

A Transition-Aware Method for the Simulation of Compliant Contact with Regularized Friction

Alejandro M. Castro^{1*}, Ante Qu^{1,2*}, Naveen Kuppuswamy¹, Alex Alspach¹, and Michael Sherman¹

Abstract—Multibody simulation with frictional contact has been a challenging subject of research for the past thirty years. Rigid-body assumptions are commonly used to approximate the physics of contact, and together with Coulomb friction, lead to challenging-to-solve nonlinear complementarity problems (NCP). On the other hand, robot grippers often introduce significant compliance. Compliant contact, combined with regularized friction, can be modeled entirely with ODEs, avoiding NCP solves. Unfortunately, regularized friction introduces high-frequency stiff dynamics and even implicit methods struggle with these systems, especially during slip-stick transitions. To improve the performance of implicit integration for these systems we introduce a Transition-Aware Line Search (TALS), which greatly improves the convergence of the Newton-Raphson iterations performed by implicit integrators. We find that TALS works best with semi-implicit integration, but that the explicit treatment of normal compliance can be problematic. To address this, we develop a Transition-Aware Modified Semi-Implicit (TAMSI) integrator that has similar computational cost to semi-implicit methods but implicitly couples compliant contact forces, leading to a more robust method. We evaluate the robustness, accuracy and performance of TAMSI and demonstrate our approach alongside relevant sim-to-real manipulation tasks.

Index Terms—contact modeling, simulation and animation, grasping, robotics manipulation, dynamics.

I. INTRODUCTION

RECENTLY robotics has experienced a dramatic boom due to the introduction of new sensor technologies, actuation, and innovative software algorithms that allow robots to reason about the world around them. These new technologies are allowing the next generation of robots to start moving from their highly structured environments in factories and research labs to less structured, richer environments such as those found in the home. There are still, however, a large variety of research problems to be solved. In particular, manipulation is one area of robotics that raises significant challenges including high-speed sensing, planning, and control.

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. Manuscript received: September 10, 2019; Revised December 16, 2019; Accepted January 13, 2020.

This paper was recommended for publication by Editor Hong Liu upon evaluation of the Associate Editor and Reviewers' comments.

*Authors have contributed equally, and names are in alphabetical order.

¹Toyota Research Institute, One Kendall Square, Binney St Building 100, Suite 1-201, Cambridge, MA 02139, USA.

²Stanford University, Computer Science, 353 Jane Stanford Way, Stanford, CA 94305, USA.

[alejandro.castro, naveen.kuppuswamy, alex.alspach, sherm]@tri.global, antequ@cs.stanford.edu

Digital Object Identifier (DOI): see top of this page.

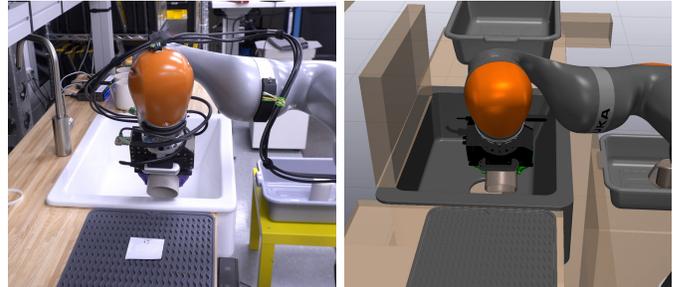


Fig. 1: Sim-to-real comparison for the task of reorienting a mug by making contact with an external surface. See supplemental video.

Simulation has proven indispensable in this new era of robotics, aiding at multiple stages during the mechanical and control design, testing, and training of robotic systems. For instance, [1] demonstrates the use of simulated data to train a robotic system to grasp new objects, while [2] studies how to transfer policies trained in simulation to the real world.

Near real-time forward simulation has found applications in contact-aware state estimation wherein predicted contacts and physical consistency [3] are used to estimate the state of a robot, manipulands, or both [4], [5].

We would like to synthesize, train, and validate controllers in simulation with the expectation that they will work well in reality. Hence simulation should present a controller with a range of *physical* models and disturbances, but must avoid significant *non-physical* simulation artifacts that have no real-world counterparts. A central challenge in simulation for manipulation is the physical modeling and numerical solution of multibody dynamics with contact and friction. Such simulations often involve high mass ratios, stiff dynamics, complex geometries, and friction. Artifacts that are unimportant for other applications are highly amplified in the simulation of a manipulation task; simulations either become unstable or predict highly unstable grasps even if stable in the real system. These characteristics impose strict requirements on robustness, accuracy, and performance to the simulation engine when compared to other robotic scenarios with contact, such as walking.

A popular approximation to the true compliant physics of contact is the mathematical limit in which bodies are rigid; however, it can lead to indeterminate systems with multiple solutions, or no solution [6]. Still, the rigid-body approximation is at the core of many simulation engines, enabling them to run at interactive rates.

Generally, rigid-body assumptions and Coulomb friction lead to a complex formulation in terms of a nonlinear complementarity problem (NCP), which can be simplified to a

linear complementarity problem (LCP) using a polygonization of the friction cone [7], [8]. LCPs are often solved using direct pivoting methods such as Lemke’s algorithm. Even though there are theoretical results for the solution of LCPs [9], the class of systems for which they apply are seldom found in practice, and robust LCP implementations are either slow or proprietary. The authors of [10] iteratively solve a quadratic program (QP) for the friction impulses and a second QP for the normal impulses. Although the results are promising, the coupled problem is non-convex and the solution is not guaranteed to exist.

Another drawback of LCP formulations is that the linearization of the friction cone might lead to preferential directions and cause bias in the solution [11]—an example of a potentially-significant non-physical simulation artifact. The computational gain from approximating the friction cone is not clear given that the LCP introduces many auxiliary variables and constraints, and thus [12], [13] propose to solve the original NCP with a non-smooth Newton method. It is common practice to relax contact constraints by introducing *regularization* [14] or *softening* [15], [16], [17], allowing objects to interpenetrate due to *numerical compliance*.

Given that the true physics of contact involves compliance and many of the approaches using rigid-body assumptions introduce numerical compliance to make the problem tractable, our work favors using compliant contact. Models in the literature include those for point contact based on Hertz theory [18], volumetric models [19] and more sophisticated approaches modeling contact patches [20], [21].

Coulomb friction can also be regularized by replacing the strict friction-cone constraint with a smooth function of the slip velocity and introducing a *regularization parameter*, or *stiction velocity*, v_s so that objects supposedly in stiction still slide with a velocity smaller than v_s ; for example see [22].

