

DiffCloth: Differentiable Cloth Simulation with Dry Frictional Contact

by

Yifei Li

B.S., Carnegie Mellon University (2020)

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 10, 2022

Certified by
Wojciech Matusik
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

DiffCloth: Differentiable Cloth Simulation with Dry Frictional Contact

by

Yifei Li

Submitted to the Department of Electrical Engineering and Computer Science
on May 10, 2022, in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering and Computer Science

Abstract

Cloth simulation has wide applications in computer animation, garment design, and robot-assisted dressing. This work presents a differentiable cloth simulator whose additional gradient information facilitates cloth-related applications. Our differentiable simulator extends a state-of-the-art cloth simulator based on Projective Dynamics (PD) and with dry frictional contact. We draw inspiration from previous work to propose a fast and novel method for deriving gradients in PD-based cloth simulation with dry frictional contact. Furthermore, we conduct a comprehensive analysis and evaluation of the usefulness of gradients in contact-rich cloth simulation. Finally, we demonstrate the efficacy of our simulator in a number of downstream applications, including system identification, trajectory optimization for assisted dressing, closed-loop control, inverse design, and real-to-sim transfer. We observe a substantial speedup obtained from using our gradient information in solving most of these applications.

Thesis Supervisor: Wojciech Matusik

Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

First of all, I would like to thank my advisor, Prof. Wojciech Matusik, for his insights, guidance, trust, and support during my Master’s thesis. Wojciech teaches me to step back and always keep the bigger picture in mind and constantly challenges me to reach for ambitious results while providing me with the resources I need for my research and the trust to explore different research directions.

I want to thank Tao Du, who mentored me extensively during this project and gave me constant guidance during my first two years of Ph.D. study. I thank Tao for sharing his knowledge of physical simulation, linear algebra, and numerical optimization, his feedback in writing and presentation, and his insights in research. Tao helped me develop the practical skills necessary for conducting graphics research and the critical thinking skills to solve complex problems and defend my research results.

I want to thank my collaborators, Kui Wu and Jie Xu, for their help, discussions, and support during this project and for helping and guiding me to put together the very first draft of my first SIGGRAPH project.

I also want to thank my undergraduate research advisor, Prof. Jessica Hodgins, who introduced me to the field of computer graphics and whose unconditional support ultimately led me to pursue a Ph.D. in graphics. Her words of encouragement still support me whenever I doubt myself.

Lastly, I would like to thank my friends and family for their accompany and emotional support. I want to thank my parents and sister for their love, trust, and support for every decision I make. I want to thank my best fluffy friend, Yihui, for coming into my life and always being by my side with meows and purrs. Finally, I would like to thank my best friend and partner, Tianyu Li, for his unconditional love and support in every possible aspect I could ask for.

Contents

1	Introduction	19
2	Related Works	23
2.1	Cloth Simulation	23
2.2	Cloth Contact and Friction	24
2.3	Inverse Dynamics	25
2.4	Differentiable Simulation	26
3	Background	29
3.1	Projective Dynamics Method	29
3.2	Dry Frictional Contact in Projective Dynamics	31
4	Method	35
4.1	Gradients without Contact	35
4.2	Gradients with Contact	37
5	Evaluations	43
5.1	Continuity and Smoothness	43
5.2	Usefulness of Gradients	47
5.3	Benefits of Dry Frictional Contact	48
5.4	Evaluation of Iterative Solver in Backpropagation	49
6	Applications	57
6.1	System Identification	58

6.2	Robot-Assisted Dressing	60
6.3	Inverse Design	63
6.4	A Real-to-Sim Example	64
6.5	Hat Controller	66
7	Conclusions, Limitations, Future Works	71
A	Appendix	73
A.1	Experiment Run Time	73
A.2	Experiment Details	74
A.2.1	System Identification	74
A.2.2	Robot-Assisted Dressing	74
A.2.3	Inverse Design	75
A.2.4	Real-to-Sim Example	76
A.2.5	Hat Controller	76
A.3	Optimization Results For All Random Seeds	77
	References	81

List of Figures

1-1	We present a differentiable cloth simulator with dry frictional contact and demonstrate its efficacy in multiple downstream applications, including designing this twirl dress (10902 degrees of freedom and 125 time steps with a step size of 1/120 seconds). The goal is to optimize the material parameters of the dress so that the apex angle of the cone it forms after a twirl reaches a desired value (100 degrees in this example). Left and middle: the optimized dress before and after a twirl. Top right and bottom right: motion sequences of the twirl dress before and after material parameter optimization using our differentiable simulator. The final apex angles before and after optimization are 39.06 and 100.30 degrees, respectively.	22
5-1	We simulate a piece of cloth sliding on a rigid sphere (right) to study the discontinuities and non-smoothness from contact conditions and surface geometry (Sec. 5.1). The two columns of plots show the velocities of three contact nodes at the 50-th and 200-th time steps. The “discretized” and “continuous” labels indicate whether these velocities are computed with a triangulated sphere (visualized as pink triangles on the spherical surface) or a smooth sphere (visualized as green sphere surface).	51

5-2	We simulate collisions between a piece of cloth and a rigid sphere at four different horizontal locations to study the effects of varying numbers of collisions on the smoothness of the simulation. Top: rendering of one of the time steps when the cloth is expected to be in contact with the sphere. The four scenes are ordered by their number of collisions. Bottom: landscapes of the loss defined as a function of the frictional coefficient μ and the stiffness k of the cloth (Sec. 5.1). The zoomed-in views of the landscapes show increasing discontinuity and non-smoothness as more collisions occur.	52
5-3	Flying Napkin. Comparisons between gradient-based and gradient-free methods for optimizing high-dimensional decision variables in differentiable cloth simulation. Top: the motion sequence of a piece of cloth with external forces parametrized with 300 variables. The goal is to find proper external forces with which the cloth ends at a target position. The loss function is defined as the position discrepancy between the simulated motion sequence and the reference motion. Bottom left: the loss vs. time step curves (300 DoF) from gradient-based (L-BFGS-B) and gradient-free (CMA-ES) methods. Both L-BFGS-B and CMA-ES use the same random seeds. Bottom right: we vary the degrees of freedom in the external force parametrization and repeat the experiment three times. For each experiment, we compute the ratio between the number of time steps used by gradient-free and gradient-based methods until they converge or use up the time budget and plot them as a speedup vs. DoFs curve.	53
5-4	Bowl. We simulate the motion sequence of the napkin ($50 \times 50 \times 2$ triangles, $h = 5\text{ms}$, 500 time steps) with the contact models from two differentiable cloth simulators: Liang et al. [2019] (top) and ours (middle). The area ratio (bottom) is the ratio between the napkin’s area at the current timestep and in the rest shape. Refer to our video for the full motions.	54

5-5	Wind and Slope. Motion sequences from the two benchmark tests for comparing iterative and direct solvers in back-propagation. Top: the “Wind” test ($h = 1/90$ s, 200 time steps) where a piece of cloth moves under synthetic wind. Bottom: the “Slope” test ($h = 1/100$ s, 300 time steps) where a ribbon slides along a slanted plane.	56
6-1	T-shirt (4278 DoFs, $h = 1/90$ s, 250 time steps). Estimating the cloth material and wind parameters based on a given synthetic motion sequence of the T-shirt. From left to right we show the simulated T-shirt at 0, 80, 160, and 240 time steps. Top: the ground-truth motion sequence; Middle: simulation with the initial guess on the cloth and wind parameters; Bottom: simulation after optimization with L-BFGS-B.	61
6-2	Sphere (1875 DoFs, $h = 1/180$ s, 350 time steps). Estimating the frictional coefficient μ between the sphere and the cloth based on the cloth’s contact with the sphere. From top to bottom we show the simulation at 0, 100, 200, and 300 time steps. Top: the ground-truth motion sequence; Middle: simulation with the initial guess on μ ; Bottom: simulation after optimization with L-BFGS-B.	62
6-3	Hat (1737 DoFs, $h = 1/100$ s, 400 time steps). Optimizing trajectories for a manipulator to move a hat onto the sphere. Top left: Initial trajectory from one random seed before optimization overlaid with intermediate hat positions in simulation. Top right: the optimized trajectory from L-BFGS-B (which shares a visually similar trajectory with the ones optimized by ES algorithms). Bottom: The loss vs. time step curves for all methods.	64

6-4 **Sock** (3165 DoFs, $h = 1/160$ s, 400 time steps). Optimizing trajectories for a manipulator to put a sock onto the foot model. Top left: One initial trajectory before optimization. We show the intermediate time steps on the left and the final state of the sock on the right. Top right: one control trajectory optimized by L-BFGS-B. The end effectors successfully put the sock onto the foot using the optimized trajectory. Bottom: the loss vs. time step curves for all methods. 65

6-5 **Flag**. A real-to-sim example that reconstructs a digital flapping flag (540 vertices, 1026 triangles, $h = 1/120$ s, 100 time steps) based on motion data captured from real-world experiments. Top: We plot the trajectories of 6 nodes on the cloth from the ground-truth motion (yellow curves) and the simulation results with guesses on the material and wind parameters before optimization (orange curves, left) and after optimization (orange curves, right). Bottom: The loss vs. time step curves for all methods. 67

6-6 **Hat Controller**. We train a closed-loop controller for the advanced “Hat” task (Sec. 6.5) which aims to move the hat from different initial positions (sampled from a fixed-radius hemisphere) onto the head. Top two rows: We visualize the control trajectories of the hat in ten initial positions denoted by their elevation and azimuth angle at the bottom of each subfigure. Bottom: The loss vs. time step curve for our gradient-based optimization and PPO. 69

6-7 We present the computation graph of the "Hat Controller" task in Sec. 6.5. We embed the neural network policy into our differentiable simulation pipeline. At time step t during forward simulation (blue arrows), the current particle states $(\mathbf{x}_t, \mathbf{v}_t)$ are transformed into a state vector $\mathbf{s}_t$, which is input to the neural network policy π_θ to produce an action vector $\mathbf{a}_t = \pi_\theta(\mathbf{s}_t)$. DiffCloth then simulates a new state $(\mathbf{x}_{t+1}, \mathbf{v}_{t+1})$ using current particle state and the generated action vector. The whole simulation sequence $\{x_i, v_i\}$ is used to compute a loss. During backward propagation (red arrows), the above computation pass is reversed. 70

List of Tables

5.1	Comparison between the iterative solver (ours) and the sparse LU direct solver in backpropagation under various mesh resolutions in the “Wind” and “Slope” tests. The numbers in the “Res.” column report the mesh resolution. The number in parentheses in the “Solver” column indicates the epsilon value controlling the convergence of the Jacobi solver. The “Backprop” column reports the net time in backpropagation and is further decomposed into the next five columns, from which the sum of ratios is 100%: “ $\Delta\mathbf{P}$ ” and “ $\Delta\mathbf{R}$ ” report the time spent on assembling $\Delta\mathbf{P}$ and $\Delta\mathbf{R}$ respectively, “Iter.” and “Direct” shows the time cost by either solver, and operations not covered by these four columns are in the “Other” column. The “Fail” column reports the ratio between the number of timesteps seeing nonconvergence in our iterative solver and the number of total timesteps. The “Speedup” column is the ratio between the direct solver’s time and our time in “Backprop. Time” in each test.	55
-----	---	----

6.1	Comparison between the performance of gradient-based and gradient-free optimization on all examples (except “Hat Controller”) in Sec. 6. The “Param #” and “Seed #” columns report the number of optimization parameters and random seeds used in the experiments, respectively. All methods start the optimization with the same initial random seed set. The “Min Init. Loss” column reports the minimum initial loss across all random seeds. The “Optimized Loss Percentage” column reports the optimized loss <i>as a percentage</i> (0-100%) of the minimum initial loss reached across all random seeds for each method: “Final” reports the loss reached when the optimization stops, and “Equal Step #” reports the loss reached by CMA-ES and (1+1)-ES using the same number of time steps for L-BFGS-B convergence. We color the values of loss percentage using a green-orange color scale: green corresponds to 0% and orange corresponds to 100%.	59
-----	---	----

6.2	The “Convergence Time Steps” column reports the number of simulation time steps used for each optimization method until convergence on all examples shown in Sec. 6. For a fair comparison, the time steps shown for L-BFGS-B include both forward simulation and backward propagation time. The “Iter. Conv. %” column reports the percentage of time steps in backward propagation where the iterative solver converges.	60
-----	--	----

A.1	Run time for our gradient-based optimization. We report the average wall-clock time for optimizing the examples shown in Sec. 6. We recall the simulation complexity of each example by reiterating its total Degrees of Freedom (“Dof”), number of time steps in each forward simulation sequence and back-propagation (“Time Steps”) and time interval (“ $h[s]$ ”). “Fwd.[s]” and “Back.[s]” report the mean wall-clock time for performing one iteration of forward simulation and back-propagation respectively, averaged across all optimization iterations for all initial seeds.	73
A.2	We report the optimization results for all random seeds of all methods in the “T-shirt” and “Flag” examples. For each random seed, we report its initial loss and final loss achieved by each optimization method. “Final Loss Percentage (%)” reports the optimized loss as a <i>percentage</i> ($0-100\%$) of the initial loss. “Convergence Time Steps” reports the number of time steps used by each method to reach its final loss. For all tasks, we also summarize the minimum, average and median statistics across all random seeds for each column. For each metric (“Final Loss”, “Final Loss Percentage (%)”, “Convergence Time Steps”) and each row, the minimum number across the three optimization methods is marked in bold.	78
A.3	Similar table as Table A.2 for the “Sphere” and “Dress” examples. . .	79
A.4	Similar table as Table A.2 for the “Hat”, “Sock” examples.	80

Chapter 1

Introduction

Clothing is ubiquitous in our daily lives. With the widespread appearance of clothing in the fashion industry, film industry, computer animation, and video games, simulating cloth has been an active research topic for more than two decades. Today, research advancement in cloth simulation has unlocked various applications such as virtual try-on [Guan et al. 2012], garment design [Bartle et al. 2016; Montes et al. 2020; Umetani et al. 2011; Wang 2018], fold design [Li et al. 2018c], garment grading [Brouet et al. 2012], sagging-free inversion [Ly et al. 2018], and robot-assisted dressing [Clegg et al. 2020, 2018].

Inspired by the recent development of differentiable physics simulation and its success in rigid-body systems [de Avila Belbute-Peres et al. 2018a; Geilinger et al. 2020], fluidic systems [Du et al. 2020; Hu et al. 2020], and deformable-body systems [Geilinger et al. 2020; Hahn et al. 2019], we argue in this paper that a large number of cloth-related applications would also benefit from a high-quality differentiable cloth simulator. The critical ingredient in previous differentiable simulators is their ability to compute gradients by backpropagating any differentiable performance metrics through simulation. Such additional gradient information unlocks gradient-based continuous optimization methods, which often bring a substantial speedup compared with their gradient-free counterparts in downstream applications.

Compared with the recent active development of differentiable simulators for rigid-body and soft-body dynamics, research work about differentiable cloth simulation is

still relatively sparse. Indeed, cloth simulation introduces unique challenges from its co-dimensional dynamics and, in particular, rich contact events. While many differentiable simulators have provided solutions to derive gradients for contact models of varying complexity, their techniques typically do not expect contact to be as frequent as in cloth simulation. Deriving gradients with frequent contact and self-collisions in cloth simulation is not fully resolved even in the state-of-the-art differentiable cloth simulators [Liang et al. 2019], and our work attempts to fill this gap.

In this work, we present a differentiable cloth simulator with extra care of its contact model. We base our method on the state-of-the-art cloth simulator from Ly et al. [2020] and employ its Projective Dynamics (PD) simulation method and dry frictional contact described by the Signorini-Coulomb law. Therefore, our differentiable cloth simulator inherits both the speedup from Projective Dynamics and the physical accuracy from the dry frictional contact model. Unlike previous papers relying on automatic differentiation tools to derive gradients, we present an iterative solver modified from Du et al. [2021] to accommodate the dry frictional contact model. We show that the modified iterative solver leads to a substantial speedup over a standard linear solver in gradient computation.

