# Implicit frictional dynamics with soft constraints

Egor Larionov, Andreas Longva, Uri M. Ascher, Jan Bender, and Dinesh K. Pai

Abstract-Dynamics simulation with frictional contacts is important for a wide range of applications, from cloth simulation to object manipulation. Recent methods using smoothed lagged friction forces have enabled robust and differentiable simulation of elastodynamics with friction. However, the resulting frictional behavior can be inaccurate and may not converge to analytic solutions. Here we evaluate the accuracy of lagged friction models in comparison with implicit frictional contact systems. We show that major inaccuracies near the stick-slip threshold in such systems are caused by lagging of friction forces rather than by smoothing the Coulomb friction curve. Furthermore, we demonstrate how systems involving implicit or lagged friction can be correctly used with higher-order time integration and highlight limitations in earlier attempts. We demonstrate how to exploit forward-mode automatic differentiation to simplify and, in some cases, improve the performance of the inexact Newton method. Finally, we show that other complex phenomena can also be simulated effectively while maintaining smoothness of the entire system. We extend our method to exhibit stick-slip frictional behavior and preserve volume on compressible and nearly-incompressible media using soft constraints.

Index Terms-dry friction, contact, elasticity, deformable object dynamics

#### I. INTRODUCTION

Modern simulation pipelines in computer graphics and engineering involve contact handling. This ensures that simulated objects do not interpenetrate each other as they interact. During this interaction, friction ensures that solid objects are held in place or are otherwise limited in how they move. This work focuses on the accuracy and effectiveness of smooth friction models.

In recent years we have seen a resurgence of promising work in developing more robust methods for simulating frictional contact in computer graphics. This problem is particularly difficult since friction and contact cannot be simultaneously described by a single energy potential [2]. This precludes formulating the frictional contact problem as a single energy minimization. However, energy minimization used for solving discretized ordinary differential equations (ODEs) over each time step has remained popular in graphics, due to its flexibility and robustness characteristics. Unfortunately, optimization based solvers used for dynamics equations require specialized algorithms for handling frictional contacts, which necessarily produces drawbacks in accuracy or robustness.

In this work, we demonstrate failure cases in popular optimization based frictional contact solvers and propose an alternative method for solving elastodynamic problems with frictional contacts that is simple to implement and accurate in comparison. By resolving the sliding frame and contact

Manuscript created December 2023; This work was supported by the Natural Sciences and Engineering Research Council Discovery Grant, and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number BE 5132/5-1.

forces implicitly when computing friction, our method can produce more accurate friction behavior, and it requires no additional iterations or specialized mechanisms for coupling friction, contact and elasticity.

1

Frictional contact is traditionally modeled as a non-smooth problem requiring sophisticated tools. In particular, nonsmooth integrators, root finding or optimization techniques are needed for handling inclusion terms in the mathematical model. This drastically complicates the problem and substantially limits the number of solution approaches. While nonsmoothness is required to guarantee absolute sticking, it is not generally necessary if simulations are limited in time. In fact, when observed on a microscale, even dry friction responds continuously to changes in velocity [3]. Accordingly, we adopt a smooth friction formulation. We show that when applied fully implicitly in the equations of motion, it can produce predictable sticking. In contrast, our study shows that the popular approach of lagging friction causes inaccurate and time step dependent sticking behavior. In particular, we demonstrate the importance of evaluating the sliding basis (defined in Section III-C) and contact forces implicitly for accurate friction simulation.

Finally, to maintain smoothness of the entire problem we employ a penalty force for contact resolution and a smooth implicit surface model proposed by Larionov et al. [4] for representing the contact surface. We also show how additional soft constraints can be added to the system for controlling the volume of an object.

Maintaining smoothness of the entire system makes it easier to solve and allows derivatives to be propagated through each step of the simulation. However, full differentiable simulation is outside the scope of this work.

In summary, our primary contributions are

- A comparison between lagged and fully implicit friction models focusing on frictional accuracy.
- A stable implementation of high-order time integration for frictional contact problems that is applicable to both lagged and fully implicit formulations.
- A novel smooth Stribeck friction model used for modeling stick-slip effects due to the microstructure of contacting surfaces.
- A physically-based volume change penalty for controlling compressibility in compressible and nearly incompressible regions.
- A simple adaptive penalty stiffening strategy for effectively resolving interpenetrations with penalty-based contact methods.

Furthermore, to better characterize the instability of single point frictional contacts, we present an eigen-analysis of a 2D point contact subject to friction in Section I of the supplemental document.



Fig. 1: An upside-down bowl is lifted using 3 soft pads via friction. The bowl is simulated using the lagged friction model from [1] (top row) and our fully implicit method (Eq. (16), bottom row). The chosen frames between 1 and 150, are frames 35, 40 and 100, at which point the bowl slips out for time steps h = 0.005 s, 0.0025 s, and 0.00125 s respectively for the lagged method. Using fully implicit integration, the bowl sticks even at the largest time step h = 0.005 s as shown. For the pads,  $\rho = 1000 \text{ kg/m}^3$  (density), E = 600 KPa (Young's modulus), and  $\nu = 0.49$  (Poisson ratio). For the bowl,  $\rho = 400 \text{ kg/m}^3$ , E = 11000 KPa and  $\nu = 0.1$ . The friction coefficient is set to  $\mu = 0.65$  between the two. See Figure 5 for the plot of the bowl height in all tested configurations.

#### II. RELATED WORK

Simulating the dynamics of deformable elastic objects [5] and cloth [6] has been an active area of research in graphics for decades. Methods requiring accuracy often reach for a finite element method (FEM) [7], while methods aiming for performance reach for position-based techniques [8], [9], [10], or projective-dynamics [11]. Our work targets simulation accuracy, and thus stays close to the tried and tested FEM.

a) Frictional Contact: In recent years, lots of attention has been dedicated towards robust and accurate contact and friction solutions. Earlier works in graphics developed nonsmooth methods to resolve contact and friction forces at the end of the time step obtaining solutions faithful to the Coulomb model. Kaufman et al. [12] proposed a predictorcorrector method to solve for friction and contact separately from elasticity equations. This was later extended to highorder time integration [13] and separately to full FEM simulations [4] while maintaining the decoupling. Other methods reformulate the problem as a non-smooth root finding problem [14], [15], [16], [17] or using proximal algorithms [18]. More recently, more attention was brought towards modeling friction as a smoothly changing force at the stick-slip limit [19], [1]. This allows each simulation step to remain differentiable. Geilinger et al. [19] favored a more traditional root-finding solver combining friction and contact forces with elastic equations of motion. In contrast, Li et al. [1] proposed a robust optimization framework to solve for contact and lagged friction forces. Unfortunately, even with multiple iterations, the lagged friction approach might not converge to an accurate friction solution (cf. Figure 1), which is especially noticeable in sticking configurations close to the slip threshold. Furthermore, their proposed method for higher-order integration is not applied to the contact solve causing instabilities. In this work, we demonstrate and address these shortcomings using a solution method that favors friction accuracy at the cost of some robustness, while maintaining smoothness of the system. While non-smooth methods for frictional contact are interesting and have been physically validated [13], our work is focused solely on smooth methods.

2

b) Higher-order integrators for contact problems: Most contact formulations, especially those formulated in terms of constraints, intrinsically rely on a particular choice of time discretization, which is usually backward Euler (BE). However, the highly dissipative characteristics of BE have motivated the use of higher-order schemes like TR-BDF2 or SDIRK2, which preserve high energy dynamics while maintaining stability [20], [21]. A benefit of smooth contact models based on penalty or barrier functions is that both normal and friction forces are defined with explicit formulas, as opposed to implicitly defined through constraints. This makes it possible to apply higher-order integrators directly, as demonstrated by Geilinger et al. [19] for BDF2. Li et al. [1] applied trapezoid rule (TR), however it is only applied to non-contact forces, which we show causes stability issues. Brown et al. [22] focus on first-order methods and apply TR-BDF2 to a nonsmooth optimization-based contact model with lagged friction, though this has unclear implications for high-accuracy secondorder methods. Here, we show how high-order methods can be applied to systems with lagged or implicit friction methods.

c) Volume preservation: Many solids exhibit volume preserving behavior. We focus primarily on inflated objects like tires and sports equipment (e.g. sports balls), as well as nearly incompressible objects like the human body. Inflated objects are typically simulated using soft constraints [23], [24], [25], [10], [26], where in effect the volume change of an object is penalized assuming Boyle's law for ideal gasses. These methods are effective, however, their physical accuracy is rarely questioned. Incompressible or nearly-incompressible materials are often modeled with stiff Poisson's ratios [27] or hard volume preservation constraints [28]. In contrast to previous work, we propose a unified physically-based penalty formulation for volume preservation that models both compressible and nearly-incompressible objects using a single penalty controlled by a physical compression coefficient. Additionally, we illustrate how the volume preservation models

differ, indicating the expected trade-offs in each model.

