

A Nonsmooth Newton Solver for Capturing Exact Coulomb Friction in Fiber Assemblies

FLORENCE BERTAILS-DESCOUBES, FLORENT CADOUX, GILLES DAVIET, and VINCENT ACARY
INRIA

We focus on the challenging problem of simulating thin elastic rods in contact, in the presence of friction. Most previous approaches in computer graphics rely on a linear complementarity formulation for handling contact in a stable way, and approximate Coulomb's friction law for making the problem tractable. In contrast, following the seminal work by Alart and Curnier in contact mechanics, we simultaneously model contact and exact Coulomb friction as a zero finding problem of a nonsmooth function. A semi-implicit time-stepping scheme is then employed to discretize the dynamics of rods constrained by frictional contact: this leads to a set of linear equations subject to an equality constraint involving a nondifferentiable function. To solve this one-step problem we introduce a simple and practical nonsmooth Newton algorithm which proves to be reasonably efficient and robust for systems that are not overconstrained. We show that our method is able to finely capture the subtle effects that occur when thin elastic rods with various geometries enter into contact, such as stick-slip instabilities in free configurations, entangling curls, resting contacts in braid-like structures, or the formation of tight knots under large constraints. Our method can be viewed as a first step towards the accurate modeling of dynamic fibrous materials.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Modeling, simulation, contact, Coulomb friction, dynamics of thin elastic rods, constraint-based method, knot tying, hair simulation

ACM Reference Format:

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, Article 6 (January 2011), 14 pages. DOI = 10.1145/1899404.1899410 <http://doi.acm.org/10.1145/1899404.1899410>

1. INTRODUCTION

1.1 Motivation

Objects composed of thin deformable rods in contact are widely spread in the real world: hair, wool, entangled ropes or wires, knots in suture strands, etc., all fall into this category. Simulating such systems is particularly challenging, for three main reasons: first, finding a robust model for an individual strand that properly captures the important modes of deformation—bending and twisting—is known to be a difficult problem, mainly due to the stiff, high-order equations that characterize such a system. Second, resolving the multiple impacts and resting contacts occurring within a single entangled rope or an assembly of fibers is complex, and made even more difficult by the slender geometry of individual fibers. This calls for the use of extremely robust methods both for collision detection and response. Third, capturing the typical stick-slip effects, or tangles and knots that often occur in fibrous materials (see Figure 1), requires a realistic, nonsmooth model for friction.

Recently, a number of successful models for the dynamics of thin elastic rods (also referred to as “strands”) were proposed in the computer graphics (CG) community [Bertails et al. 2006; Hadap 2006; Spillmann and Teschner 2007; Theetten et al. 2008; Bergou

et al. 2008; Selle et al. 2008]. In this article, we focus on the specific problem of the *contact and friction response* applied to thin elastic rods. This topic was hardly addressed in the past, because of the complexity of such a problem and the inability of classical methods to bring satisfying solutions. We propose here a first step towards the realistic modeling of dynamic rods subject to frictional contact.

1.2 Related Work

We briefly review existing models for thin elastic rods before presenting the main approaches for simulating contact and friction in the general case of interacting (rigid or deformable) bodies. Finally, we summarize the different techniques that have been employed for simulating contact and friction in the case of thin elastic rods.

1.2.1 Modeling Thin Elastic Rods. Models for thin elastic rods can be categorized into two distinct families: *maximal-coordinates* and *reduced-coordinates* models.

Maximal-coordinates models generally parameterize the centerline of the rod explicitly as a sequence of 3D space points, and formulate extra constraints to enforce the kinematics of the rod [Rosenblum et al. 1991; Lenoir et al. 2004; Choe et al. 2005; Spillmann and Teschner 2007; Bergou et al. 2008; Selle et al. 2008].

Authors' addresses: F. Bertails-Descoubes (corresponding author), F. Cadoux, G. Daviet, and V. Acary, INRIA Rhône-Alpes, 655 avenue de l'Europe, 38334 Saint-Ismier Cedex, France; email: Florence.Descoubes@inrialpes.fr.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2011 ACM 0730-0301/2011/01-ART6 \$10.00 DOI 10.1145/1899404.1899410 <http://doi.acm.org/10.1145/1899404.1899410>

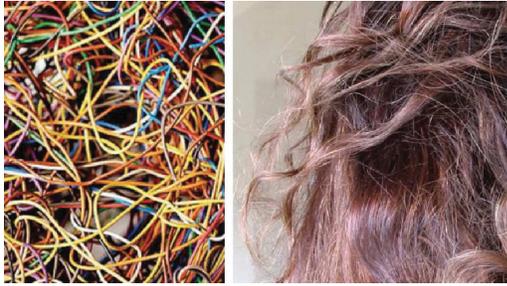


Fig. 1. Typical configurations observable in real assemblies of thin strands. Left: Tangled wires forming loops and knots. Right: hairs wisps sticking onto another hair layer due to dry friction.

These approaches are popular because finding maximal-coordinates is simple and often leads to a sparse, block diagonal mass matrix. However, formulating the extra constraints can be challenging, especially when one wants to account for the inextensibility of the rod: adding stiff terms can lead to stability issues [Rosenblum et al. 1991; Spillmann and Teschner 2007; Selle et al. 2008]. Another difficulty stems from the numerical drifts and energy loss when enforcing these constraints over time [Bergou et al. 2008].

In contrast, reduced-coordinates models parameterize a rod using its actual degrees of freedom [Hadap and Magnenat-Thalmann 2001; Hadap 2006; Bertails et al. 2006; Bertails 2009]. Such parameters often express a local rotation of the system, for example, joint quaternions in the case of the articulated rigid bodies system [Hadap and Magnenat-Thalmann 2001; Hadap 2006] and material curvatures and twist in the case of the superhelix model [Bertails et al. 2006; Bertails 2009]. The parameterization is thus guaranteed to be minimal, with no redundant variables, and the kinematic constraints are always exactly maintained over time. In this kind of formulation, however, the centerline of the rod has to be computed recursively from root to tip, as a nonlinear function of the parameters.

An argument that is often raised against reduced-coordinates models is the (apparent) difficulty to handle contact and friction with a nonexplicit centerline [Bergou et al. 2008; Selle et al. 2008]. In this article, however, we show that reduced-coordinates models perfectly fit into our frictional contact formulation, with no additional cost. We present results using the superhelix model as our rod model. To demonstrate the versatility of our approach, we have also performed simulations using maximal-coordinates rod models such as an implicit mass-spring system [Baraff and Witkin 1998] as well as the more recent CORDE model [Spillmann and Teschner 2007]. We demonstrate that the convergence of our contact algorithm is insensitive to the nature of parameterization of the rod models, but rather depends on the *level of constraining* of the dynamical system.

1.2.2 Modeling Contact and Friction. Modeling contact and friction has been an active research area in CG for several decades. Because of the large number of techniques employed and the wide spectrum of targeted applications (ranging from rigid bodies to 3D deformable objects, cloth, thin shells, and rods), it can be somewhat difficult for a nonexpert reader to find a path into this intricate field. We try here to offer a synthetic view on the topic by classifying previous work into three main categories: *penalty-based*, *explicit constraint-based*, and *implicit constraint-based* methods. Note that we focus here on the modeling of the *response* to collision. For a recent survey on collision detection methods, we refer the reader to Teschner et al. [2005].

Penalty-based methods. One common way of preventing interpenetration between colliding objects consists in adding mutual repulsive forces of the form $k\mathbf{f}(\delta)$, where δ is the penetration depth detected at current time step [Moore and Wilhelms 1988]. Though simple to implement and efficient, this method often fails to prevent excessive penetration if the stiffness factor k is not large enough, unless barrier functions are used [Kaldor et al. 2008; Harmon et al. 2009]. However, most barrier functions suffer unbounded second derivatives which ruin the stability of the fixed step integration schemes. To circumvent this problem, Harmon et al. [2009] recently proposed the use of discrete penalty layers coupled with an asynchronous symplectic solver, at the price of losing computational efficiency. Another issue inherent to penalty-based approaches is the introduction of parasitical high frequencies due to the recourse to large stiffness values for preventing interpenetration [Baraff 1989]. Penalty-based approaches are thus generally not satisfying for robust contact handling, however, they remain widely used in situations where time performance is a crucial criterion [Barbič and James 2007].

In the same vein, friction is often modeled using simple viscous forces, especially in the case of self-contact within deformable objects, for the sake of simplicity [Baraff and Witkin 1998; Choe et al. 2005; Spillmann and Teschner 2008; Kaldor et al. 2008]. Such a model cannot capture the threshold effect that characterizes friction between solids.

Explicit constraint-based methods. Instead of adding some forces that will hopefully lead to a collision-free state, an alternate way consists in deciding *in advance* what positions and velocities should be at the next time step. To enforce the desired state, some authors propose to directly alter positions and/or velocities at the end of the time step [Baraff and Witkin 1998; Müller et al. 2007] while others compute the corresponding forces that will exactly lead to the desired, collision-free state [Spillmann et al. 2007; Spillmann and Teschner 2008]. Because of its simplicity and efficiency, this kind of approach has become very popular in the CG community for resolving contact in the case of systems parameterized by sample 3D points, such as 3D point-based deformable bodies [Müller et al. 2007], cloth [Baraff and Witkin 1998; Bridson et al. 2002], and hair [Choe et al. 2005; Selle et al. 2008]. However, this kind of technique is hardly applicable to systems that are subject to multiple impacts (in this case the desired change in velocity may be unclear), as well as to sophisticated models such as reduced-coordinates models, where the 3D positions at the surface of the system are not necessarily linear functions of the degrees of freedom.

In a similar spirit, Coulomb friction is often modeled explicitly, that is, the next state of the objects (take-off, stick, slip) and the sliding direction are determined in advance by considering the contact forces and relative velocities at the current time step [Baraff 1994; Bridson et al. 2002; Selle et al. 2008]. Again, this approach is simple, but as it does not guarantee that the Coulomb conditions are satisfied at the end of each time-step, it may suffer drifts over time and thus be inaccurate. Moreover, to be able to capture perfect dry friction, this method requires the ability to directly control the change in velocity and position of the contacting points of the system. For most systems which are not controllable without altering their own dynamics, this is seldom achievable.

Implicit constraint-based methods. A last solution is to consider that the relative velocities and directions of friction at the next time step are part of the *unknowns* of the problem. The dynamic system, augmented by contact variables, has then to be solved, traditionally using a fixed point method or an optimization algorithm. Computing contact and friction implicitly is the only way of strictly enforcing

the proper conditions at the end of each time step, without the need for manual checking and readjustment.

Robust approaches have been provided for handling contact without friction, mostly relying on Linear Complementarity Problems (LCP) [Baraff 1989; Baraff and Witkin 1992]. However, simulating contact with Coulomb friction in 3D remains a challenging issue because nonlinearity has to be introduced into the equations. To overcome this issue, most approaches in CG propose to approximate Coulomb's friction law. Some authors model friction using a Tresca-like law where the sliding threshold is an arbitrary value independent of the normal force [Lötstedt 1984]. This model is not satisfying in many common situations. Consider, for example, two different solids lying on a horizontal plane. If an increasing, horizontal traction force is applied to them, both of them will start to slide exactly at the same time, no matter what their mass. In Kaufman et al. [2005], Coulomb friction is approximated by considering that the normal component of the reaction force is known from a first computational step without friction. This procedure allows them to retrieve a standard convex QP and can be viewed as a Tresca-like approximation at each time-step. To compute hair-body frictional contact, Hadap [2006] replaces the Coulomb law with a new model based on an LCP. However, the model fails to capture dry friction properly: for a contact force in the interior of the friction cone, the relative velocity is not constrained to be zero. More recent approaches start with the Coulomb law equations and linearize the friction cone to come up with a set of LCP to be solved [Erleben 2007; Kaufman et al. 2008; Otaduy et al. 2009]. While this approximate formulation captures both dynamic and static friction properly, it introduces some anisotropy in the sliding direction unless a high number of facets is used [Acary and Brogliato 2008, Section 13.3.7]. A trade-off has thus to be made between accuracy and computational cost.

