



HAL
open science

Numerical modeling of elastic slender structures subject to contact and friction: from dynamic simulation to inverse static design

Florence Bertails-Descoubes

► To cite this version:

Florence Bertails-Descoubes. Numerical modeling of elastic slender structures subject to contact and friction: from dynamic simulation to inverse static design. Modeling and Simulation. Université Grenoble Alpes, 2017. tel-01925260

HAL Id: tel-01925260

<https://hal.science/tel-01925260>

Submitted on 16 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HABILITATION À DIRIGER DES RECHERCHES

Spécialité : **Mathématiques et Informatique**

Présentée par

Florence Bertails-Descoubes

préparée au sein de l' **équipe-projet Bipop, Inria et Laboratoire Jean Kuntzmann**

et de l'école doctorale "**Mathématiques, Sciences et Technologies de l'Information, Informatique**"

Numerical Modeling of elastic slender structures subject to contact and friction

From dynamic simulation to inverse static design

Soutenue publiquement le **30 mai 2017**,
devant le jury composé de :

Pr Eitan Grinspun

Associate Professor at Columbia University, Rapporteur

Pr Pierre Alart

Professeur des Universités, Université de Montpellier II, Rapporteur

Dr Bruno Lévy

Directeur de Recherche à l'Inria, Inria Nancy-Grand Est, Rapporteur

Dr Sébastien Neukirch

Directeur de Recherche au CNRS, Laboratoire Jean le Rond d'Alembert,
Examineur

Dr Fabrice Neyret

Directeur de Recherche au CNRS, Laboratoire Jean Kuntzmann, Examineur

Dr Stéphane Redon

Chargé de Recherche à l'Inria, Laboratoire Jean Kuntzmann, Examineur

Dr Roberto Santoprete

Chercheur R&D, L'Oréal Recherche, Examineur



Acknowledgments

I would first like to thank Eitan Grinspun, Pierre Alart, and Bruno Lévy, who have kindly accepted to review this manuscript. I am also grateful to Fabrice Neyret, Sébastien Neukirch, Roberto Santoprete, and Stéphane Redon, for completing this pluridisciplinary committee. Special thanks to Sébastien for the many corrections he reported on my manuscript, which helped me improve its overall quality.

The document I am presenting here draws up a synthesis of the work which has been achieved since I was hired at Inria in 2007, mostly in collaboration with my PhD students, postdoc and engineers. I have or had the opportunity to work with talented, hard-working, open-minded and nice people, and I am thus extremely grateful to them: Alexandre Derouet-Jourdan (PhD 2013), Romain Casati (PhD 2015), Gilles Daviet (engineer, then PhD 2016), Alejandro Blumentals (ongoing PhD), Mickaël Ly (intern in 2016), and Victor Romero (ongoing postdoc).

Many thanks to the colleagues who I have or had the pleasure to collaborate with, often through student co-supervision: Florent Cadoux, Joëlle Thollot, Franck Hétroy, Bernard Brogliato, Pierre Saramito, Arnaud Lazarus, Damien Rohmer, Jean-Sébastien Franco, Rahul Narain, Laurence Boissieux. My activity has also involved industrial transfer, and I would like to thank all the industrial partners, from Néomis Animation & Beelight to L'Oréal Research, I had the opportunity to work with, as well as the Inria valorisation cell for making such collaborations possible.

I am grateful to Bernard Brogliato for having welcomed me in the BiPop group in 2007. Being the sole researcher from Computer Graphics among a set of scientists focussed on more fundamental topics such as Mechanics or Optimization was a challenging situation for me, but corresponded to the scientific environment I was looking for at that time. Beyond the scientific immersion into nonsmooth mechanics and optimization, working in BiPop brought me the opportunity to reach some scientific independence very soon. It also definitely helped me step back and get a fresh point of view on various domains, including my own community.

I precisely take this opportunity to thank my home community, Computer Graphics, for the intense creativity I found in all the people working in this field, and the overall benevolence ruling over people. Moreover, I still remain impressed by the professionalism of Siggraph and Siggraph Asia program committees, which, I believe, is truly exceptional, and had a strong impact on my evaluation practice in various other contexts.

Thinking back to my first years of research, I feel very grateful to my PhD advisors, Marie-Paule Cani and Basile Audoly, and to my postdoc supervisor Robert Bridson, for sharing with me, early in my career, their passion for research and a vision of science that goes well beyond the barriers separating different domains. While I have followed a quite independent path during the ten past years, they have always kept a benevolent and supporting view on my work.

I would also like to thank all my colleagues and friends at Inria, as well as our great team assistants, for their invaluable support.

And, above all, thank you to my dear, dear family.

Contents

1	Introduction	1
1.1	Simulating visually rich phenomena	2
1.2	Computer Graphics as my home community	4
1.3	Methodology and guidelines	6
1.4	Structure of this document	9
2	Discrete curvature-based models for dynamic fibers	11
2.1	Prior art and choice for reduced models	12
2.2	Kirchhoff equations for thin elastic rods	15
2.3	Discrete curvature-based models	17
2.3.1	Piecewise uniform curvature: the super-helix model	19
2.3.2	Piecewise linear curvature: the super-clothoid model	20
2.4	Linear-time integration of the dynamics	26
2.4.1	Recursive integration scheme	27
2.4.2	Comparison with the composite approach	28
2.5	Conclusion: Towards a curvature-based shell model	29
3	Frictional contact models for large multibody systems	31
3.1	Prior art and choice for nonsmooth models	33
3.2	Nonsmooth contact dynamics for thin elastic objects	36
3.2.1	Adding unilateral contact with Coulomb friction	36
3.2.2	The Discrete Frictional Contact Problem (DFCP)	39
3.2.3	Multiple equivalent formulations of Signorini-Coulomb	41
3.2.4	A local hybrid solving strategy	46
3.3	Continuum modeling of granular materials	48
3.3.1	Discrete vs. continuum models for granular matter	48
3.3.2	Macroscopic yield criteria	50
3.3.3	Nonsmooth rheology $\mathcal{DP}(\mu)$	51
3.3.4	Extension to a varying volume fraction	54
3.4	Conclusion: Towards a continuum-based model for fibrous materials	54
4	From geometry to mechanics: Inverse design of fibers	57
4.1	Inverse design problem: motivation and goal	58
4.2	Case of an isolated fiber	61
4.2.1	Necessary and sufficient condition for equilibrium	62
4.2.2	Sufficient condition for stability	63
4.2.3	Inverse design algorithm	64
4.2.4	From an arbitrary curve to a discrete material configuration	65
4.3	Case of a fiber assembly subject to frictional contact	66
4.3.1	Underdetermined conditions for equilibrium	68

4.3.2	A simple heuristics for a well-posed problem	69
4.4	From a geometric curve to a material curve	71
4.4.1	From a geometric curve to a material curve	72
4.4.2	Curve-angle parameterization	73
4.4.3	Direct problem: framing a Kirchhoff rod	74
4.4.4	Inverse static design problem	76
4.4.5	Discrete picture	78
4.5	Conclusion: Towards non-invasive identification	79
5	Research Perspectives: From Geometry to Mechanics	81
5.1	From geometry to mechanics, a broken pipe	81
5.2	Interpreting geometry as a mechanical state	84
5.3	Scientific problems	84
5.4	Expected Impact	86
6	Conclusion	87
	Bibliography	91

Introduction

J'ai pas mérité de jouer du rock'n'roll. Mes ghettos, mes idées ne sont pas homologués. J'ai pas le bon blouson, j'ai pas les bonnes bottes. Et en haut de mon bras je n'ai rien fait tatouer. J'ai donné aux curés du sauvetage collectif. J'ai joué les mêmes notes, swingué les mêmes riffs. Peu à peu j'ai compris les données du débat. Que rien ne bouge et l'égalité par le bas. Et tant pis si la foule gronde. Si je tourne pas dans la ronde. Papa quand je serai grand, je sais ce que je veux faire : je veux être minoritaire.

– Jean-Jacques Goldman, *Minoritaire*

My research activity deals with the mechanical modeling and numerical simulation of elastic slender structures, in particular fibers, modelled as *thin elastic rods*, but also surfaces in some more recent work, elaborating on models for *thin elastic plates* or *developable shells*. I've also got interested in the capture of *dry frictional contact* among such thin structures, as well as within granular flows. All these dynamical systems are useful to model a large range of physical phenomena at different scales, from twisted DNA to the mechanics of twining plants, sand, human hair, and cloth (see Figure 1.1). Target applications encompass visual effects for the feature film industry, as well as virtual prototyping for various domains such as cosmetology (hair), virtual try on (hair and cloth), or risk management (granular materials).




Figure 1.1 – Our main simulation results: from the numerical modeling of thin elastic rods and shells (top) to the coupling of rigid bodies, rods and shells for capturing the dynamics of complex macroscopic phenomena involving dry frictional contact (bottom).

1.1 Simulating visually rich phenomena

Deformable slender structures such as hair fibers, cloth, ribbons, paper, cables, tree branches or leaves, are ubiquitous around us. It is actually striking to note that such thin structures represent most of the deformable objects composing our direct environment, but also a wide range of physical, biological or manufactured phenomena occurring at very different scales, from macromolecules and carbon nanotubes at nanoscale to galactic filaments at the universe scale (see Figure 1.2).

At the macroscopic scale (and even smaller scales), slender structures often feature an intricate natural shape, ranging from straight to curly. For small strains, their deformation is mostly elastic (with a preference for *bending* and *twisting* when no overly constrained), yet generating large displacements due to geometric nonlinearities, and instabilities such as buckling. The diversity and complexity of the resulting shapes and motions greatly contributes to the visual richness of the real world.

When multiple thin structures are coupled together with contact and friction, which is the predominant interaction mode at the macroscopic scale, the range of emerging phenomena is even more exacerbated, giving rise to stick-slip dynamical instabilities, entangling, or spontaneous collective behavior. Human hair, which is typically composed of 150,000 thin fibers, beautifully depicts such complex mechanical behaviors when fluttering in the wind.




Figure 1.2 – Top: Slender structures such as hair, cables, plants, or cloth are widespread at the macroscopic scale, where their interactions are dominated by dry frictional contact. Bottom: it is interesting to observe that slender structures also appear at extreme scales, such as supercoiled DNA (left) and carbon nanotubes (middle) when zooming in some biological and manufactured matter, or even galactic filaments (right) when zooming out up to the universe scale. Image credits: (c) S. Paris, T. Judd, 2007, (d) J. Bahnson’s group, (e) L. Qu, U. Dayton, 2008, (f) A. Pontzen, F. Governato, 2014.

Many physicists and mathematicians have strived for centuries to understand the principles governing those complex mechanical phenomena, providing a number of continuous models for slender structures, granular matter, and frictional contact. In the XXth century, industrial applications such as process automatization and new ways of transportation have boosted the fields of Mechanical Engineering and Computer-Aided Design, where material strength, reliability of mechanisms, and safety, stand for the main priorities. Instead, large displacements of structures, buckling, tearing, or entanglement, and even dynamics, were long considered as undesirable behaviors, thus restraining the search for corresponding numerical models. In contrast, Computer Graphics, which has emerged in the 60's with the advent of modern computers, was eager to capture such peculiar phenomena, with the sole aim to produce spectacular images and create astonishing stories.

At the origin, Computer Graphics thus drastically departed from other scientific fields. Everyday-life phenomena such as cloth buckling, paper tearing, or hair fluttering in the wind, mostly ignored by other scientists at that time, became actual topics of interest, involving a large set of new research directions to be explored, both in terms of modelling and simulation.

Since a few decades, a new generation of physicists became interested again in the understanding of such visually fascinating phenomena, and started investigating the tight links between geometry and elasticity¹. Common objects such as folded or torn paper, twined plants, coiled honey threads, or human hair [Ben Amar and Pomeau, 1997; Goriely and Neukirch, 2006; Brun et al., 2012; Goldstein et al., 2012] have thus regained some popularity among the community in Nonlinear Physics². Yet, while the phenomena of interest have become remarkably close to those of Computer Graphics, the goals and employed methodologies still differ substantially from one community to the other, as explained in next section.

In parallel, the engineering industry has recently shown some new and growing interest into the modeling of dynamic phenomena prone to large displacements, contact and friction. For instance, the cosmetology industry is more and more interested in understanding the nonlinear deformation of hair and skin, with the help of simulation. Likewise, auto and aircraft manufacturers are facing new challenges involving buckling or entanglement of thin structures such as carbon or optical fibers; they clearly lack predictive, robust and efficient numerical tools for simulating and optimizing their new manufacturing processes.

¹In France this new trend was particularly stimulated by the work of Yves Pomeau, who convinced many young scientists to study the nonlinear physics of common objects such as paper, plants, or hair [ESPCI, 2016].

²It is however amusing to observe that research in these areas is quite successful in obtaining the IGNobel prize [Audoly and Neukirch, 2005; Goldstein et al., 2012], thus still being considered as an exotic research topic by physicists.

1.2 Computer Graphics as my home community

Computer Graphics is a very particular domain. From outside, one may hardly understand what this community is about if one does not know that it represents, above all, the gathering of people from various career paths (artists, industrials, scientists) who share a common *passion* for pictorial arts and somehow an aesthetic vision of the world. Most people I have met in Computer Graphics, even among scientists and industrials, do have an acute sense for artistic creativity, many of them being keen on drawing (like myself), on the making of movies and video games, or on any other creative activity. By construction, the scientific community in Computer Graphics is thus tightly linked to the digital entertainment industry, which encompasses all activities related to the making of animation feature films, special effects, and video games. It is noteworthy that a large number of researchers who are active in the field are employed by private companies such as Autodesk, Disney Research, Pixar, Weta Digital, Microsoft Research, Adobe, etc. The most prestigious (and selective, see Figure 1.3) conference of the field, ACM Siggraph, gathers every year tens thousands experts and practitioners including research scientists, engineers, and artists, fostering an extraordinary mixing between academia, private research centers and studios, and engineering and arts schools.



Figure 1.3 – The famous Siggraph bar. Nobody knows exactly where it stands, but everyone knows you should go over it to have your paper accepted. Be aware that after a first success, you will get addicted to it.

Creating and directing visual digital worlds The primary and historic goal of Computer Graphics is to digitize the world into synthetic images targeted for a *human* audience. Computer Graphics thus has the ambition to represent visually all the phenomena we may perceive through our eyes, at all relevant scales. The dream application should allow a user to navigate in real-time within an infinite virtual space, being able for instance to look up at the clouds swirling in a stormy sky, then viewing down a forest canopy shaking under the wind, next zooming into the details of the undulating tree branches and distinguishing every fluttering leaf, even this small iridescent beetle climbing onto that hairy leaf... As if this were not

enough, the user may also interact with this virtual world as intuitively as possible, for example by grasping a wreath of flowers out of the prairie, creating a new river down the hills or changing the shape of the mountains, the shininess of the ocean, or the course of the wind, just at her will.