# III. METHOD

Here we develop the equations involved in solving for the motion of deformable objects subject to frictional contacts.

A set of generalized coordinates  $\mathbf{q}(t) \in \mathbb{R}^m$  (e.g. stacked vertex positions of a FEM mesh in 3D), represents a system of solids moving through time t. A sparse symmetric positive definite (SPD)  $m \times m$  matrix M denotes generalized mass. In the following we omit the time dependence for brevity, but later reintroduce it when discussing time discretization schemes.

Using dot notation for time derivatives and with generalized velocities  $\mathbf{v} = \dot{\mathbf{q}}$  we can write the force balance equation as

$$\begin{aligned} \mathbf{M} \dot{\mathbf{v}} &= \mathbf{f}(\mathbf{q}, \mathbf{v}) \end{aligned} (1a) \\ \mathbf{f}(\mathbf{q}, \mathbf{v}) &= \mathbf{f}_e(\mathbf{q}) + \mathbf{f}_d(\mathbf{q}, \mathbf{v}) + \mathbf{f}_c(\mathbf{q}) + \mathbf{f}_f(\mathbf{q}, \mathbf{v}) + \mathbf{f}_g(\mathbf{q}), \end{aligned}$$

where  $\mathbf{f}_e$  are elastic forces,  $\mathbf{f}_d$  are damping forces,  $\mathbf{f}_c$  are contact forces,  $\mathbf{f}_f$  is friction, and  $\mathbf{f}_g$  are external forces such as gravity.

## A. Elasticity and Damping

Elastic forces are typically derived from a configuration dependent energy potential  $W(\mathbf{q})$  as  $\mathbf{f}_e(\mathbf{q}) = -\frac{\partial}{\partial \mathbf{q}}W(\mathbf{q})$ . The elastic potential W can be defined by the classic linear, neo-Hookean, StVK or Mooney-Rivlin models [29], or even a data-driven model [30]. We will focus on neo-Hookean materials for both solids and cloth. We then define the *stiffness matrix*  $\mathbf{K}(\mathbf{q}) = -\frac{\partial}{\partial \mathbf{q}}\mathbf{f}_e(\mathbf{q})$ , which dictates how resistant an object is to deformation. For nonlinear models like neo-Hookean elasticity,  $\mathbf{K}$  may be indefinite, which is important to know when picking an appropriate linear solver.

Damping forces are often defined by  $\mathbf{f}_d(\mathbf{q}, \mathbf{v}) = -\mathbf{D}(\mathbf{q})\mathbf{v}$ , where **D** is a symmetric, and usually positive semi-definite, matrix. We use the Rayleigh damping model for simplicity where  $\mathbf{D} = \alpha \mathbf{M} + \beta \mathbf{K}$  for some constants  $\alpha, \beta \ge 0$ .

# B. Contact

Traditionally contact constraints have been formulated with a nonnegativity constraint on some "gap" function d(q) that roughly determines how far objects are away from each other. This function d may be closely related to a component-wise signed distance function, however, generally it merely needs to be continuous, monotonically increasing in the direction of separation, and constant at the surface. We define  $d_i$  for each potential contact point i such that  $\mathbf{d} = (d_1, d_2, \dots, d_n)$ , where n is the total number of potential contacts. Here we use the contact model of Larionov et al. [4], where surface vertices of one object are constrained to have non-negative potential values when evaluated against a smooth implicit function d closely approximating the surface of another object. Interestingly, if we allow objects some small separation tolerance at equilibrium, we can reformulate this constraint as an equality constraint by using a soft-max [19] or a truncated log barrier [1] function. These types of equality constraints greatly

simplify the contact problem and have shown tremendous success in practice.

One disadvantage of log-barrier formulations is that the initial state must be free of collisions prior to the optimization step in order to avoid infinite energies and undefined derivatives. In the absence of thin features or risk of tunneling, it is sufficient to use a simple penalty function to resolve interpenetrating geometry. In this work we choose to use penalty-based contacts for simplicity, however, our formulation is fully compatible with a log-barrier method coupled with continuous collision detection (CCD) as proposed by Li et al. [1]. The idea is to help our solver guide interpenetrating meshes out of intersecting configurations. We define a cubic contact penalty by

$$b(x; \delta, \kappa) = \kappa \begin{cases} -\frac{1}{\delta}(x - \delta)^3 & \text{ if } x < \delta \\ 0 & \text{ otherwise} \end{cases}$$

where  $\delta > 0$  is the thickness tolerance and  $\kappa > 0$  is a contact stiffness parameter that will need to be automatically increased to ensure that no surface vertices of one object are penetrating the implicit surface of another at the end of the time step. Here *b* corresponds to the first non-zero term in the Taylor expansion of the truncated log-barrier used by Li et al. [1], but unlike the log-barrier it is well-defined also for negative arguments. The penalty is applied to each pair of contact points with distance  $d_i$  giving us an aggregate contact energy

$$W_c(\mathbf{d}) = \sum_{i=1}^n b(d_i; \delta, \kappa).$$

Now the contact force can be written simply as the negative energy derivative

$$\mathbf{f}_{c}(\mathbf{q})^{\top} = -\frac{\partial W_{c}}{\partial \mathbf{q}} = \boldsymbol{\lambda}(\mathbf{q})^{\top} \frac{\partial \mathbf{d}}{\partial \mathbf{q}}, \quad \text{where} \quad \boldsymbol{\lambda}(\mathbf{q})^{\top} = -\frac{\partial W_{c}}{\partial \mathbf{d}}$$
(2)

is the stacked vector of contact force magnitudes. In effect our contact formulation enforces the equality constraint  $W_c(\mathbf{d}(\mathbf{q})) = 0.$ 

## C. Friction

(1b)

We define the *contact Jacobian*  $\mathbf{J}_c(\mathbf{q})$  and tangent basis  $\mathbf{B}(\mathbf{q})$  over all potential contact points as in [4]. Then  $\mathbf{T}(\mathbf{q}) = \mathbf{J}_c(\mathbf{q})^\top \mathbf{B}(\mathbf{q})$  is the  $m \times n$  matrix defining the *sliding basis* [1]. In short, this matrix maps forces in contact space to generalized forces in configuration space.

We can now derive the smoothed friction force [1], [19] from first principles. For each contact i, the maximum dissipation principle (MDP) postulates that friction force ought to maximally oppose relative velocity

$$\mathbf{f}_{f,i}(\mathbf{v};\mu) = \operatorname*{argmax}_{\|\mathbf{y}\| \le \mu \lambda_i} (-\bar{\mathbf{v}}_i^\top \mathbf{y}), \tag{3}$$

where  $\mu$  is the coefficient of friction, which limits the friction force<sup>1</sup> and  $\bar{\mathbf{v}}_i \in \mathbb{R}^2$  is the relative tangential velocity at contact point *i*. The contact force magnitude  $\lambda_i$  is the *i*th element of

<sup>1</sup>Unless otherwise specified,  $\|\cdot\|$  refers to the Euclidean norm.

