Comparing Effectiveness of Relaxation Methods for Warm Starting Trajectory Optimization through Soft Contact

H.J. Terry Suh[1], Yanwei Wang[1]

***Abstract*— Planning through contact is a very difficult problem due to the non-smooth nature of contact. As a result, most methods must make a tradeoff between physical accuracy and difficulty of trajectory optimization. To mitigate this issue, we consider the framework of using indirect relaxations (methods that may not be physically accurate but have reliable and tractable solutions) to warm start direct methods (methods that respect the physics). We consider a regularized model of soft contact using Softplus as well as common techniques for relaxing complentarity constraints, and feed them into exact LCP formulation as well as single-shooting with gradient descent with ReLU representation of contact dynamics. Our results surprisingly show that even for a very simple problem, gradient descent with ReLU dynamics work tractably, reliably, and even yield lower costs compared to other methods.** [1]

I. Introduction

Contact is a crucial part of robot dynamics, as it is the most important mechanism by which robots interact with the environment. As such, the problem of planning through dynamics involving contact is widely recognized in the field of robotic manipulation [1], [2] and locomotion [3], [4]. Despite the importance of this problem, the difficulty of solving the problem tractably and reliably is astounding. The fact that contact forces exist if and only if two objects are in contact creates a notion of hybridness in the dynamics, rendering conventional trajectory optimization methods for smooth dynamics invalid. We refer to this nature of contact as *The Fundamental Nature of Contact*.

Definition. (The Fundamental Nature of Contact). Contact forces between two objects are nonzero if and only if the objects are geometrically in contact.

When solving trajectory optimization with contact dynamics, we are often faced with a tough tradeoff: we can either respect the fundamental nature of contact and deal with the resulting difficulty of trajectory optimization, or we can create a smooth approximation of contact to make the optimization problem easier, but violate the fundamental nature of contact. In this work, we label the former as *direct methods*, and the latter as *indirect methods*.

Direct methods obey the fundamental nature of contact by construction, but are often very difficult class of optimization problems. Explicitly modeling the contact mode as integer variables in the presence of nonlinear dynamics leads to Mixed-Integer Non-linear Porgramming (MINLP), which can be extremely hard and time-consuming to solve. If the dynamics are linear, the trajectory optimization problem can be transcribed as an instance of Mixed-Integer Programming (MIP) for Piecewise-Affine Dynamics which can be efficiently solved [5]. However, these techniques are often confined to linear systems, as linearizing hybrid dynamics along a trajectory can lead to large numerical approximation errors around the guard surface.

Without modeling contact modes as integer variables, the dynamics imposed by contact can be modeled into Complentarity Constraints, leading to Mathematical Programs with Complementarity Constraints (MPCC). Although some works have been successful in solving these problems with Sequential Quadratic Programming (SQP) [6], [7], [8], complementarity constraints still remain as very hard constraints to satisfy with conventional solvers. As a result, such methods still remain very time-consuming to solve, and can be unreliable in producing good solutions in practice.

Finally, a family of single-shooting methods have been applied to contact trajectory optimization through single-shooting with gradient descent or iterative Linear Quadratic Regulators (iLQR) [9], which are convenient to implement due to the fact that the cost can be formulated a function of input histories, and the gradients may be computed with backpropagation. Apart from the known numerical problems existent in single-shooting, the gradient of nonsmooth contact dynamics may be hard to define at the guard surface, and linearizing trajectories in hybrid dynamics again suffer from the aforementioned approximation.

Yet, recent advances in parameter identification of deep neural networks using ReLU networks suggest that the ill-defined gradient at the guard surface may not be as big of a practical concern as originally thought of. Thus, single shooting with gradient descent has been applied to control over learned dynamics [9], [10] quite successfully, although most of the prior work has not yet attempted to perform gradient descent on an analytical expression for contact forces using penalty methods.

Indirect methods often use a smooth and regularized contact model, and therefore make the resulting optimization problem easier [11], [12]. However, the resulting formulation may violate the fundamental nature of contact. Because the resulting dynamics are smooth under smooth contact models, conventional trajectory optimization methods such as single shooting, direct collocation, differential dynamic programming (DDP) may be utilized.

