

MIT Open Access Articles

DiffPD: Differentiable Projective Dynamics

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Du, Tao, Wu, Kui, Ma, Pingchuan, Wah, Sebastien, Spielberg, Andrew et al. 2022. "DiffPD: Differentiable Projective Dynamics." ACM Transactions on Graphics, 41 (2).

As Published: 10.1145/3490168

Publisher: Association for Computing Machinery (ACM)

Persistent URL: <https://hdl.handle.net/1721.1/143798>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International license



DiffPD: Differentiable Projective Dynamics

TAO DU, KUI WU, PINGCHUAN MA, SEBASTIEN WAH, ANDREW SPIELBERG, DANIELA RUS, and WOJCIECH MATUSIK, MIT CSAIL

We present a novel, fast differentiable simulator for soft-body learning and control applications. Existing differentiable soft-body simulators can be classified into two categories based on their time integration methods: Simulators using explicit timestepping schemes require tiny timesteps to avoid numerical instabilities in gradient computation, and simulators using implicit time integration typically compute gradients by employing the adjoint method and solving the expensive linearized dynamics. Inspired by **Projective Dynamics (PD)**, we present **Differentiable Projective Dynamics (DiffPD)**, an efficient differentiable soft-body simulator based on PD with implicit time integration. The key idea in DiffPD is to speed up backpropagation by exploiting the prefactorized Cholesky decomposition in forward PD simulation. In terms of contact handling, DiffPD supports two types of contacts: a penalty-based model describing contact and friction forces and a complementarity-based model enforcing non-penetration conditions and static friction. We evaluate the performance of DiffPD and observe it is 4–19 times faster compared with the standard Newton’s method in various applications including system identification, inverse design problems, trajectory optimization, and closed-loop control. We also apply DiffPD in a **reality-to-simulation (real-to-sim)** example with contact and collisions and show its capability of reconstructing a digital twin of real-world scenes.

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

Additional Key Words and Phrases: Projective dynamics, differentiable simulation

ACM Reference format:

Tao Du, Kui Wu, Pingchuan Ma, Sebastien Wah, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. 2021. DiffPD: Differentiable Projective Dynamics. *ACM Trans. Graph.* 41, 2, Article 13 (November 2021), 21 pages. <https://doi.org/10.1145/3490168>

1 INTRODUCTION

The recent surge of differentiable physics witnessed the emergence of differentiable simulators as well as their success in various inverse problems that have simulation inside an optimization loop. With additional knowledge of gradients, a differentiable simulator provides more guidance on the evolution of a physics system. This extra information, when properly combined with mature gradient-based optimization techniques, facilitates the quantitative study of

various downstream applications, e.g., system identification or parameter estimation, motion planning, controller design and optimization, and inverse design problems.

In this work, we focus on the problem of developing a differentiable simulator for soft-body dynamics. Despite its potential in many applications, research on differentiable soft-body simulators is still in its infancy due to the large number of degrees of freedom (DoFs) in soft-body dynamics. One learning-based approach is to approximate the true soft-body dynamics by way of a neural network for fast, automatic differentiation [Li et al. 2019b; Sanchez-Gonzalez et al. 2020]. For these methods, the simulation process is no longer physics-based but purely based on a neural network, which might lead to physically implausible and uninterpretable results and typically do not generalize well.

Another line of research, which is more physics-based, is to differentiate the governing equations of soft-body dynamics directly [Geilinger et al. 2020; Hahn et al. 2019; Hu et al. 2020, 2019]. We classify these simulators into *explicit* and *implicit* simulators based on their timestepping schemes. Explicit differentiable simulators implement explicit time integration in forward simulation and directly apply the chain rule to derive any gradients involved. While explicit differentiable simulation is fast to compute and straightforward to implement, its explicit nature requires a tiny timestep to avoid numerical instability. Moreover, when deriving the gradients, an output value (typically a reward or an error metric) needs to be backpropagated through all timesteps. Such a process requires the state at every timestep to be stored in memory regardless of the timestepping scheme. Therefore, explicit differentiable simulators typically consume orders of magnitude more memory than their implicit counterparts, and sophisticated schemes like checkpoints are needed to alleviate the memory issue [Hu et al. 2019].

Unlike explicit differentiable simulators, implicit simulation enables a much larger timestep. It is more robust numerically and much more memory-efficient during backpropagation. However, an implicit differentiable simulator typically implements Newton’s method in forward simulation and the adjoint method during backpropagation [Geilinger et al. 2020; Hahn et al. 2019], both of which require the expensive linearization of the soft-body dynamics. Even though techniques for expediting the *forward* soft-body simulation with an implicit timestepping scheme have been developed extensively, the corresponding *backpropagation* process remains a bottleneck for downstream applications in inverse problems.

In this work, we present DiffPD, an efficient differentiable soft-body simulator that implements the **finite element method (FEM)** and an implicit timestepping scheme with certain assumptions on the material and contact model. We draw inspiration from **Projective Dynamics (PD)** [Bouaziz et al. 2014], a fast and stable algorithm that can be used for solving implicit time integration

This work is sponsored by Defense Advanced Research Projects Agency (DARPA) under Grant No. FA8750-20-C-0075, Intelligence Advanced Research Projects Activity (IARPA) under Grant No. 2019-19020100001, and National Science Foundation (NSF) Award 2106962: Computational Design of Complex Fluidic Systems. Authors’ addresses: T. Du, K. Wu, P. Ma, S. Wah, A. Spielberg, D. Rus, and W. Matusik, MIT CSAIL, 32 Vassar Street, Cambridge, MA 02139, USA; emails: {taodu, kuiwu, pcma}@csail.mit.edu, sebawah@mit.edu, {aespielberg, rus, wojciech}@csail.mit.edu.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).
0730-0301/2021/11-ART13
<https://doi.org/10.1145/3490168>

with FEM when the elastic energy of the material model has a specific quadratic form. The key observation we share with PD is that the computation bottleneck in both forward simulation and backpropagation is due to the nonlinearity of soft-body dynamics. By decoupling nonlinearity in the system dynamics, PD proposes a global-local solver where the global step solves a prefactorized linear system of equations and the local step resolves the nonlinearities in the physics and can be massively parallelized. Previous work in PD has demonstrated its efficacy in forward simulation and has reported significant speedup over the classical Newton’s method. Our core contribution is to establish that with proper linear algebraic reformulation, the same idea of nonlinearity decomposition from PD can be fully extended to backpropagation as well.

To support differentiable contact handling, we revisit contact models used in previous PD papers, many of which choose to implement a soft contact force based on a fictitious collision energy [Bouaziz et al. 2014; Dinev et al. 2018a, b; Li et al. 2019a; Liu et al. 2017; Wang 2015; Wang and Yang 2016]. DiffPD naturally supports such energy-based contact and friction models as they can be seamlessly integrated into the PD framework. One notable exception is Ly et al. [2020], which solves dry frictional contact in the standard PD framework. We have also explored the possibility of making the dry frictional contact model differentiable in DiffPD. Using the fact that contact vertices must be on the soft body’s surface, which typically have much fewer DoFs than the interior vertices, we present a novel solution combining Cholesky factorization and low-rank matrix update to supporting differentiable *static* friction and non-penetration contact.

We demonstrate the efficacy of DiffPD in various 3D applications with up to nearly 30,000 DoFs. These applications include system identification, initial state optimization, motion planning, and end-to-end closed-loop control optimization. We compare DiffPD with both explicit and implicit differentiable FEM simulations and observe DiffPD’s forward and backward calculation is 4–19 times faster than Newton’s method when assumptions in PD hold. Furthermore, we embed DiffPD as a differentiable layer in a deep learning pipeline for training closed-loop neural network controllers for soft robots and report a speedup of 9–11 times in wall-clock time compared with deep **reinforcement learning (RL)** algorithms. Finally, we show a **reality-to-simulation (real-to-sim)** application that uses our differentiable simulator to reconstruct a collision event between two tennis balls from a video input, which we hope can inspire follow-up work to solve **simulation-to-reality (sim-to-real)** problems in the future.

To summarize, our article contributes the following:

- a PD-based differentiable soft-body simulator that is significantly faster than differentiable simulators using the standard Newton’s method;
- a differentiable collision handling algorithm that handles penalty-based contact and friction forces or complementarity-based non-penetration contact and static friction; and
- demonstrations of the efficacy of our method on a wide range of applications, including system identification, inverse design problems, motion planning, robotics control, and a real-to-sim experiment, using material models compatible with PD and the simplified contact model stated above.

2 RELATED WORK

We review methods on differentiable simulators and PD with contact handling, followed by a short summary of soft robot design and control methods. We summarize the differences between our work and previous papers in Table 1.

Differentiable simulation. Many recent advances in differentiable physics facilitate the application of gradient-based methods in robotics learning, control, and design tasks. Several differentiable simulators have been developed for rigid-body dynamics [de Avila Belbute-Peres et al. 2018; Degraeve et al. 2019; Macklin et al. 2020; Popović et al. 2003; Toussaint et al. 2018], soft-body dynamics [Geilinger et al. 2020; Hahn et al. 2019; Hu et al. 2020, 2019], cloth [Liang et al. 2019; Qiao et al. 2020], and fluid dynamics [Holl et al. 2020; McNamara et al. 2004; Schenck and Fox 2018; Treuille et al. 2003; Wojtan et al. 2006]. Here we mainly discuss methods for soft-body dynamics as it is the focus of our work. ChainQueen [Hu et al. 2019] and its follow-up work DiffTaichi [Hu et al. 2020] introduce differentiable physics-based soft-body simulators using the **material point method (MPM)**, which uses particles to keep track of the full states of the dynamical system and solves the momentum equations as well as collisions on a background Eulerian grid. However, the explicit time integration in their methods requires small timesteps to preserve numerical stability, leading to large memory consumption during backpropagation. To resolve this issue, ChainQueen proposes to cache checkpoint steps in memory and reconstructs the states by rerunning forward simulation during backpropagation, which increases implementation complexity and introduces extra time cost. Further, solving collisions on an Eulerian grid may introduce artifacts depending on the resolution of the grid [Han et al. 2019]. Another particle-based strategy is to approximate soft-body dynamics with graph neural networks [Li et al. 2019b; Sanchez-Gonzalez et al. 2020], which is naturally equipped with differentiability but may result in physically implausible behaviors. Our work is most similar to Hahn et al. [2019] and Geilinger et al. [2020], which use differentiable implicit FEM simulators. Compared with other differentiable deformable body simulators [Geilinger et al. 2020; Hahn et al. 2019; Liang et al. 2019; Qiao et al. 2020] that systematically apply a combination of the adjoint method and sensitivity analysis to derive gradients, DiffPD puts a special focus on a more strategic backpropagation scheme that leverages the unique structure in PD, leading to empirically faster gradient computation. Finally, DiffPD is also relevant to Li et al. [2020] that introduce a contact-aware Newton-type solver and handle contact robustly with a differentiable barrier function for soft-body simulation. Although Li et al. [2020] do not discuss gradient computation or present applications that benefit from these gradients, they show that the whole contact model is differentiable in theory, from which we believe applications using differentiable simulation in the future can benefit a lot.

PD. PD was originally proposed by Bouaziz et al. [2014] as an attractive alternative to Newton’s method for solving nonlinear system dynamics with an implicit timestepping scheme. The standard PD algorithm has been extended to support rigid bodies [Li et al. 2019a], conserve kinetic energy [Dinev et al. 2018a, b], and support a wide range of hyperelastic materials [Liu et al. 2017]. Furthermore, prior papers have also proposed

Table 1. Summary of Related Work

Method	Domain	Integration	Physics-based	Speed	Memory	Diff.	Contact handling
Hu et al. [2019]	Soft body (MPM)	Explicit	Yes	Fast	High	Yes	Solved on an Eulerian grid
Li et al. [2019b]	Soft body (Particles)	Network	No	Fast	Medium	Yes	Trained neural-network function
Liang et al. [2019]	Cloth	Implicit	Yes	Slow	Low	Yes	Non-penetration + impulse-based friction
Hahn et al. [2019]	Soft body (FEM)	Implicit	Yes	Slow	Low	Yes	Soft contact + penalty energy
Geilinger et al. [2020]	Multi-body	Implicit	Yes	Slow	Low	Yes	Penalty-based frictional contact model
Ly et al. [2020]	Cloth	Implicit	Yes	Fast	Low	No	Signorini-Coulomb contact model
DiffPD (ours)	Soft body (FEM)	Implicit	Yes	Fast	Low	Yes	Non-penetration contact + static friction

The “Diff.” column refers to the differentiability of each method. In each column, green indicates the preferred property; red means the method lacks the preferred property; yellow means the method has some, but not all aspects of the preferred property.

more advanced acceleration strategies including semi-iterative Chebyshev solvers [Wang 2015; Wang and Yang 2016], parallel Gauss-Seidel methods with randomized graph coloring [Fratrangeli et al. 2016], precomputed reduced subspace methods for the required constraint projections [Brandt et al. 2018], and multigrid solvers [Xian et al. 2019]. Macklin et al. [2020] introduce a preconditioned descent-based method of PD on GPUs with a penalty-based contact model. In DiffPD, we implement their penalty-based contact model but leave the GPU acceleration as future work. On the theoretical side, Narain et al. [2016] and Overby et al. [2017] interpret PD as a special case of the **alternating direction method of multipliers (ADMM)** from the optimization perspective. Our work inherits the standard PD framework from previous papers and augments it with gradient computation.

Contact handling. The topic of handling contact and friction has been extensively studied in soft-body simulation. There exists a diverse set of collision handling algorithms with different design consideration: physical plausibility, time cost, and implementation complexity, although only a few of them take differentiability into consideration [Chen et al. 2017; Hu et al. 2019; Li et al. 2020; Macklin et al. 2020]. In the realm of PD simulation with contact, the most widespread strategy is to treat contact as soft constraints. This is achieved by either directly projecting colliding vertices onto collision surfaces at the end of each simulation step [Dinev et al. 2018a, b; Wang 2015; Wang and Yang 2016] or imposing a fictitious collision energy [Bouaziz et al. 2014; Li et al. 2019a; Liu et al. 2017]. Such methods introduce an artificial stiffness coefficient, which is task-dependent and requires careful tuning. Alternatively, Ly et al. [2020] propose to combine the more physically plausible Signorini-Coulomb contact model with PD in forward simulation. However, using such a model creates an additional challenge during backpropagation as solving the contact forces requires nontrivial iterative optimizers which either fail to maintain differentiability or cannot exploit the prefactorized Cholesky decomposition. In our work, we consider both penalty-based and complementarity-based contact models. For complementarity-based contact, we support non-penetration contact and static friction and leave sliding friction as future work. At the core of our contact handling algorithm is the assumption that contact occurs on a small fraction of the full DoFs, leading to a low-rank update on the linear system of equations. The idea of leveraging incremental, low-rank updates during contact handling can be traced back to ArtDefo [James and Pai 1999], which simulates soft bodies with the boundary element methods and resets three columns in a matrix

when a new contact point becomes active. In our work, we use a similar idea to solve contact during forward simulation and extend it to compute gradients during backpropagation. Our work is also relevant to the recent work on sparse Cholesky updates [Herholz and Sorkine-Hornung 2020], which provides an alternative to our low-rank update strategy based on Woodbury matrix identity.

