

Differentiable Dynamics Simulation Using Invariant Contact Mapping and Damped Contact Force

Minji Lee, Jeongmin Lee and Dongjun Lee[†]

Abstract—The gradient of typical differentiable simulation is uninformative for two reasons: 1) non-smoothness in contact dynamics not considered properly, and 2) excessive local minima generated from the smoothing procedure. To tackle this issue, we first propose differentiable contact dynamics with an invariant contact set and coordinate differentiation using a signed distance function (SDF). Also, to eliminate the undesirable jittering caused by the smoothing procedure, which induces extra local minima, and to achieve a smooth and informative gradient, we further endow our framework with a novel damped contact model. Various optimization problems are implemented to demonstrate the usefulness and efficacy of our differentiable framework.

I. INTRODUCTION

Optimization problems about dynamic objects, such as model predictive control [1], trajectory optimization [2], [3], and parameter identification [4] of dynamic objects, are required in various robotics problems. However, it is challenging since consideration of the dynamics of each object and the interaction between other objects or environments is required. It is promising to use the gradient of the differentiable simulator as a way to overcome these limitations [5]–[8]. With the gradient, the gradient-based optimization methods such as the Newton method are achievable and can be used to improve the efficiency of various optimization problems.

Many techniques have been proposed to determine the gradient of the dynamics simulations. First, in the work of [9], a framework for achieving the gradient for free motion of an articulated rigid body algorithm was suggested. Differentiable rigid body simulations have been extended to include constraints such as contact [4], [10], [11]. Following that, studies to determine the gradient of soft object simulation were conducted [8], [12], [13].

Although these various differentiable simulations have been studied, little has been done about whether they can provide useful gradients for solving a variety of problems in practice. Very recently, some studies shed light on the quality of gradient of differentiable simulators, which often turns out to be uninformative (i.e. only useful in a small area) [14]–[17]. Based on our study, it can be explained by two

factors: 1) non-smoothness in contact dynamics that was not properly addressed, and 2) the occurrence of excessive local minima from the smoothing process.

First, the non-smoothness in contact dynamics is caused by discrete change of contact features [12] and non-smooth transitions between contact states [11]. The contact conditions are defined depend on the contact features (i.e. contact point, contact normal vector, and the total number of contact). Thus, as the contact appears and departs, the contact features vary, whereas the contact condition changes discontinuously. If differentiation is not performed in consideration of these changes, the gradient will not be accurate. Also, due to the complementarity nature of the contact model, the transitions between the contact states are non-smooth, thereby the gradient cannot take the transitions into account. Some methods of smoothing the contact model, such as penalty functions [8], interior point method [11], and log-barrier method [18] have been proposed. As demonstrated in this paper, these methods efficiently smooth the simulation result of a single time step; however, when applied sequentially to the simulation with time correlation, jittering occurs in the result of simulation, causing extra local minima and degrading optimization problem efficiency. Therefore, these methods cannot make the gradient to be informative.

In this paper, we suggest a differentiable simulation framework that can offer a reliable and informative gradient considering the non-smoothness in the contact dynamics. The key ideas to tackle these issues in our framework are as follows. We first propose the differentiable contact condition with two components: 1) invariant contact set and 2) contact coordinate differentiation using a signed distance function (SDF). These allow us to differentiate contact condition that was previously not differentiable due to the discontinuous changes of contact features, resulting in a more reliable gradient. Remarkably, this process also turns out to eliminate the need to take time of impact into account for the differentiation, which has a considerable impact on the accuracy of the gradient [7]. We then propose a novel contact model named damped contact force to solve the short-sightedness of the gradient. According to our analysis, the cause of the jitters is the generation of oscillation due to the energy preservation. The damped contact model allows for energy dissipation during contact, resulting in smooth behavior.

The rest of the paper is organized as follows. Sec. II introduces preliminary background about dynamics simulation for this paper. Then the analysis of the non-smoothness in the contact dynamics is presented in Sec. III, followed by proposition of our differentiable simulation framework

This research was supported by the Industrial Strategic Technology Development Program (20001045) of the Ministry of Trade, Industry & Energy (MOTIE) of Korea, the Engineering Research Center Program for Soft Robotics (2016R1A5A1938472), and the RS-2022-00144468 of the National Research Foundation (NRF) funded by the Ministry of Science and ICT (MSIT) of Korea.