To have well-defined gradients, a differentiable simulator expects all quantities computed in simulation to be sufficiently smooth. However, the non-smooth position and force changes from large numbers of contact events question this fundamental assumption. To fully understand the usefulness of a differentiable simulator in contact-rich environments, our work conducts comprehensive evaluation and analysis of the behavior of our differentiable cloth simulation with a varying number of contact events. While previous papers have provided similar discussions, they primarily focus on understanding penalty-based contact models [Geilinger et al. 2020] or discussing collisions in rigid-body dynamics [Werling et al. 2021] where contact events are not as frequent as in cloth simulation. To our best knowledge, our work is the first to present such evaluation and discussion of the usefulness of gradients in a differentiable cloth simulator with dry frictional contact.

We demonstrate the efficacy of our simulator in various applications, including

system identification of frictional coefficients in cloth simulation, inverse garment design for computer animation, and motion planning of robotic manipulators in robot-assisted dressing. Many of these applications would either be impossible or require a much longer time to solve with existing methods. With the extra gradient information from our differentiable simulator, we unlock gradient-based optimizers to solve these problems with a much higher sample efficiency than traditional gradient-free methods.

To summarize, our work contributes the following (Figure 1-1):

- We present a novel differentiable cloth simulator with dry frictional contact and an iterative solver for speeding up its gradient computation.
- We evaluate the source of non-differentiability in the dry frictional contact model and discuss the usefulness of gradients in differentiable, contact-rich cloth simulation.
- We show the efficacy of our simulator in various applications, including system identification, trajectory optimization for assisted dressing, closed-loop control, inverse design, and real-to-sim transfer.




Figure 1-1: We present a differentiable cloth simulator with dry frictional contact and demonstrate its efficacy in multiple downstream applications, including designing this twirl dress (10902 degrees of freedom and 125 time steps with a step size of $1/120$ seconds). The goal is to optimize the material parameters of the dress so that the apex angle of the cone it forms after a twirl reaches a desired value (100 degrees in this example). Left and middle: the optimized dress before and after a twirl. Top right and bottom right: motion sequences of the twirl dress before and after material parameter optimization using our differentiable simulator. The final apex angles before and after optimization are 39.06 and 100.30 degrees, respectively.

Chapter 2

Related Works

Our work is closely related to cloth simulation and its applications in computer graphics. It is also relevant to the more recent differentiable simulation methods developed in the graphics and machine learning communities.

2.1 Cloth Simulation

Physics-based cloth simulation has been a popular topic in the graphics community for decades since clothing has been widely used in our daily lives. The implicit Euler integration is used to simulate cloth robustly with large time steps [Baraff and Witkin 1998; Terzopoulos et al. 1987] while introducing excessive numerical damping. To alleviate this problem, implicit and explicit methods (IMEX) [Bridson et al. 2003; Stern and Grinspun 2009] explicitly integrate elastic forces and implicitly integrate damping forces. Researchers have also introduced other variational integrators, e.g., BDF2 [Choi and Ko 2002] and Symplectic integrator [Stern and Desbrun 2006], to conserve the total energy of the system.

For cloth modeled as a mass-spring system, Liu et al. [2013] treat the implicit Euler integration as an energy minimization problem. The global linear system then remains constant on run-time, and each spring constraint can be solved separately in a local step. That global/local solver idea is generalized as Projective Dynamics (PD) [Bouaziz et al. 2014], which supports material models whose elastic energy

has a specific quadratic form. Projective Dynamic is further extended to support more general materials [Liu et al. 2017; Overby et al. 2017]. Though a local step in PD can be processed in parallel, the global step still needs to maintain a large pre-factorized sparse matrix and do back substitutions in each step. Another relevant idea is Position-Based Dynamics (PBD) [Macklin et al. 2016; Müller et al. 2007], which iteratively projects each constraint in a non-linear Gauss-Seidel-like fashion, leading to highly parallelizable computation on GPUs.

Both PD and PBD have led to a few follow-up works on more advanced speedup techniques. Komaritzan and Botsch [2018] present techniques to speed up PD with contact for physics-based character skinning. For PD and PBD in cloth simulation, Wang [2015] and a follow-up work [Wang and Yang 2016] propose a Chebyshev acceleration technique that can be applied to both PD and PBD to speed up convergence. Fratarcangeli et al. [2016] also introduce a parallel randomized Gauss-Seidel method that re-organizes the unknowns of the sparse linear system into a few independent blocks, which can be solved in parallel with a single Gauss-Seidel step. Recently, the computation of integration is further accelerated by the multigrid method, e.g., geometric multigrid scheme [Wang et al. 2018], algebraic multigrid scheme [Tamstorf et al. 2015], and Galerkin multigrid scheme [Xian et al. 2019].

Lastly, our work is also relevant to previous attempts to matching cloth simulation with real fabric behaviors [Clyde et al. 2017; Miguel et al. 2012, 2013; Wang et al. 2011], which is typically done by fitting constitutive material models with real material properties and can benefit from extra gradient information from a differentiable simulator [Hahn et al. 2019].

2.2 Cloth Contact and Friction

Contact and friction are key ingredients in modern cloth simulation. Provot [1997] proposes a method called “impact zones” to collect the nodes involved in multiple collisions into impact zones, which are treated as rigid bodies. Harmon et al. [2008] improve the failsafe of impact zones by allowing some sliding motion of the incrimi-

nated vertices. Bridson et al. [2002] introduce a hybrid method to combine the idea of applying repulsion impulses, frictions, and impact zones to handle cloth collision robustly, which is still widely used in modern cloth simulators [Li et al. 2020b; Narain et al. 2012]. To model friction fully implicitly, Brown et al. [2018] treat friction as an additional dissipative term in optimization. Regarding contact handling, repulsive forces or penalty methods have been widely used [Bouaziz et al. 2014; Geilinger et al. 2020; Macklin et al. 2020; Wang 2015] since they are easy to implement. However, these methods often need high stiffness in the penalty energy, leading to disturbing jittering artifacts that demand careful tuning. Recently, Li et al. [2020a, 2021] define a smooth dissipative potential for friction using barrier methods. Their method can guarantee interpenetration-free states after each time step without the need for parameter tuning.

Unlike penalty-based methods, constraint-based collision handling methods formulate contact as constraints in the physics system. Otaduy et al. [2009] first formulate cloth contacts as a sparse linear complementarity problem (LCP). Their key idea is to interleave frictional contact iterations with normal contact iterations. Instead of using a pyramidal approximation of the Coulomb friction cone [Otaduy et al. 2009], Li et al. [2018b] rely upon the exact Coulomb friction cone and adaptively refine nodes to ensure an accurate treatment of frictional contact. Although constraint-based methods typically ensure the physics-based constraints characterizing contact and friction are satisfied after time integration, their computation cost is typically much more expensive than a penalty-based method. Recently, Ly et al. [2020] propose an efficient algorithm to incorporate frictional contact into Projective Dynamics so that non-penetration and Coulomb constraints are satisfied simultaneously in a semi-implicit way.

2.3 Inverse Dynamics

Inverse dynamics have been studied in robotics for decades to reconstruct internal forces or torques from the observations of robotic systems. However, existing methods

usually focus on rigid-body systems only, which have less than a hundred degrees of freedom (DoFs) [Dario Bellicoso et al. 2016; Kang et al. 2021; Mistry et al. 2010]. Inverse dynamics for high-DoF systems like soft bodies, fluids, and cloth are still under exploration due to the lack of high-quality numerical solutions in robotics for both simulation and differentiation. One noticeable distinction between inverse dynamics and differentiable simulation is that differentiable simulation computes additional gradients for initial states, system parameters, and design parameters. Therefore, differentiable simulation enables more applications like system identification, inverse design, and real-to-sim matching that traditional inverse dynamics typically do not consider.

2.4 Differentiable Simulation

Differentiable simulation is a relatively recent concept explored in the graphics and machine learning community, but its original idea can be traced back to much earlier works in graphics decades ago. Perhaps one of the earliest such papers is Witkin and Kass [1988], which shows optimizing simulation to minimize an objective. Despite the recent advances in differentiable simulators in rigid-body dynamics [de Avila Belbute-Peres et al. 2018b; Degraeve et al. 2019; Popović et al. 2003; Qiao et al. 2020; Toussaint et al. 2018; Xu et al. 2021], soft-body dynamics [Du et al. 2021; Geilinger et al. 2020; Hahn et al. 2019; Hu et al. 2020, 2019], and fluid dynamics [Holl et al. 2020; McNamara et al. 2004; Schenck and Fox 2018; Treuille et al. 2003; Wojtan et al. 2006], differentiable cloth simulation still lacks a good solution. One natural idea is to use particle-based strategies that approximate a nodal system of cloth with graph neural networks [Li et al. 2019; Sanchez-Gonzalez et al. 2020]. Although the neural networks are naturally differentiable, the physical accuracy is hard to guarantee.

To accurately predict the behavior of real-world objects, recent papers [Geilinger et al. 2020; Hahn et al. 2019] present differentiable soft-body systems and mass-spring systems with implicit Euler time integration and penalty-based contacts. Although a few recent papers [Li et al. 2020a, 2021; Macklin et al. 2020] have introduced dif-

ferentiable contact handling methods, none of them show a clothing example using the differentiability besides demonstrating the differentiability of their methods in theory. Liang et al. [2019] are the first to introduce a fully functional differentiable cloth simulation with contact, friction, and self-collision. Instead of constructing a static LCP problem, they develop a quadratic programming (QP) problem to minimize the change between the collision-free state and the original mesh state. Murthy et al. [2021] also present a differentiable cloth simulator with penalty-based frictional contact in their fully differentiable simulation and rendering pipeline. In this paper, we build upon the differentiable PD framework [Du et al. 2021] to simulate cloth dynamics with dry frictional contact [Ly et al. 2020] and augment it with gradient computation.

A significant challenge in differentiable simulation is the presence of non-smooth contact events, where non-smoothness can arise from discontinuous contact shapes [Popović et al. 2000], impulsive forces [Hu et al. 2019], or branches in contact laws [Ly et al. 2020], to name a few. Many existing differentiable simulators assume such non-smoothness does not affect the usability of gradients. Indeed, they present empirical results that observe the benefits of using gradients in downstream applications [de Avila Belbute-Peres et al. 2018a; Du et al. 2021; Geilinger et al. 2020; Liang et al. 2019]. While our work also observes the advantages of gradients over gradient-free methods in several applications, we take one step further by analyzing the source of non-smoothness and discontinuities in contact-rich cloth simulation. We hope that our experience can serve as useful heuristics for applying differentiable simulation methods in the future.

Chapter 3

Background

In this section, we briefly review the Projective Dynamics cloth simulation method with dry frictional contact described in Ly et al. [2020], on which our differentiable cloth simulator is based. Our core contribution lies in the development of gradient computation in this PD framework with dry frictional contact, which we detail in the next section.

3.1 Projective Dynamics Method

Implicit time integration We model cloth as a nodal system with m 3D nodes. Let $\mathbf{x}(t)$ and $\mathbf{v}(t)$ be two vector functions from $\mathbf{R}^+$ to $\mathbf{R}^{3m}$ indicating the positions and velocities of all nodes at time t . We consider the standard implicit time-stepping scheme discretizing $\mathbf{x}(t)$ and $\mathbf{v}(t)$ as follows:

$$\mathbf{x}_{n+1} = \mathbf{x}_n + h\mathbf{v}_{n+1}, \quad (3.1)$$

$$\mathbf{v}_{n+1} = \mathbf{v}_n + h\mathbf{M}^{-1}[\mathbf{f}_{\text{int}}(\mathbf{x}_{n+1}) + \mathbf{f}_{\text{ext}}], \quad (3.2)$$

where h is the time step size, $\mathbf{M}$ a positive diagonal mass matrix of size $3m \times 3m$, and $\mathbf{f}_{\text{int}}$ and $\mathbf{f}_{\text{ext}}$ the internal and external forces at each node stacked as two $3m$ -dimensional vectors, respectively. Note that we assume for now $\mathbf{f}_{\text{int}}$ and $\mathbf{f}_{\text{ext}}$ do not contain contact forces, which will be discussed separately in Sec. 3.2. We use the

cloth material model described in Bouaziz et al. [2014] to define $\mathbf{f}_{\text{int}}$. The implicit time-stepping scheme connects the state of the nodal system $(\mathbf{x}_n, \mathbf{v}_n)$ at time t_n to $(\mathbf{x}_{n+1}, \mathbf{v}_{n+1})$ at new time $t_{n+1} = t_n + h$.

Optimization view As discussed by Martin et al. [2011], the implicit time-stepping scheme can be rephrased as finding a saddle point of the following energy minimization problem:

$$\min_{\mathbf{x}_{n+1}} \underbrace{\frac{1}{2h^2}(\mathbf{x}_{n+1} - \mathbf{y})^\top \mathbf{M}(\mathbf{x}_{n+1} - \mathbf{y}) + W(\mathbf{x}_{n+1})}_{g(\mathbf{x}_{n+1})}, \quad (3.3)$$

where $\mathbf{y} = \mathbf{x}_n + h\mathbf{v}_n + h^2\mathbf{M}^{-1}\mathbf{f}_{\text{ext}}$ is a constant vector that can be precomputed at the beginning of the current time step. The potential energy W defines the internal force $\mathbf{f}_{\text{int}}$ by its spatial gradients: $\mathbf{f}_{\text{int}} = -\nabla W(\mathbf{x}_{n+1})$. It is easy to verify that setting the gradient of the objective function to zero leads to a system of equations identical to Eqns. (3.1) and (3.2).

Local and global solvers in PD The key assumptions in PD is that the internal energy W can be written as a sum of quadratic forms [Bouaziz et al. 2014; Liu et al. 2017]:

$$W_i(\mathbf{x}) = \min_{\mathbf{p}_i \in \mathcal{M}_i} \underbrace{\frac{w_i}{2} \|\mathbf{A}_i \mathbf{x} - \mathbf{p}_i\|_2^2}_{\widetilde{W}_i(\mathbf{x}, \mathbf{p}_i)}, \quad (3.4)$$

$$W(\mathbf{x}) = \sum_i W_i(\mathbf{x}). \quad (3.5)$$

Here, each energy W_i projects $\mathbf{A}_i \mathbf{x}$, a linear transformation of $\mathbf{x}$, to its closest point in the set $\mathcal{M}_i$ and scales its squared distance by a prespecified stiffness w_i . Both $\mathbf{A}_i$ and $\mathcal{M}_i$ are predefined and independent of $\mathbf{x}$. Here, $\mathcal{M}_i$ is the constraint manifold, and $\mathbf{p}_i$ is the auxiliary projection variable as defined in Bouaziz et al. [2014].

With such an assumption on W , PD proposes a local-global solver to minimize a surrogate objective function:

$$\widetilde{g}(\mathbf{x}_{n+1}, \mathbf{p}) = \frac{1}{2h^2}(\mathbf{x}_{n+1} - \mathbf{y})^\top \mathbf{M}(\mathbf{x}_{n+1} - \mathbf{y}) + \sum_i \widetilde{W}_i(\mathbf{x}_{n+1}, \mathbf{p}_i), \quad (3.6)$$

where $\mathbf{p}$ stacks up all $\mathbf{p}_i$ from each W_i . In the local step, PD fixes $\mathbf{x}_{n+1}$ and projects each $\mathbf{p}_i$ to its corresponding $\mathcal{M}_i$ by solving Eqn. (3.4). Such a local step can be done in parallel for each $\widetilde{W}_i$. In the global step, PD fixes $\mathbf{p}$ and minimizes $\widetilde{g}$ as a function of $\mathbf{x}_{n+1}$, which turns out to have a closed-form solution:

$$\underbrace{(\mathbf{M} + h^2 \sum_i w_i \mathbf{A}_i^\top \mathbf{A}_i)}_{\mathbf{P}} \mathbf{x}_{n+1} = \underbrace{\mathbf{M} \mathbf{y} + h^2 \sum_i w_i \mathbf{A}_i^\top \mathbf{p}_i}_{\mathbf{b}(\mathbf{p})}. \quad (3.7)$$

By alternating between the local and global steps, PD monotonically decreases the surrogate energy $\widetilde{g}$ until convergence, which can be shown to agree with the saddle point of g [Liu et al. 2017]. The source of efficiency in PD comes from the observation that $\mathbf{P}$ is a constant matrix that can be prefactorized at the beginning of the simulation, leading to an efficient global step requiring back-substitution only.

