

# Irrotational Contact Fields

Alejandro Castro, Xuchen Han, Joseph Masterjohn

**Abstract**—We present Irrotational Contact Fields (ICF), a framework for generating convex approximations of complex contact models, incorporating experimentally validated models like Hunt & Crossley coupled with Coulomb’s law of friction alongside the principle of maximum dissipation. Our approach is robust across a wide range of stiffness values, making it suitable for both compliant surfaces and rigid approximations. We evaluate these approximations across a wide variety of test cases, detailing properties and limitations. We implement a fully differentiable solution in the open-source robotics toolkit, Drake. Our novel *hybrid* approach enables efficient computation of gradients for complex geometric models by reusing factorizations from contact resolution. We demonstrate robust simulation of robotic tasks at interactive rates, with accurately resolved stiction and contact transitions, supporting effective sim-to-real transfer.

**Index Terms**—Contact Modeling, Simulation and Animation, Dexterous Manipulation, Dynamics.

## I. INTRODUCTION

**S**IMULATION of multibody systems with frictional contact is essential in robotics, aiding hardware optimization, controller design and testing, continuous integration, and data generation for machine learning. Accurate physics models enable advanced controller design and trajectory optimization algorithms. However, achieving robust, accurate, and efficient simulations for contact-rich robotics applications remains challenging.

Rigid body dynamics with frictional contact are complicated by non-smooth solutions. Acceleration-level formulations can lead to singular configurations, known as Painlevé paradoxes, where solutions may not exist. Discrete velocity-level formulations circumvent this by allowing discrete velocity jumps and impulsive forces. While many variants exist, the general form of a discrete formulation enforces balance of momentum subject to contact constraints, with additional constraints to incorporate Coulomb’s law of friction under the maximum dissipation principle. With the addition of decision variables and Lagrange multipliers, the result is a large and challenging-to-solve Non-linear Complementarity Problem (NCP) with a much larger number of variables than the original problem.

Solving NCPs robustly and efficiently has remained elusive. NCPs are equivalent to non-convex global optimization problems, which are generally NP-hard [1]. Therefore, NCPs may lack solutions or have multiple solutions. In practice, iterative solvers might incorrectly assume convergence or terminate early, leading to solutions that fail to satisfy the original equations or violate physical laws.

There are other, less frequently-discussed issues, such as numerical conditioning and implementation details, that affect

convergence properties and robustness. These are important problems when it comes to the simulation of complex systems, with either many degrees of freedom (DOFs), a large number of constraints, or both. Solutions that work for small systems often do not work when faced with real-world engineering applications.

In an attempt to make the problem tractable, Anitescu introduced a convex approximation of contact [2], effectively replacing the original NCP with a convex approximation that guarantees the existence of solutions. Later, Todorov [3] regularized this formulation to write a strictly convex formulation with a unique solution. Our previous work [4] introduced the convex Semi-Analytic Primal (SAP) formulation for modeling compliant contact. SAP embraces compliance to provide a robust and performant tool that targets robotics applications — where grippers feature compliant surfaces for stable grasps and robotic feet are often padded with compliant materials. Although the formulation is inherently compliant, rigid contact can be modeled to a good approximation using a *near-rigid* approximation. The work focused on physical correctness, numerical conditioning, and robustness.

However, the SAP formulation has several limitations. The most well-known is the artifact of *gliding* during slip, an artifact inherited from the original formulation by Anitescu [5] and shared by Todorov’s formulation [3] in MuJoCo [6]. In Anitescu’s formulation, as objects slip, they glide at a finite distance  $\delta t \mu \|v_t\|$ , proportional to the time step  $\delta t$ , the coefficient of friction  $\mu$ , and the slip speed  $\|v_t\|$ . For interactive simulation of a pushing task at  $\|v_t\| \approx 1$  m/s with  $\delta t = 10$  ms and  $\mu = 1$ , the gliding artifact is as large as 1 cm. This gliding artifact was reported inadequate for the simulation of quadrupeds in [7]. Moreover, we show that both SAP’s compliant model and Todorov’s regularized model are inconsistent. The gliding effect is also proportional to dissipation and does not vanish as the time step goes to zero. This exacerbates the nonphysical gliding artifact further, potentially increasing action at a distance orders of magnitude depending on the amount of dissipation, even at small time step sizes. Finally, SAP’s model of compliance is intrinsic to its convex formulation, and therefore, it is not possible to incorporate experimentally validated engineering-grade models of contact, such as Hertz and Hunt & Crossley [8]. Prior to this work, no convex formulation allowed the integration of such models.

Despite its limitations, the SAP formulation offers a robust solution with theoretical convergence guarantees that carry over to practical implementations. This work develops novel convex approximations that retain these guarantees while reducing (and even eliminating) *gliding* and related artifacts of existing convex formulations, enabling the use of validated contact models, enhancing friction fidelity, and improving numerical performance.

All authors are with the Toyota Research Institute, USA.  
 alejandro.castro@tri.global, xuchen.han@lbn.global,  
 joseph.masterjohn@tri.global  
 ©2026 IEEE

Our work is organized as follows. Section III summarizes our previous work and reviews compliant contact with regularized friction. Section IV presents ICF, establishing conditions for embedding arbitrary contact models into our convex formulation. Section V derives two new approximations, *Lagged* and *Similar*, by solving these conditions. Section VI completes the model using Hunt & Crossley dissipation, enabling, for the first time, integration of arbitrary contact laws in a convex setting. We analyze the approximations in Section VII, succinctly summarizing their properties in Table I. Section VIII analyzes the connection between impacts and numerical conditioning (we believe this is the first time this connection has been made) and Section IX explores the connection to barrier methods. Implementation details appear in Section X, followed by extensive benchmarks and applications in Sections XI-XII. We close with final remarks and conclusions in Section XIII.

## II. NOVEL CONTRIBUTIONS

Our main contribution is ICF, a novel mathematical framework for generating convex approximations of frictional contact, centered on the concept of *irrotational contact fields*.

While our previous work [4] draws an explicit connection between earlier convex formulations [2], [3] and the physics of compliance with regularized friction, ICF enables the integration of engineering-grade, experimentally validated, contact laws; this is something not possible with existing convex models. Moreover, the formulation imposes minimum requirements on the functional form used for the regularization of friction, allowing for the design of regularized friction models with better numerical properties.

Within this framework, we propose two new contact approximations, *Lagged* and *Similar*. Both approximations allow for the incorporation of arbitrary contact force laws. However, we show that only the Lagged approximation eliminates the *gliding* at a distance artifact of SAP [4] and MuJoCo [6] formulations.

Our key contributions are:

- 1) A mathematical framework to generate convex approximations of complex contact models.
- 2) The first convex approximation of compliant contact with Hunt & Crossley dissipation [8].
- 3) Two new convex approximations, Lagged and Similar.

In a nutshell, ICF addresses two key challenges that are impossible to resolve with existing convex formulations:

- 1) To incorporate arbitrary engineering-grade models of contact (e.g. Hunt & Crossley).
- 2) To eliminate the *gliding* artifact of existing convex formulations.

In addition to these main research contributions, our work includes a fully differentiable open-source implementation in Drake [9]; see Section X-B. We compare our convex approximations in Section VII, validate them through extensive simulation tests in Section XI, and stress test them in Section XII with engineering applications, including the complex gearing mechanism of a BarrettHand [10] (Section XII-C) and deformable FEM bodies (Section XII-E). Moreover, we

provide an extensive study on a Franka hand to assess the impact of different approximations on grasp stability (Section XII-B).

## III. MATHEMATICAL FORMULATION

### A. Preliminaries

We first introduce the notation and conventions used in our framework. We closely follow the notation in our previous work [4], [11], later extended to deformable Finite Element Method (FEM) models in [12] and Material Point Method (MPM) models in [13].

State is described by generalized positions  $\mathbf{q} \in \mathbb{R}^{n_q}$  and generalized velocities  $\mathbf{v} \in \mathbb{R}^{n_v}$ , where  $n_q$  and  $n_v$  are the number of generalized positions and velocities, respectively. Time derivatives of the generalized positions relate to the generalized velocities by the kinematic map  $\dot{\mathbf{q}} = \mathbf{N}(\mathbf{q})\mathbf{v}$ . We use joint coordinates to describe articulated rigid bodies.

Consider a system with  $n_c$  constraints introducing  $n_{eq}$  constraint equations. Constraint velocities  $\mathbf{v}_c \in \mathbb{R}^{n_{eq}}$  are given by  $\mathbf{v}_c = \mathbf{J}\mathbf{v} + \mathbf{b}$ , where  $\mathbf{J}$  is the stacked Jacobian for all constraints and  $\mathbf{b}$  is a bias term. Following [4], we write a symplectic Euler scheme with fixed time steps of size  $\delta t$  as

$$\begin{aligned} \mathbf{M}(\mathbf{q}_0)(\mathbf{v} - \mathbf{v}_0) &= \\ \delta t \mathbf{k}_1(\mathbf{q}, \mathbf{v}) + \delta t \mathbf{k}_2(\mathbf{q}_0, \mathbf{v}_0) + \mathbf{J}(\mathbf{q}_0)^T \boldsymbol{\gamma}, & \quad (1) \\ \mathbf{q} = \mathbf{q}_0 + \delta t \mathbf{N}(\mathbf{q}_0)\mathbf{v}, & \end{aligned}$$

where quantities with the naught subscript are evaluated at the previous time step  $t^n$  and quantities without subscripts are evaluated at the next time step  $t^{n+1} = t^n + \delta t$ .  $\mathbf{M}$  is the joint space mass matrix. Stiff terms, such as springs, and stabilizing terms, such as damping and rotor inertias, are treated implicitly in  $\mathbf{k}_1$ . Smooth, non-stiff terms such as gravity, gyroscopic forces, and feed-forward actuation are treated explicitly in  $\mathbf{k}_2$ . Vector  $\boldsymbol{\gamma} \in \mathbb{R}^{n_{eq}}$  corresponds to constraint impulses. Our implementation in Drake [9] includes contact constraints, holonomic constraints (e.g. weld, distance, coupler), and PD controllers with effort limits [4], [12].

We define the momentum residual  $\mathbf{m}(\mathbf{v})$  as

$$\mathbf{m}(\mathbf{v}) = \mathbf{M}(\mathbf{q}_0)(\mathbf{v} - \mathbf{v}_0) - \delta t \mathbf{k}_1(\mathbf{q}, \mathbf{v}) - \delta t \mathbf{k}_2(\mathbf{q}_0, \mathbf{v}_0),$$

and define the *free motion velocities*  $\mathbf{v}^*$  as the solution to  $\mathbf{m}(\mathbf{v}^*) = \mathbf{0}$ , that is, the velocities the system would have in the absence of constraint forces. SAP [4] linearizes (1) at  $\mathbf{v} = \mathbf{v}^*$  as

$$\mathbf{A}(\mathbf{v} - \mathbf{v}^*) = \mathbf{J}^T \boldsymbol{\gamma}, \quad (2)$$

where  $\mathbf{A}$  is a symmetric positive-definite (SPD) approximation of the gradient of the momentum residual accurate to first order, i.e.

$$\left. \frac{\partial \mathbf{m}}{\partial \mathbf{v}} \right|_{\mathbf{v}=\mathbf{v}^*} = \mathbf{A} + \mathcal{O}(\delta t).$$

While early work [14], [15] use  $\mathbf{A} = \mathbf{M}(\mathbf{q}_0)$ , the mass matrix at the previous time step, our formulation incorporates the modeling of joint springs, damping, and rotor inertia implicitly [4].