Soft robot design and control. As soft-body simulation becomes faster, more robust, and more expressive in the way of gradient information, the field of computational soft robotics is emerging as a solution to soft robot design and control problems. One of the earliest such simulators applied to computational robotics is VoxCAD [Hiller and Lipson 2014], which is not differentiable but runs quickly on CPUs. This feature makes it well-catered to genetic algorithms, leading to computer-designed soft robots that can walk, swim, and grow [Bongard et al. 2016; Cheney et al. 2013; Corucci et al. 2016]. More recently, Min et al. [2019] have presented an FEM simulator which couples soft bodies with a hydrodynamical model to enable simulation and training of swimming creatures. Like VoxCAD, this simulator is a gradient-free approach, meaning training such controllers can take hours or days.

Competing with these gradient-free approaches are gradient-based approaches, which are typically more efficient due to the additional information provided by gradients. Sin et al. [2013] and Allard et al. [2007] have presented fast, partially differentiable FEM simulation that exploits modal dynamics to accelerate simulation, at the cost of accuracy. Still, these approaches are robust enough to allow interactive simulation of soft characters [Barbič and James 2005] and control of simulated and real robots [Thieffry et al. 2018]. More recently, similar finite-element-based simulation techniques are shown to be effective in generating gaits for real-world foam quadrupeds [Bern et al. 2019]. In each of these works, only control is examined. Similar to our work, Lee et al. [2018] demonstrate how PD can be applied to dexterous manipulation using an actuation model based on the human muscular system. While their framework is flexible and can handle soft-rigid coupling, it can only be applied to quasi-static motions. By contrast, our simulator handles fully dynamical systems. A competing approach, differentiable MPM, is applied to soft robotics in Spielberg et al. [2019], co-designing soft robots over closed- and open-loop control and material distributions.

3 BACKGROUND

In this section, we review the basic concepts of the implicit timestepping scheme and PD. Let n be the number of 3D nodes in

a deformable body after FEM-discretization. After time discretization, we use $\mathbf{x}_i \in \mathcal{R}^{3n}$ and $\mathbf{v}_i \in \mathcal{R}^{3n}$ to indicate the nodal positions and velocities at the i -th timestep.

Implicit time integration. In this article, we focus on the implicit time integration:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + h\mathbf{v}_{i+1}, \quad (1)$$

$$\mathbf{v}_{i+1} = \mathbf{v}_i + h\mathbf{M}^{-1}[\mathbf{f}_{\text{int}}(\mathbf{x}_{i+1}) + \mathbf{f}_{\text{ext}}], \quad (2)$$

where h is the timestep, $\mathbf{M} \in \mathcal{R}^{3n \times 3n}$ a lumped mass matrix, and \mathbf{f}_{int} and \mathbf{f}_{ext} the sum of the internal and external forces, respectively. Substituting \mathbf{v}_{i+1} in \mathbf{x}_{i+1} gives the following nonlinear system of equations:

$$\frac{1}{h^2}\mathbf{M}(\mathbf{x}_{i+1} - \mathbf{y}_i) - \mathbf{f}_{\text{int}}(\mathbf{x}_{i+1}) = \mathbf{0}, \quad (3)$$

where $\mathbf{y}_i = \mathbf{x}_i + h\mathbf{v}_i + h^2\mathbf{M}^{-1}\mathbf{f}_{\text{ext}}$ is evaluated at the beginning of each timestep. We drop the indices from \mathbf{x} and \mathbf{y} for simplicity:

$$\frac{1}{h^2}\mathbf{M}(\mathbf{x} - \mathbf{y}) - \mathbf{f}_{\text{int}}(\mathbf{x}) = \mathbf{0}. \quad (4)$$

At each timestep, our goal is to find \mathbf{x} satisfying the equation above with the given \mathbf{y} .

As pointed out by Stuart and Humphries [1996] and Martin et al. [2011], solving \mathbf{x} from Equation (4) is equivalent to finding the critical point of the following objective g :

$$g(\mathbf{x}) = \frac{1}{2h^2}(\mathbf{x} - \mathbf{y})^\top \mathbf{M}(\mathbf{x} - \mathbf{y}) + E(\mathbf{x}), \quad (5)$$

where E is the potential energy that induces the internal force: $\mathbf{f}_{\text{int}} = -\nabla E$. It is easy to check the left-hand side of Equation (4) is ∇g :

$$\nabla g(\mathbf{x}) = \frac{1}{h^2}\mathbf{M}(\mathbf{x} - \mathbf{y}) + \nabla E(\mathbf{x}) = \frac{1}{h^2}\mathbf{M}(\mathbf{x} - \mathbf{y}) - \mathbf{f}_{\text{int}}(\mathbf{x}). \quad (6)$$

Equation (4) is typically solved with Newton's method, which iteratively solves a series of linear systems of equations. Consider the k -th iteration in Newton's method with \mathbf{x}^k being the guess on \mathbf{x} so far. Newton's method computes the next guess on \mathbf{x} as follows:

$$\mathbf{0} = \nabla g(\mathbf{x}) = \nabla g(\mathbf{x}^k + \Delta\mathbf{x}) \approx \nabla g(\mathbf{x}^k) + \nabla^2 g(\mathbf{x}^k)\Delta\mathbf{x}, \quad (7)$$

$$\nabla^2 g(\mathbf{x}^k)\Delta\mathbf{x} \approx -\nabla g(\mathbf{x}^k). \quad (8)$$

Therefore, one can let $\Delta\mathbf{x} = -[\nabla^2 g(\mathbf{x}^k)]^{-1}\nabla g(\mathbf{x}^k)$ and update their guess on \mathbf{x} at the next iteration by $\mathbf{x}^{k+1} = \mathbf{x}^k + \Delta\mathbf{x}$. In practice, Newton's method typically employs definiteness fixes or line searches when $\nabla^2 g$ is indefinite [Nocedal and Wright 2006]. For large-scale problems, solving Equation (8) at each \mathbf{x}^k requires expensive linearization and matrix factorization, which becomes the time bottleneck.

Backpropagation with implicit time integration. We sketch the main idea of backpropagation with a loss function L defined on \mathbf{x} and explain how we can compute $\frac{\partial L}{\partial \mathbf{y}}$ from $\frac{\partial L}{\partial \mathbf{x}}$. Backpropagation through multiple timesteps can be done by backpropagating through every single pair of (\mathbf{x}, \mathbf{y}) from each timestep repeatedly.

As \mathbf{x} and \mathbf{y} are implicitly constrained by $\nabla g(\mathbf{x}) = \mathbf{0}$, we can differentiate it with respect to \mathbf{y} and obtain the following equation:

$$\frac{\partial \nabla g(\mathbf{x})}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{y}} + \frac{\partial \nabla g(\mathbf{x})}{\partial \mathbf{y}} = \frac{\partial}{\partial \mathbf{y}} \mathbf{0}, \quad (9)$$

$$\frac{\partial}{\partial \mathbf{x}} \left[\frac{1}{h^2}\mathbf{M}(\mathbf{x} - \mathbf{y}) + \nabla E(\mathbf{x}) \right] \frac{\partial \mathbf{x}}{\partial \mathbf{y}} + \frac{\partial}{\partial \mathbf{y}} \left[\frac{1}{h^2}\mathbf{M}(\mathbf{x} - \mathbf{y}) \right] = \mathbf{0}, \quad (10)$$

$$\underbrace{\left[\frac{1}{h^2}\mathbf{M} + \nabla^2 E(\mathbf{x}) \right]}_{\nabla^2 g(\mathbf{x})} \frac{\partial \mathbf{x}}{\partial \mathbf{y}} - \frac{1}{h^2}\mathbf{M} = \mathbf{0}, \quad (11)$$

$$\frac{\partial \mathbf{x}}{\partial \mathbf{y}} = \frac{1}{h^2}[\nabla^2 g(\mathbf{x})]^{-1}\mathbf{M}. \quad (12)$$

We can solve $\frac{\partial \mathbf{x}}{\partial \mathbf{y}}$ from it and use the chain rule to obtain the following (assuming both $\frac{\partial L}{\partial \mathbf{x}}$ and $\frac{\partial L}{\partial \mathbf{y}}$ are row vectors):

$$\frac{\partial L}{\partial \mathbf{y}} = \frac{\partial L}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{y}} = \frac{1}{h^2} \underbrace{\frac{\partial L}{\partial \mathbf{x}} [\nabla^2 g(\mathbf{x})]^{-1} \mathbf{M}}_{\mathbf{z}^\top}. \quad (13)$$

Note that the inverse of $\nabla^2 g(\mathbf{x})$ is intentionally regrouped with $\frac{\partial L}{\partial \mathbf{x}}$ to avoid the expensive $[\nabla^2 g(\mathbf{x})]^{-1}\mathbf{M}$. The adjoint vector \mathbf{z} can be solved from the following linear system of equations:

$$\nabla^2 g(\mathbf{x})\mathbf{z} = \left(\frac{\partial L}{\partial \mathbf{x}} \right)^\top. \quad (14)$$

Note that we drop the transpose of $\nabla^2 g(\mathbf{x})$ because it is symmetric.

Putting them together, we have shown that backpropagation within one timestep can be done by Equations (13) and (14). It is now clear that $\nabla^2 g(\mathbf{x})$ plays a crucial role in both forward simulation and backpropagation, and we write its definition explicitly below:

$$\nabla^2 g(\mathbf{x}) = \frac{1}{h^2}\mathbf{M} + \nabla^2 E(\mathbf{x}). \quad (15)$$

Similar to Newton's method in forward simulation, a direct implementation of Equation (14) is computationally expensive because $\nabla^2 g(\mathbf{x})$ needs to be reconstructed and refactorized at every timestep. This motivates us to propose the novel PD-based backpropagation method in Section 4.

PD. PD considers a specific family of quadratic potential energies that decouple the nonlinearity in material models [Bouaziz et al. 2014]. Specifically, PD assumes the energy E is the sum of quadratic energies taking the following form:

$$E_c(\mathbf{x}) = \min_{\mathbf{p}_c \in \mathcal{M}_c} \underbrace{\frac{w_c}{2} \|\mathbf{G}_c \mathbf{x} - \mathbf{p}_c\|_2^2}_{\tilde{E}_c(\mathbf{x}, \mathbf{p}_c)}, \quad (16)$$

$$E(\mathbf{x}) = \sum_c E_c(\mathbf{x}), \quad (17)$$

where \mathbf{G}_c is a discrete differential operator in the form of a constant sparse matrix, w_c a scalar that determines the stiffness of the energy, and \mathcal{M}_c a constraint manifold. For example, if one wants to formulate a volume-preserving elastic energy, \mathcal{M}_c can be the set of all 3×3 matrices whose determinant is 1. E_c is defined as the distance from $\mathbf{G}_c \mathbf{x}$ to \mathcal{M}_c . Following the prevalent practice in previous work [Bouaziz et al. 2014; Liu et al. 2017; Min et al. 2019],

ALGORITHM 1: PD forward simulation in one timestep

Input: \mathbf{y} ;
Output: \mathbf{x} that satisfies Equation (4);
Initialize $\mathbf{x} = \mathbf{y}$;
while \mathbf{x} not converged **do**
 $\mathbf{p}_c = \arg \min_{\mathbf{p}_c \in \mathcal{M}_c} \tilde{E}_c(\mathbf{x}, \mathbf{p}_c)$; // Local step;
 $\mathbf{b} = \frac{1}{h^2} \mathbf{M} \mathbf{y} + \sum_c w_c \mathbf{G}_c^\top \mathbf{p}_c$;
 $\mathbf{x} = \mathbf{A}^{-1} \mathbf{b}$; // Global step;

we assume E_c is defined on each finite element with \mathbf{G}_c mapping \mathbf{x} to the local deformation gradients \mathbf{F} [Sifakis and Barbic 2012].

With the definition of E at hand, PD obtains the critical point of g by alternating between a local step and a global step, which essentially minimizes the following surrogate objective $\tilde{g}(\mathbf{x}, \mathbf{p})$:

$$\tilde{g}(\mathbf{x}, \mathbf{p}) = \frac{1}{2h^2} (\mathbf{x} - \mathbf{y})^\top \mathbf{M} (\mathbf{x} - \mathbf{y}) + \sum_c \tilde{E}_c(\mathbf{x}, \mathbf{p}_c), \quad (18)$$

where \mathbf{p} stacks up all \mathbf{p}_c from each E_c . The local and global steps in PD can be interpreted as running coordinate descent optimization on \tilde{g} . The local step fixes the current \mathbf{x} and projects $\mathbf{G}_c \mathbf{x}$ onto \mathcal{M}_c to obtain \mathbf{p}_c in each E_c , which can be massively parallelizable across all E_c . The global step fixes \mathbf{p} and minimizes \tilde{g} over \mathbf{x} , which turns out to be a quadratic function with an analytical solution solved from the following linear system of equations:

$$\underbrace{\left(\frac{1}{h^2} \mathbf{M} + \sum_c w_c \mathbf{G}_c^\top \mathbf{G}_c \right)}_{\mathbf{A}} \mathbf{x} = \frac{1}{h^2} \mathbf{M} \mathbf{y} + \sum_c w_c \mathbf{G}_c^\top \mathbf{p}_c. \quad (19)$$

It is easy to see that each local and global step ensures \tilde{g} is non-increasing. Since \tilde{g} is bounded below by 0, PD guarantees to converge to a local minimum of \tilde{g} satisfying the gradient condition $\nabla_{\mathbf{x}} \tilde{g} = \mathbf{0}$. Interestingly, Liu et al. [2017] establishes that $\nabla_{\mathbf{x}} \tilde{g} = \nabla g$ upon convergence, confirming that the solution from PD is indeed a critical point of g that solves the implicit time integration. We summarize the local-global solver in Algorithm 1, which serves as a basis for our contact handling algorithm to be described in Section 5.2.

The source of efficiency in *forward* PD simulation lies in the fact that \mathbf{A} in the global step is a constant, symmetric positive definite matrix. Therefore, the Cholesky factorization of \mathbf{A} can be precomputed, after which each global step requires back-substitution only. In the next section, we will show that we can also use \mathbf{A} in *back-propagation* to obtain significant speedup.