 $\lambda$  as defined in Eq. (2). We can solve Eq. (3) explicitly with an inclusion

$$\mathbf{f}_{f,i}(\mathbf{v};\mu) \in -\mu\lambda_i \begin{cases} \{\bar{\mathbf{v}}_i/\|\bar{\mathbf{v}}_i\|\} & \text{if } \|\bar{\mathbf{v}}_i\| > 0\\ \{\mathbf{u} \in \mathbb{R}^2 : \|\mathbf{u}\| \le 1\} & \text{otherwise.} \end{cases}$$
(4)

This is commonly referred to as Coulomb friction. Unfortunately, the non-smoothness around  $\|\bar{\mathbf{v}}_i\| = 0$  calls for nonsmooth optimization or root-finding techniques [12], [14], [18], making this problem numerically challenging. Another disadvantage of non-smoothness is that it greatly complicates differentiation of the solver, which can be critical for solving inverse problems efficiently. We opt to approximate Coulomb friction using a smoothed model [1], [19]. Since most animations call for relatively short time frames, we typically do not require absolute sticking. Interestingly, smooth friction models have been proposed in older engineering literature [31], [3], [32] to improve hysteretic behavior and alleviate numerical difficulties. A simple smoothing [19] of Eq. (4) can be written as

$$\mathbf{f}_{f,i}(\mathbf{v};\mu) \approx -\mu\lambda_i s(\|\bar{\mathbf{v}}_i\|)\boldsymbol{\eta}(\bar{\mathbf{v}}_i),\tag{5}$$

where  $oldsymbol{\eta}:\mathbb{R}^2
ightarrow\mathbb{R}^2$  defines the per-contact nonlinearity

$$\boldsymbol{\eta}(\bar{\mathbf{v}}_i) = \begin{cases} \bar{\mathbf{v}}_i / \|\bar{\mathbf{v}}_i\| & \text{if } \|\bar{\mathbf{v}}_i\| > 0\\ 0 & \text{otherwise,} \end{cases}$$
(6)

and the function s defines the pre-sliding transition. A popular  $C^1$  option for s as depicted in Figure 2a for different values of  $\epsilon$ , is

$$s(v;\epsilon) = \begin{cases} \frac{2v}{\epsilon} - \frac{v^2}{\epsilon^2} & \text{if } v < \epsilon\\ 1 & \text{otherwise.} \end{cases}$$
(7)

Simpler and smoother functions exist, however this choice is convenient since  $\epsilon$  directly controls the sliding velocity tolerance during sticking.

We can then easily express more complex friction models by substituting  $\mu\lambda_i s(\|\bar{\mathbf{v}}_i\|)$  with a velocity dependent function  $c(v, \lambda_i; \mu)$  of multiple coefficients  $\mu$ . For instance, Stribeck and viscous friction [33] can be expressed with  $\mu = (\mu_d, \mu_s, \mu_v)$  representing dynamic, static and viscous friction coefficients respectively:

$$\begin{aligned} \mathbf{f}_{f,i}(\mathbf{v};\boldsymbol{\mu}) &= -c(\|\bar{\mathbf{v}}_i\|,\lambda_i;\boldsymbol{\mu})\boldsymbol{\eta}(\bar{\mathbf{v}}_i),\\ c(v,\lambda_i;\boldsymbol{\mu}) &= (\mu_d + (\mu_s - \mu_d)g(v/v_s))s(v;\epsilon)\lambda_i + \mu_v v \quad (8)\\ g(x) &= \begin{cases} (2x+1)(x-1)^2 & \text{if } x < 1\\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

where g is a compact Gaussian approximation, and  $v_s$  is Stribeck velocity, which defines how gradually friction decays from static into dynamic as velocity increases. Figure 2b illustrates how the different parameters in Eq. (8) control the overall curve. In general,  $v_s$  should be larger than  $\epsilon$ , where values close to  $\epsilon$  may lower the observed/effective static friction force. Setting  $v_s = 10\epsilon$  works well for modeling common dry friction when  $\epsilon$  is small. The Stribeck component is a useful tool for introducing static friction into the smoothed friction model. Our example in Section V-C2 shows how this model can introduce stick-slip behavior from the real world.



(a) Pre-sliding transition s as defined in (7). The function is plotted for different values of  $\epsilon$ , which control the velocity error tolerance.

(b) Stribeck and viscous friction curve f as defined in (8). The effect of each of the material properties in  $\mu$  and parameters  $\epsilon$ and  $v_s$  are labelled.

4

Fig. 2: Components of the friction model.

We can express the nonlinearity in Eq. (6) as a function over all (stacked) relative velocities  $\bar{\mathbf{v}} \in \mathbb{R}^{2n}$  using a diagonal block matrix

$$\mathbf{H}(ar{\mathbf{v}}) = egin{bmatrix} oldsymbol{\eta}(ar{\mathbf{v}}_1) & & \ & \ddots & \ & & oldsymbol{\eta}(ar{\mathbf{v}}_n) \end{bmatrix}$$

Then the total friction force can be written compactly as

$$\mathbf{f}_f(\mathbf{q}, \mathbf{v}) = -\mathbf{T}(\mathbf{q}) \mathbf{H}(\mathbf{T}(\mathbf{q})^\top \mathbf{v}) \mathbf{c}(\mathbf{q}, \mathbf{v}), \tag{9}$$

where  $\mathbf{c}_i = c(\|\bar{\mathbf{v}}_i\|, \lambda_i; \boldsymbol{\mu})$  is the *i*th component of the stacked vector of friction force magnitudes.

#### D. Volume change penalty

In soft tissue simulation, resistance to volume change is typically controlled by Poisson's ratio. This, however, assumes that the simulated body is homogeneous and void of internal structure. For more complex structures like the human body, a zonal constraint is a more suitable method to enforce incompressibility [28]. Compressible objects, however, require a different method altogether. In this section we propose a physically-based and stable model to represent compressible and nearly incompressible objects. In particular, we want to efficiently model inflatable objects like balloons, tires and sports balls, as well as nearly incompressible objects like the human body or other organic matter.

We start from the isothermal compression coefficient [34, Section 5.3] defined by

$$\kappa_v = -\frac{1}{V} \left(\frac{\partial V}{\partial P}\right)_T,\tag{10}$$

where V is the volume of interest<sup>2</sup>, P is internal pressure and the T subscript indicates that temperature is held constant. For compressible continua like air in normal conditions, which behaves like an ideal gas,  $\kappa_v = 1/P$ . For nearly incompressible continua like water at room temperature,  $\kappa_v \approx 4.6 \times 10^{-5}$ atm<sup>-1</sup> is relatively constant. Assuming rest volume  $V_0$  and initial pressure  $P_0 = 1$  atm, we can derive the work W needed

<sup>2</sup>For instance a region occupied by FEM elements or the volume of a watertight triangle mesh.

to change the volume of the container to V. For an ideal gas PV is constant, which gives

$$W_{ig}(V) = P_0\left(V - V_0\left(1 + \ln\frac{V}{V_0}\right)\right). \tag{11}$$

For a nearly incompressible continuum,  $\kappa_v$  is constant, which yields

$$W_{ni}(V) = \frac{1}{\kappa_v} \left( V_0 - V \left( 1 - \ln \frac{V}{V_0} \right) \right). \tag{12}$$

For details of the derivation see Section II in the supplemental document.

Unfortunately, both models are undefined for negative volumes, which can easily lead to configurations with undefined penalty forces. Taking the second-order approximation of Eq. (12) gives us

$$W_2(V) = \frac{(V - V_0)^2}{2V_0\kappa_v},$$
(13)

which coincides with the second-order approximation of Eq. (11) when  $\kappa_v = 1$ . Thus, our second-order model approximates both compressible and nearly incompressible continua well for small changes in volume as shown in Figure 3. For larger changes in volume, we recommend modeling Eq. (11) directly, since it also approximates Eq. (12) well and volume changes are not significant in nearly incompressible continua.



Fig. 3: Volume change energy. The energy (negative of work) is plotted for the compressible model in Eq. (11) (dotted curve), the nearly incompressible model in Eq. (12) (dashed curve) and the 2nd order approximate model in Eq. (13) (solid curve). Here  $V_0 = 1 \text{ m}^3$ ,  $\kappa_v = 1 \text{ atm}^{-1}$ , and  $P_0 = 1 \text{ atm}$ . The quadratic model approximates both cases, but is ultimately too weak for excessive compression yet too strong during large expansion. Depending on the use case, it may be necessary to model one of Eqs. (11) or (12) directly.

To alleviate the approximation error for scenarios that involve more compression (such as in Figure 9), we recommend decreasing  $\kappa_v$  to produce stronger restorative forces.

The penalty force is then given directly by the negative derivative of Eq. (13) and controlled by the compression parameter  $\kappa_v$ :

$$\mathbf{f}_{v}(\mathbf{q}) = -\frac{(V - V_{0})}{V_{0}\kappa_{v}}\frac{\partial V}{\partial \mathbf{q}}.$$
(14)

This can then be added directly to Eq. (1b). Incidentally, the Jacobian of Eq. (14) is dense, however, it can be approximated by the sparse term involving  $\partial^2 V/\partial q^2$ , which expresses only local force changes. In matrix-free solvers where only matrix-vector products are required, the complete derivative can be computed without hindering performance since the full dense Jacobian is never stored in memory.