The authors are with the Department of Mechanical Engineering, IAMD and IER, Seoul National University, Seoul, Republic of Korea, 08826. {mingg8,ljmlgh,djlee}@snu.ac.kr. Corresponding author: Dongjun Lee.

in Sec. IV. Sec. V presents applications of the framework on optimization problems, followed by some concluding remarks in Sec. VI.

II. PRELIMINARY

A. Dynamics Discretization

The dynamics of an object with contact in continuous time are expressed as

$$M\ddot{q} = -\psi(q) + f_{ext} + J_c^T \lambda_c$$

where $M \in \mathbb{R}^{n \times n}$ is the mass matrix, $q \in \mathbb{R}^n$ is the state of the object, $-\psi(q) \in \mathbb{R}^n$ is the internal force, $f_{ext} \in \mathbb{R}^n$ are the external force, and $\lambda_c = [\lambda_1; \dots; \lambda_m] \in \mathbb{R}^{3m}$, $J_c \in \mathbb{R}^{3m \times n}$ are contact impulse/Jacobian with n, m being the dimension of the state and the number of contact points. Then we can discretize the dynamics as

$$M \frac{v_{k+1} - v_k}{T_k} = -\psi_k(q_{k+1}) + f_{ext,k} + J_{c,k}^T \lambda_{c,k}$$

$$\hat{v}_k = \frac{v_{k+1} + v_k}{2}, \quad q_{k+1} = q_k + \hat{v}_k T_k$$

where $\hat{v} \in \mathbb{R}^n$ is the representative velocity [19], T_k is the size of the time step with k denotes the step size. We determine the internal force ψ implicitly (i.e., related to \hat{v}_k) with linearized form based on the potential action model proposed in [19], to preserve the stability and reliability. The time index k is omitted from here for brevity, and when representing the next time step, superscript $+$ will be marked.

If the object is deformable, the state q will be the stack of the positions of the nodes on the object as $q = [x_1^T, \dots, x_{n_s}^T]^T \in \mathbb{R}^{3n_s}$, where n_s being the number of the nodes. On the other hand, if the object is rigid, the state can be written with the position and the rotation of a designated body coordinate.

B. Contact Models

The constraints that the contact force must meet can be expressed with various models. The Signorini-Coulomb model is widely used in physics simulation where high physical fidelity is required [13]. It can be formulated with equality constraints and complementarity conditions as:

$$0 \leq g(x_i^+) \perp \lambda_i^{(n)} \geq 0 \quad (1)$$

$$0 \leq \phi_i \perp \mu \lambda_i^{(n)} - \|\lambda_i^{(t)}\| \geq 0 \quad (2)$$

$$\lambda_i^{(n)} J_{c,i}^{(t)} \hat{v} + \phi_i \lambda_i^{(t)} = 0 \quad (3)$$

where $i = \{1, \dots, m\}$ denotes the index of each contact, $g(x)$ is the gap function between a point $x \in \mathbb{R}^3$ and the environment, the superscripts $(n), (t)$ mean the normal and tangential components respectively and $\phi_i \in \mathbb{R}$ is the auxiliary variable.

The outcomes of the Signorini-Coulomb conditions (1)-(3) can be classified into three contact states - 1) open, 2) stick, and 3) slip due to the complementarity, and the switching between each state is non-smooth.

The penalty-based contact model, in contrast to the Signorini-Coulomb condition, converts the contact constraints with compliance by introducing penalty force for the

constraint violation. The contact constraint for the penalty method can be formalized as follows:

$$\lambda_i^{(n)} = k_n \max(-g(x_i^+), 0)$$

$$\lambda_i^{(t)} = -\min(k_n \|J_{c,i}^{(t)} \hat{v}\|, \mu \lambda_i^{(n)}) \frac{J_{c,i}^{(t)} \hat{v}}{\|J_{c,i}^{(t)} \hat{v}\|}$$

where k_n is the penalty coefficient.

III. UNINFORMATIVE GRADIENT OF DIFFERENTIABLE SIMULATION

A. Non-smoothness of the Contact Condition

The inherent non-smoothness of contact dynamics is one of the reasons why gradients of differentiable simulations are not useful. Due to this non-smoothness, an informative gradient can be provided only within an area where switching does not occur. For example, if an object in a free space (i.e. contact is not detected), the gradient from simulation cannot provide information that the contact may be generated when it approaches to the ground. In other words, smooth contact conditions without such non-smoothness is required to obtain useful gradients. First, let us examine the two factors of non-smoothness.