From other scientific fields, such an ambition of *creating* and *directing* an entire digital world similar to the real one, with all possible levels of details, may look quite adventurous and even untenable. Indeed, the usual scientific methodology is first to study simplified and refined phenomena, at a specific space and time scale, before daring to tackle more complex scenarios possibly integrating multiple bodies or multiple scales. The apparent audacity of Computer Graphics may however be explained by the fact that its goals, and thus its evaluation criteria, are drastically different from those of usual scientific fields. Unlike physics or biology, the goal here is not to explain observed phenomena, but to *reproduce* them on an array of pixels, as faithfully as possible. Unlike computational mechanics, the goal is not to compute accurate internal stresses inside a structure to anticipate rupture and strengthen the material at the right locations, the goal is to compute *visible* quantities such as deformation or velocity fields at the outer surface of the body, at a good resolution and free of any visual artifact; the goal is not to predict finely the parts of the building that will be destroyed by a fire starting at a certain point, the goal is to create flames that will *bluff* the audience and *let them believe*³ they are surrounded by fire; finally, the goal is not to characterize the exact shape and magnitude of a perfectly round drop of water falling into a perfectly flat basin; the goal is to provide the user with intuitive tools to create the drop animation that will best serve the underlying storytelling and convey the right emotions, or say otherwise, that will satisfy the eye of the *art director*⁴.

Having said that, creating outstanding visual effects still demands powerful computational models which cannot be built by hand, and instead need to be automated. Following physical principles is an excellent way to achieve this. This is particularly true in the sub-field of *Computer Animation*, which deals with the representation of moving phenomena.

³The Computer Graphics community often uses the term “plausibility” to justify that its models are intended for human eyes which should be the sole judges of their validity. However I do not like this term which is often misused and may serve as a pretence to escape from validation, should it be physical, perceptual, or anything else (validation is a vast but still immature research topic in Computer Animation). I will thus attempt to define my own evaluation criteria in Section 1.3, and will avoid using the term “plausibility” in this document. I however acknowledge the need for more and better validation of my own work, which I aim at developing in future work.

⁴In some situations, strictly simulating the known equations of physics may not be effective in conveying the right emotions, even if the underlying model is realistic: the audience may simply neither be convinced nor impressed by the simulated phenomenon! In many cases, a beautification or magnification of the phenomenon is thus desired, in a way comparable to the exaggerated deformed balls drawn by cartoonists to mimick expressive rebound [Williams, 2009]. In the same way *expressive rendering* has emerged in the 90’s, *expressive animation* has recently gained some interest. Providing user control over a physics-based simulation through inverse design, as we do in the third part of this memoir (see Chapter 4), is one research direction among others of this emerging subfield.

Physics-based animation More specifically, in the case of a dynamical scene involving passive objects such as hair, cloth, fluids, etc., *physics-based simulation* has proven over the years to be a method of choice for synthesizing resulting visual effects. Unlike phenomenological methods which develop descriptive models for reproducing a given emerging phenomenon, physics-based methods provide generative models which explicitly integrate the physical causes of the phenomenon. From a set of initial conditions as well as a few physical parameters (e.g., the mass, stiffness, natural shape), a physics-based simulator may generate not just a single effect, but a wide range of emerging phenomena revealing the whole complexity of the underlying physics.

My research work precisely lies within physics-based animation. The goal in our sub-community is to design numerical models which can reproduce convincingly very complex scenes, such as detailed breaking waves, a full garment folding on a walking character, or a full head of hair suddenly shaking. Although huge progress has been achieved in visual effects since the 90's, leading today to entirely synthetic animated outdoor scenes in most adventure movies, a large number of scenarios, such as the ones mentioned just above, still resists automatic digitization and requires many hours of manual work in production studios to be properly animated. Challenging scenarios typically involve a large number of degrees of freedom, strongly nonlinear dynamics, and many constraints — in particular, contacts and friction. All these features are very difficult to capture numerically. A supplementary obstacle may be the lack of relevant (continuous) physical models in some cases. For instance, while the theory of Kirchhoff for thin elastic rods fits in well with our needs for modeling a bending rope or a tree branch at the right (human) scale, to the best of our knowledge there does not exist an integrated, continuum model for representing a full bulk of very thin fibers, like hair. *Physical modeling* thus represents an important, pluridisciplinary and long-term challenge for bridging the gap between, on the one hand, available representations of a certain class of phenomena, and on the other hand, the will to capture more complex ones, at the right scales.

In the following I explain the methodology and guidelines I have set up for tackling some of the challenges listed above, in the particular cases of elastic slender objects (fibers, cloth) and divided materials (granular and fibrous media).

1.3 Methodology and guidelines

When designing a physics-based simulator for computer graphics, I have inevitably in mind the four following criteria:

Realism: Ingredients which are necessary to capture relevant visual effects should be identified, incorporated in the model and translated numerically with as few quality loss as possible, at the right resolution, and without visual artifacts.

Robustness: The simulator should converge properly for a relevant range of parameters. This is all the more challenging as we are interested in objects prone

to large displacements, and subject to a potentially huge number of frictional contact constraints, which causes standard simulators to break down.

Efficiency: The simulator should be fast enough for allowing complex scenes to be simulated in reasonable timings on a standard PC (in our case, a few days of computation for a given scene is an upper-bound).

Control: The simulator should provide the user with some handles to control the shape and motion of the object in an intuitive way, without having to tweak manually a significant number of meaningless parameters.

Unlike the common belief according to which favoring one of these criteria is necessarily done at the price of sacrificing others, I am striving to develop *compact* numerical models satisfying *all* four criteria at the same time.

Building dedicated models To achieve this, the streamline idea is to identify *specific* emerging phenomena to simulate, and then build *dedicated* numerical models taking advantage of an *upfront reduced* mathematical formulation. For instance, instead of treating fibers and cloth as generic 3D continuum elastic models with a potentially small thickness (modeled with an ε parameter), one may directly consider equations for thin elastic rods, plates and shells, where reduction is undertaken upstream. Such dedicated models just retain the right ingredients one wishes to capture, without having to care about numerical issues inherent to the degeneracy of the model when ε becomes smaller and smaller. Similarly, so far we have considered only *inextensible* models for thin elastic rods (namely, the *Kirchhoff theory* for thin elastic rods), because when bending and twisting deformations are the main focus of interest, stretching and shearing then become irrelevant deformation modes, as proven by a simple scaling argument [Audoly and Pomeau, 2010, Section 3.7]. Finally, for modeling frictional contact as effectively as possible, we have been considering macroscopic models which directly cope with the emerging *rigidity* and *nonsmoothness* of contact and friction, instead of using locally compliant models which get stiffer and stiffer when approaching the right scale.

Targeting the full modeling pipeline Of course, on the down side, such a strategy requires that many different fields — theory for thin elastic rods and shells, non-smooth contact mechanics, etc. — be investigated in deep, before tackling merely numerical problems, which themselves involve other scientific fields such as numerical analysis, optimization, and algorithmic. However, to design effective simulators, I believe it is essential to master the full modeling pipeline, from mathematical models to numerical schemes, including algorithmic choices and implementation issues. From a personal point of view, I am also eager to investigate and learn new insights from all these different scientific fields (despite sometimes a certain frustration due to some obvious lack of time). Besides, at the scale of a person or a small group, I can testify it is quite satisfactory to build prototypes which are entirely home-made, and thus completely under control.

Developing cutting-edge, but large usage codes Finally, one may argue that the benefit of dealing with generic models is to build some generic code, which can deal with (almost) any problem, while dedicated models are limited to very specific codes, and thus to narrow applications. In the context of my research work, I partially disagree with this argument, for two reasons. First, a generic code may probably deal with many different situations (e.g., a thin rod but also a thick rod, elastic bodies but also fluids, etc.), yet instead of treating *optimally* one of the possible scenarios (w.r.t. the criteria defined above), it is likely to do a rather poor job for all of them. Second, we have noticed, interestingly, that even though we have modeled each one of our phenomena specifically, *core* portions of our codes show up some *common features*. For instance, it is striking to note that it is exactly the same nonsmooth solver, originally published in the context of dynamic hair simulation [Daviet et al., 2011], which has served to resolve our hair inverse design problem [Derouet-Jourdan et al., 2013a], and which has then been leveraged to simulate granular flows modelled as a continuum [Daviet and Bertails-Descoubes, 2016b,a]. The emergence of common *canonical* problems has thus allowed us to concentrate our efforts on their efficient implementation.

Aiming at a pluridisciplinary research Although my research has been primarily driven and motivated by Computer Graphics applications, which are the most inspiring to me, the methods that we have employed to address related scientific problems largely go beyond the sole field of Computer Graphics, sometimes addressing questions of interest for other communities.

For one thing, as commonly done in Computer Animation, we have frequently borrowed tools and concepts from more fundamental fields such as optimization, geometry and mechanics. In the study of slender elastic structures, the originality of our approaches stemmed from the use of Cosserat elastic theories — barely known in Computer Graphics before the 2000’s and likewise barely popular in Mechanical Engineering, and which I started to familiarize with during my PhD thesis. In the field of contact mechanics, some new lightnings could be brought by using the point of view of nonsmooth optimization — which I benefited from due to my immersion within the BiPop team since 2007.

Besides, we have from time to time tackled scientific issues which surpass the sole field of Computer Graphics. We have gradually realized that publishing some specific technical contributions to graphics may leave them buried under the application (the final application being necessarily put forward in a graphics paper), whereas they could, in some cases, be relevant to other fields and benefit from a more careful evaluation from experts outside graphics. In recent years, we have thus strived to better identify the nature of our contributions and attempted to conduct a more diversified publication policy. One consequence is that we have started to publish some of our work in other relevant fields such as Computer Geometric Design, Multibody System Dynamics, or Non-Newtonian Fluid Mechanics. Likewise, we have started to participate in conferences and scientific events held on various

domains, in particular in Physics and Mechanical Engineering. In parallel, we are considering diverse applications of our work, encompassing not only special effects for the movie industry and feature film animation, but also virtual prototyping for the engineering industry (past collaborations with L’Oréal Research, starting collaboration with Safran).

1.4 Structure of this document

The first part of this document provides a summary of the three main topics of research I have been investigating since 2007 (date of my arrival at Inria as a “Chargée de recherche”), and is organized according to three dedicated chapters:

1. Modeling of dynamic fibers (Chapter 2) In this first research axis, we introduce a family of high-order, reduced models for discretizing the Kirchhoff equations for thin elastic rods in both a faithful and robust way [Bertails et al., 2006; Bertails, 2009; Bertails-Descoubes, 2012; Casati and Bertails-Descoubes, 2013]. Such models are particularly well-suited for simulating inextensible fibers subject to bending and twisting, and featuring an arbitrary curly intrinsic geometry. We have recently started to extend such discrete⁵ schemes to the numerical modeling of inextensible developable shells [Blumentals et al., 2016a].

2. Modeling of frictional contact (Chapter 3) A second research axis deals with the modeling of contact and Coulomb friction for slender structures and granular flows within the nonsmooth contact dynamics framework. Beyond some recent analytical study of the problem of existence and uniqueness of solutions to the continuous contact problem [Blumentals et al., 2016b], we have designed some numerical solvers to resolve the discrete contact problem in a both robust and efficient way. Such algorithms were leveraged for robustly handling thousands packed fibers at reasonable frame rates [Daviet et al., 2011], and more recently adapted to the simulation of frictional contact in cloth [Daviet et al., 2015], as well as the exact solving of the Drucker-Prager law for granular materials [Daviet and Bertails-Descoubes, 2016b,a].

3. Inverse static design of fibers (Chapter 4) A last research axis, which will be the main starting point of my future work, is dedicated to finding new insights into the inverse modeling of static fibers [Derouet-Jourdan et al., 2010, 2011, 2013b,a]. This work consisted in taking an arbitrary curve geometry as input and inferring corresponding geometric and physical parameters of the simulator such that the input geometry corresponds to a stable configuration at equilibrium.

⁵In the whole document, the term “discrete” refers to finite-dimensional models in space (which are subsequently discretized in time). Our contributions mostly deal with this *spatial* reduction of the Kirchhoff equations, yet with some implications on the full space-time discretization scheme.

Finally, I present in **Chapter 5** my research program for the five upcoming years, focussing on the **inverse modeling of slender structures subject to contact and friction**, before concluding this memoir.

The list of my publications is given [here](#). Besides, the reference implementations accompanying our research papers are distributed, according to a dual license policy, at <http://www.inrialpes.fr/bipop/people/bertails/Papiers/Code/index.html>.

Discrete curvature-based models for dynamic fibers

Contents

2.1	Prior art and choice for reduced models	12
2.2	Kirchhoff equations for thin elastic rods	15
2.3	Discrete curvature-based models	17
2.3.1	Piecewise uniform curvature: the super-helix model	19
2.3.2	Piecewise linear curvature: the super-clothoid model	20
2.4	Linear-time integration of the dynamics	26
2.4.1	Recursive integration scheme	27
2.4.2	Comparison with the composite approach	28
2.5	Conclusion: Towards a curvature-based shell model	29

Long and thin deformable structures, called in turn *strands*, *filaments*, *threads*, or *fibers* — and we shall use the latter name in this memoir¹ —, are well-spread in plants (foliage, stems), animals (hair, coral) and human-made objects (ropes, cables). Due to their smooth curved shape and complex way of deforming, characterized by many instabilities, fibers largely participate to the world’s visual richness and aesthetics.

A large body of our research was devoted to the design of numerical models for simulating the dynamics of fibers, the so-called *super-helix* and *super-clothoid* models. The goal was to discretize in space and time the continuous mechanical equations for *inextensible thin elastic rods*, namely the *Kirchhoff* equations, which take the form of second order partial differential equations subject to boundary conditions. Noting that curvatures and twist play a major role both in the geometric and dynamic description of this model, we have come up with a spatial rod discretization based on elements that are polynomial in such quantities.

Summary of contributions Our first scheme, the *super-helix* model, developed during my PhD thesis with Basile Audoly, relied on piecewise uniform curvatures [Bertails et al., 2006]. One major advantage of such curvature-based formulations is that the kinematics of the discretized rod remains, by construction,

¹More exactly, the term *fiber* will be retained to name our phenomenon of interest, while the term *thin elastic rod* will be used to designate the mathematical model we use to represent fibers.

perfectly inextensible. Such intrinsic inextensibility thus removes the burden of adding subsequent (stiff) inextensibility constraints when solving the dynamics. The price to pay, however, is that the geometry of the rod is not readily available but has to be computed iteratively from the curvatures. Moreover, the inertia matrix of such systems turns out to be dense, leading to a N^2 complexity when solving them at each time step, where N is a number of elements.

A first strategy to resolve the latter issue consisted in exploring higher-order spatial discretizations, namely piecewise linear curvatures [Bertails-Descoubes, 2012; Casati and Bertails-Descoubes, 2013], giving rise to the so-called *super-clothoid* model, so as to reduce the number N of elements needed for a given spatial accuracy. However, we were faced with the loss of a closed-form kinematics compared to the piecewise uniform case. Our main contribution was to devise an accurate integration scheme, based on power series expansion, which proved to be orders of magnitude faster compared to classical integration methods [Casati and Bertails-Descoubes, 2013].