4 DIFFERENTIABLE PROJECTIVE DYNAMICS

We now describe our PD-based backpropagation method. Our key observation is that the bottleneck in backpropagation lies in the computation of $\nabla^2 g(\mathbf{x})$ in Equation (14). Following the same idea as in forward PD simulation, we propose to decouple $\nabla^2 g(\mathbf{x})$ into a global, constant matrix and a local, massively parallelizable nonlinear component. To see this point, we compute $\nabla^2 E$ using

ALGORITHM 2: PD backpropagation in one timestep

Input: \mathbf{y} , \mathbf{x} (already computed in forward simulation), and $\frac{\partial L}{\partial \mathbf{x}}$;
Output: $\frac{\partial L}{\partial \mathbf{y}}$;
Initialize $\mathbf{z} = \mathbf{0}$;
while \mathbf{z} not converged **do**
 $\mathbf{b} = \Delta \mathbf{A} \mathbf{z} + \left(\frac{\partial L}{\partial \mathbf{x}} \right)^\top$; // Local step parallelizing $\Delta \mathbf{A} \mathbf{z}$;
 $\mathbf{z} = \mathbf{A}^{-1} \mathbf{b}$; // Global step;
 $\frac{\partial L}{\partial \mathbf{y}} = \frac{1}{h^2} \mathbf{z}^\top \mathbf{M}$; // Equation (13);

Equations (16) and (17):

$$\nabla E(\mathbf{x}) = \sum_c w_c \mathbf{G}_c^\top (\mathbf{G}_c \mathbf{x} - \mathbf{p}_c), \quad (20)$$

$$\nabla^2 E(\mathbf{x}) = \sum_c w_c \mathbf{G}_c^\top \mathbf{G}_c - \sum_c w_c \mathbf{G}_c^\top \frac{\partial \mathbf{p}_c}{\partial \mathbf{x}}. \quad (21)$$

Note that in Equation (20), $\frac{\partial \mathbf{p}_c}{\partial \mathbf{x}}$ can be safely ignored according to the envelope theorem (see Appendix in Liu et al. [2017]). According to Equation (15), $\nabla^2 g(\mathbf{x})$ now becomes

$$\nabla^2 g(\mathbf{x}) = \frac{1}{h^2} \mathbf{M} + \sum_c w_c \mathbf{G}_c^\top \mathbf{G}_c - \underbrace{\sum_c w_c \mathbf{G}_c^\top \frac{\partial \mathbf{p}_c}{\partial \mathbf{x}}}_{\Delta \mathbf{A}} = \mathbf{A} - \Delta \mathbf{A}. \quad (22)$$

It is now clear that $\Delta \mathbf{A}$ is the source of nonlinearity in $\nabla^2 g$. The matrix splitting of $\nabla^2 g = \mathbf{A} - \Delta \mathbf{A}$ suggests the following iterative solver for Equation (14):

$$\mathbf{A} \mathbf{z}^{k+1} = \Delta \mathbf{A} \mathbf{z}^k + \left(\frac{\partial L}{\partial \mathbf{x}} \right)^\top, \quad (23)$$

where k indicates the iteration number. Therefore, we propose a local-global solver for Equation (14): at the k -th iteration, the local step computes $\Delta \mathbf{A} \mathbf{z}^k$ across all energies E_c , forming the right-hand vector in Equation (23). In the global step, we solve \mathbf{z}^{k+1} by back-substituting \mathbf{A} . Note that \mathbf{A} is the same constant matrix in forward PD simulation, so we can reuse the Cholesky factorization of \mathbf{A} . The source of efficiency in this local-global solver is similar to what PD proposes to speed up forward simulation: essentially, this local-global solver trades the expensive matrix assembly and factorization of $\nabla^2 g$ in Equation (14) with iterations on a constant, prefactorized linear system of equations. We summarize our PD backpropagation algorithm in Algorithm 2.

Convergence rate. For any iterative algorithm design, the immediate follow-up questions are whether such an algorithm is guaranteed to converge, and, if so, how fast the convergence rate is. To answer these questions, we use $\left(\frac{\partial L}{\partial \mathbf{x}} \right)^\top = \mathbf{A} \mathbf{z} - \Delta \mathbf{A} \mathbf{z}$ and Equation (23) to obtain

$$\mathbf{A} (\mathbf{z}^{k+1} - \mathbf{z}) = \Delta \mathbf{A} (\mathbf{z}^k - \mathbf{z}), \quad (24)$$

from which we conclude the error at the k -th iteration is $\|\mathbf{z}^k - \mathbf{z}\|_2 = \|(\mathbf{A}^{-1} \Delta \mathbf{A})^k (\mathbf{z}^0 - \mathbf{z})\|_2$. It follows that the iteration in Equation (23) is guaranteed to converge from any initial guess \mathbf{z}^0 if and only if $\rho(\mathbf{A}^{-1} \Delta \mathbf{A}) < 1$, where $\rho(\cdot)$ indicates the spectral radius of a matrix. It is challenging to provide more theoretical results on $\rho(\mathbf{A}^{-1} \Delta \mathbf{A})$ because it depends heavily on the specific form of

E_c , which we leave as future work. In practice, we do not observe convergence issues with Equation (23) in any of our experiments, which seems to imply $\rho(\mathbf{A}^{-1}\Delta\mathbf{A}) < 1$ is likely to be satisfied.

Further acceleration with Quasi-Newton methods. Inspired by Liu et al. [2017] which apply the quasi-Newton method to speed up forward PD simulation, we now show that a similar numerical optimization perspective can also be applied to speed up our proposed local-global solver in backpropagation. Solving Equation (14) equals finding the critical point of the following energy $s(\mathbf{z})$:

$$s(\mathbf{z}) = \frac{1}{2}\mathbf{z}^\top \nabla^2 g(\mathbf{x})\mathbf{z} - \frac{\partial L}{\partial \mathbf{x}}\mathbf{z}. \quad (25)$$

It is easy to verify that $\nabla s(\mathbf{z}) = \mathbf{0}$ is essentially Equation (14). We stress that in backpropagation both $\nabla^2 g(\mathbf{x})$ and $\frac{\partial L}{\partial \mathbf{x}}$ are known values computed at \mathbf{x} solved from forward simulation. If we apply Newton's method to this critical-point problem, the update rule will be as follows (see Equations (7) and (8) in Section 3):

$$\mathbf{z}^{k+1} = \mathbf{z}^k - [\nabla^2 s(\mathbf{z}^k)]^{-1} \nabla s(\mathbf{z}^k). \quad (26)$$

The true Hessian of $s(\mathbf{z})$ is $\nabla^2 g(\mathbf{x}) = \mathbf{A} - \Delta\mathbf{A}$ from Equation (22). If we approximate it with \mathbf{A} , we get the following quasi-Newton update rule:

$$\begin{aligned} \mathbf{z}^{k+1} &= \mathbf{z}^k - \mathbf{A}^{-1} \nabla s(\mathbf{z}^k) \\ &= \mathbf{z}^k - \mathbf{A}^{-1} \left[(\mathbf{A} - \Delta\mathbf{A})\mathbf{z}^k - \left(\frac{\partial L}{\partial \mathbf{x}} \right)^\top \right] \\ &= \mathbf{A}^{-1} \Delta\mathbf{A}\mathbf{z}^k + \mathbf{A}^{-1} \left(\frac{\partial L}{\partial \mathbf{x}} \right)^\top, \end{aligned} \quad (27)$$

which is identical to the iteration in Equation (23). As a result, the local-global solver we propose can be reinterpreted as running a simplified quasi-Newton method with a constant Hessian approximation \mathbf{A} . By applying a full quasi-Newton method, e.g., BFGS, we can reuse the Cholesky decomposition of \mathbf{A} with little extra overhead of vector products and achieve a superlinear convergence rate [Nocedal and Wright 2006]. Moreover, similar to the previous work [Liu et al. 2017], we can apply line search techniques to ensure convergence when $\rho(\mathbf{A}^{-1}\Delta\mathbf{A}) \geq 1$, even though we do not experience convergence issues in practice.

5 CONTACT HANDLING

We have described the basic framework of DiffPD in Section 4. In this section, we propose a novel method to incorporate contact handling and contact gradients into DiffPD. The challenges in developing such a contact handling algorithm are twofold: first, it must be compatible with our basic PD framework in *both* forward simulation *and* backpropagation. Second, it must support differentiability. In this section, we discuss two contact options that DiffPD supports: an explicit, penalty-based contact model with static and dynamic friction and an implicit, complementarity-based contact model supporting non-penetration conditions and static friction. Both options have their advantages and disadvantages: the penalty-based method is more straightforward to implement and easier to be integrated into a machine learning framework, e.g., as an explicit neural network layer in PyTorch. However, we

find it typically requires a careful, scene-by-scene tuning of its parameters. On the other hand, our complementarity-based method does not rely on scene-dependent parameters, but it is currently limited to static friction only. Our penalty-based method is more suitable for tasks that favor speed and simplicity over physical accuracy, and the complementarity-based method is more useful when non-penetration conditions need to be strictly enforced and slipping motions are rare, e.g., simulating a wheeled robot.

5.1 Penalty-Based Contact

Previous papers on PD simulation typically handle contact forces with a penalty-based soft contact model [Bouaziz et al. 2014; Dinev et al. 2018a, b; Liu et al. 2017; Wang 2015; Wang and Yang 2016]. One common way to model contact forces is to add an additional, fictitious energy E_c with \mathcal{M}_c being the contact surface and its exterior and \mathbf{G}_c being a matrix so that $\mathbf{G}_c\mathbf{x}$ selects contact nodes from \mathbf{x} . This way, whenever a node penetrates the contact surface, E_c exerts a contact force that attempts to push it back to the contact surface. As such a contact model can be seamlessly integrated into PD forward simulation, our backpropagation method in Section 4 naturally supports it.

Handling static and dynamic frictional forces with a penalty-based model in PD is slightly trickier. Since friction is typically related to nodal velocities instead of nodal positions \mathbf{x} , it is not straightforward to find an E_c that characterizes it. Therefore, instead of modeling friction with an additional E_c , we take the penalty-based frictional forces described in Macklin et al. [2020] and add them directly to \mathbf{f}_{ext} . Deriving gradients with respect to such frictional forces is still straightforward as we can easily compute $\frac{\partial L}{\partial \mathbf{f}_{\text{ext}}}$ using the chain rule as before.

5.2 Complementarity-Based Contact

An alternative to penalty-based contact is to model contact and friction using complementarity constraints [Ly et al. 2020; Macklin et al. 2020]. Complementarity-based contact models are suitable for applications requiring high physical fidelity but typically require extra computational cost. Ly et al. [2020] present a general framework for handling complementarity-based contact and friction in PD forward simulation. More concretely, Ly et al. [2020] model contact and friction with the Signorini-Coulomb law and focus on applications in cloth simulation. Our approach is relevant to Ly et al. [2020] but has a substantial difference because our focus in this article is on 3D volumetric deformable bodies, which typically have a sparse contact set, i.e., the nodes in contact during simulation is usually a small portion of the full set of nodes. Below, we will present a differentiable, complementarity-based contact model that leverages such sparsity to gain speedup in both forward simulation and backpropagation. Our contact model ensures non-penetration conditions and, in exchange for speedup, handles static friction only. We leave a differentiable, complementarity-based dynamic friction model as future work.

Contact model. Let $\phi(\cdot) : \mathcal{R}^3 \rightarrow \mathcal{R}$ be the signed-distance function of the contact surface with $\phi < 0$ indicating the space occupied by the obstacle. We require the solution \mathbf{x} to the implicit time integration to satisfy the following complementarity condition for

any node indexed by j :

$$\begin{cases} \phi(\mathbf{x}_j) > 0, \mathbf{r}_j = \mathbf{0}, \\ \text{or } \phi(\mathbf{x}_j) = 0, \mathbf{r}_{j|N} \geq 0, \end{cases} \quad (28a)$$

$$(28b)$$

where \mathbf{x}_j and \mathbf{r}_j are 3D vectors indicating the nodal position and contact force of node j . The notation $\mathbf{r}_{j|N} \in \mathcal{R}$ is the normal component of \mathbf{r}_j where the normal is computed from the contact surface ϕ at the contact location \mathbf{x}_j . In other words, for each node j , it must be either above the contact surface ($\phi(\mathbf{x}_j) > 0$) with zero contact force ($\mathbf{r}_j = \mathbf{0}$) or in contact ($\phi(\mathbf{x}_j) = 0$) with a positive contact force along the normal direction ($\mathbf{r}_{j|N} \geq 0$). The implicit time integration in Equation (4) now becomes

$$\begin{cases} \frac{1}{h^2} \mathbf{M}(\mathbf{x} - \mathbf{y}) - \mathbf{f}_{\text{int}}(\mathbf{x}) = \mathbf{r}, \\ (\mathbf{x}, \mathbf{r}) \text{ satisfy Equation (28)}, \end{cases} \quad (29a)$$

$$(29b)$$

where the notation \mathbf{r} stacks up all contact force \mathbf{r}_j from each node j .

Remarks on friction. Equation (29) does not fully constrain the solution \mathbf{x} because, for any \mathbf{x}_j in contact with $\phi = 0$, \mathbf{x}_j can slide on $\phi = 0$ and \mathbf{r}_j will compensate any force needed. This can be resolved by imposing additional location constraints on \mathbf{x}_j . Some common strategies include (1) in the penalty-based model before, \mathbf{x}_j is chosen as certain projection onto $\phi = 0$, (2) gluing \mathbf{x}_j to its original position at the beginning of the timestep, and (3) setting it to the contact point computed from collision detection [Chen et al. 2017], which is usually the intersection between $\phi = 0$ and the ray $\mathbf{x}_j + t\mathbf{v}_j, 0 \leq t \leq h$. Any of these strategies are compatible with DiffPD as long as they can compute a target location \mathbf{x}_j^* on $\phi = 0$ if a collision detection algorithm indicates \mathbf{x}_j is in contact with $\phi = 0$. In DiffPD, we choose the third strategy mentioned above, which essentially models a very sticky contact surface that provides infinitely large static friction once \mathbf{x}_j is in contact.

Time integration with contact. We first introduce some notations to better explain our solver to Equation (29). Let $\mathcal{I} = \{0, 1, 2, \dots, n-1\}$ be the indices of all n nodes in the system. We use \bar{S} to denote the complement of a set $S \subseteq \mathcal{I}$. In other words, S and \bar{S} is a two-set partition of \mathcal{I} . For any subsets $S_r, S_c \subseteq \mathcal{I}$, we use $\mathbf{A}_{S_r \times S_c}$ to indicate the submatrix of \mathbf{A} created by keeping entries whose row and column indices are from nodes in S_r and S_c , respectively. Similarly, for any vector \mathbf{a} , we define \mathbf{a}_S as the vector generated by keeping elements whose indices are from nodes in S . For any vectors \mathbf{a} and $\mathbf{b} \in \mathcal{R}^{3|\mathcal{I}|}$, we use $\mathbf{a}_{S=\mathbf{b}}$ to indicate that \mathbf{a} satisfies $\mathbf{a}_S = \mathbf{b}$.

The high-level idea of our time integrator with the aforementioned contact model is described in Algorithm 3, which modifies Algorithm 1 to find \mathbf{x} that satisfies Equation (29). We start with any collision detection algorithm that can propose a set of candidate contact nodes C and compute a target location \mathbf{x}_j^* for any $j \in C$. Next, we use the proposed C to split the complementarity condition in Equation (29b) and solve Equation (29a): for any $j \in C$, we set $\mathbf{x}_j = \mathbf{x}_j^*$; for any $j \notin C$, we set $\mathbf{r}_j = \mathbf{0}$. This makes Equation (29a) a balanced system with an equal number of equations and variables. Finally, we use the solved \mathbf{x} to compute \mathbf{r}_j at each

$j \in C$ and check if $\mathbf{r}_{j|N} \geq 0$ is satisfied. If \mathbf{x} results in some negative $\mathbf{r}_{j|N}$, these nodes are removed from C , and a new iteration begins with the updated C . Similarly, if $\phi(\mathbf{x}_j)$ becomes negative, such a node j is added to C . Essentially, we are running the active-set algorithm on Equation (29a) with linear constraints, and more advanced active set schemes can potentially be used to rebuild C more efficiently. Using the notations above, our algorithm attempts to solve the following reduced system at each iteration:

$$\frac{1}{h^2} \mathbf{M}_{\bar{C} \times \bar{C}} (\mathbf{x} - \mathbf{y})_{\bar{C}} - \mathbf{f}_{\text{int}}(\mathbf{x}_{C=\mathbf{x}^*})_{\bar{C}} = \mathbf{0}, \quad (30)$$

where \mathbf{x}^* stacks up \mathbf{x}_j^* for all $j \in C$. Accordingly, the definition of g is updated as follows, which we rename as g_C :

$$g_C(\mathbf{x}_{C=\mathbf{x}^*}) = \frac{1}{2h^2} (\mathbf{x} - \mathbf{y})_{\bar{C}}^\top \mathbf{M}_{\bar{C} \times \bar{C}} (\mathbf{x} - \mathbf{y})_{\bar{C}} + E(\mathbf{x}_{C=\mathbf{x}^*}). \quad (31)$$

It is easy to check that the left-hand side of Equation (30) is identical to $\nabla_{\mathbf{x}_{\bar{C}}} g_C$. Therefore, solving Equation (30) is equal to finding the critical point of this modified g function, and we can still apply Newton's method but with a slightly different definition of $\nabla^2 g$:

$$\nabla_{\mathbf{x}_{\bar{C}}}^2 g_C = \frac{1}{h^2} \mathbf{M}_{\bar{C} \times \bar{C}} + (\nabla^2 E)_{\bar{C} \times \bar{C}} = (\nabla^2 g)_{\bar{C} \times \bar{C}}. \quad (32)$$

In other words, $\nabla_{\mathbf{x}_{\bar{C}}}^2 g_C$ is a submatrix of $\nabla^2 g$ in Equation (15) created by deleting rows and columns from C .