1) *Discontinuous Change of Contact Feature*: In detection-based dynamic simulations, which conduct collision detection every time step, the contact features (i.e. contact point, normal, and the number of contact) vary discretely over time. These changes, as also mentioned in [12], generate discontinuities and reduce the reliability of the gradient.

2) *Non-smooth Transitions between Contact States*: The second factor of the non-smoothness is the switching between the contact states - stick, slip, and open [8]. Several relaxation strategies have been proposed to smooth this out. The work of [11] utilized perturbed Signorini condition to relax complementarity by introducing the central path parameter, and log barrier function is used for smoothing in [18]. Also, it can be smoothed by utilizing penalty function as in [2].

Perturbed Signorini:	$\lambda_i^{(n)} = \sigma/g$
Smooth penalty:	$\lambda_i^{(n)} = -k_n \text{smax}(-g, 0)$
Log barrier:	$\lambda_i^{(n)} = -\min(g - \hat{g}, 0)^2 \ln(\frac{g}{\hat{g}})$

(4)

where $g(x_i^+)$ is briefly written as g , \hat{g} is the target distance [18], $\text{smax}(\cdot)$ is the smooth maximum function. In this paper, we will use the smooth maximum unit [20] as one of the smooth maximum function for the smooth penalty method. The comparison between the contact models is presented in Fig. 1a.

B. Jittering in the Dynamics of Smoothed Contact Models

The non-smoothness of the contact conditions can be eliminated by solving the above two factors. However, when the above smoothing schemes are applied to successive

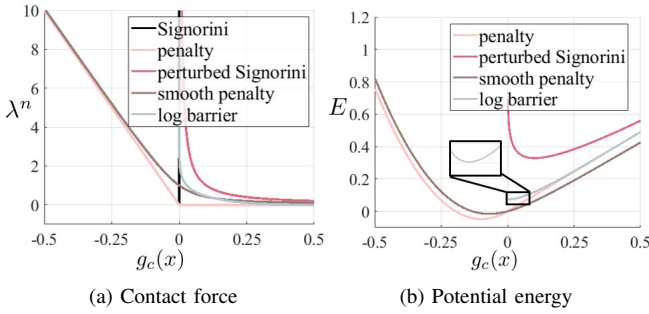


Fig. 1: Comparison of the contact models and the potential energy (6) defined from each contact models.

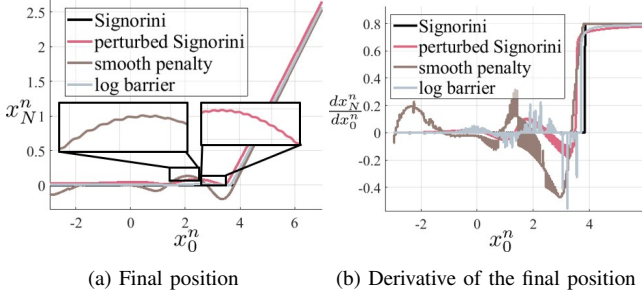


Fig. 2: Comparison of the final position and the gradient about initial position of a free-falling particle when apply each contact model

time steps, jittering may occur, resulting in excessive local minima, as shown in Fig. 2.

The reason why jittering occurs can be explained from an energy perspective. Let us assume a frictionless object and a flat environment. The indefinite integral of the smooth contact force (4) can be interpreted as potential energy, since the energy depends only on the position of the object, and not by the path traveled in moving from one position to other [21]. Then the total energy E of the system can be written as:

$$E = \dot{q}^T M \dot{q} + U \quad (5)$$

$$U := \int_{x_0}^x (-\lambda_c(g(x))) \cdot dx - G \cdot dq \quad (6)$$

where x_0 is a datum point, U is the total potential energy including gravitational energy, which shape is shown in Fig. 1b. Then the total energy is preserved, and an oscillation occurs around the local minima of the potential energy [21]. Fig. 2 shows the result of the point mass simulation with the contact models. Such jittering is even exacerbated in the gradient (Fig. 2b), necessarily degenerate the solvability of optimization problem when utilizing the gradient.