It is overall surprising that no method modeling exact Coulomb friction has become popular in the graphics community so far. Yet, such methods have been developed and successfully used for years in the contact mechanics community for accurately simulating multibody systems involving frictional contact, such as granular materials [Moreau 1994; Jean 1999]. We believe the main reason why such methods have not caught on yet in the graphics community is the apparent difficulty of implementation for people who are not experts in nonsmooth optimization. In this article, we demonstrate that the frictional contact problem can be elegantly expressed as a zero finding problem, and provided the system is not over-constrained, be simply solved using a Newton algorithm, without resorting to a complementarity formulation. Our algorithm is furthermore very easy to understand and implement, and does not rely on any black-box optimization code. To help researchers test our approach and compare their methods with ours, we freely provide on our Web page the source code for our solver, coupled with the simple interactive double loop example illustrated in Figure 9.

We note that Duriez et al. [2006] made a first attempt to model exact Coulomb friction for the real-time simulation of frictional contact for haptics. Their approach relies on a so-called Gauss-Seidel algorithm that iteratively solves each single frictional contact while other contacts are frozen: for each contact, the contact force is computed by projection onto the (nonpolyhedral) friction cone. The authors show that their method is more efficient than methods relying on a linearized friction cone for which the number of facets is sufficiently large to ensure a given precision. In their approach, however, the projection step onto the friction cone is significantly simplified as the local Delassus operator is roughly approximated by a diagonal matrix. This approximation has no incidence when applied to isotropic objects such as spheres, but it may give rise to important

drifts when dealing with general, anisotropic objects, for which the tangential velocity during dynamic friction not only depends on the applied external force, but also on the geometry of the object. Another concern of Duriez et al. [2006] is that the formulations used for contact and friction laws are not velocity based, but rely on the normal and tangential gaps between contacting elements. Such acceleration-force formulations have been shown inconsistent as they do not necessarily yield a solution [Baraff 1993; Stewart 2000]. In contrast, impulse-velocity formulations for frictional problems turn out to have better properties of convergence [Marques 1993; Stewart 2000], and as such, are often preferred.

Unlike Duriez et al. [2006], our approach provides a velocity-based formulation for the frictional contact, and relies on a simple, functional formulation of the *exact* Coulomb friction model (initially due to Alart and Curnier [1991]) which is subsequently solved using a nonsmooth Newton algorithm. Moreover, our approach *simultaneously* solves all frictional contact problems, without relying on a Gauss-Seidel-like algorithm. In our results section, we show that provided our dynamic system is not excessively constrained, we obtain satisfying convergence results that guarantee a good visual precision of our simulations.

1.2.3 Self-Contact and Friction in Rods. Due to the number of applications involved, the study and simulation of self-contact in an assembly of thin elastic rods has increasingly drawn attention from researchers in both mechanics and computer graphics.

In both fields, most Lagrangian approaches for simulating contact between strands or groups of strands are penalty based [Plante et al. 2001; Durville 2004; Choe et al. 2005; Bertails et al. 2006] or based on explicit constraints [Spillmann and Teschner 2007; 2008; Bergou et al. 2008; Selle et al. 2008]. In all these methods, friction is simply handled through a viscous term. These methods often suffer instability or inaccuracy issues, and are unable to capture subtle frictional effects with threshold, such as stick-slip effects. To simplify the problem of self-contact in large assemblies of strands (e.g., hair), some approaches resort to procedural modeling for capturing typical emerging phenomena, such as hair clumping, either by constraining hair to move as a predetermined set of wisps [Plante et al. 2001; Bertails et al. 2006], or by adding some attractive forces to capture stiction [Selle et al. 2008]. Hadap and Thalmann [2001] were the first to explore continuum laws for modeling hair self-contact. Very recently, McAdams et al. [2009] combined a Lagrangian and an Eulerian representations for hair to bring some discontinuous details to the simulation. Although spectacular, their results do not capture the typical stick-slip effects that would be visible within a hair clump coming to rest.

From the state-of-the-art, it actually turns out that whatever the scale of the targeted applications is (a knot in a single rod, hair with a hundred guide strands, hair with thousands of simulated strands), robust contact handling with accurate friction modeling remains a *constant missing element*. We propose a first solution to this challenging problem. Our method is related to implicit constraint-based models, which, to the best of our knowledge, were never explored for solving self-contacts within an assembly of fibers. Furthermore, we introduce a velocity-based model for simulating exact Coulomb friction. This allows us to design a pure physics-based model, able to capture for the first time typical contacting phenomena that occur within an assembly of strands.

1.3 Contributions

We provide a compact formulation of frictional contact for thin elastic rods. Unlike previous approaches, our formulation exactly

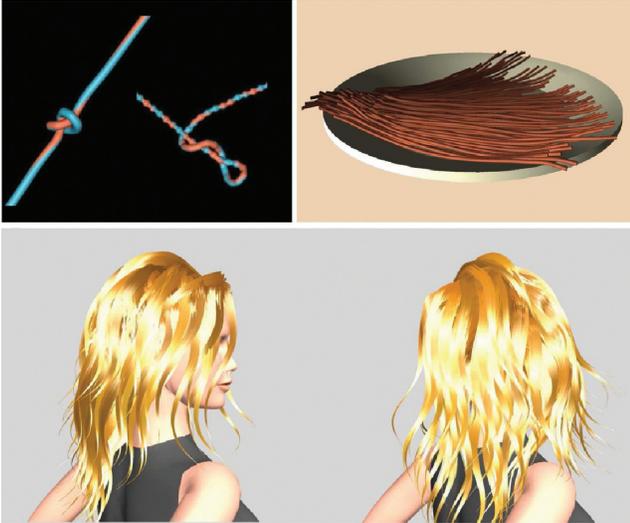


Fig. 2. Various simulations of frictional contact in thin rods. From left to right and top to bottom: interactive formation of a tight knot and a plectoneme, simulation of 76 spaghetti falling in a plate, complex hair contacts with Coulomb friction in wavy hair (84 simulated strands).

models the Coulomb friction law as a zero finding problem of a nonsmooth function, which turns out to be much simpler and efficient to solve compared to a nonlinear complementary formulation. Our approach is inspired by the theoretical work by Alart and Curnier [1991], which, to the best of our knowledge, was never exploited in CG to solve frictional contact.

We give a practical implementation for solving the optimization problem resulting from our frictional contact model. Our method combines a nonsmooth Newton step with an accurate line search in order to stabilize the algorithm and accelerate its convergence. Although theoretical convergence properties are lost in the nonsmooth case, we provide a simple experimental criterion, based on the degree of constraining of the system, to characterize the quality of convergence. When systems are not overconstrained, our solver yields fast convergence at tolerances suitable for graphics purposes and we show that many fiber systems typically lie in the favorable case.

We demonstrate the versatility of our approach by performing simulations on various representative rod models: an implicit mass-spring system, the more elaborate CORDE model (maximal-coordinates), and the superhelix model (reduced-coordinates). To some extent, we also illustrate the effectiveness of our approach on other dynamical systems, such as rigid bodies, while clearly defining the class of systems to which our contact method is the most adapted.

We show that we can robustly simulate systems with large contact forces, such as tight knot tying, as well as capture subtle collective behaviors in entangled fibrous materials, such as hair, by using a pure Lagrangian method. To the best of our knowledge, this is one of the first attempts to accurately model contact and friction within free assemblies of fibers, regarding both CG and computational mechanics fields.

1.4 Organization

Our article is organized as follows. Section 2 introduces a functional formulation for Coulomb's friction model, derived from the early

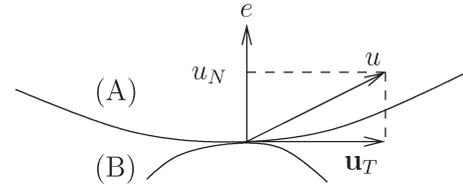


Fig. 3. Body A^i and B^i with tangent and normal spaces.

work by Alart and Curnier. In Section 3, we formulate the discrete problem of the constrained dynamics that we wish to solve at each time step, and provide a practical solving algorithm in Section 4. Finally, results of our approach are presented in Section 5 and discussed in Section 6, before concluding.

2. MODELING CONTACT WITH PERFECT COULOMB FRICTION

The friction model proposed by Charles de Coulomb around 1770–1780 is a macroscopic model which captures a crucial physical phenomenon of friction: the threshold of sliding. This phenomenon is particularly visible on fibrous materials, such as hair, where fibers often get stuck onto each other, giving rise to static configurations with multiple fibers orientations, or to the spontaneous formation of entangled wisps (see Figure 1). Though simple in appearance, the Coulomb friction model is actually difficult to take into account exactly as it is a nonsmooth, multivalued law which cannot be handled using classical mathematical tools of smooth analysis.

Formally, the Coulomb friction law can be described¹ as a set $C \subset \mathbb{R}^3 \times \mathbb{R}^3$ relating the relative velocity \mathbf{u}^i and the contact force \mathbf{r}^i at contact point i : the couple $(\mathbf{u}^i, \mathbf{r}^i)$ satisfies the Coulomb friction law if and only if $(\mathbf{u}^i, \mathbf{r}^i) \in C$. In the sequel, after defining some notations, we shall give two equivalent formulations of C : on the one hand, the classical disjunctive formulation, easy to understand but hardly tractable in practice; on the other hand, a zero finding formulation, initially due to Alart and Curnier [1991], which gives rise to a simple robust algorithm that we derive in Section 4.

2.1 Local Forces and Velocities

Let us consider a mechanical system in 3D space involving n contact points at a given instant t . Each contact, labeled by $i \in 1, \dots, n$, is assumed to involve two and only two contacting bodies, denoted by A^i and B^i . We further assume that contacting surfaces are sufficiently smooth, so that a normal vector \mathbf{e}^i (see Figure 3) can be defined. Taking body B^i as a reference, consider the relative velocity \mathbf{u}^i of A^i with respect to B^i and the force \mathbf{r}^i applied by B^i onto A^i . We note $x_N^i = \mathbf{x}^i \cdot \mathbf{e}^i$ the (scalar) component of a vector \mathbf{x}^i , and $\mathbf{x}_T^i = \mathbf{x}^i - x_N^i \mathbf{e}^i$ the tangential (vector) part of \mathbf{x}^i .

2.2 Disjunctive Formulation

Friction coefficients μ^i ($0 \leq \mu^i \leq \mu_{\max}$) being given, let us define the second-order cone with coefficient μ^i (the friction cone) by

$$K_{\mu^i} := \{\|\mathbf{x}_T\| \leq \mu^i x_N\} \subset \mathbb{R}^3. \quad (1)$$

¹This formalism actually applies to all friction laws. For example, the viscous friction law can be formulated by defining the set C as a vector space. In the case of Coulomb friction, C obviously has a much more complex structure.

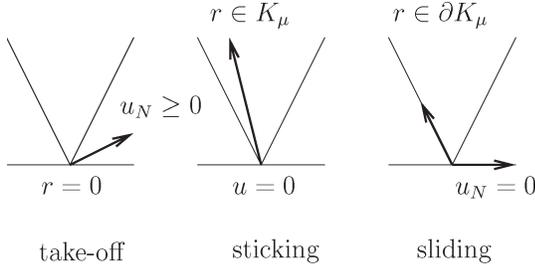


Fig. 4. The three cases of Coulomb's law.

As illustrated in Figure 4, Coulomb's law was originally formulated as a *disjunction* of three cases: that is, $(\mathbf{u}^i, \mathbf{r}^i) \in C(\mathbf{e}^i, \mu^i)$ if and only if one of the three following cases^{2,3} occurs for each $i \in 1, \dots, n$:

- take-off : $\mathbf{r}^i = \mathbf{0}$ and $u_N^i \geq 0$,
- stick : $\mathbf{r}^i \in K_{\mu^i}$ and $\mathbf{u}^i = \mathbf{0}$,
- slide : $\mathbf{r}^i \in \partial K_{\mu^i} \setminus \mathbf{0}$, $u_N^i = 0$, $\mathbf{u}_T^i \neq \mathbf{0}$ opposed to \mathbf{r}_T^i :
 $\exists \alpha^i > 0, \mathbf{r}_T^i = -\alpha^i \mathbf{u}_T^i$.