A second strategy was to propose a linear, recursive time-integration scheme for super-helices [Bertails, 2009], inspired by the Featherstone linear-time algorithm for articulated chains of rigid bodies [Featherstone, 1983]. This approach allowed us to simulate long Kirchhoff rods made of hundreds elements at a much cheaper cost, and was also leveraged to design tree-like structures of rods in a very simple way.

Organization of the chapter Section 2.1 introduces prior work regarding the modeling and simulation of dynamic fibers, and motivates our choice for reduced, curvature-based discrete models. Section 2.2 briefly recalls equations for a Kirchhoff rod, with a particular stress on the structure of the kinematic problem known as the *Darboux* problem. Section 2.3 introduces the reduction of Kirchhoff equations to finite spatial dimension, and provides two different spatial schemes relying on a piecewise constant and piecewise linear shape function for the curvature, respectively. Section 2.4 gives the main ideas of our recursive integration algorithm, which may not only apply to the super-helix model, but to curvature-based discrete models of any order. Finally, Section 2.5 concludes this part and gives some insights into a possible extension of our work to the dynamic of surfaces.

2.1 Prior art and choice for reduced models

The scientific study of fibers has a long history in various fields, tracing back to the first continuous mechanics theories a few centuries ago to their further analysis in physics and mathematics, and their recent numerical treatment in Mechanical Engineering and Computer Graphics. Motivation originates from a number of applications ranging from the understanding of DNA supercoiling [Benham and Mielke, 2005] and climbing plants [Goriely and Neukirch, 2006] to the simulation of submarine cables [Goyal et al., 2008], surgery threads and needles [Pai, 2002; Chentanez et al., 2009], or hair [Ward et al., 2007].



Figure 2.1 – Many physical fibers exhibit a smooth curled geometry with linear-like curvature profile, which is captured and deformed accurately thanks to our super-clothoid model [Casati and Bertails-Descoubes, 2013]. From left to right and top to bottom, three examples of real fibers whose shapes are synthesized and virtually deformed in real-time using a very low number of 3D clothoidal elements: a vine tendril (4 elements), a hair ringlet (2 elements), and a curled paper ribbon (1 single element with an exaggerated flat cross-section). Left photograph courtesy of Jon Sullivan, pdphoto.org.

Theories for thin elastic rods Various theories were proposed in mechanics to model the equilibria and the dynamics of fibers, depending on the type of deformation considered. Our goal here is to capture the geometric richness of typical fiber deformations such as waving hair, coiling cables, curled ribbons or twining plants. These phenomena are largely *nonlinear*, dominated by *bending* and *twisting* elastic deformations, while stretching and shearing can be neglected. To account for this regime properly, we consider *inextensible* fibers with a vanishing rotation inertia, neglect shearing, and assume moment strains to remain small — making use of an *elastic* constitutive model — while large displacements, at the origin of the desired geometric nonlinearities, are allowed. The model is thus strictly subject to finite² rotations about the cross-section axes (bending) and about the tangent of the centerline (twisting). The corresponding governing equations — a set of partial differential equations together with boundary conditions — were first developed by Kirchhoff and Clebsch in their theory of thin elastic rods under finite displacements [Dill, 1992]. Within a more general framework on shells, rods and points, the Cosserat brothers [1909] later on proposed a clever mathematical representation of the rod geometry, relying on a space curve (the *centerline*) together with a *material* frame attached to the rod cross-section and continuously rotating along the centerline about the so-called *Darboux* rotation vector. A modern description of these theories can be found in [Antman, 1995; Audoly and Pomeau, 2010]. Pai [2002] was the first to introduce them to the Computer Graphics community.

²As opposed to infinitesimal.

Discretizing material rods In Mechanical Engineering, both *finite differences* and *finite elements* approaches were developed to discretize material rods in space and time. Though finite differences schemes have in principle the advantage of being easy to set up, properly accounting for the rod boundary conditions (typically, a clamped rod with the other end free) generally requires the use of a shooting strategy, which implies the solving of multiple nonlinear problems. Moreover, the stiff nature of the Kirchhoff equations, stemming from the presence of fourth-order spatial derivatives, imposes the use of overly small steps in time and space, or sophisticated implicit integrators [Goyal et al., 2008]. In contrast, a finite elements strategy allows one to single out spatial terms from time-evolving quantities, and provides a vast choice of elements to approximate them together with the boundary conditions. A popular method is the so-called geometrically exact beam approach [Reissner, 1973; Simo and Vu-Quoc, 1986], which derives an exact weak formulation for a generalized Kirchhoff rod with stretching and shearing, and finally discretizes the displacement and rotation fields with interpolating shape functions. One important issue of this approach, which spurred many subsequent works in the finite elements community, deals with the proper interpolation of rotations for preserving *objectivity*, i.e., invariance of the strain measures under rigid motion [Crisfield and Jelenić, 1998]. Moreover, regarding our specific needs here, this method is not directly applicable to the handling of *inextensible* and *unshearable* rods.


In Computer Graphics, finite differences schemes initially proposed by Pai [2002] to solve the statics of Kirchhoff rods were subsequently superseded with more robust schemes so as to deal with the full dynamic case. *Reduced*-coordinates models, based on a minimal parametrization of the system, were proposed to account for the exact kinematics of the rod, and especially to preserve inextensibility: this was the case for the articulated rigid body approach [Hadap and Magnenat-Thalmann, 2001; Hadap, 2006], parameterized by angular joints, and then by our super-models (detailed hereafter), parameterized by curvatures and twist. In contrast, further work focused on *nodal* models in order to get an explicit, point based representation of the centerline leading to a sparse mass matrix, at the price of adding external constraints to preserve the true kinematics. In the CoRde model [Spillmann and Teschner, 2007], both positions and orientations are considered as degrees of freedom. The Lagrange equations of motion are written for discrete approximations of kinetic and potential energies — including a stretch term — and orientations are coupled back to the centerline through soft constraints. Relying on the Bishop frame, Bergou et al. [2008] use a curve-angle parameterization to reduce the number of redundant parameters and guarantee that the orientation frame naturally remains adapted to the centerline. Discrete equations of motion are then established by leveraging principles from discrete differential geometry. Due to the choice of a nodal parameterization, inextensibility however needs to be explicitly enforced, e.g., through a fast projection scheme [Bergou et al., 2008] or a stiff stretch term [Bergou et al., 2010]. Finally, to ensure proper stability at an acceptable computational cost, an implicit scheme based on Newton’s method is advocated to discretize the nonlinear stiff bending and stretching forces [Bergou et al., 2010].

Reduced Lagrangian dynamics One advantage of reduced dynamics is that, by directly considering moment strains as actual degrees of freedom (i.e., material curvatures and twist instead of positions), the rod kinematics is exactly preserved, without redundancy and without adding any further constraint. Models parameterized by curvatures³ also benefit from an inexpensive implicit handling of bending forces, as those forces are linear in curvature. Finally, while multiple collision tricks — such as the position alteration technique [Baraff and Witkin, 1998] — were specifically developed for nodal models in Computer Graphics, more sophisticated, constraint-based frictional contact solvers naturally cope with reduced Lagrangian models (see Chapter 3) without having to worry about getting intermingled with external kinematic constraints.

2.2 Kirchhoff equations for thin elastic rods

Notation In what follows, s denotes the space variable and t the time variable. Space derivatives are represented by the prime symbol, so that $\mathbf{a}'(s, t) = \frac{\partial \mathbf{a}}{\partial s}$ and time derivatives by the dot symbol, so that $\dot{\mathbf{a}}(s, t) = \frac{\partial \mathbf{a}}{\partial t}$. For the sake of clarity, we may omit the time variable when describing the geometry of the rod. The special orthogonal group of dimension 3, denoted $SO(3)$, collects finite rotations of \mathbb{R}^3 (represented as direct orthogonal matrices) and is a *non commutative* Lie group.

We consider an inextensible and unshearable material rod of length L , represented by a centerline $\mathbf{r}(s)$ together with a material frame $\mathcal{R}(s)$, both parameterized by arc length $s \in [0, L]$. At location s , the vector $\mathbf{r}(s) \in \mathbb{R}^3$ gives the 3D position of the centerline and the rotation matrix $\mathcal{R}(s) \in SO(3)$ encodes the (unitary) tangent vector⁴ $\mathbf{n}_0(s) = \mathbf{r}'(s)$ as well as the two normal vectors $\mathbf{n}_1(s)$ and $\mathbf{n}_2(s)$ attached to the cross-section of the rod.



For simplicity, we assume that the rod is clamped at $s = 0$, and that its clamped position $\mathbf{r}(0) = \mathbf{r}_{\text{cl}}$ and orientation $\mathcal{R}(0) = \mathcal{R}_{\text{cl}}$ are given. We also assume that the end $s = L$ is free, that is, its position or velocity is not prescribed kinematically. Note that these two assumptions hold in most real fibers we wish to model, e.g., plants and hair. Otherwise, any of them could easily be dropped out; the former, by releasing \mathbf{r}_{cl} and \mathcal{R}_{cl} as degrees of freedom; the latter, by adding some constraint at the $s = L$ end.

³In this document we often use the shortcut *curvatures* to refer to *material twist and curvatures*.

⁴The material frame $\mathcal{R}(s)$ is said to be *adapted* to the curve $\mathbf{r}(s)$.

Kinematics From $s = 0$ to $s = L$, the material frame $\mathcal{R}(s)$ continuously evolves along the centerline $\mathbf{r}(s)$ through infinitesimal rotations about the so-called *Darboux vector* $\Omega(s)$ which represents the instantaneous space rotation vector of the rod. This space evolution mathematically writes

$$\mathcal{R}'(s) = [\Omega(s)]_{\times} \mathcal{R}(s), \quad (2.1)$$

where $[\mathbf{u}]_{\times}$ denotes the skew symmetric matrix corresponding to the vector cross product operator, i.e., $[\mathbf{u}]_{\times} \mathbf{v} = \mathbf{u} \times \mathbf{v}$. It is noteworthy that the local coordinates of Ω in the material frame represent the *material twist* κ_0 and *material curvatures* κ_1 and κ_2 of the rod, i.e., $\Omega(s) = \mathcal{R}(s) \boldsymbol{\kappa}(s)$, where $\boldsymbol{\kappa}(s) = [\kappa_0(s), \kappa_1(s), \kappa_2(s)]$ is called the *material curvature vector* (or simply *curvature vector*) in the remainder of the document. By further using properties of rotation matrices, one can reformulate Equation (2.1) as

$$\mathcal{R}'(s) = \mathcal{R}(s) [\boldsymbol{\kappa}(s)]_{\times}. \quad (2.2)$$

Finally, by compacting the centerline and the material frame into one single variable $\mathcal{F}(s) = \{\mathbf{r}(s); \mathcal{R}(s)\}$ and assuming $\boldsymbol{\kappa}(s)$ is fixed, the full kinematics of the rod can be formulated as an explicit⁵ linear first-order Cauchy-Lipschitz problem, referred to as the *Darboux problem* (see, e.g., [Ivanova, 2000]),

$$\begin{cases} \mathcal{F}'(s) = \{\mathbf{n}_0(s); \mathcal{R}(s) [\boldsymbol{\kappa}(s)]_{\times}\} \\ \text{with } \mathcal{F}(0) = \{\mathbf{r}_{\text{cl}}; \mathcal{R}_{\text{cl}}\} \text{ as initial conditions,} \end{cases} \quad (2.3)$$

which admits a unique solution. Note that the ambient space is not a vector space but rather a *nonlinear* differentiable manifold, since the kinematic relationship for the material frame operates onto the non commutative Lie group $SO(3)$. Due to non commutativity, the solution has no formal expression in the general case.

Dynamics Let ρ be the volumetric mass of the rod and S the surface area of its cross-section. We assume the rod is subject to external forces such as gravity or contact forces. Expressing the balance of linear and angular momentums on an infinitesimal portion of the rod and neglecting inertial momentum due to the vanishing cross-section leads to the following dynamic equations for a Kirchhoff rod,

$$\begin{cases} \mathbf{T}'(s) + \mathbf{p}(s) = \rho S \ddot{\mathbf{r}}(s) & (2.4a) \\ \mathbf{M}'(s) + \mathbf{n}_0(s) \times \mathbf{T}(s) = 0 & (2.4b) \end{cases}$$

where \mathbf{p} is the the sum of external forces distributed per unit length (for instance, in the case of gravity, $\mathbf{p} = -\rho S g \mathbf{e}_z$) and $\mathbf{T}(s)$ and $\mathbf{M}(s)$ are the internal force and moment respectively, transmitted from the free part of the rod through its cross-section at s .

⁵Coefficient of the highest derivative is 1.

Elastic constitutive law Dynamic equations are completed with a constitutive law that expresses the ability of the rod to bend and twist elastically,

$$\mathbf{M}(s) = \mathbb{K}_3 (\boldsymbol{\kappa}(s) - \bar{\boldsymbol{\kappa}}(s)) \quad \text{in the local basis } \mathcal{R}(s), \quad (2.5)$$

where $\mathbb{K}_3 = \text{diag}_3(\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2)$ is a diagonal 3×3 matrix collecting the twisting and bending stiffnesses \mathcal{K}_0 , \mathcal{K}_1 , and \mathcal{K}_2 respectively, and where $\bar{\boldsymbol{\kappa}}(s) \in \mathbb{R}^3$ collects the intrinsic twist and curvatures of the rod, used to model natural curliness (i.e., the shape the rod would have in the absence of external forces). As we are modeling a homogeneous material, both the linear mass density ρS and the stiffness matrix \mathbb{K}_3 are assumed to be constant with respect to space and time⁶. In contrast, the intrinsic curvature vector $\bar{\boldsymbol{\kappa}}(s)$ may vary spatially to account for a wide range of natural shapes.

Boundary conditions As mentioned earlier, we consider the rod to be clamped at end $s = 0$ and free at the other end $s = L$. Corresponding boundary conditions read

$$\left\{ \begin{array}{ll} \mathbf{r}(0) = \mathbf{r}_{\text{cl}} & \text{Enforced clamped position} & (2.6a) \\ \mathcal{R}(0) = \mathcal{R}_{\text{cl}} & \text{Enforced clamped frame} & (2.6b) \\ \mathbf{T}(L) = 0 & \text{No external force at free end} & (2.6c) \\ \mathbf{M}(L) = 0 & \text{No external torque at free end} & (2.6d) \end{array} \right.$$

Numerical model Equations (2.3–2.5) together with the boundary conditions (2.6) at $s = 0$ and $s = L$ form a nonlinear and stiff boundary value problem, which has no explicit solution in the general case and is known to be difficult to solve numerically.