Traditional NCP formulations supplement (2) with constraint equations to model contact. Additional constraint equations, slack variables, and Lagrange multipliers lead to a large and challenging-to-solve NCP. Instead, SAP [4] follows a different approach. Inspired by the analytical inverse dynamics from Todorov [3], SAP eliminates constraints analytically to write an unconstrained convex optimization problem for the velocities at the next time step

$$\min_{\mathbf{v}} \ell_p(\mathbf{v}) = \frac{1}{2} \|\mathbf{v} - \mathbf{v}^*\|_A^2 + \ell_c(\mathbf{v}), \quad (3)$$

where  $\|\mathbf{x}\|_A^2 = \mathbf{x}^T \mathbf{A} \mathbf{x}$ , and the constraints cost  $\ell_c(\mathbf{v})$  penalizes constraint impulses. It can be shown that SAP's [4] formulation models linear compliant contact with a Kelvin-Voigt (linear) model of damping. The formulation is strongly convex, with convergence guarantees that lead to robust software implementations. Although the formulation is inherently compliant with regularized friction, [4] shows how rigid contact can be modeled to a good approximation using a *near-rigid* approximation.

To be more precise, we write  $\ell_c(\mathbf{v}; \mathbf{x}_0)$ , emphasizing this is an *incremental potential* [16]. This potential depends on the generalized velocities  $\mathbf{v}$ , given a previous state  $\mathbf{x}_0$ , making it inherently discrete—unlike continuous potentials like gravity or electric fields. The presence of the quadratic term  $\|\mathbf{x}\|_A^2$  ensures strong convexity of the optimization problem (3) when  $\ell_c$  is convex, guaranteeing a unique solution. Developing such convex potentials for frictional contact modeling is the primary aim of this work.

Taking the gradient of (3), we obtain the balance of momentum (2) and find that impulses emerge as a result of the constraint potentials

$$\boldsymbol{\gamma}(\mathbf{v}_c; \mathbf{x}_0) = -\frac{\partial \ell_c(\mathbf{v}_c; \mathbf{x}_0)}{\partial \mathbf{v}_c}, \quad (4)$$

where we use the notation  $\partial f / \partial \mathbf{x} = \nabla f \in \mathbb{R}^n$  to denote the gradient of a scalar function  $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ . As with SAP, we consider potential functions that are *separable*

$$\ell_c(\mathbf{v}_c; \mathbf{x}_0) = \sum_{i=1}^{n_c} \ell_{c,i}(\mathbf{v}_{c,i}; \mathbf{x}_0), \quad (5)$$

where  $\mathbf{v}_c$  is the stacked vector of individual constraint velocities  $\mathbf{v}_{c,i}$ . For contact constraints, we define a contact frame  $C_i$  for which we arbitrarily choose the  $z$ -axis to coincide with the contact normal  $\hat{\mathbf{n}}_i$ . In this frame, the normal and tangential components of  $\mathbf{v}_{c,i}$  are given by  $v_{n,i} = \hat{\mathbf{n}}_i \cdot \mathbf{v}_{c,i}$  and  $\mathbf{v}_{t,i} = \mathbf{v}_{c,i} - v_{n,i} \hat{\mathbf{n}}_i$  respectively, and  $\mathbf{v}_{c,i} = [v_{t,i} \ v_{n,i}]$ . By convention, we define the relative contact velocity  $\mathbf{v}_{c,i}$  and normal  $\hat{\mathbf{n}}_i$  such that  $v_{n,i} > 0$  for objects moving away from each other.

Using this separable potential (5) in (4), the impulse vector  $\boldsymbol{\gamma}$  is the stacked vector of individual constraint contributions  $\boldsymbol{\gamma}_i(\mathbf{v}_{c,i}; \mathbf{x}_0) = -\partial \ell_{c,i} / \partial \mathbf{v}_{c,i}$ . We highlight the dependence on the previous time step state  $\mathbf{x}_0$  to emphasize the discrete (or incremental) nature of these potentials and the resulting impulses. However, hereinafter, we write  $\boldsymbol{\gamma}_i(\mathbf{v}_{c,i})$  and  $\ell_{c,i}(\mathbf{v}_{c,i})$  to shorten notation, and the functional dependence on the previous time step state is implicitly assumed.

## B. Compliant Contact with Regularized Friction

When two solids come into contact, they deform—at least microscopically—to prevent interpenetration. Contact stresses arise, whose surface integral is the net contact force. *Compliant contact* models these forces as algebraic, state-dependent functions that capture quasistatic elastic effects without solving the full elasticity problem. Experimentally validated models based on Hertz theory [17] with Hunt–Crossley dissipation [8] are common examples widely adopted in engineering [18].

In this work, we consider a general force law

$$f_n(\phi, v_n), \quad (6)$$

function of the signed distance  $\phi$  (defined as negative when objects overlap), and of the normal velocity  $v_n$  (defined positive when objects separate). In the discrete setting, with a time step of size  $\delta t$ , we use a first-order approximation of the signed distance  $\phi = \phi_0 + \delta t v_n$ , implicit in the next time step velocity  $v_n$ . Using this approximation, we define a discrete normal impulse as

$$n(v_n; \phi_0) = \delta t f_n(\phi_0 + \delta t v_n, v_n), \quad (7)$$

and the associated normal potential  $\ell_n(v_n) = -N(v_n)$ , with  $N(v_n)$  the indefinite integral of  $n(v_n)$ . That is,  $\gamma_n = -\ell'_n(v_n) = N'(v_n)$ . We observe that since

$$\frac{d^2 \ell_n}{dv_n^2} = -\delta t^2 \frac{\partial f_n}{\partial \phi} - \delta t \frac{\partial f_n}{\partial v_n},$$

it follows that

$$\frac{\partial f_n}{\partial \phi} \leq 0 \text{ and } \frac{\partial f_n}{\partial v_n} \leq 0, \quad (8)$$

are sufficient conditions for the normal cost  $\ell_n(v_n)$  to be convex.

Friction can be modeled as a continuous function of state using a *regularized* approximation

$$\begin{aligned} \boldsymbol{\gamma}_t(\mathbf{v}_c) &= -\mu f(\|\mathbf{v}_t\|/\varepsilon_s) n(v_n) \hat{\mathbf{t}} \\ f(s) &= \frac{s}{\sqrt{1+s^2}} \\ \hat{\mathbf{t}} &= \frac{\mathbf{v}_t}{\|\mathbf{v}_t\|} \end{aligned} \quad (9)$$

where function  $f(s) \leq 1$  *regularizes* Coulomb friction, with  $\mu$  the coefficient of friction and  $\varepsilon_s$  the *regularization parameter*. When  $\|\mathbf{v}_t\| \ll \varepsilon_s$ , the model behaves as viscous damping with high viscosity. When  $\|\mathbf{v}_t\| \gg \varepsilon_s$ , the model approximates Coulomb's law, with friction opposing slip velocity according to the maximum dissipation principle and  $\|\boldsymbol{\gamma}_t\| \rightarrow \mu \gamma_n$ . The choice of function  $f(s)$  is somewhat arbitrary as long as  $f(s) \leq 1$  and  $f(s) = 0$  at  $s = 0$ . We choose  $f(s)$  such that (9) can be simplified to

$$\begin{aligned} \boldsymbol{\gamma}_t(\mathbf{v}_c) &= -\mu n(v_n) \hat{\mathbf{t}}_s \\ \hat{\mathbf{t}}_s &= \frac{\mathbf{v}_t}{\sqrt{\|\mathbf{v}_t\|^2 + \varepsilon_s^2}} \end{aligned} \quad (10)$$

where we define the regularized or *soft* tangent vector  $\hat{\mathbf{t}}_s$ , which can be shown to be the gradient with respect to  $\mathbf{v}_t$  of the *soft* norm  $\|\mathbf{x}\|_s = \sqrt{\|\mathbf{x}\|^2 + \varepsilon_s^2} - \varepsilon_s$  (Appendix A). Unlike

$\hat{\mathbf{t}}$ , which is not well-defined at  $\mathbf{v}_t = \mathbf{0}$ ,  $\hat{\mathbf{t}}_s$  is well-defined and continuous for all values of slip velocity. Moreover, (10) has continuous gradients, a desirable property that improves the convergence of non-linear solvers based on Newton's method.

In total, we write the contact impulse as

$$\boldsymbol{\gamma}(\mathbf{v}_c) = \begin{bmatrix} \gamma_t \\ \gamma_n \end{bmatrix} = \begin{bmatrix} -\mu \hat{\mathbf{t}}_s(\mathbf{v}_t) n(v_n) \\ n(v_n) \end{bmatrix}. \quad (11)$$

However, this model is not necessarily the result of a potential and  $\boldsymbol{\gamma}(\mathbf{v}_c) \neq -\partial \ell_c / \partial \mathbf{v}_c$  in general. In the next section we develop a general theory that allows us to write convex approximations of this model that fit within the optimization framework (3).

#### IV. IRROTATIONAL CONTACT FIELDS

Helmholtz's theorem [19] states that any vector field admits a decomposition into an irrotational field (zero curl) and a solenoidal field (zero divergence)

$$\boldsymbol{\gamma}(\mathbf{v}_c) = -\frac{\partial \ell(\mathbf{v}_c)}{\partial \mathbf{v}_c} + \nabla \times \boldsymbol{\Psi}(\mathbf{v}_c) \quad (12)$$

with  $\ell(\mathbf{v}_c)$  a scalar potential and  $\boldsymbol{\Psi}(\mathbf{v}_c)$  a *vector potential*.

Here, we neglect the solenoidal component to investigate *irrotational fields*, which satisfy the *curl-free* condition

$$\nabla \times \boldsymbol{\gamma} = \mathbf{0}. \quad (13)$$

This curl-free condition is precisely the necessary condition for the existence of a scalar potential  $\ell(\mathbf{v}_c)$  satisfying  $\boldsymbol{\gamma}(\mathbf{v}_c) = -\partial \ell(\mathbf{v}_c) / \partial \mathbf{v}_c$ , as needed in (4).

The normal component in (13)

$$\frac{\partial \gamma_{t,1}}{\partial v_{t,2}} = \frac{\partial \gamma_{t,2}}{\partial v_{t,1}}, \quad (14)$$

states that the two-dimensional field  $\boldsymbol{\gamma}_t(\mathbf{v}_t)$  is irrotational in the  $\mathbf{v}_t$  plane. Consider the generic isotropic friction model

$$\boldsymbol{\gamma}_t = g(\|\mathbf{v}_t\|, v_n) \hat{\mathbf{t}} \quad (15)$$

whose gradient is

$$\frac{\partial \boldsymbol{\gamma}_t}{\partial \mathbf{v}_t} = \frac{\partial g}{\partial \|\mathbf{v}_t\|} \mathbf{P} + g \frac{\mathbf{P}^\perp}{\|\mathbf{v}_t\|}$$

with symmetric projection matrices  $\mathbf{P}$  and  $\mathbf{P}^\perp$  (see Appendix A). Therefore,  $\partial \boldsymbol{\gamma}_t / \partial \mathbf{v}_t$  is symmetric, and condition (14) is satisfied. Thus isotropic friction fields are irrotational. While this work focuses on isotropic friction, Appendix B shows how to incorporate anisotropic friction within this same framework.

Finally, the tangential components in (13) lead to the condition

$$\frac{\partial \gamma_t}{\partial v_n} = \frac{\partial \gamma_n}{\partial \mathbf{v}_t}, \quad (16)$$

Contact models are not irrotational in the general case, and condition (16) is not satisfied. In the next section we solve for irrotational approximations of (11) that satisfy the curl condition  $\nabla \times \boldsymbol{\gamma} = \mathbf{0}$  and thus a scalar potential  $\ell(\mathbf{v}_c)$  exists.