3.2 Dry Frictional Contact in Projective Dynamics

Signorini-Coulomb law Ly et al. [2020] augment the standard PD framework described above with non-penetration collision and Coulomb friction by assuming contact applies to nodes only. At each time step, assuming a collision detection algorithm has identified a set $\mathcal{I} \subseteq \{1, 2, 3, \dots, m\}$ which describes the indices of contact nodes. For each node $j \in \mathcal{I}$, the Signorini-Coulomb law [Brogliato 2016] requires its local force $\mathbf{r}_j$ and velocity $\mathbf{u}_j$ to satisfy one of the following three conditions:

$$\left\{ \begin{array}{l} \text{Take off: } \mathbf{r}_j = \mathbf{0}, \mathbf{u}_{j|N} > 0, \end{array} \right. \quad (3.8a)$$

$$\left\{ \begin{array}{l} \text{Stick: } \|\mathbf{r}_{j|T}\| < \mu \mathbf{r}_{j|N}, \mathbf{u}_j = \mathbf{0}, \end{array} \right. \quad (3.8b)$$

$$\left\{ \begin{array}{l} \text{Slip: } \|\mathbf{r}_{j|T}\| = \mu \mathbf{r}_{j|N}, \mathbf{u}_{j|N} = 0, \mathbf{r}_{j|T} \parallel \mathbf{u}_{j|T}, \mathbf{r}_{j|T} \cdot \mathbf{u}_{j|T} \leq 0. \end{array} \right. \quad (3.8c)$$

Here, μ is the frictional coefficient, and $\mathbf{u}_j$ and $\mathbf{r}_j$ are the nodal velocity and contact force represented in the local contact frame spanned by the tangential plane and contact normal on the contact surface. The notations $j|T$ and $j|N$ represent the

tangential and normal components of a 3D vector defined in the local frame. We further define $\mathcal{C}^j \subseteq \mathbf{R}^3 \times \mathbf{R}^3$ as the set of valid $(\mathbf{r}_j, \mathbf{u}_j)$ pairs described by the conditions above, allowing us to rewrite Eqn. 3.8 compactly: $(\mathbf{r}_j, \mathbf{u}_j) \in \mathcal{C}^j$.

Implicit time integration with dry frictional contact The original implicit time integration now needs to be augmented with additional constraints describing contact conditions [Ly et al. 2020]:

$$\begin{cases} \mathbf{x}_{n+1} = \mathbf{x}_n + h\mathbf{v}_{n+1}, \\ \mathbf{v}_{n+1} = \mathbf{v}_n + h\mathbf{M}^{-1}[\mathbf{f}_{\text{int}}(\mathbf{x}_{n+1}) + \mathbf{f}_{\text{ext}} + \mathbf{J}_n^\top(\mathbf{x}_n, \mathbf{v}_n)\mathbf{r}_{n+1}], \\ \mathbf{u}_{n+1} = \mathbf{J}_n(\mathbf{x}_n, \mathbf{v}_n)\mathbf{v}_{n+1}, \\ (\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j}) \in \mathcal{C}_n^j, \forall j \in \mathcal{I}_n. \end{cases} \quad (3.9)$$

Here, we have added the subscript n in the definitions of the contact set $\mathcal{I}$ and the contact condition $\mathcal{C}^j$ described above to specify the time step they are defined from. The notation $\mathbf{r}_{n+1,j}$ and $\mathbf{u}_{n+1,j}$ select the 3D force and velocity corresponding to the j -th node from $\mathbf{r}_{n+1}$ and $\mathbf{u}_{n+1}$, respectively. The Jacobian matrix $\mathbf{J}_n$ of size $3|\mathcal{I}_n| \times 3m$ selects global vectors defined on the contact nodes and computes their coordinates in the local contact frames. Note that our definition of $\mathbf{J}_n$ implies that it is computed in an explicit manner, i.e., it relies on the last state $(\mathbf{x}_n, \mathbf{v}_n)$ entering the n -th time step. For brevity, this background section will describe simple contact events between a node and a static plane, in which case $\mathbf{J}_n$ is a constant matrix that does not need to be computed from $(\mathbf{x}_n, \mathbf{v}_n)$. Extensions to more sophisticated contact events, e.g., self-collisions between multiple contact nodes or contact events between a node and a non-planar obstacle, can be found in Ly et al. [2020] and in our implementation and experiments.

Projective Dynamics with dry frictional contact The core idea in Ly et al. [2020] is an additional local step in PD that solves each contact node independently. Noting that the contact conditions are primarily defined on nodal velocities instead of nodal positions, Ly et al. [2020] first rewrite the global step in Eqn. (3.7) as an

equation of velocities:

$$\mathbf{P}\mathbf{v}_{n+1} = \frac{1}{h} \underbrace{[\mathbf{b}(\mathbf{p}) - \mathbf{P}\mathbf{x}_n]}_{\widehat{\mathbf{b}}(\mathbf{p})}. \quad (3.10)$$

The velocity-based PD global-local steps then alternate between this new global solver and the original local step, which is unaffected by this change of variables. By comparing Eqn. (3.9) and Eqn. (3.10), we can see that incorporating the contact force adds an additional impulse $h\mathbf{J}_n^\top \mathbf{r}_{n+1}$ to the right-hand side of the global step:

$$\mathbf{P}\mathbf{v}_{n+1} = \widehat{\mathbf{b}}(\mathbf{p}) + h\mathbf{J}_n^\top \mathbf{r}_{n+1}. \quad (3.11)$$

The goal of the global solver is now to find $(\mathbf{v}_{n+1}, \mathbf{r}_{n+1})$ that satisfy contact conditions at each j . Noting that in Eqn. (3.11), $\mathbf{P}$ is the only operator that couples unknown contact forces from different j , Ly et al. [2020] propose the following iterative solver based on the decomposition of $\mathbf{P}$ to solve Eqn. (3.11):

$$\mathbf{M}\mathbf{v}_{n+1}^{k+1} = \widehat{\mathbf{b}}(\mathbf{p}) - \underbrace{(h^2 \sum_i w_i \mathbf{A}_i^\top \mathbf{A}_i)}_{\mathbf{P}-\mathbf{M}} \mathbf{v}_{n+1}^k + h\mathbf{J}_n^\top \mathbf{r}_{n+1}^{k+1}, \quad (3.12)$$

where the superscripts k and $k+1$ indicate two consecutive iterations in the iterative solver at fixed time step $n+1$. After iteration k , $\mathbf{r}_{n+1}^{k+1}$ is updated using $\mathbf{v}_{n+1}^k$ to enforce the Signorini-Coulomb law, which is then used to update $\mathbf{v}_{n+1}^{k+1}$ according to Eqn. (3.12). It is easy to see that the proposed iterative solver fully decouples the momentum equation on each node, allowing Ly et al. [2020] to enforce the Signorini-Coulomb law by adjusting $(\mathbf{v}_{n+1,j}, \mathbf{r}_{n+1,j})$ on each contact node j independently in a straightforward manner. When the iterative solver converges, Ly et al. [2020] prove that the result is indeed a solution to Eqn. (3.9). We refer the readers to the original paper for the details.

Chapter 4

Method

We now describe in detail how we extend the forward simulator in Sec. 3 to build a differentiable cloth simulator. We start by differentiating implicit time integration in PD without contact, followed by explaining how contact gradients can be added to this framework. Compared with other differentiable simulators, our simulator is unique because of its treatment of the contact-rich nature of cloth simulation: many existing differentiable simulators focus on sparse contact events in rigid-body or deformable-body dynamics [Du et al. 2021; Geilinger et al. 2020; Hu et al. 2019], and the state-of-the-art differentiable cloth simulation method [Liang et al. 2019] handles rich contact events but without physics-based contact forces or friction. To our best knowledge, our work is the first to present a differentiable cloth simulator that can handle rich contact events with Coulomb’s law of friction.

4.1 Gradients without Contact

Differentiating implicit time integration The core step in building a differentiable simulator based on implicit time-stepping scheme is to backpropagate gradients through the implicit integration described in Eqns. (3.1) and (3.2), or equivalently, to derive the Jacobian of the output $(\mathbf{x}_{n+1}, \mathbf{v}_{n+1})$ with respect to the input $(\mathbf{x}_n, \mathbf{v}_n)$. Mature techniques such as sensitivity analysis, adjoint method, and implicit function theorem have proven to be successful in computing such gradients [Du et al. 2021;

Geilinger et al. 2020; Hahn et al. 2019]. To sketch the idea, we plug $\mathbf{v}_{n+1}$ from Eqn. (3.2) into Eqn. (3.1):

$$\begin{aligned}\mathbf{x}_{n+1} &= \mathbf{x}_n + h\mathbf{v}_n + h^2\mathbf{M}^{-1}[\mathbf{f}_{\text{int}}(\mathbf{x}_{n+1}) + \mathbf{f}_{\text{ext}}] \\ &= \mathbf{y} + h^2\mathbf{M}^{-1}\mathbf{f}_{\text{int}}(\mathbf{x}_{n+1}),\end{aligned}\tag{4.1}$$

which is essentially the first-order optimality condition in Eqn. (3.3). Differentiating $\mathbf{x}_n$ on both sides gives

$$\begin{aligned}\frac{\partial \mathbf{x}_{n+1}}{\partial \mathbf{x}_n} &= \mathbf{I} + h^2\mathbf{M}^{-1}\frac{\partial \mathbf{f}_{\text{int}}(\mathbf{x}_{n+1})}{\partial \mathbf{x}_{n+1}}\frac{\partial \mathbf{x}_{n+1}}{\partial \mathbf{x}_n} \\ &= \mathbf{I} - h^2\mathbf{M}^{-1}\nabla^2 W(\mathbf{x}_{n+1})\frac{\partial \mathbf{x}_{n+1}}{\partial \mathbf{x}_n},\end{aligned}\tag{4.2}$$

or equivalently,

$$\begin{aligned}\frac{\partial \mathbf{x}_{n+1}}{\partial \mathbf{x}_n} &= [\mathbf{I} + h^2\mathbf{M}^{-1}\nabla^2 W(\mathbf{x}_{n+1})]^{-1} \\ &= [\mathbf{M} + h^2\nabla^2 W(\mathbf{x}_{n+1})]^{-1}\mathbf{M}.\end{aligned}\tag{4.3}$$

In backpropagation, such a Jacobian matrix is coupled with the gradients of a loss function L which are passed to the previous state (assuming the gradients below are both column vectors):

$$\begin{aligned}\frac{\partial L}{\partial \mathbf{x}_n} &= \left(\frac{\partial \mathbf{x}_{n+1}}{\partial \mathbf{x}_n}\right)^\top \frac{\partial L}{\partial \mathbf{x}_{n+1}} \\ &= \underbrace{\mathbf{M} [\mathbf{M} + h^2\nabla^2 W(\mathbf{x}_{n+1})]^{-1}}_{\mathbf{z}_{n+1}} \frac{\partial L}{\partial \mathbf{x}_{n+1}}.\end{aligned}\tag{4.4}$$

Here, we obtain the adjoint vector $\mathbf{z}_{n+1}$ by solving the linear system of equations with the matrix $\mathbf{M} + h^2\nabla^2 W(\mathbf{x}_{n+1})$ on the left-hand side, which avoids an explicit inversion of the large and sparse matrix. Backpropagating gradients of L with respect to $\mathbf{v}_n$ and $\mathbf{v}_{n+1}$ can be derived similarly. In fact, it solves the same linear system but with a different vector $\frac{\partial L}{\partial \mathbf{v}_{n+1}}$ on the right-hand side.

Differentiating with Projective Dynamics With assumptions on the energy form W in PD, Du et al. [2021] show that we can speed up the computation in $\mathbf{z}_{n+1}$ by exploiting the special form of $\nabla^2 W$. The first- and second-order derivatives in Eqn. (3.5) are given by the following equations [Du et al. 2021; Liu et al. 2017]:

$$\nabla W_i(\mathbf{x}) = w_i \mathbf{A}_i^\top [\mathbf{A}_i \mathbf{x} - \mathbf{p}_i^*(\mathbf{x})], \quad (4.5)$$

$$\nabla^2 W_i(\mathbf{x}) = w_i \mathbf{A}_i^\top \mathbf{A}_i - w_i \mathbf{A}_i^\top \frac{\partial \mathbf{p}_i^*}{\partial \mathbf{x}}, \quad (4.6)$$

where $\mathbf{p}_i^*(\mathbf{x}) = \arg \min_{\mathbf{p}_i \in \mathcal{M}_i} \widetilde{W}_i(\mathbf{x}, \mathbf{p}_i)$ is the projection of $\mathbf{A}_i \mathbf{x}$ onto $\mathcal{M}_i$ obtained in the local step. Interested readers can refer to the appendix of Liu et al. [2017] for their derivation details. We plug them to $\mathbf{z}_{n+1}$ and rewrite the linear system as follows:

$$\underbrace{(\mathbf{P} - h^2 \sum_i w_i \mathbf{A}_i^\top \frac{\partial \mathbf{p}_i^*}{\partial \mathbf{x}} \Big|_{\mathbf{x}_{n+1}})}_{\Delta \mathbf{P}} \mathbf{z}_{n+1} = \frac{\partial L}{\partial \mathbf{x}_{n+1}}. \quad (4.7)$$

Noting that a prefactorization of $\mathbf{P}$ exists in Projective Dynamics, Du et al. [2021] propose the following iterations to solve $\mathbf{z}_{n+1}$:

$$\mathbf{P} \mathbf{z}_{n+1}^{k+1} = \Delta \mathbf{P} \mathbf{z}_{n+1}^k + \frac{\partial L}{\partial \mathbf{x}_{n+1}}. \quad (4.8)$$

Similar to the local-global steps in Projective Dynamics, a local step can evaluate $\Delta \mathbf{P}$ across each $\nabla^2 W_i$ in parallel, and a global step, which reuses $\mathbf{P}$, can solve $\mathbf{z}_{n+1}^{k+1}$ with backsubstitution only. Du et al. [2021] show that such a local-global solver is empirically faster than directly solving Eqn. (4.4), which resembles the source of efficiency in the original PD method.

4.2 Gradients with Contact

While Du et al. [2021] have discussed extensively how to fuse backpropagation into the Projective Dynamics framework, its support of contact is limited to non-penetration

conditions without a proper treatment on friction. On the other hand, other prior papers have provided solutions to deriving gradients with contact governed by Signorini-Coulomb law, but they focus on either small-scale problems (e.g., rigid bodies in de Avila Belbute-Peres et al. [2018a]) or sparse contact events on a static ground only [Geilinger et al. 2020]. Here, we present our differentiable cloth simulator that both inherits the speedup from PD and handles gradients with complicated contact events like self-collisions, which are overlooked in many existing differentiable simulators but common in cloth simulation.