By incorporating compliant contact forces and regularized friction, we can write the system dynamics as a system of continuous ordinary differential equations (ODEs). However, we observe that in our simulations of manipulation tasks the stiffness of the model is dominated by the regularization of friction. In static equilibrium, the forces due to friction balance the pull of gravity which would otherwise accelerate the object downwards at $g \approx 9.8 \text{ m/s}^2$. Using a typical v_s value of 10^{-4} m/s , the characteristic time scale introduced by regularized friction is about $\tau \approx v_s/g \approx 10^{-5} \text{ s}$, and thus our error-controlled integrators must take steps as small as $0.1 \mu\text{s}$ to resolve these highly stiff dynamics. As a result, integrators spend most of the computational effort on resolving these artificially introduced dynamics even when common grasping tasks involve much larger time scales on the order of tenths of a second. With implicit integration, stability theory says we should be able to take large time steps once the system is in stiction, even with a very small v_s . This has proven difficult in practice, however. In this work, we analyze the cause and present methods that enable the realization of this theoretical promise in practice.

We organize our work as follows. Section II introduces the mathematical framework and notation. Section III introduces our novel Transition-Aware Line Search (TALS) in the

context of implicit integration. We systematically assess the performance of a variety of integrators using work-precision plots on a series of proposed canonical test problems in Section IV and measure the improvement in robustness and performance introduced by TALS. We show TALS performs best when *freezing* the configuration of the system and propose an original Transition-Aware Modified Semi Implicit (TAMSI) method in Section V. In section VI we show that TAMSI handles transitions robustly and outperforms our best implicit integrators for the simulation of relevant manipulation tasks. In VI-C we demonstrate our method in two sim-to-real comparisons using a Kuka arm manipulation station. Final remarks and future research directions are presented in Section VII.

II. MULTIBODY DYNAMICS WITH CONTACT

We start by stating the equations of motion and introducing our notation,

$$\dot{\mathbf{q}} = \mathbf{N}(\mathbf{q})\mathbf{v}, \quad (1)$$

$$\mathbf{M}(\mathbf{q})\dot{\mathbf{v}} = \boldsymbol{\tau}(\mathbf{q}, \mathbf{v}) + \mathbf{J}_c^T(\mathbf{q})\mathbf{f}_c(\mathbf{q}, \mathbf{v}), \quad (2)$$

where $\mathbf{q} \in \mathbb{R}^{n_q}$ and $\mathbf{v} \in \mathbb{R}^{n_v}$ are the vectors of generalized positions and velocities respectively, $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{n_v \times n_v}$ is the system’s mass matrix, $\boldsymbol{\tau}(\mathbf{q}, \mathbf{v})$ is a vector of generalized forces containing Coriolis and gyroscopic terms, gravity, externally applied forces, and actuation, and \mathbf{J}_c and \mathbf{f}_c are contact Jacobians and forces, defined in Section II-A. We explicitly emphasize the functional dependence of $\mathbf{f}_c(\mathbf{q}, \mathbf{v})$ on the state vector $\mathbf{x} = (\mathbf{q}, \mathbf{v})$ given that we use a compliant contact model with regularized friction. Finally, $\mathbf{N}(\mathbf{q}) \in \mathbb{R}^{n_q \times n_v}$ in Eq. (2) is the block diagonal mapping between generalized velocities \mathbf{v} and time derivative of generalized positions $\dot{\mathbf{q}}$. Together, Eqs. (1) and (2) describe the system’s dynamics as

$$\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}). \quad (3)$$

A. Kinematics of Contact

In point contact, contact between sufficiently smooth non-conforming surfaces can be simply characterized by a pair of witness points A_w and B_w on bodies A and B , respectively such that A_w is a point on the surface of A that lies the farthest from the surface of B and vice versa. At the i th contact we define the contact point C_i to be midway between the witness points. For a given configuration \mathbf{q} of the system each contact point is characterized by a penetration distance $\delta_i(\mathbf{q})$ and a contact normal $\hat{\mathbf{n}}_i(\mathbf{q})$ defined to point from body B into body A . We denote with $\mathbf{v}_{c,i}$ the velocity of witness point A_w relative to B_w , which can be uniquely split into normal velocity $\mathbf{v}_{n,i} = \mathbf{P}_{n,i}\mathbf{v}_{c,i}$ and tangential velocity $\mathbf{v}_{t,i} = \mathbf{P}_{n,i}^\perp\mathbf{v}_{c,i}$ where $\mathbf{P}_{n,i} = \hat{\mathbf{n}}_i \otimes \hat{\mathbf{n}}_i \succeq 0$ and $\mathbf{P}_{n,i}^\perp = \mathbf{I} - \mathbf{P}_{n,i} \succeq 0$ are projection matrices in $\mathbb{R}^{3 \times 3}$.

For a state with n_c contact points we group the velocities $\mathbf{v}_{c,i}$ of all contact points in a vector $\mathbf{v}_c \in \mathbb{R}^{3n_c}$. In Eq. (2) the contact Jacobian $\mathbf{J}_c(\mathbf{q}) \in \mathbb{R}^{3n_c \times n_v}$ maps generalized velocities to contact point velocities as $\mathbf{v}_c = \mathbf{J}_c\mathbf{v}$. We group the *scalar* separation velocities $v_{n,i} = \hat{\mathbf{n}}_i \cdot \mathbf{v}_{c,i}$ into a vector $\mathbf{v}_n \in \mathbb{R}^{n_c}$ so that $\mathbf{v}_n = \mathbf{J}_n\mathbf{v}$, where $\mathbf{J}_n = \hat{\mathbf{N}}^T\mathbf{J}_c \in \mathbb{R}^{n_c \times n_v}$

is the normal velocities Jacobian and $\hat{\mathbf{N}} = \text{diag}(\{\hat{\mathbf{n}}_i\}) \in \mathbb{R}^{3n_c \times n_c}$. For the tangential velocities we write

$$\mathbf{v}_t(\mathbf{q}, \mathbf{v}) = \mathbf{J}_t(\mathbf{q})\mathbf{v}, \quad (4)$$

with $\mathbf{J}_t(\mathbf{q}) = \mathbf{P}_n^\perp(\mathbf{q})\mathbf{J}_c(\mathbf{q}) \in \mathbb{R}^{3n_c \times n_v}$ the Jacobian of tangential velocities and $\mathbf{P}_n^\perp(\mathbf{q}) = \text{diag}(\{\mathbf{P}_{n,i}^\perp\}) \in \mathbb{R}^{3n_c \times 3n_c}$. The reader should notice the difference in sizes for $\mathbf{v}_n \in \mathbb{R}^{n_c}$ grouping the scalar separation velocities, and for $\mathbf{v}_t \in \mathbb{R}^{3n_c}$ grouping 3D tangential velocities. This choice simplifies the exposition that follows.