#### V. A FAMILY OF CONVEX APPROXIMATIONS

We present a family of convex approximations of the model in (11) satisfying (16), and establish conditions for convexity. Moreover, we present two novel convex approximations of contact: *Lagged* and *Similar*.

##### A. SAP model

While SAP [4] is convex by construction, it is a good exercise to verify that our new conditions hold for this model. SAP regularizes friction according to

$$\boldsymbol{\gamma}_t(\mathbf{v}_c) = -\min\left(\frac{\|\mathbf{v}_t\|}{R_t}, \mu \gamma_n(v_n)\right) \hat{\mathbf{t}} \quad (17)$$

where  $R_t$  is SAP's regularization parameter [4]. From [4], we know that the Hessian of the regularizer cost

$$\mathbf{G} = \frac{\partial^2 \ell}{\partial \mathbf{v}_c^2} = -\frac{\partial \boldsymbol{\gamma}}{\partial \mathbf{v}_c} = -\begin{bmatrix} \frac{\partial \gamma_t}{\partial \mathbf{v}_t} & \frac{\partial \gamma_t}{\partial v_n} \\ \frac{\partial \gamma_n}{\partial \mathbf{v}_t} & \frac{\partial \gamma_n}{\partial v_n} \end{bmatrix}$$

is symmetric positive semi-definite and therefore SAP satisfies condition (16). Moreover,  $\partial \boldsymbol{\gamma}_t / \partial \mathbf{v}_t$  is symmetric since  $\boldsymbol{\gamma}_t$  is of the isotropic form (15) and condition (14) is met also.

##### B. Lagged Model

We use the model of regularized friction (9) in which the normal impulse is *lagged* to the previous time step

$$\boldsymbol{\gamma}_t(\mathbf{v}_t) = -\mu f(s) \gamma_{n,0} \hat{\mathbf{t}} \quad (18)$$

with  $s = \|\mathbf{v}_t\| / \varepsilon_s$ , and using (6),  $\gamma_{n,0} = \delta t f_n(\phi_0, v_{n,0})$ . For a physical model of compliance for which  $\gamma_n$  is only a function of  $v_n$ , condition (16) is trivially met since

$$\begin{aligned} \frac{\partial \boldsymbol{\gamma}_t}{\partial v_n} &= \mathbf{0} \\ \frac{\partial \gamma_n}{\partial \mathbf{v}_t} &= \mathbf{0} \end{aligned}$$

which implies that the contact potential is separable as the sum of normal and friction contributions

$$\ell(\mathbf{v}_c) = \ell_t(\mathbf{v}_t) + \ell_n(v_n).$$

Notice that even though they are lagged in the friction component, the normal impulses  $\gamma_n(v_n)$  are still treated implicitly, with normal potential  $\ell_n(v_n)$ .

We verify that

$$\ell_t(\mathbf{v}_t) = \mu \gamma_{n,0} \varepsilon_s F(\|\mathbf{v}_t\| / \varepsilon_s)$$

with  $f = F'$  satisfies  $\boldsymbol{\gamma}_t = -\partial \ell_t / \partial \mathbf{v}_t$ .

The Hessian of  $\ell_t$  is given by

$$\begin{aligned} \frac{\partial^2 \ell_t}{\partial \mathbf{v}_t^2} &= -\frac{\partial \boldsymbol{\gamma}_t}{\partial \mathbf{v}_t} \\ &= \mu \gamma_{n,0} \left[ \frac{f'(s)}{\varepsilon_s} \mathbf{P}(\hat{\mathbf{t}}) + \frac{f(s)}{\|\mathbf{v}_t\|} \mathbf{P}^\perp(\hat{\mathbf{t}}) \right]. \end{aligned}$$

With a non-decreasing  $f$ ,  $F$  is convex, and the Hessian of  $\ell_t$  is the positive linear combination of two projection matrices.

Therefore,  $\ell_t$  is twice differentiable and convex. Thus the Hessian

$$\frac{\partial^2 \ell}{\partial \mathbf{v}_c^2} = -\frac{\partial \gamma}{\partial \mathbf{v}_c} = \begin{bmatrix} -\frac{\partial \gamma_t}{\partial \mathbf{v}_t} & \mathbf{0} \\ \mathbf{0}^T & -\frac{\partial \gamma_n}{\partial v_n} \end{bmatrix}$$

is positive definite for physics-based models that satisfy (8).

With the judicious choice  $F(s) = \sqrt{s^2 + 1} - 1$  from Section III-B, the cost, gradient, and Hessian involve soft norms (Appendix A) which are twice differentiable and convex, even at  $\mathbf{v}_t = \mathbf{0}$

$$\begin{aligned} \ell_t(\mathbf{v}_t) &= \mu \gamma_{n0} \|\mathbf{v}_t\|_s, \\ \gamma_t &= -\mu \gamma_{n,0} \hat{\mathbf{t}}_s, \\ \frac{\partial^2 \ell_t}{\partial \mathbf{v}_t^2} &= \mu \gamma_{n0} \frac{\mathbf{P}^\perp(\hat{\mathbf{t}}_s)}{\|\mathbf{v}_t\|_{s+\varepsilon_s}}. \end{aligned}$$

A remarkable property of this approximation is that it completely eliminates the *gliding* artifact characteristic of previous convex formulations. This is discussed in detail in Section VII-A.

### C. Similar Model

Similarity solutions to partial differential equations (PDEs) are solutions that depend on certain groupings of independent variables rather than each variable individually. In particular, self-similar solutions arise when the problem lacks a characteristic time or length scale. The Blasius solution to Prandtl's boundary layer equations in fluid mechanics is a well-known and celebrated example.

Motivated by the algebraic form of SAP impulses [4], we propose the grouping of variables  $z = v_n - \mu \|\mathbf{v}_t\|$ . Furthermore, we generalize this grouping to

$$z = v_n - \mu \varepsilon_s F(s), \quad (19)$$

where, as in Section V-B,  $\mu \varepsilon_s F(s)$  simplifies to  $\mu \|\mathbf{v}_t\|_s$  when  $F(s) = \sqrt{s^2 + 1} - 1$ . Note the consistency of units in (19), an important aspect of similar solutions. With this grouping, we propose the *similar* solution

$$\begin{aligned} \gamma_n(\|\mathbf{v}_t\|, v_n) &= n(z), \\ \gamma_t(\|\mathbf{v}_t\|, v_n) &= -\mu f(s) \gamma_n(\|\mathbf{v}_t\|, v_n) \hat{\mathbf{t}}. \end{aligned} \quad (20)$$

Unlike the Lagged model, the Similar model strongly couples friction and normal components. However, this introduces a dependency of the normal component on slip speed, an artifact we quantify in the following sections.

Differentiation of (20) leads to

$$\frac{\partial \gamma_t}{\partial v_n} = \frac{\partial \gamma_n}{\partial v_n} = -\mu f(s) n'(z) \hat{\mathbf{t}},$$

which confirms condition (16). To find the potential, we start from the normal component of the impulse

$$\gamma_n(\|\mathbf{v}_t\|, v_n) = n(z) = -\frac{\partial \ell}{\partial v_n},$$

and integrate on  $v_n$  to obtain

$$\ell(\mathbf{v}_t, v_n) = -N(z) + G(\mathbf{v}_t),$$

where  $G(\mathbf{v}_t)$  is an arbitrary function of  $\mathbf{v}_t$ . Taking the derivative with respect to  $\mathbf{v}_t$  results in

$$\frac{\partial \ell}{\partial \mathbf{v}_t} = \mu n(z) f(s) \hat{\mathbf{t}} + \frac{\partial G}{\partial \mathbf{v}_t}.$$

Comparing this equation with (20) reveals that we can set  $G = 0$  and obtain

$$\ell(\|\mathbf{v}_t\|, v_n) = -N(z), \quad (21)$$

as the desired potential function.

The Hessian of this potential is

$$\begin{aligned} \frac{\partial^2 \ell}{\partial \mathbf{v}_c^2} &= -\frac{\partial \gamma}{\partial \mathbf{v}_c} = \\ &= \mu \left[ \left( \frac{f'(s)n(z)}{\varepsilon_s} - n'(z)f^2(s) \right) \mathbf{P}(\hat{\mathbf{t}}) \right. \\ &\quad \left. + \frac{f(s)n(z)}{\|\mathbf{v}_t\|} \mathbf{P}^\perp(\hat{\mathbf{t}}) \right]. \end{aligned}$$

For a convex potential  $\ell_n$  we have  $n' = N'' = -\ell''_n \leq 0$ . With  $f(s)$  non-decreasing as in the lagged model,  $f' \geq 0$ . The Hessian  $\partial^2 \ell / \partial \mathbf{v}_c^2$  is the linear combination with positive coefficients of symmetric positive semi-definite projection matrices. Therefore, the Hessian is symmetric positive semi-definite, and the potential is convex. As with Lagged, the choice  $F(s) = \sqrt{s^2 + 1} - 1$  leads to continuously differentiable expressions in terms of soft norms (Appendix A), with no singularity at  $\mathbf{v}_t = \mathbf{0}$ .

Our *Similar* model closely relates to the primal formulation [5], [4] of Anitescu's convex approximation [2]. Following [4] we define the velocity  $\mathbf{g} = \mathbf{v}_c - \hat{\mathbf{v}}_c$ , with  $\hat{\mathbf{v}}_c = [0, 0, \hat{v}]^T$  and the *breaking velocity* from (23). Using this definition, we express  $z$  as  $z = g_n - \mu \varepsilon_s F(\|\mathbf{g}_t\|/\varepsilon_s) + \hat{v}$ . We notice that  $\ell(z)$  is constant (no contact) for  $z \geq \hat{v}$ . In terms of  $\mathbf{g}$ , this is equivalent to  $g_n - \mu \varepsilon_s F(\|\mathbf{g}_t\|/\varepsilon_s) \geq 0$ . The graph of  $g_n = \mu \varepsilon_s F(\|\mathbf{g}_t\|/\varepsilon_s)$  defines the boundary of the stick-slip transition, and therefore when  $F(s)$  is convex, its epigraph (the set of points above its graph), corresponds to a convex region (not necessarily a cone as in previous work). This is illustrated in Fig. 1. In this plot, contour levels of  $\ell(\|\mathbf{g}_t\|, g_n)$  correspond to lines of constant  $z$ , which by definition are perpendicular to the gradient  $\partial \ell / \partial \mathbf{v}_c$ . We see that the role of the potential  $\ell$  is to penalize  $\mathbf{g}$  when it lies outside the epigraph of  $g_n = \mu \varepsilon_s F(\|\mathbf{g}_t\|/\varepsilon_s)$ .

Consider  $F(s) = \sqrt{s^2 + 1} - 1$ , for which  $g_n = \mu \|\mathbf{g}_t\|_s$ . The epigraph of this function is an approximation to the dual  $\mathcal{F}^*$  of the friction cone  $\mathcal{F}$ , Fig. 1. This approximation is  $\mathcal{F}^*$  in the limit  $\varepsilon_s \rightarrow 0$ . Moreover, in the limit to rigid contact,  $\ell$  enforces  $\mathbf{g} \in \mathcal{F}^*$ , which corresponds to the cone constraint in the primal formulation [5], [4].

