

Adaptive Nonlinearity for Collisions in Complex Rod Assemblies

Danny M. Kaufman
Adobe & Columbia University

Rasmus Tamstorf
Walt Disney Animation Studios

Breannan Smith
Columbia University

Jean-Marie Aubry
Weta Digital

Eitan Grinspun
Columbia University

Abstract

We develop an algorithm for the efficient and stable simulation of large-scale elastic rod assemblies. We observe that the time-integration step is severely restricted by a strong nonlinearity in the response of stretching modes to transversal impact, the degree of this nonlinearity varying greatly with the shape of the rod. Building on these observations, we propose a collision response algorithm that adapts its degree of nonlinearity. We illustrate the advantages of the resulting algorithm by analyzing simulations involving elastic rod assemblies of varying density and scale, with up to 1.7 million individual contacts per time step.

CR Categories: I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: simulation, collisions, rods, hair, time-integration

Links:  DL  PDF  WEB  VIDEO

1 Introduction

The physical behavior of a full head of hair is dominated by rod-rod interactions in a highly dense and complex assembly generally composed in excess of one hundred thousand individual rods [Robbins 2012]. To date simulations have scaled to roughly one-tenth of these numbers [Selle et al. 2008; McAdams et al. 2009; Daviet et al. 2011; Iben et al. 2013] and clearly the remaining challenge is not in simply simulating sufficient rods. Capturing the emergent behaviors generated by rod-rod interactions such as locking and tangling requires modeling contacts at scale. In turn these interactions are challenging to simulate due to material stiffness and a time-varying contact graph that couples most of these constitutive rods in each computational step. Motivated by these observations we focus on efficient time-integration of complex rod assemblies, and in particular the critical role that *nonlinear response* plays in the collision resolution of thin rods.

Consider the collision of a discretized three-node rod as it hits a wedge. The standard approach to prevent intersections is to apply a position-based correction that enforces zero-displacement along a collision normal [Baraff 1989; Mirtich and Canny 1995; Bridson et al. 2002]. However, by resolving collisions in this way we ignore changes in internal energy, effectively applying an instantaneous impulse response [Moreau 1988].

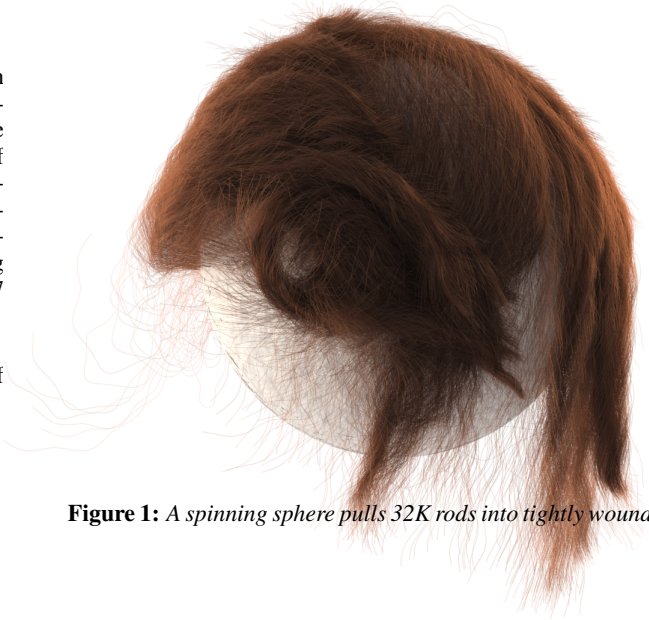
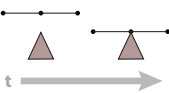


Figure 1: A spinning sphere pulls 32K rods into tightly wound locks.

An impulsive response imposes corrections restricted to just the stencil of the collision normal—in our example just the colliding vertex—so that the correction produces an unnecessarily large strain on material elements. Large, highly localized, non-physical deformations are generated, with their attendant artifacts and instabilities at both large time steps and high speeds.

In turn, notice that stretching is dominant in thin materials. Here stretching stiffness is generally at least *four orders* of magnitude greater than bending stiffness [Landau and Lifshitz 1986], while in many models, e.g., when used to approximate inextensibility, it can be even higher. While we expect such materials to potentially bend a great deal, we clearly do not wish to exercise their stretching modes unintentionally.

Impulsive models thus resolve collisions but ignore physical response at their own peril. To build this missing “physical awareness” into collision response, first-order models have been developed. *Linearly compliant* collision response methods [Baraff and Witkin 1998; Duriez et al. 2004; Duriez et al. 2006; Kaufman et al. 2008; Otaduy et al. 2009; Allard et al. 2010; Zheng and James 2011; Bertails-Descoubes et al. 2011; Daviet et al. 2011] apply *compliance* (inverse stiffness) to inexpensively communicate collision-induced strain across force stencils, effectively filtering collision response to obtain smoother post-collision configurations. Linearly compliant methods thus can balance between the difficult tradeoffs of stability and computational efficiency.

Elastic rods are modeled by a range of spatial discretizations that apply either reduced or maximal coordinate formulations [Ward et al. 2007; Hadap et al. 2007]. In this work we focus on resolving collisions in maximal coordinate rod models. To do so, however, we face a severe limitation: *first-order models are not sufficient to capture physical response for thin elastic materials in maximal coordinates*. In particular, the collision response of thin materials such as rods is strongly nonlinear so that first-order collision modeling can generate the same problematic localized collision response as an impulsive model.

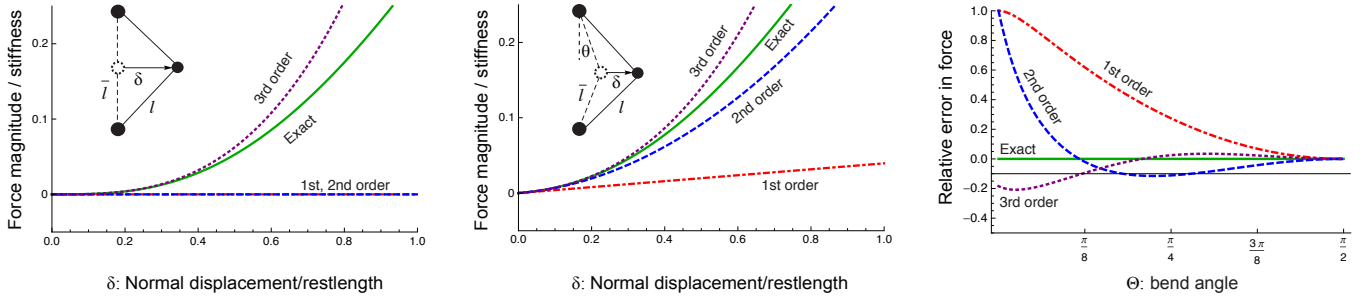


Figure 2: We plot the magnitude of rod/mass-spring stretching forces modeled by respectively first-, second- and third-order force approximations for expansions about a straight (left) and a bent (middle) configuration and compare against the ground truth stretching force evaluation. In both cases we observe that first-order modeling with respect to normal displacement δ underestimates the force while, in the case of straight configurations (left) any expansion less than third-order entirely ignores normal displacement. On the other hand, by plotting the relative error of these approximations with respect to the ground truth force evaluation at a fixed $\delta = \frac{1}{2}$, we see that as bending increases lower-order models give a correspondingly better approximation (right).

Consider the simple case of two adjacent elements, with little or no initial bend, colliding with an edge as in Fig. 3. In these configurations, when we apply a first-order, linearly compliant collision method, the resulting initial collision response varies little (to not at all) from the standard physics-oblivious method:

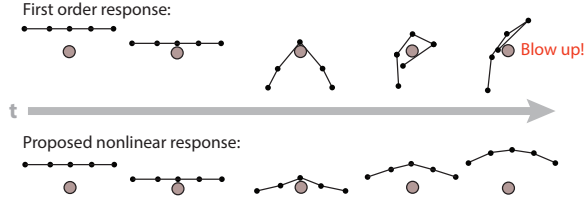
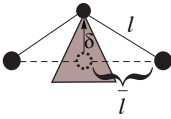


Figure 3: Left to right in time. Top: a collision resolved by first-order-modeled response remains physics-oblivious; the resulting correction generates a large, localized, non-physical deformation leading to instability. Bottom: adaptive nonlinear response obtains a stable, global response.