B. Compliant Contact with Regularized Friction

The normal component of the contact force is modeled following the functional form proposed in [23], which is continuous in both penetration distance δ_i and rate $\dot{\delta}_i$ as

$$\pi_i = k_i(1 + d_i \dot{\delta}_i)_+ \delta_{i+}, \quad (5)$$

where $(\cdot)_+ = \max(\cdot, 0)$ and k_i and d_i are stiffness and damping parameters. Thus, the normal component of the contact force on body A applied at C_i is $\mathbf{f}_{n,i} = \pi_i \hat{\mathbf{n}}_i$. These parameters can be treated as either physical parameters computed for instance according to the theory of Hertz contact as in [20] or as numerical penalty parameters as in [15], [24]. Since the penetration distance δ_i is defined positive for overlapping geometries, its time derivative relates to the separation velocity $v_{n,i}$ by $\dot{\delta}_i = -v_{n,i}$.

We approximate the Coulomb friction force on body A at C_i with a linear functional form of the tangential velocity $\mathbf{v}_{t,i}$, though smoother functions can be used

$$\mathbf{f}_{t,i} = -\tilde{\mu}_i(\|\mathbf{v}_{t,i}\|/v_s) \pi_i \hat{\mathbf{v}}_{t,i}, \quad (6)$$

$$\tilde{\mu}_i(s) = \begin{cases} \mu_i s, & 0 \leq s \leq 1, \\ \mu_i, & 1 < s, \end{cases} \quad (7)$$

where $\tilde{\mu}(s) \geq 0$ is the regularized friction coefficient with μ_i the coefficient of friction and v_s , with units of velocity, is the regularization parameter. We show in Section V-A that regularized friction with positive slopes, i.e. $\tilde{\mu}'(s) \geq 0$ leads to considerably more stable integration schemes.

As with velocities, we group contact forces $\mathbf{f}_{c,i}$ into a single vector $\mathbf{f}_c(\mathbf{q}, \mathbf{v}) \in \mathbb{R}^{3n_c}$. We split the contact forces in their normal and tangential components as $\mathbf{f}_c = \mathbf{f}_n + \mathbf{f}_t$ and define the generalized forces due to contact as $\boldsymbol{\tau}_n(\mathbf{q}, \mathbf{v}) = \mathbf{J}_c^T(\mathbf{q})\mathbf{f}_n(\mathbf{q}, \mathbf{v})$ and $\boldsymbol{\tau}_t(\mathbf{q}, \mathbf{v}) = \mathbf{J}_c^T(\mathbf{q})\mathbf{f}_t(\mathbf{q}, \mathbf{v})$. We note that, since $\mathbf{f}_t = \mathbf{P}_n^\perp \mathbf{f}_c$ we can write

$$\begin{aligned} \boldsymbol{\tau}_t(\mathbf{q}, \mathbf{v}) &= \mathbf{J}_c^T(\mathbf{q}) \mathbf{f}_t(\mathbf{q}, \mathbf{v}), \\ &= \mathbf{J}_c^T(\mathbf{q}) (\mathbf{P}_n^\perp(\mathbf{q}) \mathbf{f}_c(\mathbf{q}, \mathbf{v})), \\ &= (\mathbf{P}_n^\perp(\mathbf{q}) \mathbf{J}_c(\mathbf{q}))^T \mathbf{f}_c(\mathbf{q}, \mathbf{v}), \\ &= \mathbf{J}_t^T(\mathbf{q}) \mathbf{f}_c(\mathbf{q}, \mathbf{v}). \end{aligned} \quad (8)$$

III. IMPLICIT INTEGRATION WITH TALS

We introduce our Transition-Aware Line Search (TALS) method to improve the convergence and robustness of implicit integration methods when using large time steps. We first make a brief overview of the implicit Euler method, though

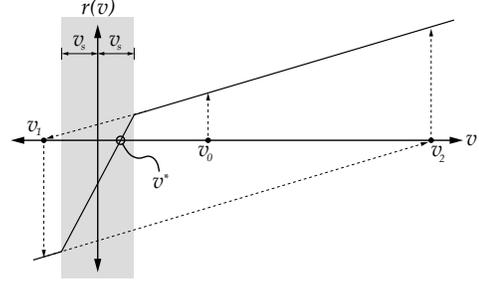


Fig. 2: Divergence of Newton-Raphson near stiction. Iterations cycle between v_2 and v_1 indefinitely.

TALS can be used with other implicit integrators. Consider a discrete step of size h from time t^{n-1} to time $t^n = t^{n-1} + h$. Implicit Euler approximates the time derivative in Eq. (3) using a first order backward differentiation formula. The resulting system of equations is nonlinear in \mathbf{x}^n and can be solved using Newton's method,

$$\begin{aligned} \Delta \mathbf{x}^k &= -(\mathbf{A}^k)^{-1} \mathbf{r}(\mathbf{x}^{n,k}); \\ \mathbf{x}^{n,k+1} &= \mathbf{x}^{n,k} + \Delta \mathbf{x}^k, \end{aligned} \quad (9)$$

where k denotes the iteration number, the residual is defined as $\mathbf{r}(\mathbf{x}) = \mathbf{x} - \mathbf{x}^{n-1} - h \mathbf{f}(\mathbf{x})$, and $\mathbf{A}^k = \nabla_{\mathbf{x}} \mathbf{r}(\mathbf{x}^k)$ is the Jacobian of the residual. We use the integrators implemented in Drake [24] which use the stopping criterion outlined in [25, §IV.8] to assess the convergence of Newton iterations.

A. Transition-Aware Line Search

We observed the Newton iteration in Eq. (9) most often fails during slip-stick transitions when using large time steps given that \mathbf{A}^k in the sliding region, $\|\mathbf{v}_{t,i}\| > v_s$, is not a good approximation of \mathbf{A}^k in the stiction region, $\|\mathbf{v}_{t,i}\| < v_s$. This problem is illustrated in Fig. 2 for a 1D system, showing the Newton residual as a function of sliding velocity. Notice the sharp gradient in the region $\|\mathbf{v}_{t,i}\| < v_s$ due to the regularization of friction. An iterate with positive velocity v_0 will follow the slope to point v_1 on the negative side. Given the transition into the stiction region is missed, next iteration then follows the slope to v_2 once again on the positive side. Subsequent iterations continue to switch back and forth between positive and negative velocities without achieving convergence. We found this can be remediated by *limiting* an iteration crossing the stiction region to fall inside of it. As soon as \mathbf{A}^k is updated with the more accurate, and larger, value within the stiction region, Newton iterations proceed without difficulties.