### **IV. NUMERICAL METHODS**

In this section we outline and motivate methods for approximately solving the non-linear force balance system (1a).

#### A. Time Integration

The equations of motion (1a) can be discretized in time using a variety of implicit time integration methods. Purely explicit integration schemes are not recommended since they prohibitively restrict the admissible time step size in stiff problems — contact and friction can produce extremely large forces causing instability that is intrinsic to the problem we are trying to solve.

Using standard notation, we assume that at time t we know  $\mathbf{q} = \mathbf{q}^t$  and  $\mathbf{v} = \mathbf{v}^t$ , and employ a step size h to proceed forward in time. The integration methods we consider can be expressed by the momentum balance equation

$$\mathbf{0} = \mathbf{r}(\mathbf{v}^{t+h}; h, \mathbf{f}, \mathbf{M}), \tag{15}$$

where we use superscripts to indicate time. Each integration scheme is characterized by one or more residual functions **r** used to determine the final velocity  $\mathbf{v}^{t+h}$ . Except for trapezoidal scheme, all integrators we consider are L-stable, indicating that they dampen high frequency error components for stiff and highly oscillatory or unstable problems. Elastodynamics with frictional contacts can exhibit instabilities (see Section I.B. in the supplemental document), however, we know that friction is naturally dissipative, and so we expect solutions to behave stably overall. L-stability ensures that any additional stiffness present in the system will not destabilize the numerical solution. For further discussion on stability see [35].

1) Backward Euler: The simplest implicit scheme is backward Euler (BE), defined by

$$\mathbf{r}_{\text{\tiny BE}}(\mathbf{v}^{t+h}) = \mathbf{M}(\mathbf{v}^{t+h} - \mathbf{v}^t) - h\mathbf{f}(\mathbf{q}^t + h\mathbf{v}^{t+h}, \mathbf{v}^{t+h}) = \mathbf{0}.$$
(16)

2) Trapezoidal Rule: A well-known method for mixing explicitly and implicitly determined forces on the system is the trapezoidal rule (TR), defined by the momentum balance

$$\mathbf{r}_{\mathrm{TR}}(\mathbf{v}^{t+h}) = \mathbf{M}(\mathbf{v}^{t+h} - \mathbf{v}^{t}) - \frac{h}{2} \left( \mathbf{f}(\mathbf{q}^{t}, \mathbf{v}^{t}) + \mathbf{f}(\mathbf{q}^{t+h}, \mathbf{v}^{t+h}) \right) = \mathbf{0},$$
  
where  $\mathbf{q}^{t+h} = \mathbf{q}^{t} + \frac{h}{2} \left( \mathbf{v}^{t} + \mathbf{v}^{t+h} \right).$  (17)

Notably, this method is equivalent to the popular implicit Newmark- $\beta$  [36], [37], [38] method with  $\beta = 1/4$  and  $\gamma = 1/2$ . The frictional contact problem we address is not stable, and can generate large stiffnesses for high elastic

moduli, large deformations or due to contact and friction. The smoothed frictional contact problem produces high frequency oscillations, which are exacerbated by TR, whereas ideally we want these to be damped away. See Section I.B. in the supplemental document for a concrete example. In spite of these flaws, TR is still used in practice, often decoupled from the frictional contact problem [1]. In Section V-A3 we demonstrate how properly coupling TR as defined in Eq. (17) can resolve some instabilities in practice.

It is also straightforward to use higher-order L-stable methods within our formulation. BDF2, TR-BDF2 and SDIRK2 are defined in the supplemental document; for further discussion of these methods see [20], [21].

# B. Damped Newton

The momentum balance Eq. (15) can be solved efficiently by second-order root-finding methods like Newton's. In the absence of constraints, this can be seen as a generalization of incremental potential optimization [39], where the merit function is set to be an energy potential<sup>3</sup>  $W(\mathbf{v})$  such that  $\partial W(\mathbf{v})/\partial \mathbf{v} = \mathbf{r}(\mathbf{v})$ . However, in that case to maintain a descent direction,  $\partial \mathbf{r}/\partial \mathbf{v}$  must be appropriately modified to remain positive definite.

Since the presence of friction forces precludes a single potential W for minimization [2], many methods relying on incremental potentials build special workarounds to solve for exact Coulomb-based friction, including staggered projections [12], fixed-point methods [18] and lagged friction [1]. Others employ non-smooth Newton to find roots of a proxy function [14], [15], [40]. We employ the penalty-based frictional contact approach promoted by Geilinger et al. [19]. We extend their approach with additional terms of the Jacobian relating to changes in the sliding basis to ensure a more accurate descent direction and parity with the method in Section IV-C. In addition to contacts against static boundaries, we also demonstrate how this method performs in full bidirectional contact between elastic solids. Empirically we found that in many common cases omitting sliding basis derivatives can speed up simulations, especially when contact Jacobians are complex, however we expect convergence can suffer in cases with complex deforming contact surfaces and large time steps. We leave a comprehensive convergence study to future work since it would require more careful contact handling. For lower resolution examples we use the damped Newton algorithm as defined in Algorithm 1, where the problem Jacobian defined by  $\mathbf{J} = \partial \mathbf{r} / \partial \mathbf{v}$  is a square and potentially non-symmetric matrix (see Section I.A. in the supplemental document). Assuming that  $\mathbf{J}$  is boundedly invertible in the neighborhood of the root, for a sufficiently good initial estimate, damped Newton is guaranteed to converge<sup>4</sup> [41]. While singular Jacobians can cause problems, in our experiments they are rare, and often can be eliminated by decreasing the time step in dynamic simulations.

 $^{4}\mathrm{If}~\mathbf{J}$  is also sufficiently regular then (undamped) Newton convergence is Q-quadratic.

Furthermore, in Section I.A. of the supplemental document we show that our method does not introduce singularities through coupling between elasticity, contact and friction on a single node.

# ALGORITHM 1: DAMPEDNEWTON Input:

 $k_{\max} \leftarrow \text{maximum number of Newton iterations}$  $\mathbf{v} \leftarrow \text{previous velocities}$ 

**Output:**  $\mathbf{v}_k \leftarrow$  velocity for the next time step  $\mathbf{1} \ \mathbf{v}_0 \leftarrow \mathbf{v}$ /\* Initialize velocity \*/ 2 for  $k \leftarrow 0$  to  $k_{max}$  do if SHOULDSTOP( $\mathbf{r}(\mathbf{v}_k), \mathbf{v}_k$ ) then 3 break 4 5 end  $\mathbf{p}_k \leftarrow -\mathbf{J}(\mathbf{v}_k)^{-1}\mathbf{r}(\mathbf{v}_k)$ 6 /\* Set search direction \*/  $\alpha \leftarrow \text{LINESEARCH}(\mathbf{v}_k, \mathbf{p}_k)$ 7  $\mathbf{v}_{k+1} \leftarrow \mathbf{v}_k + \alpha \mathbf{p}_k$ 8 9 end

#### C. Inexact Damped Newton

For large scale problems, it is often preferable to use an iterative linear solver, which can outperform a direct solver when degrees of freedom are sufficiently abundant. We use *inexact Newton* to closely couple the iterative solver with our damped Newton's method.

Since friction forces produce a non-symmetric Jacobian, we chose the biconjugate gradient stabilized (BiCGSTAB) algorithm [42] to find Newton search directions  $\mathbf{p}_k$ . For Jacobian  $\mathbf{J}_k = \partial \mathbf{r}_k / \partial \mathbf{v}$  with  $\mathbf{r}_k = \mathbf{r}(\mathbf{v}_k)$ , the search direction is determined by

$$\|\mathbf{r}_k + \mathbf{J}_k \mathbf{p}_k\| \le \sigma_k \|\mathbf{r}_k\|,$$

where  $\sigma_k = \min(\|\mathbf{r}_k\|^{\varphi}/\|\mathbf{r}_{k-1}\|^{\varphi}, \sigma)$  and  $\varphi = (1 + \sqrt{5})/2$  to maintain Q-quadratic convergence [43].