To understand why this occurs reconsider the collision of our three-node rod in a straight, unstretched configuration with edge lengths \bar{l} and l at undeformed and current configurations respectively. We now impose a normal displacement of δ due to a collision. Our first observation is that such displacements in the normal direction are *entirely* unaccounted for by expansions of stretching force up to second-order in this configuration. To see this, note that here Green's strain is quadratic in normal displacement $\epsilon_g = l^2/\bar{l}^2 - 1 = \delta^2/\bar{l}^2$ so that stretching energy with respect to normal displacement is correspondingly the quartic monomial $E_s(\epsilon_g) = \frac{1}{2}k\epsilon_g^2 = \frac{1}{2}k\delta^4/\bar{l}^4$.



Alternately consider the stretching energy used in rod and mass-spring models [Spillmann and Teschner 2007; Bergou et al. 2008; Bergou et al. 2010]. Although the stretching energy $E_s(l/\bar{l} - 1) = \frac{1}{2}k(l/\bar{l} - 1)^2 = \frac{1}{2}k(\sqrt{\delta^2 + \bar{l}^2}/\bar{l} - 1)^2$ is no longer a quartic monomial in δ , at an unstretched configuration the force remains oblivious to normal displacements up to second-order (see Figure 2, left). Nor is this restricted to collision forces; frictional forces also lead to the same stretching energy scaling. Moreover, this situation is not isolated to straight configurations. Even as the rod bends, corresponding first- and second-order force approximations continue to significantly underestimate the collision response of stretching modes (see Fig. 2, middle) and thus permit too much strain.

Summary Collision response methods using lower-order approximations are effectively physics-oblivious for small-bend collisions. On the other hand, as bending increases, first-order modeling gives a correspondingly better approximation (see Fig. 2, right).

We conclude that a full-blown nonlinear solution is overkill, while restricting ourselves to the linear model is clearly insufficient. Following these observations we seek a simple algorithm that can flexibly balance between these two extremes to enable accurate and stable progress of simulations by adaptively incorporating sufficient nonlinearity into collision response. Such an algorithm should cost no more than standard linearly compliant methods when a first-order model is sufficient, while it should introduce a minimum of additional overhead to obtain just enough nonlinearity when higher-order modeling is required. In the following sections we first develop the necessary background and then show how such an algorithm can be built by small augmentation of existing methods. We then analyze the resulting algorithm's behavior and scaling over a range of benchmark examples.

Contributions We have exposed and analyzed the strongly nonlinear behavior of thin body collision response. We develop a simple time stepping algorithm that adapts the degree of nonlinearity in the impact solve (ADONIS). By adapting on the number of constrained Newton iterations rather than on the time step size, we enable stable progress at time steps several orders of magnitude larger than previously possible. The final resulting algorithm is a simple and direct extension of Jourdan et al's [1998] "Gauss-Seidel-like" frictional contact algorithm enabling the reuse of robust existing codes for implementation.

2 Preliminaries

Discrete dynamics We have a number of independent and thus separable computational domains indexed by $x \in \{1, \dots, \ell\}$. Configuration for each domain is then \mathbf{q}_x , velocity $\dot{\mathbf{q}}_x$, masses M_x , and internal energies $V_x(\mathbf{q}_x)$. We concatenate and sum in the usual fashion, so that $\mathbf{q} = (\mathbf{q}_1^T, \dots, \mathbf{q}_\ell^T)^T \in \mathbb{R}^n$, $M = \text{diag}(M_1, \dots, M_\ell)$, and $V(\mathbf{q}) = \sum_x V_x(\mathbf{q}_x)$. We then let H denote the Hessian of V . Let δ denote the displacement over each discrete time-interval with $\delta = \mathbf{q}^{t+1} - \mathbf{q}^t$.

Implicit Euler is given by the discrete Euler-Lagrange (DEL) update equation to solve at each time step

$$M\delta = hM\dot{\mathbf{q}}^t - h^2\nabla V(\mathbf{q}^t + \delta) \quad (\text{DEL})$$

with $\dot{\mathbf{q}}^{t+1} = \frac{1}{h}\delta$.



Figure 4: A character's hair is simulated using the ADONIS algorithm.

For didactic purposes we discuss implicit Euler; analysis for other implicit time-integration methods follows similarly.

Linearly implicit integration Linearizing an implicit DEL generates a *linearly implicit* time-integration method [Hairer and Wanner 2004]. Linearizing implicit Euler (see e.g., Baraff and Witkin [1998]) yields

$$[M + h^2 H(q^t)] \delta = h M \dot{q}^t - h^2 \nabla V(q^t). \quad (\text{LIE})$$

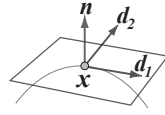
Inverting the left-hand side matrix in (LIE) gives the discrete compliance matrix

$$C(q) = [M + h^2 H(q)]^{-1},$$

which maps given forces to induced displacements.

3 Contact Model

We focus on *discretely sampled* contact models where for each identified contact k between material point pairs \mathbf{x}_i and $\mathbf{x}_j \in \mathbb{R}^3$ we construct the contact normal, \mathbf{n}_k and two orthogonal vectors of the tangent-plane, \mathbf{d}_{1k} and \mathbf{d}_{2k} .



We then set G_k to the contact basis given by the lift of the local orthonormal contact frame $\Phi_k = (\mathbf{n}, \mathbf{d}_1, \mathbf{d}_2)_k \in \mathbb{R}^{3 \times 3}$ to generalized coordinates via the relative deformation gradient $\Gamma_k(q) = \nabla_q(\mathbf{x}_i(q) - \mathbf{x}_j(q))$ so that

$$G_k(q) = (\mathbf{n}, \mathbf{d}_1, \mathbf{d}_2)_k = \Gamma_k(q)^T \Phi_k \in \mathbb{R}^{n \times 3}.$$

Note that $G_k^T \dot{q}$ then gives the relative velocity magnitudes along basis directions at contact k .

Contact constraints For m detected contacts we concatenate associated generalized contact normals in the basis $N = (\mathbf{n}_1, \dots, \mathbf{n}_m) \in \mathbb{R}^{n \times m}$. We then enforce zero- or scripted-displacements $\mathbf{s} = (s_1, \dots, s_m)$ along them by requiring

$$N^T \delta \geq \mathbf{s}.$$

Setting $\alpha = \{\alpha_1, \dots, \alpha_m\}$ to the associated vector of discrete contact force magnitudes, the global discrete contact force $\mathbf{f}_C = h N \alpha$ is then added to the right-hand side (*rhs*) of the DEL.

We then impose discrete Signorini-Fischer [Kikuchi and Oden 1988] conditions between contact constraints and forces as

$$0 \leq \alpha \perp N^T \delta - \mathbf{s} \geq 0. \quad (\text{SFC})$$

Here and henceforth $\mathbf{x} \perp \mathbf{y}$ denotes the *complementarity condition* $x_k y_k = 0, \forall k \in \{1, \dots, m\}$.

Friction We similarly concatenate the generalized tangent bases as $D_k = (\mathbf{d}_1, \mathbf{d}_2)_k \in \mathbb{R}^{n \times 2}$ and $D = (D_1, \dots, D_m)$ and form the corresponding vector of discrete frictional force coefficients as $\beta = (\beta_1, \dots, \beta_m)^T = ((\beta_1, \beta_2)_1, \dots, (\beta_1, \beta_2)_m)^T$ where each $\beta_k = (\beta_1, \beta_2)_k \in \mathbb{R}^2$ gives the frictional response coefficients at contact k along \mathbf{d}_{1k} and \mathbf{d}_{2k} respectively. The global, discrete friction force is then $\mathbf{f}_F = h D \beta$. An implicit discretization of the maximal dissipation principle [Goyal et al. 1991] requires that discrete friction satisfies the minimization

$$\min_{\beta} \{ \beta^T D^T \dot{q}^{t+1} : \|\beta_k\| \leq \mu_k \alpha_k, \forall k \}, \quad (\text{DMD})$$

where μ_k is the coefficient of friction for contact k .

Contact force coupling We organize contact force unknowns per contact in

$$\lambda_k = (\alpha_k, \beta_k)^T \in \mathbb{R}^3.$$

Concatenating the total response in $\lambda \in \mathbb{R}^{3m}$ and the response basis in $G = (G_1, \dots, G_m) \in \mathbb{R}^{n \times 3m}$ we obtain the total discrete frictional contact force $G \lambda$.

With the above discretizations the combined satisfaction of (SFC) and (DMD) per contact k is equivalently determined by seeking a λ_k satisfying the inclusion

$$\lambda_k \in \mathcal{R}_k(\delta) \iff \begin{cases} \min_{\beta_k} \{ \frac{1}{h} \beta_k^T D_k^T \delta : \|\beta_k\| \leq \mu_k \alpha_k \}, \\ 0 \leq \alpha_k \perp \mathbf{n}_k^T \delta - s_k \geq 0. \end{cases}$$

Here we use \mathcal{R}_k to denote the solution set to the single-point frictional contact problem.

