \quad (29)$$

Unfortunately, this integrability does not hold anymore for the Drucker-Prager return mapping. It can be checked that the Jacobian field of the above vector field is not symmetric. Even worse, $\delta\gamma$ is undefined in the region with $\text{tr}(\boldsymbol{\epsilon}) > 0$, because the projection there is not volume-preserving.

Fig. 11. Yield surface of the **Drucker-Prager plasticity** model in the principal strain stress under our extrapolated St. Venant-Kirchhoff model.

4.3.1 Extrapolating St. Venant-Kirchhoff. To solve the issue of $\delta\gamma$ for the area defined by $\text{tr}(\boldsymbol{\epsilon}) > 0$, we extrapolate the St. Venant-Kirchhoff constitutive model in this area as:

$$\hat{\Psi}^E(\boldsymbol{\Sigma}) = \begin{cases} \mu \|\hat{\boldsymbol{\epsilon}}\|^2 & \text{tr}(\boldsymbol{\epsilon}) \geq 0 \\ \mu \|\hat{\boldsymbol{\epsilon}}\|^2 + (\frac{\lambda}{2} + \frac{\mu}{d}) (\text{tr}(\boldsymbol{\epsilon}))^2 & \text{tr}(\boldsymbol{\epsilon}) < 0 \end{cases} \quad (30)$$

When $\text{tr}(\boldsymbol{\epsilon}) < 0$, $\hat{\Psi}^E$ is just equivalent to the St. Venant-Kirchhoff strain energy density, which separates the deviatoric term and the pressure term. When $\text{tr}(\boldsymbol{\epsilon}) \geq 0$, we extrapolate the energy only with the deviatoric term and define the yield stress to be zero. This extrapolation does not change the yield surface in the principal stress space. Instead, the yield surface in the principal strain space is extended to include the diagonal line of the first quadrant, and all the points on this ray correspond to the tip of the yield surface in the principal stress space (see Figure 11). With this extrapolated model, the volume-preserving projection can be done as well in the area of $\text{tr}(\boldsymbol{\epsilon}) \geq 0$, and $\delta\gamma$ is well-defined.

In summary, with our extrapolation, the return mapping is simplified as

$$\mathbf{Z}(\mathbf{F}^{E,tr}) = \begin{cases} \mathbf{F}^{E,tr}, & \delta\gamma \leq 0 \\ \mathbf{U} \exp(\boldsymbol{\epsilon} - \delta\gamma \frac{\hat{\boldsymbol{\epsilon}}}{\|\hat{\boldsymbol{\epsilon}}\|}) \mathbf{V}^\top, & \text{otherwise} \end{cases} \quad (31)$$

where

$$\delta\gamma = \begin{cases} \|\hat{\boldsymbol{\epsilon}}\|, & \text{tr}(\boldsymbol{\epsilon}) > 0 \\ \|\hat{\boldsymbol{\epsilon}}\| + \alpha \frac{d\lambda + 2\mu}{2\mu} \text{tr}(\boldsymbol{\epsilon}), & \text{otherwise} \end{cases} \quad (32)$$

4.3.2 Recover Integrability. To resolve the non-integrability, we update the yield stress iteratively during integration (Figure 13). At each time step, we solve a series of optimization problems with constant yield stresses. The yield stress $\tau_{Y,p}^{tr}$ for each particle p is computed from its elastic predictor $\mathbf{F}_p^{E,tr}$ at the beginning of the optimization:

$$\tau_Y^{tr} = \begin{cases} 0, & \text{tr}(\boldsymbol{\epsilon}) > 0 \\ -\alpha(d\lambda + 2\mu) \text{tr}(\boldsymbol{\epsilon}), & \text{otherwise} \end{cases} \quad (33)$$

and the corresponding $\delta\gamma$ is defined with a fixed yield stress:

$$\delta\gamma = \begin{cases} \|\hat{\boldsymbol{\epsilon}}\|, & \text{tr}(\boldsymbol{\epsilon}) > 0 \\ \|\hat{\boldsymbol{\epsilon}}\| - \frac{\tau_Y^{tr}}{2\mu}, & \text{otherwise} \end{cases} \quad (34)$$




Fig. 14. The viscoelastic decomposition of the deformation gradient and its relationship to the elastoplastic decomposition.

Although the return mapping for viscoelasticity is totally different from the one for elastoplasticity, the vector field

$$\frac{\partial \Psi^N}{\partial \mathbf{F}^N}(\mathbf{Z}(\mathbf{F})) \mathbf{Z}(\mathbf{F})^\top \mathbf{F}^{-\top}$$

turns out to be integrable if Ψ^N is from the St. Venant-Kirchhoff constitutive model, and the augmented ECI energy for this vector field is

$$\Psi^{\text{Visco}}(\Sigma) = \hat{\mu} \text{tr}((\log \Sigma)^2) + \frac{\hat{\lambda}}{2} (\text{tr}(\log \Sigma))^2, \quad (42)$$

where $\hat{\mu} = A\mu_N$ and $\hat{\lambda} = A\lambda_N - AB(2\mu_N + d\lambda_N)$.

Without plasticity, $\mathbf{F}^P \equiv \mathbf{I}$, and then the strain energy density for a viscoelastic material is simply

$$\Psi(\mathbf{F}) = \Psi^E(\mathbf{F}) + \Psi^{\text{Visco}}(\mathbf{F}). \quad (43)$$

With MPM discretization, each particle p independently tracks the evolutions of $\mathbf{F}_p^N$ and $\mathbf{F}_p^E$ and independently updates them accordingly at the end of each time step.