Using BiCGSTAB as the iterative solver additionally allows one to use forward automatic differentiation to efficiently compute products Jp [44].

The final inexact Newton algorithm is presented in Algorithm 2.

#### D. Contact

To ensure that no penetrations remain (i.e.  $d_i > 0$  for each contact i) at the end of the time step we measure the deepest penetration depth  $\mathbf{d}_{\text{deepest}} = \min_i(\mathbf{d}_i(\mathbf{q}^{t+h}))$ , and bump the contact stiffness parameter  $\kappa$  by the factor  $\frac{b'(d_{\text{deepest}})}{b'(0.5\delta)}$  whenever  $d_{\text{deepest}} < 0$ , where b' is the scalar derivative of b. The same step is then repeated with the new  $\kappa$ . This scheme sets the optimal contact penalty value found by the Newton scheme to appear for contacts  $0.5\delta$  outside the contact surface. As such, in most cases one time step with an additional contact iteration is sufficient before all contacts are resolved. Furthermore,  $\kappa$  is

<sup>&</sup>lt;sup>3</sup>While the original incremental potential is intended to be optimized over positions, the velocity derivatives of all integrators we consider are a constant multiple of positional derivatives. Thus optimizing over velocity here is equivalent to optimizing over positions.

This article has been accepted for publication in IEEE Transactions on Visualization and Computer Graphics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2024.3437417

7

# ALGORITHM 2: INEXACTDAMPEDNEWTON

Input:

 $k_{\max} \leftarrow maximum number of Newton iterations$ 

 $\mathbf{v} \leftarrow \text{previous velocities}$  $c_1 \leftarrow 10^{-4}$  $\sigma \leftarrow 0.01$  $\rho \leftarrow 0.5$ 

**Output:**  $\mathbf{v}_k \leftarrow$  velocity for the next time step  $\mathbf{v}_0 \leftarrow \mathbf{v}$ /\* Initialize velocity \*/ **2** for  $k \leftarrow 0$  to  $k_{max}$  do if SHOULDSTOP( $\mathbf{r}(\mathbf{v}_k), \mathbf{v}_k$ ) then 3 break 4 5 end  $\sigma_k \leftarrow \min(\|\mathbf{r}_k\|^{\varphi}/\|\mathbf{r}_{k-1}\|^{\varphi},\sigma)$ 6 Find  $\mathbf{p}_k$  such that  $\|\mathbf{r}_k + \mathbf{J}_k \mathbf{p}_k\| \leq \sigma_k \|\mathbf{r}_k\|$ 7  $\alpha \leftarrow 1$ 8 /\* Backtracking \*/ while  $\|\mathbf{r}(\mathbf{v} + \alpha \mathbf{p})\| > (1 - c_1 \alpha (1 - \sigma_k)) \|\mathbf{r}(\mathbf{v}_k)\|$  do 9 10  $\alpha \leftarrow \rho \alpha$ end 11  $\sigma_k \leftarrow 1 - \alpha(1 - \sigma_k)$ 12  $\mathbf{v}_{k+1} \leftarrow \mathbf{v}_k + \alpha \mathbf{p}_k$ 13 14 end

never decreased so long as there are active contacts to avoid oscillations at the contact surface.

The downside of this technique is that it compromises the smoothness of the simulation. We postulate that in practice, this may not be problematic in a differentiable pipeline since  $\kappa$  is not changed frequently and subsequent differentiable iterations can carry forward the maximal  $\kappa$  to maintain smoothness.

#### E. Correct TR integration in IPC

Li et al. [1] introduced incremental potential contact (IPC), a robust method to resolve contacts by minimizing incremental potentials with friction being evaluated using lagged positional estimates from the previous time step. Here we propose a simple fix for handling higher-order time integrators in the IPC framework, and show how our formulation relates to this method. With the lagged friction approach, the BE and TR discretizations are given by

$$\mathbf{r}_{\text{BE,IPC}}(\mathbf{v}^{t+h}) = \mathbf{v}^{t+h} - \mathbf{v}^t - h\mathbf{M}^{-1}\mathbf{f}_{lag}(\mathbf{q}^t, \mathbf{q}^{t+h}, \mathbf{v}^{t+h}),$$
(18)

$$\mathbf{r}_{\text{TR,IPC}}(\mathbf{v}^{t+h}) = \mathbf{v}^{t+h} - \mathbf{v}^{t} \\ -\frac{h}{2}\mathbf{M}^{-1}\left(\mathbf{f}_{lag}(\mathbf{q}^{t}, \mathbf{q}^{t+h}, \mathbf{v}^{t+h}) + \mathbf{f}_{lag}(\mathbf{q}^{t}, \mathbf{q}^{t}, \mathbf{v}^{t})\right),$$
(19)

respectively, where we write

$$\mathbf{f}_{lag}(\mathbf{q}^{t}, \mathbf{q}^{t+h}, \mathbf{v}^{t+h}) = \mathbf{f}_{edcg}(\mathbf{q}^{t+h}, \mathbf{v}^{t+h}) + \mathbf{f}_{f}(\mathbf{q}^{t}, \mathbf{v}^{t+h}).$$
(20)

Here  $\mathbf{f}_{edcg}$  is the sum of elastic, damping, contact and external forces and  $\mathbf{f}_{f}$  is the friction force as before. Notice that in (19)

the entire net force  $f_{lag}$  is split into implicit and explicit parts, whereas the original proposal [1] for TR in IPC is to apply this splitting to non-contact forces only. In Section V-A3 we show how solving (19) will generate more stable results when compared to the original IPC implementation.

Note that  $\mathbf{f}_{lag}$  has a well-defined antiderivative with respect to  $\mathbf{v}^{t+h}$ , which can be minimized using common optimization tools. In this view, IPC effectively solves Eqs. (18) or (19) using the proposed log-barrier potential as a merit function, CCD aided line search and Hessian projection. Although this can be done iteratively with better estimates for the lagged friction force, this approach has limitations as demonstrated in Section V-A1.

#### V. RESULTS

All experiments were run on the AMD Ryzen Threadripper 1920X CPU with 12 cores, 24 threads at 3.7 GHz boost clock and 32 GB RAM. We used Blender 3.1 [45] for all 3D renderings. For Algorithm 1 we used the Intel MKL sparse LU solver to solve the square non-symmetric linear system on line 6. We used dual numbers for forward automatic differentiation [46] and a custom BiCGSTAB implementation for the inexact Newton Algorithm 2. In the following results Algorithms 1 and 2 are dubbed "Direct" and "Iterative", respectively, since the former uses a direct linear solver and the latter uses an iterative linear solver. We did not evaluate viscous friction effects here, so  $\mu_v = 0$  in all examples. For examples using the same dynamic and static friction coefficients  $\mu_d$  and  $\mu_s$ , we omit the subscript.

#### A. Friction accuracy

With the following examples we demonstrate two scenarios where lagged friction causes large deviations from an expected accurate and stable friction response.

1) Block slide: In this example we let a stiff elastic block slide down a 10 degree slope expecting it to stop for  $\mu = 0.177 > \tan(10^\circ)$  after sliding for a total of  $x_T = 0.769$  m for T = 15.38 seconds (see supplemental document for details on the experiment).

Figure 4 demonstrates that our method produces consistent stopping across a variety of time step sizes using BE and TR time integration. We compare against IPC [1], a state-of-the-art smoothed friction method using a lagged friction approach to show that it fails to establish consistent stopping with BE, and fails to stop with TR altogether after 50 seconds. We reproduce the lagged friction method in our simulator to demonstrate that TR can be used to generate reliable stopping if the equations of motion are correctly integrated as in Eq. (19). Our method produces a more accurate stopping distance using TR than IPC does using BE even after using multiple fixed point iterations.

2) *Bowl grasp:* Control over the friction coefficient is particularly important in grasping scenarios since grasped objects are often delicate. This means that friction forces involved in lifting are often close to the sliding threshold.

As shown in Figure 1, an upside down bowl is lifted using 3 soft pads to compare sticking stability of *lagged* friction given in Eq. (18) against a fully implicit method from Eq. (16). The



(a) Total displacement travelled by the block before coming to a stop. With BE and multiple fixed point iterations the lagged friction approach can produce a reasonable approximation of the stopping behavior, however, the approximation does not converge to the true solution. Our method produces an accurate estimate of the true behavior with BE, while TR produces a better estimate than lagged friction at h = 0.1 s and lower. With lagged friction, the TR method presented in Eq. (19) produces a more accurate result than IPC since contact is handled together with other implicit forces.