2.3 Discrete curvature-based models

Realizing that *curvature* plays a key role in both the kinematics and the dynamics of the rod, an interesting idea consists in approximating the curvature vector with a simple, *polynomial* expression that is function of s . The coefficients of the polynomial are then taken as primary variables of the discrete model. One immediate consequence is that bending forces, which are linear in curvature, become linear in the discrete variables. Being stiff in nature, those forces can thus be treated implicitly in a straightforward manner, without having to solve a nonlinear problem. Furthermore, the kinematic (Darboux) problem becomes numerically *tractable*, and even solvable with a closed-form solution in the piecewise constant case.

⁶This assumption is debatable in some cases, like for instance vegetal stems for which it has been shown that the stiffness may vary spatially during the plant growing process [Beusmans and Silk, 1988]. Note however that such an assumption is not mandatory for building our discrete rod models in Section 2.3. Yet, it will be leveraged by our inversion process presented in Chapter 4.

Discrete kinematics In the following, we shall denote by $\mathbf{q} \in \mathbb{R}^m$ our state variable collecting the discrete curvatures and twists⁷, and $\bar{\mathbf{q}}$ the vector of same size storing the discrete intrinsic curvatures and twists. For now, we assume that the centerline of the rod \mathbf{r} can be computed as a function of s , \mathbf{q} , \mathbf{r}_{cl} and \mathcal{R}_{cl} , by solving the Darboux problem (2.3) with an accurate numerical method. The latter point is specifically addressed in Sections 2.3.1 and 2.3.2 in the case of the super-helix and the super-clothoid model, respectively. Formally differentiating the centerline twice with respect to time leads to the following expression for acceleration,

$$\ddot{\mathbf{r}}(s, t) = \ddot{\mathbf{r}}^*(s, t) + \dot{\mathbf{q}}(t) \frac{\partial^2 \mathbf{r}}{\partial \mathbf{q}^2}(s, t) \dot{\mathbf{q}}(t) + \frac{\partial \mathbf{r}}{\partial \mathbf{q}}(s, t) \ddot{\mathbf{q}}(t), \quad (2.7)$$

where $\ddot{\mathbf{r}}^*(s, t)$ is the acceleration generated by the clamping motion, which can be dropped when the clamped end is static. Expression (2.7) puts in evidence the linear dependence of the centerline acceleration $\ddot{\mathbf{r}}$ with respect to the generalized acceleration $\ddot{\mathbf{q}}$, a property that will be leveraged when designing a linear-time integration scheme in Section 2.4.

Discrete dynamics Discrete equations of motion result from a weak formulation of the strong Kirchhoff equations (2.4), where the trial functions are deduced from the constrained, piecewise polynomial kinematics. Consider an infinitesimal virtual displacement $\delta \mathbf{q}$ of our discrete degrees of freedom. This translates into a perturbation $\delta \boldsymbol{\kappa}(s)$ in the curvatures and twist function, which causes an infinitesimal rotation of the material frame about a virtual rotation vector $\delta \theta(s)$, such that $\delta \mathcal{R}(s) = [\delta \theta(s)]_{\times} \mathcal{R}(s)$, as well as an infinitesimal displacement $\delta \mathbf{r}(s)$ of the centerline. Applying the principle of virtual work [Reissner, 1973] while considering an inextensible rod as well as the boundary conditions (2.6), leads to the following weak formulation,

$$\int_0^L (\mathbf{M}'(s) + \mathbf{n}_0(s) \times \mathbf{T}(s)) \cdot \delta \theta(s) ds = 0,$$

where $\mathbf{T}(s) = \int_s^L (\mathbf{p}(s') - \rho S \ddot{\mathbf{r}}(s')) ds'$. Integrating by parts and noting that $\delta \boldsymbol{\kappa}(s) = \delta \theta'(s)$ and $[\delta \theta(s)]_{\times} \mathbf{n}_0(s) = (\delta \mathbf{r})'(s)$, we get

$$\int_0^L \mathbf{M}(s) \cdot \delta \boldsymbol{\kappa}(s) ds + \int_0^L \mathbf{p}(s) \cdot \delta \mathbf{r}(s) ds = \rho S \int_0^L \ddot{\mathbf{r}}(s) \cdot \delta \mathbf{r}(s) ds.$$

Finally, relating perturbed quantities to the virtual displacement $\delta \mathbf{q}$ and using the kinematic equation (2.7) yields the discrete dynamic equations

$$\mathbb{M}(\mathbf{q}) \ddot{\mathbf{q}} + \mathbb{K}(\mathbf{q} - \bar{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) + \mathbf{A}(\mathbf{q}, \dot{\mathbf{q}}) = 0 \quad (2.8)$$

⁷Discrete curvature and twists \mathbf{q} corresponds to per-element curvatures and twist for the super-helix model, and per-joint curvatures and twist for the super-clothoid model.

where

$$\begin{aligned}
\mathbb{M}(\mathbf{q}) &= \rho S \int_0^L \left(\frac{\partial \mathbf{r}}{\partial \mathbf{q}} \right)^\top \frac{\partial \mathbf{r}}{\partial \mathbf{q}} ds \\
\mathbf{G}(\mathbf{q}) &= - \rho S \int_0^L \left(\frac{\partial \mathbf{r}}{\partial \mathbf{q}} \right)^\top \mathbf{g} ds \\
\mathbf{A}(\mathbf{q}, \dot{\mathbf{q}}) &= \rho S \int_0^L \left(\frac{\partial \mathbf{r}}{\partial \mathbf{q}} \right)^\top \left(\dot{\mathbf{q}} \frac{\partial^2 \mathbf{r}}{\partial \mathbf{q}^2} \dot{\mathbf{q}} + \ddot{\mathbf{r}}^* \right) ds,
\end{aligned} \tag{2.9}$$

and where the stiffness matrix \mathbb{K} is a constant function linear in \mathbb{K}_3 , with a pattern depending on the order of discretization (see Sections 2.3.1 and 2.3.2).

The main challenge consists in evaluating vectors \mathbf{G} and \mathbf{A} and matrix \mathbb{M} in a both *accurate* and *fast* way. Sections 2.3.1 and 2.3.2 propose two different spatial discretizations — one piecewise constant and one piecewise linear with respect to s — and show how each of them addresses this problem.

Semi-implicit time-integration Finally, once the discrete dynamic equation (2.8) is assembled, we may use a semi-implicit Euler scheme to find \mathbf{q} at next time step, by simply taking linear terms as implicit and others as explicit, i.e., such that the time-discrete dynamic equation reads

$$\mathbb{M}(\mathbf{q}^t) \ddot{\mathbf{q}}^{t+1} + \mathbb{K}(\mathbf{q}^{t+1} - \bar{\mathbf{q}}) + \mathbf{G}(\mathbf{q}^t) + \mathbf{A}(\mathbf{q}^t, \dot{\mathbf{q}}^t) = 0. \tag{2.10}$$

As the elastic term of the dynamics appears to be the stiffer term of the equation, this cheap integration scheme proves sufficient to guarantee a good stability of the algorithm in most situations, even those involving some strong motion of the rod.

2.3.1 Piecewise uniform curvature: the super-helix model

The super-helix element We discretize the rod into N elements with arc lengths at joints denoted by s_i , $i \in \{0..N\}$. On each element E_i between two successive joints s_i and s_{i+1} , the curvature vector $\boldsymbol{\kappa}(s)$ is assumed to be constant with respect to arc length s , i.e.,

$$\boldsymbol{\kappa}(s) = \boldsymbol{\kappa}_i \quad \forall s \in [s_i, s_{i+1}].$$

The generalized coordinate vector \mathbf{q} of the system thus collects the $3N$ curvatures and twist values $\kappa_{i,j}$ with $j \in \{0..3\}$. The stiffness matrix is simply defined as

$$\mathbb{K} = \begin{pmatrix} \ell_0 \mathbb{K}_3 & 0 & \cdots & 0 \\ 0 & \ell_1 \mathbb{K}_3 & \ddots & \vdots \\ 0 & \ddots & \ddots & 0 \\ \vdots & \ddots & \ell_{N-1} \mathbb{K}_3 & 0 \\ 0 & \cdots & 0 & \ell_{N-1} \mathbb{K}_3 \end{pmatrix} = \mathbb{K}_3 \otimes \text{diag}_N(\boldsymbol{\ell}), \tag{2.11}$$

where $\boldsymbol{\ell} = (\ell_0, \dots, \ell_{N-1})^\top$ is the vector of lengths, $\text{diag}_N(\boldsymbol{\ell})$ is the diagonal matrix of the lengths of size N , and \otimes is the tensor (or Kronecker) product.

Closed-form kinematics As the curvature function $\kappa(s)$ is constant on each element E_i , it can be shown that the Darboux vector $\Omega(s)$ is constant too on each element, and that the Darboux problem (2.3) has an explicit solution on each element E_i , given by

$$\mathcal{R}(s) = \mathcal{R}(s_i) \exp(s [\kappa_i]_{\times}) \quad \text{and} \quad \mathbf{r}(s) = \mathbf{r}(s_i) + \int_{s_i}^{s_{i+1}} \mathcal{R}(s) \mathbf{e}_x \quad \text{with} \quad \mathbf{e}_x = [1 \ 0 \ 0]^{\top},$$

which can be computed in *closed-form*, i.e., using only basic and cheap unitary operations. In our case, the above expressions can be evaluated using only cosine and sine functions, and it is noteworthy that the centerline $\mathbf{r}(s)$ takes the form of a perfect *circular helix* on each element E_i . Starting integration at $s = 0$ with clamping conditions (2.6a – 2.6b) and then using continuity conditions for $\mathcal{R}(s)$ and $\mathbf{r}(s)$ at each joint s_i , we can integrate the full kinematics of the rod from $s = 0$ to $s = L$. The super-helix model then features a C^1 -smooth piecewise helical centerline, and all the terms of the kinematics (material frame, centerline, but also deformation gradients), and finally all the terms of the dynamics (2.8), can be computed in closed-form⁸.

Limitations A discrete rod with piecewise constant curvature may still represent a rather rough approximation of the continuous case, with an improper degree of continuity at the joints (see Figure 2.2). Instead of using an excessively refined primitive, one may think it would be worth designing a richer, higher-order element with *linear* curvatures, that would better stick to the actual curvature profile of real fibers and guarantee visually pleasing smoothness (i.e., at least C^2 -continuity) of the centerline *at any resolution*. One becomes unfortunately faced with the loss of a formal expression for the kinematics. Yet, observing that the Darboux problem still possesses a lot of structure, we show in the following that such a *space clothoid* element can be conveniently derived. The key is to introduce a fast and accurate integration scheme based on an adaptive power series summation algorithm. This numerical algorithm is then used as a formal computation tool to evaluate the spatial terms of the dynamics at a high precision.

2.3.2 Piecewise linear curvature: the super-clothoid model

In the piecewise linear case, where each element takes the form of a so-called *3D clothoid*⁹, we were able to build an accurate integration scheme which proved to be

⁸Computing each term of the dynamics (2.8) amounts to evaluating the integral over $[0, L]$ of cosine and sine functions, and products of cosine and sine functions.

⁹A *clothoid* or *Euler spiral* is a planar curve characterized by a linear curvature w.r.t. the arc length. By extension to 3D, the term *3D Euler spiral* was coined by Harary and Tal [2012] to name a space curve characterized by linearly *geometric* curvature and torsion. In our case we rather use the term *space clothoid* to name the centerline of a rod characterized by linear *material* curvatures and twist. The centerline of our rod element is actually more general than the so-called 3D Euler spiral, the entire class of so-called 3D Euler spirals being obtained by cancelling the first material curvature. Note that this difference will be discussed again in Chapter 4.




Figure 2.2 – Comparison of *fairness* (or smoothness) between (a) the super-helix model and (b) the super-clothoid model. Whereas joints between elements are particularly visible for the super-helix model (5 elements here), the super-clothoid model generates a smooth, visually pleasing shape even at a very coarse resolution (2 elements here).

orders of magnitude faster compared to classical integration methods [Casati and Bertails-Descoubes, 2013]. The key of our approach was to leverage the form of the solution as a power series expansion, while avoiding the pitfall of catastrophic cancellation through an adaptive integration strategy. With this tool in hand, we were able to demonstrate that the super-clothoid model could capture intricate shapes both robustly and efficiently, with better spatial accuracy and geometric fairness compared to state-of-the-art methods (see Figure 2.1).

The super-clothoid element As before we discretize the rod into N elements. However, now the discrete curvature variables $\hat{\boldsymbol{\kappa}}_i$ are located at joints s_i , and the curvature vector \boldsymbol{q} collects $3(N + 1)$ degrees of freedom. On each element E_i of length ℓ_i , the curvature vector $\boldsymbol{\kappa}(s)$ is assumed to vary linearly with arc length s , so that its expression reads

$$\boldsymbol{\kappa}(s) = \left(1 - \frac{s - s_i}{\ell_i}\right) \hat{\boldsymbol{\kappa}}_i + \frac{s - s_i}{\ell_i} \hat{\boldsymbol{\kappa}}_{i+1} \quad \forall s \in [s_i, s_{i+1}].$$

The stiffness matrix is not diagonal anymore but becomes tridiagonal,

$$\mathbb{K} = \begin{pmatrix} \frac{\ell_0}{3} \mathbb{K}_3 & \frac{\ell_0}{6} \mathbb{K}_3 & 0 & \cdots & 0 \\ \frac{\ell_0}{6} \mathbb{K}_3 & \frac{\ell_0 + \ell_1}{3} \mathbb{K}_3 & \ddots & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & \frac{\ell_{N-2} + \ell_{N-1}}{3} \mathbb{K}_3 & \frac{\ell_{N-1}}{6} \mathbb{K}_3 \\ 0 & \cdots & 0 & \frac{\ell_{N-1}}{6} \mathbb{K}_3 & \frac{\ell_{N-1}}{3} \mathbb{K}_3 \end{pmatrix} = \mathbb{K}_3 \otimes \text{tridiag}_N(\boldsymbol{\ell}), \quad (2.12)$$

where $\text{tridiag}_N(\boldsymbol{\ell})$ is a tridiagonal matrix of size N linearly depending on the lengths.