## VI. CONVEX APPROXIMATION OF COMPLIANT CONTACT WITH HUNT & CROSSLEY DISSIPATION

So far, we have considered a generic functional form of the compliant law in (6). To close our model, we consider a linear elastic law in the signed distance  $\phi$  with Hunt & Crossley [8] dissipation. As long as conditions (8) are met, other alternatives such as the Hertz model can be used. In practice,




Fig. 1: The *Similar* model penalizes  $\mathbf{g}$  when it lies outside the region  $g_n \geq \mu \varepsilon_s F(\|\mathbf{g}_t\|/\varepsilon_s)$ . With  $F(s) = \sqrt{s^2 + 1} - 1$  and  $\varepsilon_s \rightarrow 0$ , the model enforces  $\mathbf{g} \in \mathcal{F}^*$  in the limit to rigid contact, consistent with the primal formulations in [5], [4].

however, users often find the linear model satisfactory. We write this model within framework (6) as

$$f_n(\phi, v_n) = k(-\phi)_+ (1 - dv_n)_+, \quad (22)$$

where  $k$  is the linear contact stiffness,  $d$  is the Hunt & Crossley dissipation parameter and  $(a)_+ = \max(0, a)$ . This force is zero whenever  $v_n \geq \hat{v}$ , with

$$\hat{v} = \min(-\phi_0/\delta t, 1/d), \quad (23)$$

the minimum normal velocity for contact to break. For  $v_n < \hat{v}$  we define the (indefinite) antiderivative  $N^+(v_n)$

$$N^+(v_n; \phi_0) = \delta t k \left[ -v_n \left( \phi_0 + \frac{1}{2} \delta t v_n \right) + d \frac{v_n^2}{2} \left( \phi_0 + \frac{2}{3} \delta t v_n \right) \right].$$

Since the impulse is zero for  $v_n \geq \hat{v}$ , its antiderivative must be constant. Therefore, we write

$$N(v_n) = N^+(\min(v_n, \hat{v}); \phi_0),$$

resulting in a continuous function for all values of  $v_n$ .

Though inherently compliant, this model effectively handles very stiff contact due to the robustness of our convex formulation. In Section XI, we use Hertz theory to estimate steel stiffness, and Section XII demonstrates that our approach can manage stiffnesses up to five orders of magnitude higher.

Finally, we observe that this model is easily extended to incorporate *hydroelastic contact* [20] — a modern rendition of *Elastic Foundation* [17, §4.3] to model continuous contact patches. We use the approach in [11].

## VII. COMPARATIVE ANALYSIS OF THE CONVEX APPROXIMATIONS

All models in Section V are convex approximations of contact, each with its own strengths and limitations. Hereinafter, we refer to each of these formulations as SAP [4], *Lagged* (section V-B), and *Similar* (section V-C). We summarize their properties and artifacts in Table I and analyze them in detail in the following subsections. For the Lagged and Similar models, we use the compliant model with Hunt & Crossley dissipation from Section VI.

### A. Gliding During Slip

SAP, as well as the formulations from Anitescu and Todorov [2], [3], introduce a number of artifacts as the result of the convex approximation of contact. The most well-known artifact is that objects *glide* at a finite offset during slip [5], [21], [4]. For SAP, this offset is [4]

$$\phi_{\text{offset}} = \mu(\delta t + \tau_d)\|\mathbf{v}_t\|. \quad (24)$$

With zero dissipation ( $\tau_d = 0$ ), this reduces to the gliding distance  $\mu\delta t\|\mathbf{v}_t\|$  in Anitescu's formulation of rigid contact [2]. There are no artifacts during stiction.

While the term  $\mu\delta t\|\mathbf{v}_t\|$  can be negligible for small time steps and slip speeds, the term  $\mu\tau_d\|\mathbf{v}_t\|$  can dominate when dissipation is large. Therefore, users must face the trade-off between accurate dissipation modeling and the magnitude of this gliding artifact, making the choice of this parameter cumbersome.

The Similar approximation eliminates this coupling with dissipation, though it still introduces the gliding artifact. This is caused by the non-physical coupling of the tangential velocity into the normal component of the contact impulse in (20). Replacing  $v_n$  with  $z$  from (19) into (22), we find that the impulse goes to zero at  $\phi_{\text{offset}} = \mu\delta t\|\mathbf{v}_t\|$  rather than at  $\phi = 0$ . Gliding during slip for SAP can be traced back to this same artificial coupling of the tangential velocity into the normal force; see [4] for details.

In contrast, the Lagged approximation does not introduce gliding, since it completely eliminates the contribution of the tangential velocity in the normal force component.

### B. Compliance Modulation

For SAP, the effective compliance during slip is  $k_{\text{eff}} = k/(1 + \tilde{\mu}^2)$ , where  $\tilde{\mu} = \mu(R_t/R_n)^{1/2}$  [4] depends on the ratio of the tangential to normal regularization parameters. SAP chooses this ratio so that  $\tilde{\mu} \approx 0$ , but it is still non-zero. We refer to this artifact as *compliance modulation*, and it was first reported in [4].

Replacing  $v_n$  with  $z$  in (22) and factoring out the term  $1 + \mu d\|\mathbf{v}_t\|$ , we see that during slip, the Similar approximation has effective stiffness  $k_{\text{eff}} = k(1 + \mu d\|\mathbf{v}_t\|)$  and dissipation  $d_{\text{eff}} = d/(1 + \mu d\|\mathbf{v}_t\|)$ . That is, the effective stiffness increases while the effective dissipation decreases during slip.

The Lagged approximation does not introduce compliance modulation.

### C. Strong Coupling

We use the term *strong coupling* to describe the tight, algebraic, coupling of the friction and normal components. While SAP and Similar approximations do introduce strong coupling, the Lagged approximation lags the normal force in Coulomb's law. Our numerical studies (Sections XI and XII) show that this lag effect is noticeable only in high-energy impacts where normal forces change rapidly. Otherwise, the approximation is  $\mathcal{O}(\delta t)$ , consistent with the first-order Taylor expansion of the signed distance introduced in Section III-B. Moreover, our grasp stability analysis (Section XII-B) confirms that the

Model	Consistent	Strongly coupled	Gliding during slip	Compliance modulation		Transition slip
				Stiffness	Dissipation	
SAP	No	Yes	$\mu(\delta t + \tau_d)\ \mathbf{v}_t\ $	$k/(1 + \tilde{\mu}^2)$	$\tau_d$	$\sigma w \mu \gamma_n$
Lagged	Yes	No	—	—	—	$v_s$
Similar	No	Yes	$\mu \delta t \ \mathbf{v}_t\ $	$k(1 + \mu d \ \mathbf{v}_t\ )$	$d/(1 + \mu d \ \mathbf{v}_t\ )$	$v_s$

TABLE I: Properties of each convex approximation in Section V.

lagged approximation is suitable, even for highly dynamic manipulation tasks.

#### D. Consistency

In the context of ordinary differential equations (ODEs), a scheme is *consistent* when the original set of ODEs is recovered in the limit  $\delta t \rightarrow 0$ . This is not the case for SAP and Similar models, due to the existence of *compliance modulation* (Section VII-B), regardless of the time step size. In particular for SAP, the gliding artifact does not vanish in (24) unless dissipation is zero. Conversely, Lagged is a consistent first order approximation of (11).

#### E. Stick-Slip Transition

Since the models in this work *regularize* friction, the transition between stiction and sliding occurs at a finite slip velocity. For Lagged and Similar models, this transition is parameterized by  $\varepsilon_s$  in (9), and it is fixed by the user to a value  $v_s$  which we call the *stiction tolerance*.

For SAP, we see from (17) that stick-slip transitions occur at a magnitude of slip equal to

$$v_s = \sigma w \mu \gamma_n, \quad (25)$$

where we used the fact that SAP uses  $R_t = \sigma w$ , with  $w$  a diagonal approximation of the Delassus operator [4] (the *effective inverse mass* of the contact). The dimensionless parameter  $\sigma$  is set to  $\sigma = 10^{-3}$ , a good trade-off between numerical conditioning and a tight approximation of stiction, as required for the simulation of manipulation applications in robotics. Refer to [4] for a thorough study of SAP's properties.

### VIII. IMPACTS, STIFFNESS AND CONDITIONING

We are interested in the numerical *stiffness* introduced by the regularization of friction as it is critical for a robust numerical implementation. We define  $G_t = d\|\gamma_t\|/d\|\mathbf{v}_t\|$  and examine its value at  $\|\mathbf{v}_t\| = 0$ . For all approximations we find

$$G_t = \frac{\mu}{v_s} \tilde{\gamma}_n,$$

with  $v_s$  the stiction tolerance (Table I),  $\tilde{\gamma}_n = \gamma_n$  for SAP and Similar models, and  $\tilde{\gamma}_n = \gamma_{n,0}$  for Lagged.

As seen in (25), SAP's stiction tolerance depends on impulse, often much higher during impacts than in sustained contact (e.g., a robot grasping an object). This effectively increases SAP's regularization during impacts, making the problem better-conditioned but reducing stiction accuracy. In contrast, Lagged and Similar solve a much stiffer problem given the stiction tolerance  $v_s$  is constant. We confirm this in Section XII-A in a case with impacts.

The situation reverses during sustained contact, as the normal impulse depends mainly on object's weight, and external forces, not rapid velocity changes. Since impulse scales with time step size,  $G_t = \mathcal{O}(\delta t)$  for Similar and Lagged, while SAP's  $G_t = 1/R_t$  remains constant (with  $R_t$  from (17)), leading to a stiffer problem. This is confirmed in Section XII-A as we analyze the numerical conditioning of the problem and the solver's number of iterations.

It could be argued that, for most robotics applications, accurately resolving stiction during impacts is unnecessary. Therefore, we propose a *regularized* version of the Lagged model, where the regularization parameter is

$$\varepsilon_s = \max(v_s, \sigma w \mu \gamma_{n,0}). \quad (26)$$

In this formulation, regularization is *softened* during strong impacts when values of  $\gamma_{n,0}$  are higher, but a tight value bounded by  $v_s$  is used for sustained contacts. We study the effect of this formulation in Section XII-A.

### IX. BARRIER FUNCTIONS AS COMPLIANT CONTACT

While in this work we find that compliance provides a reliable approximation of contact, we cannot neglect the fact that most simulation software in the robotics and graphics communities implement rigid approximations of contact. It is important to point out that rigid contact, as well as compliant contact, are both approximations of the real physics, each with their strengths and limitations.

In this section, we provide a discussion on the connection between these two apparently unrelated technologies, in particular when barrier or similarly smooth functions are used to enforce the non-penetration and complementarity conditions appearing in rigid approximations [22], [23], [24]. We have not implemented any of these barriers in our framework, though this could be an interesting avenue of future research.

#### A. Logarithmic Barriers for Interior Point Methods

Interior-point methods (IP) are the standard for solving optimization problems with inequality constraints, essential for rigid contact modeling. IP solvers can also handle a richer set of constraints, like conic constraints, as needed to model Coulomb friction. Notable solvers include open-source Ipopt [25] and proprietary options like Gurobi [26] and Mosek [27]. Fundamentally, IP methods replace constraints with a barrier function that penalizes infeasible solutions. Logarithmic functions are commonly used, which for rigid contact penalize distances near zero

$$\ell_n(\phi) = -\kappa \ln(\phi),$$

while penetration states ( $\phi \leq 0$ ) are infeasible. The barrier parameter  $\kappa > 0$  is iteratively reduced to approach a rigid approximation, though it can never reach zero. Therefore, in practice the solution is effectively compliant, with a force law that fits the framework (6)

$$f_n(\phi, v_n) = \frac{\kappa}{\phi} \quad \phi \geq 0. \quad (27)$$

As a physical model of contact, users may struggle to interpret (27), which has a parameter  $\kappa$  with the unit of energy and nonphysical action at infinite distance. In practice,  $\kappa$  is not exposed, but hidden as part of the solver internals. Users *believe* they are working with a true model of rigid contact when, in reality, the solver is using a compliant model approximation. Even if  $\kappa$  can be reduced to very small values (according to some hidden metric), the solver often ends up solving a much more challenging problem than needed since, in reality, physical materials have finite stiffness.