Similar to other system-specific line search approaches [12], [26], for three dimensional problems TALS limits the iteration update in Eq. (9) according to $\mathbf{x}^{n,k+1} = \mathbf{x}^{n,k} + \alpha \Delta \mathbf{x}^k$ where $\alpha \in [0, 1]$. If we *freeze* the configuration \mathbf{q} in Eq. (4), this is equivalent to limiting the tangential velocity of each contact point according to $\mathbf{v}_{t,i}^{n,k+1} = \mathbf{v}_{t,i}^{n,k} + \alpha \Delta \mathbf{v}_{t,i}^k$, with $\Delta \mathbf{v}_{t,i}^k$ the change predicted by Eq. (9), see Fig. 3. We monitor transition to stiction by detecting the moment at which the line connecting two iterations intersects the stiction region,

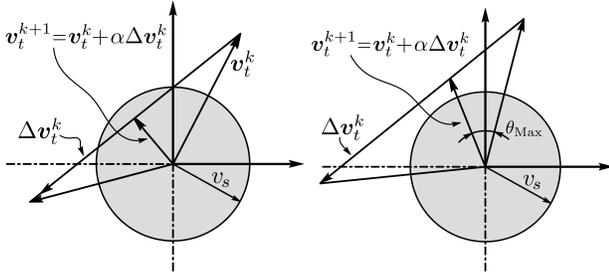


Fig. 3: TALS limits the update of tangential velocities for the next Newton iteration when it detects transitions through the stiction region (left). We found that limiting updates to a maximum angle θ_{Max} improves robustness (right).

inspired from the idea used in [27] to compute incremental impulses for impact problems.

TALS proceeds as follows: given a pair of iterates $\mathbf{x}^{n,k}$ and $\mathbf{x}^{n,k+1}$ from Eq. (9), we update tangential velocities according to Eq. (4). At the i th contact point, if the line connecting $\mathbf{v}_{t,i}^{n,k}$ and $\mathbf{v}_{t,i}^{n,k+1}$ crosses the stiction region, we compute α_i so that $\mathbf{v}_{t,i}^{n,k+1} \cdot \Delta \mathbf{v}_{t,i}^k = 0$. Even if transition is not detected, we also limit large angular changes between $\mathbf{v}_{t,i}^{n,k}$ and $\mathbf{v}_{t,i}^{n,k+1}$ to a maximum value θ_{Max} , see Fig. 3. In practice we found the angular limit to increase robustness and use $\theta_{\text{Max}} = \pi/3$. We finally compute the global TALS limiting parameter as $\alpha = \min(\{\alpha_i\})$.

B. Implementation Details

We incorporate TALS in the implicit integrators implemented in Drake [24]. Drake offers *error-controlled* integration to a user-specified accuracy a , which translates roughly to the desired number of significant digits in the results. We can also control whether to use a full- or quasi-Newton method with Jacobian update strategies as outlined in [25, §IV.8].

We make the distinction between *error-controlled* and *convergence-controlled* integration, which retries smaller time steps when Newton iterations fail to converge, as described in [25, §IV.8]. All of our fixed-step integrators use convergence control.

IV. INTEGRATION PERFORMANCE

We evaluate TALS with fixed-step and error-controlled implicit Euler (IE) integrators. The precision obtained with fixed-step integrators is controlled via the time step size h . In contrast, error-controlled integrators adjust the time step size to meet a user-specified accuracy a , with additional overhead for local error estimation.

Our objective is to evaluate the trade-off between performance and precision for a variety of integration methods and understand if it is worth paying the additional cost of error-controlled integration. We accomplish this by creating so-called *work-precision* plots that measure cost vs. precision in the solution. We use the number of evaluations of the system's dynamics $\mathbf{f}(t, \mathbf{x})$ as a metric of *work*, including the evaluations used to approximate the Jacobian of $\mathbf{f}(t, \mathbf{x})$ through forward differencing. For *precision* we want a metric

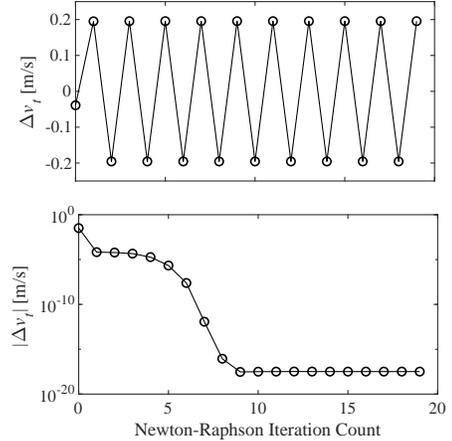


Fig. 4: Newton's method misses stiction, oscillates, and does not converge (top). TALS detects the transition to stiction, limits the iteration update, and guides Newton's method to convergence (bottom).

that measures global error while avoiding undesirable drifts observed over large periods, especially when using first order methods. Therefore we *localize* our global error metric by quantifying it within a time window Δ_w . To be more precise, we introduce the *flow map* $\mathbf{y} = \Phi_\tau(\mathbf{x}_0)$ such that $\mathbf{y} = \int_0^\tau \mathbf{f} dt$ with initial condition $\mathbf{y}(0) = \mathbf{x}_0$. We compute a solution \mathbf{x}^m with the method under test at discrete times t^m spaced by intervals of duration Δ_w . We then compute a *localized* reference solution defined by $\mathbf{x}_r^m = \Phi_{\Delta_w}(\mathbf{x}^{m-1})$, where notice we use the test solution \mathbf{x}^{m-1} as the initial condition of the reference solution \mathbf{x}_r^m for the next time window Δ_w . In practice we integrate \mathbf{x}_r^m numerically with a much higher precision than the errors in the solution we want to measure. Finally, we make an error estimate by comparing \mathbf{x} with the reference \mathbf{x}_r , see below.

A. Performance Results with a Small System

We choose a simple 2D box system that exhibits periodic stick-slip transitions. The box of mass $m = 0.33$ kg lies on top of a horizontal surface with friction $\mu = 1.0$ and is forced to move sideways by an external harmonic force $f(t)$ of amplitude 4 N and a frequency of 1 Hz. The system's dynamics for this case reduces to $m\dot{v} = f(t) - \tilde{\mu}(v)W\text{sgn}(v)$, with W the weight of the box in Earth's gravity of $g = 9.8$ m/s² and $\tilde{\mu}(v)$ as defined in Eq. (7) with $v_s = 10^{-4}$ m/s.

Using a fixed time step $h = 10$ ms, much larger than the time scale introduced by regularized friction, we observe Newton iterations not to converge as described in Section III-A. Figure 4 shows the Newton residual history the first time, at $t = 0.160$ s, the box transitions from sliding into stiction. TALS is able to detect transition, properly limit the iteration update and even recover the second order convergence rate of Newton's method.

Figure 5 shows a work-precision plot with an error estimate e_v in the horizontal velocity v defined as the L^2 -norm of the difference between the solution $\{v^m\}$ and the reference solution $\{v_r^m\}$. We use a localized global error estimate computed

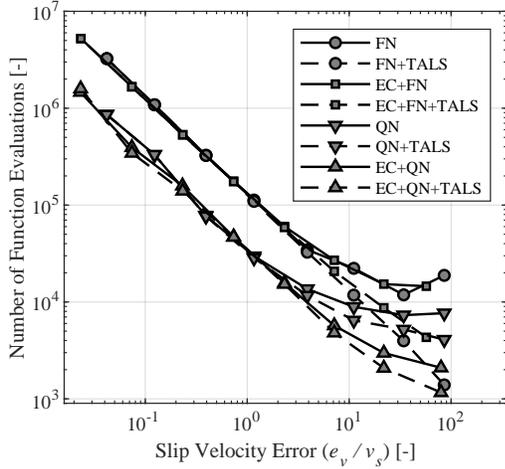


Fig. 5: We evaluate implicit Euler integration performance with full- (FN) and quasi- (QN) Newton updates, each with and without error control (EC), and with and without TALS.

with $\Delta_w = 25$ ms and a reference solution computed by a 3rd order Runge-Kutta with fixed step of $h = 10^{-7}$ s. Each point in this figure corresponds to a different accuracy a when using error control, or a different time step h when using fixed steps. The same error metric in the horizontal axis allows for a fair comparison.