Finally, the stencils of contacts' generalized friction and normal bases will overlap. Thus we consider the globally-coupled, joint satisfaction of (SFC) and (DMD) over all contacts. This is denoted by a total response λ satisfying the inclusion

$$\lambda \in \mathcal{R}(\delta) \iff \lambda_k \in \mathcal{R}_k(\delta) \quad \forall k \in \{0, \dots, m\},$$

where \mathcal{R} is the set of all *response* solutions globally satisfying frictional contact conditions [Stewart 2011].

Time-integration Adding frictional contact forces and the global contact conditions discussed above to (DEL) then forms the implicit frictional contact integrator

$$M \delta = h M \dot{q}^t - h^2 \nabla V(q^t + \delta) + G \lambda, \quad \lambda \in \mathcal{R}(\delta). \quad (\text{CDEL})$$

Contact groups In the absence of contacts each separable domain comprises an independent system that we integrate individually. In the presence of contacts, at each time step we decompose the contact graph formed by detected contacts into its connected components. Each of these maximally connected subgraphs similarly forms an independent system, a *contact group*, that we can integrate separately, independently, and (when computational resources allow) in parallel.

4 Sufficient Nonlinearity

In analogy to the linearly implicit unconstrained case in (LIE), linearizing the update equation in (CDEL) at start of step generates the standard *linearly compliant* frictional contact method

$$\lambda \in \mathcal{R}(\delta_l + C(q^t)G\lambda), \quad (1)$$

where δ_l is the solution to (LIE) and λ is the unknown linear estimate of frictional contact force coefficients.

This system encodes a linearly compliant model of *constraint anticipation* [Otdady et al. 2009] in that it models global changes in position due to constraint and friction forces up to first order. In many cases this model thus provides a physics-aware, global collision response. However, for thin materials straight collisions are largely hidden from this linear analysis, resulting in a response that is highly localized and instantaneously identical to the standard zeroth-order response discussed above.

4.1 An Adaptive Algorithm

If we instead linearize (CDEL) about the unconstrained solution of (DEL), δ_u , we obtain the linearly compliant update method

$$\lambda \in \mathcal{R}(\delta_u + C(q^t + \delta_u)G\lambda). \quad (2)$$

In turn we can then identify this linearly compliant model in (2) as exactly a single Newton-Raphson iteration:

Lemma The frictional contact, linearly implicit integrator (2) is exactly the first Newton iteration towards the solution of the nonlinear frictional contact integrator (CDEL) when initialized by the unconstrained solution of implicit Euler (DEL) as the starting seed.

(See Appendix B for proof.)

Additional iterations, appropriately computed, applied, and adaptively terminated, then provide a natural path towards sufficient nonlinearity.

A simple algorithm now takes shape: we start by stepping our system forward by a full solve of the nonlinear unconstrained time step (DEL), providing a stable guess as a preconditioner. We then detect collisions and solve linear compliance in (2). If we have sufficiently approximated collision response we stop. Otherwise, we proceed with further iterations to obtain an improved estimate of compliance for frictional contact.

We now elaborate on the requisite details to flesh out this algorithm. We first show how to construct these additional iterations (§5); we then describe a simple geometric termination criterion (§5.1); and then finally show how to construct an algorithm that enables our opportunistic, just enough nonlinearity to work within the popular framework of iterative contact algorithms (§6).



Figure 6: Collisions and tangling, as in this combing stress test, exercise the strongly nonlinear collision response of rods.

5 Iterating Towards Compliant Response

Solving the nonlinear (CDEL) system via Newton iteration amounts to sequentially solving a sequence of approximations on the combined primal (displacement) and dual (frictional contact force) unknowns δ and λ . Each iteration i linearizes (CDEL) about (δ^i, λ^i) and solves the resulting frictional contact problem to find the next update $(d\delta^i, d\lambda^i)$ with $\delta^{i+1} = \delta^i + d\delta^i$.

At each iteration, i , we then compute the tangent compliance

$$C^i = C(q^t + \delta^i)$$

to obtain the linear-estimated displacement from (CDEL)

$$d\delta^i = C^i \left[hM\dot{q}^t - h^2 \nabla V(q^t + \delta^i) + Gd\lambda^i + G\lambda^i - M\delta^i \right].$$

Noting from (CDEL) that the residual at iteration i is given by

$$r^i = hM\dot{q}^t - h^2 \nabla V(q^t + \delta^i) + G\lambda^i - M\delta^i$$

and that the dual unknown is the contact force update quantity $d\lambda^i$ we equivalently have displacement given by $d\delta^i = C^i r^i + C^i G d\lambda^i$. Observe then that the first quantity on the *rhs* simply gives the default full Newton step. Roughly, this describes the local, first-order approximation of displacement due to imbalance between response and force approximations in the previous iterate. The second term is simply the local, first-order estimate of displacement due to changes in frictional contact forces required after the last iteration's state update. Thus each iterate locally improves its compliance estimate, updates contact forces, and obtains an improved response estimate.

The unconstrained displacement estimate is then

$$d\delta_u^i = C^i \left[hM\dot{q}^t - h^2 \nabla V(q^t + \delta^i) - M\delta^i \right],$$

so that it remains to solve for unknown λ^{i+1} in

$$\lambda^{i+1} \in \mathcal{R}(\delta^i + d\delta_u^i + C^i G \lambda^{i+1}) \quad (3)$$

as we will describe in §6 below.



Figure 5: We simulate 16K rods affixed to a sphere scripted through a series of rapid rotations.

5.1 Termination

Motivated by our observations in §1 we focus on the *geometric* objective of reducing non-physical stretch. For close to inextensible materials such as hair, we thus judge sufficient nonlinearity by considering axial strain. At iteration i we compute a geometric ∞ -norm measure of the maximum stretch-factor over all edges $e_{x,j}$ belonging to rods indexed by x

$$\text{sf}^i = \max_{x,j} |e_{x,j}^i| / |\bar{e}_{x,j}| - 1,$$

where \bar{e} indicates edge rest length. We adaptively stop each time step solve when a sufficient reduction in stretch indicates that collision response has been communicated across the domain. Note that standard practice [Bridson et al. 2002] suggests $\text{sf}^i < 10^{-1}$.

6 Solution of \mathcal{R}

At each iteration we require the solution of the implicit contact system in (3). A direct solve of this nonlinear optimization is reserved for future work; we found a direct solve to be intractable for practical-size systems given current available solvers. Instead, we focus on how to solve these systems within the popular framework of iterative contact-collision methods, specifically Gauss-Seidel methods.

We extend the standard contact iteration approach by observing that we can iterate through the full contacting system (3) in multiple Gauss-Seidel passes, while updating the discrete compliance matrix until the stretch tolerance is sufficiently small.

For convenience we designate \bar{G}_k and $\bar{\lambda}_k$ as respectively the complement matrix and vector formed by zeroing out the three columns in G and the three entries in λ corresponding to the frame basis vectors of contact k . We then overload addition for complement pairs to ensure consistent block-vector addition so that $\lambda = \lambda_k + \bar{\lambda}_k$.

To begin we then first note that displacement at the next Newton iterate is equivalently given by

$$\delta^{i+1} = \delta_u^{i+1} + C^i [G_k \lambda_k^{i+1} + \bar{G}_k \bar{\lambda}_k^{i+1}],$$

where

$$\delta_u^{i+1} = C^i [h M \dot{q}^t - h^2 \nabla V(q^t + \delta^i) + h^2 H(q^t + \delta^i) \delta^i]$$

can be precomputed at the beginning of each Newton iteration.

To solve (3) at *each Newton iteration* i we then iterate Gauss-Seidel fashion *per contact* k , computing directly for λ_k^{i+1} (and thus implicitly for δ_k^{i+1}) by solving

$$\lambda_k^{i+1} \in \mathcal{R}_k(\delta_u^{i+1} + C^i [G_k \lambda_k^{i+1} + \bar{G}_k \bar{\lambda}_k^{i+1}]) \quad (4)$$

and then setting $\lambda^{i+1} \leftarrow \bar{\lambda}_k^{i+1} + \lambda_k^{i+1}$.