We note importantly that indirect methods do not always consider smooth models of contact. For instance, a relaxation

[1]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, hjsuh@mit.edu, felixw@mit.edu

[1]Project Codebase: https://github.com/hjsuh94/soft_contact

of complementarity constraint by a bilinear inequality constraint with slack variables [6], [7], [8] can allow forces from a distance and make the optimization problem easier compared to strict complementarity constraints. Such methods can also be classified as indirect, as the problem is allowed to violate the fundamental nature of contact.

From a practical point of view, one might question how important it is that our trajectory respects the fundamental nature of contact. Indeed, one might argue that this is a modeling problem that is apparent in every engineering task. However, violating the fundamental nature of contact is different from that of truncation error or parameter error, which smoothly grows in time. For instance, consider a task where a two finger gripper must lift up a cup. If the trajectory obtained by a method allows forces from a distance, the robot may try to lift up the cup when not in contact, completely invalidating the solution. This strongly suggests that regardless of the method, the final trajectory *must* respect the fundamental nature of contact throughout its entire timestep.

Therefore, we seem to be at an impasse: we cannot obtain fast and reliable methods that obeys the crucial physics of contact at the same time. How can we possibly overcome this impasse? We believe that the key solution lies in combining the benefits of both direct methods and indirect methods by iteratively tuning the solution of an indirect method to obey the fundamental nature of contact. A powerful framework for doing so comes from warm starting a direct method with the solution of an indirect method. (i.e. providing the solution of an indirect method as an initial guess to a direct method). Similar strategies were taken in prior work involving trajectory optimization with contact [11].

A natural question that follows is: which combination of indirect methods and direct methods provide the best performance? From an optimization perspective, it would be plausible that better initial guesses makes convergence much faster, especially if the initial guess is closer to the feasible set of a difficult constraint. Understanding the synergy of a particular indirect method and a direct method may reveal a way to synthesize better indirect methods, such as regularized contact models or relaxation of constraints, that serve as better initial guesses for direct methods. Thus, our work seeks to understand the strength of relaxations that particular indirect methods have on the direct methods.

In our work, we compare two class of indirect methods: using traditional trajectory optimization methods (direct collocation and single shooting) for softplus approximation of smooth contact, and LCP relaxation through bilinear inequality constraints. We study how these methods affect warm starting of direct methods: exact LCP, and single shooting with gradient descent with contact forces represented as ReLU functions. Surprisingly, our empirical study showed that single-shooting gradient descent methods achieve very tractable and physically accurate performances, especially when warm-started with regularized contact models. On the other hand, LCP methods and their relaxations often tend to be unreliable and costly.

II. Preliminaries

A. Time-Stepping Semi-Implicit Integration

The equations of motions of body configurations $q \in \mathbb{R}^n$ and velocities $v \in \mathbb{R}^n$ subject to external forces $\lambda \in \mathbb{R}^c$ can be described using Lagrangian dynamics in continuous form:

$$\mathbf{M}(q)\dot{v} + \mathbf{C}(q, v)v = \mathbf{B}(q)u + \mathbf{J}(q)^T\lambda, \tag{1}$$

where $\mathbf{M}$ is the positive definite generalized inertia matrix, $\mathbf{C}$ the Coriolis matrix, $\mathbf{G}$ the gravitational terms, $\mathbf{B}(q)$ the actuation matrix, and $\mathbf{J}$ the Jacobian term resulting from linearizing forward dynamics. Throughout the work, we also utilize the notation x to denote the Markovian state of the system, $x = [p, v]^T \in \mathbb{R}^{2n}$.

We discretize the equations of motions in time using semi-implicit time-stepping, as used in [13], due to its accuracy in modeling energy conservation as well as convenience in forward simulation. Given a timestep of h, the above equations are discretized as

$$\begin{cases} q_{t+1} - q_t = hv_{t+1} \\ \mathbf{M}_t(v_{t+1} - v_t) = -h\big(\mathbf{C}_t v_t - \mathbf{B}_t u_t + \mathbf{J}_t\lambda_t\big), \end{cases} \tag{2}$$

where the subscripts denotes time. The subscript on the matrices denote the fact that the matrices are calculated with respect to values of quantities at time t (e.g. $\mathbf{M}_t = \mathbf{M}(q_t)$).

B. Soft Contact Approximation

Before optimizing trajectories with contact dynamics, obtaining the contact forces themselves during forward simulation may sometimes be an ill-posed problem due to non-uniqueness such as static interminancy [13]. Other works have therefore solved linear complementarity problems (LCP) at each timestep to identify the value of contact forces [13], implicitly defined contact forces inside the optimization constraint [6], or solved QPs to find out the minimum value of contact forces for feasible dynamics [12].