To alleviate this jittering, damping can be added to the dynamics. However, a significant quantity of damping is necessary to suppress the oscillation, and the object can not move dynamically. To avoid this, for the penalty method, damping can be added only when penetration happens [8]. However, there are also problems with this method: 1) the process of detecting penetration is discontinuous, resulting



Fig. 3: Comparison between penalty method with damping on dynamics (yellow), proposed method (blue), and position-based Signorini condition (pink). The damping coefficient is chosen to be critically damped with $k_n = 1000$, and damping parameter selected as $\alpha = 1.5$ for proposed method.

in discontinuity; 2) when penetration occurs, the movement is suppressed by strong damping and cannot move, thereby rolling or lifting from the floor cannot be implemented well. Fig. 3 is a snapshot of the Stanford bunny simulation with an initial angular velocity, and it can be seen that the bunny cannot roll.

C. Summary

The gradient offered from simulations are not typically informative due to the non-smoothness of contact dynamics. Some differentiable simulators have suggested different contact models to solve the issue of the non-smooth transitions between the states (Sec. III-A.2). However, the problems that the contact features change discontinuously (Sec. III-A.1), or jittering may exist (Sec. III-B) have not been investigated. In this paper, we aim to obtain a smooth contact model without such discontinuous or uneven changes by solving the problem of discontinuous change of contact features and jittering.

IV. PROPOSED DIFFERENTIABLE CONTACT SIMULATION

A. Differentiable Contact Feature

The contact Jacobian J_c consists of a *contact point mapping* $J_p \in \mathbb{R}^{3m \times n}$ that maps a state to contact points and a *rotation mapping* $R_c \in \mathbb{R}^{3m \times 3m}$ that rotates the frame to the local coordinates as:

$$J_c = R_c J_p \in \mathbb{R}^{3m \times n}$$

$$R_c = \begin{bmatrix} R_{c,1} & \cdots & 0 \\ 0 & \ddots & 0 \\ 0 & \cdots & R_{c,m} \end{bmatrix}, J_p = \begin{bmatrix} J_{p,1} \\ \vdots \\ J_{p,m} \end{bmatrix}$$

where $J_{p,i}$ is the Jacobian matrix for i -th contact point, $R_{c,i}$ is $SO(3)$ matrix that transforms coordinate to the contact frame $\{c_i\}$ from global frame $\{g\}$, with m being the number of the contact. The contact point mapping is dependent on the contact point x_i , and the rotation mapping is dependent on the contact normal vector n_i .

These contact features, however, vary discretely as the contact occurs and vanishes discontinuously. Even the number of the contacts m , which determines the dimension of the contact condition changes discretely. To this end, we suggest 1) invariant contact set to make the number of contact point constant and 2) SDF-based coordinate differentiation to calculate the derivative the rotation mapping.




Fig. 4: Schematic illustration of the invariant contact mapping using the collision proxies and SDF-based coordinate differentiation. (Purple dots : collision proxies)

1) *Invariant Contact Set*: The key idea is to pre-specify collision proxies on the surface of an object as illustrated in Fig. 4, and apply the contact condition (1)-(3) for all proxies in the contact solving process. When using a detection-based method, the points with penetration is selected every time step and contact condition is enforced to the detected points. In this selection process, a discontinuity occurs, and the mapping matrix J_p changes discretely. Our method, on the other hand, does not require selection because contact conditions are enforced to all collision proxies and the mapping matrix is constant.

For the case of soft objects, we can define the collision proxies as the mesh nodes. Then the contact mapping becomes a selection matrix for the collision proxies, for example:

$$J_{p,i} = [0_{3 \times 3} \quad \cdots \quad I_{3 \times 3} \quad \cdots \quad 0_{3 \times 3}] \in \mathbb{R}^{3 \times n}$$

which makes the contact point mapping J_p constant.

Unlike deformable objects, rigid objects are expressed by the position and the rotation of the origin of the body coordinate. Therefore, the mapping to a point on the body becomes nonlinear. However, if the rotation of the rigid object is represented with all nine components of the rotation matrix [22], the contact point mapping becomes constant as:

$$q = [x_r^T \quad R_x^T \quad R_y^T \quad R_z^T]^T \in \mathbb{R}^{12}$$

$$J_{p,i} = \begin{bmatrix} r_{c,i}^T & 0_{1 \times 3} & 0_{1 \times 3} \\ I_{3 \times 3} & 0_{1 \times 3} & r_{c,i}^T & 0_{1 \times 3} \\ 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & r_{c,i}^T \end{bmatrix} \in \mathbb{R}^{3 \times 12}$$

where $x_r, R_x, R_y, R_z \in \mathbb{R}^3$ are the origin and the three basis vectors of the body coordinate, and $r_{c,i} \in \mathbb{R}^3$ is the position of the i -th collision proxy represented in the body coordinate, which is constant, as shown in Fig. 4. The contact point mapping for a rigid object also is constant for every time step, meaning that it is invariant, and is also exact (i.e., no linearization error).