Kinematic integration using power series expansion We first focus on a single element with index $i = 0$. When the curvature vector is polynomial of degree ≥ 1 , the key idea is to formulate the solution of the Darboux problem (2.3) as a *power series expansion* (PSE). Indeed, invoking Cauchy's theorem on analytic solutions of linear ODEs with analytic coefficients (see, e.g., [Poole, 1936, §2]), we deduce that the solution \mathcal{F} of (2.3) is C^∞ and admits a power series expansion

$\mathcal{F}(s) = \left\{ \sum_{n=0}^{\infty} \mathbf{r}_n s^n; \sum_{n=0}^{\infty} \mathcal{R}_n s^n \right\}$ on \mathbb{R} . In our particular case where the curvature is a polynomial of degree 1 (space clothoid element), the general term of the series $\left\{ \tilde{\mathbf{r}}_n(s) = \mathbf{r}_n s^n; \tilde{\mathcal{R}}_n(s) = \mathcal{R}_n s^n \right\}$ is recursively defined as

$$\begin{cases} \tilde{\mathcal{R}}_0 = \mathcal{R}_{\text{cl}} \\ \tilde{\mathcal{R}}_1(s) = s \tilde{\mathcal{R}}_0 [\hat{\boldsymbol{\kappa}}_0]_{\times} \\ \tilde{\mathcal{R}}_{n+2}(s) = \frac{s}{n+2} \left(\tilde{\mathcal{R}}_{n+1}(s) [\hat{\boldsymbol{\kappa}}_0]_{\times} + s \tilde{\mathcal{R}}_n(s) [\boldsymbol{\gamma}]_{\times} \right) & \forall n \in \mathbb{N} \\ \tilde{\mathbf{r}}_0 = \mathbf{r}_{\text{cl}} \\ \tilde{\mathbf{r}}_{n+1} = \frac{s}{n+1} \tilde{\mathcal{R}}_n \mathbf{e}_{\mathbf{x}} & \forall n \in \mathbb{N}, \end{cases} \quad (2.13a)$$

$$\begin{cases} \tilde{\mathbf{r}}_0 = \mathbf{r}_{\text{cl}} \\ \tilde{\mathbf{r}}_{n+1} = \frac{s}{n+1} \tilde{\mathcal{R}}_n \mathbf{e}_{\mathbf{x}} & \forall n \in \mathbb{N}, \end{cases} \quad (2.13b)$$

where ℓ_0 is the length of the element, $\hat{\boldsymbol{\kappa}}_0$ and $\hat{\boldsymbol{\kappa}}_1$ are its initial and final curvatures respectively, and $\boldsymbol{\gamma}$ is the curvature slope, $\boldsymbol{\gamma} = \frac{\hat{\boldsymbol{\kappa}}_1 - \hat{\boldsymbol{\kappa}}_0}{\ell_0}$. Note that according to (2.13b), computing the centerline follows from that of the material frame, the latter involving the recursive sequence (2.13a) of second order.

Furthermore, one important property of the above power series expansion is that only the very first terms of the series are relevant, the following ones rapidly decreasing in norm and falling below machine precision. This means that the solution \mathcal{F} can be evaluated accurately through a truncated series containing a small number of terms (in practice, around 100 terms). This nice property is due to the simple structure of our kinematic problem (2.3), which formulates as an explicit linear ODE with polynomial coefficients. In this case indeed, we can prove that the general term of the series decreases super-linearly to zero as n tends to infinity [Neher, 1999].

Catastrophic cancellation issue Unfortunately, when numerically evaluating the sum of the relevant terms in finite precision, one is inevitably faced with round-off issues leading to huge approximation errors. Consider expression $(1 + y) - y$ which should be equal to 1 whatever the value of y . In floating-point arithmetic, this equality only holds if y is close enough to 1. In double precision for example, take $y = 10^{16}$ and compute the expression above. The numerical result is 0.0, yielding a relative error of 100%. This error is the consequence of first, an *absorption* phenomenon when computing the sum $1 + 10^{16}$, which, due to machine overflow when aligning mantissa, is approximated to 10^{16} . Then, a *cancellation* phenomenon when subtracting 10^{16} . Such unfortunate combination of absorption and cancellation leads to erroneous results and for this reason is called *catastrophic cancellation*. Details on floating-point arithmetic can be found in, e.g., [Goldberg, 1991].




Figure 2.3 – 2.3a: Dramatic loss of precision when naively summing power series of the kinematics. 2.3b: In contrast, our piecewise summation algorithm guarantees high precision of the summation. 2.3c: A long and highly curved space clothoid integrated with our piecewise computation method, using 109 subdivisions.

In our case, catastrophic cancellation problems precisely occur when naively computing the sum of our power series $\sum \tilde{\mathcal{R}}_n(s)$ for a long and/or curly rod (see Figure 2.3a). To explain this phenomenon, we have plotted in Figure 2.4 (in blue) the norm of the general term $\tilde{\mathcal{R}}_n(s)$ as a function of n , for different values of s . The resulting “hillock”-like profile implies that when computing the sum of the series, one actually adds very small values together with very large ones in norm, the widest range being obtained when getting to the top of the hillock. Note that the larger s , the higher the top of the hillock. More precisely, we have shown in [Casati and Bertails-Descoubes, 2013] that the top of the hillock $\mathcal{H}(s)$ grows quasi-exponentially with the increasing function $\lambda(s)$ defined as

$$\lambda(s) = 2s(\|\hat{\boldsymbol{\kappa}}_0\|_\infty + s \|\boldsymbol{\gamma}\|_\infty).$$

Moreover, looking back at Recursion (2.13a), one notes that entries of the matrices to be added are of *alternating* sign, due to the product with skew symmetric matrices. This results in cancellation when computing the sum. All this combined together, it is then not surprising that we are faced with a catastrophic cancellation issue when $\lambda(s)$ (and thus s) becomes *too large*. As $\lambda(s)$ increases with s as well as with intrinsic curvatures, we now understand why numerical issues show up for a long and/or curly rod. Luckily enough, as shown in the following, we can avoid catastrophic cancellation by devising an *adaptive* summation algorithm.

Fast and Precise Power Series Summation To avoid catastrophic cancellation, a natural idea consists in upper-bounding x by a value M depending on the machine precision, so that the top of the hillock remains within the range where additions between two numbers can be safely performed, i.e., with no absorption of their leading digit. More precisely, if the machine has precision 10^{-d} ($d = 7$ for a floating number encoded on 32 bits, $d = 16$ on 64 bits), then the top of the hillock




Figure 2.4 – Hillock-like profile of the general term $\tilde{\mathcal{R}}_n(s)$ in norm. In blue: Evolution of $\|\tilde{\mathcal{R}}_n(s)\|_\infty$ function of n (at fixed s), in log scale. As expected [Neher, 1999], the decreasing towards 0 appears to be super-linear. In red: Evolution, function of s , of the upper-bound $\mathcal{H}(s)$, in log scale. Note that the plot of this upper-bound visually matches the maximum function $\max_n \|\tilde{\mathcal{R}}_n(s)\|_\infty$, meaning that the top of the hillock grows quasi-exponentially with s .

should be bounded by $10^{\frac{d}{2}}$ so as to be able to safely cover additions on the range $[10^{-\frac{d}{2}}, 10^{\frac{d}{2}}]$. We have proved that a sufficient upper-bound for M is

$$M \leq \max \left\{ n \in \mathbb{N} \text{ s.t. } (n+1)^n \leq 10^{\frac{d}{2}} n! \right\}. \quad (2.14)$$

One obtains $M \leq 19$ for $d = 16$. In practice, we set M to 10 to maintain good precision across summation. This choice allowed us to reach high precision for all the summations we have computed.

Adaptive piecewise summation Our full summation method then relies on an automatic subdivision of the integration domain into subintervals, on which integration can be safely performed. More precisely, our adaptive summation algorithm consists in splitting the evaluation domain $[0, \ell_0]$ into p adaptive subintervals $[0, \sigma_1], [\sigma_1, \sigma_2], \dots, [\sigma_{p-1}, \ell_0]$ such that $\sigma_{i+1} = \sigma_i + s_{\max}(\sigma_i)$, where $s_{\max}(\sigma_i)$ is provided by this upper-bound function,

$$s_{\max}(\sigma) = \begin{cases} \frac{\sqrt{\|\kappa(\sigma)\|_\infty^2 + 2M\|\gamma\|_\infty} - \|\kappa(\sigma)\|_\infty}{2\|\gamma\|_\infty} & \text{if } \gamma \neq 0 \\ \frac{M}{2\|\kappa(\sigma)\|_\infty} & \text{else if } \kappa(\sigma) \neq 0 \\ +\infty & \text{otherwise.} \end{cases} \quad (2.15)$$

On each subinterval, summation is thus guaranteed to be performed with good accuracy (see an illustration in Figure 2.5).




Figure 2.5 – Our piecewise summation algorithm applied to the rod’s kinematics, for one clothoidal element. The length of each subinterval which guarantees a safe evaluation of the geometry is automatically provided by our method.

We show in [Casati and Bertails-Descoubes, 2013] that our adaptive piecewise summation algorithm reaches a good precision (up to the machine precision) *orders of magnitude faster* compared to classical integration schemes such as Euler or Runge Kutta of order 2 and 4.

Finally, using continuity conditions at joints, the kinematics of the full rod can be accurately integrated from $s = 0$ to $s = L$.

Propagation to the terms of the dynamics To compute all the terms of the discrete dynamic equation (2.8), we need to perform several operations on our kinematic terms, such as linear combination, scalar product, differentiation with respect to \mathbf{q} , or integration with respect to s . All these operations are actually easy to perform using power series expansion; furthermore, we show that our precision guarantee resists all these operations, thus providing us with a fast and accurate method for computing all the terms of the dynamics.

Model comparison We have compared the super-clothoid model against the super-helix model and the discrete elastic rod model [Bergou et al., 2008], by performing a simple dynamic experiment: a naturally curly rod unfolds and falls under gravity. At high resolution, we have first observed that the three models converge towards exactly the same equilibrium. Then we have measured, for each model, the spatial accuracy reached at equilibrium, function of the resolution used; and similarly, the computational time made to reproduce the full dynamic experiment, function of the resolution. Results are compiled and cross-viewed in Figure 2.6. They show that for curly configurations, the two super-models provide a better trade-off in terms of accuracy/efficiency compared to the discrete elastic rod model. For high accuracy ranges, the super-clothoid model also offers a better alternative compared to the super-helix model. Moreover, as mentioned before, the great advantage of the super-clothoid model is to yield a visually pleasing, C^2 -smooth geometry, at any resolution.

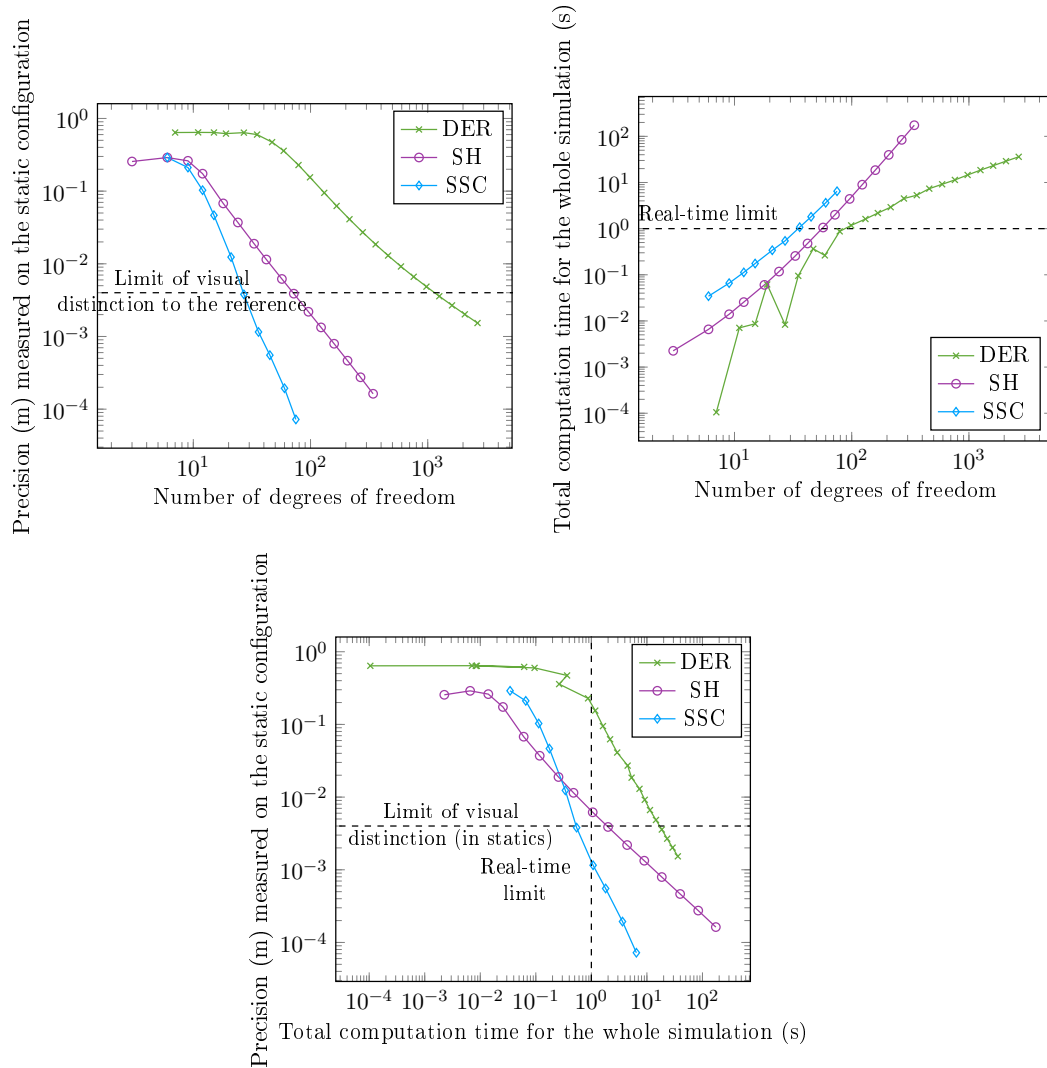


Figure 2.6 – Comparisons between the super-clothoid model (*SSC*), the super-helix model (*SH*), and the discrete elastic rod model (*DER*), in terms of accuracy and efficiency.

2.4 Linear-time integration of the dynamics

One major advantage of the aforementioned curvature-based formulations is that the kinematics of the discretized rod remains, by construction, perfectly inextensible. Such intrinsic inextensibility thus removes the burden of adding subsequent (stiff) inextensibility constraints when solving the dynamics. The price to pay, however, is that the geometry of the rod is not readily available but has to be computed iteratively from the curvatures. Moreover, the inertia matrix of such systems turns out to be dense, leading to a N^2 complexity when solving them at each time step, where N is a number of elements. Exploring high-order discrete models such as the clothoidal element helped us reduce the number of elements required for a given

precision; however, for a very long rod requiring high spatial precision, the N^2 computation would remain a bottleneck.

One strategy to deal with the latter issue was to propose a linear, recursive scheme for curvature-based discrete models, inspired by the Featherstone algorithm for articulated chains of rigid bodies [Featherstone, 1983]. We have investigated such a linear strategy in [Bertails, 2009] in the case of a super-helix, but our algorithm could be similarly applied to higher-order discrete models such as the super-clothoid model. Below we explain the general principle of the method.