Coulomb's law can be interpreted as follows: if the first case, the normal relative velocity is nonnegative and the contact breaks (take-off case), then there can be no force between the two bodies; this models dry friction without adherence. Note that rebound is handled separately through a restitution coefficient (Section 3). The two bodies can also remain in resting contact, then the contact force can lie anywhere in K_{μ^i} . Finally, if the two bodies are sliding with respect to each other with tangent relative velocity, then the contact force must belong to the boundary of K_{μ^i} and the tangent force must be collinear to the relative velocity with the opposite direction (according to the "maximum dissipation principle" [Moreau 1988]). The value of μ^i depends on the characteristics of the contacting surfaces (from $\mu^i = 0$ for perfect contact without friction to $\mu^i = \mu_{\max}$ for rough surfaces).

This disjunctive formulation is intuitive, but not very practical because of its combinatorial nature (there are 3^n cases to check, if the system contains n contacts); in the next subsection, we give an equivalent but more tractable formulation of Coulomb's law.

2.3 Functional Formulation

In this subsection we shall omit the upper script i of contact variables, for the sake of clarity. Alart and Curnier [1991] were the first to demonstrate that it is actually possible to give a functional characterization of the Coulomb friction law, that is, to find a function \mathbf{f}^{AC} of $\mathbb{R}^3 \times \mathbb{R}^3$ in \mathbb{R}^3 such that

$$(\mathbf{u}, \mathbf{r}) \in C(\mathbf{e}, \mu) \iff \mathbf{f}^{\text{AC}}(\mathbf{u}, \mathbf{r}) = \mathbf{0}. \quad (2)$$

²Note that, unlike most approaches, we compactly formulate a frictional contact law that both includes the pure contact law (namely the *Signorini conditions* [Moreau 1988]) and the pure Coulomb friction law.

³The Signorini law classically relates contact forces to interpenetration distances. Here we model contact at a higher order, using a velocity-based formulation. Moreau proved that in the continuous case, the velocity-based formulation of constraints implies the Signorini conditions (result known as the "viability lemma") [Moreau 1988]. As we already mentioned in the Introduction, it has been experienced that in the numerical case, and when friction has to be accounted for, the velocity-based formulation conducts to more stable simulations [Stewart 2000].

The function \mathbf{f}^{AC} proposed by Alart and Curnier to satisfy this equivalence is defined as

$$\mathbf{f}^{\text{AC}}(\mathbf{u}, \mathbf{r}) := \begin{bmatrix} \mathbf{f}_N^{\text{AC}}(\mathbf{u}, \mathbf{r}) \\ \mathbf{f}_T^{\text{AC}}(\mathbf{u}, \mathbf{r}) \end{bmatrix}, \quad (3)$$

where the so-called *normal* and *tangential* parts of \mathbf{f} respectively read

$$\mathbf{f}_N^{\text{AC}}: \begin{cases} \mathbb{R}^3 \times \mathbb{R}^3 \longrightarrow \mathbb{R} \\ (\mathbf{u}, \mathbf{r}) \longmapsto \mathbf{P}_{\mathbb{R}^+}(\mathbf{r}_N - \rho_N u_N) - \mathbf{r}_N \end{cases} \quad (4a)$$

and

$$\mathbf{f}_T^{\text{AC}}: \begin{cases} \mathbb{R}^3 \times \mathbb{R}^3 \longrightarrow \mathbb{R}^2 \\ (\mathbf{u}, \mathbf{r}) \longmapsto \mathbf{P}_{\mathbf{B}(\mathbf{0}, \mu \rho_N)}(\mathbf{r}_T - \rho_T \mathbf{u}_T) - \mathbf{r}_T \end{cases}, \quad (4b)$$

where ρ_N and ρ_T are two positive constants ($\rho_N = \rho_T = 1$ for instance), $\mathbf{B}(\mathbf{0}, \delta) \subset \mathbb{R}^2$ is the ball centered in $\mathbf{0}$ of radius δ , and \mathbf{P}_K is the projection function (applying to any vector of dimension 1 or 2) onto the convex set K ($\dim K \leq 2$), with the convention that

$$\forall \mathbf{y} \in \mathbb{R}^2, \mathbf{P}_K(\mathbf{y}) = \mathbf{0}_{\mathbb{R}^2}.$$

We give the proof for equivalence (2) in Appendix A. It is fairly simple and we invite the reader to check the equivalence by himself so that he can fully understand the meaning and the power of the Alart-Curnier function with respect to Coulomb's friction law.

3. FORMULATION OF THE ONE-STEP PROBLEM

Let us consider a set of N_s strands (modeled as implicit mass-spring chains, CORDES, or superhelices). The mass-matrix \mathbf{M} of this system is diagonal-block where each block i contains the mass-matrix \mathbf{M}_i of the i^{th} strand. Let m be the total number of degrees of freedom of the system, and $\dot{\mathbf{q}} \in \mathbb{R}^m$ the generalized velocities of the mechanical system, resulting from the concatenation of the generalized velocities of the N_s individual strands. Let \mathbf{f} collect the set of (internal and external) forces applied onto the total system. \mathbf{J} denotes the gradient that relates the local velocity of the mechanical system to its global velocity, and its calculation is explained in Section 4.1.

We wish to solve the following problem in the unknowns $(\dot{\mathbf{q}}, \mathbf{u}, \mathbf{r})$, where $\mathbf{u} \in \mathbb{R}^{3n}$ gathers all the relative 3D velocities \mathbf{u}^i at contact points, and $\mathbf{r} \in \mathbb{R}^{3n}$ collects the 3D contact forces \mathbf{r}^i :

$$\begin{cases} \mathbf{M}(\mathbf{q}) \frac{d\dot{\mathbf{q}}}{dt} + \mathbf{f} = \mathbf{J}^\top \mathbf{r} \\ \mathbf{u} = \mathbf{J}\dot{\mathbf{q}} + \mathbf{w} \\ \forall i = 1 \dots n, (\mathbf{u}^i - c^i \mathbf{e}^i, \mathbf{r}^i) \in C(\mathbf{e}^i, \mu^i). \end{cases} \quad (5)$$

The set $C(\mathbf{e}^i, \mu^i)$ contains the couples $(\mathbf{u}^i, \mathbf{r}^i)$ satisfying Coulomb's law for a given value of the friction coefficient μ^i and a normal vector \mathbf{e}^i at the i^{th} contact point. The constant $c^i \geq 0$ accounts for impacts: after a shock, Coulomb's law imposes that the normal part of $\mathbf{u}^i - c^i \mathbf{e}^i$ be zero, which implies that the norm of the normal relative velocity u_N^i takes the value c^i . By setting c^i equal to a constant k (with $k \in [0, 1]$) times the normal velocity *before* the shock, we obtain Newton's impact law with $k = 0$ modeling inelastic impacts and $k = 1$ modelling perfect rebound. In practice we model impacts between thin strands as merely inelastic, thus choosing $c^i = 0$. For impacts between a heavy rod and a rigid body, however, it may be suitable to set c^i to a small nonzero value to allow for some rebound.

Solving the initial problem (5) amounts to solving

$$\begin{cases} \mathbf{M}(\mathbf{q}) \frac{d\dot{\mathbf{q}}}{dt} - \mathbf{J}^\top \mathbf{r} + \mathbf{f} = \mathbf{0} \\ \mathbf{f}^{\text{AC}}(\mathbf{J}\dot{\mathbf{q}} + \mathbf{w} - \mathbf{c}_e, \mathbf{r}) = \mathbf{0} \end{cases} \quad (6)$$

where \mathbf{c}_e is a constant vector containing the terms $c^i \mathbf{e}^i$ from (5) and \mathbf{f}^{AC} collects the Alart-Curnier function (3) at all contact points. Using an arbitrary step length δt , we follow Moreau's time-stepping scheme for integrating nonsmooth problems [Moreau 1988]: the velocity derivative $\frac{d\mathbf{q}}{dt}$ is replaced with the velocity jump between the end and the start of the timestep, the position derivative $\dot{\mathbf{q}}$ is approximated at first order, \mathbf{r} and \mathbf{u} are discretized using an *implicit* scheme (since brutal changes in their values are expected due to impacts), and the other terms⁴ are discretized explicitly (since they are assumed to be smooth). We denote by $(\mathbf{v}, \underline{\lambda})$ the discrete-time approximations of $(\dot{\mathbf{q}}, \mathbf{r})$, where $\underline{\lambda}$ is the discrete impulse $\mathbf{r}_{t+\delta t} - \mathbf{r}_t$, and by $\underline{\mathbf{M}}, \underline{\mathbf{J}}, \underline{\mathbf{f}}$ and $\underline{\mathbf{w}}$ (accounting for $\mathbf{w} - \mathbf{c}_e$) the discretized data, where $\underline{\mathbf{f}} = \mathbf{f}_t \delta t - \underline{\mathbf{M}} \mathbf{v}_t$. We thus get the following one-step problem in the unknowns $(\mathbf{v}, \underline{\lambda})$:

$$\begin{cases} \underline{\mathbf{M}}\mathbf{v} - \underline{\mathbf{J}}^T \underline{\lambda} + \underline{\mathbf{f}} = \mathbf{0} \\ \mathbf{f}^{\text{AC}}(\underline{\mathbf{J}}\mathbf{v} + \underline{\mathbf{w}}, \underline{\lambda}) = \mathbf{0}. \end{cases} \quad (7)$$

Our method consists in searching a solution to system (7) using Newton's algorithm (despite the fact that \mathbf{f}^{AC} is nonsmooth). We have formulated our initial problem in terms of impulses and velocities instead of forces and displacements. It is well known that the impulse/velocity formulation is the only way to consistently integrate time systems with impacts. The price to pay is a formulation of unilateral constraints and Coulomb friction on the impulse level, which means with a time integral of forces when the evolution is smooth enough. For more details on the slight differences between an impulse model of friction and the standard one in forces, we refer to Frémond [2002] and Acary and Brogliato [2008].

4. COMPUTATION OF THE CONTACT FORCES

At each time step, we wish to solve the system of equations (7) in the unknowns $(\mathbf{v}, \underline{\lambda})$. We first briefly explain the computation of $\underline{\mathbf{J}}$ before reducing the system as a zero finding problem of the function \mathbf{f}^{AC} . The solving by Newton's method requires the computation of the Jacobian matrix of \mathbf{f}^{AC} , which is described in detail in the following. The final algorithm for solving the frictional contacting problem, together with interpenetration correction and optimization, is finally presented.

4.1 Computation of $\underline{\mathbf{J}}$

As expressed in Eq. (5), the gradient matrix $\underline{\mathbf{J}}$ of size $(3n, m)$ linearly relates the 3D relative velocities \mathbf{u} at the n contact points to the m generalized velocities $\dot{\mathbf{q}}$ of the system. Let us consider the 3D row-block $\underline{\mathbf{J}}_i$ corresponding to the i^{th} contact point. As mentioned in Section 2.1, we assume that each contact involves no more than two bodies (or body parts), denoted A^i and B^i . Let $\mathbf{r}^{A^i}(s^{A^i}, \mathbf{q}^{A^i})$ ($\mathbf{r}^{B^i}(s^{B^i}, \mathbf{q}^{B^i})$) respectively be the centerline of the rod corresponding to body A^i (respectively, to body B^i): it gives the spatial coordinates of the point located at the curvilinear abscissa s^{A^i} (respectively s^{B^i}) on the rod A^i (respectively, B^i) with generalized-coordinates \mathbf{q}^{A^i} (respectively, \mathbf{q}^{B^i}). At instant t_c , $\underline{\mathbf{J}}_i$ thus reads

$$\underline{\mathbf{J}}_i(t_c) = \frac{\partial \mathbf{r}^{A^i}}{\partial \mathbf{q}^{A^i}}(s^{A^i}, \mathbf{q}^{A^i}(t_c)) - \frac{\partial \mathbf{r}^{B^i}}{\partial \mathbf{q}^{B^i}}(s^{B^i}, \mathbf{q}^{B^i}(t_c)),$$

⁴Actually, in the case of superhelices and implicit mass-spring chains, internal elastic forces are computed implicitly. The actual matrix $\underline{\mathbf{M}}$ of our discretized system (7) is thus $\underline{\mathbf{M}} + dt^2 \underline{\mathbf{K}}$, where $\underline{\mathbf{K}}$ is the total stiffness matrix of the system. We skip this detail in the description of the algorithm, for the sake of simplicity.

where $s_c^{A^i}$ and $s_c^{B^i}$ respectively denote the curvilinear abscissa at contact point for rod A^i and rod B^i . Note that in the case of self-contact, we have $A^i = B^i$ and $s_c^{A^i} \neq s_c^{B^i}$.