Each such iteration subproblem can then be most usefully viewed as simply solving a single point frictional contact problem where $\delta_u^{i+1} + \bar{G}_k \bar{\lambda}_k^{i+1}$ is the predicted displacement and λ_k^{i+1} is the contact's unknown response force. As such (4) can then be solved or approximated by a plethora of available numerical methods [Alart and Curnier 1991; Stewart 2001; Duriez et al. 2006; Bonnefon and Daviet 2011; Bertails-Descoubes et al. 2011; Stewart 2011] customized for the resolution of systems subject to a single contact. In our implementation, discussed below in §7, we employ the robust hybrid solver of Daviet et al. [2011], for which code has been publicly released.

Our approach thus applies a simple extension of the standard contact iteration [Jean and Moreau 1992] by observing that we can iterate through the full contacting system in multiple Gauss-Seidel passes while updating the discrete compliance matrix until the stretch tolerance is sufficiently small.

After each Gauss-Seidel solve we then update

$$\delta^{i+1} \leftarrow \delta_u^{i+1} + C^i G \lambda^{i+1}.$$

The stretch measure (§5.1) is tested and, if too large, we update the nonlinear terms as described below and proceed to solve the next Newton iteration by multiple Gauss-Seidel passes. The resulting ADONIS method is summarized in Algorithm 1 below. Additional details are included in the supplemental material.

Finally, it follows directly from §4 that the Gauss-Seidel solve of the first Newton iteration retains our guarantee that it corresponds to the exact solution of the linearly compliant contact problem.

A localized modified-Newton strategy Consider that each of the above Newton iterations is applied with the specified goal of building a better distribution of contact responses to reduce stretch. At the start of each such Newton iteration we thus restrict our update of Hessian terms to rods within each contacting system whose stretch is larger than the requested tolerance. All other Hessians are left unchanged from the prior iteration. This *localized* modified-Newton step allows us to focus computation on improving nonlinear terms of domains that are poorly modeled while the following Gauss-Seidel iterations then redistribute forces and displacements over the whole contacting system via the improved model.

7 Implementation

As discussed above we have engineered our implementation so that we run, by construction, exactly a single Gauss-Seidel-based linearly compliant solve for the first iteration of our method. The overhead beyond a linearly compliant solve for all time steps where a single iteration is sufficient to satisfy tolerances is one query per edge to compute the stretch-factor. As a starting point, we adapt an implementation of Daviet et al.'s [2011] contact solver applied to Discrete Elastic Rods [Bergou et al. 2008; Bergou et al. 2010]. We

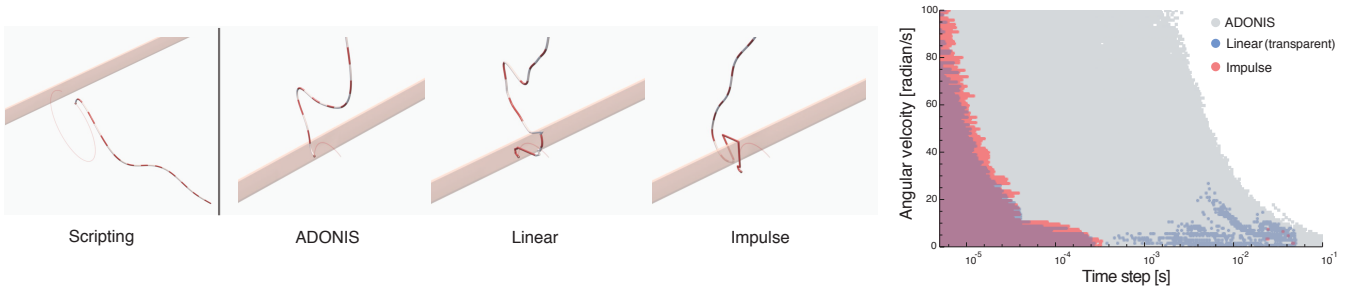


Figure 7: Whip-it stability test. Left: To test the relative stability of response methods we rotate a scripted handle connected to a rod so that it repeatedly whips the rod against the edge of a thin wall obstacle. At a time step size of 3 ms this results in, from left to right, a smoothly varying collision response from ADONIS, while the linearly compliant and impulse methods both obtain large, localized, non-physical deformations. Right: To understand the overall stability behavior of these algorithms we plot their respective stability regions as we vary time step (x-axis) in log-scale from 5 μ s to 100 ms and rotational whipping speeds (y-axis). For each successful simulation we plot a corresponding grey marker for ADONIS; transparent blue for linearly compliant response; and red for impulse response. Here we observe a generally two orders of magnitude gain in maximum stable time step size for ADONIS.

Algorithm 1 ADONIS(q^t, \dot{q}^t, h)

```

1:  $\delta \leftarrow \text{solve DEL}(q^t, \dot{q}^t, h)$ 
2:  $\mathbb{K} \leftarrow \text{collision detection}(\delta, q^t)$ 
3: while  $\text{sf} > \text{stretch.tol}$  do
4:    $C \leftarrow [M + h^2 H(q^t + \delta)]^{-1}$ 
5:    $\delta_u \leftarrow C[hM\dot{q}^t - h^2 \nabla V(q^t + \delta) + h^2 H(q^t + \delta)\delta]$ 
6:    $\text{gs.itr} \leftarrow 0$ 
7:   while  $\text{contact.err} > \text{contact.tol} \ \& \ \text{gs.itr} < \text{gs.max}$  do
8:     for  $k$  in  $\mathbb{K}$  do
9:        $\lambda_k \leftarrow \text{solve} : \lambda_k \in \mathcal{R}_k(\delta_u + C[G_k \lambda_k + \bar{G}_k \bar{\lambda}_k])$ .
10:       $\lambda \leftarrow \bar{\lambda}_k + \lambda_k$ .
11:       $\text{gs.itr} \leftarrow \text{gs.itr} + 1$ 
12:     end for
13:   end while
14:    $\delta \leftarrow \delta_u + CG\lambda$ .
15: end while
16:  $\dot{q}^{t+1} \leftarrow \frac{1}{h}\delta$ 
17:  $q^{t+1} \leftarrow q^t + \delta$ 
18: return ( $q^{t+1}, \dot{q}^{t+1}$ )

```

do not explicitly compute the inverses required for discrete compliance matrices in our implementation; rather we factor terms block-wise and back solve as updates and solves require.

For collision detection we employ a bounding volume hierarchy for broad-phase pruning combined with continuous collision detection (CCD) for narrow-phase collision determination [Provot 1997; Bridson et al. 2002]. We also utilize a second collision detection system employing a spatial hash for broad-phase with proximity queries for narrow-phase. Throughout our implementation we use the former CCD based algorithm for all rod-mesh collision checks while we experiment with both the CCD and proximity approach for rod-rod collision detection as detailed per-example below.

A number of steps in our algorithm are parallelized. Unconstrained predictor solutions (DEL) are computed trivially in parallel per rod. Broad-phase updates, collision detection queries, and contact point processing are performed in parallel to determine contact constraints. The resulting contact graph is then subdivided into its contact groups with each group solved as an independent contact problem. Finally, each Gauss-Seidel loop is likewise parallelized by graph-color partitioning.

Importantly, Gauss-Seidel iteration applied to frictional contact problems converges slowly or not at all as problem size grows. In

order to balance the tradeoff between cost and accuracy we have set a hard upper limit on the number of Gauss-Seidel iterations in our code. We specify this limit, `gs_max`, per problem below with an ∞ -norm convergence tolerance set to `contact_tol` = 10^{-6} .

8 Evaluation

Road map To understand the performance and behavior of our proposed algorithm we run a range of benchmark examples, described in detail below, over a variety of collision scenarios. As we also wish to better understand how the behavior of ADONIS compares with that of both standard zeroth-order, impulse response (obtained by setting $C^i = M^{-1}$) and linearly compliant response we consider these methods below as well.

We first consider and quantify stability gains obtained by adaptive nonlinearity (Fig. 7). This leads us to an analysis of our runtime performance where we observe that contact solves, which scale in the number of contacts, dominate the cost of our simulations so that, in turn, our algorithm likewise scales in the number of contacts processed (Fig. 9).

However, this is only the beginning of the story—the number of contacts processed in each scene clearly affect the simulation output. We thus explore the tradeoff between the final cost and quality of simulations, via the number of contacts processed, by considering variations in both the seeding density of rooted rods (rods/cm²), and the number of contacts sampled by collision detection.

In the following we vary seeding density by increasing the number of rods rooted over a fixed surface area. As we do so we see that the number of contacts resolved grows in a superlinear fashion (Fig. 12, middle) suggesting that we face a severe computational challenge as we scale towards contacting rod assemblies at reported human hair densities.

Likewise for a fixed seeding density we can increase the number of contacts sampled at the cost of lengthening runtime. What are the advantages to increasing contact sampling? We investigate this below and observe that increased contact sampling leads to the capture of more local features in our simulations.