### B. Incremental Potential Contact

IPC [23] is an optimization-based framework to model rigid contact and guarantee intersection-free solutions. The method attains strong robustness given its implicit time-stepping scheme and line search augmented with continuous collision detection (CCD) to maintain feasibility. However, IPC formulates a non-convex optimization problem and can fall into local minima, not satisfying the original physical laws. Moreover, the method lacks convergence guarantees, as properly pointed out by the original authors.

Similar to the logarithmic barrier functions of IP methods, IPC proposes a  $C^2$  potential

$$\ell_c(\phi) = -\kappa_{\text{IPC}} (\hat{d} - \phi)_+^2 \ln(\phi/\hat{d})$$

where  $\kappa_{\text{IPC}}$  is a parameter automatically adjusted to improve numerical conditioning, and  $\hat{d}$  is a user parameter. Typical values of  $\hat{d}$  used by the original authors in their extensive simulations test cases are in the range 0.1 mm to 1 mm. This potential is proposed to achieve intersection-free solutions, eliminate nonphysical action at infinite distances, and maintain smoothness for better numerics. We observe that this method again fits the framework (6), with a *compliant* contact force law of the form

$$f_n(\phi, v_n) = \kappa_{\text{IPC}} (\hat{d} - \phi)_+ \left[ \frac{(\hat{d} - \phi)_+}{\phi} - 2 \ln(\phi/\hat{d}) \right],$$

which is only non-zero for  $\phi \in (0, \hat{d}]$ . Performing Taylor expansion around  $\phi = \hat{d}$ , we see that  $f_n \approx 3\kappa_{\text{IPC}}(\hat{d} - \phi)_+^2/\hat{d}$ , and the force models a quadratic spring of stiffness  $\kappa_{\text{IPC}}$  (with units of N/m). In the limit to  $\phi \rightarrow 0^+$ , the force approximates  $f_n \approx \kappa_{\text{IPC}}\hat{d}^2/\phi$ , the interior point force (27).

In summary, this method models a thin, compliant layer around a rigid core instead of the strict non-penetration conditions largely favored in the literature. We do not view this as a flaw, as the authors have demonstrated their method's effectiveness through extensive simulation studies. Instead, we see this as an indication of the levels of rigidity that can be achieved in practice.

## X. IMPLEMENTATION

Our work is implemented in Drake [9], a robotics toolkit that provides modeling abstractions and optimization tools for the modeling, simulation, and analysis of robotics systems. Our implementation includes support for deformable Finite Element Models (FEM), holonomic constraints, PD controllers with effort limits, reflected inertia, and the modeling of continuous contact patches with *Hydroelastic Contact*. Refer to our previous work for further details [4], [11], [12].

SAP [4] uses Newton's method to compute a search direction  $\Delta \mathbf{v}$  for (3) according to

$$\mathbf{H}(\mathbf{v})\Delta \mathbf{v} = -\mathbf{r}(\mathbf{v}), \quad (28)$$

where  $\mathbf{r}(\mathbf{v}) = \partial \ell / \partial \mathbf{v}$  is the *residual* of the primal cost  $\ell(\mathbf{v})$ , and  $\mathbf{H}(\mathbf{v})$  is its Hessian. Upon convergence, the residual  $\mathbf{r}(\mathbf{v})$  is zero, which is equivalent to the linearized balance of momentum (2). The solution at iteration  $m$  is updated as

$$\mathbf{v}_{m+1} = \mathbf{v}_m + \alpha \Delta \mathbf{v},$$

with  $\alpha$  determined via a line search along  $\Delta \mathbf{v}$  to minimize  $\ell(\mathbf{v})$ . Our line search uses a Newton-Raphson method based on [28, §9.4], augmented with bracketing and bisection to guarantee convergence. The derivatives required for the line search are computed with  $\mathcal{O}(n)$  complexity [4], enabling convergence to machine precision with negligible computational overhead. This strategy has proven to be very robust in practice.

Multibody *tree* structures create distinct cliques, while degrees of freedom (DOFs) for modeling a deformable body are also grouped into their own cliques [12]. Similarly, constraints involving the same pair of cliques are grouped into *clusters*. We exploit this structure using a supernodal Cholesky factorization [29, §9] of the Hessian  $\mathbf{H}$  in (28) that takes advantage of dense algebra optimizations. We compute the elimination ordering of the supernodes using approximate minimum degree (AMD) ordering [30] to minimize fill-ins.

### A. Geometric Queries

Unlike other approaches, our method performs only a single geometric query at the start of each time step, avoiding repeated queries during Newton iterations. Changes in contact configuration are captured using the first-order approximation of the signed distance discussed in Section III-B. This approximation is extended in [11] to model continuous contact surfaces with hydroelastic contact [20], implemented as part of Drake and accelerated with an oriented bounding box hierarchy. Point contact queries use FCL [31]. All simulations in this work include contact pairs with signed distance below a fixed threshold,  $\phi < 10$  cm, including contact pairs even before physical contact is established.

### B. Differentiation Through Contact

Our implementation in Drake provides an end-to-end solution for the computation of gradients through contact for applications such as system identification, reinforcement learning, and trajectory optimization. While differentiation through contact is not new [32], [22], [33], our approach is unique

in that we use automatic differentiation to compute gradients through geometry queries and the dynamics ( $\partial \mathbf{r}/\partial \boldsymbol{\theta}$  below) and the *implicit function theorem* to propagate these gradients to the next time solution.

We denote differentiation parameters with  $\boldsymbol{\theta} \in \mathbb{R}^{n_\theta}$ , which can include physical quantities such as mass and inertias, contact parameters, actuation, and even the previous state of the multibody system. Using this notation, SAP's optimality condition can be written in terms of the residual in (28) as  $\mathbf{r}(\mathbf{v}; \boldsymbol{\theta}) = \mathbf{0}$ . In this notation, the residual is a function of the generalized velocities  $\mathbf{v}$  given a set of parameters  $\boldsymbol{\theta}$ .

We use the implicit function theorem on  $\mathbf{r}(\mathbf{v}; \boldsymbol{\theta}) = \mathbf{0}$

$$\frac{\partial \mathbf{r}(\mathbf{v}; \boldsymbol{\theta})}{\partial \mathbf{v}} \frac{d\mathbf{v}}{d\boldsymbol{\theta}} + \frac{\partial \mathbf{r}(\mathbf{v}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{0},$$

and note that  $\partial \mathbf{r}/\partial \mathbf{v} = \mathbf{H}(\mathbf{v})$  to write

$$\mathbf{H} \frac{d\mathbf{v}}{d\boldsymbol{\theta}} = -\frac{\partial \mathbf{r}(\mathbf{v}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}. \quad (29)$$

In our approach, the expensive-to-compute Cholesky factorization of  $\mathbf{H}$  is only computed at each Newton iteration in (28) during forward simulation. Upon convergence,  $\mathbf{H}$  is already assembled and factorized and is thus reused in (29) an additional  $n_\theta$  times to propagate derivatives through the solver into gradients  $d\mathbf{v}/d\boldsymbol{\theta}$  of the generalized velocities.

In our hybrid approach,  $\partial \mathbf{r}/\partial \boldsymbol{\theta}$  in (29) is computed with automatic differentiation. This enables the computation of gradients through arbitrarily complex geometric models, while the implicit function theorem propagates derivatives accurately and efficiently through the contact resolution phase.

## XI. TEST CASES

We analyze a series of two-dimensional cases to assess accuracy, quantify artifacts introduced by the convex approximations, and gain intuition into the physics and numerics. Here, we focus on comparing our convex formulations. We refer the reader to [21], [7] for previous work comparing other approaches, including Anitescu's convex formulation.

For all cases, we use  $v_s = 10^{-4}$  m/s for Lagged and Similar, and  $\sigma = 10^{-3}$  for SAP, leading to very tight stiction modeling as required for simulating manipulation tasks.

We estimate contact stiffness using Hertz theory. For a sphere of mass  $m$  and radius  $R$ , Hertz theory predicts a penetration  $\delta = (3mg/(4ER^{1/2}))^{2/3}$ . For steel with Young's modulus  $E = 200$  GPa and using the radii and masses from Sections XI-B and XI-C, we obtain penetrations around  $\delta \approx 2.5 \times 10^{-7}$  m and stiffnesses  $k \approx 1 \times 10^7 - 2 \times 10^7$  N/m. We use  $k = 10^7$  N/m for all cases in this section.

For some cases, we perform a convergence study where we compute the error in the positions  $\mathbf{q}_{\delta t}$  obtained using step size  $\delta t$  against a reference  $\mathbf{q}_{\text{ref}}$  as a function of the time step size

$$e_q(\delta t) = \left( \frac{1}{T} \int_0^T dt \|\mathbf{q}_{\delta t}(t) - \mathbf{q}_{\text{ref}}(t)\|^2 \right)^{1/2}$$

where  $T$  is simulation duration. The reference solution is obtained numerically using a time step 10 times smaller than the smallest time step in the convergence study. Since Lagged

is the only approximation that is consistent (Section VII-D), we use it to compute the reference solution.

### A. Oscillating Conveyor Belt

This test illustrates artifacts in the strongly coupled SAP and Similar approximations. A 1 kg box with 5 cm sides is placed on a conveyor belt oscillating at 1 Hz with 0.2 m amplitude (Fig. 2). Friction is  $\mu = 0.7$ . Even though dissipation models are different, using  $d = 500$  s/m for Similar and Lagged and  $\tau_d = 10^{-3}$  s for SAP yields comparable dissipation.




Fig. 2: Oscillating Conveyor Belt. SAP and Similar introduce artificial *gliding* during slip phases (Section VII-A).

Figure 3 shows contact velocity and force computed using  $\delta t = 0.01$  s. Contact between the box and belt transitions back and forth between stiction and sliding. The Lagged approximation predicts no vertical motion, with the normal force balancing the box's weight as expected. However, SAP and Similar models show artifacts in the normal direction during sliding, with non-zero normal velocity due to the *gliding* artifact, which disappears in stiction. Additionally, we observe spurious transients in the normal force during sliding — as slip speed changes so does *gliding*, causing vertical acceleration and thus normal force fluctuations. Finally, SAP and Similar introduce normal force spikes during the abrupt transition from sliding to stiction, when *gliding* vanishes causing a sudden normal velocity change (Fig. 3).




Fig. 3: Contact velocity (left) and force (right).  $\delta t = 0.01$  s.

Figure 4 shows a convergence study with step sizes  $\delta t \in \{2 \times 10^{-3}, 10^{-2}, 5 \times 10^{-2}\}$ . Both Lagged and Similar exhibit first order convergence, as expected, though the artifacts introduced by Similar (Fig. 3), cause higher errors than Lagged. Finally, SAP's error plateaus at the smallest time step due to model inconsistency, where the term  $\tau_d \mu \|v_t\|$  in (24) does not vanish as time step decreases.




Fig. 4: Convergence of the box trajectory with time step size. The dashed line is a reference for first order convergence.

### B. Falling Sphere

The conveyor belt case from the previous section favors the Lagged model due to the steady-state normal force, where the Lagged model is exact. Here, we introduce a collision test to evaluate approximations under sudden contact force changes.

In this test, a 0.5 kg steel sphere 5 cm in diameter falls from a height of 5 cm with an initial horizontal velocity  $U_0 = 2$  m/s, see Fig. 5. Friction with the ground is  $\mu = 0.5$ . Upon impact, the sphere slides, and then transitions to rolling as friction induces angular momentum. After this transition, friction ceases to dissipate energy. We model compliant contact with stiffness  $k = 10^7$  N/m and dissipation constants  $d = 500$  s/m and  $\tau_d = 10^{-3}$  s.