Differentiating implicit time integration with contact Consider the n -th time step in simulation with contact which takes as input $(\mathbf{x}_n, \mathbf{v}_n)$ and computes $(\mathbf{x}_{n+1}, \mathbf{v}_{n+1}, \mathbf{r}_{n+1}, \mathbf{u}_{n+1})$ that satisfy Eqn. (3.9). In particular, for each contact node j , the forward simulation identifies which one of the contact conditions in Eqn. (3.8) applies to $(\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j})$. As an example, assuming a contact node j satisfies Eqn. (3.8c), its constraints can be summarized as follows:

$$\|\mathbf{r}_{n+1,j|T}\| - \mu \mathbf{r}_{n+1,j|N} = 0, \quad (4.9a)$$

$$\mathbf{u}_{n+1,j|N} = 0, \quad (4.9b)$$

$$(\mathbf{u}_{n+1,j|T})_x (\mathbf{r}_{n+1,j|T})_y - (\mathbf{u}_{n+1,j|T})_y (\mathbf{r}_{n+1,j|T})_x = 0, \quad (4.9c)$$

$$\mathbf{u}_{j|T} \cdot \mathbf{r}_{j|T} \leq 0, \quad (4.9d)$$

where $(\cdot)_x$ and $(\cdot)_y$ extract the x and y components of a two-dimensional vector, respectively. The other two cases of Eqn. (3.8) similarly enforce three equality constraints ($\mathbf{r}_{n+1,j} = 0$ for taking off and $\mathbf{u}_{n+1,j} = 0$ for sticking) and a number of inequality constraints on $(\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j})$. If the inequality constraint is inactive, slightly perturbing the inputs to the simulator will keep $(\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j})$ inside its interior. Therefore, we can remove the inactive inequality when deducing the gradients during backpropagation. When the inequality constraint is active, the gradients of the simulation are not well defined because it represents corner cases that can be categorized into more than one contact types. These corner cases introduce non-smoothness

to the simulator, but it is worth mentioning that they do not create discontinuities, just like the standard rectifier (ReLU) activation function is still continuous despite the non-smoothness at its turning point.

After removing the inequality constraint, we further define a nonlinear vector function $\mathbf{C}_n^j(\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j})$ for the left-hand side of the three equality constraints and compactly represent the constraints as $\mathbf{C}_n^j(\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j}) = 0$. It is convenient that $\mathbf{C}_n^j$ for all three cases in Eqn. (3.8) have three dimensional outputs. We can stack $\mathbf{C}_n^j$ from all contact nodes $j \in \mathcal{I}_n$ into a nonlinear function $\mathbf{C}_n(\cdot, \cdot) : \mathbf{R}^{3|\mathcal{I}_n|} \times \mathbf{R}^{3|\mathcal{I}_n|} \rightarrow \mathbf{R}^{3|\mathcal{I}_n|}$ (note that we only evaluate the function at valid input pairs that lie in $\mathcal{C}^j$), allowing us to restate Eqn. (3.9) as follows:

$$\begin{cases} \mathbf{v}_{n+1} = \mathbf{v}_n + h\mathbf{M}^{-1}[\mathbf{f}_{\text{int}}(\mathbf{x}_n + h\mathbf{v}_{n+1}) + \mathbf{f}_{\text{ext}} + \mathbf{J}_n^\top \mathbf{r}_{n+1}], \\ \mathbf{C}_n(\mathbf{r}_{n+1}, \mathbf{J}_n \mathbf{v}_{n+1}) = \mathbf{0}. \end{cases} \quad (4.10)$$

Note that we choose $\mathbf{v}_{n+1}$ instead of $\mathbf{x}_{n+1}$ as our variable in order to be consistent with the forward PD simulation with contact described by Ly et al. [2020]. Differentiating both sides of the first equation with respect to $\mathbf{v}_n$ leads to the following result:

$$\frac{\partial \mathbf{v}_{n+1}}{\partial \mathbf{v}_n} = \mathbf{I} + h\mathbf{M}^{-1}[-h\nabla^2 W(\mathbf{x}_n + h\mathbf{v}_{n+1}) + \mathbf{J}_n^\top \frac{\partial \mathbf{r}_{n+1}}{\partial \mathbf{v}_{n+1}}] \frac{\partial \mathbf{v}_{n+1}}{\partial \mathbf{v}_n} \quad (4.11)$$

$$\frac{\partial \mathbf{v}_{n+1}}{\partial \mathbf{v}_n} = [\mathbf{I} + h^2\mathbf{M}^{-1}\nabla^2 W(\mathbf{x}_n + h\mathbf{v}_{n+1}) - h\mathbf{M}^{-1}\mathbf{J}_n^\top \frac{\partial \mathbf{r}_{n+1}}{\partial \mathbf{v}_{n+1}}]^{-1} \quad (4.12)$$

$$= [\mathbf{M} + h^2\nabla^2 W(\mathbf{x}_n + h\mathbf{v}_{n+1}) - \underbrace{h\mathbf{J}_n^\top \frac{\partial \mathbf{r}_{n+1}}{\partial \mathbf{v}_{n+1}}}_{\Delta \mathbf{R}^\top}]^{-1} \mathbf{M}. \quad (4.13)$$

Comparing it with Eqn. (4.3), we see the matrix to be inverted now has an additional component dependent on $\frac{\partial \mathbf{r}_{n+1}}{\partial \mathbf{v}_{n+1}}$, which we obtain from differentiating the constraint $\mathbf{C}_n = \mathbf{0}$ in Eqn. (4.10):

$$\frac{\partial \mathbf{C}_n}{\partial \mathbf{r}} \Big|_{\mathbf{r}_{n+1}} \frac{\partial \mathbf{r}_{n+1}}{\partial \mathbf{v}_{n+1}} + \frac{\partial \mathbf{C}_n}{\partial \mathbf{u}} \Big|_{\mathbf{J}_n \mathbf{v}_{n+1}} \mathbf{J}_n = \mathbf{0}, \quad (4.14)$$

$$\frac{\partial \mathbf{r}_{n+1}}{\partial \mathbf{v}_{n+1}} = -\left(\frac{\partial \mathbf{C}_n}{\partial \mathbf{r}} \Big|_{\mathbf{r}_{n+1}}\right)^{-1} \frac{\partial \mathbf{C}_n}{\partial \mathbf{u}} \Big|_{\mathbf{J}_n \mathbf{v}_{n+1}} \mathbf{J}_n. \quad (4.15)$$

We stress that computing $\frac{\partial \mathbf{C}_n}{\partial \mathbf{r}}$ and $\frac{\partial \mathbf{C}_n}{\partial \mathbf{u}}$ is trivial because both partial derivatives are 3×3 block-diagonal matrices. Therefore, $\Delta \mathbf{R}$ can be parallelized among all contact nodes. Backpropagation through $\mathbf{v}_{n+1}$ to $\mathbf{v}_n$ can be implemented with the same adjoint method before, which we give in the equation below for completeness with the notation $\mathbf{z}_{n+1}$ overloaded:

$$\frac{\partial L}{\partial \mathbf{v}_n} = \mathbf{M} \underbrace{[\mathbf{M} + h^2 \nabla^2 W(\mathbf{x} + h\mathbf{v}_{n+1}) - \Delta \mathbf{R}]^{-1}}_{\mathbf{z}_{n+1}} \frac{\partial L}{\partial \mathbf{v}_{n+1}}. \quad (4.16)$$

Speedup with Projective Dynamics With additional information about W from PD, we can rewrite the adjoint vector $\mathbf{z}_{n+1}$ in Eqn. (4.16) by comparing it with Eqn. (4.7):

$$(\mathbf{P} - \Delta \mathbf{P} - \Delta \mathbf{R})\mathbf{z}_{n+1} = \frac{\partial L}{\partial \mathbf{v}_{n+1}}, \quad (4.17)$$

which naturally leads to the following iterative solver:

$$\mathbf{P}\mathbf{z}_{n+1}^{k+1} = (\Delta \mathbf{P} + \Delta \mathbf{R})\mathbf{z}_{n+1}^k + \frac{\partial L}{\partial \mathbf{v}_{n+1}}. \quad (4.18)$$

Comparing it with Eqn. (4.8), we see the role of $\Delta \mathbf{P}$ is replaced with $\Delta \mathbf{P} + \Delta \mathbf{R}$. Since we have shown that computing $\Delta \mathbf{R}$ can be massively parallelized among contact nodes, it is suitable for contact-rich cloth simulation and preserves the efficiency of the local-global solver.

Convergence In theory, such an iterative solver is guaranteed to converge from any initial guesses of $\mathbf{z}_{n+1}$ when the spectral radius $\rho[\mathbf{P}^{-1}(\Delta \mathbf{P} + \Delta \mathbf{R})] < 1$. Empirically, we notice in our cloth simulation that divergence is uncommon, especially when high-precision back-propagation is not required (Sec. 6). When the iterative solver fails to converge, we switch back to a direct sparse matrix solver to solve Eqn. (4.17).

Extensions We deliberately skipped the full definition of $\mathbf{J}_n$ in all equations above for a clearer presentation of the main idea behind our differentiable cloth simulation.

We now elaborate on the role of $\mathbf{J}_n$ and discuss its implications on more complex collisions. When the contact surface is static but non-planar with spatially varying surface normals, $\mathbf{J}_n$ is dependent on the positions of the nodes where the contact events occur. In this case, we estimate the contact normal based on the positions where the contact events occur, which are given by any collision detection algorithms. Replacing $\mathbf{J}_n$ with $\mathbf{J}_n(\mathbf{x}_n, \mathbf{v}_n)$ requires very minimal extra work to the gradients in Sec. 4.2, as it contributes to the gradients with respect to $\mathbf{v}_n$ in a straightforward manner. Another type of collisions we consider is self-collisions between nodes. In this case, contact occurs between pairs of nodes where the contact normals are defined by the relative position between two nodes in the pair. Therefore, we can still define $\mathbf{J}_n$ as a function of $\mathbf{x}_n$ with each row block now consisting of two blocks of nonzero elements corresponding to the two contact nodes in the pair. Similarly, the gradient derivation remains unchanged except that the dependencies between $\mathbf{J}_n$ and $(\mathbf{x}_n, \mathbf{v}_n)$ need to be added to Eqn. (4.11). As we inherit from Ly et al. [2020] the contact model on nodes only, we also inherit its limitation of not handling other types of self-collisions like edge-edge collisions or vertex-face collisions. We leave the derivation of gradients with such cases as future work.

Chapter 5

Evaluations

This section evaluates and discusses numerical properties of the proposed differentiable cloth simulation method in Sec. 4. We first evaluate its gradients by analyzing their source of non-smoothness and studying their usefulness in high-dimensional optimization problems. Next, we compare the dry frictional model with the contact model in the state-of-the-art differentiable cloth simulation method [Liang et al. 2019] and discuss their differences. Finally, we analyze the numerical properties of our iterative solver in back-propagation and compare its performance with a direct linear solver.

5.1 Continuity and Smoothness

A fundamental assumption in any differentiable simulators is that the underlying physics system is smooth so that gradients can be well-defined. Eqn. (3.9) contains three possible sources of discontinuities and non-smoothness in the order of their occurrences: determining the contact set $\mathcal{I}_n$, computing $\mathbf{J}_n$ that represents the local contact frames, and choosing between three branches from $\mathcal{C}_n^j$ in Eqn. (3.8). Below we discuss the effects from each of them in detail.

Continuity of branches in contact conditions First, we empirically demonstrate through an experiment below that switching between the three branches in Eqn. (3.8) does not create discontinuity in Eqn. (3.9). In particular, consider the

n -th time step and assume that $\mathcal{C}_n^j$ is fixed and that $\mathbf{J}_n$ is a continuous function of $(\mathbf{x}_n, \mathbf{v}_n)$, if we treat Eqn. (3.9) as a function that takes as input $(\mathbf{x}_n, \mathbf{v}_n)$ and returns $(\mathbf{x}_{n+1}, \mathbf{v}_{n+1})$, we will validate that this function is continuous. In other words, perturbing $(\mathbf{x}_n, \mathbf{v}_n)$ a little will not cause jumps in the resultant $(\mathbf{x}_{n+1}, \mathbf{v}_{n+1})$ even though the corresponding $\mathbf{r}_{n+1}$ may need to switch between branches during the perturbation. The intuition is that the three branches in Eqn. (3.8) together define a connect set $\mathcal{C}_n^j$ in which $(\mathbf{r}_{n+1,j}, \mathbf{u}_{n+1,j})$ from one branch can smoothly transition to another.

We empirically validate the continuity of branch switching with the following experiment. We simulate a piece of cloth on a rigid, static, and frictional sphere for 200 timesteps. The sphere has one frictional coefficient μ , and we repeat the experiments by varying μ from 0 to 0.35 (Fig. 5-1). When μ is large, all nodes are fixed on the sphere due to their large static friction. When μ is close to 0, the cloth slides on the sphere under gravity and takes off from the sphere near the end of the simulation. As μ changes from 0.35 to 0, each node on the cloth undergoes the transition from sticking to slipping, and eventually takes off. However, since each node has a different contact normal, the turning point for each of them to switch between these branches is different. Overall, when we gradually change μ , the ratio among the numbers of nodes with static friction, dynamic friction, and no friction at the end of the simulation also changes gradually, allowing us to observe how switching between these branches affects the continuity of the physical quantities in simulation.

We summarize the quantitative results from this simulation experiment in Fig. 5-1 (green curves). Specifically, we plot the velocity of three nodes A , B , and C as a function of μ at an intermediate time step (50) and near the end of simulation (200). We select three nodes from the corners, edge centers, and the center of the cloth, respectively. All velocities converge to 0 when μ becomes large, which is as expected because a large μ leads to static friction that freezes these nodes. We also notice a turning point in the velocity curve for each node, indicating a switch between static and dynamic friction. The turning points are located at slightly different μ values for each node because their normals on the sphere surface are different. We observe from the figure that the branches in Eqn. (3.8) do not cause discontinuities.

While the above experiment confirms that these branches do not create discontinuities, we note that branch switching does introduce non-smoothness due to corner cases that can be arbitrarily classified into either branch, e.g., when a node is static but about to slip. Gradients at these corner cases are not well-defined, but the subset these corner cases occupy in $\mathcal{C}_n^j$ has measure zero. Therefore, we still expect gradient-based optimization to be functional, just like we have observed the success of gradient-based optimization in modern deep learning with non-smooth but continuous operators, e.g., ReLU, max pooling, and so on.

Continuity of local frames at contact nodes A second source of possible discontinuity and non-smoothness comes from computing $\mathbf{J}_n$, which consists of two steps: determining a contact point on the contact surfaces and calculating the local normal and tangent vectors. Both steps depend on the geometric representation of the contact surfaces. As pointed out by previous papers in rigid body dynamics [Popović et al. 2000; Werling et al. 2021], contact surface discretization causes jumps in surface normals, and therefore they create discontinuities in velocity and position calculation. To confirm this observation in cloth simulation, we repeat the previous experiment by replacing the analytically described sphere with a triangulated one (pink triangulated sphere surface in Fig. 5-1 right). After such replacement, we clearly observe jumps in the intermediate and final velocity from the selected contact nodes (pink curves in Fig. 5-1). Note that the jumps in velocities from the final velocities are less evident than from the intermediate velocities because the vertical axes have different ranges. This observation agrees with similar experiments from previous papers about rigid body dynamics and suggests that one should favor analytical surfaces in differentiable cloth simulation whenever possible.

We end our discussion on the discontinuities from $\mathbf{J}_n$ with two more remarks. First, we notice the contact node velocity curves in Fig. 5-1 are partitioned into a small number of continuous segments, which is consistent with the result reported by previous work [Popović et al. 2000] discussing contact events on rigid body dynamics. Second, if the scene contains multiple objects the cloth can be in contact with, it is

not uncommon to see jumps in the locations of contact events, e.g., from one object to another, even with a continuous representation of each object. Such jumps naturally lead to large discontinuities in simulation. While we do not provide solutions to it in this work, a closely related issue has been extensively studied in differentiable rendering [Li et al. 2018a] from which we may borrow inspirations in the future.