2.4.1 Recursive integration scheme

Non-local dynamics We consider again the time-and-space discrete dynamic equations (2.10) for a super-model, which are recalled here, with the two unknowns (implicit variables) marked at time $t + 1$,

$$\underbrace{\mathbb{M}(\mathbf{q})\ddot{\mathbf{q}}^{t+1}}_{\text{non-local}} + \underbrace{\mathbb{K}(\mathbf{q}^{t+1} - \bar{\mathbf{q}})}_{\text{local}} + \mathbf{G}(\mathbf{q}) + \mathbf{A}(\mathbf{q}, \dot{\mathbf{q}}) = 0. \quad (2.16)$$

Let E_i be the i^{th} element of the rod. While the i^{th} block of the elastic term is local to the element and only depends on the i^{th} block of generalized coordinates \mathbf{q}_i , the block of the inertial term depends on *all* generalized accelerations $\ddot{\mathbf{q}}$, as $\mathbb{M}(\mathbf{q})$ is obviously dense from its expression in (2.9), knowing that the term $\frac{\partial \mathbf{r}}{\partial \mathbf{q}}$ evaluated at $s \in E_i$ depends on all previous coordinates $\mathbf{q}_{j,j \leq i}$ due to the recursive kinematics of super-models.

Building local dynamic equations from tip to root The key of our approach, detailed in [Bertails, 2009], is to show that Equation (2.16) can be rewritten *locally* on each element E_i as

$$\mathbb{K}_i(\mathbf{q}_i - \bar{\mathbf{q}}_i) = \mathbb{A}_i \ddot{\mathbf{q}}_i + \mathbb{R}_i \ddot{\mathbf{r}}_i(0) + \mathbb{C}_i \ddot{\boldsymbol{\theta}}_i(0) + \mathbf{b}_i, \quad (2.17)$$

where $\ddot{\mathbf{r}}_i(0)$ and $\ddot{\boldsymbol{\theta}}_i(0)$ are the linear and angular accelerations of the first end of element E_i , respectively, and the 3×3 matrices \mathbb{A}_i , \mathbb{R}_i and \mathbb{C}_i are called *cumulate inertias*, and the vector \mathbf{b}_i *cumulate force* of the rod.

As in [Featherstone, 1983], the cumulate inertias and the cumulate force of the rod can be computed recursively from the free end ($i = N - 1$) to the clamped end ($i = 0$) of the rod. This can be proved by induction, using the kinematic property (2.7) according to which the linear (and also angular) acceleration of the rod is linear w.r.t. the generalized acceleration $\ddot{\mathbf{q}}$.

Recursive solving from root to tip Finally, once each local dynamic equation (2.17) is built, one may simply integrate the dynamics recursively from root to tip, using continuity conditions at the joints between elements.

2.4.2 Comparison with the composite approach

The *recursive* approach presented above brings several advantages compared to the so-called *composite* approach detailed in Section 2.3, in which full terms of the dynamics are assembled.

First, we have observed a computational gain of about $\times 20$ for rods made of more than 20 elements. In practice, this allowed us to simulate very long rods in real-time (see Figure 2.7(a)). Then, by simply adapting boundary conditions at the joints between elements, we were able to extend the recursive approach quite straightforwardly to the handling of *tree-like* structures (see Figure 2.7(b)). Note that tree-like structures would have been much less trivial to model using the composite approach.

On the downside, the recursive approach may require a smaller timestep to be stable. This is due to the explicit computation of $\ddot{\mathbf{q}}_i$ in the first pass of our algorithm, which makes it slightly less stable than the composite method for an equivalent time step. In the case of articulated rigid bodies, the explicit solving of Featherstone’s algorithm was pointed out by Hadap [2003] who proposed an interesting framework for enforcing an implicit solving [Hadap, 2006]. Note however that in our case, the difference in stability between the two methods keeps on being small.

Finally, we shall not use the recursive formalism of super-models in the next chapter, which is dedicated to the introduction of frictional contact for which composite models are better suited for. However, we believe that the recursive model would be particularly useful for modeling botanical or biological structures. In the latter case, a promising extension would be to design an adaptive scheme as in [Redon et al., 2005], by relying upon Featherstone’s divide-and-conquer algorithm [Featherstone, 1999a,b]. So far, to the best of our knowledge, such adaptive algorithms were leveraged for articulated rigid bodies only. Adaptive super-models could be very useful for manipulating in real-time long and complex biological filaments (e.g., macro-molecules) prone to bending and twisting, by automatically tuning their number of active degrees of freedom.




Figure 2.7 – (a) Interactive manipulation of a long super-helix made of 30 helical elements. (b) Interactive animation of a weeping willow swept by wind, represented as a tree-like structure of super-helices (composed of 242 helical elements) with various stiffness values.

2.5 Conclusion: Towards a curvature-based shell model

We have investigated new discrete models for solving the Kirchhoff dynamic equations for thin elastic rods. All our models share a curvature-based spatial discretization, allowing them to capture inextensibility of the rod intrinsically, without the need for adding any kinematic constraint. Moreover, elastic forces boil down to linear terms in the dynamic equations, making them well-suited for implicit integration. We shall furthermore show in Chapter 4 that this property will be leveraged for solving inverse design problems.

Interestingly, our discretization methodology can be interpreted from two different points-of-view. From the finite-elements point-of-view, our strain-based discrete schemes can be seen as discontinuous Galerkin methods of zero and first orders. From the multibody system dynamics point-of-view, our discrete models can be interpreted as deformable Lagrangian systems in finite dimension, for which a dedicated community has started to grow recently [Sugiyama et al., 2013]. We note that adopting the multibody system dynamics point-of-view helped us formulate a linear-time integration scheme, which had only been investigated in the case of multibody rigid bodies dynamics so far.

Towards a curvature-based shell model Recently, we have started to investigate similar high-order modeling strategies for *surfaces*, in particular for the case of inextensible and developable shells. In [Blumentals et al., 2016a] we have built an inextensible shell patch by taking as degrees of freedom the curvatures of its mid-surface, expressed in the local frame. As in the super-helix model, we show that when taking curvatures uniform over the element, each term of the equations of motion may be computed in closed-form; besides, the geometry of the element corresponds to a cylinder patch at each time step (see Figure 2.8). Compared to the 1D (rod) case however, some difficulties arise in the 2D (plate/shell) case, where compatibility conditions are to be treated carefully. Moreover, the question of assembling such shell patches to build a full inextensible and developable shell model remains an open problem.




Figure 2.8 – (a) Cylindrical geometry of our shell patch with uniform material curvatures. (b-d) Dynamical oscillations of our inextensible shell patch clamped at bottom right corner and subject to gravity, with an initially curved intrinsic shape.

Frictional contact models for large multibody systems

Contents

3.1	Prior art and choice for nonsmooth models	33
3.2	Nonsmooth contact dynamics for thin elastic objects	36
3.2.1	Adding unilateral contact with Coulomb friction	36
3.2.2	The Discrete Frictional Contact Problem (DFCP)	39
3.2.3	Multiple equivalent formulations of Signorini-Coulomb	41
3.2.4	A local hybrid solving strategy	46
3.3	Continuum modeling of granular materials	48
3.3.1	Discrete vs. continuum models for granular matter	48
3.3.2	Macroscopic yield criteria	50
3.3.3	Nonsmooth rheology $\mathcal{DP}(\mu)$	51
3.3.4	Extension to a varying volume fraction	54
3.4	Conclusion: Towards a continuum-based model for fibrous materials	54

The second core topic of my research activity deals with the modeling and simulation of dry (Coulomb) frictional contact within (potentially large) collections of thin structures, such as rigid grains, thin elastic rods or plates/shells. Target applications include the simulation of fibrous and granular materials (such as human hair and sand, respectively), and more recently, cloth. For all these dynamical systems, dry frictional contact plays a major role in their visual appearance (see Figure 3.1).

Summary of contributions Our initial motivation for conducting research in frictional contact algorithms was to improve the realism of *hair* simulations, where self-interactions were simply neglected in the worse case [Anjyo et al., 1992], and regularized in the best case [McAdams et al., 2009], yielding an overly smooth appearance of hair. We have started our study by modeling hair as a collection of individual thin elastic rods interacting through contact and friction — a method known as the *discrete element method*. However, instead of using common compliant and regularized models for contact and friction, we have followed the *nonsmooth contact dynamics* approach developed by Moreau and Jean a few decades ago [Moreau, 1988; Moreau and Jean, 1996; Jean, 1999], in which emerging effects of



Figure 3.1 – Taking into account dry frictional contact is of utmost importance for capturing significant visual features such as a stable sand pile or complex hair patterns.

contact and solid friction are integrated at the macroscopic scale within a *rigid* and *nonsmooth* law. In this framework, the discrete frictional contact problem (named hereafter **DFCP**) possesses various possible formulations, from the root-finding of complementarity merit functions to a sequence of parameterized convex problems. We have first studied the numerical counter-parts for each main formulation and evaluated how each of them performs in terms of efficiency, robustness, and scalability [Bertails-Descoubes et al., 2011]. Facing robustness and scalability issues when using state-of-the-arts approaches, we have then designed a new frictional contact solver by coupling an iterative Gauss-Seidel strategy together with an extremely robust one-contact solver relying upon a hybrid local solver [Daviet et al., 2011]. This new solver proved to converge well in scenarios involving thousands fibers subject to tens thousands frictional contact points, and thus allowed us to enhance considerably the realism of hair simulations (see Figure 3.8). Recently, we have adapted this solver for coping with frictional contact in *cloth* modeled as a *nodal* system [Daviet et al., 2015] (see Figure 3.9).

In parallel to our work on the discrete element modeling of fiber assemblies, we have started to investigate *continuum* modeling strategies so as to scale up simulations to hundreds thousands fibers — the actual order of magnitude of a full head of hair. Faced with the glaring lack of macroscopic models for dry fibrous materials in the literature, we have started this long-term investigation by considering the simpler case of *granular* matter, where each element is not a deformable fiber but a rigid and isotropic grain. We have modeled a large assembly of

grains as a dilatible viscoplastic fluid combined with a Drucker-Prager yield criterion (which stands for a generalization of the Coulomb friction law to a continuum) and a unilateral compressibility constraint, first in the dense regime [Daviet and Bertails-Descoubes, 2016b], then accounting for a varying volume fraction [Daviet and Bertails-Descoubes, 2016a]. Thanks to a semi-implicit time-stepping scheme and a careful spatial discretization of our rheology, we were able to preserve at each time step the exact coupling between normal and tangential stresses, in a stable way. This contrasts with previous approaches that either regularize or linearize the yield criterion for implicit integration, leading to artificial sliding motion or visible grid artifacts [Chambon et al., 2011; Narain et al., 2010]. Remarkably, our discrete problem turned out to be very similar to the discrete frictional contact problem (**DFCP**) that we had formerly encountered when dealing with discrete element modeling, which allowed us to leverage again our robust nonsmooth solver [Daviet et al., 2011]. Our method was qualitatively validated by successfully capturing typical macroscopic features of some classical experiments, such as the discharge of a silo or the collapse of a granular column.

Organization of the chapter Section 3.1 motivates our choice for modeling frictional contact using rigid and nonsmooth frictional contact laws, in the context of both multibody and continuum dynamics. Section 3.2 first formulates the discrete element problem where a large number of deformable structures are interacting through non-penetration constraints combined with the nonsmooth Coulomb friction law, then spells out the specificity of our resulting discrete problem (**DFCP**), and discusses robust and efficient numerical solutions for solving it. Following a similar outline, Section 3.3 presents the nonsmooth Drucker-Prager law for pressure-dependent yield-stress fluids, combined with a non-compressibility condition, and provides insights into the semi-implicit discretization of the full dynamical problem, which turns out to share a structure similar to that of our **DFCP**.

3.1 Prior art and choice for nonsmooth models

Most popular approaches in Computer Graphics and Mechanical Engineering consist in assuming that the objects in contact are locally *compliant*, allowing them to slightly penetrate each other. This is the principle of *penalty-based methods* (or *molecular dynamics*), which consists in adding mutual repulsive forces of the form $k\mathbf{f}(\delta)$, where δ is the penetration depth detected at current time step [Cundall, 1971; Moore and Wilhelms, 1988]. Though simple to implement and computationally efficient, the penalty-based method often fails to prevent excessive penetration of the contacting objects, which may prove fatal in the case of thin objects as those may just end up traversing each other (see Figure 3.2(a)). One solution might be to set the stiffness factor k to a large enough value, however this causes the introduction of parasitical high frequencies and calls for very small integration steps [Baraff, 1989]. Penalty-based approaches are thus generally not satisfying for ensuring robust contact handling.




Figure 3.2 – Simulation of complex geometrical configurations involving self-contact, such as a knot (a) and a plectoneme (b). With a classical penalty-based method, tying the knot causes the contact to break, whereas contact resists arbitrarily large user forces when using an implicit constraint-based method, without having to tune an extra stiffness parameter.

In the same vein, the friction law between solid objects, or within a yield-stress fluid (used to model foam, sand, or cement, which, unlike water, cannot flow beyond a certain threshold), is commonly modeled using a *regularized* friction law (sometimes even with simple viscous forces), for the sake of simplicity and numerical tractability (see e.g., [Spillmann and Teschner, 2008; Frigaard and Nouar, 2005]). Such a model cannot capture the threshold effect that characterizes friction between contacting solids or within a yield-stress fluid. The nonsmooth transition between sticking and sliding is however responsible for significant visual features, such as the complex patterns resting on the outer surface of hair, the stable formation of sand piles, or typical stick-slip instabilities occurring during motion (see Figure 3.1).

After having employed such regularized strategies during my PhD, the search for a realistic, robust and stable frictional contact method encouraged me to depart from those, and instead to focus on *rigid* contact models coupled to the *exact* nonsmooth Coulomb law for friction (and respectively, to the *exact* nonsmooth Drucker-Prager law in the case of a fluid), which better integrate the effects of frictional contact at the macroscopic scale. Those modeling choices are summarized in Figure 3.3.

For stability and consistency purposes, nonsmooth models require *implicit* integration schemes¹ when being simulated numerically. Resulting discrete problems then take the form of algebraic equations subject to *complementarity set-valued constraints*, which are better understood and manipulated using convex analysis tools, the latter having been mainly developed in the early 60’s by Jean-Jacques Moreau and R. Tyrrel Rockafellar. Such nonsmooth methods being disregarded in

¹This means that at each time step, contact forces \mathbf{r} (and the stress tensor $\boldsymbol{\sigma}$ respectively) are unknowns of the problem, together with the relative velocities \mathbf{u} of the problem (and the strain rate tensor $\dot{\boldsymbol{\epsilon}}$ respectively).




Figure 3.3 – Summary of our nonsmooth modeling choices for dealing with frictional contact (right column), and comparison with their most simple regularized counterparts (left column). Notation: $h(\mathbf{q})$ is the gap function between the two contacting objects, r_N and r_T are the normal and tangential components of the contact force, respectively, \mathbf{u}_T is the tangential component of the relative velocity between the two contacting objects, $\boldsymbol{\sigma}_T$ the tangential stress tensor of the continuum, and $\dot{\boldsymbol{\epsilon}}_T$ the strain rate tensor.

most Computer Graphics groups² a decade ago (in favor of regularized or explicit

²It is interesting however to note that nonsmooth contact mechanics have been introduced in Computer Graphics by David Baraff as soon as in the late 80's [Baraff, 1989, 1991, 1993, 1994]. But it has not been followed up by the community at that time, then considerably losing ground on the Mechanical Engineering community in this area. In Computer Graphics, nonsmooth mechanics has slowly regained popularity in the context of rigid body dynamics, notably with the work of Kauman and colleagues [2005; 2008].

methods), their assessment was my main motivation for integrating in 2007 the Inria project-team BiPop, specialized in the *implicit* treatment of frictional contact within the nonsmooth mechanics framework [Acary and Brogliato, 2008].