8.1 Case Study 1: Single Rod Collisions

We begin with a simple example to consider how the stability behavior of an isolated colliding rod varies as we change collision response method. We designed this example to specifically exercise

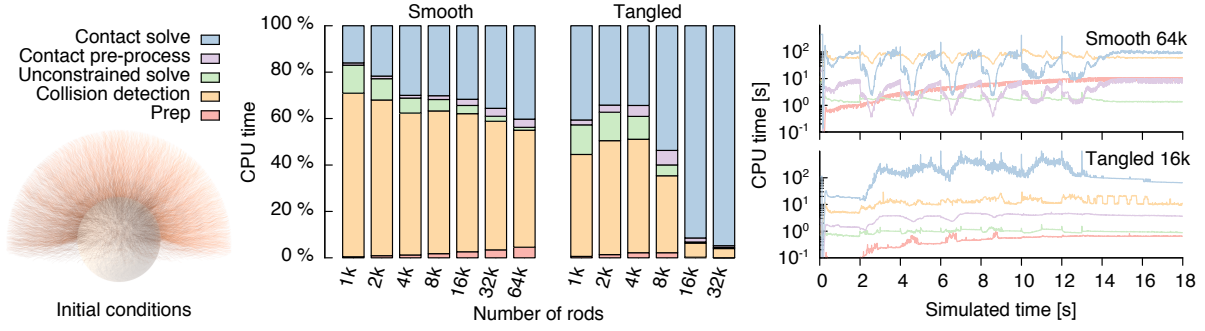


Figure 8: CPU breakdown statistics. Left: Initial conditions for the hair ball examples (8K rods). Middle: across all hair ball simulations as we increase the number of rods/seeding density. Right: By time step for the 64K rod smooth (top) and 16K rod tangled (bottom) simulations.

the worst-case scenario of repeated small-bend collisions to better understand the extremity of nonlinear collision response behavior. We rotate a scripted handle connected to an elastic rod so that it repeatedly and vigorously whips the rod against the edge of a thin wall obstacle. See Fig. 7, left, for set-up and simulation snapshots. The physical parameters for the rod used in this example are: material density $\rho = 1.3 \text{ g/cm}^3$, elastic modulus $Y = 10^{10} \text{ g/(cm} \cdot \text{s}^2)$, shear modulus $S = 3.4 \times 10^9 \text{ g/(cm} \cdot \text{s}^2)$, and radius $r = 60 \text{ } \mu\text{m}$.

To understand stability gains we plot the stability regions of all three response algorithms over varying time step size and rotational whipping speed. To determine stability we stipulate success as a completed run over a time-period of five seconds during which axial extension does not exceed an unusually forgiving tolerance of 50% rest length (our usual tolerance for simulation examples is 1% rest length). Fig. 7, right, depicts the stable regions. Observe the gain of two orders of magnitude in maximum stable time step size for ADONIS. The stability regions of the linear and impulse response algorithms largely overlap, confirming that in practice for small-bend collisions linear compliance obtains the same effective response as that of an unfiltered impulse.

8.2 Case Study 2: Hair Balls

We now shift our focus from a single to a dense assembly of rods. Starting with a sphere of roughly human head proportions (18 cm diameter), we uniformly seed curly rods over 50% of the surface, and script the sphere through a sequence of rotations about three orthogonal axes, alternating rotations with rest phases (see Fig. 10, bottom, for scripting details). This scenario exercises the rod assembly through a full range of tossing, tumbling, and spinning. Fast collisions are initiated at both the start and end of each rotation phase, and the pauses in between are sufficient for settling into slower contacting behavior.

We run our simulations in geometrically increasing sequences ranging from 1K up to 64K rods in our largest example. This corresponds to a maximum seeding density of $\sim 125 \text{ rods/cm}^2$. In comparison the average full head of human hair has $175 - 300 \text{ hairs/cm}^2$ [Robbins 2012]. We select reported human hair parameters: material density $\rho = 1.32 \text{ g/cm}^3$, viscosity $\mu_v = 5 \times 10^7 \text{ g/(cm} \cdot \text{s)}$, gravity $g = 981 \text{ cm/s}^2$, elastic modulus $Y = 3.9 \times 10^9 \text{ g/(cm} \cdot \text{s}^2)$, shear modulus $S = 3.4 \times 10^9 \text{ g/(cm} \cdot \text{s}^2)$, rod-rod frictional coefficient $\mu = 0.2$, rod-sphere frictional coefficient $\mu = 0.1$, and rod radius $r = 37 \text{ } \mu\text{m}$.

In the following hair ball examples each rod is discretized to 119 DoFs total, 30 vertices and an additional 29 twist DoFs per rod. Except where otherwise noted, we apply proximity based collision detection for rod-rod contact sampling and employ a time step of

$h = 10 \text{ ms}$. To better understand the effects of increased contact sampling we have applied two different proximity radii, $2.5 \text{ } \mu\text{m}$ and $25 \text{ } \mu\text{m}$ respectively. As we explore below, varying proximity radii has some clear tradeoffs. In the following sections, based on the resulting behaviors we obtain we distinguish between these two simulation types as respectively *smooth* and *tangled*. For the smooth and tangled simulations we set `gs_max` to 1050 and 150 respectively. Here and in the following the convergence criterion for ADONIS limits stretching to $\text{sf} < 1\%$.

Statistics for the smooth and tangled simulations were obtained respectively with Intel Xeon E5-4650 @ 2.7GHz (8 core Sandy Bridge-EP, 4 sockets) and Intel Xeon E5-2680v2 @ 2.8GHz (10 core Ivy Bridge-EP, 2 sockets) systems, with the exception of the 32K tangled simulation, which ran on an Intel Xeon E5-2650 @ 2GHz (8 core Sandy Bridge-EP, 2 sockets).

Timing breakdown and scaling Fig. 8, middle, shows the breakdown of CPU time for these simulations. The cost of collision resolution becomes dominant as we move towards increased contact sampling and higher rod densities. Indeed, for the most complex scenarios collision resolution is dominant at every time step. See Fig. 8, right. Recall that at each time step we can have multiple contact problems, each corresponding to the resolution of an independent contact group. In turn, the total cost of contact resolution at each time step depends on the number and size of these groups. In Fig. 9 we plot the time required to solve individual contact groups as a function of the number of contacts in the group, across all groups encountered in the smooth (top) and tangled (bottom) simulation sequences, and observe close to linear scaling.

At each Newton iteration we attempt to solve the contact problem in (3). A solution for this optimization should be expected to scale nonlinearly in the number of contact variables so that it is initially surprising that we observe linear scaling overall. However, recall that to solve (3) we employ the Gauss-Seidel solver which regularly saturates at its upper iteration limit. We conjecture that this saturation is the source of the observed linear scaling.

Sufficient nonlinearity We return to the question of exactly how much additional work is required when linear compliance is insufficient to better understand the role this plays in our algorithm's performance. Across our simulations we find that the average number of constrained Newton iterations, per contact solve, remains close to one. However, this does not offer a sufficiently detailed view of how our algorithm adapts to nonlinearity over time. Nor does it help us to understand the potential cost of the resulting additional Newton iterations. In Fig. 10 we instead plot the average number of constrained Newton iterations required by our algorithm *weighted*

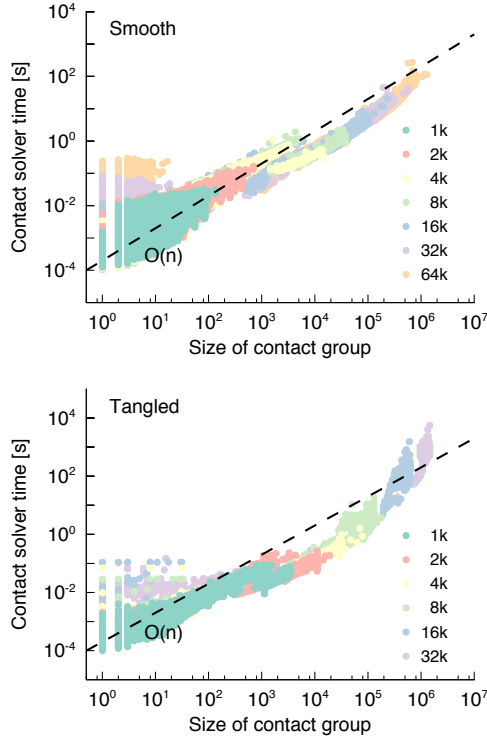


Figure 9: Contact solver scaling statistics. Here we plot the time required to solve a contact group, as a function of the number of contacts, across all groups encountered in the smooth (top) and tangled (bottom) simulation sequences. Linear scaling is plotted for reference.

by contact problem size per time step since, as we have seen above, larger contact problems are more costly. We correspondingly observe here and in Fig. 12, left, that the vast majority of contact problems encountered require small numbers of iterations.