5 SPATIAL-TEMPORAL INTEGRATION

In this section, we present the detailed pipeline of ECI applied to MPM. The algorithm stages from t^n to t^{n+1} are listed as follows:

- (1) **Particles-to-grid transfer.** Grid mass m_i^n and velocity $\mathbf{v}_i^n$ are transferred from particle mass m_p , velocity $\mathbf{v}_p^n$, and angular velocity information $\mathbf{C}_p^n$ with APIC [Jiang et al. 2015].
- (2) **Optimize new grid velocity.** A series of optimization problems in the form of Equation 7 using ECI augmented energies are solved until the fixed-point iteration converges or the maximal number of iterations is reached. See Section 5.1.
- (3) **Grid-to-particles transfer.** The grid velocity $\hat{\mathbf{v}}_i^{n+1}$ from the time integration are transferred back to particles to update particle velocity $\mathbf{v}_p^{n+1}$ and angular velocity information $\mathbf{C}_p^{n+1}$.
- (4) **Particle strain update.** The elastic strain $\mathbf{F}^E$ or $\mathbf{F}^N$ are updated according to return mappings.
- (5) **Particle advection.** Particles are advected via particle velocity: $\mathbf{x}_p^{n+1} = \mathbf{x}_p^n + \mathbf{v}_p^{n+1} \Delta t$.

We only elaborate on the second stage in the following section. The other stages are the same as the standard explicit MPM simulation pipeline [Jiang et al. 2016].

Algorithm 1: Iterative Stress Optimization Time Integration

```

1: procedure MPMTIMEINTEGRATION( $\Delta \mathbf{v}^{\text{init}}, \mathbf{M}^n, \mathbf{v}^n, \Delta t, \epsilon$ )  $\triangleright \mathbf{M}^n, \mathbf{v}^n$  are
   stacked grid mass and velocity,  $\Delta \mathbf{v}^{\text{init}}$  is the initial guess
2:    $\Delta \mathbf{v} = \Delta \mathbf{v}^{\text{init}}$ 
3:   do  $\triangleright$  Iterative Stress Iteration
4:     for each particle  $p$ 
5:       Evaluate trial hardening state  $q_p^{tr}$   $\triangleright$  Equation 38
6:       Evaluate trial yield stress  $(\tau_Y^{tr})_p$   $\triangleright$  Equation 35 38
7:     end
8:     do  $\triangleright$  Solve Equation 7
9:        $\mathbf{r} \leftarrow -\nabla E(\Delta \mathbf{v})$   $\triangleright E$  as in Equation 7
10:       $\delta \Delta \mathbf{v} \leftarrow \text{InexactMINRES}(\text{ProjectPD}(\nabla^2 E(\Delta \mathbf{v})), \mathbf{r})$   $\triangleright$  Section
11:      5.1.2
12:       $\alpha \leftarrow \text{InversionFreeFilter}(\delta \Delta \mathbf{v})$   $\triangleright$  Section 5.1.3
13:       $E_{\text{init}} \leftarrow E(\Delta \mathbf{v})$ 
14:      while  $E(\Delta \mathbf{v} + \alpha \delta \Delta \mathbf{v}) > E_{\text{init}}$   $\triangleright$  Line search
15:         $\alpha \leftarrow \frac{\alpha}{2}$ 
16:      end
17:       $\Delta \mathbf{v} \leftarrow \Delta \mathbf{v} + \alpha \delta \Delta \mathbf{v}$ 
18:       $\hat{\mathbf{r}} = \text{Residual}(\mathbf{r})$   $\triangleright$  Section 5.1.4
19:      while  $\|\hat{\mathbf{r}}\|_\infty > \epsilon$ 
20:        while yield stress not converged
21:          for each particle  $p$   $\triangleright$  Advance hardening state
22:             $q^{n+1} = q_p^{tr}$ 
23:          end
24:        end
25:      end procedure

```

5.1 Iterative Stress Optimization Time Integration

To make the internal force of implicit plasticity integrable, the yield stress is viewed as constant in the force formulation, i.e., each particle sees a local cylinder-like yield surface with a different yield stress. Multiple optimizations with updated yield stresses are needed to make the final computed stresses consistent with the true yield surface. Each optimization problem is solved robustly using the projected Newton method with backtracking line search [Wang et al. 2020], where the Hessian matrix is projected to a nearby positive definite form [Teran et al. 2005]. See Algorithm 1 for the pseudo-code.

The update procedure of τ_Y^{tr} can be viewed as a fixed point iteration:

$$\tau_Y^{tr,j+1} = \Gamma_Y^{tr}(\mathbf{F}^{E,tr}(\Delta \mathbf{v}(\tau_Y^{tr,j}))). \quad (44)$$

Here j is the index of stress iteration, $\Delta \mathbf{v}(\tau_Y^{tr,j})$ is the equilibrated grid velocity field returned by a single optimization based on the yield stress vector $\tau_Y^{tr,j}$, and the bold symbol represents the stacked stress vector from all particles or all grid nodes. Since the Jacobian of this iteration has a scalar Δt^2 (see the supplemental document for details), the convergence of this fixed-point iteration is guaranteed if Δt and the residual for the equilibrium are both small enough. In practice, we find that even with large time steps, only several fixed-point iterations are required to produce visually high-quality results.

5.1.1 Boundary Conditions. The boundary conditions in our simulations are all from rigid collision objects. At the beginning of each time step, we detect the set of grid nodes colliding with the collision objects and directly enforce the velocity continuity condition across

Table 1. Simulation Statistics.