Implications on forward PD. In the PD framework, g_C also induces a modified surrogate function \tilde{g} , which we rename as \tilde{g}_C :

$$\begin{aligned} & \tilde{g}_C(\mathbf{x}_{C=\mathbf{x}^*}, \mathbf{p}) \\ &= \frac{1}{2h^2} (\mathbf{x} - \mathbf{y})_{\bar{C}}^\top \mathbf{M}_{\bar{C} \times \bar{C}} (\mathbf{x} - \mathbf{y})_{\bar{C}} + \sum_c \tilde{E}_c(\mathbf{x}_{C=\mathbf{x}^*}, \mathbf{p}_c). \end{aligned} \quad (33)$$

It is still true that the original local-global solver will ensure \tilde{g}_C is non-increasing and converge to a critical point of g_C . With the constraint $\mathbf{x}_C = \mathbf{x}^*$, the local step can project each $\mathbf{G}_c \mathbf{x}$ to obtain \mathbf{p}_c as before. The global step, on the other hand, requires some modification, as can be best seen after computing $\nabla_{\mathbf{x}_{\bar{C}}} \tilde{g}_C$:

$$\nabla_{\mathbf{x}_{\bar{C}}} \tilde{g}_C = \frac{1}{h^2} \mathbf{M}_{\bar{C} \times \bar{C}} (\mathbf{x} - \mathbf{y})_{\bar{C}} + \sum_c w_c (\mathbf{G}_c^\top)_{\bar{C} \times \mathcal{I}} (\mathbf{G}_c \mathbf{x} - \mathbf{p}_c). \quad (34)$$

Setting $\nabla_{\mathbf{x}_{\bar{C}}} \tilde{g}_C = \mathbf{0}$ and using the fact that $\mathbf{x}_C = \mathbf{x}^*$, we obtain the new global step with a linear system modified from Equation (19):

$$\mathbf{A}_{\bar{C} \times \bar{C}} \mathbf{x}_{\bar{C}} = \left[\frac{1}{h^2} \mathbf{M} \mathbf{y} + \sum_c w_c \mathbf{G}_c^\top (\mathbf{p}_c - \mathbf{G}_c \mathbf{x}_{C=\mathbf{x}^*, \bar{C}=\mathbf{0}}) \right]_{\bar{C}}, \quad (35)$$

where $\mathbf{x}_{C=\mathbf{x}^*, \bar{C}=\mathbf{0}}$ is a vector satisfying $\mathbf{x}_C = \mathbf{x}^*$ and $\mathbf{x}_{\bar{C}} = \mathbf{0}$. Although the right-hand side seems complicated, it can still be parallelized across all E_c . It is the left-hand side matrix $\mathbf{A}_{\bar{C} \times \bar{C}}$ that deserves more attention: since C is a set that changes dynamically between each timestep, $\mathbf{A}_{\bar{C} \times \bar{C}}$ randomly erases different rows and columns from \mathbf{A} , which means the Cholesky factorization of \mathbf{A} no longer applies. Our key observation is that C is usually a small subset of full nodes in 3D volumetric deformable bodies. This allows us to formulate row and column deletions on \mathbf{A} as a *low-rank update*, from which we derive efficient solvers that can reuse the Cholesky factorization of \mathbf{A} .

ALGORITHM 3: PD forward simulation with contact

```

Input:  $y$ ;
Output:  $x$  that satisfies Equation (29);
Run a collision detection algorithm to get  $C$  and  $x^*$ ;
while  $C$  not converged do
  Initialize  $x = y$  and set  $x_C = x^*$ ;
  while  $x$  not converged do
     $p_c = \arg \min_{p_c \in M_c} \tilde{E}_c(x, p_c)$ ; // Local step;
     $b = \frac{1}{h^2} M y + \sum_c w_c G_c^T (p_c - G_c x_{C=x^*, \bar{C}=0})$ ;
    // Global step;
    Run Algorithm 4 to solve  $x_{\bar{C}} = (A_{\bar{C} \times \bar{C}})^{-1} b_{\bar{C}}$ ;
     $r = \frac{1}{h^2} M(x - y) - f_{\text{int}}(x)$ ; // Equation (29a);
    Update  $C$  based on  $r_j, \phi(x_j)$ , and Equation (28);
  
```

Low-rank update. Define a permutation σ on \mathcal{I} with the following property: σ shuffles \mathcal{I} so that indices from C come before those in \bar{C} and the internal orders inside C and \bar{C} are preserved. Define \mathbf{P} as the corresponding permutation matrix: $\mathbf{P}_{ij} = 1$ if $\sigma(i) = j$ and 0 otherwise. Now $\mathbf{A}\mathbf{P}$ shuffles all columns of \mathbf{A} so that the i -th column in \mathbf{A} becomes the $\sigma(i)$ -th column in $\mathbf{A}\mathbf{P}$. Similarly, $\mathbf{P}^T \mathbf{A}$ shuffles all rows of \mathbf{A} in the same way. We now rewrite $\mathbf{P}^T \mathbf{A}\mathbf{P}$ as a 2×2 block matrix:

$$\mathbf{P}^T \mathbf{A}\mathbf{P} = \begin{pmatrix} \mathbf{A}_{C \times C} & \mathbf{A}_{C \times \bar{C}} \\ \mathbf{A}_{\bar{C} \times C} & \mathbf{A}_{\bar{C} \times \bar{C}} \end{pmatrix}. \quad (36)$$

Let $c = |C|$ and define $\mathbf{U} \in \mathcal{R}^{3n \times 2c}$ as follows:

$$\mathbf{U} = (\mathbf{U}_L \quad \mathbf{U}_R) = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\bar{C} \times C} \end{pmatrix}, \quad (37)$$

where $\mathbf{U}_L, \mathbf{U}_R \in \mathcal{R}^{3n \times c}$ represent the left and right halves of \mathbf{U} and \mathbf{I} the identity matrix of a proper size. Similarly, we define $\mathbf{V} \in \mathcal{R}^{2c \times 3n}$ as follows:

$$\mathbf{V} = \begin{pmatrix} \mathbf{U}_R^T \\ \mathbf{U}_L^T \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathbf{A}_{C \times \bar{C}} \\ \mathbf{I} & \mathbf{0} \end{pmatrix}. \quad (38)$$

It is now easy to verify that the product of $\mathbf{U}\mathbf{V}$ is the following low-rank matrix:

$$\mathbf{U}\mathbf{V} = \begin{pmatrix} \mathbf{0} & \mathbf{A}_{C \times \bar{C}} \\ \mathbf{A}_{\bar{C} \times C} & \mathbf{0} \end{pmatrix}, \quad (39)$$

and subtracting it from $\mathbf{P}^T \mathbf{A}\mathbf{P}$ results in a block-diagonal matrix:

$$\mathbf{P}^T \mathbf{A}\mathbf{P} - \mathbf{U}\mathbf{V} = \mathbf{P}^T \underbrace{(\mathbf{A} - \mathbf{P}\mathbf{U}\mathbf{V}\mathbf{P}^T)}_{\mathbf{A}_p} \mathbf{P} = \begin{pmatrix} \mathbf{A}_{C \times C} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\bar{C} \times \bar{C}} \end{pmatrix}. \quad (40)$$

Therefore, we can obtain $(\mathbf{A}_{\bar{C} \times \bar{C}})^{-1}$ by inverting $\mathbf{P}^T \mathbf{A}\mathbf{P} - \mathbf{U}\mathbf{V}$. Since inverting \mathbf{P} is trivial ($\mathbf{P}^{-1} = \mathbf{P}^T$), we focus on computing \mathbf{A}_p^{-1} using the Woodbury matrix identity:

$$\mathbf{A}_p^{-1} = \mathbf{A}^{-1} + \mathbf{A}^{-1} \mathbf{P}\mathbf{U}(\mathbf{I} - \mathbf{V}\mathbf{P}^T \mathbf{A}^{-1} \mathbf{P}\mathbf{U})^{-1} \mathbf{V}\mathbf{P}^T \mathbf{A}^{-1}. \quad (41)$$

Since \mathbf{A} is prefactorized, operations using \mathbf{A}^{-1} in the matrix identity above can be executed efficiently. Moreover, with the assumption that $c \ll n$, $\mathbf{I} - \mathbf{V}\mathbf{P}^T \mathbf{A}^{-1} \mathbf{P}\mathbf{U} \in \mathcal{R}^{2c \times 2c}$ is a small matrix compared to \mathbf{A} , and inverting it (solving a linear system whose left-hand side is this matrix) can be done efficiently. Putting everything

together, we transform the problem of factorizing \mathbf{A}_p into factorizing a much smaller linear system of equations.

Time complexity. We now consider a brute-force implementation of Equation (41) and analyze its time complexity. The time cost is dominated by computing $\mathbf{A}^{-1} \mathbf{P}\mathbf{U}$ which takes $O(n^2 c)$ time. While this is still asymptotically smaller than the cost of factorizing the modified matrix, which generally takes $O(n^3)$ time, the speedup in practice may not be as much as predicted due to the sparsity of \mathbf{A}_p . Therefore, further simplification on Equation (41) would still be desirable.

Further acceleration. To reduce the time cost of computing $\mathbf{A}^{-1} \mathbf{P}\mathbf{U}$, we notice that $\mathbf{P}\mathbf{U}$ shuffles all rows of \mathbf{U} with the inverse mapping σ^{-1} . As a result, $\mathbf{P}\mathbf{U}_L$, the left part of $\mathbf{P}\mathbf{U}$, is effectively $\mathbf{I}_{\mathcal{I} \times C}$, i.e., a collection of one-hot column vectors $\mathbf{e}_j, j \in C$, where the j -th entry in \mathbf{e}_j is 1. This means that we can precompute $\mathbf{A}^{-1} \mathbf{I}_{\mathcal{I} \times C}$ using a maximum possible C (e.g., all surface nodes) before the whole simulation begins and look up $\mathbf{A}^{-1} \mathbf{e}_j, j \in C$ on the fly.

It turns out that the same idea can also be used for computing $\mathbf{A}^{-1} \mathbf{P}\mathbf{U}_R$, the right half of the solution, with a slight modification. Notice that $\mathbf{P}\mathbf{U}_R$ can be obtained from \mathbf{A} by fetching $\mathbf{A}_{\mathcal{I} \times C}$ and zeroing out corresponding rows in C :

$$\mathbf{P}\mathbf{U}_R = \mathbf{A}_{\mathcal{I} \times C} - \mathbf{I}_{\mathcal{I} \times C} \mathbf{A}_{C \times C}. \quad (42)$$

We can, therefore, compute $\mathbf{A}^{-1} \mathbf{P}\mathbf{U}_R$ as follows:

$$\begin{aligned} \mathbf{A}^{-1} \mathbf{P}\mathbf{U}_R &= \mathbf{A}^{-1} \mathbf{A}_{\mathcal{I} \times C} - \mathbf{A}^{-1} \mathbf{I}_{\mathcal{I} \times C} \mathbf{A}_{C \times C} \\ &= \mathbf{I}_{\mathcal{I} \times C} - \mathbf{A}^{-1} \mathbf{I}_{\mathcal{I} \times C} \mathbf{A}_{C \times C}. \end{aligned} \quad (43)$$

Since $\mathbf{A}^{-1} \mathbf{I}_{\mathcal{I} \times C}$ has been precomputed, the time complexity will be bounded by the matrix multiplication $O(nc^2)$. Moreover, noting that \mathbf{A}^{-1} is symmetric and \mathbf{V} can be obtained from \mathbf{U} by swapping and transposing block matrices \mathbf{U}_L and \mathbf{U}_R , the results derived here can also be reused to assemble $\mathbf{V}\mathbf{P}^T \mathbf{A}^{-1}$.

In conclusion, we have reduced the time complexity of computing $\mathbf{A}^{-1} \mathbf{P}\mathbf{U}$ from $O(n^2 c)$ to $O(nc^2)$. Since the remaining operations, excluding solving \mathbf{A}^{-1} in Equation (41), are also bounded by $O(nc^2)$, we now have reduced the *overhead* of applying Equation (41) from $O(n^2 c)$ to $O(nc^2)$, with the overhead defined as the extra cost brought by Equation (41) in addition to one linear solve \mathbf{A}^{-1} with any right-hand side vector. We present the complete algorithm in pseudocode in Algorithm 4, which serves as a subroutine in Algorithm 3. We use \mathbf{B}_k and \mathbf{a}_k to denote intermediate matrices and vectors, respectively, with the subscript k indicating the order of their first occurrence.

Backpropagation. With a contact set C and the corresponding x^* , the backpropagation scheme in Section 4 also needs modifications. Backpropagating from $\frac{\partial L}{\partial x}$ to $\frac{\partial L}{\partial y}$ now becomes trickier due to the existence of C and x^* from a collision detection algorithm, which splits both x and y into two vectors $x_C, x_{\bar{C}}, y_C$, and $y_{\bar{C}}$. Here, we will sketch the core idea by showing how gradients can be backpropagated from $\frac{\partial L}{\partial x_{\bar{C}}}$ to $\frac{\partial L}{\partial y_{\bar{C}}}$. Backpropagation through other dependencies is easier to derive and therefore skipped.