The term $\underline{\mathbf{J}}_i$ can thus be computed by calculating the formal gradient $\frac{\partial \mathbf{r}}{\partial \mathbf{q}}$ of a rod with centerline \mathbf{r} and generalized-coordinates \mathbf{q} , and evaluating it at the point of current contact i . Such a computation is straightforward for the mass-spring chain and the CORDE model. In the case of the superhelix model, we relied on the Maple software [MapleSoft 2010] for deriving the gradient analytically.

4.2 Elimination of \mathbf{v}

The problem (7) can be reduced by eliminating \mathbf{v} from the equations. Indeed, from (7) there holds

$$\mathbf{v} = \underline{\mathbf{M}}^{-1} \underline{\mathbf{J}}^T \underline{\lambda} - \underline{\mathbf{M}}^{-1} \underline{\mathbf{f}}$$

which implies

$$\underline{\mathbf{u}} = \underbrace{\underline{\mathbf{J}} \underline{\mathbf{M}}^{-1} \underline{\mathbf{J}}^T}_{=\underline{\mathbf{W}}} \underline{\lambda} + \underbrace{\underline{\mathbf{w}} - \underline{\mathbf{J}} \underline{\mathbf{M}}^{-1} \underline{\mathbf{f}}}_{=\underline{\mathbf{b}}} = \underline{\mathbf{W}} \underline{\lambda} + \underline{\mathbf{b}},$$

where $\underline{\mathbf{W}}$ is the so-called Delassus operator. Its computation involves $\underline{\mathbf{M}}^{-1}$: it can be performed column by column using a conjugate gradient algorithm (then $\underline{\mathbf{M}}$ and $\underline{\mathbf{J}}$ need not be assembled), or using the Choleski decomposition of $\underline{\mathbf{M}}$ (which can be reused from a time-step to another if the mass matrix does not change).

After eliminating \mathbf{v} , the problem becomes

$$\mathbf{f}^{\text{AC}}(\underline{\mathbf{W}} \underline{\lambda} + \underline{\mathbf{b}}, \underline{\lambda}) = \mathbf{0}. \quad (8)$$

On the one hand, this new system is smaller than system (7) (size $3n$ instead of $m + 3n$) and we exploit the fact that the linear part of system (7) does not change during Newton's iterations by computing $\underline{\mathbf{W}}$ once for all. On the other hand the computation of $\underline{\mathbf{W}}$ may be costly and $\underline{\mathbf{W}}$ may not be as sparse as system (7). The choice of performing this elimination or not is probably problem dependent. In our experiments, we chose to use it since $\underline{\mathbf{W}}$ was sparse enough. Indeed, despite the fact that the mass matrix and the gradient $\underline{\mathbf{J}}$ of individual rods are dense, they are sparse and block-structured for the whole system due to the connectivity of the contacts and the bodies (see Section 4.5).

4.3 Computing the Gradient of \mathbf{f}^{AC}

The normal direction \mathbf{e}^i at i^{th} contact point being known, choose \mathbf{e}_s^i and \mathbf{e}_t^i so that $(\mathbf{e}^i, \mathbf{e}_s^i, \mathbf{e}_t^i)$ form an orthonormal basis. Define the matrix of the projection onto the normal direction by $\mathbf{P}_N := (\mathbf{e}^i)^T$ and the matrix of the projection onto the tangent direction by $\mathbf{P}_T := [\mathbf{e}_s^i, \mathbf{e}_t^i]^T$. We are now able to differentiate \mathbf{f}^{AC} . For the sake of simplicity, detailed calculations that follow are given in the case of a single contact.

Normal part. Function f_N^{AC} defined by (4a) is continuous and affine by pieces : if $r_N - \rho_N u_N < 0$, then

$$f_N^{\text{AC}}(\mathbf{u}, \mathbf{r}) = -\mathbf{P}_N \mathbf{r}, \quad \frac{\partial f_N^{\text{AC}}}{\partial \mathbf{u}} = \mathbf{0}_{1 \times 3}, \quad \frac{\partial f_N^{\text{AC}}}{\partial \mathbf{r}} = -\mathbf{P}_N$$

whereas, if $r_N - \rho_N u_N > 0$

$$f_N^{\text{AC}}(\mathbf{u}, \mathbf{r}) = -\rho_N \mathbf{P}_N \mathbf{u}, \quad \frac{\partial f_N^{\text{AC}}}{\partial \mathbf{u}} = -\rho_N \mathbf{P}_N, \quad \frac{\partial f_N^{\text{AC}}}{\partial \mathbf{r}} = \mathbf{0}_{1 \times 3}.$$

Tangential part. Define the following function g whose Jacobian matrix, where it exists, is computed in Appendix B:

$$\mathbf{g}: \begin{cases} \mathbb{R} \times \mathbb{R}^2 \longrightarrow \mathbb{R}^2 \\ (\delta, \mathbf{y}) \longmapsto \mathbf{P}_{B(0, \delta)}(\mathbf{y}). \end{cases}$$

Then the tangential part \mathbf{f}_T^{AC} of the Alart-Curnier function (4b) is

$$\mathbf{f}_T^{\text{AC}}(\mathbf{u}, \mathbf{r}) = g(\underbrace{\mu \mathbf{P}_N \mathbf{r}}_{=: \delta}, \underbrace{\mathbf{P}_T \mathbf{r} - \rho_T \mathbf{P}_T \mathbf{u}}_{=: \mathbf{y}}) - \mathbf{P}_T \mathbf{r}.$$

Differentiating \mathbf{f}_T^{AC} using the chain rule, we obtain

$$\begin{cases} \frac{\partial \mathbf{f}_T}{\partial \mathbf{u}} = \frac{\partial g}{\partial \delta} \frac{\partial \delta}{\partial \mathbf{u}} + \frac{\partial g}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \\ \frac{\partial \mathbf{f}_T}{\partial \mathbf{r}} = \frac{\partial g}{\partial \delta} \frac{\partial \delta}{\partial \mathbf{r}} + \frac{\partial g}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \mathbf{r}} - \mathbf{P}_T \end{cases}$$

that is to say, using the formulas of Appendix B:

—If $\|\mathbf{y}\| \leq \delta$, then

$$\begin{cases} \mathbf{f}_T^{\text{AC}}(\mathbf{u}, \mathbf{r}) = -\rho_T \mathbf{P}_T \mathbf{u}, \\ \frac{\partial \mathbf{f}_T^{\text{AC}}}{\partial \mathbf{u}}(\mathbf{u}, \mathbf{r}) = -\rho_T \mathbf{P}_T, \\ \frac{\partial \mathbf{f}_T^{\text{AC}}}{\partial \mathbf{r}}(\mathbf{u}, \mathbf{r}) = \mathbf{0}_{2 \times 3}. \end{cases}$$

—If $\|\mathbf{y}\| > \delta > 0$, then

$$\begin{cases} \mathbf{f}_T^{\text{AC}}(\mathbf{u}, \mathbf{r}) = \lambda \frac{\mathbf{y}}{\|\mathbf{y}\|} - \mathbf{P}_T \mathbf{r}, \\ \frac{\partial \mathbf{f}_T^{\text{AC}}}{\partial \mathbf{u}}(\mathbf{u}, \mathbf{r}) = -\rho_T \frac{\delta}{\|\mathbf{y}\|} (I_2 - \frac{\mathbf{y} \mathbf{y}^T}{\|\mathbf{y}\|^2}) \mathbf{P}_T, \\ \frac{\partial \mathbf{f}_T^{\text{AC}}}{\partial \mathbf{r}}(\mathbf{u}, \mathbf{r}) = \mu \frac{\mathbf{y}}{\|\mathbf{y}\|} \mathbf{P}_N + \frac{\delta}{\|\mathbf{y}\|} (I_2 - \frac{\mathbf{y} \mathbf{y}^T}{\|\mathbf{y}\|^2}) \mathbf{P}_T - \mathbf{P}_T. \end{cases}$$

—If $\|\mathbf{y}\| \geq 0 \geq \delta$, then

$$\begin{cases} \mathbf{f}_T^{\text{AC}}(\mathbf{u}, \mathbf{r}) = -\mathbf{P}_T \mathbf{r}, \\ \frac{\partial \mathbf{f}_T^{\text{AC}}}{\partial \mathbf{u}}(\mathbf{u}, \mathbf{r}) = \mathbf{0}_{2 \times 3}, \\ \frac{\partial \mathbf{f}_T^{\text{AC}}}{\partial \mathbf{r}}(\mathbf{u}, \mathbf{r}) = -\mathbf{P}_T. \end{cases}$$

In limit cases, such as $\|\mathbf{y}\| = \delta$, the computation of the Jacobian matrix does not make sense. The notion of generalized gradients has to be introduced, which numerically amounts to arbitrarily choosing one of the two possible formulas. For the normal part, for instance, one of the condition $r_N - \rho_N u_N > 0$ (or $r_N - \rho_N u_N > 0$) includes the equality case $r_N - \rho_N u_N = 0$. For more details on the computation of the generalized gradient and the theoretical aspects of generalized Newton's method, we refer to Qi and Sun [1993]. The formulation of the function \mathbf{f}^{AC} slightly differs from the original one presented in Alart and Curnier [1991]. In particular, the tangential part \mathbf{f}_T^{AC} uses the projection on the ball $\mathbf{B}(\mathbf{0}, \mu r_N)$ rather than the projection on the modified ball $\mathbf{B}(\mathbf{0}, \mathbf{P}_N(r_N - \rho_N u_N))$. In Christensen et al. [1998], the same modification is done to simplify the computation of the function and its gradients; according to our experience, this variant does not change the behavior of the Newton method.