In contrast, we choose to approximate contact with a spring-like soft-contact model, which facilitates much easier forward simulation [14], [15]. This comes at the cost of stiffness in the resulting ODE during forward simulation, but the formulation can remain general enough to be applied to physically soft contact.

Given the penetration depth between two bodies $\phi(q)$, a signed distance function, we approximate the contact forces between the two bodies as

$$\lambda_t = k \max(\phi(q_t), 0) = k \cdot \text{ReLU}(\phi(q_t)), \tag{3}$$

where k is the stiffness of the contact, and ReLU is the rectified linear unit. Damping can also be added to this formulation to create a viscoelastic model, such as those motivated by Hertzian contact models [15].

C. Trajectory Optimization with Contact

As usual, we begin by formulating a cost function in Bolza form, defined on the trajectory $\mathcal{T}$ with timesteps of integer indices, $t \in [0, T]$:

$$J(\mathcal{T}) = x_T^T \mathbf{Q}_T x_T + \sum_{t=0}^{T-1} (x_t^T \mathbf{Q}_t x_t + u_t^T \mathbf{R}_t u_t). \quad (4)$$

We formulate a trajectory optimization problem as finding a sequence of states $x_0, x_1, \cdots, x_T$ and inputs $u_0, u_1, \cdots, u_T$ to minimize the cost while obeying the time-discretized dynamics in (2).

III. Indirect Methods

We first discuss two class of indirect methods that facilitate solving the optimization problem at the cost of violating the fundamental nature of contact: contact smoothing with softplus, and LCP relaxation into bilinear inquality constraints.

A. Smooth Contact Modeling

Following the proposed soft contact model in (3), we can additionally relax the ReLU function with a Softplus, which can be modeled as

$$\lambda_t = k \log_\varepsilon (1 + \varepsilon^{\phi(q_t)}), \quad (5)$$

which has the property that convergence to ReLU is guaranteed as $\epsilon \to 0$, and has an analytical derivative (logsitic function). Such form of soft smoothing contact has been utilized in [11], and has empirically shown good advantages as an indirect method. This model is visualized in Fig.2.

With this contact model, we formulate the trajectory optimization problem

$$\begin{aligned} \underset{x_t, u_t}{\text{minimize}} \quad & J(x_0, \cdots, x_T, u_0, \cdots, u_{t-1}) \\ \text{subject to} \quad & q_{t+1} - q_t = h v_{t+1}, v_{t+1} - v_t = h a_t, \\ & \mathbf{M}_t a_t = -h(\mathbf{C}_t v_t - \mathbf{B}_t u_t - \mathbf{J}_t^T \lambda_t), \\ & \lambda_t = k \log_\epsilon (1 + \varepsilon^{\phi(q_t)}), \\ & x_0 = x_i \quad x_T = x_f \qquad \forall t, \end{aligned} \quad (6)$$

which can be considered as a nonlinear direct collocation method with a fixed timestep. As the smooth contact constraint is a convex equality constraint relating penetration depth and contact force, this problem can converge quickly if the dynamics are linear with respect to the states and the contact force.

B. LCP Relaxation

We also consider forms of LCP relaxation methods considered in [6], [7], [8], by noting that the ReLU function in (3) can be encoded in the optimization problem using the following constraints:

$$\lambda_t = k \max(\phi(q), 0)) \iff \begin{cases} \lambda_t \geq 0 \\ \lambda_t - k\phi(q_t) \geq 0 \\ \lambda_t^T (\lambda_t - k\phi(q_t)) = 0. \end{cases} \quad (7)$$

The intuition behind these constraints can be obtained by observing the LCP constraints: either the contact forces are zero, at which there is no contact, or it obeys $\lambda_t = k(\phi(q_t))$, at which there is contact. Because the last constraint, known as complementarity, is often a difficult to impose, [6] proposed to relax the complementarity constraint with a bilinear inequality constraint, allowing the complementarity constraint to be positive:

$$\begin{cases} \lambda_t \geq 0 \\ \lambda_t - k\phi(q_t) \geq 0 \\ \lambda_t^T (\lambda_t - k\phi(q_t)) \leq \varepsilon. \end{cases} \quad (8)$$

The relaxation effectively allows forces at a distance, and allows iteratively tuning ϵ to meet the exact LCP constraints. This relaxation is visualized in Fig.2.