ALGORITHM 4: Global step in Algorithm 3.

```

Input:  $C \subseteq \mathcal{I}$ ,  $\mathbf{x}^* \in \mathcal{R}^{3|C|}$ , and  $\mathbf{b} \in \mathcal{R}^{3n}$ ;
Output:  $\mathbf{x}$  such that  $\mathbf{x}_C = \mathbf{x}^*$  and  $\mathbf{A}_{\overline{C} \times \overline{C}} \mathbf{x}_{\overline{C}} = \mathbf{b}_{\overline{C}}$ ;
Collect  $\mathbf{B}_1 = \mathbf{A}^{-1} \mathbf{I}_{\mathcal{I} \times C}$  from precomputed data;
 $\mathbf{B}_2 = \mathbf{I}_{\mathcal{I} \times C} - \mathbf{B}_1 \mathbf{A}_{C \times C}$ ;
 $\mathbf{B}_3 = (\mathbf{B}_1, \mathbf{B}_2)$ ; //  $\mathbf{B}_3 = \mathbf{A}^{-1} \mathbf{P} \mathbf{U}$ ;
//  $\mathbf{V} \mathbf{P}^\top$  can be fetched from  $\mathbf{A}$  without permutation;
// No need to compute  $\mathbf{P}$ ;
 $\mathbf{B}_4 = \mathbf{I} - \mathbf{V} \mathbf{P}^\top \mathbf{B}_3$ ;
Solve  $\mathbf{a}_1$  from  $\mathbf{A} \mathbf{a}_1 = \mathbf{b}$ ;
 $\mathbf{a}_2 = (\mathbf{b}^\top \mathbf{B}_2, \mathbf{b}^\top \mathbf{B}_1)$ ; // Row vector;
Solve  $\mathbf{a}_3$  from  $\mathbf{B}_4 \mathbf{a}_3 = \mathbf{a}_2^\top$ ;
 $\mathbf{x} = \mathbf{a}_1 + \mathbf{B}_3 \mathbf{a}_3$ ;
Set  $\mathbf{x}_C = \mathbf{x}^*$ ;

```

From Equations (30) and (31), we see that $\mathbf{x}_{\overline{C}}$ and $\mathbf{y}_{\overline{C}}$ are constrained by $\nabla_{\mathbf{x}_{\overline{C}}} g_C = \mathbf{0}$. By differentiating Equation (30), we obtain

$$\nabla_{\mathbf{x}_{\overline{C}}}^2 g_C \frac{\partial \mathbf{x}_{\overline{C}}}{\partial \mathbf{y}_{\overline{C}}} - \frac{1}{h^2} \mathbf{M}_{\overline{C} \times \overline{C}} = \mathbf{0}, \quad (44)$$

which is a reduced version of Equation (9). The chain rule still applies in a similar way:

$$\frac{\partial L}{\partial \mathbf{y}_{\overline{C}}} = \frac{\partial L}{\partial \mathbf{x}_{\overline{C}}} \frac{\partial \mathbf{x}_{\overline{C}}}{\partial \mathbf{y}_{\overline{C}}} = \frac{1}{h^2} \underbrace{\frac{\partial L}{\partial \mathbf{x}_{\overline{C}}} [(\nabla^2 g)_{\overline{C} \times \overline{C}}]^{-1}}_{\mathbf{z}^\top} \mathbf{M}_{\overline{C} \times \overline{C}}, \quad (45)$$

where a new adjoint vector \mathbf{z} is defined. It should now become very clear that \mathbf{z} is obtained from the following linear system of equations:

$$(\nabla^2 g)_{\overline{C} \times \overline{C}} \mathbf{z} = \left(\frac{\partial L}{\partial \mathbf{x}_{\overline{C}}} \right)^\top. \quad (46)$$

Now using Equation (22), we see the iterative solver in Section 4 becomes

$$\mathbf{A}_{\overline{C} \times \overline{C}} \mathbf{z}^{k+1} = \Delta \mathbf{A}_{\overline{C} \times \overline{C}} \mathbf{z}^k + \left(\frac{\partial L}{\partial \mathbf{x}_{\overline{C}}} \right)^\top, \quad (47)$$

from which we see a similar issue we experience in forward simulation: $\mathbf{A}_{\overline{C} \times \overline{C}}$ changes dynamically, so the Cholesky factorization of \mathbf{A} is not directly applicable. This is exactly where we can use the same global solver in Algorithm 4 to retain the source of efficiency in our PD backpropagation algorithm. We summarize this new backpropagation method in Algorithm 5.

Summary. In summary, we have presented a differentiable contact handling algorithm that ensures non-penetration conditions and imposes infinitely large static friction. Moreover, we have also discussed its implementation in forward simulation and backpropagation that can still benefit from the Cholesky factorization of \mathbf{A} . We stress that there exist more physically accurate contact handling algorithms that satisfy not only non-penetration conditions but also the Coulomb's law of friction [Chen et al. 2017; Li et al. 2020; Ly et al. 2020]. However, our contact handling algorithm

ALGORITHM 5: PD backpropagation with contact

```

Input:  $\mathbf{y}$ ,  $\mathbf{x}$ , and  $C$  (from forward simulation), and  $\frac{\partial L}{\partial \mathbf{x}_{\overline{C}}}$ ;
Output:  $\frac{\partial L}{\partial \mathbf{y}_{\overline{C}}}$ ;
Initialize  $\mathbf{z} = \mathbf{0}$ ;
while  $\mathbf{z}$  not converged do
   $\mathbf{b} = (\Delta \mathbf{A})_{\overline{C} \times \overline{C}} \mathbf{z} + \left( \frac{\partial L}{\partial \mathbf{x}_{\overline{C}}} \right)^\top$ ; // Local step;
  Run Algorithm 4 to solve  $\mathbf{z} = (\mathbf{A}_{\overline{C} \times \overline{C}})^{-1} \mathbf{b}$ ; // Global
  step;
 $\frac{\partial L}{\partial \mathbf{y}_{\overline{C}}} = \frac{1}{h^2} \mathbf{z}^\top \mathbf{M}_{\overline{C} \times \overline{C}}$ ; // Equation (13);

```

achieves a good tradeoff between differentiability, physical plausibility, and compatibility with our differentiable PD framework.

6 EVALUATION

In this section, we compare DiffPD with a few baseline differentiable simulation methods and conduct ablation studies on the acceleration techniques in Section 4 and Section 5. We start by discussing the difference between implicit and explicit timestepping schemes in backpropagation. Next, we compare our simulator with two other fully implicit simulators implemented with the Newton's method. We end this section with a discussion on the two contact models implemented in DiffPD. The end goal of this section is to evaluate the difference between different timestepping methods and understand the source of efficiency in DiffPD. We implement both baseline algorithms and DiffPD in C++ and use Eigen [Guennebaud et al. 2010] for sparse matrix factorization and linear solvers. We run all experiments in this section and the next section on a virtual machine instance from Google Cloud Platform with 16 Intel Xeon Scalable Processors (Cascade Lake) @ 3.1 GHz and 64 GB memory. We use OpenMP for parallel computing and eight threads by default unless otherwise specified.

6.1 Comparisons with Explicit Method

Compared with explicit timestepping methods used in previous papers on differentiable simulation [Hu et al. 2020, 2019; Spielberg et al. 2019], implicit time integration brings two important changes to a differentiable simulator: first, implicit methods enable a much larger timestep during simulation, resulting in much fewer number of frames. This is particularly beneficial for solving an inverse problem with a long time horizon as we store fewer states (nodal positions and velocities) in memory during backpropagation. Second, due to implicit damping, we can expect the landscape of the loss function defined on nodal states and their derived quantities to be smoother.

To demonstrate the memory consumption, we consider a soft cantilever discretized into $12 \times 3 \times 3$ hexahedral elements (a low-resolution version of the ‘‘Cantilever’’ example in Figure 1, left). We impose Dirichlet boundary conditions on one end of the cantilever and simulate its vibration after twisting the other end of the cantilever under gravity for 0.2 seconds. We define a loss function L as a randomly generated weighted average of the final nodal positions and velocities. The implicit time integration in our simulator allows us to use a timestep as large as 10 milliseconds,




Fig. 1. The “Cantilever” and “Rolling sphere” examples in Section 6 designed for comparing DiffPD to the Newton’s method. The “Cantilever” example starts with a twisted cantilever (left), oscillates, and bends downwards eventually due to gravity. In the “Rolling sphere” example, we roll a soft sphere on the ground (right) which constantly breaks and reestablishes contact.




Fig. 2. The relative changes in both the loss (top left) and the magnitude of the gradient (bottom left) for the explicit method (cyan) and our method (green) for 5 out of the 16 random directions in the neighborhood of the initial nodal positions \mathbf{x}_0 . Also shown are the means (solid curves) and standard deviations (shaded) of the percent change in loss (top right) and the magnitude of the gradient (bottom right) for all 16 random directions.

while an explicit implementation is only numerically stable in both forward simulation and backpropagation for timesteps of 0.5 milliseconds. Since memory consumption during backpropagation is proportional to the number of frames, we can expect a 20× increase in memory consumption for the explicit method, requiring additional techniques like checkpoint states [Hu et al. 2019; Spielberg et al. 2019] before the problem size can be scaled up.

To demonstrate the influences of timestepping schemes on the smoothness of the energy landscape, we visualize in Figure 2 the loss function L and its gradient norm $|\nabla L|$ sliced along 16 random directions in the neighborhood of the cantilever’s initial nodal positions. Specifically, let \mathbf{x}_0 be the initial nodal positions, and let $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{16}$ be the random directions. We plot $L(\mathbf{x}_0 + \alpha \mathbf{r}_i)$ and $|\nabla L(\mathbf{x}_0 + \alpha \mathbf{r}_i)|$ for each \mathbf{r}_i (Figure 2, left) with α being the step size, which is uniformly sampled between -0.3% and 0.3% of the cantilever beam length. The standard deviations from 16 random directions (Figure 2, right) indicate that small perturbations in \mathbf{x}_0 lead to much smoother loss and gradients when implicit time integration is applied, which is not surprising due to numerical damping. From the perspective of differentiable simulation, a smoother energy landscape can be more favorable as it induces more well-defined gradients to be used by gradient-based optimization techniques.

6.2 Comparisons with Other Implicit Methods

We now compare our simulator with other implicit timestepping schemes to evaluate its speedup in both forward and backward modes. We choose Newton’s method with two standard sparse linear solvers: an iterative solver using preconditioned conjugate gradient (Newton-PCG) and a direct solver using Cholesky decomposition (Newton-Cholesky), as our baseline solvers for implicit time integration in Equation (3). First, we compare with the Newton’s method in Section 4 without contact. Then, we benchmark the performance of the contact handling algorithm in Section 5. We reiterate that just like in the standard PD framework, any resultant speedup from DiffPD over Newton’s method is under the assumption that the material model has a quadratic energy function. We extend our discussion to general hyperelastic materials at the end of the article and leave it as future work.

Simulation without contact. We benchmark our method, Newton-PCG, and Newton-Cholesky using a cantilever with $32 \times 8 \times 8$ elements, 8019 DoFs, and 243 Dirichlet boundary constraints (“Cantilever” in Figure 1 and Table 3). The example runs for 25 frames with timesteps of 10 milliseconds. We define the loss L as a randomly generated weighted sum of the final nodal positions and velocities.

In terms of the running-time comparison, we report results in Figure 3 from running all three methods with two, four, and eight threads and a range of convergence threshold (from $1e-1$ to $1e-7$) on the relative error in solving Equation (3). The speedup from parallel computing is less evident in the Newton’s method because the majority of their computation time is spent on matrix refactorization—a process that cannot be trivially parallelized in Eigen. We conclude that our simulator has a clear advantage over Newton’s method on the time cost of both forward simulation and backpropagation. For forward simulation, the speedup is well understood and discussed in many previous PD papers [Bouaziz et al. 2014; Liu et al. 2017]. For moderate tolerances ($1e-3$ to $1e-5$), we observe a speedup of 9–16 times in forward simulation with eight threads and note that it becomes less significant as precision increases. Both of these observations agree with previous work on PD for forward simulation. In backpropagation, DiffPD method is 6–13 times faster than the Newton’s method for moderate tolerances due to the reuse of the Cholesky decomposition and the quasi-Newton update. Specifically, we point out that without the proposed acceleration technique with quasi-Newton methods in Section 4, PD backpropagation is faster than the Newton’s method only for very low precision (orange in Figure 3, right), confirming the necessity of the quasi-Newton updates.

Since PD is an iterative method whose result is dependent on the convergence threshold, it is necessary to justify which threshold




Fig. 3. Top: the net wall-clock times (left), forward times (middle), and backpropagation times (right) for different convergence thresholds tested on the “Cantilever” example in Section 6.2. The results are obtained from simulating this example using each of the three methods: Newton-PCG (PCG), Newton-Cholesky (Cholesky), and DiffPD (Ours). The number following the method name denotes the number of threads used. Also shown are the results from running our method without applying the quasi-Newton method (orange, right). Bottom: the loss (left) and magnitude of the gradient of the loss (right) for different convergence thresholds used to terminate iterations in Newton’s method and our simulator.

is the most proper. To analyze the influence of the choice of thresholds, we use results from Newton-Cholesky as the oracle because it is a direct solver whose solution is computed with the machine precision in Eigen. We then compare both our method and Newton-PCG with the oracle by computing the loss and gradients of the “Cantilever” example with varying convergence thresholds and analyze when the results from the three methods start to coincide. This comparison provides quantitative guidance on the choice of convergence threshold and reveals the range in which our method can be a reliable alternative to the Newton’s method in optimization tasks. We report our findings in Figure 3. As Newton-PCG and our method are iterative methods, their accuracy improves when the convergence threshold becomes tighter. It can be seen from the figure that our method agrees with the Newton’s method on the numerical losses and gradients when using a threshold as large as $1e-4$. Therefore, we use $1e-4$ as our default threshold in all applications to be discussed below unless otherwise specified. Referring back to Figure 3, using eight threads and with a convergence threshold of $1e-4$, our method achieves significant speedup (12–16 times faster in forward simulation and 6.5–9 times faster in backpropagation) compared with Newton-PCG and Newton-Cholesky.

Simulation with contact. To create a benchmark scene that requires contact handling constantly, we roll a soft sphere on a hori-

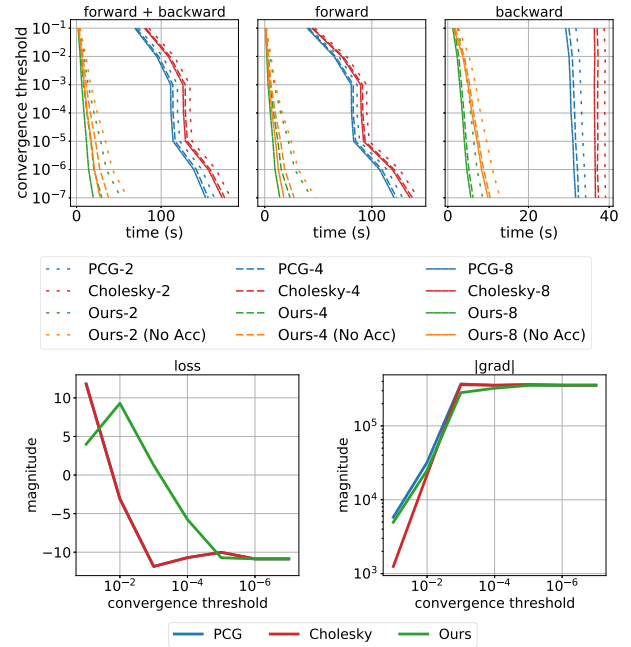


Fig. 4. Top: the net wall-clock times (left), forward times (middle), and backpropagation times (right) for different convergence thresholds tested on the “Rolling sphere” example for contact handling. The results are obtained from three methods: Newton-PCG (PCG), Newton-Cholesky (Cholesky), and DiffPD (Ours). The number following the method name denotes the number of threads used. Also shown are the results from running our method without applying the further acceleration technique (Alg. 4) in Section 5.2 (orange). Bottom: the loss (left) and magnitude of the gradient of the loss (right) for different convergence thresholds used to terminate iterations in the Newton’s method and our simulator.

zontal collision plane for 100 frames with a timestep of 5 milliseconds (“Rolling sphere” in Figure 1 and Table 3). The sphere is voxelized into 552 elements with 2,469 DoFs, and the maximum possible contact set C we consider consists of 72 nodes (216 DoFs) on the surface of the sphere. Similar to the “Cantilever” example, we define the loss function L as a randomly generated weighted average of the final nodal positions and velocities. We implement the contact handling algorithm in Section 5.2 with Newton-PCG, Newton-Cholesky, and our method, and we report their time cost as well as their loss and gradients in Figure 4. It can be seen from Figure 4 that the results from three methods start to converge when the convergence threshold reaches $1e-6$, with which our method is 10 times faster than the Newton’s method in both forward and backward mode (Figure 4). Such speedup mainly comes from the low-rank update algorithm (Algorithm 4) which avoids the expensive matrix factorization from scratch. Additionally, by comparing the orange and green curves in Figure 4, we conclude that the acceleration technique of caching $A^{-1}I_{I \times C}$ further speeds up DiffPD by 25% in forward mode and 44% in backward mode when measured with eight threads and a convergence threshold of $1e-6$.