4.4 Newton Step Computation, Line Search and Convergence

Once the computation of the Alart-Curnier function \mathbf{f}^{AC} and its Jacobian matrix is implemented, the algorithm is simple: the unknowns are initialized (in the case of nonvanishing contacts, we use the values of the previous time-step), and the Newton step is computed by solving the linearization of system (7). A simple Goldstein-Price line search [Bonnans et al. 2003] is then performed in the Newton direction to ensure that the least-square criterion $\|\mathbf{f}^{\text{AC}}\|^2$ decreases. We loop until $\|\mathbf{f}^{\text{AC}}\|^2$ goes below a given tolerance or, alternatively, until the maximum number of iterations is reached. The nonsmooth Newton method and the Goldstein-Price line search are respectively described with pseudocode in Algorithms 1 and 2. An important feature of our algorithm, compared

Algorithm 1. Nonsmooth Newton Method

Require: Initial guess \mathbf{r}_0

Require: tol tolerance, itermax maximum number of iteration

Ensure: Solution of $\mathbf{f}^{\text{AC}} = \mathbf{0}$

```

/* Initialization */
k ← 0
rk ← r0
/* Start Newton loop */
while  $\frac{1}{2} \|\mathbf{f}^{\text{AC}}(\mathbf{r}_k)\|^2 > \text{tol}$  and  $k \leq \text{itermax}$  do
  /* Compute  $\mathbf{f}^{\text{AC}}$  and choose a sub-gradient */
  uk ← Wrk + b
  bk ←  $\mathbf{f}^{\text{AC}}(\mathbf{u}_k, \mathbf{r}_k)$ 
  Ak ←  $\frac{\partial \mathbf{f}^{\text{AC}}}{\partial \mathbf{u}}(\mathbf{u}_k, \mathbf{r}_k) \mathbf{W} + \frac{\partial \mathbf{f}^{\text{AC}}}{\partial \mathbf{r}}(\mathbf{u}_k, \mathbf{r}_k)$ 
  /* Compute Newton's direction */
  Solve Akdk = bk
  /* Compute length with a line search procedure. See Algorithm 2. */
  Compute αk.
  /* Update rk */
  rk ← rk + αkdk
end while

```

Algorithm 2. Goldstein-Price line search procedure

Require: α_{init} initial value of line length

Require: $[\alpha_l, \alpha_u]$ initial interval for α

Require: α_{max} maximum value of α

Require: $m_1 = 0.1, m_2 = 0.9$ user defined parameters

Ensure: Optimal value of α

```

α ← αmin
/* Compute the initial value of the merit function */
q0 ←  $\frac{1}{2} \|\mathbf{f}^{\text{AC}}(\mathbf{u}_k, \mathbf{r}_k)\|^2$ 
/* Compute the initial value of the slope */
q'0 ←  $(\mathbf{f}^{\text{AC}}(\mathbf{u}_k, \mathbf{r}_k))^T \frac{\partial \mathbf{f}^{\text{AC}}}{\partial \mathbf{r}}(\mathbf{u}_k, \mathbf{r}_k) \mathbf{d}_k$ 
success ← FALSE
while success = FALSE do
  rtrial ← rk + αdk
  utrial ← Wrtrial + b
  qk ←  $\mathbf{f}^{\text{AC}}(\mathbf{u}_{\text{trial}}, \mathbf{r}_{\text{trial}})$ 
  slope ←  $\frac{q - q_0}{\alpha}$ 
  if (slope ≥ m2q'0 and slope ≤ m1q'0) then
    success ← TRUE
  else
    if slope < m2q'0 then
      αl ← α
    else
      αu ← α
    end if
    if αu < αmax then
      α ←  $\frac{1}{2}(\alpha_l + \alpha_u)$ 
    else
      α ← 10α
    end if
  end if
end while

```

to some previous works limited to the search for stationary points of optimization problems (not necessarily solutions of the frictional contact problem) is that our convergence error $\|\mathbf{f}^{\text{AC}}\|^2$ is a measure of the distance to an *actual* solution of our initial problem (7). Indeed, recall that cancelling \mathbf{f}^{AC} is strictly equivalent to solving problem (7).

4.4.1 *Convergence, in Theory.* Since $\|\mathbf{f}^{\text{AC}}\|^2$ decreases along the iterations, it eventually converges (unless the maximum number of iterations is reached), but not necessarily to a solution: $\|\mathbf{f}^{\text{AC}}\|^2$ may converge to a value above the given tolerance. Some further hints on the theoretical expectations of convergence are given in Cadoux’s Ph.D. thesis [2009]. On the one hand, the dimension $nd \times m$ of matrix \mathbf{J} (where n is the total number of contacts, m is the total number of degrees of freedom of the system, and d is the dimension of space, typically $d = 3$) is deeply linked to the existence of a solution to the one-step problem (7): when $nd > m$, \mathbf{J} is not surjective anymore, which does not guarantee the existence of a solution. On the other hand, the non-surjectivity of \mathbf{J} implies \mathbf{W} ’s singularity, and consequently, is likely to make the problem harder to solve. For more details, we refer the reader to Cadoux [2009] and Acary et al. [2010].

4.4.2 *Convergence, in Practice.* In our experiments presented in Section 5, we shall see that the fraction

$$\nu = \frac{nd}{m}$$

plays an important role with respect to the quality and rate of convergence. In practice, the range of simulations that we are interested in mostly fall into the favorable case, namely the case when the criterion $\nu < 1$ is roughly satisfied. Provided the time-step is chosen small enough (to overcome the difficulty when \mathbf{f}^{AC} is nonsmooth close to the current iterate), convergence with reasonable speed will then be observed. Actually, ν can be viewed as a conditioning number for our simulations.

4.5 Collision Detection and Constraints Partitioning

This article focuses on the response due to impact and resting contact, and not on collision detection. However, in order to optimize the solving of the system resulting from the active constraints, we provide a constraints partitioning technique, described next.

The geometry of a rod is approximated by a set of bounding cylinders and for each pair of objects subject to collide we track the pairs of closest points that lie on the axes of the corresponding boundary cylinders over time. One contact is declared active when two distinct boundary cylinders are found to intersect at a given time step, that is, once the distance between the two closest points is below the diameter of the rod. In practice we used around 1 cm-long cylinders, that is, 30 cylinders for a typical strand of 30 cm simulated in most of our examples. This simple algorithm proved sufficiently robust for our tests, even when using large time steps. For very high-speed motion, however, it may be preferable to use continuous detection algorithms [Redon et al. 2002] in order to capture every single occurring collision.

To save computing time, we partition the set of active constraints and the set of strands into $N_{\mathcal{P}}$ groups of *coupled* constraints and strands at each time step. This allows us to solve $N_{\mathcal{P}}$ independent systems of the form (7), instead of one single large system accounting for all strands of the scene. To determine the partition \mathcal{P} , let us represent the sets of constraints and strands as a graph, where nodes stand for strands and edges for constraints. An example is presented in Figure 5. Two different strands (respectively, constraints) are said to be coupled if there exists a constraints path in the graph (respectively, a strands path in the dual graph) connecting their corresponding nodes (respectively, their corresponding nodes in the dual graph). Partitioning our initial set of constraints thus amounts to partitioning a graph into connected sets, which is achieved efficiently.

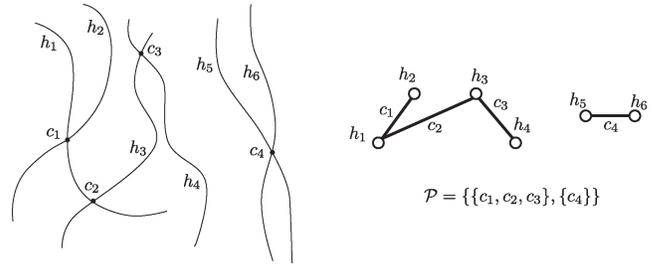


Fig. 5. Example of constraints partitioning using graph connectivity. Left: contacting strands; right: corresponding graph and partition ($N_{\mathcal{P}} = 2$).

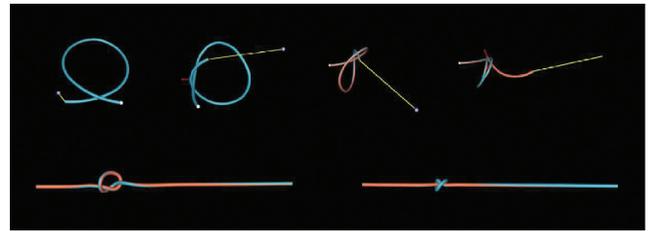


Fig. 6. Interactive making of a knot from a superhelix of 20 elements ($\mu = 0$). With our method the user can tighten the knot as strongly as he wants, without causing contact to break.

Finally, we have accelerated the search for contacting pairs by employing a spatial hashing of the 3D space with a uniform grid, similarly to Teschner et al. [2003]. In our simulation, collision detection took less than 10% of the total computational time, and thus never appeared to be the computational bottleneck in our simulations. See Section 5.3 for the detailed performance.

5. RESULTS

In this section, we accurately analyze our method in terms of (1) *realism* (i.e., its ability to capture relevant emerging phenomena), (2) *numerical stability* (i.e., its convergence properties), and (3) *computational efficiency* (i.e., its time performance). Limitations and discussion are provided in the next section.

5.1 Realism: Capture of Relevant Physical Properties

5.1.1 *Contact Robustness.* Figures 2 (left), 6, and the accompanying video (accessible at the ACM Digital Library) demonstrate the ability of our method to address the challenging scenarios of plectoneme formation and tight knot tying. Because of the slender geometry of a thin rod, simulating self-contact and especially the formation of plectonemes or tight knots requires an extremely robust algorithm for contact. Unlike basic penalty-based approaches which fail to prevent interpenetration beyond a certain force/torque magnitude and thus may cause the self-crossing of the rod, our constraint-based algorithm robustly maintains contact whatever the amplitude of the force/torque applied onto the free end of the rod. With our method, the violation of constraints is prevented whatever the magnitude of the applied forces.

5.1.2 *Capture of Dynamic and Static Friction.* The accompanying video shows the effect of increasing the coefficient of friction μ on the simple example of a strand contacting a rotating sphere. When μ is large enough, dry friction is captured as expected. In the

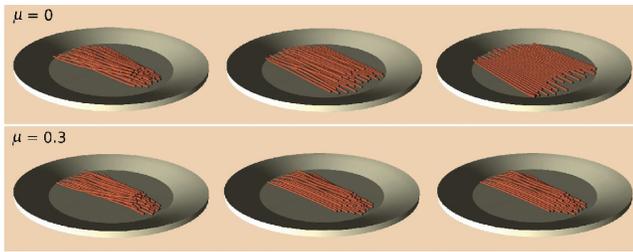


Fig. 7. Dense packing of 31 spaghetti in a plate. Top: without friction ($\mu = 0$), the rods just spread out over the plate. Bottom: in the presence of friction ($\mu = 0.3$), the rods stack in a stable way.

knot example, the user can easily perceive the influence of μ when virtually tying a knot: in the absence of friction this task is actually challenging since contact points are always sliding and thus cannot guide the user. With a nonzero friction coefficient, however, static friction occurs and greatly helps the user as it maintains some fixed points of contact that can be used as pivots.

The typical stick-slip instabilities occurring inside fiber assemblies are illustrated in the following examples.

5.1.3 Packing of Fibers. We have simulated the fall of a dense packing of 31 parallel spaghetti⁵ (modeled as mass-spring chains with both ends free) on to a plate. Figure 7 and the video show the effects of varying the friction coefficient. In the case where there is no friction, the rods spread over the plate. When friction is activated between the rods and the plate, but not between the rods, the spaghetti form a stack but continue to slide against each other. When friction is active everywhere, the stack is higher and remains perfectly still.

In a second experiment depicted in the video, we have simulated a bunch of 7 spaghetti as CORDE rods with a straight shape at rest, but initialized with a nonzero curvature. After impact, the rods tend to recover their natural straight shape when there is no friction, whereas they remain curved in the frictional case as tangential friction forces compensate for elastic forces.

5.1.4 Complex Interactions in Hair Dynamics. We have tested our method to resolve interactions in hair dynamics in the presence of friction. Figures 2 (right), 8, and the accompanying video show a wavy head of hair under motion where thin wisps are modeled as 84 contacting superhelices, in the presence of friction ($\mu = 0.3$) or not ($\mu = 0$). Compared to previous approaches, our method is the first to properly capture subtle hair contacts occurring between hair curls, thus correctly preserving the hair volume as well as capturing sticking and coherence of motion in the presence of friction. In Figure 8, observe how hair wisps get entangled, forming complex and discontinuous patterns. On the right, note how friction plays a major role in the realism of motion: without friction, hair wisps just slide against each other and against the body, whereas in the presence of friction, typical stick-slip instabilities (especially visible on the shoulder of the character and on top of the head) occur.

5.1.5 Contact in Rigid Bodies. In order to explore the versatility of our method we have attempted to solve frictional contact within other dynamic systems, such as rigid bodies. Figure 9 illustrates the loop benchmark (introduced by Kaufman et al. [2005]),

⁵As a matter of interest, modeling and animating spaghetti was amusingly listed by Blinn [1998] as one of the top 10 unsolved problems in computer graphics.