The full form of the trajectory optimization problem can be described as

$$\begin{aligned} \underset{x_t, u_t}{\text{minimize}} \quad & J(x_0, \cdots, x_T, u_0, \cdots, u_{t-1}) \\ \text{subject to} \quad & q_{t+1} - q_t = h v_{t+1}, v_{t+1} - v_t = h a_t, \\ & \mathbf{M}_t a_t = -h(\mathbf{C}_t v_t - \mathbf{B}_t u_t - \mathbf{J}_t^T \lambda_t), \\ & \lambda_t \geq 0, \quad \lambda_t - k\phi(q_t) \geq 0, \\ & \lambda_t^T (\lambda_t - k\phi(q_t)) \leq \epsilon, \\ & x_0 = x_i \quad x_T = x_f \qquad \forall t. \end{aligned} \quad (9)$$

However, the question of whether this is truly the right relaxation to consider remains. First, if we are going to allow forces from a distance anyways, why do we pose it as a nonconvex bilinear inequality constraint that are still difficult to impose? Given that convex relaxations of nonconvex bilinear constraints often involve spatial branch and bound through McCormick envelopes [16], this relaxation does not seem like big improvement over branch and bound through Mixed-Integer methods.

Second, does this method truly provide better initial guesses to the exact LCP formulation in (7)? Although the method seems intuitive as a relaxed formulation of (7), the comparative strength of this relaxation over more tractable contact models (such as softplus) has yet to be theoretically justified.

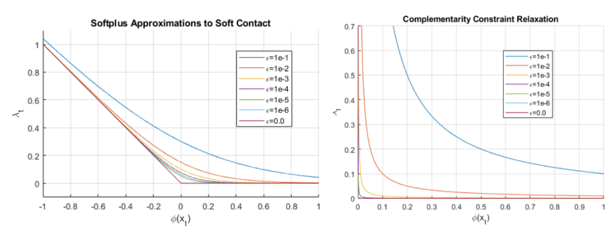

Fig. 1. Left: Softplus relaxation of penalthy contact forces. Right: LCP relaxation into bilinear inequality constraints.

IV. Warm Starting Direct Methods

Assuming that we have a trajectory $\mathcal{T}$ obtained from an indirect method, we warm start a direct method using this trajectory as an initial guess. We numerically study which of the two methods above are better at warm starting direct methods.

A. Warm Starting Exact LCP

The exact LCP formulation obeys the fundamental nature of contact in (7), as contact forces are zero if and only if penetration distance is positive. Thus, the trajectory optimization problem can be formulated as in (9) with $\varepsilon = 0$.

In order to study the effectiveness of relaxations in Sec.III, we formulate a metric that represents the worst-case violation of the exact LCP constraint, formulated as a function:

$$D_t = \max_{q_t} |\lambda_t \cdot (\lambda_t - k\phi(q_t))|. \quad (10)$$

Since the exact LCP formulation requires that $D(\lambda_t) = 0$, we can analyze the theoretical complexity of the behavior of $D(\lambda_t)$ as we decrease ε. Note that softplus satisfies the first two constraints of (8) as an outer approximation.

For the relaxed LCP, it is not hard to see that $D \in O(\varepsilon)$, as the maximum violation of $D(\lambda_t)$ is exactly ϵ. In the case of Softplus, the analytical form for $D(\lambda_t)$ can be lengthy, and arguing about its complexity class is not straightforward. Instead, we illustrate it with a numerical comparison in Fig.2. We observe that the bilinear inequality constraint of the relaxed LCP has a linear convergence rate of $\varepsilon \to 0$, while the convergence of Softplus initially descends faster than the LCP relaxation, and rapidly degrades as $\varepsilon \to 0$.

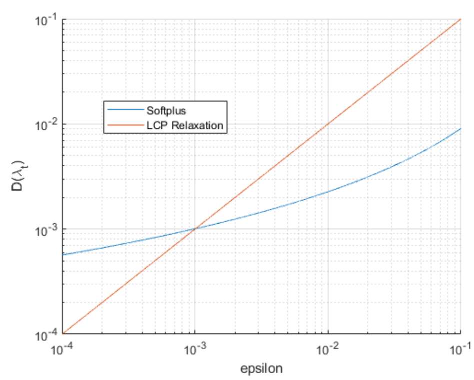

Fig. 2. Log-Lot plot of worst-case constraint violation of two relaxation schemes.