6.3 Ablation Study

We end this section with an ablation study on multiple components in our algorithm. We start with an empirical analysis on the

iterative solver and the line search algorithm in our backpropagation algorithm (Section 4), followed by an evaluation on the penalty-based and the complementarity-based contact models.

Spectral radius and line search. One key assumption we have made in our backpropagation solver is that the spectral radius of $\rho(\mathbf{A}^{-1}\Delta\mathbf{A}) < 1$, which is also one of the primary reasons why we have employed the line search algorithm as a safeguard when the assumption does not hold. Here, we use the “Cantilever” example check if this assumption holds empirically. We explicitly calculate $\rho(\mathbf{A}^{-1}\Delta\mathbf{A})$ we experience in “Cantilever” and observe a maximum value of 0.996, indicating that we can expect convergence in the iterative solver, which we further confirm by testing the iterative solver with 100 randomly generated, artificial right-hand side vector $\frac{\partial L}{\partial \mathbf{x}}$. We observe similar results about the convergence of the backpropagation solver in the “Rolling sphere” example as well as in our applications to be described in Section 7, indicating that it seems safe to expect the iterative solver to converge in practice despite the lack of a theoretical guarantee on it.

As employing line searches in our algorithm serves as a safeguard to cases when $\rho(\mathbf{A}^{-1}\Delta\mathbf{A}) > 1$, an implication from the observations on the spectral radius is that we rarely trigger line searches to reduce the step size in practice. In fact, in this “Cantilever” example, and in almost all applications below, we notice that the default step size (1 in the Newton’s and quasi-Newton methods) almost always allows us to skip the line search stage. Still, we precautionarily set the maximum number of line search iterations to be 10 for all examples.

Penalty-based contact. We implement the penalty-based contact and frictional forces from Macklin et al. [2020] in DiffPD and analyze them in both forward simulation and backpropagation. First, we use a standard “Slope” test with varying frictional coefficients in the penalty-based model to understand the expressiveness of this contact model in forward simulation. Second, we use a “Duck” example which optimizes frictional coefficients using the gradients of this contact model in backpropagation.

To show the capacity of the penalty-based contact model in forward simulation, we consider the “Slope” test visualized in Figure 5. We place a squishy rubber duck (16,776 DoFs and 24,875 tetrahedrons) on four slopes with varying frictional coefficients from the penalty form in Macklin et al. [2020] and let it slide for 2 seconds under gravity. We can see from the figure that with decreasing sliding friction from the left slope to the right slope, the implementation of Macklin et al. [2020] in DiffPD generates different sliding distances that match our expectation qualitatively.

Backpropagating a penalty-based contact model is straightforward because it only requires a procedural application of chain rules to differentiate the penalty energy. To show the penalty-based model is fully compatible with DiffPD’s backpropagation and can be useful in optimization problems, we design a “Duck” example (Figure 6) with the same rubber duck but on a curved slide with frictional coefficients to be optimized (three DoFs in total). The duck slides off the curved surface and aims to land on a target location (indicated by the white circle). The frictional coefficients affect the stickiness of the curve surface and control the exiting velocity of the duck when it leaves the slide, which further




Fig. 5. **Slope.** A rubber duck (16,776 DoFs and 24,875 tetrahedrons) slides on slopes with varying frictional coefficients implemented in DiffPD using a penalty-based contact and friction model [Macklin et al. 2020]. Left: the initial position of the duck. Middle left to right: the final positions of the duck after 2 seconds with a decreasing frictional coefficient.




Fig. 6. **Duck.** The same rubber duck in Figure 5 now slides on a curved surface with trainable frictional coefficients. The goal in this test is to optimize the frictional coefficients so that the duck’s final position after 1 second of simulation reaches the center of the white circle as closely as possible. We overlay the intermediate positions of the rubber duck at 0 s, 0.25 s, 0.5 s, 0.75 s, and 1 s in simulation with an initial guess of the frictional coefficients before optimization (left) and the final coefficients after gradient-based optimization (right).

determines its movement under gravity afterwards. From the two motion sequences in Figure 6 before and after gradient-based optimization, we observe a substantial improvement that eventually leads the duck to the target position. This confirms the usefulness of gradients computed in DiffPD using the penalty-based contact method.

Complementarity-based contact. For the contact model described in the complementarity form, our backpropagation algorithm assumes the contact set is a small subset of full DoFs. Specifically, Algorithm 4 requires a relatively small size of C at each timestep to gain substantial speedup over directly solving the modified linear system without leveraging the low-rank update. Given that C is a subset of surface vertices, whose number is much fewer than the number of interior vertices in a typical 3D volumetric deformable body, such an assumption can be easily satisfied in many applications. Indeed, in the next section, we will present various 3D examples involving contact, none of which have more than 6% active contact nodes throughout simulation.

The assumption that $|C|$ is relatively small is much more likely to be violated when we simulate a co-dimensional object, e.g., a one-dimensional rope or a piece of cloth in 3D, in which case it is entirely possible to have all nodes in C at some point. Although simulating co-dimensional objects is beyond the scope of this work, it can be a good test to reveal a critical ratio where the speedup

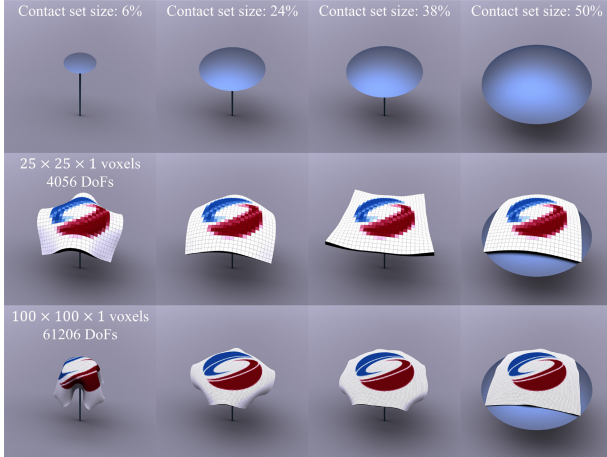


Fig. 7. Ablation study on the relative size of the active contact set $|C|$ and its influence on DiffPD’s speed. We simulate a one-layer “Napkin” with resolutions from $25 \times 25 \times 1$ voxels (middle row) to $100 \times 100 \times 1$ voxels (bottom row) falling onto a spherical obstacle (blue, top row) with a varying solid angle. We report the relative size of $|C|$ and the degrees of freedom of the napkin. Please refer to Table 2 for the detailed running time and our video for the full motion of the falling napkin.

from Algorithm 4 starts to diminish. To mimic a co-dimensional object, we engineer a “Napkin” example consisting of one-layer voxels (Figure 7) falling onto a spherical obstacle with an adjustable solid angle to control the size of $|C|$. The relative size of $|C|$ is capped by 50% when all the bottom nodes are in contact with the spherical obstacle (Figure 7, right column). We vary the mesh resolution from $25 \times 25 \times 1$ voxels (4,056 DoFs) to $100 \times 100 \times 1$ voxels (61,206 DoFs) and report the running time of Newton’s method and DiffPD in Table 2 for each resolution and contact set size. We can use Table 2 to decide between using our low-rank update method and directly solving the modified matrix in a downstream application. For example, for around 15k DoFs, Table 2 suggests that the low-rank update method is faster until the relative size of C reaches around 40%.

7 APPLICATIONS

In this section, we show various tasks that can benefit from DiffPD and classify them into five categories: system identification, inverse design, trajectory optimization, closed-loop control, and real-to-sim applications. Although prior efforts on differentiable simulators have demonstrated their capabilities on almost all these examples, we highlight that DiffPD is able to achieve comparable results but reduce the time cost by almost an order of magnitude. We provide a summary of each example in Table 3. For examples with actuators, we implement the contractile fiber model as discussed in Min et al. [2019]. Regarding the optimization algorithm, we use L-BFGS in our examples by default unless otherwise specified. We report the time cost and the final loss after optimization in Table 4. For fair comparisons, we use the same initial guess and termination conditions when running L-BFGS with different simulation methods. When reporting the loss in Table 4, we linearly normalize it so that a loss of 1 represents the average performance

Table 2. Average Running Time Per Step in the Napkin Example With Various Mesh Resolution

Res.	Method	6%	24%	38%	50%
$25 \times 25 \times 1$ (4,056 DoFs)	Newton-PCG	0.8	1.6	1.6	1.5
	DiffPD	0.1	0.6	1.2	1.4
$50 \times 50 \times 1$ (15,606 DoFs)	Newton-PCG	5.4	8.4	10.2	8.5
	DiffPD	1.2	6.0	11.3	10.5
$75 \times 75 \times 1$ (34,656 DoFs)	Newton-PCG	14.7	25.1	25.2	22.5
	DiffPD	7.1	29.0	42.8	48.1
$100 \times 100 \times 1$ (61,206 DoFs)	Newton-PCG	44.6	65.0	47.3	48.5
	DiffPD	32.0	169.8	158.7	163.4

(“res”) and relative contact set size from 6% to 50% (Figure 7, top row). The reported time is averaged over all steps when the napkin is in contact with the obstacle. All times are in seconds. For each mesh resolution and each relative contact set, we report the running time from both the Newton’s method and DiffPD with the shorter time in bold.




Fig. 8. **System identification.** Motion sequences of the “Plant” example sampled at the 1st, 100th, 150th, and 200th (final) frames (left to right). We generated three motion sequences with a random initial guess of the material parameters (top row), optimized material parameters (middle row), and the ground truth (bottom row). The colored boxes highlight the motion differences before and after optimization. The goal is to optimize the material parameters so that the motion of the plant matches the ground truth.

from 16 randomly sampled solutions and a loss of 0 maps to a desired solution. For examples using a bounded loss, we map zero loss to an oracle solution that achieves the lower bound of the loss (typically 0). For unbounded losses used in the walking and swimming robots (Sections 7.3 and 7.4), we map zero loss to the performance of solutions obtained from DiffPD. The full details about each experiment can be found in our Supplemental Material and our source code.

7.1 System Identification

In this section, we discuss two examples that aim to estimate the material parameters (Young’s modulus and Poisson’s ratio) from dynamic motions of soft bodies: the “Plant” example estimates material parameters from its vibrations, and the “Bouncing ball” example predicts its parameters from its interaction with the ground. We generate the ground truth using our forward PD simulator with a set of predefined material parameters.

Plant. We first initialize an elastic, 3D house plant model with 3, 863 hexahedral elements and 29, 763 DoFs (Figure 8). We impose

Table 3. The Setup of All Examples in the Article

Sec.	Task name	# of elements	# of DoFs	h (ms)	# of steps	Gravity	Dirichlet	Contact	Hydrodynamics	Actuation
6.2	Cantilever	2,048	8,019	10	25	✓	✓			✓
	Rolling sphere	552	2,469	5	100	✓		✓		
7.1	Plant	3,863	29,763	10	200		✓			
	Bouncing ball	1,288	9,132	4	125	✓		✓		
7.2	Bunny	1,601	7,062	1	100	✓		✓		
	Routing tendon	512	2,475	10	100	✓	✓			✓
7.3	Torus	568	3,204	4	400	✓		✓		✓
	Quadruped	648	3,180	10	100	✓		✓		✓
	Cow	475	2,488	1	600	✓		✓		✓
7.4	Starfish	1,492	7,026	33.3	200				✓	✓
	Shark	2,256	9,921	33.3	200				✓	✓
7.5	Tennis balls	640	978	5.6	150	✓		✓		

The right five columns report whether the example has gravity as an external force, imposes Dirichlet boundary conditions on nodal positions, handles contact, requires hydrodynamical forces, and has actuators.

Table 4. The Performance of DiffPD on All Examples

Sec.	Task name	Newton-PCG				Newton-Cholesky				DiffPD (Ours)				
		Fwd.	Back.	Eval.	Loss	Fwd.	Back.	Eval.	Loss	Fwd.	Back.	Eval.	Loss	Speedup
6.2	Cantilever	118.2	39.4	-	-	160.1	55.9	-	-	10.5	5.5	-	-	10×
	Rolling sphere	107.3	31.3	-	-	135.6	36.6	-	-	14.0	5.7	-	-	8×
7.1	Plant	1,089.5	530.5	10	1.9e-3	929.6	525.2	10	1.9e-3	71.6	94.7	28	5.9e-7	9×
	Bouncing ball	269.3	90.9	43	7.9e-2	262.6	102.5	22	8.4e-2	15.8	14.2	12	9.6e-2	12×
7.2	Bunny	277.7	88.0	21	7.0e-3	358.2	126.9	29	5.1e-3	24.0	17.3	11	2.3e-2	9×
	Routing tendon	108.2	56.7	36	6.0e-4	107.3	58.7	38	4.9e-4	8.3	9.9	30	9.6e-4	9×
7.3	Torus	751.9	210.3	47	-2.3e-3	719.9	212.4	43	-2.4e-2	84.3	81.9	27	0	6×
	Quadruped	289.2	51.5	69	-1.8e0	246.3	47.8	54	-1.1e0	50.2	15.8	30	0	4×
	Cow	771.7	141.7	14	9.7e-1	620.1	140.2	20	9.8e-1	105.3	43.7	31	0	5×
7.4	Starfish	217.7	105.1	100	4.8e-1	244.0	129.4	100	1.4e-1	5.7	10.8	100	0	19×
	Shark	260.7	159.3	100	9.8e-1	599.4	241.8	100	-9.0e-3	35.5	15.3	100	0	8×
7.5	Tennis balls	54.6	6.4	14	7.2e-2	26.8	5.8	12	7.2e-2	24.1	15.9	41	6.9e-2	0.8×

For each method and each example, we report the time cost of evaluating the loss function once (Fwd.) and its gradients once (Back.) in the unit of seconds. Also shown are the number of evaluations of the loss and its gradients (Eval.) in BFGS optimization. The "Loss" column reported the normalized final loss after optimization (lower is better) with the best one in bold. Finally, we report the speedup computed as the ratio between the forward plus backward time of the Newton's method and of DiffPD, i.e., ratio between the sum of "Fwd." and "Back." columns. Note that no speedup is gained from DiffPD in the real-to-sim example "Tennis balls" because of its too small number of DoFs (Table 3).

Dirichlet boundary conditions at the root of the plant such that it is fixed to the ground. We apply an initial horizontal force at the start of simulation, causing the plant to oscillate. Starting from an initial guess using randomly picked material parameters, we deform a new plant in the same manner as the ground truth and optimize the logarithm of the Young's modulus and Poisson's ratio of the new plant to match that of the ground truth plant. The loss at each timestep is determined as the squared sum of the elementwise difference in positions between the new plant and the reference plant.