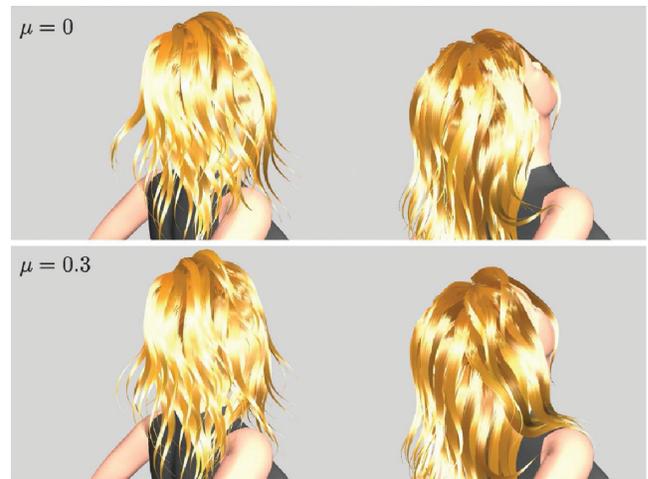


Fig. 8. Simulating frictional contact in a full head of hair (84 simulated strands). Top: without friction ($\mu = 0$). Bottom with friction ($\mu = 0.3$). On every picture, observe the retrieval of a realistic hair volume as well as the presence of discontinuous patterns caused by the proper solving of self-contacts. Left: during motion, hair subject to friction (bottom) exhibits less volume than the frictionless case (top) due to an enhanced coherence of motion. Right: at rest, when friction is applied (bottom), some hair wisps remain stuck over the shoulder thanks to the correct modeling of dry friction, whereas unnatural sliding occurs in the frictionless case (top). Observe the hair wisps that rest on top of another hair layer, similarly as in Figure 1. See also the accompanying video.



Fig. 9. Real-time simulation of a ball in a double loop. The ball initially has some spin, and different friction coefficients between the ball and the rails are tested. When friction is high, the spin energy is converted into translational kinetic energy, which makes the ball roll through the loops.

extended to a double loop, where a ball rolls over two rails ($n = 2$) with a curly shape. The ball initially has some spin, but no translational velocity. Increasing the friction coefficient reduces sliding phases in favor of pure rolling (adherence), which causes the ball to go further through the looping rails (see the accompanying video).

In this simple example, convergence is very fast (less than 5 iterations) even when using a large time-step (11 ms). However, as we shall see in Section 6.2, our method may not be adapted to handle a set of rigid bodies with stacking as the conditioning of the method will rapidly decrease due to the large number of contacts compared to the low number of degrees of freedom of the system (see also next section).

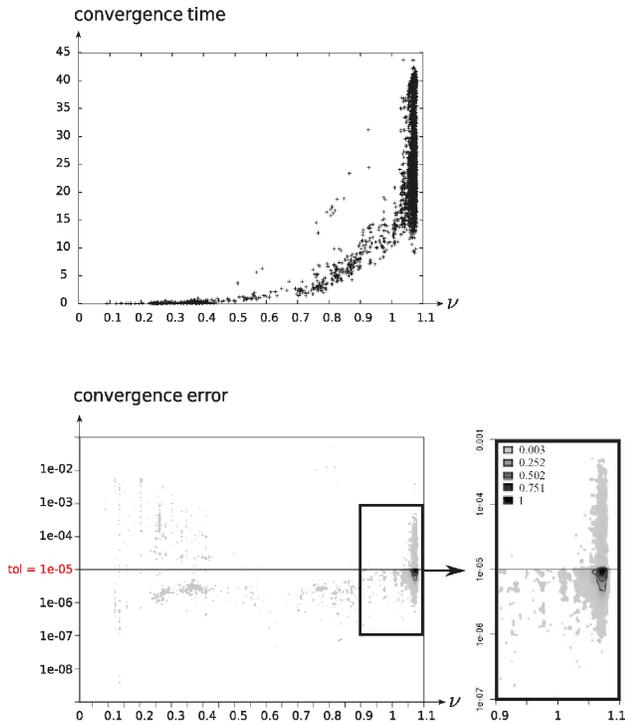


Fig. 10. Analysis of convergence for the 76 spaghetti demo. Top: convergence time of the solver (in seconds), as a function of ν . Bottom: relative density of the convergence error of the solver (measured as $\|f^{AC}\|^2$), as a function of ν . The tolerance error is set to 10^{-5} in our internal, dimensionless units. The dense zone inside the rectangle (in which most error values concentrate) is zoomed on the right to better depict the density distribution of error values (represented as greyscale values). Note that most error values lie beneath the tolerance.

5.2 Stability: A Criterion for Fast Convergence

We have carefully analyzed the convergence properties of our algorithm on a spaghetti simulation consisting of 76 flexible rods (modeled as mass-spring chains composed of 22 nodes with both ends free) subject to external contacts (with a plate) as well as self-contacts, in the presence of friction ($\mu = 0.3$). After a series of impacts, the assembly of rods rapidly stabilizes as expected, forming a stack subject to more than 1800 contacts (see Figure 2, right, and the accompanying video).

To evaluate the quality of convergence of the solver, we have plotted the time required for the solver to converge as a function of the number $\nu = \frac{nd}{m}$ introduced in Section 4.4. In this particular case, $d = 3$ and $m = 76 \times 22 \times 3 = 5016$, thus $\nu = \frac{n}{1672}$, where n is the total number of contacts. Figure 10, top, gives the results of this plot. It clearly shows the influence of ν on the rate of convergence and identifies that the value $\nu = 1$ critically defines a threshold for proper convergence: when the criterion $\nu < 1$ is satisfied, convergence remains reasonably fast, whereas it grows exponentially as soon as $\nu > 1$. In practice, we have noted that the solver almost always manages to converge until ν reaches a value close to 1.5; however, the time taken to converge grows larger as ν increases. Figure 10, bottom, displays the relative density of convergence error as a function of ν . It is noticeable that the solver converges in most cases to the desired precision (here, fixed to 10^{-5}), even when ν is slightly superior to 1. Convergence occurred in roughly 100 iterations on

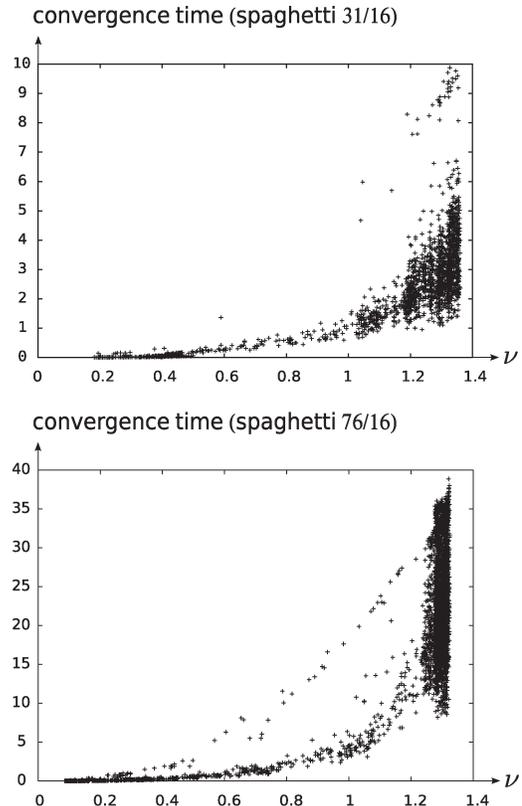


Fig. 11. Comparison of the solver convergence time profile (in seconds) between 2 different simulations: 31 mass-springs of 16 nodes (top) vs. 76 mass-springs of 16 nodes (bottom). Note that the profiles are similar up to a scaling factor directly related to the mean number of contacts.

average (the maximum number of iterations was set to 200). In 10% of the cases we studied, the solver did not strictly reach the required precision. Note, however, that the solver *always* converged to an acceptable solution, close to the target value 0 (the maximum recorded error was only of 10^{-2}). In the $\nu < 1$ zone, we noticed that the solver happened to converge to such approximate solutions in situations involving large impacts (typically, when the spaghetti collide with the plate at maximum speed). We strongly suspect this issue may be due to the lack of a proper warm start in such cases, as contact forces and postimpact velocities are initialized with zero values when new contacts are set active. The solver thus starts from a configuration that is far from the solution when impacts are large, causing the Newton algorithm to fall into some local minima. In the near future, we plan to investigate how to properly estimate an initial guess of the solution in the specific case of impacts, based on the knowledge of relative velocities before impact. We note, however, that in all cases, such an approximate convergence did not produce any visually disturbing effects nor did it cause the solver to fail in the following time steps.

We have plotted similar data for other rods experiments, and observed exactly the same properties for the convergence of the solver, as depicted in Figure 11: the solver quickly converges to a solution when $\nu < 1$, whereas convergence still occurs for ν slightly superior to 1 (but is slowed down), and finally the solver fails to converge when the criterion is far from being satisfied (typically, when $\nu > 1.5$). This result confirms the observation previously made by Cadoux [2009] in the context of simpler experiments with

Table I. Performance Results for Our Simulations on Various Models Composed of N_s Rods and N_e Elements (or nodes) per Rod

Rod model	Example (N_s, N_e)	μ	dt (ms)	n max	ν max	mean time (s)
SH	Plectoneme (1, 20)	0	11	18	0.9	0.04
SH	Tight knot (1, 20)	0	11	8	0.4	0.033
SH	Wavy hair (84, 8)	0	11	392	0.58	2.6
SH	Wavy hair (84, 8)	0.3	11	380	0.57	41.2
C	Self-stacking (1, 256)	0.3	0.25	303	1.18	0.31
C	Spaghetti (7, 16)	0.3	0.1	121	1.08	0.36
MS	Spaghetti (31, 16)	0.3	1	653	1.32	2.26
MS	Spaghetti (76, 22)	0.3	1	1806	1.08	19.1

SH: Super-helix; C: CORDE; MS: Mass-spring.

a lower number of contacts, and enhances the fact that ν plays the role of a generic condition number.

In addition, by comparing the results obtained on different simulations, we have observed that for a constant ν , the time of convergence grows as the number n of contacts increases. Actually, the convergence time profile as a function of ν observed for two different simulations appears to be the same up to a scaling factor related to the mean number of contacts (see Figure 11).

Finally, although our problem is nonsmooth, we observe typical Newton quadratic convergence in some favorable cases where a high precision ($\|f^{AC}\|^2 < 10^{-18}$) is reached by the solver. This is, for example, the case for the simple simulation presented in Figure 9, where the solver reaches a precision of 10^{-35} in less than 5 iterations, at each time-step! This observation is actually in line with the theoretical study made by Qi and Sun [1993] in the nonsmooth case (though they make more restrictive assumptions than we do).

5.3 Performance

Although time performance was not our primary goal, it turns out that our method is fast enough to be used for handling hundreds of frictional contacts within a few minutes per frame. This is mainly due to our sparse implementation of the algorithm, which is all the more valuable as the number of contacting strands is high. However, computational time quickly becomes the main bottleneck of our simulations, which currently limits the use of our technique to systems subject to a few thousands contacts.

All the simulations presented in the article as well as in the accompanying video were processed on a single threaded application running on an Intel Core 2 Duo CPU at 2.1 GHz. Table I gives a summary of our computational results. The indicated mean computational time is based on computations during one elementary time-step, including the dynamics of thin rods, collision detection, and collision response. For large systems ($n > 100$) the computational time taken by our solver represents 80% to 90% of the total computational time, depending on the rod model used.

For every simulation, we used a tolerance on $\|f^{AC}\|^2$ of 10^{-5} in dimensionless units (i.e., physical quantities were divided by typical values of the rods system, in order to bring matrices coefficients close to 1), which in practice corresponds to a good trade-off between visual accuracy and computational cost.⁶ We chose the time step between 0.1 and 11 ms, depending on the rod model used.

⁶Alternatively, a relative nonsquare error measure such as $\epsilon = \frac{\|f^{AC}\|}{1+\|f\|}$ could be used. In our experiments, we have measured that our 10^{-5} tolerance is equivalent to a maximal relative error of $\epsilon = 10^{-3}$.

Typically, the CORDE model, based on an explicit integration scheme, required a smaller time step to remain stable, compared to other rods models.

6. DISCUSSION

In this section, we provide comparisons with previous methods and discuss the strengths and current limitations of our approach.

6.1 Comparison with Previous Methods

6.1.1 Smooth Friction Models. Our results clearly demonstrate the advantage of using a nonsmooth solver. Dry friction is indeed a predominant phenomenon in contacting rods systems. Classical methods based on viscous friction would typically not be able to recover the stick-slip effects depicted by our simulations. Hair is a crucial example where previous models (e.g. Plante et al. [2001], Bertails et al. [2006], Hadap [2006], Selle et al. [2008], and McAdams et al. [2009]) fail in capturing this typical, nonsmooth behavior.