(b) Time taken by each block to come to a stop. With BE, our method produces an accurate approximation to the true stopping time for all but the largest time step. In contrast, lagged friction fails to converge to the true stopping time under multiple fixed point iterations. With TR, the approximation is not as accurate as with BE, although still closer to the true value than lagged friction using BE. With TR, lagged friction does not produce a reliable stopping time, even under refinement. Using IPC's TR implementation, the box does not stop after 50 s for any time step.

Fig. 4: Block slide. Comparisons of analytic stopping conditions of a sliding block to numerical results.

bowl is successfully picked up and stuck to the pads for a range of time steps when using the *implicit* method, however it slips for different time step values with *lagged* friction. The height of the bowl is plotted in Figure 5 for each method and time step combination.

3) Ball in a box: A rubber ball placed inside an elastic box has an initial spin of 4800 rotations per minute and an initial velocity set to  $\vec{v}_0 = (-0.923, -0.385, 0)$ . In Figure 6, we demonstrate how our fully coupled TR integrator produces more stable dynamic simulations compared to the decoupled TR as proposed by Li et al. [1]. This scenario is simulated with both methods for 800 frames at h = 0.01 s with comparable damping parameters as shown in Figure 6b. The TR implementation used by IPC blows up, whereas in our formulation the energy is eventually dissipated as expected.

#### B. Performance

In this section, we show how various combinations of volume preservation and frictional contact constraints can affect the performance of the simulation. In addition, we compare Algorithms 1 and 2 in performance and memory usage.

1) Tube cloth bend: The inexact Newton Algorithm 2 shines particularly in scenarios with numerous contacts, such as with tight-fitting garments, where the entire garment is in contact with a body. Here we simulate a simplified scenario of a tube cloth wrapped around a bending soft object resembling an elbow or knee as depicted in Figure 7. Table I shows the corresponding timing results, which indicate that inexact New-



Fig. 5: The height of the bowl in Figure 1 is plotted against frame number as time step h is varied. Here the bowl slips for the lagged friction method, whereas the fully implicit method maintains stable sticking for every time step, hence all plotted lines overlap.

ton performs much better than the damped Newton algorithm employing a direct solver. Furthermore, the performance gap becomes large when the number of elements is increased. The higher resolution example in Table I does not include the "Direct" solver because it is intractable at that resolution due to the scalability limitations of the direct solve. In this case exact Newton can use an iterative solver, however a full comparison of the trade-offs between different solvers is outside the scope of the paper since this type of solver would require an additional fixed tolerance, which would affect performance and convergence characteristics. Interestingly this data also indicates that larger friction coefficients cause a bigger bottleneck for the solve compared even to stiff volume change penalties (indicated by small  $\kappa_v$ ).

# Elements	$\mu$	Solver Type	$\kappa_v$	Time	Memory	Volume Loss
3K Tris 5K Tets	0.2	Iterative Direct	N/A N/A	4.81 22.1	1.84 GB 6.51 GB	0.962% 0.962%
	0.8	Iterative Direct	N/A N/A	12.7 25.6	1.91 GB 7.09 GB	0.962% 0.962%
8K Tris	0.2	Iterative Iterative	N/A 4.6e-5	8.59 12.2	830 MB 858 MB	2.55% 2.44e-4%
30K Tets	0.8	Iterative Iterative	N/A 4.6e-5	51.6 53.2	625 MB 918 MB	2.57% 2.44e-4%

TABLE I: *Tube cloth bend performance data.* Time is measured in seconds per frame and memory refers to total memory growth during the simulation as reported by Houdini's performance monitor. Simulations without volume change penalty have no  $\kappa_v$ .

2) Ball Squish: A hollow ball at various resolutions (1K, 55K and 160K elements) is pressed between two flat planes as demonstrated in Figure 8. As a result the ball experiences volume loss. To preserve some of the volume we simulate the compression with compression coefficients  $\kappa_v = 1$  (e.g. a ball filled with air) and  $\kappa_v = 0.01$  (e.g. a ball filled with water). In the latter case we expect significantly less volume loss, which is reflected in our experiments as shown in Table II. Furthermore, we note from Table IIb that scenarios with small



9

(a) The TR formulation proposed in IPC is unstable (top row) whereas our method produces a dissipative scenario (bottom row) even with a slightly weaker damping.



(b) The ball (top) and the box (bottom) are simulated individually with both methods and centroid heights are plotted to ensure that the oscillation amplitudes with TR as implemented by IPC do not exceed those generated by our implementation. This test ensures that the reason for blowup in IPC is not due to integration of elasticity equations alone, but indeed due to loose coupling between elastic and contact terms.

Fig. 6: *Ball in a box.* A spinning ball bounces inside an elastic box. Here  $\rho = 10000 \text{ kg/m}^3$ , E = 500 KPa, and  $\nu = 0.1$  for the box and  $\rho = 500 \text{ kg/m}^3$ , E = 50 KPa and  $\nu = 0.45$  for the ball, with friction coefficient set to  $\mu = 0.1$  between them. For IPC, damping ratio is set to 0.02, and for our model, the damping parameter is set to 0.1 Hz. In both cases Rayleigh damping is used.

 $\kappa_v$  favor the "Iterative" method. In Table IIc we see that this is true whether **J** is sparsely approximated ("Inexact") or not ("Exact"). In contrast, stiffer scenarios prefer the "Direct" method due to worse system conditioning.

#### C. Real world phenomena

In the following example we show how our simulator can reproduce deformations captured in the real world.

1) Tennis ball: Tennis ball dynamics is a prime example for all methods proposed in this paper. Volume preservation and higher-order integration enables accurate bounce behavior,



Fig. 7: *Tube cloth bend.* A cylindrical garment is wrapped around a bending capsule. Lower resolution example is shown at the top row and higher resolution at the bottom row. The first column shows the initial configuration, second column shows the end result with low friction ( $\mu = 0.2$ ), and the last column shows the end result with high friction ( $\mu = 0.8$ ).



Fig. 8: *Ball squish.* A hollow ball of various resolutions is squished between two parallel rigid plates. The same simulations are performed for varying Young's moduli E and compression coefficients  $\kappa_v$ . As expected, the squished ball occupies a larger volume for smaller  $\kappa_v$  (last row), and, possibly due to locking artifacts, occupies a smaller volume for larger E.

whereas accurate friction is needed for predicting the bounce trajectory. Here we demonstrate the need for the former.

First, we launch a tennis ball at a wall at 100 mph (44.7 m/s) to reproduce accurate slow motion deformation. Tennis balls are typically pressurized to approximately 1 atm above atmospheric pressure to produce a livelier bounce during play. In Figure 9 we show how deformation changes when the ball is pressurized and compare the result with live footage. The simulation contains 100K tetrahedra and 92K vertices. The ball is hollow with a stiff inner layer (E = 6 MPa,  $\nu = 0.4995$  and  $\rho = 934$  kg/m<sup>3</sup>) and a light outer felt material (E = 5.4 MPa,  $\nu = 0.3$  and  $\rho = 4.69$  kg/m<sup>3</sup>). The volume change penalty is applied to the interior of the ball. This 1000 frame simulation took 4.89 seconds per frame and a total of 11.33 GB in memory.

Next, a tennis ball is dropped from a 254 cm height to evaluate its bounce with and without pressurization. We show how pressurization and choice of integrator can drastically affect the height of the bounce in Figure 10.

# Elements	1K	55K	160K						
Direct Iterative	0.354 (0.228) 0.391 (0.081)	15.9 (1.69) 31.9 (0.719)	61.0 (4.26) 111 (1.92)						
Volume loss	33.6%	27.8%	27.7%						
(a) $E = 100$ KPa, $\kappa_v = 1$ .									
# Elen	nents   1H	K 55	K						
Direct	/e   1.08 (0 0.463 (	0.478) 89.1 ( 0.214) 42.6 (0	2.03) ).884)						
Volum	e loss   0.49	0.36	5%						
(b) $E = 100$ KPa, $\kappa_v = 0.01$ .									
$\kappa_v =$		1 (	0.01						
Direct Ex Direct In Iterative	xact J $0.832$ exact J $0.354$ 0.391	(0.393)0.871(0.228)1.08(0.081)0.463	(0.672) (0.478) 3 (0.214)						
(c) $E = 100$ KPa, for 1K elements.									
# Elements	1K	55K	160K						
Direct Iterative	0.312 (0.388) 0.593 (0.097)	8.28 (1.52) 28.3 (0.732)	25.5 (4.14) 56.2 (2.62)						
Volume loss	74.3%	74.6%	75.0%						

(d) E = 1000 KPa with no volume preservation constraint.