This lesson tells us two things: theoretically, the relaxed LCP has better convergence properties to the feasible set of exact LCP compared to softplus relaxations. Second, softplus might offer faster convergence during the initial stage of the optimization process if we begin the method with a relatively high tolerance.

B. Single Shooting with Gradient Descent (GD)

We also consider warm starting a single shooting method, and additionally fine-tuning it with with gradient descent. Unlike previous formulations, we directly formulate a symbolic expression for the state history as a function of input history, and try to minimize a cost function. Previous formulations achieved stability with the use of the terminal state constraint, while we aim to acehive stability by penalizing how much the final trajectory deviates away from the goal state in the cost function. Thus, the unconstrained optimization problem may be defined as

$$\underset{u_t}{\text{minimize}} \quad J(x_0, \cdots, x_T, u_0, \cdots, u_{t-1}), \quad (11)$$

where $x_1, \cdots, x_T$ are results of forward simulating the dynamics with the input sequence $u_0, \cdots, u_{T-1}$.

We solve this problem using gradient descent, which takes advantage of the fact that the dynamics imposed by contact can be formulated as ReLU functions, which are common in deep learning pipelines. The update rule is illustrated in (12).

$$u_t \leftarrow u_t - \eta \nabla_{u_t} J(u_t), \quad (12)$$

where η refers to the 'learning rate', or rate of descent. Similar ideas have been proposed in the field of MIP, where deep ReLU networks have been used as surrogate models of Mixed-Integer optimization [17].

To warm start this method, we take a trajectory from the indirect method and use the input trajectory of the indirect method as an initial guess, and perform gradient descent from the initial guess.

V. Case Study: 3 Cart System

To study the effect of warm starting different direct methods with indirect methods, we set up a toy problem with contact, motivated by [18], illustrated in Fig.3. The two carts on the left and right are actuated, while the cart in the middle is unactuated, and can only moved through contact. We also assume a global damping on the velocity terms.

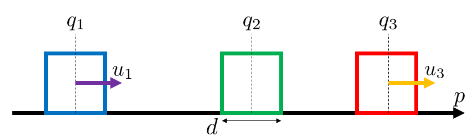

Fig. 3. Free Body Diagram of the three cart system

The forward dynamics of this system is given by

$$\begin{cases} q_1^{t+1} = q_1^t + hv_1^{t+1} \\ q_2^{t+1} = q_2^t + hv_2^{t+1} \\ q_3^{t+1} = q_3^t + hv_3^{t+1} \\ v_1^{t+1} = v_1^t + h(-cv_1^t - \lambda_1^t + u_1) \\ v_2^{t+1} = v_2^t + h(-cv_2^t + \lambda_1^t - \lambda_3^t) \\ v_3^{t+1} = v_3^t + h(-cv_3^t + \lambda_3^t + u_3), \end{cases} \quad (13)$$

and the contact forces are given by

$$\begin{cases} \lambda_1^t = \max\{k(q_2 - q_1 - d), 0\} \\ \lambda_3^t = \max\{k(q_3 - q_2 - d), 0\}. \end{cases} \quad (14)$$

A. Implementation Details

For implementation, we use a timestep of $h = 0.01$ with $T = 500$. The fixed-timestep direct collocation, relaxed and exact LCP is implemented using a SNOPT [19] wrapper in Drake [14]. Throughout the results, we only report trajectories where the optimization has succeeded (optimality conditions are satisfied in the SNOPT solver). The single shooting implementation is done in PyTorch [20] without the use of GPU. We use the Stochastic Gradient Descent (SGD) solver with a learning rate of $\eta = 500$. We simulate the parameters with each mass of 1kg, spring constant of $k = 30$, and viscous damping of $c = 2$.

B. Comparing Trajectories of Relaxations

We compare trajectories of different relaxation methods as we iteratively tune ε, and illustrate how the trajectories change using different relaxation methods. Fig.4 and Fig.5 utilizes regularized contact with the softplus model, respectively with direct collocation and single shooting with gradient descent. The tolerance in the single shooting example can be tuned to be exact, where $\epsilon = 0$ corresponds to gradient descent on ReLU dynamics (Fig.5). Finally, Fig.6 illustrates change of trajectories as the bounds on LCP relaxation is tightened.