Fig. 5: Falling sphere. After free fall, the sphere slides until friction with the ground establishes a rolling contact.

Figure 6 shows contact velocity and force computed with  $\delta t = 2 \times 10^{-3}$  s. In the force plots, both SAP and Similar initiate contact earlier due to the *action at a distance* artifact in these convex models, with SAP engaging even earlier due to the non-vanishing term  $\tau_d \mu \|v_t\|$  in (24) (Section VII-A). The sphere slides from initial contact until about  $t = 1.1$  s, when it transitions to rolling. While Lagged brings normal velocity to zero almost instantly, SAP and Similar models link normal velocity to slip velocity, only reaching zero at stiction. During the sliding-to-stiction transition, we observe a rapid normal force spike, as with the conveyor belt problem. This artifact is absent with the Lagged approximation. In all cases, the transition from sliding to pure rolling is predicted without the need for specialized rolling constraints. A convergence study with step sizes  $\delta t \in \{4 \times 10^{-4}, 2 \times 10^{-3}, 10^{-2}\}$  is shown in

Fig. 7. While all of these schemes are first order, SAP exhibits a constant error at convergence due to the  $\tau_d \mu \|v_t\|$  term in (24).




Fig. 6: Contact velocity (left) and force (right) with  $\delta t = 2 \times 10^{-3}$  s.




Fig. 7: Convergence of the sphere trajectory with time step size. The dashed line is a reference for first order convergence.

### C. Sliding Rod

This case is particularly interesting as it leads to impact without collision [34, §5.3]. A rod initially angled with the ground makes single-point contact with a horizontal velocity (see Fig. 8). As it slides, friction rotates the rod into the ground, increasing the normal force. Under specific conditions, both normal and frictional forces intensify, potentially leading to a singularity in acceleration-level formulations with Coulomb friction known as Painlevé's paradox, where forces become infinite. This problem is resolved in the discrete setting, where finite impulses and discrete velocity changes are allowed. Physically, bodies aren't perfectly rigid — they deform, vibrate, and may even undergo plastic (permanent) deformations. Nonetheless, a rapidly increasing contact force develops an impact that makes the rod jam into the ground and jump into the air. The rod measures 0.5 m in length, 1 cm in diameter, and has a mass of 0.3 kg.

Analytical analysis [34, §5.3] shows that the singularity occurs when  $\mu > 4/3$  and the initial kinetic energy overcomes potential energy as the rod's center of gravity rises and friction dissipates energy. We set  $\mu = 2.3$ , a  $30^\circ$  initial angle (Fig. 8), and an initial horizontal speed  $U_0 = 10$  m/s.

Using a reference solution with a time step of  $\delta t = 10^{-7}$  s and no normal force dissipation, we observe that the rod rotates




Fig. 8: Sliding rod. Initially forming an angle  $\varphi_0$  with the ground and with horizontal velocity  $U_0$ . Friction makes the rod rotate clockwise. The contact force increases until the rod jams into the ground, causing the rod to jump into the air.

upward and jams into the ground upon contact as expected. All three models yield similar results, as the compliant model is identical in the absence of dissipation, differing only in friction regularization. Pre-impact forces oscillate based on ground compliance, and impact location is nearly identical across models ( Fig. 9) despite being very sensitive to model parameters.




Fig. 9: Contact forces in the case with zero dissipation. The figure on the right shows a close-up near the impact.

Figure 10 shows the contact point moves into the ground due to large contact forces until the tangential component of the velocity goes to zero and the rod jams into the ground. During stiction  $\|v_t\| < v_s$  for a finite period of about 0.2 ms.




Fig. 10: Contact velocities in the case with zero dissipation. The figure on the right shows a close-up near the impact.

We run the simulation with Hunt & Crossley dissipation  $d = 0.2$  s/m and relaxation time  $\tau_d = 4.0 \times 10^{-6}$  s. These low dissipation values have minimal effect on the time of impact for the Lagged model, our reference solution, as seen in Figs. 11 and 12. However, Similar and SAP models predict shifted impact times—earlier for Similar, later for SAP. Although their contact forces and velocities differ significantly, we choose  $\tau_d$  for SAP to match the shift in time of impact observed in the

Similar model, albeit in the opposite direction. Compliance modulation in Similar becomes evident in Fig. 11, where we observe a frequency shift on the force oscillations. This is caused by the larger effective stiffness of the model during sliding,  $k_{\text{eff}} = k(1 + \mu\|v_t\|d)$ .




Fig. 11: Contact forces in the case with dissipation. The figure on the right shows a close-up near the impact.

A convergence analysis with and without dissipation is shown in Fig. 13, using time steps  $\delta t = \{6.4 \times 10^{-4}, 1.6 \times 10^{-4}, 4.0 \times 10^{-5}, 1.0 \times 10^{-5}\}$  s. Without dissipation, SAP and Similar model solutions are indistinguishable. The Lagged model exhibits the largest error at the largest time step, revealing a limitation: its lagged normal force affects Coulomb friction modeling in rapidly changing scenarios, even missing impacts when  $\delta t > 10^{-3}$  s. However, it achieves first-order convergence with smaller steps. Similar trends occur with non-zero dissipation, though the curves diverge as time of impact predictions shift.




Fig. 12: Contact velocities in the case with dissipation. The figure on the right shows a close-up near the impact.




Fig. 13: Convergence in positions with time step. Case without dissipation (left) and with dissipation (right).

## XII. APPLICATIONS

We simulate various robotics-relevant cases to evaluate model usefulness, accuracy, numerical conditioning, and solver

performance. In all simulations, the solver iterates fully to convergence with no early termination, using a relative tolerance  $\varepsilon_r = 10^{-5}$  [4].

### A. Clutter

We reproduce the clutter experiment from [4] to evaluate solver performance in cluttered environments, common in robotic manipulation. We drop 40 objects, arranged in four columns of 10, into an  $80 \times 80 \times 80$  cm box, see Fig. 14. Each column has a mix of 10 cm diameter spheres and 10 cm boxes, with masses calculated using water density: 0.524 kg for spheres and 1.0 kg for boxes. A high stiffness of  $k = 10^7$  N/m models steel, with Hunt & Crossley dissipation set to  $d = 10$  s/m and SAP dissipation at  $\tau_d = 10^{-4}$  s. Lagged and Similar models use a stiction tolerance of  $v_s = 10^{-4}$  m/s, while SAP and the regularized Lagged model use  $\sigma = 10^{-3}$ . All surfaces have a friction coefficient  $\mu = 1.0$ . We let objects fall and simulate for 3 seconds.




Fig. 14: Initial condition (left) and final steady state at  $t = 3$  seconds (right).

We analyze each method's performance with a time step  $\delta t = 2 \times 10^{-3}$  s. Figure 15 shows Newton solver iterations and Hessian condition numbers over time. An intense initial transient occurs as objects collide upon falling, with impacts subsiding around  $t = 2$  s as objects settle. During this early phase, the solver requires more iterations, and condition numbers are higher. Notably, Lagged and Similar approximations show higher condition numbers and iterations during the initial transient, whereas the situation is reversed for SAP. This behavior relates to the stiffness  $G_t$  with regularized friction (Section VIII). Using the *regularized* stiction tolerance (26) in the Lagged approximation reduces iterations and improves conditioning during the initial phase. Though not included here, the same regularization can be used with Similar.

Figure 16 shows the effective stiction tolerance. For Lagged and Similar approximations, stiction tolerance is fixed at  $v_s = 10^{-4}$  m/s (dashed black line). In contrast, SAP and the regularized Lagged approximation have a tolerance that varies with normal impulse (Section VIII). This aligns with prior observations: during the initial transient, Lagged and Similar models enforce a tighter friction approximation. Past this transient, SAP unnecessarily solves a much tighter approximation.




Fig. 15: Iterations (left) and condition number (right) as a function of time.  $\delta t = 2 \times 10^{-3}$  s.




Fig. 16: Effective stiction tolerance as a function of time.  $\delta t = 2 \times 10^{-3}$  s.

An informative metric is the mean iteration count and condition number as a function of time step, Fig. 17. For SAP, the condition number remains almost unchanged across time step sizes since  $G_t = 1/R_t$  (see Section VIII) is constant. In contrast, the conditioning of Lagged and Similar models improves as the time step decreases, because  $G_t$  is proportional to impulse and time step size. This reduction explains why the number of iterations decreases for Lagged and Similar as the time step decreases, while for SAP, it remains almost constant. Figure 17 also highlights the benefit of regularization for Lagged at large time steps. At small time steps,  $v_s$  dominates in (26), making Lagged with and without regularization perform similarly.




Fig. 17: Mean number of iterations (left) and condition number (right) per time step as a function of time step size.

We examine compliance as a method for approximating *rigid* contact by measuring mean penetration distance over the last 1.25 seconds, when objects settle at the bottom of the box. Figure 18 shows steady-state penetration across stiffness

values spanning eight orders of magnitude. For reference, Hertz theory predicts a stiffness of  $10^7$  N/m for steel. We stress-test with stiffness up to five orders higher, confirming the robustness of the convex formulation. For a time step  $\delta t = 0.005$  s, SAP’s *near-rigid* [4] stiffness estimate is  $10^6$  N/m, with penetration at only tenths of microns. At an extreme, nonphysical  $k = 10^{13}$  N/m, the solver fails due to round-off errors, while  $10^5 - 10^6$  N/m suffices for approximating rigid contact in typical robotics applications.




Fig. 18: Mean penetration distance in steady state during the last 1.25 s of the simulation.

Finally, figure 19 shows the effect of stiffness on performance. As expected, higher stiffness values degrade conditioning and ultimately impair performance. However, we note that the performance degradation is minimal — only within 20%, even for stiffness values as high as those of steel.




Fig. 19: Mean number of iterations (left) and condition number (right) per time step as a function of stiffness.

### B. Grasp Stress Test

To evaluate grasp stability for Lagged and Similar approximations, we simulate a Franka hand holding a bronze rod (density  $8000$  kg/m<sup>3</sup>) following a 3 cm circular horizontal trajectory (Fig. 20). At low frequencies, the rod is secure, but near the rod’s compound pendulum frequency, it begins to rock, slide, and ultimately fall. We refer to this time as *time to failure*  $T_f$ .

We identify three error sources:  $\mathcal{O}(\delta t)$  truncation error in the symplectic Euler scheme (1), gliding error ( $\mu \delta t \|v_i\|$ ), and compliance modulation. Tests are designed to evaluate their significance and determine if weak coupling in the Lagged approximation (Section VII-C) has a measurable impact compared to these errors.




Fig. 20: Franka hand holding a rod. Initial condition (left) and near grasp failure (right).

We compute reference solutions with Lagged (recall it is *consistent*, Section VII-D) at a small time step of 0.2 ms to minimize truncation errors (validated via a refinement study). To assess compliance modulation, we compare solutions (Fig. 21) with dissipation ( $d = 50$  s/m) and without ( $d = 0$ ). Below 1.4 Hz, the grasp remains stable, but  $T_f$  decreases rapidly near the rod’s compound pendulum frequency. Dissipation slightly improves stability, increasing  $T_f$  by up to 10%.




Fig. 21: Time to Failure  $T_f$  with and without Hunt & Crossley dissipation. Reference solutions with  $\delta t = 0.2$  ms.