Continuity of contact sets The last and most common source of discontinuity comes from deciding the contact set at each time step, i.e., whether a node should be added to the contact set or not. Obviously, this selection process is not continuous. It is worth noting that having a constant contact set does not cause too much trouble for gradient computation regardless of its size (Fig. 5-1). Instead, it is the *change* in the contact set from time to time that brings in discontinuities because whenever a new node is put in the contact set, it adds an impulsive force to the right-hand side of Eqn. (3.9).

To better understand the effects of changes in contact sets, we hang a piece of cloth above a static sphere in simulation and let the lower half of the cloth fall and slide on the sphere due to gravity (Fig. 5-2 top). We equip this experiment with a system identification task of the frictional coefficient μ and the stiffness parameter k of the cloth: given a motion sequence of the cloth generated from a pair of unknown μ and k , we define a loss function that sums over all time steps the squared distance between each node position in simulation and its corresponding location in the given motion sequence. We repeat the task with four settings of the sphere at different horizontal offsets, leading to a varying frequency of contact events among them.

We plot the landscape of the loss function in Fig. 5-2 (middle) and compare its smoothness among the four settings with different frequencies of contact events. At first glance, it seems that all four landscapes are equally smooth. However, magnifying a small region of each landscape shows that there is a profound distinction between their continuity and smoothness (Fig. 5-2 bottom): as establishing and breaking contact becomes more frequent, the local landscape tends to be bumpier.

Summary We have discussed the three sources of potential discontinuities and non-smoothness in our differentiable cloth simulator ordered by their damage to the gradients: the branches in the contact conditions only introduce non-smoothness and maintain continuity; contact surface discretization creates discontinuities due to the jumps of normals across adjacent triangles, but we still expect continuity almost everywhere; adding or deleting nodes in the contact set creates frequent and the most severe discontinuities due to the introduction or removal of impulsive forces.

5.2 Usefulness of Gradients

One advantage of gradient-based approaches over their gradient-free counterparts is their faster convergence rate: typically, gradient-free methods explore the local landscape of an objective function by evaluating it with massive samples in the neighborhood. However, such a sampling approach quickly becomes less efficient as the dimension of the decision variables grows higher. Due to this reason, we hypothesize that using gradients from our differentiable simulator will become most beneficial when we have a large number of decision variables. We verify this hypothesis using a control optimization problem: we simulate a piece of cloth with time-invariant external forces applied to each node. The goal is to design the force at each node to pull the center of the cloth to a target position in the end (Fig. 5-3 top). We simulate the cloth with four settings on its DoFs: 4×4 nodes, 5×5 nodes, 7×7 nodes, and 10×10 nodes. Therefore, the number of variables in the optimization problem is 48, 75, 147, and 300, respectively.

Fig. 5-3 reports in this problem the convergence rates of L-BFGS-B [Liu and Nocedal 1989] and CMA-ES [Hansen 2006], two representative gradient-based and gradient-free approaches, with varying DoFs. We start with the largest setting of the cloth (300 DoFs) and run both L-BFGS-B and CMA-ES with five randomly chosen initial guesses on the force values in parallel. We then plot the loss vs. time step curves for both methods in Fig. 5-3 bottom left. Here, the loss is the minimum

loss from each method’s five parallel runs as a function of the time steps they have consumed during the optimization process. It is clear to see the obvious speedup of L-BFGS-B over CMA-ES in this 300-DoF optimization problem, just as we expected.

Additionally, we repeat the experiment with three other settings of the cloth and plot the speedup vs. DoF curves in Fig. 5-3 bottom right, where the speedup is the ratio between the number of time steps used by CMA-ES and L-BFGS-B when their losses reach 0.01. Again, we observe significant speedup from L-BFGS-B across the board, with the largest speedup from the cloth with the highest DoFs. This observation confirms the benefits of using gradients from our differentiable cloth simulator in inverse design problems.

5.3 Benefits of Dry Frictional Contact

We compare the contact model in our differentiable simulator with that in the state-of-the-art differentiable cloth simulator from Liang et al. [2019]. Both models ensure penetration-free simulation and detect self-collisions (node-node collisions in ours and vertex-face and edge-edge collisions in Liang et al. [2019]). However, the contact model in Liang et al. [2019] does not take into account complementarity conditions on either contact forces or frictional forces, which may lead to undesirable artifacts. To visualize this issue, we consider a test scenario in which a napkin falls freely onto the inner surface of a bowl (Fig. 5-4). The differentiable simulators from both Liang et al. [2019] and our work manage to simulate the napkin without numerical explosion. However, the dry frictional contact model in our differentiable simulator shows more physically realistic motion (Fig. 5-4 middle), whereas the contact model from Liang et al. [2019] leads to more drastic changes in the size of the napkin with a popping artifact after its contact with the bowl (Fig. 5-4 top and bottom). These artifacts are explainable because the collision handling algorithm in Liang et al. [2019] modifies node positions after penetration without verifying whether such an update requires sticky contact forces. When the napkin becomes in contact with the concave inner surface of the bowl, such a direct modification injects extra elastic energies.

This experiment confirms supporting a more advanced contact and friction model in differentiable cloth simulation is both viable and beneficial.

5.4 Evaluation of Iterative Solver in Backpropagation

Another key component in our differentiable cloth simulation is the iterative solver in Eqn. (4.17) that utilizes the prefactorized $\mathbf{P}$ in PD. In a standard differentiable simulator, solving Eqn. (4.17) would be done with a sparse matrix solver treating $(\mathbf{P} - \Delta\mathbf{P} - \Delta\mathbf{R})$ as a whole. To understand the performance of the proposed iterative solver, we design two benchmark tests: a “Wind” test where a hanging napkin moves under synthetic wind and a “Slope” test where a ribbon slides on a slanted plane (Fig. 5-5). These two tests represent two extreme cases in contact handling: the napkin in “Wind” barely has any contact or self-collisions, whereas every single node of the ribbon in “Slope” is in contact with the plane. We then compare the time cost of our iterative solver with a sparse LU solver in both tests. We implement both solvers using Eigen [Guennebaud et al. 2010] and choose LU factorization because $(\mathbf{P} - \Delta\mathbf{P} - \Delta\mathbf{R})$ is usually not a symmetric or positive definite matrix, preventing us from using more specialized sparse matrix solvers like Cholesky factorization or Conjugate Gradient methods. For the iterative solver, we report two results from low-precision ($1e-4$) and high-precision ($1e-6$) convergence thresholds that control the termination condition in the iterations. We find these thresholds by varying it from $1e-1$ to $1e-9$ until the gradients computed from the iterative solver start to agree with the direct solver, which we treat as the ground-truth gradients. We repeat the experiments with three mesh resolutions to test the scalability of our method.

We summarize the statistics from both tests in Table 5.1. Overall, the iterative solver in backpropagation is faster than the direct solver, and the speedup becomes more substantial as we increase the mesh resolution. The speedup is also more evident with low-precision threshold, which is consistent with the results reported in

previous PD papers [Bouaziz et al. 2014; Du et al. 2021; Liu et al. 2017]. A further decomposition of time confirms that solving the linear system takes up the majority of backpropagation time, so any improvement in the choice of solver has a dominating positive effect. Although the low-precision results may not match the ground-truth gradients as closely as their high-precision counterparts, their backpropagation is significantly faster and can still benefit gradient-based optimization. This is because even an imperfect gradient can still guide gradient descent algorithms to converge as long as it is along the descending direction. This is confirmed empirically by our experiments in Sec. 6, in most of which we use low-precision convergence threshold and still observe successful gradient-based optimization results. It is also worth mentioning that contact-related operators, whose time cost is included in the “Other” and “ $\Delta \mathbf{R}$ ” columns in Table 5.1, add very little extra overhead to the backpropagation. Finally, we observe uncommon but non-negligible failure cases of the iterative solvers in one experiment (48×48 with $1e-6$ as the threshold in Table 5.1) “Wind”. For the 15% failed timesteps, we switch to the sparse LU solver, adding an extra 11.5% time in backpropagation.




Figure 5-1: We simulate a piece of cloth sliding on a rigid sphere (right) to study the discontinuities and non-smoothness from contact conditions and surface geometry (Sec. 5.1). The two columns of plots show the velocities of three contact nodes at the 50-th and 200-th time steps. The “discretized” and “continuous” labels indicate whether these velocities are computed with a triangulated sphere (visualized as pink triangles on the spherical surface) or a smooth sphere (visualized as green sphere surface).








Figure 5-2: We simulate collisions between a piece of cloth and a rigid sphere at four different horizontal locations to study the effects of varying numbers of collisions on the smoothness of the simulation. Top: rendering of one of the time steps when the cloth is expected to be in contact with the sphere. The four scenes are ordered by their number of collisions. Bottom: landscapes of the loss defined as a function of the frictional coefficient μ and the stiffness k of the cloth (Sec. 5.1). The zoomed-in views of the landscapes show increasing discontinuity and non-smoothness as more collisions occur.




Figure 5-3: **Flying Napkin.** Comparisons between gradient-based and gradient-free methods for optimizing high-dimensional decision variables in differentiable cloth simulation. Top: the motion sequence of a piece of cloth with external forces parametrized with 300 variables. The goal is to find proper external forces with which the cloth ends at a target position. The loss function is defined as the position discrepancy between the simulated motion sequence and the reference motion. Bottom left: the loss vs. time step curves (300 DoF) from gradient-based (L-BFGS-B) and gradient-free (CMA-ES) methods. Both L-BFGS-B and CMA-ES use the same random seeds. Bottom right: we vary the degrees of freedom in the external force parametrization and repeat the experiment three times. For each experiment, we compute the ratio between the number of time steps used by gradient-free and gradient-based methods until they converge or use up the time budget and plot them as a speedup vs. DoFs curve.




Figure 5-4: **Bowl**. We simulate the motion sequence of the napkin ($50 \times 50 \times 2$ triangles, $h = 5\text{ms}$, 500 time steps) with the contact models from two differentiable cloth simulators: Liang et al. [2019] (top) and ours (middle). The area ratio (bottom) is the ratio between the napkin’s area at the current timestep and in the rest shape. Refer to our video for the full motions.

Table 5.1: Comparison between the iterative solver (ours) and the sparse LU direct solver in backpropagation under various mesh resolutions in the “Wind” and “Slope” tests. The numbers in the “Res.” column report the mesh resolution. The number in parentheses in the “Solver” column indicates the epsilon value controlling the convergence of the Jacobi solver. The “Backprop” column reports the net time in backpropagation and is further decomposed into the next five columns, from which the sum of ratios is 100%: “ $\Delta\mathbf{P}$ ” and “ $\Delta\mathbf{R}$ ” report the time spent on assembling $\Delta\mathbf{P}$ and $\Delta\mathbf{R}$ respectively, “Iter.” and “Direct” shows the time cost by either solver, and operations not covered by these four columns are in the “Other” column. The “Fail” column reports the ratio between the number of timesteps seeing nonconvergence in our iterative solver and the number of total timesteps. The “Speedup” column is the ratio between the direct solver’s time and our time in “Backprop. Time” in each test.

Test	Res.	Solver	Backprop (s)	$\Delta\mathbf{P}$ (%)	$\Delta\mathbf{R}$ (%)	Iter. (%)	Direct (%)	Other (%)	Fail (%)	Speedup
Wind	12x12	Direct	0.606	13.9		-	84.4	1.7		-
		Ours (1e-4)	0.212	39.9	0	55.5	-	4.6	0	2.9x
		Ours (1e-6)	0.424	19.2		78.5	-	2.3		1.4x
	24x24	Direct	3.360	9.9		-	89.1	1.0		-
		Ours (1e-4)	0.591	58.2	0	36.3	-	5.5	0	5.7x
		Ours (1e-6)	2.858	11.7		87.2	-	1.1		1.2x
	48x48	Direct	24.001	6.8		-	92.5	0.7	0	-
		Ours (1e-4)	2.262	63.2	0	29.7	-	7.1	0	10.6x
		Ours (1e-6)	29.255	5.4		82.6	11.5	0.6	15	0.8x
Slope	12x12	Direct	0.978	13.0	3.4	-	82.4	1.2		-
		Ours (1e-4)	0.312	70.1	9.4	16.1	-	4.5	0	3.1x
		Ours (1e-6)	0.643	18.9	2.7	76.8	-	1.6		1.5x
	24x24	Direct	5.331	10.4	1.5	-	87.4	0.7		-
		Ours (1e-4)	0.781	70.1	9.4	16.1	-	4.5	0	6.8x
		Ours (1e-6)	1.507	34.7	4.9	58.1	-	2.4		3.7x
	48x48	Direct	37.296	6.5	0.7	-	92.5	0.4		-
		Ours (1e-4)	3.095	68.1	8.2	19.8	-	3.8	0	12.0x
		Ours (1e-6)	6.825	31.2	3.5	63.6	-	1.8		5.5x




Figure 5-5: **Wind** and **Slope**. Motion sequences from the two benchmark tests for comparing iterative and direct solvers in back-propagation. Top: the “Wind” test ($h = 1/90\text{s}$, 200 time steps) where a piece of cloth moves under synthetic wind. Bottom: the “Slope” test ($h = 1/100\text{s}$, 300 time steps) where a ribbon slides along a slanted plane.

Chapter 6

Applications

In this section, we demonstrate a variety of cloth-related applications that can benefit from our proposed differentiable cloth simulation with dry frictional contact. We repeat all experiments with multiple random seeds and use the same random seed set (therefore same initial parameter set) for all methods. We report the optimization results and comparisons with other gradient-free algorithms in Table 6.1. In Table 6.2, We report the time steps used by each optimization method to reach minimum final loss as well as the convergence ratio of our iterative solver during back-propagation. For gradient-free algorithms, the number of time steps counts the total steps of forward simulation (each steps the simulation forward in time by h). For our gradient-based method, we double this number to include the number of back-propagation steps. Note that in practice, back-propagation takes much less time than forward simulation. We have included the wall clock time for running forward simulation and back-propagation of all examples in Appendix A. In Table 6.1 and all loss vs. time steps plot below, we report the minimum loss or plot the minimum loss envelop achieved across all random seeds. See Appendix A.3 for the complete optimization results for each random seed. Below we describe each application and highlight major results. More information of each example is detailed in Appendix A.2.

Implementations. We write the backbone of our simulator in C++ and use Eigen for matrix and vector operations. Each application defines an optimization problem

which we solve with L-BFGS-B, a classic gradient-based optimizer that can leverage the differentiability of our cloth simulator while limiting parameters to physically-plausible ranges. We use the implementation of *LBFGS++* [Yixuan 2021] for L-BFGS-B, which implements the Moré-Thuente line search [Moré and Thuente 1994].

6.1 System Identification

We start by showing two system identification examples: “T-shirt” (Fig. 6-1) and “Sphere” (Fig. 6-2).

T-shirt In the “T-shirt” example, we are given a sequence of motions of a hanging T-shirt under synthetic wind generated from a parameterized sinusoidal function. The goal is to estimate a material parameter in the cloth (1 DoF) and identify the wind model parameters (5 DoFs controlling the amplitude, phase, and frequency of the sinusoidal function in three dimensions) from the motion data. We define the loss function as the L2-distance between the nodal positions of the T-shirt from the simulation and the given motion sequence.

Unlike the “Wind” example in Sec. 5, contacts and frictions are much more frequent in this example due to self-collision between the front and back layer of the T-shirt under the wind. We show the simulation of the T-shirt with parameters before and after optimization in Fig. 6-1 and compare the three optimizers in Table 6.1 and Table 6.2: L-BFGS-B, CMA-ES, and (1+1)-ES. Both CMA-ES and (1+1)-ES are standard gradient-free evolutionary strategies (ES). We can conclude from Fig. 6-1 and Table 6.1 and 6.2 that all three methods manage to optimize system parameters leading to motion sequences visually identical to the given input, but L-BFGS-B converges much faster due to the extra knowledge of gradients.