TABLE II: Ball squish timings, memory usage (in parentheses) and volume loss. A hollow ball is squished between two flat rigid plates at time step h = 0.001 s. Timings are given in seconds per frame and averaged over 804 frames. The memory usage measured in GB over the entire simulation sequence is shown in parentheses. Volume loss is computed for the change in volume between frames 1 and 804 as a percentage of initial volume inside the ball. Each table specifies the number of elements, Young's modulus E and compression coefficient  $\kappa_v$  where applicable. Here we compare how the performance characteristics of our "Direct" and "Iterative" methods change when problem stiffness and  $\kappa_v$  are varied for different mesh resolutions.



Fig. 9: *Highspeed tennis ball collision*. Simulated hollow tennis ball collides against a wall at 100mph. The deformation here closely resembles that of a real tennis ball as captured by Anderson [47].

2) *Tire wrinkling:* We model an inflated tire used by top fuel dragsters to show the folding phenomenon at the start of the race. The tires are deliberately inflated at a low pressure of 0.68 atm above atmospheric pressure, which allows them to better grip the asphalt for a better head start. As a result the soft tire tends to wrinkle as the wheels start to turn. This phenomenon allows for a larger contact patch between the tire



Fig. 10: *Tennis ball drop.* A pressurized and non-pressurized tennis ball dropped from a 254 cm height, is simulated using different integrators. As expected, the pressurized ball bounces higher than the corresponding non-pressurized ball. The highly damping BE and BDF2 produce a much lower bounce than SDIRK2 across a range of different time steps [21]. Higher-order integrators are defined in the supplemental document.

and the ground for better traction, which translates to a larger acceleration. In Figure 11 we demonstrate this phenomenon in simulation with a shell model tire inflated using our volume change penalty. The outer side of the tire is initially stuck to the ground and then dragged while maintaining consistent contact. Accurate simulation of stick-slip transitions of the tire is critical in determining its performance since traction transfers torque into forward acceleration of the vehicle, which ultimately determines the outcome of a race. The tire is simulated using 93K triangles and 46K vertices. The tire mesh is split into an outer stiffer part that is in contact with the ground and a softer inner part where the label is printed. The inner rim sets the Dirichlet boundary condition generating the rotation. Here  $\mu_d = 0.5$ ,  $\mu_s = 1.5$ , and  $\nu = 0.49$  everywhere, while E = 400 KN/m,  $\rho = 200$  kg/m<sup>2</sup> and bending stiffness set at 0.56 on the outer part, and E = 200 KN/m,  $\rho = 50$ kg/m<sup>2</sup>, and bending stiffness set at 0.01 on the inner part. The simulation ran with h = 0.00125 s for 20 seconds per frame on average using the damped Newton solver. This example shows how the stick-slip phenomenon can be modeled with distinct static and dynamic friction coefficients using the smoothed Stribeck model proposed in Section III-C.

#### VI. CONCLUSIONS AND LIMITATIONS

We have presented a fully implicit method for simulating hyperelastic objects subject to frictional contacts. This method generalizes the popular optimization framework for simulating hyperelastics with contact and lagged friction potentials. We demonstrate how higher-order integrators can be applied in our method as well as in IPC-style frameworks. Our method addresses the lack of friction convergence in lagged friction formulations by evaluating contacts, friction forces as well as tangential bases implicitly.

We have extended this approach to include static friction and physically-based volume preservation that can be used to simulate compressible as well as nearly incompressible media in a single solve.

*a) Local minima:* Solving for roots of nonlinear momentum equations allows one to resolve friction forces accurately,



Fig. 11: *Dragster tire wrinkle*. A soft tire spins against the ground with a large friction coefficient, causing the rubber to wrinkle. Left: simulation of the tire being rotated in place causing the wrinkle due to a high coefficient of friction with the ground. Right: a real dragster tire exhibiting the wrinkle phenomenon at the start of the race [48].

however, this comes with a trade-off. Optimization theory allows one to reliably find a descent direction even when the objective Hessian is indefinite via projection or filtering techniques. Although computing the descent direction for finding roots of nonlinear equations allows one to use the entire unfiltered Jacobian, global convergence can only be theoretically guaranteed when the Jacobian is bounded on the neighborhood of the initial point. For stiff systems, this assumption can become problematic, although the practical implications are unclear.

b) Hydrostatic equilibrium: Our volume change penalty model expects hydrostatic equilibrium, which may not always be a good approximation. For quickly deforming objects like in the tennis ball and tire examples, some details of the deformation may be missing due to this approximation. This is because the object deforms faster than the air moves inside the volume, creating a non-uniform pressure distribution. The comparison of our hydrostatic model to a fully dynamic fluid simulation remains as future work.

#### REFERENCES

- M. Li, Z. Ferguson, T. Schneider, T. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D. M. Kaufman, "Incremental potential contact: Intersection- and inversion-free large deformation dynamics," ACM Transactions on Graphics, vol. 39, no. 4, 2020.
- [2] G. De Saxcé and Z. Q. Feng, "The bipotential method: A constructive approach to design the complete contact law with friction and improved numerical algorithms," *Mathematical and Computer Modelling*, vol. 28, no. 4, pp. 225–245, Aug. 1998. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0895717798001198
- [3] J. Wojewoda, A. Stefański, M. Wiercigroch, and T. Kapitaniak, "Hysteretic effects of dry friction: modelling and experimental studies," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 366, no. 1866, pp. 747–765, Mar. 2008. [Online]. Available: https://royalsocietypublishing.org/doi/ 10.1098/rsta.2007.2125
- [4] E. Larionov, Y. Fan, and D. K. Pai, "Frictional contact on smooth elastic solids," ACM Trans. Graph., vol. 40, no. 2, Apr. 2021. [Online]. Available: https://doi.org/10.1145/3446663
- [5] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically deformable models," *SIGGRAPH Comput. Graph.*, vol. 21, no. 4, p. 205–214, aug 1987. [Online]. Available: https://doi.org/10.1145/37402. 37427
- [6] D. Baraff and A. Witkin, "Large steps in cloth simulation," in *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. ACM, 1998, pp. 43–54.