Scene	Figure	Model	Δt (s)	Δx (m)	Geometry	Elasticity	Plasticity/Viscosity	s/Step
Sand castle	19	Drucker-Prager	0.0004 ^a	0.007	2.26M particles	$E = 5 \times 10^6$	$\phi_f = 30^\circ, c = 0.0025$	18
Snow ball	2 18	Drucker-Prager	0.001	0.01	1.00M particles	$E = 10^6$	$\phi_f = 30^\circ, c = 0.0025$	10
Noddle	1	Von-Mises	0.001	0.01	2.07M particles	$E = 10^6$	$\tau_Y = 7.7 \times 10^2, \xi = 0$	31
Hydraulic test (Can)	1 21 22	Von-Mises	0.01	/	156K elements	$E = 7 \times 10^9$	$\tau_Y = 3 \times 10^7, \xi = 0.5$	5.6
Hydraulic test (Cylinder)	23	Von-Mises	0.01	/	299K elements	$E = 7 \times 10^9$	$\tau_Y = 3 \times 10^7, \xi = 0.1$	15
Hydraulic test (Square)	24	Von-Mises	0.01	/	230K elements	$E = 7 \times 10^9$	$\tau_Y = 3 \times 10^7, \xi = 0.1$	8.1
Armadillo	7	Von-Mises	0.01	/	121K elements	$E = 10^6$	$\tau_Y = 10^5, \xi = 0.5$	99
Car crash	4	Von-Mises	0.005 ^b	/	152K elements	$E = 2 \times 10^9$	$\tau_Y = 2.5 \times 10^6, \xi = 0.1$	51
Memory foam	12	Viscoelasticity	0.01	/	212K elements	$E = 10^3$	$E_N = 2 \times 10^5, v = 0.01$	17

^aUsing a smaller time step to satisfy the CFL limit for particle advection ($\Delta t \leq CFL \cdot \Delta x / v_{\max}$).

^bUsing a smaller time step to decrease numerical damping for a more vivid animation.

the collision interface. In each Newton iteration, the linear solver is projected so that the solved search direction remains tangent to the constraint manifold.

5.1.2 Inexact Newton-Krylov Methods. Following Wang et al. [2020], we use an inexact Newton-Krylov method. The tolerance for the linear systems is set relatively large in an adaptive way. Although more Newton iterations are needed, the reduced linear solve cost can still improve the world-clock timing of Newton convergence. Specifically, we use matrix-free Minimal Residual Method (MINRES) to solve the linear systems and the relative tolerance of each MINRES solve is set to $\min(0.5, \max(0.1, \sqrt{\mathbf{r}^\top \mathbf{P} \mathbf{r}}))$, where $\mathbf{r}$ is the right-hand side vector and $\mathbf{P}$ is the preconditioning matrix.

5.1.3 Inversion-free Line Search. The Hencky strain requires that the deformation gradient is not inverted, i.e., $\det(\mathbf{F}^{E,tr}) > 0$. Following [Li et al. 2020, 2021c; Smith and Schaefer 2015], before the line search, we first compute a large admissible step size α for the search direction such that the energy is well-defined for any step size $t \in [0, \alpha]$, and then the backtracking procedure starts with the filtered step size α .

5.1.4 Stopping Criteria. To terminate the Newton iterations early while ensuring visually high-quality simulation results, we normalize the grid residual vector $\mathbf{r}$, the gradient of the system energy, by the grid mass vector. This gives a residual in the unit of velocity (m/s), which is associated with a physical meaning. However, due to numerical rounding errors, small-mass nodes sometimes can have large residuals but contribute little to the particle advection. Therefore, we use grid-to-particle transfer to transfer the grid residual vector onto the particles to get the final residual vector. All our examples are running with tolerance 10^{-2} m/s based on the infinity norm of the velocity-unit residual vector on particles.

5.1.5 Timestep Size Restriction. The time step size of MPM is bounded by the advection CFL condition [Gast et al. 2015]. For those without stress-iterations, no further restrictions are required for our optimization integrator. For those with stress-iterations, theoretically, there is indeed a timestep size restriction for the stress iteration to fully converge, but we have not observed non-converging cases.

6 DISCRETIZATION WITH FEM

ECI is independent of spatial discretization choices. Hence it can also be conveniently applied in Finite Element Methods (FEM).

In FEM, the conservation-of-momentum equation (Equation 3) is directly discretized and solved in the material space. For FEM with linear tetrahedral elements, the discretized nodal internal force is

$$\mathbf{f}_i = - \sum_e V_e^0 \mathbf{P}_e \nabla N_{ie}, \quad (45)$$

where e indices all tetrahedral elements, V_e^0 is the rest volume of element e , and ∇N_{ie} is the gradient of the shape function on node i evaluated at the barycenter of element e [Irving et al. 2006].

Considering implicit plasticity, the internal force can be written as (see the supplemental document for details)

$$\mathbf{f}_i^{n+1} = - \sum_e V_e^0 \frac{\partial \Psi^E}{\partial \mathbf{F}^E} (\mathbf{Z}(\mathbf{F}_e^{E,tr})) \mathbf{Z}(\mathbf{F}_e^{E,tr})^\top \mathbf{F}_e^{E,tr-\top} \mathbf{F}_e^{P,n-\top} \nabla N_{ie}. \quad (46)$$

The integrability of the vector field $\frac{\partial \Psi^E}{\partial \mathbf{F}^E} (\mathbf{Z}(\mathbf{F}^{E,tr})) \mathbf{Z}(\mathbf{F}^{E,tr})^\top \mathbf{F}^{E,tr-\top}$ leads us to the integrable internal force from the augmented elasto-plastic energy density Ψ :

$$\mathbf{f}_i^{n+1} = - \frac{\partial}{\partial x_i} \left(\sum_e \Psi(\mathbf{F}_e^{E,tr}) V_e^0 \right), \quad (47)$$

where x_i is the world space coordinate of node i .