Sphere To highlight the effect of dry frictional contact in our simulator, we create the “Sphere” example with the goal of matching the motion sequence of a cloth interacting with a sphere by estimating the frictional coefficient between the sphere and the cloth. In this example, we let a piece of square cloth fall freely on a sphere,

Table 6.1: Comparison between the performance of gradient-based and gradient-free optimization on all examples (except “Hat Controller”) in Sec. 6. The “Param #” and “Seed #” columns report the number of optimization parameters and random seeds used in the experiments, respectively. All methods start the optimization with the same initial random seed set. The “Min Init. Loss” column reports the minimum initial loss across all random seeds. The “Optimized Loss Percentage” column reports the optimized loss *as a percentage* (0-100%) of the minimum initial loss reached across all random seeds for each method: “Final” reports the loss reached when the optimization stops, and “Equal Step #” reports the loss reached by CMA-ES and (1+1)-ES using the same number of time steps for L-BFGS-B convergence. We color the values of loss percentage using a green-orange color scale: green corresponds to 0% and orange corresponds to 100%.

Name	Param #	Seed #	Min Init. Loss	Optimized Loss Percentage					
				L-BFGS-B		CMA-ES		(1+1)-ES	
				Final (%)	Final (%)	Final (%)	Equal Step # (%)	Final (%)	Equal Step # (%)
T-shirt	6	5	6.62	0.53	0.60	15.78		1.74	2.43
Sphere	1	10	0.46	0.43	0.00	0.43		0.65	2.37
Hat	18	5	21.83	0.42	0.55	37.51		0.51	5.87
Sock	36	5	17.08	11.59	17.39	64.70		18.43	52.80
Dress	2	16	0.90	76.91	79.02	80.02		77.25	82.57
Flag	8	10	2.88	4.77	10.57	38.89		25.57	40.14

Table 6.2: The “Convergence Time Steps” column reports the number of simulation time steps used for each optimization method until convergence on all examples shown in Sec. 6. For a fair comparison, the time steps shown for L-BFGS-B include both forward simulation and backward propagation time. The “Iter. Conv. %” column reports the percentage of time steps in backward propagation where the iterative solver converges.

Name	Convergence Time Steps			Iter. Conv. %
	L-BFGS-B	CMA-ES	(1+1)-ES	
T-shirt	4500	53500	47750	99.75
Sphere	2450	18900	11200	100.00
Hat	11200	159200	106000	90.58
Sock	17600	176800	91600	99.41
Dress	2750	3125	10000	84.80
Flag	6800	17300	24800	96.00

whose motion after being in contact with the sphere is largely controlled by the frictional coefficient. This example involves both self-collisions between nodes on the cloth and external contacts with the sphere. Similar to the example above, we run both L-BFGS-B and two ES baselines and report their statistics in Table 6.1 and 6.2. All methods can optimize to a frictional coefficient that generates a motion sequence visually identical to the given input. However, this optimization problem is special in that it only has one parameter and the landscape of the loss function is bumpy due to the large variation in the collision set as a function of the frictional coefficient, as suggested by the *Continuity of contact sets* experiment in Sec. 5.1. In this case, evolutionary strategies can achieve a lower final loss because the samples can freely explore the full range of frictional contact, while L-BFGS-B can get trapped in a local optimum.

6.2 Robot-Assisted Dressing

Another line of research that can benefit from a differentiable cloth simulator is robot-assisted dressing. The mainstream solution to these tasks is typically gradient-free methods like evolutionary strategies, reinforcement learning, or inverse dynamics before a differentiable cloth simulator becomes available. We present two examples to




Figure 6-1: **T-shirt** (4278 DoFs, $h = 1/90$ s, 250 time steps). Estimating the cloth material and wind parameters based on a given synthetic motion sequence of the T-shirt. From left to right we show the simulated T-shirt at 0, 80, 160, and 240 time steps. Top: the ground-truth motion sequence; Middle: simulation with the initial guess on the cloth and wind parameters; Bottom: simulation after optimization with L-BFGS-B.




Figure 6-2: **Sphere** (1875 DoFs, $h = 1/180$ s, 350 time steps). Estimating the frictional coefficient μ between the sphere and the cloth based on the cloth's contact with the sphere. From top to bottom we show the simulation at 0, 100, 200, and 300 time steps. Top: the ground-truth motion sequence; Middle: simulation with the initial guess on μ ; Bottom: simulation after optimization with L-BFGS-B.

demonstrate the usage of gradients in robot-assisted dressing: “Hat” (Fig. 6-3) and “Sock” (Fig. 6-4). In both examples, the goal is to find trajectories for a kinematic robotic manipulator to put on the hat or the sock. The end effectors of the manipulator pick a few prespecified vertices on the cloth meshes and pull them along the kinematic trajectories, which are parametrized as B-splines. By optimizing the parameters of the B-splines (18 DoFs in “Hat” and 36 DoFs in “Sock”), we can direct the manipulator to move the hat until it reaches the target location on top of the sphere. We define the loss function in “Hat” as the L2-distance between the hat’s final position at the last time step and a predefined target position, which we generate by translating the hat’s rest shape onto the top of the sphere. The loss function for “Sock” is defined as the L2-distance between the desired and simulated locations of a few key points manually chosen on the sock and evaluated at the middle and the end of the simulation.

With the gradient information at hand, we run the L-BFGS-B optimizer to tune the parameters of the trajectories and compare its performance with CMA-ES and (1+1)-ES (Table 6.1 and 6.2). We notice that within the same time steps, L-BFGS-B converges substantially faster than the gradient-free baselines to a better solution (Fig. 6-3 and Fig. 6-4). We can safely conclude that the fast convergence of L-BFGS-B unlocked by our differentiable cloth simulator is a clear advantage over gradient-free methods.

6.3 Inverse Design

The next application we present is “Dress”, an inverse design example that aims to optimize cloth material parameters in a dress so that its dynamic motion can satisfy certain design intents. Specifically, we optimize the material parameters of a twirl dress so that after the dress spins, the apex angle of the cone-like dress agrees with the target value (100 degrees in our experiment). We define the loss function as the difference between the hemline height corresponding to the target apex angle and the estimated apex angle from points on the hemline of the dress at the last frame of the




Figure 6-3: **Hat** (1737 DoFs, $h = 1/100$ s, 400 time steps). Optimizing trajectories for a manipulator to move a hat onto the sphere. Top left: Initial trajectory from one random seed before optimization overlaid with intermediate hat positions in simulation. Top right: the optimized trajectory from L-BFGS-B (which shares a visually similar trajectory with the ones optimized by ES algorithms). Bottom: The loss vs. time step curves for all methods.

simulation. We report the optimization results from L-BFGS-B and two ES baselines in Table 6.1 and visualize the simulation results before and after optimization in Fig. 1-1. Similar to the previous tasks, we notice that L-BFGS-B achieves better optimized results using fewer time steps.

6.4 A Real-to-Sim Example

In this section, we present a real-to-sim “Flag” example (Fig. 6-5). In this example, we use the real-world motion sequence captured on a flag flapping in the wind from previous work [White et al. 2007] and aim to reconstruct a digital twin of the scene in simulation. This includes not only estimating the material parameters of the flag but




Figure 6-4: **Sock** (3165 DoFs, $h = 1/160s$, 400 time steps). Optimizing trajectories for a manipulator to put a sock onto the foot model. Top left: One initial trajectory before optimization. We show the intermediate time steps on the left and the final state of the sock on the right. Top right: one control trajectory optimized by L-BFGS-B. The end effectors successfully put the sock onto the foot using the optimized trajectory. Bottom: the loss vs. time step curves for all methods.

also modeling the unknown wind condition at the capture time, which is particularly challenging due to its intricate stochastic model with unknown degrees of freedom. We model the wind force at each time step as a 3D force applied near the center of the scene and spatially decaying proportional to the inverse distance to the center. To model the transient nature of the wind force, we modulate magnitude of the 3D vector of a sinusoidal wave as a function of time with parameterized frequency and phase offset. Together, the material and wind model define an eight-dimensional parameter space to be optimized. We define the loss function as the L2-distance between the positions of all nodes at each time step in simulation and the ground-truth motion sequence.

We solve this task using L-BFGS-B and the two ES methods and report their performance in Table 6.1, 6.2, and Fig. 6-5. All three methods substantially reduce the loss after optimization, but L-BFGS-B achieves a lower final loss. We plot the trajectories of 6 key nodes before and after optimization (orange) along with the ground-truth reference trajectories (yellow) from the motion-captured data in Fig. 6-5. By comparing the left and right images in Fig. 6-5, we can see that the L-BFGS-B optimizer reduces the discrepancy substantially between the simulated and actual trajectories after optimization. The real-to-sim matching is still imperfect as indicated by the nonzero final loss, which we suspect could be due to the simplistic nature of our synthetic wind model. A more sophisticated and expressive wind model, e.g., a neural network, may serve a better role in modeling and matching the real-world physics, which we leave as future work.

6.5 Hat Controller

We end this section with an advanced “Hat” task. Unlike the previous open-loop trajectory optimization with a fixed starting position, the goal of this new task is to train a generalizable closed-loop controller that can put on the hat from a *random* starting position sampled from a fixed-radius hemisphere around the head. Specifically, we train a closed-loop control policy $\mathbf{a}_t = \pi_\theta(\mathbf{s}_t)$, which takes as input the current state $\mathbf{s}_t$ of the task and outputs an action vector $\mathbf{a}_t$ at each time step t . The state vector $\mathbf{s}_t$ includes the hat node positions, the orientation of the hat, and the distance between the two end effectors, and the action vector $\mathbf{a}_t$ represents the position of the two end effectors at the next time step. We represent the control policy π_θ as a neural network parametrized by θ consisting of two fully-connected hidden layers with 64 neurons and `tanh` activation functions (117126 parameters in total). To train the controller with gradient-based optimization, we integrate the neural network policy with our differentiable simulator as described in Fig. 6-7.

In each epoch during training, we randomly sample 20 starting positions of the hat and compute a loss averaged from all simulation sequences. The loss function of




Figure 6-5: **Flag.** A real-to-sim example that reconstructs a digital flapping flag (540 vertices, 1026 triangles, $h = 1/120$ s, 100 time steps) based on motion data captured from real-world experiments. Top: We plot the trajectories of 6 nodes on the cloth from the ground-truth motion (yellow curves) and the simulation results with guesses on the material and wind parameters before optimization (orange curves, left) and after optimization (orange curves, right). Bottom: The loss vs. time step curves for all methods.

each sequence is defined by $L = L_{\text{deform}} + L_{\text{target}} + L_{\text{dir}}$, where L_{deform} measures the stretching of the hat using the distance change between the two end-effectors, L_{target} measures the L2-distance between the last-time-step pose of the hat and the target pose, and L_{dir} is the orientation difference between the last-time-step pose and the target pose of the hat. The gradient of the loss is then computed by our differentiable simulation framework and used by a gradient-based optimizer (Adam [Kingma and Ba 2017]) to update the policy parameters. For testing, we evaluate the controller from 20 fixed starting position configurations uniformly sampled on the hemisphere surface.

To compare our method with gradient-free methods, we also solve the task with Reinforcement Learning (RL), which has been widely used to train complex neural network policies for robot-assisted dressing problems [Clegg et al. 2018]. Specifically, we compare our gradient-based method with PPO [Schulman et al. 2017], a state-of-the-art RL algorithm. We similarly sample a starting position from the hemisphere in each iteration when training PPO. For a fair comparison, we design the reward function r_t to be the sum of the negative counterpart of L plus a constant to avoid negative rewards, and observation of the environment to be $\mathbf{s}_t$. We evaluate Adam and PPO using L as the common metric and plot the optimization curve at the bottom of Fig. 6-6. We see that both methods reach a similar final loss, but with our differentiable simulation framework, the gradient-based method reaches its final loss with a 85x speedup (Adam uses 23200 time steps; PPO uses 1978000 time steps).

After the training process converges, both our gradient-based method and PPO successfully move the hat onto the head from all 20 testing positions. We visualize the trajectories generated by Adam’s trained policy from 10 testing starting positions at the top of Fig. 6-6. Unlike existing differentiable simulation papers [Du et al. 2021; Hu et al. 2020, 2019; Liang et al. 2019] that train a closed-loop network controller only for a fixed state, we highlight that we use differentiable simulation to train the network from multiple, random states and study its generalizability in a test set of unseen states. This allows us to conduct a fairer comparison between gradient-based optimization method and RL methods, which is typically overlooked in previous papers.




Figure 6-6: Hat Controller. We train a closed-loop controller for the advanced “Hat” task (Sec. 6.5) which aims to move the hat from different initial positions (sampled from a fixed-radius hemisphere) onto the head. Top two rows: We visualize the control trajectories of the hat in ten initial positions denoted by their elevation and azimuth angle at the bottom of each subfigure. Bottom: The loss vs. time step curve for our gradient-based optimization and PPO.




Figure 6-7: We present the computation graph of the "Hat Controller" task in Sec. 6.5. We embed the neural network policy into our differentiable simulation pipeline. At time step t during forward simulation (blue arrows), the current particle states $(\mathbf{x}_t, \mathbf{v}_t)$ are transformed into a state vector $\mathbf{s}_t$, which is input to the neural network policy π_θ to produce an action vector $\mathbf{a}_t = \pi_\theta(\mathbf{s}_t)$. DiffCloth then simulates a new state $(\mathbf{x}_{t+1}, \mathbf{v}_{t+1})$ using current particle state and the generated action vector. The whole simulation sequence $\{x_i, v_i\}$ is used to compute a loss. During backward propagation (red arrows), the above computation pass is reversed.

Chapter 7

Conclusions, Limitations, Future Works

In this paper, we presented a differentiable cloth simulator built on PD with Signorini-Coulomb frictional contact. Our differentiable simulator is different from existing papers in its simultaneous accommodation of rich and frequent (self-)contact, Signorini-Coulomb contact law, and differentiability in cloth simulation. We analyzed the numerical properties of gradients from our differentiable simulator, including the source of discontinuities and its empirical speedup over gradient-free approaches in high-dimensional problems. We additionally presented an iterative solver that exploits the contact gradients to speed up the backpropagation and observed a substantial speedup (up to an order of magnitude with low-precision simulation) over direct solvers. Our differentiable cloth simulator enabled gradient-based optimization methods in a diverse set of applications, for which traditional gradient-free methods are generally much less sampling efficient. In particular, we presented a preliminary study on training a generalizable closed-loop controller using differentiable simulation, in which our approach and PPO achieved comparable performance and generalizability, but we used much less time.

There are still quite a few limitations in our method that are worth further investigation. First, since our method is built on PD, it also limits the choice of material models. It would be useful to generalize the current framework and support more

physically accurate cloth material models, e.g., the piecewise linear elastic model described in Wang et al. [2011].

Second, the contact model we use from Ly et al. [2020] does not take into account vertex-face or edge-edge self-collisions. While we empirically observed that handling only vertex-vertex self-collisions managed to produce plausible results with medium mesh resolution, vertex-face and edge-edge collision detection and handling is still highly desirable for a more physically realistic cloth simulator.

Third, although we have identified some possible sources of discontinuities and non-smoothness in our differentiable simulator with empirical experiments, their effects on gradient-based optimization still require a thorough investigation. In particular, the locally bumpy energy landscape we observed in Fig. 5-2 due to changes in contact sets makes us question both the necessity and the usefulness of exact gradients in optimization, although Sec. 6 implies that gradients were still helpful in many downstream applications. Noting that the bumpiness in Fig. 5-2 is local and the global view of the energy landscape is still smooth, we hypothesize an inexact but smoothed gradients would be more powerful, which we leave as future work.

Fourth, there is no theoretical guarantee on the convergence of the iterative solver we implemented in backpropagation. Although non-convergence is uncommon in our experiments, it introduces a costly switch to the slower direct solver, which we hope to fully resolve in future work.

Finally, many of our applications were in simulation only. It would be more exciting if these results could be replicated in real-world settings. We consider connecting our differentiable cloth simulator to more real-world applications, including real-world robot-assisted dressing, material parameter identification for real-world fabric samples, computational design for sports suits, and so on. Closing the sim-to-real gap is a nontrivial problem, in which we believe our differentiable simulator could play a beneficial role.