- [7] E. Sifakis and J. Barbic, "Fem simulation of 3d deformable solids: A practitioner's guide to theory, discretization and model reduction," in ACM SIGGRAPH 2012 Courses, ser. SIGGRAPH '12. New York, NY, USA: ACM, 2012, pp. 20:1–20:50. [Online]. Available: http://doi.acm.org/10.1145/2343483.2343501
- [8] M. Müller, B. Heidelberger, M. Hennix, and J. Ratcliff, "Position based dynamics," *Journal of Visual Communication and Image Representation*, vol. 18, no. 2, pp. 109–118, 2007.
- [9] M. Macklin, M. Müller, and N. Chentanez, "Xpbd: Position-based simulation of compliant constrained dynamics," in *Proceedings of the* 9th International Conference on Motion in Games, ser. MIG '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 49–54. [Online]. Available: https://doi.org/10.1145/2994258.2994272
- [10] J. Bender, M. Müller, and M. Macklin, "A survey on position based dynamics, 2017," in *EUROGRAPHICS 2017 Tutorials*. Eurographics Association, 2017.
- [11] S. Bouaziz, S. Martin, T. Liu, L. Kavan, and M. Pauly, "Projective dynamics: Fusing constraint projections for fast simulation," ACM *Trans. Graph.*, vol. 33, no. 4, Jul. 2014. [Online]. Available: https://doi.org/10.1145/2601097.2601116
- [12] D. M. Kaufman, S. Sueda, D. L. James, and D. K. Pai, "Staggered projections for frictional contact in multibody systems," ACM Transactions on Graphics (SIGGRAPH Asia 2008), vol. 27, no. 5, pp. 164:1–164:11, 2008.
- [13] D. Chen, D. I. W. Levin, W. Matusik, and D. M. Kaufman, "Dynamics-aware numerical coarsening for fabrication design," ACM Trans. Graph., vol. 36, no. 4, jul 2017. [Online]. Available: https://doi.org/10.1145/3072959.3073669
- [14] F. Bertails-Descoubes, F. Cadoux, G. Daviet, and V. Acary, "A nonsmooth newton solver for capturing exact coulomb friction in fiber assemblies," ACM Transactions on Graphics (TOG), vol. 30, no. 1, p. 6, 2011.
- [15] G. Daviet, F. Bertails-Descoubes, and L. Boissieux, "A hybrid iterative solver for robustly capturing coulomb friction in hair dynamics," ACM *Trans. Graph.*, vol. 30, no. 6, p. 1–12, dec 2011. [Online]. Available: https://doi.org/10.1145/2070781.2024173
- [16] 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.
- [17] M. Macklin, K. Erleben, M. Müller, N. Chentanez, S. Jeschke, and V. Makoviychuk, "Non-smooth newton methods for deformable multibody dynamics," *ACM Transactions on Graphics*, vol. 38, no. 5, pp. 1–20, Oct 2019. [Online]. Available: http://dx.doi.org/10.1145/3338695
- [18] K. Erleben, "Rigid body contact problems using proximal operators," in *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, ser. SCA '17. New York, NY, USA: ACM, 2017, pp. 13:1–13:12. [Online]. Available: http://doi.acm.org/10.1145/ 3099564.3099575
- [19] 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, vol. 39, no. 6, pp. 1–15, Nov. 2020. [Online]. Available: https://dl.acm.org/doi/10.1145/3414685.3417766
- [20] F. Löschner, A. Longva, S. Jeske, T. Kugelstadt, and J. Bender, "Higher-Order Time Integration for Deformable Solids," *Computer Graphics Forum*, vol. 39, no. 8, pp. 157–169, Dec. 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1111/cgf.14110
- [21] U. M. Ascher, E. Larionov, S. H. Sheen, and D. K. Pai, "Simulating deformable objects for computer animation: A numerical perspective," *Journal of Computational Dynamics*, vol. 9, no. 2, pp. 47–68, 2022.
- [22] G. E. Brown, M. Overby, Z. Forootaninia, and R. Narain, "Accurate dissipative forces in optimization integrators," ACM Trans. Graph., vol. 37, no. 6, dec 2018. [Online]. Available: https://doi.org/10.1145/ 3272127.3275011
- [23] J. Bonet, R. D. Wood, J. Mahaney, and P. Heywood, "Finite element analysis of air supported membrane structures," *Computer Methods in Applied Mechanics and Engineering*, vol. 190, no. 5, pp. 579–595, 2000. [Online]. Available: https://www.sciencedirect.com/ science/article/pii/S0045782599004284
- [24] M. Skouras, B. Thomaszewski, B. Bickel, and M. Gross, "Computational design of rubber balloons," *Computer Graphics Forum*, vol. 31, no. 2pt4, pp. 835–844, 2012. [Online]. Available: https://onlinelibrary.wiley. com/doi/abs/10.1111/j.1467-8659.2012.03064.x
- [25] M. Skouras, B. Thomaszewski, P. Kaufmann, A. Garg, B. Bickel, E. Grinspun, and M. Gross, "Designing inflatable structures," ACM

*Trans. Graph.*, vol. 33, no. 4, jul 2014. [Online]. Available: https://doi.org/10.1145/2601097.2601166

- [26] J. Panetta, F. Isvoranu, T. Chen, E. Siéfert, B. Roman, and M. Pauly, "Computational inverse design of surface-based inflatables," ACM Trans. Graph., vol. 40, no. 4, Jul. 2021. [Online]. Available: https://doi.org/10.1145/3450626.3459789
- [27] B. Smith, F. D. Goes, and T. Kim, "Stable neo-hookean flesh simulation," ACM Trans. Graph., vol. 37, no. 2, mar 2018. [Online]. Available: https://doi.org/10.1145/3180491
- [28] S. H. Sheen, E. Larionov, and D. K. Pai, "Volume preserving simulation of soft tissue with skin," *Proceedings of the ACM on computer graphics* and interactive techniques, vol. 4, no. 3, pp. 1–23, 2021.
- [29] P. G. Ciarlet, *Lectures on three-dimensional elasticity*. Berlin: published for the Tata Institute of Fundamental Research [by] Springer, 1983, vol. 71.;71;.
- [30] H. Wang, J. F. O'Brien, and R. Ramamoorthi, "Data-driven elastic models for cloth: modeling and measurement," in ACM SIGGRAPH 2011 papers on - SIGGRAPH '11. Vancouver, British Columbia, Canada: ACM Press, 2011, p. 1. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1964921.1964966
- [31] R. Kikuuwe, N. Takesue, A. Sano, H. Mochiyama, and H. Fujimoto, "Fixed-step friction simulation: from classical Coulomb model to modern continuous models," in 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems. Edmonton, Alta., Canada: IEEE, 2005, pp. 1009–1016. [Online]. Available: http: //ieeexplore.ieee.org/document/1545579/
- [32] J. Awrejcewicz, D. Grzelczyk, and Y. Pyryev, "404. a novel dry friction modeling and its impact on differential equations computation and lyapunov exponents estimation." *Journal of Vibroengineering*, vol. 10, no. 4, 2008.
- [33] B. Armstrong-Helouvry, "Stick slip and control in low-speed motion," *IEEE Transactions on Automatic Control*, vol. 38, no. 10, pp. 1483– 1496, 1993.
- [34] F. Mandl, Statistical physics. London; New York;: Wiley, 1971.
- [35] U. M. Ascher and L. R. Petzold, Computer methods for ordinary differential equations and differential-algebraic equations. Siam, 1998, vol. 61.
- [36] N. M. Newmark, "A method of computation for structural dynamics," *Journal of the Engineering Mechanics Division*, vol. 85, no. 3, pp. 67–94, 1959. [Online]. Available: https://ascelibrary.org/doi/abs/10. 1061/JMCEA3.0000098
- [37] T. J. R. Hughes and G. E. P. A. Knovel, Academic, *The finite element method: linear static and dynamic finite element analysis.* Mineola, NY: Dover Publications, 2000.
- [38] M. Géradin, D. Rixen, M. Géradin, and D. Rixen, *Théorie des vibrations: application à la dynamique des structures*. Masson Paris, 1993, vol. 2.
- [39] C. Kane, J. E. Marsden, M. Ortiz, and M. West, "Variational integrators and the newmark algorithm for conservative and dissipative mechanical systems," *International Journal for Numerical Methods in Engineering*, vol. 49, no. 10, pp. 1295–1325, 2000.
- [40] D. M. Kaufman, R. Tamstorf, B. Smith, J.-M. Aubry, and E. Grinspun, "Adaptive nonlinearity for collisions in complex rod assemblies," ACM Transactions on Graphics (TOG), vol. 33, no. 4, p. 123, 2014.
- [41] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York, NY: Springer New York, 2006.
- [42] H. A. van der Vorst, "Bi-cgstab: A fast and smoothly converging variant of bi-cg for the solution of nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.*, vol. 13, no. 2, p. 631–644, Mar. 1992. [Online]. Available: https://doi.org/10.1137/0913035
- [43] S. C. Eisenstat and H. F. Walker, "Choosing the forcing terms in an inexact newton method," *SIAM Journal on Scientific Computing*, vol. 17, no. 1, pp. 16–32, 1996. [Online]. Available: https: //doi.org/10.1137/0917003
- [44] A. Griewank and A. Walther, *Evaluating Derivatives*, 2nd ed. Society for Industrial and Applied Mathematics, 2008. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9780898717761
- [45] B. Online Community, Blender a 3D modelling and rendering package, Blender Foundation, Stichting Blender Foundation, Amsterdam, 2021. [Online]. Available: http://www.blender.org
- [46] E. Larionov, "autodiff," 2022. [Online]. Available: https://github.com/ elrnv/autodiff
- [47] M. Anderson, "Tennis ball hitting the wall at 100 mph," Feb. 2018. [Online]. Available: https://www.youtube.com/watch?v=FC8Tpi3U0H0
- [48] Goodyear, "Nhra drag race tire wrinkle in slow motion top fuel and funny car," 2020. [Online]. Available: https://www.youtube.com/watch? v=gXp2QgY10B8