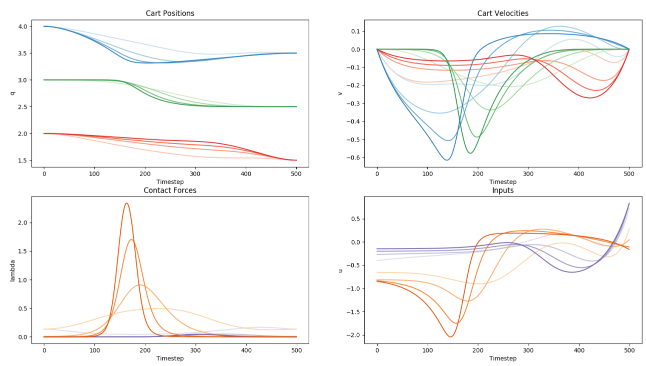

Fig. 4. Change of trajectories while iteratively tightening bounds for smooth contact relaxation with direct collocation. The tolerances used are $\varepsilon = [1e-3, 1e-5, 1e-7, 1e-9]$

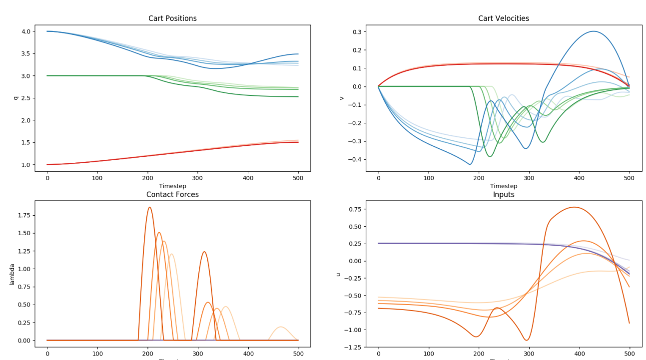

Fig. 5. Change of trajectories while iteratively tightening bounds for smooth contact relaxation with gradient descent. The tolerances used are $\varepsilon = [1e-3, 1e-5, 1e-7, 0.0]$, where 0.0 corresponds to gradient descent with exact ReLU. Darker colors indicate lower tolerances. Note the colors match with the colors of each quantity in Fig.3

C. Quantitative Comparison of Warm Starting

To quantitatively compare the results of different indirect methods with direct methods, we evaluate the resulting final trajectories using the following criteria:

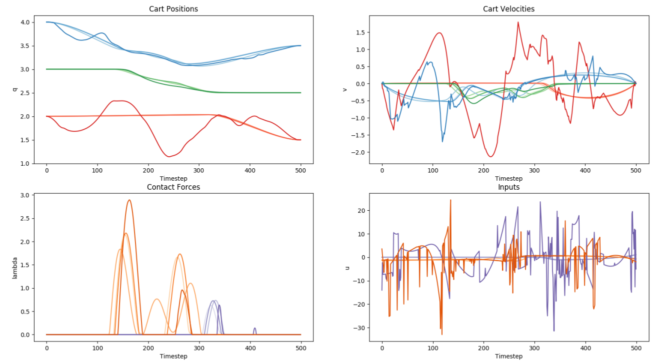

Fig. 6. Change of trajectories while iteratively tightening bounds for LCP relaxation. The tolerances used are $\varepsilon = [1e-1, 1e-2, 1e-3, 0.0]$, where 0.0 corresponds to the exact LCP constraint. Darker colors indicate lower tolerances.

1) Resulting cost of the final trajectory
2) Time it took to obtain the trajectory
3) Stability cost of the final trajectory
4) Physical Violations

We utilize sum of input norm for the cost of the final trajectory ($\sum_t u^T u$), and the stability cost is formulated as the 2-norm difference between the desired final state and the actual final state. Finally, the physical violation is evaluated as a sum of LCP violations in (10). The resulting comparison of the four items are tabulated in Table I.

D. Discussion of Quantiative Results

From the quantitative results in Table I, we hope to discuss the following lessons from the empirical study.

1) Regularized Contact Models Help Convergence of Direct Methods: Regardless of the choice of the direct method, we saw that regularized contact models create a very favorable condition to warm start the optimization problem. This is illustrated by how both LCP and ReLU Gradient Descent (GD) have lower costs when warm started by Softplus models of contact.

2) LCP Relaxation is rarely a good idea: In practice, we see that warm starting the optimization with LCP relaxations take painfully long as nonconvex bilinear constraints. However, it is even more surprising that strict LCP, when warm-started with relaxed LCP, yields *higher cost* compared to just doing strict LCP from the beginning. In comparison, strict LCP warm-started with Softplus models yielded much lower cost compared to no warm starting.