We observe the expected theoretical first order slope for errors near v_s or smaller. However, for errors larger than v_s , integrators without TALS depart from the theoretical result as their performance degrades. Figure 5 shows that TALS improves performance by extending the range in which the first order behavior is valid, up to a factor of 10 (FN) or 3 (QN).

B. Performance with Larger Systems

When we applied TALS with implicit Euler to systems with more degrees of freedom, we found TALS to be highly sensitive to changes in the configuration \mathbf{q} of the system. Specifically, we observed iterations during which tangential velocities fall outside the region $\|\mathbf{v}_t\| < v_s$ even after TALS detects a transition and limits the Newton update. We traced this problem to changes in the tangential velocities Jacobian $\mathbf{J}_t(\mathbf{q})$ in Eq. (4) caused by small changes to the configuration \mathbf{q} . These iterations during transition steps often result in the same convergence failure as experienced by integrators without TALS (Figure 4), leading to a performance deterioration.

V. TRANSITION-AWARE SEMI-IMPLICIT METHOD

We draw from the lessons learned in the previous section to design a scheme customized for the solution of multibody dynamics with compliant contact and regularized Coulomb friction. Based on the observation that TALS is sensitive to changes in the system configuration \mathbf{q} , we decided to *freeze* the configurations in Eq. 2 so that we could iterate on the generalized velocities without changes in the configuration affecting the stability of TALS. An important computational

advantage of this approach is that the typically demanding geometric queries only need to be performed once at the beginning of the time step.

This is the same idea behind the semi-implicit Euler scheme. This scheme, however, becomes unstable for stiff contact forces in the normal direction since the position dependent terms are treated explicitly as in the conditionally stable explicit Euler scheme. Our TAMSI scheme deals with this problem by introducing an implicit first order approximation of the penetration distances with generalized velocities.

We introduce TAMSI in stages to analyze the properties of different contributions separately. We start with the traditional semi-implicit Euler and highlight the differences as we progressively introduce TAMSI.

A. Semi-Implicit Euler: One-Way Coupled TAMSI Scheme

The semi-implicit Euler scheme applied to Eqs. (1)-(2) effectively *freezes* the normal contact forces to $\mathbf{f}_{n,0}$. In this regard the scheme is one-way coupled, meaning that normal forces couple in the computation of friction forces but not the other way around.

Using a time step of length h , the semi-implicit Euler scheme applied to Eqs. (1)-(2) reads,

$$\mathbf{q} = \mathbf{q}_0 + \mathbf{N}_0 \mathbf{v}, \quad (10)$$

$$\mathbf{M}_0 \mathbf{v} = \mathbf{p}^* + h \mathbf{J}_{c,0}^T \mathbf{f}_c(\mathbf{q}_0, \mathbf{v}), \quad (11)$$

where the naught subscript in \mathbf{q}_0 , \mathbf{v}_0 , $\mathbf{M}_0 = \mathbf{M}(\mathbf{q}_0)$, $\mathbf{J}_{c,0}^T = \mathbf{J}_c^T(\mathbf{q}_0)$, and $\mathbf{N}_0 = \mathbf{N}(\mathbf{q}_0)$ denotes quantities evaluated at the previous time step. To simplify notation we use bare \mathbf{q} and \mathbf{v} to denote the state at the next time step. We defined $\mathbf{p}^* = \mathbf{M}_0 \mathbf{v}_0 + h \boldsymbol{\tau}(\mathbf{q}_0, \mathbf{v}_0)$ as the momentum the system would have on the next time step if contact forces were zero. Since we are interested in low-speed applications to robotic manipulation, the gyroscopic terms in \mathbf{p}^* are treated explicitly. It is known however that for highly dynamics simulations, more robust approaches should treat these terms implicitly [28].

A semi-implicit Euler scheme solves Eq. (11) for \mathbf{v} at the next time step and uses it to advance the configuration vector \mathbf{q} to the next time step with Eq. (10). However this approach often becomes unstable because the stiff compliant normal forces are explicit in the configuration \mathbf{q} .

To solve Eq. (11) we define the residual

$$\mathbf{r}(\mathbf{v}) = \mathbf{M}_0 \mathbf{v} - \mathbf{p}^* - h \boldsymbol{\tau}_{n,0} - h \boldsymbol{\tau}_t(\mathbf{q}_0, \mathbf{v}), \quad (12)$$

and use its Jacobian in Newton iterations. The Jacobian is

$$\nabla_{\mathbf{v}} \mathbf{r}(\mathbf{v}) = \mathbf{M}_0 - h \nabla_{\mathbf{v}} \boldsymbol{\tau}_t(\mathbf{q}_0, \mathbf{v}). \quad (13)$$

Substituting in Eq. (8),

$$\nabla_{\mathbf{v}} \boldsymbol{\tau}_t(\mathbf{q}_0, \mathbf{v}) = \mathbf{J}_{t,0}^T \nabla_{\mathbf{v}} \mathbf{f}_t(\mathbf{q}_0, \mathbf{v}), \quad (14)$$

$$= \mathbf{J}_{t,0}^T \nabla_{\mathbf{v}_t} \mathbf{f}_t(\mathbf{q}_0, \mathbf{v}) \mathbf{J}_{t,0}, \quad (15)$$

where we used Eq. (4) and $\nabla_{\mathbf{v}_t} \mathbf{f}_t = \text{diag}(\{\nabla_{\mathbf{v}_{t,i}} \mathbf{f}_{t,i}\}) \in \mathbb{R}^{3n_c \times 3n_c}$, with each block computed as the gradient with respect to $\mathbf{v}_{t,i}$ in Eq. (6):

$$\nabla_{\mathbf{v}_{t,i}} \mathbf{f}_{t,i} = -\pi_{i,0} \left[\frac{\tilde{\mu}(s_i)}{\|\mathbf{v}_{t,i}\|} \mathbf{P}_n^\perp(\mathbf{q}) + \frac{\tilde{\mu}'(s_i)}{v_s} \mathbf{P}_n(\mathbf{q}) \right], \quad (16)$$

where $\tilde{\mu}'(s) = d\tilde{\mu}/ds$. We introduce the tangential direction Delassus operator as $\mathbf{W}_{tt,0} = -\mathbf{J}_{t,0}^T \nabla_{\mathbf{v}_t} \mathbf{f}_t(\mathbf{q}_0, \mathbf{v}) \mathbf{J}_{t,0}$ and write the Jacobian of the residual as

$$\nabla_{\mathbf{v}} \mathbf{r}(\mathbf{v}) = \mathbf{M}_0 + h \mathbf{W}_{tt,0}. \quad (17)$$

We note that when $\tilde{\mu}'(s) \geq 0$, $-\nabla_{\mathbf{v}_{t,i}} \mathbf{f}_{t,i}$ is a linear combination of positive semi-definite matrices and thus $\mathbf{W}_{tt,0} \succeq 0$. Therefore $\nabla_{\mathbf{v}} \mathbf{r}(\mathbf{v}) \succ 0$, and its inverse always exists.