Appendix A

Appendix

A.1 Experiment Run Time

We run all optimizations on a workstation of 80 CPU cores and 80G memory. Depending on the problem complexity, the wallclock time of running these optimizations varies from less than 30 minutes to 2 hours for all methods. We report the run time for optimizing the examples shown in Sec. 6 using our gradient-based method in Table A.1.

Table A.1: Run time for our gradient-based optimization. We report the average wall-clock time for optimizing the examples shown in Sec. 6. We recall the simulation complexity of each example by reiterating its total Degrees of Freedom (“Dof”), number of time steps in each forward simulation sequence and back-propagation (“Time Steps”) and time interval (“ $h[s]$ ”). “Fwd.[s]” and “Back.[s]” report the mean wall-clock time for performing one iteration of forward simulation and back-propagation respectively, averaged across all optimization iterations for all initial seeds.

Name	Dof	Time Steps	h [s]	Fwd.[s]	Back.[s]
T-shirt	4278	250	1/90	141.5	13.2
Sphere	1875	350	1/180	5.1	4.2
Hat	1737	400	1/100	57.9	20.7
Sock	3165	400	1/160	89.9	14.3
Dress	10902	125	1/120	266.3	84.7
Flag	1620	100	1/120	13.7	1.6

A.2 Experiment Details

We provide detailed information for the examples shown in Sec. 6, including their setup, the exact form of their loss function, and their decision variables in optimization.

A.2.1 System Identification

T-shirt The loss function is defined as

$$L = \sum_{n=1}^N \|\mathbf{x}_n^{\text{current}} - \mathbf{x}_n^{\text{target}}\|_2, \quad (\text{A.1})$$

where $N = 240$ is the total number of time steps, $\mathbf{x}_n^{\text{current}}$ the position of the mesh during optimization, and $\mathbf{x}_n^{\text{target}}$ the position of the ground-truth simulation generated using our system. There are 4 parameters (6 DoFs) to be optimized: $\theta_{\text{T-shirt}} = (k_{\text{stretch}}, \phi, \omega, \mathbf{d})$, where k_{stretch} controls the stretching stiffness of the cloth material and $(\phi, \omega, \mathbf{d})$ controls the parameterized external wind force: each node receives a three-dimensional wind force $0.5[\sin(\omega t + \phi) + 1.0]\mathbf{d}$.

Sphere The loss function is defined the same as in the “T-shirt” example above except that the total time $N = 300$. The decision variable is the 1-DoF frictional coefficient of the sphere.

A.2.2 Robot-Assisted Dressing

Hat The loss function is defined as the L2 norm between the last-time-step position of the hat $\mathbf{x}_N^{\text{current}}$ and a target position $\mathbf{x}_N^{\text{target}}$ generated by translating the initial pose of the hat onto the head:

$$L = \|\mathbf{x}_N^{\text{current}} - \mathbf{x}_N^{\text{target}}\|_2, \quad (\text{A.2})$$

where $N = 400$ is the total number of time steps. The trajectory of each of the two end effectors is controlled by a cubic Hermite spline. Each spline has 3 parameters (9

DoFs): $\theta_{\text{spline}} = (\mathbf{t}_1, \mathbf{t}_2, \mathbf{p}_{\text{end}})$, where $\mathbf{t}_1$ and $\mathbf{t}_2$ are the two tangents of the spline and $\mathbf{p}_{\text{end}}$ is the endpoint of the spline.

Sock The goal is to optimize for the trajectory of the three end effectors holding onto a sock so that the sock can be put onto a synthetic foot model from an initial starting position. To guide the end effectors to first hook the opening of the sock onto the tip of the foot, then slide the sock upward onto the leg, the loss function is designed as

$$L = \sum_{n \in \{\frac{N}{2}, N\}} \sum_{(p_{\text{foot}}, p_{\text{sock}}) \in \mathcal{P}_n} \|\mathbf{x}_{N, p_{\text{foot}}}^{\text{target}} - \mathbf{x}_{n, p_{\text{sock}}}^{\text{current}}\|_2, \quad (\text{A.3})$$

where $\mathcal{P}_t$ defines a set of manually selected keypoint correspondence pairs $(p_{\text{foot}}, p_{\text{sock}})$ between the sock mesh $\mathbf{x}_n^{\text{current}}$ and the foot model $\mathbf{x}_N^{\text{target}}$ at halfway $t = \frac{N}{2}$ and the end of the simulation $t = N$, respectively. We sum up the L-2 norm of the position difference between each correspondence. The optimization parameters are the Hermite spline parameters for each of the four end effectors defined similarly as in the "Hat" example (36 DoFs in total).

A.2.3 Inverse Design

Dress The loss function is defined as

$$L = \sum_{p \in \mathcal{P}} (h_p - h)^2, \quad (\text{A.4})$$

where h is the calculated target height of the bottom of the dress when the desired apex angle is reached. $\mathcal{P}$ is the set of all points located at the bottom of the dress, and h_p is the height of the point p (corresponding to the y coordinate of the point in our implementation). There are 2 parameters (2 DoFs) that are being optimized: $\theta_{\text{Dress}} = (d, k_{\text{bend}})$ where d is the density of the fabric and k_{bend} is the bending stiffness of the fabric.

A.2.4 Real-to-Sim Example

In this task, we optimize for the material parameters of the flag and the parameters of a simple wind model to match the motion trajectory of a flag to real-world captured data. The wind model is defined as

$$\mathbf{f}_{\text{ext}} = \frac{\sin(\omega t + \phi) + 1.0}{2} \mathbf{k} \mathbf{d}^\top, \quad (\text{A.5})$$

where $\mathbf{d}$ is a 3D vector to be optimized and $\mathbf{k} \in \mathbf{R}^m$ a constant coefficient vector with each entry equal to a node’s inverse distance to the flag center evaluated at the first time step. There are 8 parameters to be optimized $\theta_{\text{flag}} = (k_{\text{stretch}}, k_{\text{bend}}, \rho, \phi, \omega, \mathbf{d})$, where k_{stretch} and k_{bend} are the stretching and bending stiffness of the fabric, ρ the density of the fabric, and $\phi, \omega, \mathbf{d}$ the parameters of the wind model defined above.

A.2.5 Hat Controller

In this task, the goal is to optimize for the neural network parameters of the hat controller so that the two end effectors put a hat on a head model from any starting position defined on a fixed-radius hemisphere centered at the head model. During training, we uniformly sample 20 starting positions on the hemisphere for each epoch, and compute the loss averaged from all simulation sequences. We define the loss function as

$$L = L_{\text{deform}} + L_{\text{target}} + L_{\text{dir}}. \quad (\text{A.6})$$

L_{deform} minimizes the distance change between the two end effectors and is defined as $L_{\text{deform}} = \sum_{n=1}^N \|\mathbf{x}_{n,e_1} - \mathbf{x}_{n,e_2}\|_2$ where e_1 and e_2 are the indices of the nodes pulled by the two end effectors. L_{target} minimizes the L2-distance between the poses of the hat and the target pose at the last few frames (20 in our implementation) and is defined as

$$L_{\text{target}} = \sum_{n=N-20}^N \|\mathbf{x}_n^{\text{current}} - \mathbf{x}_n^{\text{target}}\|_2. \quad (\text{A.7})$$

L_{dir} minimizes the orientation difference between the last-time-step pose and the target pose of the hat and is defined as

$$L_{\text{dir}} = \mathbf{d}_{\text{target}} \cdot \mathbf{d}_{\text{current}} \quad (\text{A.8})$$

where $\mathbf{d}_{\text{target}}$ is the up vector of the hat at the target pose and $\mathbf{d}_{\text{current}}$ is defined similarly to the hat at the last time step.

A.3 Optimization Results For All Random Seeds

In this section we report the optimization results for all random seeds in Table A.2, Table A.3 and Table A.4. For most experiments, L-BFGS-B achieves lower or comparable optimized loss than the gradient-free methods using a fraction (often to an order of magnitude) of time steps, thanks to the gradient information provided by our differentiable simulator. For some random seeds, L-BFGS-B converges to a relatively large final loss percentage, possibly due to being stuck in a local minimum. In practice, it is common and recommended to run gradient-based optimizations with several initial seeds to alleviate this problem, and is the rationale behind why we choose to report the minimum loss achieved across all random seeds in the results shown in Table 6.1 and plot the minimum loss envelop in Fig 5-3, 6-3, 6-4, 6-5. It is also worth mentioning that the examples shown Sec. 6 have a relative low number of design variables, while the speed-up of gradient-based methods becomes more evident with more design variables as demonstrated by the “Flying Napkin” experiment in Fig. 5-3.

Table A.2: We report the optimization results for all random seeds of all methods in the “T-shirt” and “Flag” examples. For each random seed, we report its initial loss and final loss achieved by each optimization method. “Final Loss Percentage (%)” reports the optimized loss as a *percentage (0-100%)* of the initial loss. “Convergence Time Steps” reports the number of time steps used by each method to reach its final loss. For all tasks, we also summarize the minimum, average and median statistics across all random seeds for each column. For each metric (“Final Loss”, “Final Loss Percentage (%)”, “Convergence Time Steps”) and each row, the minimum number across the three optimization methods is marked in bold.

	Initial Loss	Final Loss		Final Loss Percentage (%)		Convergence Time Steps				
		L-BFGS-B	CMAES (1+1)	L-BFGS-B	CMAES (1+1)	L-BFGS-B	CMAES (1+1)			
T-shirt	22.49	0.042	0.052	1.416	0.2	0.2	6.2	2500	41750	42000
	60.00	0.079	0.269	0.248	0.1	0.4	0.4	7250	46750	9750
	6.62	6.560	0.169	0.115	99.0	2.7	1.8	3250	17500	47750
	30.93	0.035	0.286	0.213	0.1	0.9	0.7	2250	41000	36750
	10.32	0.035	0.078	0.159	0.3	0.8	1.6	6250	27000	19750
Flag	MIN	0.035	0.052	0.115	0.1	0.2	0.4	2250	17500	9750
	AVERAGE	1.350	0.171	0.430	20.0	1.0	2.1	4300	34800	31200
	MEDIAN	0.042	0.169	0.213	0.2	0.8	1.6	3250	41000	36750
	2.88	0.137	1.118	1.154	4.8	38.9	40.1	3400	17300	24800
	3.84	1.136	0.945	1.079	29.6	24.6	28.1	1100	900	27600
Flag	4.07	1.175	0.304	1.021	28.9	7.5	25.1	5200	1200	28500
	5.08	0.595	0.958	1.043	11.7	18.9	20.5	3800	4500	26600
	4.08	1.980	0.997	1.053	48.6	24.5	25.8	200	29200	27600
	3.26	0.190	0.743	0.952	5.8	22.8	29.2	4500	27500	28400
	3.59	0.863	0.697	1.134	24.0	19.4	31.5	800	9900	28100
	3.68	0.156	0.711	1.089	4.2	19.3	29.6	9800	22500	24600
	3.73	1.095	1.077	1.109	29.4	28.9	29.8	3100	12600	28900
	9.98	0.861	1.086	0.715	8.6	10.9	7.2	1800	5400	9200
	MIN	0.137	0.304	0.715	4.2	7.5	7.2	200	900	9200
	AVERAGE	0.819	0.864	1.035	19.6	21.6	26.7	3370	13100	25430
	MEDIAN	0.862	0.952	1.066	17.9	21.1	28.7	3250	11250	27600

Table A.3: Similar table as Table A.2 for the “Sphere” and “Dress” examples.

	Initial Loss	Final Loss		Final Loss Percentage (%)		Convergence Time Steps				
		L-BFGS-B	CMA-ES (1+1)	L-BFGS-B	CMA-ES (1+1)	L-BFGS-B	CMA-ES (1+1)			
Sphere	0.46	0.002	0.000	0.003	0.4	0.0	0.6	2450	18900	11200
	0.90	0.064	0.000	0.000	7.2	0.0	0.0	3500	36050	5600
	0.58	0.002	0.000	0.000	0.3	0.0	0.0	1400	24850	10150
	2.20	0.904	0.000	0.000	41.1	0.0	0.0	700	31850	8050
	4.03	0.904	0.000	0.031	22.4	0.0	0.8	700	15050	26250
	0.90	0.903	0.000	0.020	100.0	0.0	2.0	350	26600	44100
	0.90	0.002	0.002	0.000	0.2	0.2	0.0	6650	5600	18200
	0.89	0.893	0.000	0.008	100.0	0.0	0.8	350	48300	46550
	0.88	0.861	0.000	0.001	97.6	0.0	0.1	1400	49700	9100
	0.85	0.514	0.000	0.000	60.4	0.0	0.0	3150	35700	15400
Dress	MIN	0.002	0.000	0.000	0.2	0.0	0.0	350	5600	5600
	AVERAGE	0.505	0.000	0.006	43.0	0.0	0.4	2065	29260	19460
	MEDIAN	0.688	0.000	0.001	31.7	0.0	0.1	1400	29225	13300
	2.41	1.820	0.712	0.845	75.4	33.3	39.6	375	3125	12750
	1.35	0.716	0.830	0.696	53.2	60.4	50.6	1125	49875	10000
	1.57	1.406	0.824	0.782	89.3	52.3	49.6	1625	31875	9750
	1.03	0.841	0.825	0.822	82.0	80.3	80.0	1625	29375	19750
	0.90	0.880	0.824	0.823	97.7	90.3	90.1	750	19750	16250
	1.49	1.490	0.832	0.698	99.9	55.4	46.5	875	44000	20750
	1.09	0.875	0.828	0.830	80.3	75.2	75.4	250	35000	17250
Dress	1.88	0.693	0.861	0.792	37.0	45.4	41.8	1375	25875	40625
	1.83	1.305	0.823	0.858	71.2	56.1	58.5	1125	50000	10250
	1.27	1.219	0.824	0.854	96.1	64.4	66.8	500	49500	14000
	1.31	1.178	0.814	0.872	90.1	62.2	66.6	2375	4125	10250
	1.26	0.856	0.824	0.823	68.2	65.0	65.0	1000	47125	11000
	MIN	0.693	0.712	0.696	37.0	33.3	39.6	250	3125	9750
	AVERAGE	1.107	0.818	0.808	78.4	61.7	60.9	1083	32469	16052
	MEDIAN	1.029	0.824	0.823	81.1	61.3	61.7	1063	33438	13375

Table A.4: Similar table as Table A.2 for the “Hat”, “Sock” examples.