3) LCP is unreliable and costly: It was interesting to observe that LCP methods often take much longer time to compute compared to gradient descent methods, but yield costs that are orders of 2 magnitudes higher than GD methods. This is also illustrated in Fig.6, where trajectories obtained by LCP methods tend to be jerky and numerically ill-conditioned. We remind the reader that optimality conditions were still satisfied.

4) Gradient Descent with ReLU dynamics is surprisingly effective: Finally, the success of single shooting gradient descent methods with ReLU model of contact was extremely

Strict LCP (No Warm Start)					ReLU GD (No Warm Start)				
Tol.(ε)	Cost	Stability	Physics	Time(s)	Tol.(ε)	Cost	Stability	Physics	Time(s)
0.0	40779.28	0.00	0.00	31.56	0.0	542.56	0.04	0.16	49.59
LCP Relaxation + Strict LCP					**LCP Relaxation + ReLU GD**				
Tol.(ε)	Cost	Stability	Physics	Time(s)	Tol.(ε)	Cost	Stability	Physics	Time(s)
1e-2	484.21	0.00	2.99	1407.90	Same as left				
1e-3	561.00	0.00	0.34	580.77	Same as left				
1e-4	469.94	0.00	0.03	430.16	Same as left				
0.0	46415.20	0.00	0.00	20.73	0.0	492.06	0.00	0.20	0.27
Softplus Direct Collocation + Strict LCP					**Softplus Direct Collocation + ReLU GD**				
Tol.(ε)	Cost	Stability	Physics	Time(s)	Tol.(ε)	Cost	Stability	Physics	Time(s)
1e-3	197.86	0.00	2445.98	11.64	Same as left				
1e-5	261.23	0.00	625.11	9.78	Same as left				
1e-7	339.53	0.00	239.50	12.40	Same as left				
0.0	16619.35	0.00	0.00	100.07	0.0	360.31	0.00	0.20	1.43
Softplus GD + Strict LCP					**Softplus GD + ReLU GD**				
Tol.(ε)	Cost	Stability	Physics	Time(s)	Tol.(ε)	Cost	Stability	Physics	Time(s)
1e-3	149.18	0.41	0.01	4.07	Same as left				
1e-5	166.45	0.27	0.11	3.77	Same as left				
1e-7	215.14	0.20	0.12	2.94	Same as left				
0.0	3748.30	0.00	0.00	151.54	0.0	309.42	0.05	0.21	2.38

TABLE I
COMBINING INDIRECT AND DIRECT METHODS

surprising, given the conventional doubts about numerical losses encountered by the method. Whether or not this success will continue on to more complicated systems is questionable, but the effectivess of the method on our particular example is beyond doubt, as it consistently yielded two orders of magnitude lower cost with the two magnitude lower computing time.

VI. CONCLUSION & FUTURE WORKS

We have considered a framework of warm starting direct methods of contact with less physically accurate, but easier-to-solve methods in order to produce tractable and physically accurate trajectories for planning through contact. We introduced two relaxations: a Softplus approximation for contact forces, and relaxation of complentarity constraints into bilinear inequality constraints. Using these relaxations, we warm started two methods that respect the physics: the strict LCP method, and the single shooting method with gradient descent. We have attempted a combination of different methods on the three cart example of Fig.3.

Through our results, we were able to show that smooth and regularized models of contact often serve as effective methods to warm start exact formulations of contact, strengthening the result obtained in [11].

In contrast, although we made theoretical claims on how solving the relaxed LCP converges faster to the true LCP solution compared to the regularized model, our studies showed that solving the relaxed LCP itself is often unreliable and computationally costly in practice. Not only is complementarity a difficult condition to satisfy, but relaxing it to a nonconvex bilinear inequality constraints also does not seem to help much.

A particularly surprising result that we demonstrated was the effectiveness of approximating contact with ReLU dynamics, and performing single shooting on this dynamics with gradient descent. For our simple example, this method outperformed LCP by orders of magnitude, both in final cost and computation time. We also note that even in this case, performing trajectory optimization using smooth models of contact with SoftPlus and using the solution to warm start an exact ReLU model helped decrease the cost by a large margin. This particularly goes against the accepted consensus about the shortcomings of single-shooting methods, and we plan to study the reason for its effectiveness in the future.