2) *SDF-based Coordinate Differentiation*: To differentiate the local contact coordinate, we express the environment as a SDF. Then, the basis vectors R_i can be written with the

first-order derivative of the SDF g as:

$$n_i = \nabla g(x_i) \in \mathbb{R}^3$$

$$[t_i^1 \quad t_i^2] = \text{null}(n_i^T) \in \mathbb{R}^{3 \times 2}$$

$$R_i = [t_i^1 \quad t_i^2 \quad n_i]^T \in \text{SO}(3)$$

where n_i, t_i^1 and t_i^2 are the normal and two tangential basis of the local contact coordinate each, and g is the SDF function. Thus, the first derivative of R_i can in turn be written by using the second derivative of SDF g .

A derivative of a rotation matrix can be written as an multiplication of skew matrix of an screw axis. Let $w_{\xi,i}$ be the screw axis for the derivative of R_i about a parameter or state ξ :

$$\frac{d(R_i)^T}{d\xi} = S(w_{\xi,i})(R_i)^T \quad (7)$$

Then, what we need to calculate is the axis $w_{\xi,i}$. For the normal vector n_i , the differentiation (7) can be rewritten as:

$$\frac{dn_i}{d\xi} = w_{\xi,i} \times n_i \quad (8)$$

where, $n_i := \nabla g_i, \frac{dn_i}{d\xi} = \nabla_\xi \nabla g_i$

Since n_i and $\frac{dn_i}{d\xi}$ can be calculated from the SDF, from (8), $w_{\xi,i}$ can be obtained as follows:

$$w_{\xi,i} = n_i \times \frac{dn_i}{d\xi} + \beta n_i \quad \forall \beta \in \mathbb{R}$$

where β can be chosen arbitrarily, which can be proved from the fact that $J_c^T \lambda$ is invariant of the tangential basis. In this paper, β is chosen as zero. Then the derivative of R_i can be obtained as:

$$\frac{dR_i}{d\xi} = \begin{bmatrix} (w_{\xi,i} \times t_i^1)^T \\ (w_{\xi,i} \times t_i^2)^T \\ \nabla_\xi \nabla g_i(x_i) \end{bmatrix}$$

With the calculated derivative of the contact coordinate, the more reliable gradient of the simulation can be obtained.

Therefore, the contact point mapping function can be differentiated as 1) the contact set is invariant and 2) the local contact coordinate R_c can be differentiated.

B. Damped Contact Force

With the invariant contact mapping and smooth contact conditions (4), the non-smoothness of the contact dynamics can be eliminated. However, as pointed in Sec. III-B, jittering of the contact dynamics occur when smooth contact force is utilized in conjunction of position-based Signorini condition. Since jittering may induce excessive local minima, the efficiency of optimization problems may be degenerated.

In real physics, the energy (5) is dissipated at the moment of the contact, but when a smooth contact force is used, the contact force functions as a conservative force, and the oscillation takes place. Based on this analysis, we propose a novel contact model named damped contact force that

dissipates energy to achieve beneficial gradients without jittering:

$$0 \leq \underbrace{g^n(x_i^+) + \alpha J_{c,i}^n \hat{v} T \perp \lambda_i^n}_{=:g_d(x_i^+, \hat{v})} \geq 0 \quad (9)$$

where $\alpha \geq 0$ is a damping parameter. If $\alpha = 0$, the proposed contact model is equivalent with the original Signorini-Coulomb condition. If α goes to infinity, then the object will not fall down and stop in the air, as the energy is excessively dissipated. To have an intuitive grasp of the damped contact force, let us assume the SDF $g(\cdot)$ is a linear mapping implying the environment is flat. The contact condition can therefore be rewritten as:

$$(x_i^n)^+ \geq \frac{\alpha}{\alpha + 1} x_i^n$$

The contact force bounds the normal directional position of the collision proxy above $\alpha x_i^n / (\alpha + 1)$, while Signorini-Coulomb condition bounds it above zero. If the object is moving away from the contact point or the velocity of approaching is not sufficient, the damping force is not generated, which makes a difference from adding the damping explicitly in the dynamics. As mentioned above, explicit damping causes the problem of being stuck in the floor, whereas the proposed method solves the problem and is able to implement rolling (Fig. 3).