After optimization, DiffPD, Newton-PCG, and Newton-Cholesky converge to local minima with a final Young's modulus of 1.00 MPa, 0.96 MPa, and 0.96 MPa, respectively. Regarding the Poisson's ratio, DiffPD converge to 0.4 while both Newton's methods converged to 0.44. The reference plant is initialized with a Young's modulus of 1 MPa and a Poisson's ratio of 0.4. While the three methods all reach solutions that are similar to the ground truth, the optimization process is highly expedited by

a factor of 9 for loss and gradient evaluation using our method (Table 4). We observe that DiffPD converged to a solution closer to the ground truth but used more function evaluations due to the numerical difference between DiffPD and the Newton's method. However, if DiffPD terminated after the same number of function evaluations (10) as the Newton's method, the optimized Young's modulus and Poisson's ratio would be almost identical to results from the Newton's method (0.97 MPa for Young's modulus and 0.44 for Poisson's ratio), implying the 9× speedup indeed comes from DiffPD's improvements over the Newton's method on the simulation side.

Bouncing ball. In this example, we consider a ball with 1,288 hexahedral elements and 9,132 DoFs thrown at the ground from a known initial position (Figure 9). The ball has three cylindrical holes extruded through the faces in order to produce more complex deformation behavior than a fully solid ball. This example uses the complementarity-based contact model in Section 5.2. We can




Fig. 9. **System identification.** Motion sequences of the “Bouncing ball” example sampled at the 1st frame (left), the 19th frame when collision occurs (middle), and the 125th (final) frame (right). We generated three motion sequences with a random initial guess of the material parameters (top row), optimized material parameters (middle row), and the ground truth (bottom row). The goal is to optimize the material parameters so that the motion of the ball matches the ground truth.

estimate the material parameters of a bouncing ball by observing its behavior after it collides with the ground. The loss definition is the same as in the parameter estimation of the “Plant” example. Regarding the optimization process, all three methods converge to a Young’s modulus of 1.78 MPa and Poisson’s ratio of 0.2. The ground truth values for the Young’s modulus and Poisson’s ratio are 2 MPa and 0.4, respectively. While the optimized material parameters are significantly different from the ground truth values, the motion sequences are very similar as reflected by the final loss in Table 4 and Figure 9. Since the loss function is defined on the motion only, there could exist many material parameters that result in close-to-zero loss. As in the “Plant” example, our method enjoys a substantial speedup (12 \times) in computation time even with collisions in simulation.

7.2 Initial State Optimization

We present two examples demonstrating the power of using gradient information to optimize the initial configuration of a soft-body task. In the “Bunny” example, we optimize the initial position and velocity of a soft Stanford bunny so that its bounce trajectory ends at a target position. In the “Routing tendon” example, we optimize a constant actuation signal applied to each muscle in a soft cuboid with one face sticky on the ground so that the corner at the opposite face reaches a target point at the end of simulation.

Bunny. For this example, we optimize the initial pose and velocity of a Stanford bunny (1,601 elements and 7,062 DoFs) so that its center of mass (red dots in Figure 10) reaches a target position (blue dot in Figure 10) when the simulation finishes. This example uses the complementarity-based contact model, and we add 251 surface vertices (753 DoFs) to the set of possible contact nodes—approximately 10.7% of the 2,354 vertices. Figure 10 illustrates the trajectory of the bunny before and after optimization: the initial guess generates a trajectory almost to the opposite direction of the target, and the optimized trajectory ends much closer to the target.




Fig. 10. **Inverse design.** Initial (left) and optimized (right) trajectories of the “Bunny” example. The red dots indicate the location of the center of mass of the bunny, and the blue dot is the target location. The goal is to adjust the initial position, velocity, and orientation of the bunny so that its final center of mass can reach the target location.

Note that none of the three methods solve this task perfectly: the trajectory does not reach the target even after optimization. This is because the target is chosen arbitrarily rather than generated from simulating a ground truth bounce trajectory, so it is not guaranteed to be reachable. Table 4 shows that the final loss from DiffPD is larger than from the Newton’s method, but the increase in performance makes up for it. Using eight threads, our method achieves a speedup of nine times overall with a large set of potential contact points.

Routing tendon. We initialize a soft cuboid with 512 elements and 2,475 DoFs and impose Dirichlet boundary conditions such that its bottom face is stuck to the ground. We also add actuators to each element and group them into 16 muscle groups. The level of actuator activation is a scalar between 0 and 1, indicating muscle contraction and expansion, respectively. The elements within a specific actuation group all share the same, time-invariant actuation signal to be optimized in order to manipulate the endpoint (red dot in Figure 11) of the soft body to reach a target point (blue dot in Figure 11). The normalized losses at the final iteration for each of the methods (Table 4) are all close to zero, indicating that the task is solved almost perfectly. Using DiffPD, we observe a 9 \times speedup over the Newton’s methods.

7.3 Trajectory Optimization

Torus. In our first trajectory optimization example, a torus is tasked with rolling forward as far as possible in 1.6 seconds, simulated as 400 steps of 4 milliseconds in length (Figure 12). To achieve this, we set the objective to be the negation of the robot’s center of mass at the final step of its trajectory. Eight muscle tendons are routed circumferentially along the center of the torus, combined, creating a circle that can be actuated along any of eight segments. The optimization variables are the actuation of each muscle at each of 20 linearly spaced knot points, and the actual action sequences are generated by linearly interpolating variables at these knot points. Since there are eight muscles, this results in 160 decision variables overall. We use a convergence threshold of $1e-6$ as indicated by the evaluation experiment in Section 6.2.

The major challenge in optimizing the sequence of actions of this rolling torus lies in the fact that it constantly breaks and reestablishes contact with the ground. When running L-BFGS on this example, we noticed more local minima than previous examples and L-BFGS often terminated prematurely without making




Fig. 11. **Inverse design.** The initial and final states of the “Routing tendon” example before and after optimization. The goal is to let the end effector (red dot) hit a target position (blue dot) at the end of simulation. Left: the initial configuration of the tendon with randomly generated actuation signals. The red (muscle expansion) and cyan (muscle contraction) colors indicate the magnitude of the action. Middle left: final state of the tendon with random actuation. Middle right: initial configuration of the tendon with optimized actuation. Right: final state with optimized actuation.

significant progress. To alleviate this issue, we randomly sampled 16 initial solutions and selected the best among them to initialize L-BFGS optimization, which eventually learned a peristaltic contraction pattern that allows it to start rolling forward and make considerable forward progress (Figure 12); further, DiffPD provides a 6× speedup using eight threads compared to the Newton’s method.

In order to demonstrate the applicability of our system’s differentiability to solving complex trajectory optimization tasks, we apply our simulator to three locomotion tasks: a “Torus,” a “Quadruped,” and a “Cow.” All three robots are equipped with muscle fibers whose sequences of actions are to be optimized, and the goal for all three robots is to walk forward without losing balance or drifting sideways. All examples use the complementarity-based contact model in Section 5.2.

Quadruped. For our second trajectory optimization example, a rectangular quadruped is tasked with moving forward as far as possible in 1 second. The same performance objective is applied here as in the “Torus” example; however, in this example a simpler control scheme is implemented. This robot has eight muscles, routed vertically along the front and back face of each leg, allowing the legs to bend forward or backward. For each leg, the front and back muscle groups are paired antagonistically; however, they are allowed a different maximum actuation strength—a parameter to be optimized. Finally, the entire quadruped is provided a single sinusoidal control signal, whose frequency is to be optimized, that actuates each leg synchronously. These front and back actuation strengths, combined with the frequency of the input signal, provide three parameters to be optimized. After optimization, the quadruped was able to walk forward several body lengths (Figure 13). In terms of the speedup, DiffPD accelerates loss and gradient evaluation by a factor of 4 compared with the Newton’s method.

Cow. For our third and final trajectory optimization example, a cow quadruped based off Spot [Crane 2020] is tasked with walking forward as far as possible in 0.6 seconds (Figure 14). This is a particularly difficult task, as Spot’s oversized head makes it front-heavy, and prone to falling forward. In order to compensate, we regularized the objective to promote a more upright gait, adding an addi-

tional -0.3 times the center of mass in z to regularize the forward objective. Spot uses the same controller and muscle arrangement as the “Quadruped” example, and a convergence threshold of $1e-6$ is used during optimization. Similar to previous examples, the cow optimizes to walk forward and DiffPD provides a 5 times speedup compared to the Newton’s method.

Discussion. Locomotion tasks generally involve significant contact, which limits the speedups (4–6 times in the examples above) DiffPD is able to achieve compared with contact-free problems. Given the complexity of planning the motion of walking robots with contacts (optimizing each of the three examples above took hours to converge with the Newton’s method), a 4–6 times speedup is still favorable. It is also worth noting that for the “Quadruped” and “Cow” examples, optimization with the Newton’s method led to solutions significantly different from DiffPD, as indicated by the final loss reported in Table 4. We believe this is mostly due to the algorithmic difference between the Newton’s method and DiffPD: as discussed in Liu et al. [2017] and Section 4, DiffPD is essentially running the quasi-Newton method (as opposed to the Newton’s method) to minimize the objective in Equation (5) which is typically not convex. Therefore, multiple critical points may exist especially when contacts are involved. For the three locomotion tasks in this section, it is possible that DiffPD and two Newton’s methods each explored different critical points individually and led to different solutions.

7.4 Closed-Loop Control

Inspired by Min et al. [2019], we consider designing a closed-loop neural network controller for two marine creatures: “Starfish” and “Shark” (Figure 15). For each example, we specify muscle fibers as internal actuators similar to Min et al. [2019] in the arms of the starfish and the caudal fin of the shark. We manually place velocity sensors on the body of each example serving as the input to the neural network controller. The goal of these examples is to optimize a swimming controller so that each fish can advance without drifting sideways. To achieve this, we define the loss function as a weighted sum of forward velocities and linear velocities at each timestep. In terms of the neural network design, we choose a three-layer multilayer perceptron network with 64 neurons in each layer (30,788 parameters in “Starfish” and 22,529 parameters in “Shark”). We use the hyperbolic tangent function as the activation function in the neural network. Unlike prior examples for which L-BFGS is used for optimization, we follow the common practice of using gradient descent with Adam [Kingma and Ba 2015] to optimize the neural network parameters. During optimization, we use a convergence threshold of $1e-3$ in DiffPD and the Newton’s method. Table 4 provides the final loss after optimization, and we observe a speedup of 8–19 times for the two examples, respectively.

Comparisons with reinforcement learning. We compare our gradient-based optimization method to PPO [Schulman et al. 2017], a state-of-the-art reinforcement learning algorithm. In particular, we use the forward simulation of DiffPD as the simulation environment for PPO. For a fair comparison, we construct and initialize the network for both DiffPD and PPO with the same random seed. We also implement code-level optimization techniques as suggested




Fig. 12. **Trajectory optimization.** The motion sequence of the “Torus” example with random actions (top) and after optimizing the action sequences (160 parameters to be optimized) with DiffPD (bottom). The goal is to maximize the rolling distance of the torus while maintaining its balance. The red and cyan color indicates the magnitude of the action signal. In particular, the expansion and contraction pattern (middle left and middle right) allows the torus to roll forward.




Fig. 13. **Trajectory optimization.** The motion sequences of the “Quadruped” example with sinusoidal control signals whose three parameters are to be optimized. The goal is to maximize the walking distance of the quadruped. Top: the motion sequence with a random sinusoidal wave of actions. Bottom: the motion sequence after optimization with DiffPD.




Fig. 14. **Trajectory optimization.** The motion sequences of the “Cow” example with sinusoidal control signals whose three parameters are to be optimized. The goal is to maximize the walking distance of the cow. Top: the motion sequence with a random sinusoidal wave of actions. Bottom: the motion sequence after optimization with DiffPD.

in Engstrom et al. [2019] and tune PPO hyperparameters toward its best performance. Please refer to our Supplemental Material for more implementation details.

When comparing the performance of a gradient-free algorithm like PPO to gradient-based algorithms like Adam or L-BFGS, we expect gradient-based optimization to be more sampling efficient than PPO as gradients expose more information about the soft-body dynamics that are not accessible to PPO. Note that this does not ensure gradient-based methods are always faster than PPO when measured by their wall-clock time, because each sample in a gradient-based method requires additional gradient computation time. Furthermore, the sampling scheme in PPO is massively parallelizable. We report the optimization progress of PPO and our method in Figure 16. Note that unlike other examples, we follow

the convention in reinforcement learning of maximizing a reward as opposed to minimizing a loss. In particular, a zero reward indicates the average performance of randomly selected unoptimized neural networks, and a unit reward is the result from DiffPD after optimization. We conclude from Figure 16 that Adam and DiffPD achieves comparable results to PPO but is more sampling efficient by one or two orders of magnitude. Regarding the wall-clock time, we observe a speedup of 9–11 times for both examples, respectively, although each sample in DiffPD is more expensive due to the gradient computation.

7.5 A Real-to-Sim Experiment

We end this section with a real-to-sim experiment “Tennis balls.” In this example, we capture a video clip of two colliding tennis balls




Fig. 15. **Control optimization.** Motion sequences from the optimized closed-loop, neural network controllers of “Starfish” (first row) and “Shark” (third row), with the corresponding muscle fibers plotted in the second and fourth rows. The goal is to optimize a controller so that the marine creatures can rise (“Starfish”) or swim forward (“Shark”). The gray and cyan colors on the surface of the muscle fibers indicate the magnitude of the actuation, with gray being zero actuation and blue the maximum contraction.

on a flat terrain, from which we aim to estimate the camera information, the initial position and velocity of each ball, and the parameters in the contact model. We model each ball in simulation using a mesh of sphere with 320 tetrahedrons and 489 DoFs (Table 3). We model the ball-ground contact with the complementarity-based method and use the following penalty-based model to compute the ball-ball contact: when the two balls are in contact, we add a restitution force computed as the product of a stiffness parameter to be optimized and the difference between the ball diameter and the actual distance between the two ball centers. Essentially, the restitution force can be treated as a spring model with a rest length equal to the ball diameter. Additionally, we add a frictional force whose direction is orthogonal to the restitution force and whose magnitude is controlled by a frictional coefficient to be optimized.

To define a loss function that measures the discrepancy between the simulated and actual motions of the two balls, we first extract the pixel location of two balls’ centers in each frame of the video clip. Next, we compute in simulation the position of each ball and project them to the same image space through a pinhole camera model. We define the loss function as the difference between the pixel locations of the balls in simulation and from the video clip. By minimizing this loss, we get our estimation of the camera information, the initial state of each ball, and the parameters in the contact model.

We summarize our optimization results in Table 4 and Figure 17. We randomly sample multiple sets of parameters and pick those with the smallest loss as the initial guess to our optimization (Figure 17, left), which shows motion sequences similar to those in

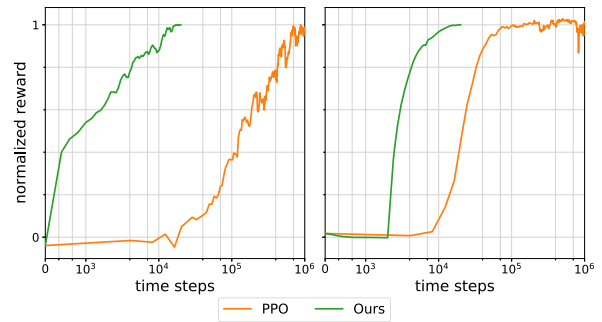


Fig. 16. The optimization progress of Adam plus DiffPD (green) and PPO (orange) for “Starfish” (left) and “Shark” (right). Note that the axis of timesteps spent during training or optimization is plotted on a logarithmic scale.

the video clip but still with substantial visual differences. The optimization process refines our estimation of the parameters and manages to further suppress the loss and mimics the motions in the video more closely (Figure 17, middle). The results can be further improved if we take into account camera lens distortion or replace the penalty-based collision model between two balls with a more accurate one, which we leave as future work.