Turning the contact numbers knob Observing that our algorithm scales well in the number of contacts we next explore the tradeoffs between final runtime cost and the resulting simulation by considering variations in seeding density and contact sampling.

Varying seeding density in the hair ball examples while keeping seed area fixed we see in Fig. 12, middle, that the total number of contacts in the system grows superlinearly with the number of rods. Correspondingly in our supplemental video we show a comparison sequence to illustrate how simulation quality of the hair ball examples shifts with increased seeding density.

Simply counting hairs, however, is insufficient to determine the number of contacts we must process and the quality of the resulting simulation. We can also vary the amount of contacts we detect by changing our rod discretization, our collision detection method, or even parameters within our collision detection method. Here we consider the latter and look at the change in behavior and contact numbers for the two applied proximity radii we use in rod-rod collision detection. As we see in a side by side comparison of tangled (left) and smooth (right) simulations of the 32K scene in Fig. 11, increased contact sampling leads to the capture of local features, e.g., lock and ringlet-like structures. See also our supplemental video. However, due to the extremely tight confines of these twists, this comes at the cost of a 10-fold and greater increase in the number of contacts we process as we scale to denser systems and

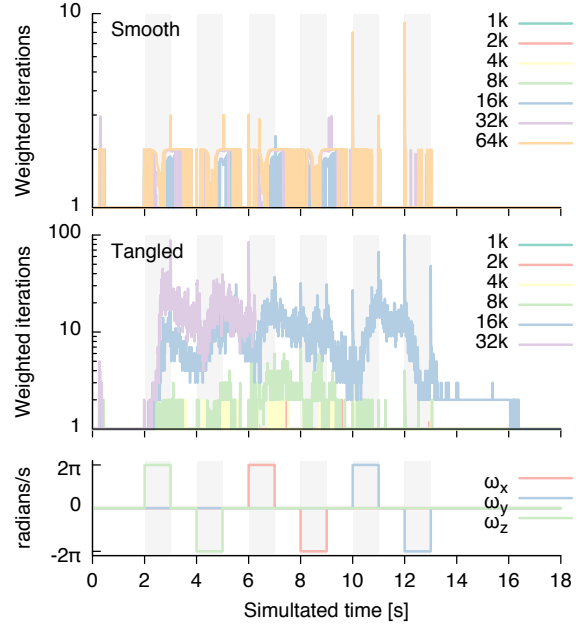


Figure 10: Sufficient nonlinearity. To understand the potential cost of the resulting additional Newton iterations we plot the average number of constrained Newton iterations applied by ADONIS, weighted by contact problem size per time step for smooth (top) and tangled (middle) simulations. Spikes in iteration counts correspond to instants in time when the scripting (bottom) exhibits discontinuities. Either by initiating new rotations or abruptly coming to a halt.



Figure 11: Differing contact numbers and resulting simulated behaviors are obtained by changing the proximity radii we use in rod-rod collision detection from 25 μm (left) to 2.5 μm (right).

a corresponding increase in the total run time and memory usage. See Fig. 12, right. We have documented both simulation types in our supplemental video, where we see that the level of contact complexity we reach begins to involve large portions (often the entire assembly) of constituent rods tightly intertwined together in locks.

Robustness against jitters: choice of unconstrained guess

It is tempting to try initializing each of our contact solves with the inexpensive solution to (LIE) rather than the more expensive, fully nonlinear guess we generate from solving (DEL). However, doing so introduces large jitters and popping to the simulation, as documented in the accompanying video. These are especially distracting as simulations try to come to a rest. Moreover, as discussed above, the *unconstrained* solutions to (DEL) are not a bottleneck to compute, as each rod can be solved independently in parallel.

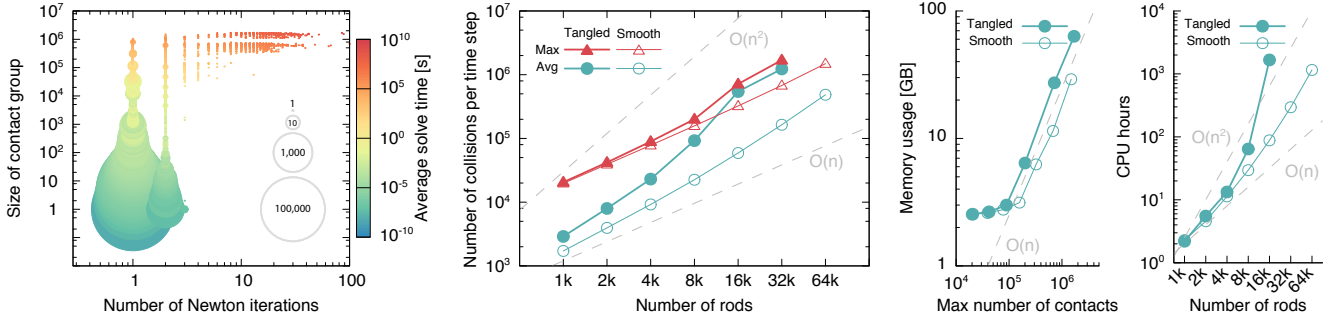


Figure 12: Left: The distribution of Newton iterations per contact group size and corresponding average solve times, over all hair ball simulation sequences. Disc areas indicate numbers of occurrences of problems at each group size and iteration count; grey circles on the right demonstrate the scaling between disc size and number. The color of all plotted discs shows the average solve time on a logarithmic scale from which we note that the vast majority of problems are small and inexpensive to solve. Middle: The number of contacts grows superlinearly in the number of rods/seeding density. Right: A summary of maximum memory utilization and CPU scaling across hair ball simulations.

Stability in rod assemblies We have seen above in §8.1 that when stressed by high-speed collisions ADONIS enables progress at stable time steps orders of magnitude larger than existing methods. Is this stability advantage maintained when we consider less violent motions in rod assemblies? To answer this question we revisit the scripted hair ball example, instrumenting it to examine the stability behavior of all three resolution methods as we vary both time step and seeding density. Since we wish to compare stability and performance across time step sizes we employ CCD for rod-rod contact sampling. This ensures that larger time step simulations do not gain the unfair advantage of smaller contact groups due to missed collisions, or “tunneling.”

We observe in Figure 15 that ADONIS is consistently stable at larger time step sizes in these examples, at a minimum an order of magnitude larger than linear compliance and two orders of magnitude larger than impulsive response. Furthermore this increased stability does not impose additional cost. Rather we see that non-linear adaptivity maintains the fastest runtimes across simulations when allowed to take large, stable time steps.

8.3 Combing, Flinging and Tangling

In the hair ball tests we consider rod assemblies at slower speeds. Here we start to stress ADONIS with a trio of examples where the simulation of combing, flinging, and tangling behaviors are made possible by adaptive nonlinearity. In all of the following examples we once again employ rod-rod CCD.

Comb out : A combing stress test subjects thin rods to collisions and tangling that exercise the strongly nonlinear collision response of rods. We comb through rods as they are tightly rotated and coiled about two rotors. See Fig. 6. Simulation details: $h = 4$ ms; 5,200 rods, 119 DoFs each; runtime: 24h28m on a Xeon E7-8870.

Debris fling Dropped debris is entrained and thrown by rapidly rotating stiff bristles. See Fig. 13. Simulation details: $h = 8$ ms; 1040 bristle rods, 399 DoFs each; 26 dropped rods, 79 DoFs each; runtime: 44m on a MacBook Pro 2011, Intel Core i7 @ 2GHz.

Rod catch Thin rods are caught and pulled into two separate hanks by stiff, rotating elastic bristles. The rods are then wound about each other so that when pulled back out of the rotating bristles they are braided together. See Fig. 14. Simulation details:

$h = 8$ ms; 1040 bristle rods, 399 DoFs each; 1000 thin rods, 239 DoFs each; runtime: 30h14m on a Xeon E7-8870.