6.1.2 Ad Hoc, Visually Acceptable Friction Models. Some approaches have attempted to modify Coulomb's friction model so as to stick to an easily solvable contact formulation [Milenkovic and Schmidl 2001; Kaufman et al. 2005]. In Milenkovic and Schmidl [2001] Coulomb's friction model is adapted to fit into the Quadratic Program (QP) framework for resolving nonpenetration. Frictional impacts are treated by restricting contact impulses to the linearized Coulomb friction cone, while resting contact is modelled by enforcing the relative acceleration to lie in the dual linearized friction cone and by subsequently projecting boundary-located accelerations onto the tangent space. While the former formulation is a reasonable approximation of Coulomb's model at impacts (see Section 6.1.3), the latter is more questionable: objects have to reach a certain normal acceleration before they are allowed to take off, which appears to be unphysical. Moreover, the authors do not fully resolve the resulting QP but rather stop after a few iterations only, with no bounded tolerance. Kaufman et al. [2005] derive an approximation of Coulomb friction based on the maximum dissipation principle, leading to the formulation of two separable QPs. Similarly as in Milenkovic and Schmidl [2001], the problem is only partially solved for the sake of efficiency, using a limited number of iterations regardless of any convergence error.

Such simplified methods are appealing due to the fact they may scale favorably with respect to the number of contacts. In Kaufman et al. [2005] computational time grows linearly with the number of contacts, which allows them to simulate a large number of rigid bodies subject to tens of thousands of contacts, in a reasonable amount of time. In contrast, due to its higher complexity, our approach is currently limited to the handling of a few thousand contacts. However, simplifications made on Coulomb friction and on the accuracy of resolution have a price to be paid, in terms of physical realism and stability: with our solver, as illustrated on the spaghetti and hair demos, all contacting systems remain *perfectly still* when they have come to rest, as one would expect. This is hardly achieved by the simulations presented in Kaufman et al. [2005].

6.1.3 Faceted Coulomb's Friction Cone. The current trend in graphics to model Coulomb friction consists of solving an LCP generated by linearizing the Coulomb friction cone [Erleben 2007; Kaufman et al. 2008; Otaduy et al. 2009]. The main advantage of this approach is a theoretical result due to Stewart and Trinkle [1996], which both provides a criterion for the existence of at least one solution, and proves its computability via a pivoting technique

such as the Lemke algorithm. While using an alternative technique based on staggered projections Kaufman et al. [2008] also report good convergence properties, even when using a very low convergence tolerance (down to 10^{-8} , compared to a 10^{-1} tolerance said to be “sufficient to generate convincing frictional behavior”). Note that such quantitative measures of convergence are very sparse in the graphics literature when dealing with the frictional contact problem.

The faceting process suffers a number of issues. Firstly, the faceting introduces an artificial anisotropy in the sliding plane yielding some artifacts in the mechanical behavior of the whole system. This induced anisotropy is generally difficult to avoid and to forecast with respect to the number of facets. The main reason is the nonmonotonicity of the convergence of the error with the number of facets compared to the original second-order friction cone. As an illustration, a simple example of this issue where a sphere on a table is controlled by forces to draw a triangle is presented in Acary and Brogliato [2008, Section 13.3.7]. Such a lack of precision may especially be disturbing for critical applications such as haptics or virtual prototyping. The second drawback of faceting the friction cone is the poor numerical efficiency beyond 1000 contacts. In most favorable cases, a pivoting method for an LCP such as Lemke’s method is empirically of complexity $O(n^3)$, where n is the number of unknowns. When the number of facets increases at each contact point, the number of unknowns grows linearly leading quickly to intractable problems. Last but not least, one major drawback of the faceting approach is the difficulty of its implementation compared to a straightforward Newton method.

6.1.4 Global and Local Approaches. Our contact algorithm attempts to compute the contact forces by changing all forces at all contact points simultaneously. It can solve, in particular, problems involving only one contact, and this property can be used to devise another algorithm for the same problem, which solves a local problem contact-per-contact and loops over all contacts. This method is often called “Gauss-Seidel” due to its resemblance with the Gauss-Seidel algorithm to solve linear systems. The main interests of this algorithm are its good scalability with respect to the number of contacts, and the fact that it may be able, in some cases, to solve massive stacking problems of systems with a low number of degrees of freedom. Actually, this method requires very little implementation once the global method is available. In our rods experiments, however, this method did not necessarily prove faster nor more robust than the global method: the method suffered slow convergence together with cycling issues and sometimes failed to converge even in the favorable case where $\nu < 1$; moreover, its stopping criterion is less clear since there is no parameter (such as our $\|f^{AC}\|^2$) which monotonically controls convergence. As a consequence, we advocate the use of the global method when the criterion $\nu < 1$ is roughly satisfied (which was typically the case in our experiments), and of the local approach otherwise (typically, in the case of densely packed fibers or nonconvex rigid bodies stacking).

6.2 Limitations and Future Work

As mentioned before, our global approach is not adapted to over-constrained systems, that is, systems where the number of degrees of freedom is low compared to the number of contacts. This is typically the case in assemblies of rigid bodies. Note that in Figure 9 we have $d = 3$, $m = 6$ and $n = 2$, which makes this example fall into the favorable case where $\nu = 1$. However, the criterion is generally not verified in stacking systems of rigid bodies. Take, for example, a pile of cubes in 3D: if the number of contacts per cube is $n = 4$,

then $\nu = 2$ and the criterion is not satisfied. Even after filtering the number of active contacts at the collision detection stage (leading to at least 3 contacts per face), ν would still be out of range. In practice, we indeed observed slow convergence and frequent occurrence of large convergence errors when trying to simulate this kind of problem. The worst case is actually obtained when simulating rigid bodies ($m = 6$) with a high number of self-contacts: typically, concave and convex rigid bodies that get entangled.

Our convergence criterion $\nu < 1$ is defined independently from the parameterization of the rod model used in the simulations. However, this criterion is obviously more favorable to systems whose kinematics is loosely defined, for example, computed through energy functions (penalization terms) rather than hard constraints (reduced-coordinates or explicit hard constraints), since loose kinematics do not exactly remove degrees of freedom. However, in the case of deformable bodies, one acceptable solution may consist in increasing the resolution of the model at locations where the number of contacts is high. Note that if this is required in order to get good properties of convergence, this is also often desirable for increasing the quality of simulation, as high deformations are likely to occur near the contacting zones.

In the future, we would like to better address the difficult case where a high number of contacts is applied onto deformable objects without causing high deformations (such as a straight rod, with a low number of degree of freedom, lying on a flat surface): in this case, it would of course be preferable to model the rod using a low number of degrees of freedom (for instance, reduced-coordinates), which for the moment would be hardly solvable by our approach in the case when a high number of contacts is applied. Our goal would be to find a global approach that scales better in the number of contacts, while possessing better convergence properties than classical Gauss-Seidel like approaches. We note that in the specific case of dense structured stacking of rigid bodies with a faceted Coulomb cone, Erleben [2007] proposed a heuristic based on velocity-shock propagation to accelerate the rate of convergence of the Gauss-Seidel algorithm. In the same spirit, exploiting the peculiar structure of a dense packing of fibers may help us design a dedicated, optimized iterative scheme that would properly scale up for the efficient handling of thousands of fibers in contact.

7. CONCLUSION

We have introduced a new method for computing self-contacts in an assembly of thin elastic rods, in the presence of friction. Our method is simple to implement and can be applied to any rod model—maximal or reduced-coordinates. The only required information to provide, besides the dynamics, is the gradient \mathbf{J} . Furthermore, we provide the user with a simple criterion able to predict the quality of convergence for a given simulation. This criterion is often favorable in the case of deformable bodies. By testing our algorithm on various rods models, we showed our ability to capture subtle, nonsmooth effects due to contact and friction.

In the future, we would like to further investigate the scalability of our method to simulate very large assemblies of strands. Our main motivation is to design a reference model for hair material in order to better understand the process of wisps formation during motion, for example, to analyze the influence of hair geometry and hair friction on this process, such studies being very challenging to conduct experimentally [Bertails et al. 2006]. Such an accurate simulator would serve as a first step before designing a more efficient, macroscopic hair model that could further be validated against this reference model.

We would also like to extend our method to the robust handling of self-contact in cloth, which is crucial for capturing important features such as folds and wrinkles.

APPENDIX

A. EQUIVALENCE BETWEEN DISJUNCTIVE AND ZERO-FINDING FORMULATIONS FOR COULOMB'S LAW

We show here how to prove the equivalence $(\mathbf{u}, \mathbf{r}) \in C \iff \mathbf{f}^{\text{AC}}(\mathbf{u}, \mathbf{r}) = \mathbf{0}$ where $\mathbf{f}^{\text{AC}}(\mathbf{u}, \mathbf{r})$ is the Alart-Curnier function defined by Eq. (4). This can be easily achieved by considering each case separately. Let us first demonstrate that $(\mathbf{u}, \mathbf{r}) \in C \implies [\mathbf{f}_N^{\text{AC}}, \mathbf{f}_T^{\text{AC}}](\mathbf{u}, \mathbf{r}) = \mathbf{0}$. In the taking off case, $\mathbf{r} = \mathbf{0}$ and $u_N \geq 0$ therefore $\mathbb{P}_{\mathbb{R}^+}(r_N - \rho_N u_N) = \mathbb{P}_{\mathbb{R}^+}(-\rho_N u_N) = 0 = r_N$ and $\mathbb{P}_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T - \rho_T \mathbf{u}_T) = \mathbb{P}_{\mathbf{B}(0,0)}(\dots) = \mathbf{0} = \mathbf{r}_T$. In the sticking case, $\mathbf{u} = \mathbf{0}$ and $\mathbf{r} \in K$ therefore $\mathbb{P}_{\mathbb{R}^+}(r_N - \rho_N u_N) = \mathbb{P}_{\mathbb{R}^+}(r_N) = r_N$ and $\mathbb{P}_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T - \rho_T \mathbf{u}_T) = \mathbb{P}_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T) = \mathbf{r}_T$. Finally, in the sliding case, $\mathbb{P}_{\mathbb{R}^+}(r_N - \rho_N u_N) = \mathbb{P}_{\mathbb{R}^+}(r_N) = r_N$ (because $u_N = 0$ and $r_N \geq 0$) and $\mathbb{P}_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T - \rho_T \mathbf{u}_T) = \mathbf{r}_T$ (because $\mathbf{r}_T \in \partial K$ and $-\rho_T \mathbf{u}_T \in N_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T)$).

Conversely, if $[\mathbf{f}_N^{\text{AC}}, \mathbf{f}_T^{\text{AC}}](\mathbf{u}, \mathbf{r}) = \mathbf{0}$, let us show that $(\mathbf{u}, \mathbf{r}) \in C$ holds. As $\mathbf{f}_T^{\text{AC}}(\mathbf{u}, \mathbf{r}) = \mathbf{0}$ one has $\mathbf{r}_T \in \mathbf{B}(0, \mu r_N)$ therefore $\mathbf{r} \in K$. As $\mathbf{f}_N^{\text{AC}}(\mathbf{u}, \mathbf{r}) = \mathbf{0}$, one has $r_N \geq 0$ and we treat the two cases $r_N = 0$ (*) and $r_N > 0$ (***) separately. If $r_N = 0$ (*) then $\mathbf{r}_T = \mathbf{0}$ (because $\mathbf{r}_T = \mathbb{P}_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T - \rho_T \mathbf{u}_T) = \mathbb{P}_{\mathbf{B}(0,0)}(\dots) = \mathbf{0}$) and $u_N \geq 0$ (because $0 = r_N = \mathbb{P}_{\mathbb{R}^+}(-\rho_N u_N)$) and one is in the taking off case. If $r_N > 0$ (***) then $r_N = \mathbb{P}_{\mathbb{R}^+}(r_N - \rho_N u_N) = r_N - \rho_N u_N$ therefore $u_N = 0$. Again two different cases appear: $\mathbf{r}_T \in \text{int } \mathbf{B}(0, \mu r_N)$ and $\mathbf{r}_T \in \partial \mathbf{B}(0, \mu r_N)$. In the first case, $\mathbf{r}_T = \mathbb{P}_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T - \rho_T \mathbf{u}_T) = \mathbf{r}_T - \rho_T \mathbf{u}_T$ therefore $\mathbf{u}_T = \mathbf{0}$, this corresponds to the sticking case. In the second case, $\mathbf{r} \in \partial K$ and $\mathbf{u}_T \in N_{\mathbf{B}(0, \mu r_N)}(\mathbf{r}_T)$ therefore $\exists \alpha > 0$, $\mathbf{r}_T = -\alpha \mathbf{u}_T$, this is the sliding case. This ends the proof.