B. Implicit Approximation to Normal Forces

We present an approximation to treat normal forces implicitly while still keeping the configurations *frozen*. For simplicity, we consider frictionless contact first

$$\mathbf{r}(\mathbf{v}) = \mathbf{M}_0 \mathbf{v} - \mathbf{p}^* - h \mathbf{J}_{c,0}^T \mathbf{f}_n(\mathbf{q}, \mathbf{v}), \quad (18)$$

where notice we decided to *freeze* the contact Jacobian but not the compliant contact forces.

To treat the normal contact forces implicitly while still freezing the configuration at \mathbf{q}_0 , we propose the following first-order estimation of the penetration distance δ_i at the i th contact point,

$$\delta_i \approx \delta_{i,0} - h v_{n,i}, \quad (19)$$

where we used the fact that $\dot{\delta}_i = -v_{n,i}$. We can then substitute Eq. (19) into Eq. (5) to express the compliant forces, $\pi_i(v_{n,i})$, as function of the normal velocities and group them into a vector, $\boldsymbol{\pi}(\mathbf{v}_n) \in \mathbb{R}^{n_c}$. We then express the normal forces as $\mathbf{f}_n(\mathbf{v}_n) = \hat{\mathbf{N}}_0 \boldsymbol{\pi}(\mathbf{v}_n)$. This can then be used in Eq. (18) to compute the Jacobian of the residual,

$$\begin{aligned} \nabla_{\mathbf{v}} \mathbf{r}(\mathbf{v}) &= \mathbf{M}_0 - h \mathbf{J}_{c,0}^T \nabla_{\mathbf{v}} \mathbf{f}_n(\mathbf{v}_n), \\ &= \mathbf{M}_0 - h \mathbf{J}_{c,0}^T \nabla_{\mathbf{v}_n} \mathbf{f}_n(\mathbf{v}_n) \mathbf{J}_{n,0}, \\ &= \mathbf{M}_0 - h \mathbf{J}_{c,0}^T \hat{\mathbf{N}}_0 \nabla_{\mathbf{v}_n} \boldsymbol{\pi}(\mathbf{v}_n) \mathbf{J}_{n,0}, \\ &= \mathbf{M}_0 + h \mathbf{W}_{nn,0}, \end{aligned} \quad (20)$$

where we introduced the normal direction Delassus operator $\mathbf{W}_{nn,0} = -\mathbf{J}_{n,0}^T \nabla_{\mathbf{v}_n} \boldsymbol{\pi}(\mathbf{v}_n) \mathbf{J}_{n,0}$ and $\nabla_{\mathbf{v}_n} \boldsymbol{\pi}(\mathbf{v}_n) = \text{diag}(\{d\pi_i/dv_{n,i}\})$. Eq. (5) implies $d\pi_i/dv_{n,i} \leq 0$, and thus $\mathbf{W}_{nn,0} \succeq 0$. Therefore in Eq. (20) $\nabla_{\mathbf{v}} \mathbf{r}(\mathbf{v}) \succ 0$, and its inverse always exists, explaining the high stability of this scheme for contact problems without friction.

C. TAMSIs: Two-Way Coupled Scheme

TAMSIs is a semi-implicit Euler scheme in the friction forces as introduced in Section V-A modified to implicitly couple the compliant contact forces using the approximation in Eq. (19). It is a *two-way coupled* scheme in that, in addition to normal forces feeding into the computation of the friction forces through Eq. (6), friction forces feedback implicitly into the normal forces.

Freezing the position kinematics to \mathbf{q}_0 and using the approximation in Eq. (19), the full TAMSIs residual becomes

$$\mathbf{r}(\mathbf{v}) = \mathbf{M}_0 \mathbf{v} - \mathbf{p}^* - h \mathbf{J}_{n,0}^T \boldsymbol{\pi}(\mathbf{v}_n) - h \mathbf{J}_{t,0}^T \mathbf{f}_t(\mathbf{v}_t). \quad (21)$$

We can then use the results from the previous sections, except we also need to take into account how the friction

forces in Eq. (6) change due to changes in the normal forces. Instead of the result in Eq. (15), we now have

$$\begin{aligned} \nabla_{\mathbf{v}} \boldsymbol{\tau}_t &= \mathbf{J}_{t,0}^T \nabla_{\mathbf{v}} \mathbf{f}_t, \\ &= \mathbf{J}_{t,0}^T \nabla_{\mathbf{v}_t} \mathbf{f}_t \mathbf{J}_{t,0} + \mathbf{J}_{t,0}^T \nabla_{\mathbf{v}_n} \mathbf{f}_t \mathbf{J}_{n,0}, \\ &= -\mathbf{W}_{tt} - \mathbf{W}_{tn}, \end{aligned} \quad (22)$$

where $\nabla_{\mathbf{v}_n} \mathbf{f}_t = -\text{diag}(\{\tilde{\mu}_i \hat{\mathbf{v}}_{t,i} d\pi_i/dv_{n,i}\}) \in \mathbb{R}^{3n_c \times n_c}$. We define the generally non-symmetric operator $\mathbf{W}_{tn} = -\mathbf{J}_{t,0}^T \nabla_{\mathbf{v}_n} \mathbf{f}_t \mathbf{J}_{n,0} \in \mathbb{R}^{n_v \times n_v}$ which introduces the additional two-way coupling between compliance in the normal direction and friction in the tangential direction.

Using these results we can write the Jacobian of the TAMSIs scheme as

$$\nabla_{\mathbf{v}} \mathbf{r}(\mathbf{v}) = \mathbf{M}_0 + h (\mathbf{W}_{nn,0} + \mathbf{W}_{tt,0} + \mathbf{W}_{tn,0}). \quad (23)$$

Notice that, due to $\mathbf{W}_{tn,0}$, the Jacobian is not symmetric and in general not invertible; however, since $\mathbf{M}_0 \succ 0$, the Jacobian is invertible for sufficiently small h .

After computing the next Newton iteration using this Jacobian, we use TALS to selectively backtrack the iteration.