A big question that remains is whether or not the conclusion from our work can be extended to more complicated dynamics involving contact. For instance, friction has not been considered in our example, although [11] proposed regularizing friction with a logistic-like function. Other examples include whole-body motion planning of humanoids or multi-finger manipulation, which involve systems of much larger dimensions and multiple nonlinearities [21].

We are also very interested in whether single shooting with gradient descent will still remain effective for adversarial settings, where unstable systems create exploding/vanishing gradient problems. We believe that performing direct collocation with regularized contact models, and then doing gradient descent 'fine tuning' might mitigate some of the problems associated with directly performing gradient descent.

ACKNOWLEDGEMENT

The authors would like to thank Tobia Marcucci, Wilko Schwarting, and Russ Tedrake for the helpful discussion on the material, as well as for a wonderful term of class!

REFERENCES

[1] M. T. Mason, "Mechanics of robotic manipulation," 2001.

[2] F. R. Hogan and A. Rodriguez, *Feedback Control of the Pusher-Slider System: A Story of Hybrid and Underactuated Contact Dynamics*. Cham: Springer International Publishing, 2020, pp. 800–815. [Online]. Available: https://doi.org/10.1007/978-3-030-43089-4_51

[3] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, "Optimization-based locomotion planning, estimation, and control design for atlas," *Autonomous Robots*, vol. 40, no. 3, p. 429–455, 2016.

[4] H. Zhao, W. Ma, A. D. Ames, and M. B. Zeagler, "Human-inspired multi-contact locomotion with amber2," in *2014 ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*, 2014, pp. 199–210.

[5] T. Marcucci and R. Tedrake, "Mixed-integer formulations for optimal control of piecewise-affine systems," in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 230–239. [Online]. Available: https://doi.org/10.1145/3302504.3311801

[6] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014. [Online]. Available: https://doi.org/10.1177/0278364913506757

[7] Z. Manchester, N. Doshi, R. J. Wood, and S. Kuindersma, "Contact-implicit trajectory optimization using variational integrators," *The International Journal of Robotics Research*, vol. 38, no. 12-13, pp. 1463–1476, 2019. [Online]. Available: https://doi.org/10.1177/0278364919849235

[8] A. Önol, P. Long, and T. Padır, "Contact-implicit trajectory optimization based on a variable smooth contact model and successive convexification," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 2447–2453.

[9] J. Fu, S. Levine, and P. Abbeel, "One-shot learning of manipulation skills with online dynamics adaptation and neural network priors," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 4019–4026.

[10] Y. Li, J. Wu, J.-Y. Zhu, J. B. Tenenbaum, A. Torralba, and R. Tedrake, "Propagation networks for model-based control under partial observation," in *ICRA*, 2019.

[11] T. Marcucci, M. Gabiccini, and A. Artoni, "A two-stage trajectory optimization strategy for articulated bodies with unscheduled contact sequences," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 104–111, 2017.

[12] E. Todorov, "A convex, smooth and invertible contact model for trajectory optimization," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 1071–1076.

[13] P. Song, J.-S. Pang, and V. Kumar, "A semi-implicit time-stepping model for frictional compliant contact problems," *International Journal for Numerical Methods in Engineering*, vol. 60, no. 13, pp. 2231–2261, 2004. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.1049

[14] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: https://drake.mit.edu

[15] A. M. Castro, A. Qu, N. Kuppuswamy, A. Alspach, and M. Sherman, "A transition-aware method for the simulation of compliant contact with regularized friction," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1859–1866, 2020.

[16] P. Castro, "Tightening piecewise mccormick relaxations for bilinear problems," *Computers Chemical Engineering*, vol. 72, 04 2014.

[17] B. Grimstad and H. Andersson, "Relu networks as surrogate models in mixed-integer linear programs," 2019.

[18] A. Aydinoglu, V. M. Preciado, and M. Posa, "Contact-aware controller design for complementarity systems," *ArXiv*, vol. abs/1909.11221, 2019.

[19] P. E. Gill, W. Murray, and M. A. Saunders, "Snopt: An sqp algorithm for large-scale constrained optimization," *SIAM Rev.*, vol. 47, no. 1, p. 99–131, Jan. 2005. [Online]. Available: https://doi.org/10.1137/S0036144504446096

[20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.

[21] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," *ACM Trans. Graph.*, vol. 31, no. 4, July 2012. [Online]. Available: https://doi.org/10.1145/2185520.2185539