B. JACOBIAN MATRIX OF FUNCTION G

Distinguishing between three cases, the following formulas are easily derived.

If $\|y\| < \lambda$, then $g(\lambda, y) = y$, $\frac{\partial g}{\partial \lambda}(\lambda, y) = \mathbf{0}_{2 \times 1}$, $\frac{\partial g}{\partial y}(\lambda, y) = I_2$.

If $\|y\| > \lambda > 0$, then $g(\lambda, y) = \lambda \frac{y}{\|y\|}$, $\frac{\partial g}{\partial \lambda}(\lambda, y) = \frac{y}{\|y\|}$, $\frac{\partial g}{\partial y}(\lambda, y) = \frac{\lambda}{\|y\|} (I_2 - \frac{yy^T}{\|y\|^2})$.

If $\|y\| > 0 > \lambda$, then $g(\lambda, y) = \mathbf{0}_{2 \times 1}$, $\frac{\partial g}{\partial \lambda}(\lambda, y) = \mathbf{0}_{2 \times 1}$, $\frac{\partial g}{\partial y}(\lambda, y) = \mathbf{0}_{2 \times 2}$.

REFERENCES

ACARY, V. AND BROGLIATO, B. 2008. Numerical Methods for Nonsmooth Dynamical Systems. Lecture Notes in Computational and Applied Mechanics, vol. 35. Springer.

ACARY, V., CADOUX, F., LEMARECHAL, C., AND MALICK, J. 2010. A formulation of the linear discrete Coulomb friction problem via convex optimization. *ZAMM angew Math Mech; Zeitschrift für Angewandte Mathematik und Mechanik*.

ALART, P. AND CURNIER, A. 1991. A mixed formulation for frictional contact problems prone to Newton like solution methods. *Comput. Methods Appl. Mech. Engin.* 92, 3, 353–375.

BARAFF, D. 1989. Analytical methods for dynamic simulation of nonpenetrating rigid bodies. In *Proceedings of the ACM SIGGRAPH' Conference*. ACM, New York, 223–232.

BARAFF, D. 1993. Issues in computing contact forces for non-penetrating rigid bodies. *Algorithmica* 10, 292–352.

BARAFF, D. 1994. Fast contact force computation for nonpenetrating rigid bodies. In *Proceedings of the ACM SIGGRAPH' Conference*. ACM, New York, 23–34.

BARAFF, D. AND WITKIN, A. 1992. Dynamic simulation of nonpenetrating flexible bodies. In *Proceedings of the ACM SIGGRAPH' Conference*. 303–308.

BARAFF, D. AND WITKIN, A. 1998. Large steps in cloth simulation. In *Proceedings of the ACM SIGGRAPH' Conference*. 43–54.

BARBIČ, J. AND JAMES, D. 2007. Time-Critical distributed contact for 6-dof haptic rendering of adaptively sampled reduced deformable models. In *Proceedings of the ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'07)*. Eurographics Association, 171–180.

BERGOU, M., WARDETZKY, M., ROBINSON, S., AUDOLY, B., AND GRINSPUN, E. 2008. Discrete elastic rods. *ACM Trans. Graph.* 27, 3, 1–12.

BERTAELS, F. 2009. Linear time super-helices. *Comput. Graph. Forum* 28, 2.

BERTAELS, F., AUDOLY, B., CANI, M.-P., QUERLEUX, B., LEROY, F., AND LÉVÊQUE, J.-L. 2006. Super-helices for predicting the dynamics of natural hair. In *Proceedings of the ACM SIGGRAPH*. ACM, 1180–1187.

BLINN, J. 1998. Ten more unsolved problems in computer graphics. *IEEE Comput. Graph. Appl.* 18, 5, 86–89.

BONNANS, J. F., GILBERT, J. C., LEMARÉCHAL, C., AND SAGASTIZÁBAL, C. A. 2003. *Numerical Optimization*. Springer, Berlin.

BRIDSON, R., FEDKIW, R., AND ANDERSON, R. 2002. Robust treatment of collisions, contact and friction for cloth animation. *ACM Trans. Graph.* 21, 3, 594–603.

CADOUX, F. 2009. Méthodes d'optimisation pour la dynamique non-régulière. Ph.D. thesis, Université Joseph Fourier.

CHOE, B., CHOI, M., AND KO, H.-S. 2005. Simulating complex hair with robust collision handling. In *Proceedings of the ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'05)*. 153–160.

CHRISTENSEN, P., KLARBRING, A., PANG, J., AND STROMBERG, N. 1998. Formulation and comparison of algorithms for frictional contact problems. *Int. J. Numer. Methods Engin.* 42, 1, 145–173.

DURIEZ, C., DUBOIS, F., KHEDDAR, A., AND ANDRIOT, C. 2006. Realistic haptic rendering of interacting deformable objects in virtual environments. *IEEE Trans. Visualiz. Comput. Graph.* 1, 12, 36–47.

DURVILLE, D. 2004. Modelling of contact-friction interactions in entangled fibrous materials. In *Proceedings of the 6th World Congress on Computational Mechanics (WCCM VI)*.

ERLEBEN, K. 2007. Velocity-Based shock propagation for multibody dynamics animation. *ACM Trans. Graph.* 26, 2.

FRÉMOND, M. 2002. *Non-Smooth Thermo-Mechanics*. Springer, Berlin.

HADAP, S. 2006. Oriented strands - Dynamics of stiff multi-body system. In *Proceedings of the ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'06)*. ACM, 91–100.

HADAP, S. AND MAGNENAT-THALMANN, N. 2001. Modeling dynamic hair as a continuum. *Comput. Graph. Forum* 20, 3, 329–338.

HARMON, D., VOUGA, E., SMITH, B., TAMSTORF, R., AND GRINSPUN, E. 2009. Asynchronous contact mechanics. *ACM Trans. Graph. Conference*.

JEAN, M. 1999. The non smooth contact dynamics method. *Comput. Methods Appl. Mech. Engn.* 177, 235–257.

KALDOR, J., JAMES, D., AND MARSCHNER, S. 2008. Simulating knitted cloth at the yarn level. *ACM Trans. Graph. Conference*. 1–9.

KAUFMAN, D., EDMUNDS, T., AND PAI, D. 2005. Fast frictional dynamics for rigid bodies. *ACM Trans. Graph.* 24, 3, 946–956.

- KAUFMAN, D., SUEDA, S., JAMES, D., AND PAI, D. 2008. Staggered projections for frictional contact in multibody systems. *ACM Trans. Graph.* 27, 5, 164:1–164:11.
- LENOIR, J., MESEURE, P., GRISONI, L., AND CHAILLOU, C. 2004. A suture model for surgical simulation. In *Proceedings of the 2nd International Symposium on Medical Simulation (ISMS'04)*. 105–113.
- LÖTSTEDT, P. 1984. Numerical simulation of time-dependent contact friction problems in rigid body mechanics. *SIAM J. Sci. Stat. Comput.* 5, 2, 370–393.
- MAPLESOFT. 2010. Maple 14. <http://www.maplesoft.com/>.
- MARQUES, M. M. 1993. *Differential Inclusions in Nonsmooth Mechanical Problems. Shocks and Dry Friction*. Progress in Nonlinear Differential Equations and their Applications. 9. Birkhäuser, Basel.
- MCADAMS, A., SELLE, A., WARD, K., SIFAKIS, E., AND TERAN, J. 2009. Detail preserving continuum simulation of straight hair. *ACM Trans. Graph. Conference* 28, 3, 1–6.
- MILENKOVIC, V. AND SCHMIDL, H. 2001. Optimization-Based animation. In *Proceedings of the ACM SIGGRAPH' Conference*. ACM, New York, 37–46.
- MOORE, M. AND WILHELMS, J. 1988. Collision detection and response for computer animation3. In *Proceedings of the ACM SIGGRAPH' Conference*. 289–298.
- MOREAU, J. 1994. Some numerical methods in multibody dynamics: Application to granular materials. *Euro. J. Mechan. A/Solids supp.*, 4, 93–114.
- MOREAU, J.-J. 1988. Unilateral contact and dry friction in finite freedom dynamics. *Nonsmooth Mechanics and Applications, CISM Courses Lectures* 302, 1-82 (1988).
- MÜLLER, M., HEIDELBERGER, B., HENNIX, M., AND RATCLIFF, J. 2007. Position based dynamics. *J. Vis. Comm. Image Represent.* 18, 2, 109–118.
- OTADUY, M., TAMSTORF, R., STEINEMANN, D., AND GROSS, M. 2009. Implicit contact handling for deformable objects. *Comput. Graph. Forum* 28, 2.
- PLANTE, E., CANI, M.-P., AND POULIN, P. 2001. A layered wisp model for simulating interactions inside long hair. In *Proceedings of the EG Workshop on Computer Animation and Simulation (EGCAS'01)*. Springer, 139–148.
- QI, L. AND SUN, J. 1993. A nonsmooth version of Newton's method. *Math. Program.* 58, 353–367.
- REDON, S., KHEDDAR, A., AND COQUILLART, S. 2002. Fast continuous collision detection between rigid bodies. *Comput. Graph. Forum.* 21, 279–288.
- ROSENBLUM, R., CARLSON, W., AND TRIPP, E. 1991. Simulating the structure and dynamics of human hair: Modeling, rendering, and animation. *J. Visualiz. Comput. Anim.* 2, 4, 141–148.
- SELLE, A., LENTINE, M., AND FEDKIW, R. 2008. A mass spring model for hair simulation. *ACM Trans. Graph.* 27, 3, 1–11.
- SPILLMANN, J., BECKER, M., AND TESCHNER, M. 2007. Non-Iterative computation of contact forces for deformable objects. *J. WSCG 15*, 33–40.
- SPILLMANN, J. AND TESCHNER, M. 2007. Corde: Cosserat rod elements for the dynamic simulation of one-dimensional elastic objects. In *Proceedings of the ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'07)*. ACM, 63–72.
- SPILLMANN, J. AND TESCHNER, M. 2008. An adaptive contact model for the robust simulation of knots. *Comput. Graph. Forum* 27, 2.
- STEWART, D. AND TRINKLE, J. 1996. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and Coulomb friction. *ijnme* 39, 15.
- STEWART, D. E. 2000. Rigid-body dynamics with friction and impact. *SIAM Rev.* 42, 1, 3–39.
- TESCHNER, M., HEIDELBERGER, B., MÜLLER, M., POMERANERTS, D., AND GROSS, M. 2003. Optimized spatial hashing for collision detection of deformable objects. In *Proceedings of the Conference on Vision, Modeling, Visualization (VMV'03)*. 47–54.
- TESCHNER, M., KIMMERLE, S., HEIDELBERGER, B., ZACHMANN, G., RAGHUPATHI, L., FUHRMANN, A., CANI, M.-P., FAURE, F., MAGNENAT-THALMANN, N., STRASSER, W., AND VOLINO, P. 2005. Collision detection for deformable objects. *Comput. Graph. Forum* 24, 1, 61–81.
- THEETTEN, A., GRISONI, L., ANDRIOT, C., AND BARSKY, B. 2008. Geometrically exact splines. *J. Comput. Aid. Des.* 40, 1, 35–48.

Received February 2010; revised October 2010; accepted November 2010