Initial Loss	Final Loss			Final Loss Percentage (%)		Convergence Time Steps				
	L-BFGS-B	CMA-ES	(1+1)	L-BFGS-B	CMA-ES	(1+1)	L-BFGS-B	CMA-ES	(1+1)	
Hat	54.60	10.391	0.754	0.330	19.0	1.4	0.6	4400	47600	50000
	33.81	0.402	0.665	0.105	1.2	1.9	0.3	2000	46000	36400
	43.45	0.091	0.723	0.236	0.2	1.6	0.5	5600	42800	44000
	48.63	0.105	0.283	1.290	0.2	0.6	2.6	4000	48000	45600
	21.83	0.096	0.946	0.314	0.4	4.3	1.4	2800	38800	30000
AVERAGE MEDIAN	MIN	0.091	0.283	0.105	0.2	0.6	0.3	2000	38800	30000
		2.217	0.674	0.455	4.2	2.0	1.1	3760	44640	41200
		0.105	0.723	0.314	0.4	1.6	0.6	4000	46000	44000
Sock	46.02	1.980	7.267	5.335	4.3	15.8	11.6	8800	42800	47600
	41.22	3.243	6.131	8.396	7.9	14.9	20.4	16000	44000	32400
	39.10	8.856	10.547	3.047	22.6	27.0	7.8	4400	38400	46000
	17.08	2.589	4.149	7.830	15.2	24.3	45.8	15200	31200	24000
	33.29	2.652	4.126	5.791	8.0	12.4	17.4	5200	49200	48000
AVERAGE MEDIAN	MIN	1.980	4.126	3.047	4.3	12.4	7.8	4400	31200	24000
		3.864	6.444	6.080	11.6	18.9	20.6	9920	41120	39600
		2.652	6.131	5.791	8.0	15.8	17.4	8800	42800	46000

References

- David Baraff and Andrew Witkin. 1998. Large Steps in Cloth Simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 43–54.
- Aric Bartle, Alla Sheffer, Vladimir G. Kim, Danny M. Kaufman, Nicholas Vining, and Floraine Berthouzoz. 2016. Physics-Driven Pattern Adjustment for Direct 3D Garment Editing. *ACM Trans. Graph.* 35, 4, Article 50 (July 2016), 11 pages.
- 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, Article 154 (July 2014), 11 pages. <https://doi.org/10.1145/2601097.2601116>
- Robert Bridson, Ronald Fedkiw, and John Anderson. 2002. Robust Treatment of Collisions, Contact and Friction for Cloth Animation. *ACM Trans. Graph.* 21, 3 (July 2002), 594–603.
- R. Bridson, S. Marino, and R. Fedkiw. 2003. Simulation of Clothing with Folds and Wrinkles. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (San Diego, California) (SCA '03)*. Eurographics Association, Goslar, DEU, 28–36.
- Bernard Brogliato. 2016. *Nonsmooth Lagrangian Systems*. Springer International Publishing, Cham, 241–370.
- Remi Brouet, Alla Sheffer, Laurence Boissieux, and Marie-Paule Cani. 2012. Design Preserving Garment Transfer. *ACM Trans. Graph.* 31, 4, Article 36 (July 2012), 11 pages.
- George E. Brown, Matthew Overby, Zahra Forootaninia, and Rahul Narain. 2018. Accurate Dissipative Forces in Optimization Integrators. *ACM Trans. Graph.* 37, 6, Article 282 (Dec. 2018), 14 pages.
- Kwang-Jin Choi and Hyeong-Seok Ko. 2002. Stable but Responsive Cloth. *ACM Trans. Graph.* 21, 3 (July 2002), 604–611.
- Alexander Clegg, Zackory Erickson, Patrick Grady, Greg Turk, Charles C. Kemp, and C. Karen Liu. 2020. Learning to Collaborate From Simulation for Robot-Assisted Dressing. *IEEE Robotics and Automation Letters* 5, 2 (2020), 2746–2753.
- Alexander Clegg, Wenhao Yu, Jie Tan, C. Karen Liu, and Greg Turk. 2018. Learning to Dress: Synthesizing Human Dressing Motion via Deep Reinforcement Learning. *ACM Trans. Graph.* 37, 6, Article 179 (Dec. 2018), 10 pages.

- David Clyde, Joseph Teran, and Rasmus Tamstorf. 2017. Modeling and Data-Driven Parameter Estimation for Woven Fabrics. In *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (Los Angeles, California) (*SCA '17*). Association for Computing Machinery, New York, NY, USA, Article 17, 11 pages.
- C. Dario Bellicoso, Christian Gehring, Jemin Hwangbo, Péter Fankhauser, and Marco Hutter. 2016. Perception-less terrain adaptation through whole body control and hierarchical optimization. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. 558–564. <https://doi.org/10.1109/HUMANOIDS.2016.7803330>
- Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. 2018a. End-to-end differentiable physics for learning and control. *Advances in neural information processing systems* 31 (2018), 7178–7189.
- Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and Zico Kolter. 2018b. 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.
- 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 (Oct. 2021), 21 pages. <https://doi.org/10.1145/3490168>
- Tao Du, Kui Wu, Andrew Spielberg, Wojciech Matusik, Bo Zhu, and Eftychios Sifakis. 2020. Functional Optimization of Fluidic Devices with Differentiable Stokes Flow. *ACM Trans. Graph.* 39, 6, Article 197 (Dec. 2020), 15 pages. <https://doi.org/10.1145/3414685.3417795>
- Marco Fratarcangeli, Valentina Tibaldo, and Fabio Pellacini. 2016. Vivace: A Practical Gauss-Seidel Method for Stable Soft Body Dynamics. *ACM Trans. Graph.* 35, 6, Article 214 (Nov. 2016), 9 pages.
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. 2020. ADD: analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.
- Peng Guan, Loretta Reiss, David A. Hirshberg, Alexander Weiss, and Michael J. Black. 2012. DRAPE: DRessing Any PErson. *ACM Trans. Graph.* 31, 4, Article 35 (July 2012), 10 pages.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.

- David Hahn, Pol Banzet, James M Bern, and Stelian Coros. 2019. Real2sim: Viscoelastic parameter estimation from dynamic motion. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.
- Nikolaus Hansen. 2006. The CMA Evolution Strategy: A Comparing Review.
- David Harmon, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. 2008. Robust Treatment of Simultaneous Collisions. In *ACM SIGGRAPH 2008 Papers* (Los Angeles, California) (*SIGGRAPH '08*). Association for Computing Machinery, New York, NY, USA, Article 23, 4 pages.
- 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 Frédo 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. In *2019 International conference on robotics and automation (ICRA)*. IEEE, 6265–6271.
- Dongho Kang, Simon Zimmermann, and Stelian Coros. 2021. Animal Gaits on Quadrupedal Robots Using Motion Matching and Model-Based Control. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.
- Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs.LG]
- Martin Komaritzan and Mario Botsch. 2018. Projective Skinning. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1, Article 12 (July 2018), 19 pages. <https://doi.org/10.1145/3203203>
- Cheng Li, Min Tang, Ruofeng Tong, Ming Cai, Jieyi Zhao, and Dinesh Manocha. 2020b. P-Cloth: Interactive Complex Cloth Simulation on Multi-GPU Systems Using Dynamic Matrix Assembly and Pipelined Implicit Integrators. *ACM Trans. Graph.* 39, 6, Article 180 (Nov. 2020), 15 pages.
- Jie Li, Gilles Daviet, Rahul Narain, Florence Bertails-Descoubes, Matthew Overby, George E. Brown, and Laurence Boissieux. 2018b. An Implicit Frictional Contact Solver for Adaptive Cloth Simulation. *ACM Trans. Graph.* 37, 4, Article 52 (July 2018), 15 pages.
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020a. Incremental Potential Contact: Intersection-and Inversion-Free, Large-Deformation Dynamics. *ACM Trans. Graph.* 39, 4, Article 49 (July 2020), 20 pages.

- Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021. Codimensional Incremental Potential Contact. *ACM Trans. Graph. (SIGGRAPH)* 40, 4, Article 170 (2021).
- Minchen Li, Alla Sheffer, Eitan Grinspun, and Nicholas Vining. 2018c. FoldsSketch: Enriching Garments with Physically Reproducible Folds. *ACM Trans. Graph.* 37, 4, Article 133 (July 2018), 13 pages.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018a. Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 37, 6 (2018), 222:1–222:11.
- Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua Tenenbaum, and Antonio Torralba. 2019. 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*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2019/file/28f0b864598a1291557bed248a998d4e-Paper.pdf>
- Dong C. Liu and Jorge Nocedal. 1989. On the Limited Memory BFGS Method for Large Scale Optimization. *MATHEMATICAL PROGRAMMING* 45 (1989), 503–528.
- Tiantian Liu, Adam W. Bargteil, James F. O'Brien, and Ladislav Kavan. 2013. Fast Simulation of Mass-Spring Systems. *ACM Trans. Graph.* 32, 6, Article 214 (Nov. 2013), 7 pages.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-newton methods for real-time simulation of hyperelastic materials. *Acm Transactions on Graphics (TOG)* 36, 3 (2017), 1–16.
- Mickaël Ly, Romain Casati, Florence Bertails-Descoubes, Mélina Skouras, and Laurence Boissieux. 2018. Inverse Elastic Shell Design with Contact and Friction. *ACM Trans. Graph.* 37, 6, Article 201 (Dec. 2018), 16 pages.
- Mickaël Ly, Jean Jouve, Laurence Boissieux, and Florence Bertails-Descoubes. 2020. Projective Dynamics with Dry Frictional Contact. *ACM Trans. Graph.* 39, 4, Article 57 (July 2020), 8 pages.
- M. Macklin, K. Erleben, M. Müller, 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* (Virtual Event, Canada) (*SCA '20*). Eurographics Association, Goslar, DEU, Article 9, 12 pages.

- Miles Macklin, Matthias Müller, and Nuttapong Chentanez. 2016. XPBD: Position-Based Simulation of Compliant Constrained Dynamics. In *Proceedings of the 9th International Conference on Motion in Games* (Burlingame, California) (*MIG '16*). Association for Computing Machinery, New York, NY, USA, 49–54.
- Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. 2011. Example-Based Elastic Materials. *ACM Trans. Graph.* 30, 4, Article 72 (July 2011), 8 pages. <https://doi.org/10.1145/2010324.1964967>
- Antoine McNamara, Adrien Treuille, Zoran Popović, and Jos Stam. 2004. Fluid Control Using the Adjoint Method. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 449–456.
- E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner. 2012. Data-Driven Estimation of Cloth Simulation Models. *Comput. Graph. Forum* 31, 2pt2 (May 2012), 519–528.
- Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C. Schvartzman, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Steve Marschner, and Miguel A. Otaduy. 2013. Modeling and Estimation of Internal Friction in Cloth. *ACM Trans. Graph.* 32, 6, Article 212 (Nov. 2013), 10 pages.
- Michael Mistry, Jonas Buchli, and Stefan Schaal. 2010. Inverse dynamics control of floating base systems using orthogonal decomposition. In *2010 IEEE International Conference on Robotics and Automation*. 3406–3412. <https://doi.org/10.1109/ROBOT.2010.5509646>
- Juan Montes, Bernhard Thomaszewski, Sudhir Mudur, and Tiberiu Popa. 2020. Computational Design of Skintight Clothing. *ACM Trans. Graph.* 39, 4, Article 105 (July 2020), 12 pages.
- Jorge J. Moré and David J. Thuente. 1994. Line Search Algorithms with Guaranteed Sufficient Decrease. *ACM Trans. Math. Softw.* 20, 3 (Sept. 1994), 286–307. <https://doi.org/10.1145/192115.192132>
- Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position Based Dynamics. *J. Vis. Comun. Image Represent.* 18, 2 (April 2007), 109–118.
- J. Krishna Murthy, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Breandan Considine, Jérôme Parent-Lévesque, Kevin Xie, Kenny Erleben, Liam Paull, Florian Shkurti, Derek Nowrouzezahrai, and Sanja Fidler. 2021. gradSim: Differentiable simulation for system identification and visuomotor control. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=cE8kFWfhp0>
- Rahul Narain, Armin Samii, and James F. O’Brien. 2012. Adaptive Anisotropic Remeshing for Cloth Simulation. *ACM Trans. Graph.* 31, 6, Article 152 (Nov. 2012), 10 pages.

- Miguel A. Otaduy, Rasmus Tamstorf, Denis Steinemann, and Markus Gross. 2009. Implicit Contact Handling for Deformable Objects. *Computer Graphics Forum* 28, 2 (2009), 559–568.
- Matthew Overby, George E. Brown, Jie Li, and Rahul Narain. 2017. ADMM $\supseteq$ Projective Dynamics: Fast Simulation of Hyperelastic Models with Dynamic Constraints. *IEEE Transactions on Visualization and Computer Graphics* 23, 10 (Oct 2017), 2222–2234.
- Jovan Popović, Steven Seitz, and Michael Erdmann. 2003. Motion Sketching for Control of Rigid-Body Simulations. *ACM Trans. Graph.* 22, 4 (Oct. 2003), 1034–1054.
- Jovan Popović, Steven M. Seitz, Michael Erdmann, Zoran Popović, and Andrew Witkin. 2000. Interactive Manipulation of Rigid Body Simulations. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., USA, 209–217. <https://doi.org/10.1145/344779.344880>
- Xavier Provot. 1997. Collision and self-collision handling in cloth model dedicated to design garments. In *Computer Animation and Simulation '97*, Daniel Thalmann and Michiel van de Panne (Eds.). Springer Vienna, Vienna, 177–189.
- 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)* (2018).
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- Ari Stern and Mathieu Desbrun. 2006. Discrete Geometric Mechanics for Variational Time Integrators. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) (*SIGGRAPH '06*). Association for Computing Machinery, New York, NY, USA, 75–80.
- Ari Stern and Eitan Grinspun. 2009. Implicit-Explicit Variational Integration of Highly Oscillatory Problems. *Multiscale Modeling & Simulation* 7, 4 (2009), 1779–1794.
- Rasmus Tamstorf, Toby Jones, and Stephen F. McCormick. 2015. Smoothed Aggregation Multigrid for Cloth Simulation. *ACM Trans. Graph.* 34, 6, Article 245 (Oct. 2015), 13 pages.

- Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. 1987. Elastically Deformable Models. *SIGGRAPH Comput. Graph.* 21, 4 (Aug. 1987), 205–214.
- 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.
- Nobuyuki Umetani, Danny M. Kaufman, Takeo Igarashi, and Eitan Grinspun. 2011. Sensitive Couture for Interactive Garment Modeling and Editing. *ACM Trans. Graph.* 30, 4, Article 90 (July 2011), 12 pages.
- Huamin Wang. 2015. A Chebyshev Semi-Iterative Approach for Accelerating Projective and Position-Based Dynamics. *ACM Trans. Graph.* 34, 6, Article 246 (Oct. 2015), 9 pages.
- Huamin Wang. 2018. Rule-Free Sewing Pattern Adjustment with Precision and Efficiency. *ACM Trans. Graph.* 37, 4, Article 53 (July 2018), 13 pages.
- Huamin Wang, James F. O’Brien, and Ravi Ramamoorthi. 2011. Data-Driven Elastic Models for Cloth: Modeling and Measurement. In *ACM SIGGRAPH 2011 Papers* (Vancouver, British Columbia, Canada) (*SIGGRAPH ’11*). Association for Computing Machinery, New York, NY, USA, Article 71, 12 pages.
- Huamin Wang and Yin Yang. 2016. Descent Methods for Elastic Body Simulation on the GPU. *ACM Trans. Graph.* 35, 6, Article 212 (Nov. 2016), 10 pages.
- Zhendong Wang, Longhua Wu, Marco Fratarcangeli, Min Tang, and Huamin Wang. 2018. Parallel Multigrid for Nonlinear Cloth Simulation. *Computer Graphics Forum* 37, 7 (2018), 131–141.
- Keenon Werling, Dalton Omens, Jeongseok Lee, Ioannis Exarchos, and C. Karen Liu. 2021. Fast and Feature-Complete Differentiable Physics Engine for Articulated Rigid Bodies with Contact Constraints. In *Proceedings of Robotics: Science and Systems*. Virtual. <https://doi.org/10.15607/RSS.2021.XVII.034>
- Ryan White, Keenan Crane, and David Forsyth. 2007. Capturing and Animating Occluded Cloth. In *ACM Transactions on Graphics (SIGGRAPH)*.
- Andrew Witkin and Michael Kass. 1988. Spacetime Constraints. In *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’88)*. Association for Computing Machinery, New York, NY, USA, 159–168. <https://doi.org/10.1145/54852.378507>
- 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* (Vienna, Austria) (*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, Article 162 (Nov. 2019), 13 pages.
- Jie Xu, Tao Chen, Lara Zlokapa, Michael Foshey, Wojciech Matusik, Shinjiro Sueda, and Pulkit Agrawal. 2021. An End-to-End Differentiable Framework for Contact-Aware Robot Design. In *Proceedings of Robotics: Science and Systems*. Virtual. <https://doi.org/10.15607/RSS.2021.XVII.008>
- Qiu Yixuan. 2021. LBFGS++. <https://github.com/yixuan/LBFGSpp/>.