At the end of each time step, we need to track and update $\mathbf{F}^P$ on each element with

$$\mathbf{Z}(\mathbf{F}^{E,tr}) \mathbf{F}^{P,n+1} = \mathbf{F}^{E,tr} \mathbf{F}^{P,n}. \quad (48)$$

ECI combined with Incremental Potential Contact (IPC) [Li et al. 2020] allows us to simulate various scenarios where both accurate frictional contacts and inelastic responses are essential.

Timestep Size Restriction. Similar to MPM, there is also a timestep size restriction for the stress iterations to fully converge in FEM. Other than that, no further restrictions are needed. However, there is certainly a tradeoff between the number of timesteps and the accuracy and overall efficiency of the simulation [Li et al. 2020], which holds for all time discretized numerical schemes.




Fig. 15. **Sand Column Collapse.** (a) The explicit method explodes with $\Delta t = 10^{-4}s$. The implicit method [Klár et al. 2016] with vanilla Newton fails at a time step where the scene almost becomes static $\Delta t = 10^{-4}s$ (the frame right before the failure is rendered here). Our method works with all three time step sizes $\Delta t = 10^{-5}s, 10^{-4}s, 10^{-3}s$, and produces consistent results. The semi-implicit method produces artificial elastic behaviors even with a small time step size. (b) With the explicit method as the ground truth, our method has a smaller error (larger IoU score) than the semi-implicit method.

Table 2. Iteration statistics of 2D sand colume collapse.

Δt	# Stress Iter. (Avg. / Max)	# Newton Iter. (Avg. / Max)	# Line search (Avg. / Max)
0.01	8.8 / 13	112.3 / 186	202.0 / 476
0.005	6.8 / 9	45.2 / 75	48.8 / 201
0.0025	5.1 / 7	18.4 / 34	9.6 / 60
0.00125	3.7 / 6	9.2 / 14	1.3 / 9
0.000625	2.6 / 4	4.7 / 8	0.0 / 0

7 EVALUATION

We demonstrate the versatility of ECI with both MPM and FEM simulations. Among these examples, the ones that do not contain topological changes are simulated with FEM, and the frictional contact is modeled with IPC [Li et al. 2020]. For our MPM simulations, we use a CFL number of 0.6 [Gast et al. 2015]. The world-clock timing and the simulation setup are reported in Table 1. The statistics are based on Intel Core i9-10920X 3.5-GHz CPU with 12 cores.

7.1 Unit Tests

Convergence of Stress Iteration.

We test the convergence of the stress iteration on a 2D sand column collapse experiment. We use a direct solver to solve linear systems in the optimization time integrator, to avoid complicating the experiments with possibly inexact Krylov solves. The convergence criteria of the stress iteration is $\|(\tau_Y^{j+1} - \tau_Y^j)\|_2 < 10^{-9}(2\mu\sqrt{N})$, where N is the number of particles, and the Newton tolerance is 10^{-5} . Note that these tolerances are much tighter than needed so that we can verify that our method can converge with high accuracy. We consecutively halve the time step size from $\Delta t = 0.01s$. All these tests successfully converge with the given convergence criteria and have consistent results (see the inset figure). The iteration statistics





Fig. 16. With lower and upper bounds, the stored elastic energy in the soda can changes periodically over time during the compressing–stretching cycles, which demonstrates the long-time stability of our simulation.

are listed in Table 2, which shows that as the time step size decreases, the required number of stress iterations, Newton iterations, and line searches all decrease as expected.

Long-Time Stability. To test the long-time stability of our method, we simulate a soda can being periodically compressed and stretched 10 cycles with $\Delta t = 10^{-2}s$ (Figure 16). The Young’s modulus of the soda can is 7 GPa. The stored elastic energy over time is always bounded and it oscillates along with the compressing–stretching cycles, demonstrating the strong long-time stability property of our method.

7.2 Comparisons to Explicit and (Semi-)Implicit Plasticity

We compare our variational method with both explicit and implicit methods proposed in Klár et al. [2016] on a 3D *sand column collapse*




Fig. 17. ECI achieves 2× speedup compared to the explicit method in a sand column collapse experiment and is with comparable speed to the implicit method in Klár et al. [2016].

experiment. The time step size Δt for these two methods both need fine-tuning to avoid numerical explosion. For explicit integration, the time step is bounded by the sound-speed CFL [Sun et al. 2020], which is small in general, especially for stiff materials and at high resolution. With Klár et al. [2016]’s implicit method based on the non-integrable implicit force (Equation 12, with asymmetric force Jacobian), the convergence of time integration can only be reached if the initial guess is sufficiently close to the local optimum. Furthermore, the search performed by the Newton-Raphson iterations (we refer it as the vanilla Newton method) can result in deformation gradients with non-positive determinants that cause simulation failure.

We experiment under three different time step sizes $\Delta t = 10^{-3}s$, $10^{-4}s$ and $10^{-5}s$. Explicit MPM can run with $\Delta t = 10^{-5}s$, but it explodes with $\Delta t = 10^{-4}s$ (Figure 15a top left). The vanilla Newton method can run with $\Delta t = 10^{-5}s$, but fails at a step when the simulation almost becomes static with $\Delta t = 10^{-4}s$ and at the first step with $\Delta t = 10^{-3}$ (Figure 15a top middle). Our method, on the other hand, works well with all these three time step sizes and produces consistent results (Figure 15a right).

A common heuristic treatment is to directly replace the grid update step in the explicit MPM simulation with implicit time integration without plasticity and only conduct return mappings at the end of the time steps. We refer to this elasticity-plasticity-decoupled scheme as the semi-implicit method in this paper. Although its stability and convergence can be guaranteed by the optimization time integration, the semi-implicit method can lead to severe artifacts as shown in Figure 15a bottom left, where the forces provided by the stresses outside the yield surface make the continuum behave more like a purely elastic body. This is due to the ignorance of the plasticity by the implicit solve, which in turn overestimates the material’s resistance to tensile deformation. Our method, on the other hand, fully resolves plasticity in the implicit solve and does not suffer from any such artifacts. We use the explicit method as the ground truth to quantitatively measure errors. Figure 15b shows that our method has a smaller error than the semi-implicit method, where we compute the Intersection over Union (IoU) metrics between MPM grid mass distributions (computed as the ratio of the number of common grid nodes to the number of union grid nodes).