D. Implementation Details

We implemented a single-threaded TAMSIs in the open-source robotics toolbox Drake [24]. At this moment, our initial implementation did not focus on performance but rather on proving the stability and robustness of the method. Therefore our initial implementation forms the Newton-Raphson Jacobian in Eq. (23) explicitly. Forming \mathbf{M}_0 is an $\mathcal{O}(n_v^2)$ operation, while computing the Delassus operators in Eq. (23) via explicitly multiplying the Jacobians is $\mathcal{O}(n_v^2 \cdot n_c)$. The factorization of $\nabla_{\mathbf{v}} \mathbf{r}$ is an $\mathcal{O}(n_v^3)$ operation and we expect it to dominate for large scenarios. However, even in our largest simulations with hundreds of DOFs and hundreds of contacts [29], forming this Jacobian is still the most expensive operation, followed by its factorization. Performance can be greatly improved by using matrix-free operators instead of explicitly forming the contact Jacobians, allowing the assembly of $\nabla_{\mathbf{v}} \mathbf{r}$ in $\mathcal{O}(n_v^2)$. Further improvements include exploiting the sparsity pattern of $\nabla_{\mathbf{v}} \mathbf{r}$ during both assembly and factorization as well as parallelization for large scenes with hundreds of objects.

VI. RESULTS AND DISCUSSION

We present a series of simulation test cases to assess robustness, accuracy and performance. For all cases we use a regularized friction parameter of $v_s = 10^{-4}$ m/s, which results in negligible sliding during stiction.

A. Parallel Jaw Gripper

To assess the robustness of our method in a relevant manipulation task with large external disturbances, we simulate a parallel jaw gripper holding a mug and forced to oscillate vertically with a period of $T = 0.5$ s and an amplitude of $A = 15$ cm, see Fig. 6. To stress the solver in a situation with slip-stick transitions, we chose a low coefficient of friction of $\mu = 0.1$ and a grip force of 10 N. The mug is 10 cm tall with a radius of 4 cm and a mass of 100 g. There is no gravity.

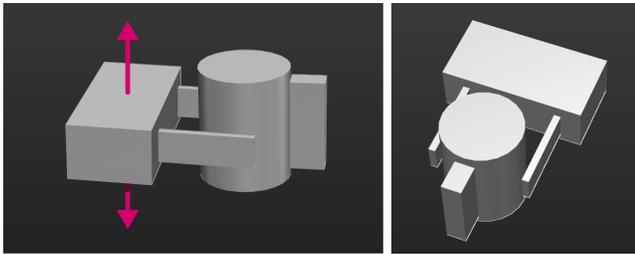


Fig. 6: Parallel jaw gripper model. We shake the mug vertically to assess robustness.

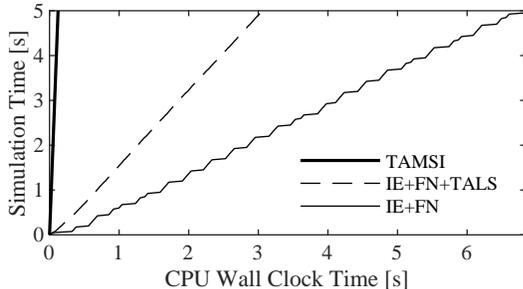


Fig. 7: Simulated time vs. wall-clock time for the parallel jaw gripper case. All three runs use $h = 3$ ms. TAMSI is the fastest, on the far left. The horizontal plateaus in the IE+FN run indicate that without TALS, the integrator slows down since it takes smaller step sizes for convergence control.

With a step $h = 3$ ms TAMSI completes 5 seconds of simulation time with a single thread on an Intel i7-6900K 3.2 GHz CPU in 123 ms of wall-clock time, or $40\times$ real-time rate. In contrast, IE+FN with TALS is $25\times$ slower than TAMSI and IE+FN without TALS is $55\times$ slower than TAMSI. Simulated time vs wall-clock time for each is shown in Fig. 7. Next, we perform a convergence study of TAMSI running with different time steps and estimate errors against a reference solution using $h = 10^{-7}$ s, Fig. 8. As expected, TAMSI exhibits first-order convergence with step size.

B. Allegro Hand

We simulate a 16 DOF Allegro hand controlled in open-loop to perform a periodic reorientation of a mug, see Fig. 9 and the supplemental video accompanying this manuscript. This interesting system includes multiple points of contact, complex geometry and a large number of DOFs.

We observe that TALS does not improve the performance of implicit Euler — as the system configuration \mathbf{q} changes, the tangent space reorients, and TALS is unable to properly control iterations with transitions into stiction. TAMSI however, with a time step $h = 0.7$ ms, runs $7.8\times$ faster than our fastest integrator setup, the error-controlled implicit Euler using quasi-Newton when solving to a (loose) accuracy of $a = 0.01$ m/s. Our single-threaded TAMSI completes 15 seconds simulated time on an Intel i7-6900K 3.2 GHz CPU in 7.0 s of wall-clock time, or $2.14\times$ real-time rate.

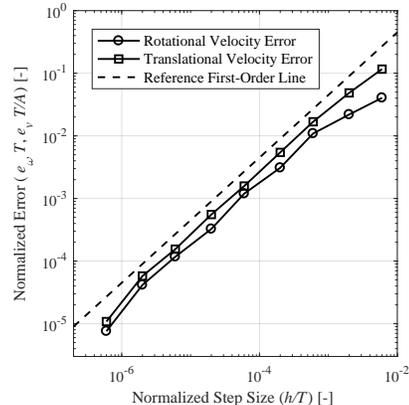


Fig. 8: L^2 -norm of the generalized velocity errors against time step, nondimensionalized with the problem's period T and amplitude A . Both translational and rotational DOFs exhibit first-order convergence as expected.

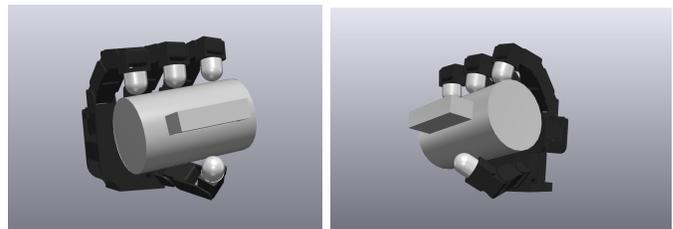


Fig. 9: Model of a 16 DOF Allegro hand holding a mug. As the hand reorients the mug, multiple simultaneous points of contact securing the grasp are made.

C. Sim-to-Real Experiments

We demonstrate our simulation capability in two manipulation tasks with closed-loop control involving force-feedback and manipuland-world contact.

A Kuka IIWA arm (7 DOF) is outfitted with a Schunk WSG 50 gripper. We use an inverse dynamics controller with gains in the simulation set to best match reality, even though the specifics of Kuka's proprietary joint-level controller are unavailable. The controller process tracks a prescribed sequence of Cartesian end-effector keyframe poses and computes desired joint-space trajectories using Jacobian IK. We use force feedback to regulate the grasp force and judge for placement success.