In conjunction with the damped contact force and (4), we can smooth the proposed damped contact condition as:

$$\begin{aligned} \text{Perturbed Signorini:} \quad & \lambda_i^n = \sigma^n / g_d \\ \text{Penalty:} \quad & \lambda_i^n = -k_n \text{smax}(-g_d, 0) \end{aligned} \quad (10)$$

where $g_d(x_i^+, \hat{v})$ is abbreviated as g_d . With the damped contact force, the total energy defined in (5) is dissipated during contact. The energy change can be written as follow:

$$\begin{aligned} \dot{E} &= \dot{x}^T (M\ddot{x} - J_c^n \lambda(g(x)) - G) \\ &= \dot{x}^T (J_c^n)^T (\lambda(g(x)) + \alpha J_c^n \dot{x} - \lambda(g(x))) \\ &\approx \alpha \|J_c^n \dot{x}\|_{\nabla \lambda}^2 < 0 \end{aligned}$$

Since $\lambda(\cdot)$ is a decreasing function as shown in Fig. 1a, \dot{E} is negative, energy dissipation occurs, and as the value of α increases, the amount of the dissipation increases. In the case of the penalty function, $\nabla \lambda = -k_n$, therefore $\dot{E} = -\alpha k_n \|J_c^n \dot{x}\|^2$. With the damped contact force, the energy dissipates which suppresses the oscillation. Finally, with the damped contact forces and smooth contact conditions (i.e., invariant contact set), smooth, not jittery, and informative gradients can be obtained.

V. EVALUATIONS

A. Comparison of the Cost Landscape: Rigid Box Flipping

In order to verify that the dynamics are effectively smoothed with our framework, an ablation study of the landscape of the cost with a box flipping task is conducted. Fig. 5 shows the simulation snapshots of the task. The angle and height of lifting the floor are parameterized, and the norm of the difference between the final pose and the initial




Fig. 5: Snapshots of the rigid box flipping simulation.




Fig. 6: Comparison of the cost and gradient of the cost of rigid box flipping when applying each contact model. model 1: without invariant contact set, model 2: without damped contact force (perturbed complementarity), model 3: without damped contact force (smooth penalty).

pose is set as cost. The simulation is performed for 0.8 s, the size of the time step is 10 ms, and the damping parameter α is chosen to be 9.

The landscape of the cost and the derivative of the cost about the parameter is presented in Fig. 6. As shown in the figure, when detection is conducted, which means that the contact set is not invariant (model 1), the cost itself is jittery (i.e. varying contact mapping, perturbed Signorini, no damped contact force). Also without our damped contact force methods, the cost and the gradient of both contact models are shown to be jittery (model 2, model 3). On the other hand, with the proposed method, jittering is suppressed and a smooth landscape is induced.

B. Comparison of the Reliability: Rigid Ball Rolling

For a gradient to be reliable, it must be possible to accurately obtain the differentiation for an arbitrarily cost using the gradient. To show the reliability, we compare the exact cost and the reconstructed cost calculated by numerically integration using the gradient. The reconstructed cost must be equal to the exact cost.

We present a simulation of a rigid ball rolling inside a sphere, and the cost is the distance from the final pose to the goal pose. For intuitive comparison via 2D graph, we parameterize initial position and velocity with one-dimensional parameter p . The following equation shows how




Fig. 7: Comparison of the cost and gradient of the cost of rigid box flipping when applying each contact model. Proposed model : $\alpha = 0.25$, model 1: without SDF-based coordinate differentiation, model 2: without invariant contact set.