3.2 Nonsmooth contact dynamics for thin elastic objects

In this section we present a *discrete element model* for an assembly of thin elastic objects (typically, fibers), where frictional contacts are modeled and time-discretized using the nonsmooth contact dynamics framework [Moreau and Jean, 1996]. The resulting Discrete Frictional Contact Problem (**DFCP**) is specifically solved for by a new robust and scalable solver in the case of a large assembly of rods (≈ 2000) modeled as super-helices [Bertails et al., 2006], and for which state-of-the-art nonsmooth solvers were shown to fail. We have also started to adapt our solver to the interaction handling between large nodal systems and rigid bodies, such as a full garment worn by a walking character.

3.2.1 Adding unilateral contact with Coulomb friction

As individual components, we consider discrete Lagrangian systems; that is, spatially discretized systems whose kinematics is fully described by the choice of a finite set of reduced (or generalized) coordinates, without the need for adding any constraint. This is typically the case of the discrete curvature-based models derived for inextensible rods in Chapter 2, but also of the nodal systems classically used for modelling extensible plates and shells in Computer Graphics [Baraff and Witkin, 1998; Grinspun et al., 2003].

Contact-free dynamics Let $\mathbf{q} \in \mathbb{R}^m$ be the m generalized coordinates of such a spatially discrete system and $\mathbf{v} = \frac{d\mathbf{q}}{dt}$ its generalized velocities. Equations of motion simply read as an ODE,

$$\mathcal{M}(t, \mathbf{q}) \frac{d\mathbf{v}}{dt} = \mathbf{f}(t, \mathbf{q}, \mathbf{v}), \quad (3.1)$$

where the inertia matrix $\mathcal{M}(t, \mathbf{q})$ is symmetric positive definite, and the generalized force \mathbf{f} collects all forces applying onto the system, including internal elastic forces, nonlinear inertial forces, and external forces such as gravity or air friction.

Now, when considering multiple such systems simultaneously by concatenating their respective contributions, the structure of Equation (3.1) is globally preserved. If we look more into details — this will be relevant when designing solving algorithms in Section 3.2.4 — we note that the inertia matrix \mathcal{M} naturally becomes sparse-block as several individual components are added. In the remainder of the text, we shall equally refer to Equation (3.1) for the unconstrained dynamics of an individual or multiple components, keeping in mind the sparse pattern of \mathcal{M} in the case of multiple components.

Local contact basis When coupling one or various components together with contact and friction (including contact with an external object, mutual contacts between components, and self-contact), we make several important simplifying assumptions. First, we only consider a finite set of contact points. Such a simplification is debatable, as it is well-known that some contacting configurations, for instance a plectoneme (see Figure 3.2(b)) may exhibit a continuous line of contact beyond a certain number of applied turns, a so-called *supercoiled* configuration [Neukirch, 2004]. However, concentrating contact forces into an adequate³ set of points often proves sufficient to avoid penetration and capture the Coulomb friction law, while considerably simplifying the numerical settings. We further assume that one contact always occurs between exactly two bodies, and that the surface of contacting objects is sufficiently smooth so that we can define a tangent surface \mathcal{S} and a contact normal \mathbf{e} , from which we build a local basis (see Figure 3.4). We can then express in this basis the relative velocity \mathbf{u} between the two bodies, and the force \mathbf{r} at the contact point applying from one body to the other.




Figure 3.4 – Local contact basis, with normal and tangent subspaces.

Non-penetration conditions The first law that we wish to enforce is a *strict non-penetration* constraint between the two objects. Let $h(\mathbf{q})$ be the gap function between the two bodies, which corresponds to a signed distance function. We assume impacts to be purely inelastic (that is, kinetic energy instantaneously dissipates when the two objects collide), and thus enforce a vanishing post-impact normal relative velocity $u_N(t^+)$. When the two objects are in contact at time t (i.e., such that $h(\mathbf{q}) = 0$), some velocity-level condition, named *Signorini* condition, can be formulated to express the non-penetration constraint at time $t > 0$,

$$0 \leq u_N \perp r_N \leq 0, \quad (3.2)$$

where the $x \perp y$ notation means that the two scalar variables x and y should be orthogonal, i.e., $xy = 0$ (more generally, in \mathbb{R}^d , $\mathbf{x} \perp \mathbf{y}$ means $\mathbf{x}^\top \mathbf{y} = 0$).

The *complementarity* condition (3.2) means that either the contact ceases ($u_N > 0$) and necessarily the normal contact force vanishes ($r_N = 0$), or the contact persists ($u_N = 0$) and necessarily the normal contact force, if not vanishing, becomes (or remains) active ($r_N > 0$) to prevent penetration.

³The set of contact points is determined on the fly using a collision detection process.

Dry (solid) friction law The dry friction law, often attributed to Coulomb [1781] (but which a number of other authors have contributed to, such as de Vinci, Amontons, and Euler), results from the macroscopic observation that the normal component r_N and the tangential component \mathbf{r}_T of a contact force \mathbf{r} between two solids satisfy the following law,

$$\begin{cases} \|\mathbf{r}_T\| \leq \mu_S r_N & \text{if } \mathbf{u}_T = \mathbf{0} \quad (\text{static regime}) \\ \|\mathbf{r}_T\| = \mu_D r_N & \text{if } \mathbf{u}_T \neq \mathbf{0} \quad (\text{dynamic regime}), \end{cases} \quad (3.3)$$

where μ_S and μ_D are the static and dynamic friction coefficients, respectively. Although it has been observed that μ_S is generally slightly higher than μ_D , we shall assume in the following, for the sake of simplicity, that they are equal and correspond to the same parameter μ . The value of μ shall depend on the characteristics of the contacting surfaces (from $\mu^i = 0$ for perfect contact without friction to $\mu^i = \mu_{\max}$ for rough surfaces).

One may observe that the Coulomb law is a *threshold* model, since the tangential contact force should reach a certain magnitude before the objects are able to slide on top of each other. Moreover, this magnitude directly depends on the normal contact force, meaning that a heavy object requires some higher pulling force to be dragged on a surface compared to a light object. This *normal-dependent threshold*, which is key to realism⁴, is also responsible for the fact that Coulomb's law does not fit in with the so-called *associated rules* or *generalized standard materials* [De Saxcé and Feng, 1998], which are easier to manipulate formally and to handle numerically. As a practical consequence, it will not be possible to view the Coulomb friction constraint as the optimality condition of a convex minimization problem (see 3.2.3). A full discussion about the inclusion of Coulomb's law in the family of implicit standard materials (non-associated flow rules) is provided in [Daviet, 2016].

The full Signorini-Coulomb disjunctive law Combining non-penetration constraint (3.2) together with Coulomb friction (3.3) is commonly referred to as the *Signorini-Coulomb* law. It can be written in a compact and disjunctive manner, as

- Taking-off case: $\mathbf{r} = \mathbf{0}$ and $u_N \geq 0$
- Sticking case: $\mathbf{r} \in K_\mu$ and $\mathbf{u} = \mathbf{0}$
- Sliding case: $\mathbf{r} \in \partial K_\mu$, $u_N = 0$ and $\exists \alpha \in \mathbb{R}_+^*$, $\mathbf{u}_T = -\alpha \mathbf{r}_T$,

where K_μ is the second-order cone (**SOC**), defined in dimension d as

$$K_\mu := \left\{ \mathbf{x} = [x_N, \mathbf{x}_T] \in \mathbb{R} \times \mathbb{R}^{d-1}, \mu x_N \geq \|\mathbf{x}_T\| \right\},$$

⁴Indeed, imagine a stack of fiber layers, or a sand pile under gravitational field. If the friction threshold is constant, sliding will occur at any arbitrary depth of the pile, leading to an unrealistic avalanching behavior. In contrast, if the threshold is correctly set proportional to the normal force, the top surface will be, as expected, more prone to sliding avalanches compared to deeper layers.

with x_N the (scalar) normal component of $\mathbf{x} \in \mathbb{R}^d$, and \mathbf{x}_T its tangential component.

The Signorini-Coulomb law, illustrated in Figure 3.5, can be interpreted as follows: if the first case, the normal relative velocity is nonnegative and the contact breaks (taking-off case). Then, in the case when the two bodies are in resting contact, either they are sticking to each other and the contact force can lie anywhere in \mathcal{K}_μ , or they are sliding with respect to each other, then the contact force must belong to the boundary of \mathcal{K}_μ and the tangent force must be collinear to the relative velocity with the opposite direction (according to the “maximum dissipation principle” [Moreau, 1988]).




Figure 3.5 – Illustration in 3D of the three cases of the Signorini-Coulomb law.

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 Section 3.2.3, we give equivalent formulations of Signorini-Coulomb’s law which prove to be more tractable numerically in the general case.

In the remainder of this document we shall denote by \mathcal{C}_μ the set of vectors (\mathbf{r}, \mathbf{u}) in $\mathbb{R}^d \times \mathbb{R}^d$ satisfying Signorini-Coulomb’s law.

3.2.2 The Discrete Frictional Contact Problem (DFCP)

We now assume that our system (3.1) is subject to n frictional contacts, supposed to be purely inelastic. To simplify notation, the n relative velocities \mathbf{u}^i and contact forces \mathbf{r}^i are collected into \mathbf{u} and \mathbf{r} , respectively. We define $\mathcal{H}(\mathbf{q}) := \frac{\partial \mathbf{u}}{\partial \mathbf{v}}(\mathbf{q})$ the deformation gradient matrix⁵ of dimension (nd, m) relating the relative spatial velocities at contact points \mathbf{u} to the generalized velocities \mathbf{v} . Let \mathbf{u}_f be the value of \mathbf{u} when $\mathbf{v} = \mathbf{0}$, which can be non-zero in the case of forced motion⁶. The nonsmooth

⁵For a super-model, the deformation gradient matrix reads $\mathcal{H}(\mathbf{q}, s) = \frac{\partial \mathbf{r}}{\partial \mathbf{q}}(s)$ and can be evaluated efficiently at any contact point $s = s_c$, either formally (super-helix) or using adaptive power series summation (super-clothoid). For a nodal system, the gradient deformation matrix simply consists of identity blocks.

⁶Unlike nodal systems where $\mathbf{u}_f = 0$, for a super-model this term may be nonzero when the clamped end is moving; in that case an extra velocity $\dot{\mathbf{r}}^*(s)$, independent of $\mathbf{v} = \dot{\mathbf{q}}$, shows up, and we have for each contact block $\mathbf{u}_f = \dot{\mathbf{r}}^*(s_c)$, where s_c is the location of contact.

system including contact forces \mathbf{r} can be written as

$$\begin{cases} \mathcal{M}(\mathbf{q}) \frac{d\mathbf{v}}{dt} = \mathbf{f}(t, \mathbf{q}, \mathbf{v}) + \mathcal{H}(\mathbf{q})^\top \mathbf{r} \\ \mathbf{u} = \mathcal{H}(\mathbf{q})\mathbf{v} + \mathbf{u}_f(t, \mathbf{q}) \\ \forall i = 1 \dots n, (\mathbf{r}^i, \mathbf{u}^i) \in \mathcal{C}_{\mu_i}. \end{cases} \quad (3.4)$$

The notation $\frac{d\mathbf{v}}{dt}$ is actually misused as the inclusion of contacts implies that we might observe jumps in the generalized velocities \mathbf{v} . To circumvent this difficulty, we follow Moreau's time-stepping scheme and integrate the equations over an arbitrary timestep δt [Moreau, 1988; Acary and Brogliato, 2008]. We thus get a formulation over mechanical impulses instead of forces, and a velocity jump instead of an acceleration.

Let us set $\mathbf{r} := \int_{\delta t} \mathbf{r} dt$, and let \mathbf{v} and \mathbf{u} be the discrete approximations of the generalized velocity and of the relative velocity, respectively, at the end of the timestep. The constraint-free dynamics (3.1) is discretized using a θ -scheme and can be written at a given timestep as $\mathbf{M}\mathbf{v} = \mathbf{f}$. We further assume that $\mathcal{H}(\mathbf{q})$ and $\mathbf{u}_f(t, \mathbf{q})$ are almost constant during the timestep, and approximate them with \mathbf{H} and \mathbf{u}_f , respectively. The discretized velocity/impulse formulation of (3.4), with unknowns $(\mathbf{v}, \mathbf{u}, \mathbf{r})$, then reads

$$\begin{cases} \mathbf{M}\mathbf{v} - \mathbf{f} = \mathbf{H}^\top \mathbf{r} \\ \mathbf{u} = \mathbf{H}\mathbf{v} + \mathbf{u}_f \\ \forall i = 1 \dots n, (\mathbf{r}^i, \mathbf{u}^i) \in \mathcal{C}_{\mu_i}. \end{cases} \quad (3.5)$$

Problem (3.5), referred to as the *one-step problem* in [Cadoux, 2009], is the core problem we wish to solve in this chapter. In the remainder of this document, we shall call it *Discrete Frictional Contact Problem (DFCP)*. Note that we may eliminate \mathbf{v} in (3.5) by introducing the *Delassus* operator $\mathbf{W} := \mathbf{H}\mathbf{M}^{-1}\mathbf{H}^\top$ and by setting $\mathbf{b} := \mathbf{u}_f + \mathbf{H}\mathbf{M}^{-1}\mathbf{f}$, and finally obtain the following reduced **DFCP**,

$$\begin{cases} \mathbf{W}\mathbf{r} + \mathbf{b} = \mathbf{u} \\ \forall i = 1 \dots n, (\mathbf{r}^i, \mathbf{u}^i) \in \mathcal{C}_{\mu_i}, \end{cases} \quad (3.6)$$

which may again be compacted as a force (or impulse) based problem,

$$\forall i = 1 \dots n, (\mathbf{r}^i, (\mathbf{W}\mathbf{r})^i + \mathbf{b}^i) \in \mathcal{C}_{\mu_i}. \quad (3.7)$$

There is a number of difficulties when attempting to solve the **DFCP**:

- First, the inclusion of $(\mathbf{r}^i, \mathbf{u}^i)$ in the set \mathcal{C}_{μ_i} makes the problem *nonconvex* and *nonsmooth*. Standard convex optimization algorithms, which have well-established convergence guarantees, can thus not be applied to our problem, and specific algorithms need to be developed.