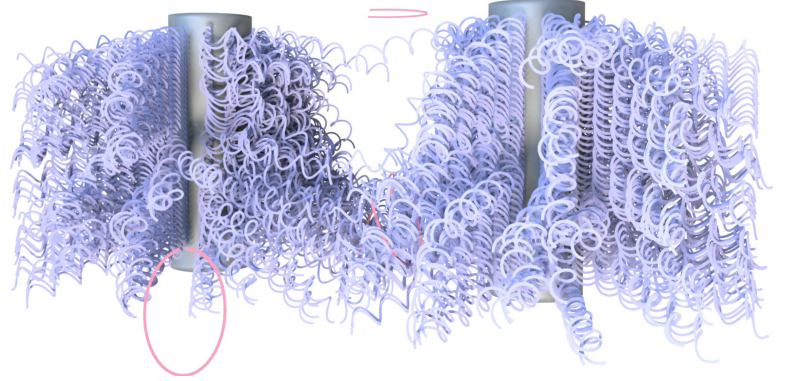


Figure 13: Debris is entrained and thrown by rotating bristles.

rods	method	100ms	10ms	5ms	1ms	0.1ms
1K	ADONIS	x	3h42m	3h28m	4h09m	10h58m
1K	Linear	x	x	x	6h38m	11h09m
1K	Impulse	x	x	x	x	13h27m
2K	ADONIS	x	6h29m	3h46m	5h11m	23h49m
2K	Linear	x	x	x	8h30m	23h35m
2K	Impulse	x	x	x	x	28h51m
4K	ADONIS	x	43h55m	21h47m	19h40m	55h19m
4K	Linear	x	x	x	x	55h37m
4K	Impulse	x	x	x	x	76h34m

Figure 15: A comparison of stable time step sizes and runtimes for response methods on the hair ball example. As we scale to larger seeding densities we see a stability gain for ADONIS of one to two orders of magnitude. Entries in the table give either the runtime to completion or an x to indicate a failed simulation. These simulations were all run on an Intel Xeon X5650 @ 2.67GHz (4 core Westmere-EP, 1 socket).

8.4 Limitations

Capturing the full range of rod assembly behavior requires consideration of both rapid impact between slender bodies—our focus here—as well as stable, persistent contact. In these latter,

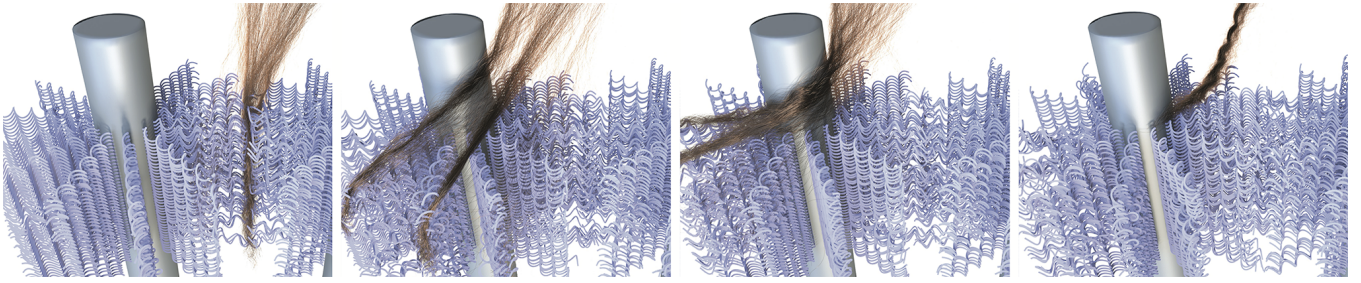


Figure 14: Thin rods are caught and pulled into two separate hanks by stiff, rotating elastic bristles. The rods are wound about each other so that when pulled out they are braided together.

slow moving contact phases, we have shown that ADONIS retains comparable computational advantages of existing contact resolution methods. However, it also retains the same weaknesses: the Gauss-Seidel solver typically does not converge. We conjecture that this lack of convergence is responsible for a notable artifact. We observe that locks formed in our simulations may fall apart on their own, over time (refer to supplemental video). The simulation of stable frictional contact assemblies depends on the accurate resolution of the underlying contact problem [Kaufman et al. 2008]. We conjecture that the unraveling we observe could be addressed in future work by replacing the current Gauss-Seidel solver employed in steps 7-13 of ADONIS with a convergent solver for (3).

To resolve and store contact graphs during solves we utilize a large amount of memory that grows linearly in the number of contacts and thus superlinearly in seeding density; see Fig. 12, right. Our method offers robust stability for low cost. We have been able to run our simulations at large time steps relative to existing methods; but of course, even the Newton method can be unstable, or require too many iterations, if the time step applied is too large.

Thus far, our focus has been on identifying challenging contact scenarios, giving us confidence that our method is robust enough to perform well within a broad range of parameters. As we begin to work with hair scenarios there are many complexities that need to be accounted for. Already, those examples that are hair-like in appearance were configured with rods of typical human hair diameter (as opposed to the diameter of a wisp of hair). Achieving a realistic look for human hairs will likely involve tuning parameters such as the friction coefficient, plasticity, and damping, and possibly introducing physics for adhesion, lubrication, and electrostatic forces.

9 Discussion

Considering the geometry of thin body collisions we have observed that the degree of nonlinearity in collision response for stretching modes varies greatly with configuration. Noting that capturing this nonlinearity is essential for the stable progress of simulations at practical time step sizes, we have constructed a simple yet parsimonious algorithm that applies first-order modeling in most solves. When linearity is sufficient our proposed method incurs marginal cost over existing linearly compliant methods. It identifies instances where a frugal application of additional computation is most needed, enabling us to take time steps several orders of magnitude larger than previously possible. Importantly, this does not mean that *every* response is resolved with a small number of Newton iterations: as we have seen above in §8.2, in a *small but critical* subset of collision events, large numbers of iterations are necessary, without which the time step restriction is severe. The crucial point is that these instances are infrequent, and so when amortized across a simulation we incur minimal additional cost.

In our development we have focused on implicit Euler for didactic purposes and noted that our algorithm can be based on other implicit time-integration methods in a similar fashion. There is one caveat in that not all methods are guaranteed to generate an inelastic impact without modification. We cover the necessary details in Appendix A. Likewise, we have so far considered close to inextensible materials. For stiff materials further from the inextensible limit we can instead apply a bound on a discrete strain-rate measure $sr^i = \max_{x,j} (|e_{x,j}^i| - |e_{x,j}^t|)/h$ as a termination criterion.

Looking forward this work illuminates a number of potential future research directions. Most immediately we note that our analysis in §1 of the geometry of nonlinear collision response for stretching modes applies directly to shells and thus to improving collision resolution in cloth simulation. We also note that it is tempting to consider how adaptive nonlinearity can be formulated as a pairwise-iteration method [Bridson et al. 2002]; however, in our experience so far this approach is subject to highly challenging convergence issues. We have also observed that the primary source of nonlinear scaling in our solves comes from the size of the largest connected components in the contact graph. On the other hand, decomposing the graph by simply omitting contacts leads to entanglement and visual artifacts. Research towards an adequate decomposition of the contact graph is an exciting and impactful future avenue of exploration.

Acknowledgements

We thank Jorge Obregon, Justin Kern, Ying Liu, David Aguilar, and Heather Pritchett (WDAS); Alasdair Coull, David Clayton, David Gouge, John Martin, and Marco Relevant (Weta); Keenan Crane and Michelle Ming-Yen Lee (Columbia); Gilles Daviet (INRIA); and the Bedford Hill Coffee Bar. This work was supported in part by the NSF (Grants IIS-1319483, CMMI-1331499, IIS-1217904, IIS-1117257, CMMI-1129917, IIS-0916129), the Israel-US Binational Science Foundation, Intel, The Walt Disney Company, Weta Digital, Side Effects, Autodesk, and NVIDIA.