Fig. 8: Result of trajectory optimization in cloth flipping scenario with the proposed framework ($\alpha = 0.3$). The goal is to place the edges marked with the green balls at the goal positions marked with the purple balls. Position control is conducted on the nodes with red balls.

the reconstructed cost is calculated:

$$\begin{aligned} \mathcal{O}(p_k) &= \|q_N - q_g\|^2 \\ \frac{d\mathcal{O}}{dp} \Big|_{p=p_k} &= 2(q_N - q_0)^T \frac{dq_N}{dp} \Big|_{p=p_k} \\ \mathcal{O}_r(p_k) &= \mathcal{O}_r(p_{k-1}) + \frac{d\mathcal{O}}{dp} \Big|_{p=p_k} (p_k - p_{k-1}) \end{aligned}$$

where $q_0, q_N, q_g \in \mathbb{R}^6$ are a initial, final, goal pose respectively, $\mathcal{O}(p_k)$ is the exact cost with parameter p_k , and $\mathcal{O}_r(p_k)$ is the reconstructed cost using the calculated gradient $\frac{d\mathcal{O}}{dp}$.

Compare the exact cost and the reconstructed cost obtained from the simulator using different contact models. As shown in Fig. 7, with our proposed contact model, the reconstructed cost \mathcal{O}_r well match with the exact cost \mathcal{O} . On the other hand, as the environment is not flat, without contact coordinate differentiation (model 1), the reconstructed cost significantly differs from the exact cost. Additionally, without the invariant contact point set (model 2), discontinuity due to the detection remains, and the cost jitters significantly as shown in the figure. In addition, the gradient achieved without the invariant contact point mapping does not match the cost.

C. Trajectory Optimization: Cloth Manipulation

The derivative of the dynamics can be to solve trajectory optimization or optimal control effectively by enabling gradient-based optimization. Here, we consider a trajectory optimization that optimizes control input for flipping the cloth. We construct the optimization using the direct shooting

method as:

$$\begin{aligned} \min_u \sum_{k=1}^{N_s} l_k(x_k, u_k) + h(x_{N_s}, u_{N_s}) \\ \text{s.t. } x_{k+1} = f(x_k, u_k) \end{aligned}$$

where N_s is the length of the time step, u_k is the control input, l_k is the running cost function of k -th time step, h is the terminal cost function, and f is the dynamics with contact conditions. We adopt the quasi-newton method for our cloth manipulation. The cloth is modeled with the finite element method (FEM) with plane stress [23]. The flipping simulation is conducted for 0.8 s, and the length of the time step is 10 ms. The objective function is defined as a distance to the goal position of two selected nodes and a roughly tuned nominal trajectory is given.

As can be seen from Fig. 8, the cost is well reduced during optimization and the cloth is successfully flipped. Furthermore, we compare the optimization results when the damped contact force/invariant contact set is not used. We test by changing the initial trajectory/goal position and averaging the result. The damping parameter α is chosen to be 0.3. The mean cost after 25 iterations is in Table I, and the result validates the efficacy of our reliable gradient on the actual optimization problem.

Model	Proposed	w.o. damped contact force	w.o. invariant contact set
Cost	0.0081	0.0125	0.0311

TABLE I: Comparison between average final cost of proposed framework and without damped contact force/invariant contact set.

VI. CONCLUSION

We present a novel differentiable simulator with invariant contact point mapping and damped contact force. With the proposed methods, the non-smoothness of the contact dynamics is smoothed, and a reliable and informative gradient can be obtained.

Collisions in this article are restricted to collisions between dynamics objects and static objects, and cannot handle self-collisions or collisions between dynamics objects. However, by representing each tetrahedral mesh as a set of SDFs, it will be able to deal with contact in various situations with the proposed methods. Furthermore, it increases simulation time since all collision proxies are considered as contact points. This can be solved by modifying the contact condition modeling to enable contact poing culling. From this, some future improvement directs include: 1) contact conditions that can be applied to self-collisions, 2) fast forward / backward simulation of the differentiable simulation. Furthermore, the degree of smoothing is affected by the smoothing parameter α for the damped contact force method. It would be useful to establish a criterion for determining this parameter based on the optimization problem and scenario of the simulation.