Fig. 18. Larger cohesion strength increases the chunkiness of the snow. From left to right, $c = 0.00125, 0.0025, 0.005$.

In practice, we can limit the number of Newton iterations and Krylov iterations. On a *sand column collapse* experiment with the same physical parameters and initial setup as above, our method with $\Delta t = 2 \times 10^{-5}s$ achieves 2× speedup compared to the explicit method with $\Delta t = 10^{-5}s$, as shown in Figure 17. With 2 stress iterations per time step, 1 Newton iteration per stress iteration, and 5 MINRES iterations per Newton iteration, our method can still generate physically plausible results. To make it a fair comparison, the maximal numbers of Newton iterations and GMRES iterations are set to 2 and 10 respectively for Klár’s implicit method with the same $\Delta t = 2 \times 10^{-5}s$ as ours. The simulation using Klár’s method does not go unstable in this setting, and its computational cost is similar to ours, as expected.

7.3 Druker-Prager Plasticity with Cohesion

Snow Castle. To further demonstrate the artifacts caused by fully decoupling elasticity and plasticity, we simulate a snow castle hit by a high-speed elastic fish. The snow is modeled with wet soil by Druker-Prager plasticity with cohesion. With our variational model, the fish smashes the snow castle into pieces completely. However, with the semi-implicit method, the castle behaves like an elastic body and ends up holding the fish in an unrealistic way.

Our extrapolated StVK constitutive model combined with the volume-preserving return mapping plays a vital role in generating fractures in this example. Intuitively, our scheme mimics the cohesion behavior better because it allows particles to be compressed a little before exerting resisting force. Under the same time step size ($\Delta t = 5 \times 10^{-5}$), we use the result from the explicit method with the extrapolated StVK model as the ground truth to compare the accuracy between our method and semi-implicit methods (with/without the extrapolated StVK model). The visual and quantitative comparisons in Figure 20 both show that our method is more accurate.

Snow Ball. We use the Druker-Prager plasticity model with cohesion to simulate a snow ball hitting a static dragon (Figure 2). We also simulate with different cohesion strengths to show the controllability of our method on simulating different levels of chunkiness (Figure 18).

7.4 Von-Mises Plasticity (with Hardening)

Play-Doh Noodle. MPM can automatically handle topology changes. By leveraging this feature, we simulate a Play-Doh modeled by the von-Mises plasticity pressed through a cylindrical noodle mold (Figure 1 (top row)).

Hydraulic Tests on Metals. Hardening is widely observed in metals. We simulate hydraulic tests on soda cans with different hardening




Fig. 19. **Snow Castle.** With our variational inelasticity model, the castle can be smashed into pieces after hitting by the fish, while with the semi-implicit method the castle behaves like an elastic body, holding the fish in an unrealistic way.




Fig. 20. Our method is more accurate visually and quantitatively than the semi-implicit methods with/without the extrapolated StVK constitutive model.




Fig. 21. Different hardening coefficients lead to varying restorations towards the rest shape and generate different crushing patterns. From left to right, the hardening coefficient $\xi = 0.5, 0.3, 0$.

coefficients and compare with the simulation without hardening (Figure 1 bottom row). The hardening mechanism makes plastic deformations harder to happen as the yield surface expands. This lets the object restore its original rest shape partially when all boundary conditions are released. As shown in the last frame when the upper press withdraws (Figure 21), the red can with the largest hardening coefficient restores the most, and the orange can with no hardening almost does not restore at all. Furthermore, different hardening coefficients generate different crushing patterns. As shown in Figure 22, the deformation patterns in one of our compressed cans can match that from a real experiment. Modeling hardening also allows us to successfully capture the snap-through instability of metal, which can be observed in real experiments (see our video demonstration). When we swap in long steel pipes for the hydraulic tests (one cylindrical, one square), the crushing patterns also match real experiments well (Figure 23, 24).




Fig. 22. One of our hydraulic test simulations on metal cans generate a crushing pattern well matching that in a real video footage [Youtube 2021].




Fig. 23. **Hydraulic Test on a Cylinder Pipe.** The crushing pattern matches the result of a real-world experiment [Youtube 2021].

Car Crash and Crushed Armadillo. To further demonstrate the hardening behaviors of metals, we simulate a high-speed car crashing into another stationary car (Figure 4) and an armadillo rolling



Fig. 24. **Hydraulic Test on a Square Pipe.** The crushing pattern matches the result of a real-world experiment [Youtube 2018].




Fig. 25. The semi-implicit von-Mises plasticity (right) overestimates the resistance response and results in a large error compared to the ground truth (left). Ours (middle) is much more accurate.

through a metal crusher driven by frictions (Figure 7). Both examples show realistic denting effects with sufficient restoration towards the rest shape enabled by hardening.

Comparison to Semi-Implicit Plasticity. We simulate a stiff elastic ball hitting a wall modeled by the von-Mises plasticity to compare our method with the semi-implicit method. As shown in Figure 25, the permanent deformations of the wall clearly show that the semi-implicit plasticity overestimates the material’s resistance. We use the explicit method as the ground truth to compare the position error of the wall (computed as the average squared norm of vertex position differences), which shows that our method is more accurate than the semi-implicit method.