We compute the relative error in  $T_f$  against the reference solutions (Fig. 22), with positive values indicating overestimation. At  $\delta t = 0.2$  ms and zero dissipation, Lagged and Similar solutions differ by less than 0.5% (Fig. 22, left), indicating that Similar’s *gliding* contribution  $\mu \delta t \|v_i\|$  is negligible.

At large  $\delta t = 10$  ms errors increase up to 50% due to how incredibly sensitive  $T_f$  is. The goal however is to understand the relative importance of each error contribution rather than a precise determination of  $T_f$ . Without dissipation, truncation errors dominate, causing both approximations to underpredict  $T_f$ .  $T_f$  with Similar is slightly larger compared to Lagged (error is less negative), as the gliding effect introduces transient penetrations  $\delta t \mu \|v_i\|$  that increase the mean grasp force.

From Fig. 22 (left), Similar’s gliding is negligible at 0.2 ms. Thus, the 30% error increase in Fig. 22 (right) is fully attributed to compliance modulation, comparable to truncation error.




Fig. 22: Relative error in  $T_f$  without (left) and with (right) dissipation. Positive values indicates  $T_f$  is over predicted.

At  $\delta t = 10$  ms, Lagged’s error remains nearly unchanged with or without dissipation. Similar’s error decreases due to the cancellation of two effects: truncation errors causing  $T_f$  underestimation (Fig. 22, left) and compliance modulation leading to overprediction (Fig. 22, right).

At large time steps for interactive simulation, both models exhibit similar error magnitudes. Even in this highly dynamic case, Lagged’s weak coupling does not degrade grasp performance. Since Lagged eliminates *gliding* and *compliance modulation*, we recommend it over Similar and SAP.

### C. BarrettHand

We demonstrate our method’s predictive capability by simulating the BarrettHand’s unique gearing mechanism [10], which relies on friction and stick-slip transitions. The proximal and distal links of a finger are driven by a single motor through a coupled gear system. The proximal gear (blue in Fig. 23) rides freely on internal threads along the shaft of the adjacent worm gear. Belleville washers at the end of the threads act as a clutch for the proximal gear. When compressed by the proximal gear (Fig. 23 left), the washers’ stiction holds the proximal gear stationary relative to the worm gear, transmitting torque from the motor to the proximal link. When the proximal link reaches an external force limit, the torque causes the proximal gear to slip and break away from the washers (Fig. 23 right). At this point, the motor torque transfers solely to the distal link. Stiction between the proximal wheel (purple) and worm gear (green) makes the system non-backdrivable, holding the proximal link in place.

We model a single BarrettHand finger along with its entire motor and gear system. Our method captures the characteristic loading, driving, and breakaway modes of the finger. Contact geometries are modeled to specification with hydroelastic meshes [20] generated from CAD drawings. The Belleville washers are modeled with a single compliant hydroelastic cylinder (pink in Fig. 23). The distal wheel (orange in Fig. 23) connects to the distal link via a pulley, modeled with a holonomic constraint. The motor is driven by a PD controller with effort limits to rotate at 200 rad/s. We set  $\delta t = 0.5$  ms to limit tooth travel to 25% of their width per time step. The friction coefficient between clutch and proximal gear is 1.0 (effectively rough). For the worm gears, the friction coefficient is estimated as  $\mu = 1.05 \cdot \tan(\alpha)$ , with  $\alpha$  the lead angle, to ensure non-backdrivability. All other surfaces are frictionless.




Fig. 23: Barrett hand finger with clutch engaged (left) and disengaged (right). Gears are color-coded in the proximal drivetrain as: proximal gear (blue), clutch (pink), proximal worm (green), proximal wheel (purple), and in the distal drivetrain as: distal gear (yellow), distal worm (red), distal wheel (orange).

Overall, the system has 7 degrees of freedom. On average, there are approximately 1100 contacts per time step.

Figure 24 shows the simulated contact torques on the proximal worm gear during operation. In the initial configuration, the proximal gear and clutch are not in contact. Thus, we drive the motor into a “Loading Phase” such that the proximal gear is driven into contact with the clutch. Joint limits and friction lock the proximal drivetrain while the proximal gear winds down to meet the clutch until the motor reaches its effort limit (0.6 N-cm). We then reverse the motor direction at  $t = 0.05$  s to enter a “Driving Phase” using a higher effort limit (0.66 N-cm) — proximal gear and clutch are engaged as the resulting contact torques from the motor do not exceed the magnitude of the loaded torque. We observe a characteristic “hammering” effect in the contact torques as loads transfer from one tooth to the next. A frequency analysis of the torques using a Fast Fourier Transform (FFT) verifies this, containing harmonics consistent with the gear ratios, tooth widths, and angular velocities. At  $t = 0.26$  s the proximal gear comes into contact with a fixed obstacle. The clutch torque builds, resisted by contact between the now fixed proximal wheel and worm gear, until it exceeds the limit reached during the “Loading Phase”. At this point, the proximal gear transitions into slip, eventually winding up its threads until it is completely out of contact with the clutch. The clutch disengages, and the net torque on the worm gear is zero; the drivetrain is completely disconnected and stiction prevents the worm gear from backdriving.

Finally, we point out that, with the Similar model, we




Fig. 24: Torques on the proximal worm gear transmitted through contact from the proximal wheel and clutch.

observe a spurious “locking” of the proximal wheel’s teeth as they glide within the threads of the worm gear. This is similar to the phenomenon reported in Fig. 2 of [21]. Therefore, these simulation results use the Lagged model, which does not introduce any of these artifacts during sliding.

#### D. Trajectory Optimization

We demonstrate end-to-end differentiation through contact using an implementation of contact-implicit iLQR with Drake [35]. To compute a local feedback controller and optimal trajectory, iLQR for a system with discrete dynamics  $x_{k+1} = f(x_k, u_k)$  must compute derivatives with respect state  $\partial f / \partial x_k$  and control inputs  $\partial f / \partial u_k$ . In this demonstration, the task is for the Kinova Gen3 robot arm to move a ball on a plane from a starting position to a target position (Fig. 25). We use the exact model and environment setup described in [35], with all geometries modeled with hydroelastic contact [20] and a time step of 10 ms over a 0.5 s time horizon. The method is able to efficiently solve for an optimal trajectory involving contact between the arm and the ball without explicitly dictating a contact sequence.




Fig. 25: Keypoints of an optimal trajectory found by iLQR for the arm moving the ball to the target position (red).

#### E. Deformable Bodies

Our method naturally extends to support frictional contact involving deformable bodies. In this experiment, we simulate deformable FinRay gripper fingers attached to a Panda arm in a peg-in-hole task (Fig. 26). Each finger is discretized as a tetrahedral mesh with 1,009 vertices and 2668 tetrahedra and simulated with the linear corotational model described in [12]. The fingers are attached to the Panda hand using holonomic constraints between the mesh vertices the rigid hand. Material properties of the fingers include a Young’s modulus of  $E = 2.5 \times 10^6$  Pa, Poisson’s ratio  $\nu = 0.49$ , density  $\rho = 1000$  kg/m<sup>3</sup>, and Rayleigh stiffness damping coefficient  $\zeta = 0.01$  s.

The robot is teleoperated using a 6-DoF space mouse to control the end-effector pose, while the arm’s joint positions are solved through differential inverse kinematics. The simulation employs the Lagged approximation with a 10 ms time step. The robot is commanded to place a rigid cylinder into a rigid utensil holder welded to the table. The average simulation time per time step is 24.5 ms, fast enough for interactive teleoperation, and the maximum and average number of contact constraints are 114 and 73.9, respectively.




Fig. 26: Simulation of deformable Finray grippers in a teleoperation task. (Left) the mesh used to model the FinRay gripper. (Center) the peg-in-hole teleoperation task. (Right) the characteristic caging deformation induced by frictional contact with the manipulant.

### XIII. CONCLUSIONS

We presented ICF, a novel theory for the convex approximation of contact. Our mathematical framework establishes a family of convex approximations of frictional contact, and we show that previous approaches [2], [3], [4] are members of this family. This framework enables us to incorporate complex physics-based models of contact, such as the Hunt & Crossley [8] model, within a convex formulation for the first time. These models, grounded in physics and experimentally validated, have the potential to narrow the *sim2real* gap. Within this framework, we develop two convex approximations of regularized friction: *Similar* and *Lagged*.

This work presents a thorough characterization of these approximations in terms of consistency, the coupling between normal and tangential components, and artifacts introduced by the convex approximation. While previous work has documented *gliding* during slip over a distance of  $\delta t \mu \|v_t\|$ ,

we identify previously unrecognized artifact characteristics of compliant contact [3], [4]. We validate these findings with a rich set of test cases designed to expose these problems and gain insight into the new formulations. Moreover, our analysis led to new understandings of the coupling between normal and frictional components of the contact, allowing us to design a *regularized* scheme that improves numerical conditioning during difficult-to-resolve impact events.

Our investigation concludes that, even though normal and friction forces are weakly coupled, the Lagged approximation is well-suited for the modeling of most robotic tasks. Moreover, the Lagged approximation completely eliminates artifacts associated with previous convex approximations.

Our work is implemented in the open-source robotics toolkit Drake [9]. We rigorously tested our implementation on various robotics-relevant problems, including an iLQR application to highlight our differentiable pipeline and a deformable body simulation to demonstrate compatibility with FEM-based methods.

One of the most significant limitations of our simulation pipeline is related to *tunneling* or *passthrough* problems, where objects can bypass each other without registering contact due to the nature of discrete contact detection. This issue is particularly pronounced with large time step sizes and thin objects. We are currently investigating a solution based on *speculative constraints* [36] for hydroelastic contact [20], [11] that we believe can help mitigate this issue.

#### APPENDIX A SOFT NORM

We define the *soft* norm of a vector  $\mathbf{x}$  as

$$\|\mathbf{x}\|_s = \sqrt{\|\mathbf{x}\|^2 + \varepsilon_x^2} - \varepsilon_x,$$

where  $\varepsilon_x > 0$  has units of  $\mathbf{x}$ . Notice that  $\|\mathbf{0}\|_s = 0$ .

The gradient of the soft norm is

$$\frac{\partial \|\mathbf{x}\|_s}{\partial \mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|_s + \varepsilon_x} = \hat{\mathbf{x}}_s,$$

where we defined the *soft unit vector*  $\hat{\mathbf{x}}_s$ .

The Hessian of the soft norm, the gradient of  $\hat{\mathbf{x}}_s$ , is

$$\begin{aligned} \frac{\partial \hat{\mathbf{x}}_s}{\partial \mathbf{x}} &= \frac{\partial^2 \|\mathbf{x}\|_s}{\partial \mathbf{x}^2} = \frac{1}{\|\mathbf{x}\|_s + \varepsilon_x} (\mathbf{I} - \mathbf{P}(\hat{\mathbf{x}}_s)) \\ &= \frac{\mathbf{P}^\perp(\hat{\mathbf{x}}_s)}{\|\mathbf{x}\|_s + \varepsilon_x} \end{aligned}$$

where the projection matrix is defined as  $\mathbf{P}(\hat{\mathbf{v}}) = \hat{\mathbf{v}} \otimes \hat{\mathbf{v}}$ . Note that  $\mathbf{P}(\hat{\mathbf{v}}) \succeq 0$  and  $\mathbf{P}^\perp(\hat{\mathbf{v}}) \succeq 0$  for all unit vectors  $\hat{\mathbf{v}} \in \mathbb{R}^n$ . Therefore, the soft norm is twice differentiable with positive semi-definite Hessian, and thus it is convex.

These expressions for the gradient and Hessian of the norm of a vector are still valid in the limit  $\varepsilon_x \rightarrow 0$ , but they are not well-defined at  $\mathbf{x} = \mathbf{0}$ . However, the *soft* versions have the nice property that they are numerically well-behaved near and at  $\mathbf{x} = \mathbf{0}$ , and they are continuously differentiable.