REFERENCES

- [1] S. Le Cleac’h, T. Howell, M. Schwager, and Z. Manchester. Fast contact-implicit model-predictive control. *arXiv preprint arXiv:2107.05616*, 2021.
- [2] J. Carius, R. Ranftl, V. Koltun, and M. Hutter. Trajectory optimization with implicit hard contacts. *IEEE Robot. Autom. Lett.*, 3(4):3316–3323, 2018.
- [3] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard. Crocodyl: An efficient and versatile framework for multi-contact optimal control. In *IEEE Int. Conf. Robot. Autom.*, pages 2536–2542, 2020.
- [4] Q. Le Lidec, I. Kalevatykh, I. Laptev, C. Schmid, and J. Carpentier. Differentiable simulation for physical system identification. *IEEE Robot. Autom. Lett.*, 6(2):3413–3420, 2021.
- [5] J. Degraeve, M. Hermans, J. Dambre, et al. A differentiable physics engine for deep learning in robotics. *Frontiers in neurorobotics*, 13:6, 2019.
- [6] Y. Hu, J. Liu, A. Spielberg, J. B. Tenenbaum, W. T. Freeman, J. Wu, D. Rus, and W. Matusik. Chainqueen: A real-time differentiable physical simulator for soft robotics. In *IEEE Int. Conf. Robots. Autom.*, pages 6265–6271, 2019.
- [7] Y. Hu, L. Anderson, T. M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand. DiffTaichi: Differentiable programming for physical simulation. In *Int. Conf. on Learning Representations*, 2020.
- [8] M. Geilinger, D. Hahn, J. Zehnder, M. Bächer, B. Thomaszewski, and S. Coros. Add: analytically differentiable dynamics for multi-body systems with frictional contact. *Transactions on Graphics*, 39(6):1–15, 2020.
- [9] J. Carpentier and N. Mansard. Analytical derivatives of rigid body dynamics algorithms. In *Robotics: Science and systems*, 2018.
- [10] J.P. Sleiman, J. Carius, R. Grandia, M. Wermelinger, and M. Hutter. Contact-implicit trajectory optimization for dynamic object manipulation. In *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 6814–6821, 2019.
- [11] K. Werling, D. Omens, J. Lee, I. Exarchos, and C. K. Liu. Fast and feature-complete differentiable physics engine for articulated rigid bodies with contact constraints. In *Robotics: Science and Systems*, 2021.
- [12] Y. Li, T. Du, K. Wu, J. Xu, and W. Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Trans. Graph.*, 42(1):1–20, 2022.
- [13] T. Du, K. Wu, P. Ma, S. Wah, A. Spielberg, D. Rus, and W. Matusik. Diffpd: Differentiable projective dynamics. *ACM Trans. Graph.*, 41(2):1–21, 2021.
- [14] R. Antonova, J. Yang, K. M. Jatavallabhula, and J. Bohg. Rethinking optimization with differentiable simulation from a global perspective. In *6th Annual Conference on Robot Learning*.
- [15] H. J. Suh, M. Simchowitz, K. Zhang, and R. Tedrake. Do differentiable simulators give better policy gradients? In *Int. Conf. Machine Learning*, pages 20668–20696. PMLR, 2022.
- [16] J. Xu, V. Makoviychuk, Y. Narang, F. Ramos, W. Matusik, A. Garg, and M. Macklin. Accelerated policy learning with parallel differentiable simulation. In *Int. Conf. on Learning Representations*, 2021.
- [17] L. Metz, C. D. Freeman, S. S. Schoenholz, and T. Kachman. Gradients are not all you need. *arXiv preprint arXiv:2111.05803*, 2021.
- [18] M. Li, Z. Ferguson, T. Schneider, T. R. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D. M. Kaufman. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM Transactions on Graphics*, 39(4):49, 2020.
- [19] M. Kim, Y. Lee, Y. Lee, and D.J. Lee. Haptic rendering and interactive simulation using passive midpoint integration. *Int. J. Robot. Res.*, 36(12):1341–1362, 2017.
- [20] K. Biswas, S. Kumar, S. Banerjee, and A. K. Pandey. Smu: smooth activation function for deep networks using smoothing maximum technique. *arXiv preprint arXiv:2111.04682*, 2021.
- [21] P. A. Dourmashkin. *Classical mechanics: Mit 8.01 course notes*. Wiley Custom Learning Solutions, 2014.
- [22] L. Lan, D. M. Kaufman, M. Li, C. Jiang, and Y. Yang. Affine body dynamics: Fast, stable & intersection-free simulation of stiff materials. *arXiv preprint arXiv:2201.10022*, 2022.
- [23] P. Volino, N. Magnenat-Thalmann, and F. Faure. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Trans. Graph.*, 28(4):Article–No, 2009.