The gripper surfaces are slightly compliant. We identify the stiffness parameters of the point contact model by matching the level of compliance observed in the real hardware.

In the first task the robot is commanded to grab a water bottle, perform a vigorous shaking motion, and place the bottle at a new location, demonstrating the robustness of the method to strong disturbances; see the supplemental video.

In the second task a Kuka arm in a kitchen scenario grabs a mug and uses external contact with the sink to help reorient the mug, relying on slip within the fingers of the gripper; see Fig. 1 and the accompanying supplemental video. TAMSI robustly handles multiple points of contact, changes in the contact configuration, and stick-slip transitions.

TAMSI enabled us to prototype this task entirely in simulation with controllers that transitioned seamlessly to reality.

VII. CONCLUSIONS

In this work we systematically analyzed implicit integration for multibody problems with compliant contact and regularized friction. Error-controlled implicit Euler with quasi-Newton iterations performs best, though it spends a significant amount of time resolving the high-frequency transition dynamics introduced by regularized friction. Our new TALS method helps to address this problem, though its performance degrades as the configuration of the system changes and contact surfaces reorient. This observation led us to develop our novel TAMSI method. TAMSI approximates penetration distances to first-order to couple contact forces implicitly while performing only a single geometric query at the beginning of a time step, resulting in improved performance. We demonstrate the added robustness and performance of TAMSI with simulations of relevant manipulation tasks. Sim-to-real comparisons show the usage of TAMSI to prototype controllers in simulation that transfer effectively to reality.

Additional examples can be found in the “examples” source directory of the open-source robotics toolbox Drake [24] and in the “Drake Gallery” section of the documentation.

Ongoing work conducted at the Toyota Research Institute is leveraging the proposed method for prototyping and validating controllers for robot manipulation in dense cluttered environments [29] and extending TAMSI to work with more sophisticated contact models [21].

ACKNOWLEDGMENT

We thank the reviewers for their feedback, Allison Henry and Siyuan Feng for technical support, and Evan Drumwright and Russ Tedrake for helpful discussions. The work of Ante Qu was supported in part by the National Science Foundation under grant DGE-1656518. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, *et al.*, “Using simulation and domain adaptation to improve efficiency of deep robotic grasping,” in *IEEE International Conference on Robotics and Automation*, 2018, pp. 4243–4250.
- [2] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, “Closing the sim-to-real loop: Adapting simulation randomization with real world experience,” in *IEEE International Conference on Robotics and Automation*, 2019, pp. 8973–8979.
- [3] S. Kolev and E. Todorov, “Physically consistent state estimation and system identification for contacts,” in *IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, 2015, pp. 1036–1043.
- [4] D. J. Duff, T. Mörwald, R. Stolkin, and J. Wyatt, “Physical simulation for monocular 3d model based tracking,” in *IEEE International Conference on Robotics and Automation*, 2011, pp. 5218–5225.
- [5] S. Li, S. Lyu, J. Trinkle, and W. Burgard, “A comparative study of contact models for contact-aware state estimation,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015, pp. 5059–5064.
- [6] S. Hogan and K. U. Kristiansen, “On the regularization of impact without collision: the painlevé paradox and compliance,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 473, no. 2202, p. 20160773, 2017.
- [7] 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.
- [8] 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.
- [9] R. Cottle, J. Pang, and R. Stone, *The Linear Complementarity Problem*, ser. Computer science and scientific computing. Academic Press, 1992.
- [10] D. M. Kaufman, S. Sueda, D. L. James, and D. K. Pai, “Staggered projections for frictional contact in multibody systems,” *ACM Transactions on Graphics (TOG)*, vol. 27, no. 5, p. 164, 2008.
- [11] J. Li, G. Daviet, R. Narain, F. Bertails-Descoubes, M. Overby, G. E. Brown, and L. Boissieux, “An implicit frictional contact solver for adaptive cloth simulation,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, p. 52, 2018.
- [12] E. Todorov, “Implicit nonlinear complementarity: A new approach to contact dynamics,” in *IEEE International Conference on Robotics and Automation*, 2010, pp. 2322–2329.
- [13] M. Macklin, K. Erleben, M. Müller, N. Chentanez, S. Jeschke, and V. Makoviychuk, “Non-smooth newton methods for deformable multi-body dynamics,” *ACM Transactions on Graphics (TOG)*, vol. 38, no. 5, p. 140, 2019.
- [14] C. Lacoursiere and M. Linde, “Spook: a variational time-stepping scheme for rigid multibody systems subject to dry frictional contacts,” Umeå University, UMINF report 11.09, 2011.
- [15] E. Catto, “Soft constraints: Reinventing the spring,” in *Game Developer Conference*, 2011.
- [16] E. Coumans, “Bullet physics library,” 2015. [Online]. Available: <http://bulletphysics.org>
- [17] R. Smith, “Open dynamics engine.” [Online]. Available: <http://www.ode.org>
- [18] L. Luo and M. Nahon, “A compliant contact model including interference geometry for polyhedral objects,” *Journal of Computational and Nonlinear Dynamics*, vol. 1, no. 2, pp. 150–159, 2006.
- [19] Y. Gonthier, “Contact dynamics modelling for robotic task simulation,” Ph.D. dissertation, University of Waterloo, 2007.
- [20] M. A. Sherman, A. Seth, and S. L. Delp, “Simbody: multibody dynamics for biomedical research,” *Procedia Iutam*, vol. 2, pp. 241–261, 2011.
- [21] 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 *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 8238–8245.
- [22] D. Baraff, “Coping with friction for non-penetrating rigid body simulation,” *ACM Computer Graphics (SIGGRAPH 91)*, vol. 25, no. 4, pp. 31–41, 1991.
- [23] 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, 1975.
- [24] R. Tedrake and the Drake Development Team, “Drake: Model-based design and verification for robotics,” 2019. [Online]. Available: <https://drake.mit.edu>
- [25] E. Hairer, S. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations II. Stiff and Differential-Algebraic Problems*, 2nd ed. Springer, 01 1996, vol. 14.
- [26] Y. Zhu, R. Bridson, and D. M. Kaufman, “Blended cured quasi-newton for distortion optimization,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, p. 40, 2018.
- [27] T. K. Uchida, M. A. Sherman, and S. L. Delp, “Making a meaningful impact: modelling simultaneous frictional collisions in spatial multibody systems,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 471, no. 2177, p. 20140859, 2015.
- [28] C. Lacoursiere, “Stabilizing gyroscopic forces in rigid multibody simulations,” Umeå University, UMINF report 06.05, 2006.
- [29] R. Tedrake, “TRI taking on the hard problems in manipulation research toward making human-assist robots reliable and robust,” June 2019. [Online]. Available: <https://www.tri.global/news/tri-taking-on-the-hard-problems-in-manipulation-re-2019-6-27>