References

- ALART, P., AND CURNIER, A. 1991. A mixed formulation for frictional contact problems prone to Newton like solution methods. *Computer Methods in Applied Mechanics and Engineering* 92, 3 (Nov.), 353–375.
- ALLARD, J., FAURE, F., COURTECUISSÉ, H., FALIPOU, F., DURIEZ, C., AND KRY, P. G. 2010. Volume Contact Constraints at Arbitrary Resolution. *ACM Trans. Graph.* 29, 4 (July), 82:1–82:10.
- BARAFF, D., AND WITKIN, A. 1998. Large Steps in Cloth Simulation. In *Proceedings of SIGGRAPH 98*, Annual Conference Series, 43–54.
- BARAFF, D. 1989. Analytical Methods for Dynamic Simulation of Non-penetrating Rigid Bodies. *Computer Graphics* 23, 223–232.
- BERGOU, M., WARDETZKY, M., ROBINSON, S., AUDOLY, B., AND GRINSUN, E. 2008. Discrete Elastic Rods. *ACM Trans. Graph.* 27, 3 (Aug.), 63:1–63:12.
- BERGOU, M., AUDOLY, B., VOUGA, E., WARDETZKY, M., AND GRINSUN, E. 2010. Discrete Viscous Threads. *ACM Trans. Graph.* 29, 4 (July), 116:1–116:10.
- BERTAILS-DESCOUBES, F., CADOUX, F., DAVIET, G., AND ACARY, V. 2011. A Nonsmooth Newton Solver for Capturing Exact Coulomb Friction in Fiber Assemblies. *ACM Trans. Graph.* 30, 1 (Feb.), 6:1–6:14.
- BONNEFON, O., AND DAVIET, G. 2011. Quartic formulation of Coulomb 3D frictional contact. Tech. Rep. RT-0400, INRIA, Jan.
- BRIDSON, R., FEDKIW, R., AND ANDERSON, J. 2002. Robust Treatment of Collisions, Contact, and Friction for Cloth Animation. *ACM Trans. Graph.* 21, 3 (July), 594–603.
- DAVIET, G., BERTAILS-DESCOUBES, F., AND BOISSIEUX, L. 2011. A Hybrid Iterative Solver for Robustly Capturing Coulomb Friction in Hair Dynamics. *ACM Trans. Graph.* 30, 6 (Dec.), 139:1–139:12.
- DURIEZ, C., ANDRIOT, C., AND KHEDDAR, A. 2004. Signorini’s contact model for deformable objects in haptic simulations. In *IEEE/RSJ IROS*, vol. 4, 3232–3237.
- DURIEZ, C., DUBOIS, F., KHEDDAR, A., AND ANDRIOT, C. 2006. Realistic Haptic Rendering of Interacting Deformable Objects in Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics* 12, 1 (Jan.), 36–47.
- GOYAL, S., RUINA, A., AND PAPADOPOULOS, J. 1991. Planar sliding with dry friction Part 2. Dynamics of motion. *Wear* 143, 2, 331–352.
- HADAP, S., CANI, M.-P., LIN, M., KIM, T.-Y., BERTAILS, F., MARSCHNER, S., WARD, K., AND KAČIĆ-ALESIĆ, Z. 2007. Strands and Hair: Modeling, Animation, and Rendering. In *ACM SIGGRAPH Courses*, 1–150.
- HAIRER, E., AND WANNER, G. 2004. *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, second ed. Springer.
- IBEN, H., MEYER, M., PETROVIC, L., SOARES, O., ANDERSON, J., AND WITKIN, A. 2013. Artistic Simulation of Curly Hair. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 63–71.
- JEAN, M., AND MOREAU, J. J. 1992. Unilaterality and dry friction in the dynamics of rigid body collections. In *Proceedings of Contact Mechanics International Symposium*, vol. 1, 31–48.
- JOURDAN, F., ALART, P., AND JEAN, M. 1998. A Gauss-Seidel like algorithm to solve frictional contact problems. *Computer Methods in Applied Mechanics and Engineering* 155, 1 (Mar.), 31–47.
- KAUFMAN, D. M., SUEDA, S., JAMES, D. L., AND PAI, D. K. 2008. Staggered Projections for Frictional Contact in Multibody Systems. *ACM Trans. Graph.* 27, 5 (Dec.), 164:1–164:11.
- KIKUCHI, N., AND ODEN, J. T. 1988. *Contact Problems in Elasticity: A Study of Variational Inequalities and Finite Element Methods*, vol. 8 of *SIAM Studies in Applied and Numerical Mathematics*. Society for Industrial and Applied Mathematics.
- LANDAU, L. D., AND LIFSHITZ, E. 1986. *Theory Of Elasticity, Course of Theoretical Physics, Vol. 7*. Pergamon Press, Oxford.
- MCADAMS, A., SELLE, A., WARD, K., SIFAKIS, E., AND TERAN, J. 2009. Detail Preserving Continuum Simulation of Straight Hair. *ACM Trans. Graph.* 28, 3 (July), 62:1–62:6.
- MIRTICH, B., AND CANNY, J. 1995. Impulse-based Simulation of Rigid Bodies. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics*, 181–ff.
- MOREAU, J. J. 1988. Unilateral Contact and Dry Friction in Finite Freedom Dynamics. *Nonsmooth Mechanics and Applications, CISM Courses and Lectures*, 302, 1–82.
- OTADUY, M. A., TAMSTORF, R., STEINEMANN, D., AND GROSS, M. 2009. Implicit Contact Handling for Deformable Objects. *Computer Graphics Forum* 28, 2, 559–568.
- PROVOT, X. 1997. Collision and self-collision handling in cloth model dedicated to design garments. In *Computer Animation and Simulation*, Eurographics, 177–189.
- ROBBINS, C. R. 2012. *Chemical and Physical Behavior of Human Hair*, fifth ed. Springer.
- SELLE, A., LENTINE, M., AND FEDKIW, R. 2008. A Mass Spring Model for Hair Simulation. *ACM Trans. Graph.* 27, 3 (Aug.), 64:1–64:11.
- SPILLMANN, J., AND TESCHNER, M. 2007. CORDE: Cosserat Rod Elements for the Dynamic Simulation of One-Dimensional Elastic Objects. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 63–72.
- STEWART, D. E. 2001. Finite-dimensional contact mechanics. *Phil. Trans. R. Soc. Lond. A* 359, 2467–2482.
- STEWART, D. E. 2011. *Dynamics with Inequalities: Impacts and Hard Constraints*. Society for Industrial and Applied Mathematics.
- WARD, K., BERTAILS, F., KIM, T.-Y., MARSCHNER, S. R., CANI, M.-P., AND LIN, M. C. 2007. A Survey on Hair Modeling: Styling, Simulation, and Rendering. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (Mar.), 213–234.
- ZHENG, C., AND JAMES, D. L. 2011. Toward High-Quality Modal Contact Sound. *ACM Trans. Graph.* 30, 4 (July), 38:1–38:12.

A Implicit restitution in time-integration

Lemma A fully *inelastic impact* will only be obtained by one-step numerical integrators that satisfy an implicit-Euler-type velocity update of the form $\dot{\mathbf{q}}^{t+1} = a\delta$, $a \in \mathbb{R}$.

Proof. Recall that we have $\mathbf{q}^{t+1} = \mathbf{q}^t + \delta$ and that imposing a linearized non-negative displacement constraint requires $0 \leq \alpha_k \perp \mathbf{n}_k^T \delta \geq 0$. Whenever a contact force is applied we have $\alpha_k > 0$ and thus $\mathbf{n}_k^T \delta = 0$. A fully inelastic impact is given by a response satisfying the velocity-level condition $\mathbf{n}_k^T \dot{\mathbf{q}}^{t+1} = 0$. By substitution this can only be satisfied by an implicit-Euler-type update where $\dot{\mathbf{q}}^{t+1}$ is given by a scaling of δ . \square

As concrete examples consider that the implicit Euler and implicit midpoint velocity updates are respectively $\dot{\mathbf{q}}_E^{t+1} = 1/h \delta$ and $\dot{\mathbf{q}}_M^{t+1} = 2/h \delta - \dot{\mathbf{q}}^t$. Then, by substitution, we have a *fully inelastic impact* for implicit Euler since $\mathbf{n}_k^T \dot{\mathbf{q}}_E^{t+1} = 0$ and, on the other hand, a *fully elastic impact* for implicit Midpoint, corresponding to a coefficient of restitution equal to one, since $\mathbf{n}_k^T \dot{\mathbf{q}}_M^{t+1} = -\mathbf{n}_k^T \dot{\mathbf{q}}^t$.

B First iterate equivalence

Proof. Recall that each Newton iteration is given by

$$\lambda^{i+1} \in \mathcal{R}(\delta^i + d\delta_u^i + \mathbf{C}^i \mathbf{G} \lambda^{i+1}).$$

Or equivalently

$$\lambda^{i+1} \in \mathcal{R}(\delta^i + \mathbf{C}^i [h\mathbf{M}\dot{\mathbf{q}}^t - h^2 \nabla V(\mathbf{q}^t + \delta^i) - \mathbf{M}\delta^i] + \mathbf{C}^i \mathbf{G} \lambda^{i+1}).$$

If the initial seed guess is the solution of the unconstrained time step problem (DEL) we have

$$h\mathbf{M}\dot{\mathbf{q}}^t - h^2 \nabla V(\mathbf{q}^t + \delta^0) - \mathbf{M}\delta^0 = 0,$$

so that substitution into the iteration gives

$$\lambda^1 \in \mathcal{R}(\delta^0 + \mathbf{C}^0 \mathbf{G} \lambda^1).$$

\square