8 LIMITATIONS AND FUTURE WORK

Differentiable soft-body simulation with proper contact handling is a challenging problem due to its large number of DoFs and




Fig. 17. **A Real-to-Sim Experiment.** Motion sequences of the “Tennis balls” example before (left) and after optimization (middle). The corresponding video clip is shown on the right. Transparent balls indicate the balls’ intermediate locations. To visualize the difference between motion sequences in simulation and reality, we use yellow squares in the rendered images (left and middle) to denote the corresponding pixel locations of the balls’ centers from the video clip.




Fig. 18. Twisting Armadillo (a 44337-DoF tetrahedral mesh) with the Neo-hookean material model for 1 second at 30 frames per second. We use DiffPD (top) and the Newton’s method (bottom) to compute the forward simulation and backpropagation. Left to right: the intermediate state of Armadillo at 0, 10, 20, and 30 frames. The visually identical motions from the top and bottom rows confirm the correctness of DiffPD’s implementation of the Neo-hookean material model. We report the time cost of DiffPD and the Newton’s method at the lower right corner of the images and no longer witness a speedup from DiffPD in backpropagation over a direct solver in the Newton’s method.

complexity in resolving contact forces. We believe one ambitious direction along this line of research is to provide a physically realistic simulator that can facilitate the design and control optimization of *real* soft robots. To close the sim-to-real gap, some nontrivial but rewarding enhancements need to be integrated into our current implementation. First and foremost, similar to other PD papers, a major limitation in DiffPD is its assumption on the energy function of the material model. Even though technical solutions to supporting general hyperelastic materials in PD exist [Liu et al. 2017], it turns out that supporting such materials in DiffPD is not straightforward. This is because the derivation in Section 4 starts to fall apart from Equations (20) and (21) when hyperelastic material models are used, forcing DiffPD to reassemble the Hessian matrix $\nabla^2 E$ in each timestep during backpropagation. Although we can still apply the iterative solver from Section 4 in this case, we no longer observe a speedup over a direct solver (Figure 18). Therefore, we switch to the direct solver for backpropagation in DiffPD when hyperelastic materials are used and leave speeding it up as future work.

Second, our contact methods do not fully resolve differentiable, complementarity-based contact and friction. Due to the focus of

this article, our choice of the contact model was intentionally biased toward ensuring differentiability and compatibility with PD. It would be more accurate and useful to upgrade the contact models for both static and sliding frictional forces [Ly et al. 2020] or to apply a more realistic contact model [Li et al. 2020] while maintaining its efficiency and differentiability in PD.

A third direction that is worth exploring is to improve the scalability of our algorithm. Currently, the largest example in this article contains thousands of elements and tens of thousands of DoFs. It would be desirable to scale problems up by at least one or two orders of magnitude in order to explore the effects of more complex geometry. This would obviously be computationally expensive; it would therefore be interesting to explore possible GPU implementations of DiffPD method to unlock large-scale applications.

Fourth, although DiffPD is substantially faster than the standard Newton’s method when assumptions in PD hold, the speedup is less significant for locomotion tasks (4–6 times in our examples). We suspect it is the inclusion of contact that slows down DiffPD both in forward simulation and in backpropagation. Therefore, a more comprehensive analysis on the assumption of sparse contact in Section 5 would possibly reveal the source of inefficiency. Specifically, removing such an assumption would be much desired to unlock contact-rich applications, e.g., cloth simulation or manipulation.

Finally, there is room for improving the optimization strategies that can better leverage the benefits of gradients. In all our examples, we couple gradient information with gradient-based continuous optimization methods. Being inherently local, such methods inevitably suffer from terminating at local minima prematurely especially when the loss function has a non-convex landscape. It is worth exploring the field of global optimization methods or even combining ideas from gradient-free strategies, e.g., genetic algorithms or reinforcement learning, to present a more robust global optimization algorithm specialized for differentiable simulation.

ACKNOWLEDGMENTS

We thank Desai Chen, David I. W. Levin, Bo Zhu, and Eftychios Sifakis for their feedback and suggestions on this article. The duck and cow mesh models in Figures 5, 6, and 14 are obtained from Keenan Crane’s 3D model repository [Crane 2020] under the CC0 1.0 Universal license.

REFERENCES

- J r mie Allard, St phane Cotin, Fran ois Faure, Pierre-Jean Bensusoussan, Fran ois Poyer, Christian Duriez, Herv  Delingette, and Laurent Grisoni. 2007. –An open source framework for medical simulation. *Stud Health Technol Inform* 125 (2007), 13–8. <https://pubmed.ncbi.nlm.nih.gov/17377224/>.
- Jernej Barbi  and Doug James. 2005. Real-time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph.* 24, 3 (July 2005), 982–990.
- James Bern, Pol Banzet, Roi Poranne, and Stelian Coros. 2019. Trajectory optimization for cable-driven soft robot locomotion. In *Robotics: Science and Systems*.
- Josh Bongard, Cecilia Laschi, Hod Lipson, Nick Cheney, and Francesco Corucci. 2016. Material properties affect evolutions ability to exploit morphological computation in growing soft-bodied creatures. In *Artificial Life Conference Proceedings 13*. MIT Press, 234–241.
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: Fusing constraint projections for fast simulation. *ACM Trans. Graph.* 33, 4 (July 2014), Article 154, 11 pages.
- Christopher Brandt, Elmar Eiseemann, and Klaus Hildebrandt. 2018. Hyper-reduced projective dynamics. *ACM Trans. Graph.* 37, 4 (July 2018), Article 80, 13 pages.
- Desai Chen, David Levin, Wojciech Matusik, and Danny Kaufman. 2017. Dynamics-aware numerical coarsening for fabrication design. *ACM Trans. Graph.* 36, 4 (July 2017), Article 84, 15 pages.
- Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. 2013. Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*. 167–174.
- Francesco Corucci, Nick Cheney, Hod Lipson, Cecilia Laschi, and Josh Bongard. 2016. Evolving swimming soft-bodied creatures. In *ALIFE XV, the 15th International Conference on the Synthesis and Simulation of Living Systems, Late Breaking Proceedings*, Vol. 6.
- Keenan Crane. 2020. *Keenan’s 3D Model Repository*. Retrieved Aug. 18, 2020 from <https://www.cs.cmu.edu/~kmc Crane/Projects/ModelRepository/>.
- Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and Zico Kolter. 2018. End-to-end differentiable physics for learning and control. In *Advances in Neural Information Processing Systems*. 7178–7189.
- Jonas Degraeve, Michiel Hermans, Joni Dambre, and Francis Wyffels. 2019. A differentiable physics engine for deep learning in robotics. *Frontiers in Neurorobotics* 13 (2019), 6.
- Dimitar Dinev, Tiantian Liu, and Ladislav Kavan. 2018a. Stabilizing integrators for real-time physics. *ACM Trans. Graph.* 37, 1 (Jan. 2018), Article 9, 19 pages.
- Dimitar Dinev, Tiantian Liu, Jing Li, Bernhard Thomaszewski, and Ladislav Kavan. 2018b. FEPR: Fast energy projection for real-time simulation of deformable objects. *ACM Trans. Graph.* 37, 4 (July 2018), Article 79, 12 pages.
- Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. 2019. Implementation matters in deep RL: A case study on PPO and TRPO. In *International Conference on Learning Representations*.
- Marco Fratarcangeli, Valentina Tibaldo, and Fabio Pellacini. 2016. Vivace: A practical Gauss-Seidel method for stable soft body dynamics. *ACM Trans. Graph.* 35, 6 (Nov. 2016), Article 214, 9 pages.
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz B cher, Bernhard Thomaszewski, and Stelian Coros. 2020. ADD: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph.* 39, 6 (Nov. 2020), Article 190, 15 pages.
- Ga l Guennebaud, Beno t Jacob, et al. 2010. Eigen v3. Retrieved Jun. 3, 2020 from <http://eigen.tuxfamily.org>.
- David Hahn, Pol Banzet, James Bern, and Stelian Coros. 2019. Real2Sim: Visco-elastic parameter estimation from dynamic motion. *ACM Trans. Graph.* 38, 6 (Nov. 2019), Article 236, 13 pages.
- Xuchen Han, Theodore Gast, Qi Guo, Stephanie Wang, Chenfanfu Jiang, and Joseph Teran. 2019. A hybrid material point method for frictional contact with diverse materials. *Proc. ACM Comput. Graph. Interact. Tech.* 2, 2 (July 2019), Article 17, 24 pages.
- Philipp Herholz and Olga Sorkine-Hornung. 2020. Sparse Cholesky updates for interactive mesh parameterization. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–14.
- Jonathan Hiller and Hod Lipson. 2014. Dynamic simulation of soft multimaterial 3D-Printed objects. *Soft Robotics* 1, 1 (2014), 88–101.
- Philipp Holl, Nils Thuerey, and Vladlen Koltun. 2020. Learning to control PDEs with differentiable physics. In *International Conference on Learning Representations*.
- Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Fr do Durand. 2020. DiffTaichi: Differentiable programming for physical simulation. *International Conference on Learning Representations* (2020).
- Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B. Tenenbaum, William T. Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. 2019. Chainqueen: A real-time differentiable physical simulator for soft robotics. *International Conference on Robotics and Automation* (2019).
- Doug James and Dinesh Pai. 1999. ArtDefo: Accurate real time deformable objects. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH’99)*. ACM Press/Addison-Wesley Publishing Co., 65–72.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Seunghwan Lee, Ri Yu, Jungnam Park, Mridul Aanjaneya, Eftychios Sifakis, and Jehee Lee. 2018. Dexterous manipulation and control with volumetric muscles. *ACM Trans. Graph.* 37, 4 (July 2018), Article 57, 13 pages.
- Jing Li, Tiantian Liu, and Ladislav Kavan. 2019a. Fast simulation of deformable characters with articulated skeletons in projective dynamics. In *Proceedings of the 18th Annual ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA’19)*. Association for Computing Machinery, New York, NY, Article 1, 10 pages.
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny Kaufman. 2020. Incremental potential contact: Intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.* 39, 4 (July 2020), Article 49, 20 pages.
- Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua Tenenbaum, and Antonio Torralba. 2019b. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *International Conference on Learning Representations*.
- Junbang Liang, Ming Lin, and Vladlen Koltun. 2019. Differentiable cloth simulation for inverse problems. In *Advances in Neural Information Processing Systems*. 772–781.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-Newton methods for real-time simulation of hyperelastic materials. *ACM Trans. Graph.* 36, 4 (May 2017), Article 116a, 16 pages.
- Micka l Ly, Jean Jouve, Laurence Boissieux, and Florence Bertails-Descoubes. 2020. Projective dynamics with dry frictional contact. *ACM Trans. Graph.* 39, 4 (July 2020), Article 57, 8 pages.
- M. Macklin, K. Erleben, M. M ller, N. Chentanez, S. Jeschke, and T. Y. Kim. 2020. Primal/Dual descent methods for dynamics. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA’20)*. Eurographics Association, Goslar, DEU, Article 9, 12 pages.
- Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. 2011. Example-based elastic materials. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH’11)*. Association for Computing Machinery, New York, NY, Article 72, 8 pages.
- Antoine McNamara, Adrien Treuille, Zoran Popovi , and Jos Stam. 2004. Fluid control using the adjoint method. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 449–456.
- Sehee Min, Jungdam Won, Seunghwan Lee, Jungnam Park, and Jehee Lee. 2019. Soft-con: Simulation and control of soft-bodied animals with biomimetic actuators. *ACM Trans. Graph.* 38, 6 (Nov. 2019), Article 208, 12 pages.
- Rahul Narain, Matthew Overby, and George E. Brown. 2016. ADMM \supseteq Projective Dynamics: Fast Simulation of General Constitutive Models. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA’16)*. Eurographics Association, Goslar, DEU, 21–28.
- Jorge Nocedal and Stephen Wright. 2006. *Numerical Optimization*. Springer Science & Business Media.
- Matthew Overby, George Brown, Jie Li, and Rahul Narain. 2017. ADMM \supseteq projective dynamics: Fast simulation of hyperelastic models with dynamic constraints. *IEEE Trans. Vis. Comput. Graph.* 23, 10 (2017), 2222–2234.
- Jovan Popovi , Steven Seitz, and Michael Erdmann. 2003. Motion sketching for control of rigid-body simulations. *ACM Trans. Graph.* 22, 4 (Oct. 2003), 1034–1054.
- Yi-Ling Qiao, Junbang Liang, Vladlen Koltun, and Ming Lin. 2020. Scalable differentiable physics for learning and control. In *International Conference on Machine Learning*.
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. 2020. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning*.
- Connor Schenck and Dieter Fox. 2018. SPNets: Differentiable fluid dynamics for deep neural networks. *Conference on Robot Learning (CoRL’18)*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. arXiv:1707.06347. <https://arxiv.org/abs/1707.06347>.
- Eftychios Sifakis and Jernej Barbic. 2012. FEM simulation of 3D deformable solids: A practitioner’s guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH’12)*. Association for Computing Machinery, New York, NY, Article 20, 50 pages.
- Fun Shing Sin, Daniel Schroeder, and Jernej Barbi . 2013. Vega: Non-linear FEM deformable object simulator. *Computer Graphics Forum* 32, 1 (2013), 36–48.
- Andrew Spielberg, Allan Zhao, Yuanming Hu, Tao Du, Wojciech Matusik, and Daniela Rus. 2019. Learning-in-the-loop optimization: End-to-end control and co-design of soft robots through learned deep latent representations. In *Advances in Neural Information Processing Systems*. 8284–8294.
- Andrew Stuart and Tony Humphries. 1996. *Dynamical systems and numerical analysis*. Cambridge University.
- Maxime Thieffry, Alexandre Kruszewski, Christian Duriez, and Thierry-Marie Guerra. 2018. Control design for soft robots based on reduced-order model. *IEEE Robotics and Automation Letters* 4, 1 (2018), 25–32.

- Marc Toussaint, Kelsey Allen, Kevin Smith, and Joshua Tenenbaum. 2018. Differentiable physics and stable modes for tool-use and manipulation planning. In *Robotics: Science and Systems*, Vol. 2.
- Adrien Treuille, Antoine McNamara, Zoran Popović, and Jos Stam. 2003. Keyframe control of smoke simulations. *ACM Trans. Graph.* 22, 3 (July 2003), 716–723.
- Huamin Wang. 2015. A Chebyshev semi-iterative approach for accelerating projective and position-based dynamics. *ACM Trans. Graph.* 34, 6 (Oct. 2015), Article 246, 9 pages.
- Huamin Wang and Yin Yang. 2016. Descent methods for elastic body simulation on the GPU. *ACM Trans. Graph.* 35, 6 (Nov. 2016), Article 212, 10 pages.
- Chris Wojtan, Peter Mucha, and Greg Turk. 2006. Keyframe control of complex particle systems using the adjoint method. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA'06)*. Eurographics Association, Goslar, DEU, 15–23.
- Zangyueyang Xian, Xin Tong, and Tiantian Liu. 2019. A scalable Galerkin multigrid method for real-time simulation of deformable objects. *ACM Trans. Graph.* 38, 6 (Nov. 2019), Article 162, 13 pages.

Received October 2020; revised August 2021; accepted October 2021