7.5 Viscoelasticity

Memory foam is a typical material demonstrating the viscoelastic behaviors in the real world. We simulate a pillow made by memory foam pressed down by a hand for a while, and then we lift the hand suddenly. We successfully capture the intricate process where the pillow slowly recovers its rest shape, completely removing the imprint of the hand (Figure 12).

8 DISCUSSION

In summary, we developed ECI, a new formulation that augments hyperelastic energy density functions to enable variational forms for a wide range of elastoplastic and viscoelastic materials. Our

algorithm enables the fully implicit simulation of inelasticity in recently advanced optimization-based time integrators, embracing advantages of long-time stability, global convergence, large time step sizes, and high accuracy.

Our method is most naturally “plug and play” when applied to J2 von Mises materials and finite strain viscoelastic materials. However, when extended to pressure-dependent plasticity or strain hardening mechanisms, additional iterations on the stress are necessary to achieve final convergence. In our examples, usually 1-2 stress iterations are sufficient to generate convergent or visually high-quality results. It is promising future work to devise theoretical and algorithmic improvements to guarantee and accelerate the convergence, particularly for accuracy-demanding applications.

The integrability of the implicit elastoplastic force depends on both the elastic model and the plastic model. For instance, although the combination of St. Venant-Kirchhoff elasticity with von-Mises plasticity adopted by ECI leads to a symmetric force Jacobian, neo-Hookean elasticity with von-Mises plasticity does not. It is an interesting future work to explore integrable approximations to other combinations.

ECI assumes the full-dimensional volumetric deformation gradient. Accordingly, our metal cans and pipes are all simulated with thin single-layer linear tetrahedral elements, which could potentially suffer from shear locking. It would be interesting to extend ECI to codimensional geometries like shells and rods [Narain et al. 2013].

Finally, our augmentation to the strain energy density function changes the conditioning of the global stiffness matrix. It is interesting future work to study its effect on the linear solve, and strategies to precondition the ECI-augmented system.

ACKNOWLEDGMENTS

We would like to thank Pingying Chen for narrating the video, assisting with the video editing, and her sublime finishing touches on renderings, especially the teaser figure. We would also like to thank the anonymous reviewers for their valuable comments. This work has been supported in part by NSF CAREER 2153851, CCF-2153863, ECCS-2023780, DOE ORNL contract 4000171342.

REFERENCES

- Iván Alduán and Miguel A Otaduy. 2011. SPH granular flow with friction and cohesion. In *Proceedings of the 2011 Symp. Computer animation*. 25–32.
- A. W. Bargteil, J. K. Hodgins C. Wojtan, and G. Turk. 2007. A finite element method for animating large viscoplastic flow. *ACM Trans. on Graph.* 26, 3 (2007).
- Christopher Batty, Florence Bertails, and Robert Bridson. 2007. A fast variational framework for accurate solid-fluid coupling. *ACM Trans. on Graph.* 26, 3 (2007).
- Jan Bender, Matthias Müller, and Miles Macklin. 2017. A survey on position based dynamics, 2017. In *European Association for Computer Graphics: Tutorials*. 1–31.
- Javier Bonet and Richard D Wood. 1997. *Nonlinear continuum mechanics for finite element analysis*. Cambridge university press.
- S. Bouaziz, S. Martin, T. Liu, L. Kavan, and M. Pauly. 2014. Projective dynamics: Fusing constraint projections for fast simulation. *ACM Trans. on Graph.* 33, 4 (2014).
- J. U. Brackbill and H. M. Ruppel. 1986. FLIP: A method for adaptively zoned, particle-in-cell calculations of fluid flows in two dimensions. *J. Comput. Phys.* 65, 2 (1986).
- G.E. Brown, M. Overby, Z. Forootaninia, and R. Narain. 2018. Accurate dissipative forces in optimization integrators. *ACM Trans. on Graph.* 37, 6 (2018), 1–14.
- George E Brown and Rahul Narain. 2021. WRAPD: weighted rotation-aware ADMM for parameterization and deformation. *ACM Trans. on Graph.* 40, 4 (2021), 1–14.
- Jagabandhu Chakraborty and WJ Drugan. 1988. Theory of plasticity. (1988).
- Wei Chen, Fei Zhu, Jing Zhao, Sheng Li, and Guoping Wang. 2018. Peridynamics-Based Fracture Animation for Elastoplastic Solids. In *Computer Graphics Forum*, Vol. 37.