#### APPENDIX B ANISOTROPIC COULOMB FRICTION

When friction is anisotropic, the friction force can have a component perpendicular to the slip velocity and, even in the absence of external forces, objects follow curved paths [37]. The ellipsoidal friction cone  $\mathcal{F} = \{\mathbf{x}_t, x_n\} \in \mathbb{R}^3 \mid \|\boldsymbol{\mu}^{-1} \mathbf{x}_t\| \leq x_n\}$ , with  $\boldsymbol{\mu} \succeq 0$  the friction tensor, is a popular approximation.  $\boldsymbol{\mu}$  is diagonal when expressed in a frame aligned with its principal axes, and  $\boldsymbol{\mu} = \mu \mathbf{I}$  for isotropic friction.

We write an anisotropic model of Coulomb friction that satisfies the principle of maximum dissipation

$$\gamma_t = \arg \max_{\boldsymbol{\xi} \in \mathcal{F}} -\mathbf{v}_t \cdot \boldsymbol{\xi}. \quad (30)$$

We can solve this problem analytically by a change of variables that maps the ellipsoidal section of the friction cone into a circular section. The result is

$$\gamma_t = -\boldsymbol{\mu} \gamma_n \hat{\mathbf{t}}(\boldsymbol{\mu} \mathbf{v}_t), \quad (31)$$

where we define  $\hat{\mathbf{t}}(\boldsymbol{\mu} \mathbf{v}_t) = \boldsymbol{\mu} \mathbf{v}_t / \|\boldsymbol{\mu} \mathbf{v}_t\|$ . This friction model opposes slip when tensor  $\boldsymbol{\mu}$  is isotropic, while it introduces a component perpendicular to the line of motion, as experimentally confirmed in [37].

As with (15), we write a generic form of this model as

$$\gamma_t = g(\|\tilde{\mathbf{v}}_t\|, v_n) \tilde{\mathbf{t}}, \quad (32)$$

where we defined the *tilde* quantities as  $\tilde{\mathbf{v}}_t = \boldsymbol{\mu} \mathbf{v}_t$  and  $\tilde{\mathbf{t}} = \boldsymbol{\mu} \hat{\mathbf{t}}(\tilde{\mathbf{v}}_t)$ , consistent with (31). We verify that

$$\frac{\partial \gamma_t}{\partial \mathbf{v}_t} = \frac{\partial g}{\partial \|\tilde{\mathbf{v}}_t\|} \tilde{\mathbf{P}} + g \frac{\tilde{\mathbf{P}}^\perp}{\|\tilde{\mathbf{v}}_t\|},$$

with  $\tilde{\mathbf{P}} = \mathbf{P}(\tilde{\mathbf{t}})$  and  $\tilde{\mathbf{P}}^\perp = \boldsymbol{\mu} \mathbf{P}^\perp(\hat{\mathbf{t}}(\tilde{\mathbf{v}}_t)) \boldsymbol{\mu} = \boldsymbol{\mu}^2 - \tilde{\mathbf{P}}$ . Therefore,  $\partial \gamma_t / \partial \mathbf{v}_t$  for (32) is symmetric and condition (14) is met.

Using these *tilde* variables, we define a potential for Lagged

$$\ell_t(\mathbf{v}_t) = \gamma_{n0} \varepsilon_s F(\|\tilde{\mathbf{v}}_t\| / \varepsilon_s),$$

and a potential for Similar is obtained by updating (19) to  $z = v_n - \varepsilon_s F(\|\tilde{\mathbf{v}}_t\| / \varepsilon_s)$ .

#### ACKNOWLEDGMENT

The authors would like to thank the Dynamics & Simulation and Large Behavior Model teams at TRI for their continuous patience and support. We thank William Townsend and Barrett for providing us with technical specifications for the BarrettHand as well as advice and discussions on our modeling approach. We thank Cody Simpson for modeling the deformable FinRay gripper.


#### REFERENCES

- [1] D. M. Kaufman, S. Sueda, D. L. James, and D. K. Pai, "Staggered projections for frictional contact in multibody systems," *ACM Trans. Graph.*, vol. 27, no. 5, Dec. 2008.
- [2] M. Anitescu, "Optimization-based simulation of nonsmooth rigid multibody dynamics," *Mathematical Programming*, vol. 105, no. 1, pp. 113–143, 2006.
- [3] E. Todorov, "A convex, smooth and invertible contact model for trajectory optimization," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1071–1076.

- [4] A. M. Castro, F. N. Permenter, and X. Han, "An unconstrained convex formulation of compliant contact," *IEEE Transactions on Robotics*, 2022.
- [5] H. Mazhar, D. Melanz, M. Ferris, and D. Negrut, "An analysis of several methods for handling hard-sphere frictional contact in rigid multibody dynamics," Citeseer, Tech. Rep., 2014.
- [6] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
- [7] Q. Le Lidec, W. Jallet, L. Montaut, I. Laptev, C. Schmid, and J. Carpentier, "Contact models in robotics: a comparative analysis," *IEEE Transactions on Robotics*, 2024.
- [8] K. H. Hunt and F. R. E. Crossley, "Coefficient of Restitution Interpreted as Damping in Vibroimpact," *Journal of Applied Mechanics*, vol. 42, no. 2, pp. 440–445, 06 1975. [Online]. Available: <https://doi.org/10.1115/1.3423596>
- [9] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," <https://drake.mit.edu>, 2019.
- [10] W. Townsend, "The barretthand grasper—programmably flexible part handling and assembly," *Industrial Robot: an international journal*, vol. 27, no. 3, pp. 181–188, 2000.
- [11] J. Masterjohn, D. Guoy, J. Shepherd, and A. Castro, "Velocity level approximation of pressure field contact patches," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 593–11 600, 2022.
- [12] X. Han, J. Masterjohn, and A. Castro, "A convex formulation of frictional contact between rigid and deformable bodies," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6219–6226, 2023.
- [13] Z. Zong, C. Jiang, and X. Han, "A convex formulation of frictional contact for the material point method and rigid bodies," *arXiv preprint arXiv:2403.13783*, 2024.
- [14] D. E. Stewart and J. C. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction," *International Journal for Numerical Methods in Engineering*, vol. 39, no. 15, pp. 2673–2691, 1996.
- [15] M. Anitescu and F. A. Potra, "Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems," *Nonlinear Dynamics*, vol. 14, no. 3, pp. 231–247, 1997.
- [16] A. Pandolfi, C. Kane, J. E. Marsden, and M. Ortiz, "Time-discretized variational formulation of non-smooth frictional contact," *International Journal for Numerical Methods in Engineering*, vol. 53, no. 8, pp. 1801–1829, 2002.
- [17] K. L. Johnson and K. L. Johnson, *Contact mechanics*. Cambridge university press, 1987.
- [18] P. Flores, "Contact mechanics for dynamical systems: a comprehensive review," *Multibody System Dynamics*, pp. 1–51, 2021.
- [19] H. Bhatia, G. Norgard, V. Pascucci, and P.-T. Bremer, "The helmholtz-hodge decomposition—a survey," *IEEE Transactions on visualization and computer graphics*, vol. 19, no. 8, pp. 1386–1404, 2012.
- [20] R. Elandt, E. Drumwright, M. Sherman, and A. Ruina, "A pressure field model for fast, robust approximation of net contact force and moment between nominally rigid objects," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 8238–8245.
- [21] P. C. Horak and J. C. Trinkle, "On the similarities and differences among contact models in robot simulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 493–499, 2019.
- [22] T. A. Howell, S. Le Cleac'h, J. Z. Kolter, M. Schwager, and Z. Manchester, "Dojo: A differentiable simulator for robotics," *arXiv preprint arXiv:2203.00806*, vol. 9, no. 2, p. 4, 2022.
- [23] 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 Trans. Graph.*, vol. 39, no. 4, p. 49, 2020.
- [24] M. Macklin, K. Erleben, M. Müller, N. Chentanez, S. Jeschke, and V. Makovychuk, "Non-smooth newton methods for deformable multi-body dynamics," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 5, pp. 1–20, 2019.
- [25] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, pp. 25–57, 2006.
- [26] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2024. [Online]. Available: <https://www.gurobi.com>
- [27] M. ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 10.1.*, 2024. [Online]. Available: <http://docs.mosek.com/latest/toolbox/index.html>
- [28] W. H. Press, H. William, S. A. Teukolsky, A. Saul, W. T. Vetterling, and B. P. Flannery, *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- [29] T. A. Davis, S. Rajamanickam, and W. M. Sid-Lakhdar, "A survey of direct methods for sparse linear systems," *Acta Numerica*, vol. 25, pp. 383–566, 2016.
- [30] P. R. Amestoy, T. A. Davis, and I. S. Duff, "An approximate minimum degree ordering algorithm," *SIAM Journal on Matrix Analysis and Applications*, vol. 17, no. 4, pp. 886–905, 1996.
- [31] J. Pan, S. Chitta, and D. Manocha, "FCL: A general purpose library for collision and proximity queries," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3859–3866.
- [32] M. Geilinger, D. Hahn, J. Zehnder, M. Bächer, B. Thomaszewski, and S. Coros, "ADD: Analytically differentiable dynamics for multi-body systems with frictional contact," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–15, 2020.
- [33] Q. L. Lidec, L. Montaut, Y. de Mont-Marin, and J. Carpentier, "End-to-end and highly-efficient differentiable simulation for robotics," *arXiv preprint arXiv:2409.07107*, 2024.
- [34] F. Pfeiffer and C. Glocker, *Multibody Dynamics with Unilateral Contacts*, ser. Wiley Series in Nonlinear Science. Wiley, 1996.
- [35] V. Kurtz and H. Lin, "Contact-implicit trajectory optimization with hydroelastic contact and iLQR," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 8829–8834.
- [36] E. Catto, "Continuous collision," 2013, game Developers Conference [https://box2d.org/files/ErinCatto\\_ContinuousCollision\\_GDC2013.pdf](https://box2d.org/files/ErinCatto_ContinuousCollision_GDC2013.pdf).
- [37] S. Walker and R. Leine, "Set-valued anisotropic dry friction laws: formulation, experimental verification and instability phenomenon," *Nonlinear Dynamics*, vol. 96, pp. 885–920, 2019.

**IEEE Transactions on Robotics (T-RO) paper, presented at ICRA 2026, Vienna, Austria. Cite as T-RO paper.**

**Alejandro M. Castro** is Senior Manager at the Toyota Research Institute (TRI), Cambridge, MA, USA. He received the B.S.E. degree in Nuclear Engineering in 2006 and the M.Sc. in Physics in 2007 from Balseiro Institute, Bariloche, Río Negro, Argentina and the Ph.D. degree in Mechanical Engineering from The University of Iowa, Iowa City, IA, USA in 2011. His current research focus is on numerical methods for simulation, planning and control of rigid and deformable mechanical systems with frictional contact for robotics applications.



**Xuchen Han** received the B.A. degree in Mathematics and Economics from Northwestern University, Evanston, IL, USA, in 2015 and the Ph.D. degree in Mathematics from University of California, Los Angeles (UCLA), Los Angeles, CA, USA in 2020. He is now a research scientist at Toyota Research Institute (TRI), Los Altos, CA, USA.



**Joseph Masterjohn** is a senior research scientist on the Dynamics & Simulation team at Toyota Research Institute (TRI) Cambridge, MA. He holds a B.S (2014), M.S. (2017) and Ph.D. (2020) in Computer Science from the University of Miami, Coral Gables, FL, USA.