- Simon Clavet, Philippe Beaudoin, and Pierre Poulin. 2005. Particle-based viscoelastic fluid simulation. In *Proceedings of the 2005 Symp. Computer animation*. 219–228.
- G. Daviet and F. Bertails-Descoubes. 2016. A semi-implicit material point method for the continuum simulation of granular materials. *ACM Trans. on Graph.* 35, 4 (2016).
- Dimitar Dinev, Tiantian Liu, and Ladislav Kavan. 2018. Stabilizing integrators for real-time physics. *ACM Trans. on Graph.* 37, 1 (2018), 1–19.
- Michael Falkenstein, Ben Jones, Joshua A Levine, Tamar Shinar, and Adam W Bargteil. 2017. Reclustering for large plasticity in clustered shape matching. In *Proceedings of the Tenth International Conference on Motion in Games*. 1–6.
- Y. Fang, M. Li, M. Gao, and C. Jiang. 2019. Silly rubber: an implicit material point method for simulating non-equilibrated viscoelastic and elastoplastic solids. *ACM Trans. on Graph.* 38, 4 (2019), 1–13.
- Y. Fei, C. Batty, E. Grinspun, and C. Zheng. 2019. A multi-scale model for coupling strands with shear-dependent liquid. *ACM Trans. on Graph.* 38, 6 (2019), 1–20.
- 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* 40, 4 (2021).
- M. Gao, X. Wang, K. Wu, A. Pradhana, E. Sifakis, C. Yuksel, and C. Jiang. 2018. GPU optimization of material point methods. *ACM Trans. on Graph.* 37, 6 (2018).
- T. F. Gast, C. Schroeder, A. Stomakhin, C. Jiang, and J.M. Teran. 2015. Optimization integrator for large time steps. *trans. on vis. and comp. graph.* 21, 10 (2015).
- J. Gaume, T. Gast, J. Teran, A. van Herwijnen, and C. Jiang. 2018. Dynamic anticrack propagation in snow. *Nature Communications* 9, 1 (2018), 3047.
- Dan Gerszewski, Haimasree Bhattacharya, and Adam W Bargteil. 2009. A point-based method for animating elastoplastic solids. In *Symp. Computer animation*. 133–138.
- C. Gissler, A. Henne, S. Band, A. Peer, and M. Teschner. 2020. An implicit compressible SPH solver for snow simulation. *ACM Trans. on Graph.* 39, 4 (2020).
- Francis H Harlow. 1964. The particle-in-cell computing method for fluid dynamics. *Methods Comput. Phys.* 3 (1964), 319–343.
- Xiaowei He, Huamin Wang, and Enhua Wu. 2017. Projective peridynamics for modeling versatile elastoplastic materials. *trans. on vis. and comp. graph.* 24, 9 (2017).
- Jan Hegemann, Chenfanfu Jiang, Craig Schroeder, and Joseph M Teran. 2013. A level set method for ductile fracture. In *Symp. Computer animation*. 193–201.
- Y. Hu, Y. Fang, Z. Ge, Z. Qu, Y. Zhu, A. Pradhana, and C. Jiang. 2018. A moving least squares material point method with displacement discontinuity and two-way rigid body coupling. *ACM Trans. on Graph.* 37, 4 (2018), 150.
- Tiffany Inglis, M-L Eckert, James Gregson, and Nils Thuerey. 2017. Primal-dual optimization for fluids. In *Computer Graphics Forum*, Vol. 36.
- Geoffrey Irving, Joseph Teran, and Ronald Fedkiw. 2004. Invertible finite elements for robust simulation of large deformation. In *Symp. Computer animation*. 131–140.
- Geoffrey Irving, Joseph Teran, and Ronald Fedkiw. 2006. Tetrahedral and hexahedral invertible finite elements. *Graphical Models* 68, 2 (2006), 66–89.
- Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. 2015. The Affine Particle-in-Cell Method. *ACM Trans. Graph.* 34, 4 (2015).
- Chenfanfu Jiang, Craig Schroeder, Joseph Teran, Alexey Stomakhin, and Andrew Selle. 2016. The material point method for simulating continuum materials. In *ACM SIGGRAPH 2016 Courses*. 1–52.
- Ben Jones, April Martin, Joshua A Levine, Tamar Shinar, and Adam W Bargteil. 2016a. Ductile fracture for clustered shape matching. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 65–70.
- Ben Jones, Nils Thuerey, Tamar Shinar, and Adam W Bargteil. 2016b. Example-based plastic deformation of rigid bodies. *ACM Trans. on Graph.* 35, 4 (2016).
- Ben Jones, Stephen Ward, Ashok Jallepalli, Joseph Perenia, and Adam W Bargteil. 2014. Deformation embedding for point-based elastoplastic simulation. *ACM Trans. on Graph.* 33, 2 (2014), 1–9.
- I. Karamouzas, R. Sohre, N. and Narain, and S.J. Guy. 2017. Implicit crowds: Optimization integrator for robust crowd simulation. *ACM Trans. on Graph.* 36, 4 (2017).
- G. Klár, T. Gast, A. Pradhana, C. Fu, C. Schroeder, C. Jiang, and J. Teran. 2016. Drucker-prager elastoplasticity for sand animation. *ACM Trans. on Graph.* 35, 4 (2016).
- Lei Lan, Yin Yang, Danny Kaufman, Junfeng Yao, Minchen Li, and Chenfanfu Jiang. 2021. Medial IPC: accelerated incremental potential contact with medial elastics. *ACM Trans. on Graph.* 40, 4 (2021), 1–16.
- Jing Li, Tiantian Liu, and Ladislav Kavan. 2018. Laplacian damping for projective dynamics. In *Proceedings of the 14th Workshop on Virtual Reality Interactions and Physical Simulations*. 29–36.
- M. Li, Z. Ferguson, T. Schneider, T. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D.M. Kaufman. 2020. Incremental potential contact: Intersection- and inversion-free, large-deformation dynamics. *ACM transactions on graphics* (2020).
- M. Li, M. Gao, T. Langlois, C. Jiang, and D. M. Kaufman. 2019. Decomposed optimization time integrator for large-step elastodynamics. *ACM Trans. on Graph.* 38, 4 (2019).
- Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021b. Codimensional Incremental Potential Contact. *ACM Transactions on Graphics* 40, 4 (2021).
- X. Li, Y. Fang, M. Li, and C. Jiang. 2021a. BFEMP: Interpenetration-free MPM–FEM coupling with barrier contact. *Comp. meth. applied mech. eng.* (2021).
- Yue Li, Xuan Li, Minchen Li, Yixin Zhu, Bo Zhu, and Chenfanfu Jiang. 2021c. Lagrangian–Eulerian multidensity topology optimization with the material point method. *Internat. J. Numer. Methods Engrg.* 122, 14 (2021), 3400–3424.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-newton methods for real-time simulation of hyperelastic materials. *ACM Trans. on Graph.* 36, 3 (2017).
- Neil Molino, Zhaosheng Bao, and Ron Fedkiw. 2004. A virtual node algorithm for changing mesh topology during simulation. *ACM Trans. on Graph.* 23, 3 (2004).
- Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position based dynamics. *J. Visual Communication and Image Repres.* 18, 2 (2007), 109–118.
- M. Müller, B. Heidelberger, M. Teschner, and M. Gross. 2005. Meshless deformations based on shape matching. *ACM Trans. on Graph.* 24, 3 (2005).
- Matthias Müller, Richard Keiser, Andrew Nealen, Mark Pauly, Markus Gross, and Marc Alexa. 2004. Point based animation of elastic, plastic and melting objects. In *Proceedings of the 2004 Symp. Computer animation*. 141–151.
- Rahul Narain, Abhinav Golas, and Ming C Lin. 2010. Free-flowing granular materials with two-way solid coupling. In *ACM SIGGRAPH Asia 2010 papers*. 1–10.
- R. Narain, M. Overby, and G.E. Brown. 2016. ADMM $\supseteq$ projective dynamics: fast simulation of general constitutive models.. In *Symp. Comput. Animat.*, Vol. 1. 2016.
- Rahul Narain, Tobias Pfaff, and James F O'Brien. 2013. Folding and crumpling adaptive sheets. *ACM Trans. on Graph.* 32, 4 (2013), 1–8.
- Rahul Narain, Armin Samii, and James F O'Brien. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. on Graph.* 31, 6 (2012), 1–10.
- James F O'Brien, Adam W Bargteil, and Jessica K Hodgins. 2002. Graphical modeling and animation of ductile fracture. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. 291–294.
- Michael Ortiz and Laurent Stainier. 1999. The variational formulation of viscoplastic constitutive updates. *Comp. meth. applied mech. eng.* 171, 3-4 (1999), 419–444.
- R. Radovitzky and M. Ortiz. 1999. Error estimation and adaptive meshing in strongly nonlinear dynamic problems. *Comp. meth. applied mech. eng.* 172, 1-4 (1999).
- Stewart A Silling. 2000. Reformulation of elasticity theory for discontinuities and long-range forces. *Journal of the Mechanics and Physics of Solids* 48, 1 (2000), 175–209.
- Juan C Simo. 1992. Algorithms for static and dynamic multiplicative plasticity that preserve the classical return mapping schemes of the infinitesimal theory. *Comp. meth. applied mech. eng.* 99, 1 (1992), 61–112.
- Juan C Simo and Thomas JR Hughes. 1998. *Computational inelasticity*. Springer-Verlag.
- Jason Smith and Scott Schaefer. 2015. Bijective parameterization with free boundaries. *ACM Trans. on Graph.* 34, 4 (2015), 1–9.
- 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.
- A. Stomakhin, C. Schroeder, L. Chai, J. Teran, and A. Selle. 2013. A material point method for snow simulation. *ACM Trans. on Graph.* 32, 4 (2013), 102.
- D. Sulsky, Z. Chen, and H. I. Schreyer. 1994. A particle method for history-dependent materials. *Comp. meth. applied mech. eng.* 118, 1-2 (1994).
- Yunxin Sun, Tamar Shinar, and Craig Schroeder. 2020. Effective time step restrictions for explicit MPM simulation. In *Computer Graphics Forum*, Vol. 39.
- A. P. Tampubolon, T. Gast, G. Klár, C. Fu, J. Teran, C. Jiang, and K. Museth. 2017. Multi-species simulation of porous sand and water mixtures. *ACM Trans. on Graph.* 36, 4 (2017), 105.
- Joseph Teran, Efthychios Sifakis, Geoffrey Irving, and Ronald Fedkiw. 2005. Robust quasistatic finite elements and flesh simulation. In *Proceedings of the 2005 Symp. Computer animation*. 181–190.
- Demetri Terzopoulos and Kurt Fleischer. 1988. Modeling inelastic deformation: viscoelasticity, plasticity, fracture. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques*. 269–278.
- Huamin Wang and Yin Yang. 2016. Descent methods for elastic body simulation on the GPU. *ACM Trans. on Graph.* 35, 6 (2016), 1–10.
- Xinlei Wang, Minchen Li, Yu Fang, Xinxin Zhang, Ming Gao, Min Tang, Danny M Kaufman, and Chenfanfu Jiang. 2020. Hierarchical optimization time integration for cfl-rate mpm stepping. *ACM Trans. on Graph.* 39, 3 (2020), 1–16.
- Chris Wojtan, Nils Thuerey, Markus Gross, and Greg Turk. 2009. Deforming meshes that split and merge. In *ACM SIGGRAPH 2009 papers*.
- Chris Wojtan and Greg Turk. 2008. Fast viscoelastic behavior with thin features. In *ACM SIGGRAPH 2008 papers*. 1–8.
- Tao Yang, Jian Chang, Ming C Lin, Ralph R Martin, Jian J Zhang, and Shi-Min Hu. 2017. A unified particle system framework for multi-phase, multi-material visual simulations. *ACM Trans. on Graph.* 36, 6 (2017), 1–13.
- Youtube. 2018. *Crushing Long Steel Pipes with Hydraulic Press*. <https://www.youtube.com/watch?v=TM5dyY8zfxs&t=215s>
- Youtube. 2021. *Satisfying steel pipe crush video*. <https://www.youtube.com/watch?v=1s53ejdJgY&t=4s>
- Y. Yue, B. Smith, C. Batty, C. Zheng, and E. Grinspun. 2015. Continuum foam: A material point method for shear-dependent flows. *ACM Trans. on Graph.* 34, 5 (2015), 160.
- Y. Yue, B. Smith, P. Y. Chen, M. Chantharayukhonthorn, K. Kamrin, and E. Grinspun. 2018. Hybrid grains: Adaptive coupling of discrete and continuum simulations of granular media. In *SIGGRAPH Asia 2018 Technical Papers*. ACM, 283.