- Then, the **DFCP** may not always possess a solution⁷. One typical example is the (discrete) Painlevé paradox⁸, described for instance in [Baraff, 1991; Cadoux, 2009]. Yet, in our case where we wish to simulate objects that are not overly constrained (avoiding, e.g., the scenario of a hand strongly compressing a hair wisp), our problem is likely to possess a solution most of the time [Cadoux, 2009]. When designing a splitting solving method however, the non-existence issue may show up again for each subproblem. Thus, *robust* schemes, able to find an approximate solution in the absence of a theoretical one, have to be designed.
- Finally, our scenarios may rapidly involve thousands or even tens thousands of contacts, making the problem very *large* in size. Solving algorithms should thus be able to *scale up* well with the size of the problem, that is, maintain *good convergence* properties at *acceptable* timings. Typically, in our scenarios where the time step might be as small as 1 ms, we aim at solving the **DFCP** in a few seconds up to a few minutes, in order to keep a reasonable (i.e., about one hour-long) frame rate.

To face up all these challenges at the same time, we have been opting for a *practical* study, where we have been first testing a wide range of nonsmooth solvers, and then designing new ones targeted at the scenario to be simulated. The result of our study is not a mathematical analysis of a given algorithm with some proof of convergence⁹. Instead, for each scenario of interest (hair, cloth), we have been characterizing the main features of the resulting **DFCP** and we have built a *dedicated* method or heuristics, leading to some fast and robust solver.

3.2.3 Multiple equivalent formulations of Signorini-Coulomb

We have first been extensively playing with the many different equivalent formulations of the Signorini-Coulomb law, each one of them leading to various solving strategies. An extensive study of these different strategies is proposed in Gilles Daviet’s PhD thesis [Daviet, 2016]. Here we only mention the formulations and methods we have been frequently making use in our different applications, or comparing against.

⁷Non-existence typically occurs when $\text{rank } \mathbf{H} < nd$, in which case \mathbf{W} is not necessarily strongly positive-definite [Cadoux, 2009]. Moreover, if the **DFCP** admits a solution, it might not be unique. In this chapter, where we wish to solve the *direct* problem (i.e., to retrieve generalized positions and velocities from the dynamic equations, a set of parameters, and initial conditions), we will not care too much about non-uniqueness of the solution. However, this will become an actual concern when attempting to solve the *inverse* problem and to identify parameters, as discussed in Chapter 4.

⁸Note that the Painlevé paradox, as well as more general cases of nonexistence of solution, do not only occur in the discrete settings, but also in the time-continuous problem of Lagrangian systems subject to frictional contact. This was our very objet of study in [Blumentals et al., 2016b].

⁹In the last decades there has been many attempts to solve the **DFCP**, but the literature still remains very sparse when dealing with convergence guarantees [Acary and Brogliato, 2008]

Linearized friction cone One popular approach, which is actually not equivalent to the Signorini-Coulomb law but stands for a *linear approximation* of it, consists in approaching the second-order cone \mathcal{K}_μ with a convex polyhedron (a faceted cone). This approach, initially proposed by Klarbring [1987], has been widely used in Mechanics, Robotics and Computer Graphics [Stewart and Trinkle, 1996; Miller and Christensen, 2003; Erleben, 2007; Kaufman et al., 2008; Otaduy et al., 2009], as it advantageously yields a Linear Complementarity Problem¹⁰ (**LCP**), for which many solvers have been proposed in the literature. While this approximate formulation captures both dynamic and static friction regimes properly, it introduces some anisotropy in the sliding direction (see Figure 3.6) unless a high number of facets is used [Acary and Brogliato, 2008, Sec. 13.3.7]. A trade-off thus has to be made between accuracy and computational cost. In our applications we have not found this approach to be more efficient compared to an exact treatment of the friction cone, even with a low number of facets, and thus have rapidly discarded it. All the formulations presented in the sequel consider the *exact* Coulomb cone \mathcal{K}_μ .




Figure 3.6 – Exact *vs.* approximate model for Coulomb friction. (a) Periodic orbit of the free end of a rod resting on top of a rotating sphere. Frictional contact is simulated using our eight different solvers. Solvers that model exact Coulomb friction all reach the same orbit, no matter the choice of the error function. Other models (dashed lines) reach completely different ones. (b) Visual comparison of a full hair simulation with exact (left) and linearized (right) Coulomb friction, for equal computation time. With exact Coulomb friction, spontaneous hair clumping emerges in a more visible way.

Disjunctive formulation We have already mentioned that solving the disjunctive Signorini-Coulomb formulation by enumerating each case is intractable for an arbitrary number of contacts n . However, when $n = 1$ (and we shall see the interest of considering the one-contact problem when dealing with splitting algorithms), then an exact algorithm can be designed [Bonnefon and Daviet, 2011]. Actually, the only difficult case is the sliding case, which requires computing the roots of a quartic polynomial. This analytic one-contact solver has the advantage of providing a certificate of (non)-existence of solutions, and when a solution exists, to give its value down to the machine precision.

¹⁰A Linear Complementarity Problem (**LCP**) reads $0 \leq \mathbf{x} \perp \mathbf{y} \geq 0$ with $\mathbf{y} = \mathbf{Ax} + \mathbf{b}$.

Complementarity formulation The Signorini-Coulomb law implies that \mathbf{r}^i lies inside the cone \mathcal{K}_{μ_i} and \mathbf{u}^i in the half-space $\mathbb{R}_+ \times \mathbb{R}^{d-1}$. Actually, this can be expressed as a cone complementarity problem by mapping this half-space to the dual cone of \mathcal{K}_{μ_i} , which is $\mathcal{K}_{\perp}^{\mu_i}$ (see Figure 3.7, (a) and (b)). Such a change of variables has been described by De Saxcé and Feng [1998],

$$\tilde{\mathbf{u}}^i := \mathbf{u}^i + \mu^i \|\mathbf{u}_T^i\| \mathbf{e}, \quad (3.8)$$

and leads to the following expression of the Coulomb law,

$$(\mathbf{r}^i, \mathbf{u}^i) \in \mathcal{C}_{\mu_i} \iff \mathcal{K}_{\perp}^{\mu_i} \ni \tilde{\mathbf{u}}^i \perp \mathbf{r}^i \in \mathcal{K}_{\mu_i}. \quad (3.9)$$

Note that for frictionless contacts ($\mu^i = 0$), the complementarity problem (3.9) simply boils down to an **LCP** on the normal parts of \mathbf{u}^i and \mathbf{r}^i .

The interpretation of the Signorini-Coulomb law as a complementarity condition on dual cones (3.9) is at the origin of a myriad of other equivalent formulations, each one of them giving rise to different solving algorithms. In the following, we shall focus on two new formulations inspired by (3.9): the so-called De Saxcé functional formulation, and Cadoux's fixed-point algorithm.




Figure 3.7 – Changes of variables applied onto (a) the relative velocity \mathbf{u} and the contact force \mathbf{r} , in order to exhibit (b) the complementarity of $\tilde{\mathbf{u}}$ and \mathbf{r} proposed by De Saxcé and Feng [1998].

Functional formulation It is often practical to express the Signorini-Coulomb law as a root-finding problem, i.e., in the form

$$(\mathbf{r}^i, \mathbf{u}^i) \in \mathcal{C}_{\mu_i} \iff f(\mathbf{r}^i, \mathbf{u}^i) = 0, \quad (3.10)$$

where f is a nonsmooth merit function from $\mathbb{R}^d \times \mathbb{R}^d$ to \mathbb{R}^d .

One classical example of such a nonsmooth function was notably provided by Alart and Curnier [1991], and defined as

$$f_{AC} : \mathbb{R}^d \times \mathbb{R}^d \longrightarrow \mathbb{R}^d$$

$$(\mathbf{r}^i, \mathbf{u}^i) \longmapsto \begin{pmatrix} \Pi_{\mathbb{R}_+} (r_N^i - \xi_N u_N^i) - r_N^i \\ \Pi_{\mathcal{B}^{d-1}(\mu r_N)} (\mathbf{r}_T^i - \xi_T \mathbf{u}_T^i) - \mathbf{r}_T^i \end{pmatrix},$$

where $\Pi_{\mathcal{C}}$ is the orthogonal projection operator on the convex space \mathcal{C} , ξ_T and ξ_N are positive real numbers, and $\mathcal{B}^{d-1}(a) \subset \mathbb{R}^{d-1}$ is the ball of radius $a \geq 0$ centered at the origin.

As suggested above, another well-known function satisfying (3.10) may be derived based upon De Saxcé and Feng’s change of variable (3.8), by using the projection theorem (see, e.g., [Hiriart-Urruty and Lemaréchal, 2001, Proposition A.5.3.3]),

$$\begin{aligned} f_{\text{DS}} : \mathbb{R}^d \times \mathbb{R}^d &\longrightarrow \mathbb{R}^d \\ (\mathbf{r}^i, \mathbf{u}^i) &\longmapsto \Pi_{\mathcal{K}_{\mu_i}}(\mathbf{r}^i - \xi \tilde{\mathbf{u}}^i) - \mathbf{r}^i, \end{aligned}$$

where ξ is a positive real number.

Functional formulations may be solved for by using a *generalized Newton* method, in which the usual gradient is replaced with a generalized gradient at non-differentiable points [Qi and Sun, 1993].

In our work [Daviet et al., 2011], summarized in Section 3.2.4, we have been deriving yet another merit function, using the equivalence between a Second-Order Cone Complementarity Problem (**SOCCP**) and a modified Fischer-Burmeister¹¹ functional formulation [Fukushima et al., 2002]. Due to better smoothness properties compared to the two other formulations above, this modified Fischer-Burmeister formulation proves in practice easier to solve using a generalized Newton solver, compared to the two formulations above.

However, in all cases, when attempting to solve our large **DFCP** using any of the functional method in a global way (that is, solving all the contacts simultaneously), our solver always failed to scale up to more than one hundred contacting fibers. More precisely we have found on our range of problems that the convergence success of the generalized Newton method was directly related to the conditioning number $\nu = \frac{nd}{m}$ [Bertails-Descoubes et al., 2011]. For $\nu < 1$ (loosely connected fibers), the method would converge properly. However, for $\nu \geq 1$ (tight fiber packing, i.e., our typical scenario of interest), convergence would be very slow and eventually the method would fail to converge. For this reason we have thus changed strategies and instead used a splitting method, where contacts are solved one by one in an iterative fashion (see 3.2.4).

Still, one advantage of functional formulations is that they give a natural stopping criterion for *any* iterative method attempting to solve the reduced **DFCP** (3.7). Indeed, the norm $\frac{1}{2} \|\Phi(\mathbf{r})\|^2$ where $\Phi(\mathbf{r})$ is the \mathbb{R}^{dn} -value function such that $\Phi^i(\mathbf{r}) = f(\mathbf{r}^i, (\mathbf{W}\mathbf{r})^i + \mathbf{b}^i)$, gives the global error of the iterative method. In our splitting algorithm, we have evaluated the error provided by the global Alart-Curnier function for stopping iterations of our algorithm (see Section 3.2.4).

Fixed-point sequence of convex optimization problems Another class of methods attempts to formulate the **DFCP** as a sequence of *convex* optimization

¹¹For complementarity problems in \mathbb{R} , a well-known merit function is the so-called *Fischer-Burmeister* function defined as $f_{\text{FB}}(x, y) = \sqrt{x^2 + y^2} - x - y$. One may easily check that $0 \leq x \perp y \geq 0 \iff f_{\text{FB}}(x, y) = 0$. Fukushima [2002] has extended this formulation to second-order conical complementarity constraints.

problems [Haslinger, 1983; Cadoux, 2009]. In particular, Florent Cadoux has shown in his PhD thesis [Cadoux, 2009; Acary et al., 2011] that solving the **DFCP** is equivalent to performing a fixed point loop on either a *primal* (i.e., velocity-based) or *dual* (i.e., impulse-based) convex optimization problem.

Let us give some insights to this interesting method. Starting from De Saxé and Feng’s change of variable (3.8) and concatenating it for every contact, one gets

$$\tilde{\mathbf{u}} := \mathbf{u} + \mathbf{E} \mathbf{s} \in L^*,$$

where $\mathbf{s} = [\|\mathbf{u}_T\|_1, \dots, \|\mathbf{u}_T\|_n]^\top = [\|\tilde{\mathbf{u}}_T\|_1, \dots, \|\tilde{\mathbf{u}}_T\|_n]^\top$, $\mathbf{E} = \text{blockdiag}(\mu_i \mathbf{e}_i)$, and $L^* = \prod_i \mathcal{K}_{\mu_i}^*$. The **DFCP** (3.5) then reads

$$\begin{cases} \mathbf{M} \mathbf{v} - \mathbf{f} = \mathbf{H}^\top \mathbf{r} & (a) \\ \tilde{\mathbf{u}} = \mathbf{H} \mathbf{v} + \mathbf{u}_f + \mathbf{E} \mathbf{s} & (b) \\ L^* \ni \tilde{\mathbf{u}} \perp \mathbf{r} \in L & (c) \\ \mathbf{s} = [\|\tilde{\mathbf{u}}_T\|_1, \dots, \|\tilde{\mathbf{u}}_T\|_n]^\top & (d) \end{cases}$$

The key of the Cadoux approach is to note that if s is *fixed*, then (a), (b), and (c) are exactly the *optimality conditions* of a *convex* optimization problem subject to conical constraints, which can be equivalently formulated in a *primal* form,

$$\begin{cases} \min \frac{1}{2} \mathbf{v}^\top \mathbf{M} \mathbf{v} - \mathbf{f}^\top \mathbf{v} & \text{(quadratic, strict. convex)} \\ \mathbf{H} \mathbf{v} + \mathbf{u}_f + \mathbf{E} \mathbf{s} \in L^* & \text{(conical constraints)} \end{cases} \quad (3.11)$$

or in a *dual* form,

$$\begin{cases} \min \frac{1}{2} \mathbf{r}^\top \mathbf{W} \mathbf{r} + \mathbf{b}^\top \mathbf{r} & \text{(quadratic, convex)} \\ \mathbf{r} \in L & \text{(conical constraints)} \\ \mathbf{W} = \mathbf{H} \mathbf{M}^{-1} \mathbf{H}^\top \\ \mathbf{b} = \mathbf{H} \mathbf{M}^{-1} \mathbf{f} + \mathbf{u}_f + \mathbf{E} \mathbf{s} \end{cases} \quad (3.12)$$

Finally, the full **DFCP** can be solved thanks to a global fixed-point loop,

$$F(\mathbf{s}) = \mathbf{s} \quad \text{with} \quad F^i(\mathbf{s}) := \|\mathbf{u}_T^i(\mathbf{s})\|, \quad (3.13)$$

where at each iteration, $\mathbf{u}^i(\mathbf{s})$ is computed by solving either the primal problem (3.11), or the dual problem (3.12).

From this algorithm, it becomes now clear that the **DFCP** cannot be simply cast into a convex optimization problem. Interestingly, previous approaches which have approximated the **DFCP** with a convex optimization problem [Anitescu, 2005; Mazhar et al., 2015] can be interpreted as a single fixed-point iteration of the procedure above. One shortcoming inherent to this approximation is that the relative velocity \mathbf{u} , which should be tangent to the contacting surface when contact is active, is biased by the adding of a non-zero normal component, which in practice may cause artificial jumps of sliding objects during simulation.

We have been testing the two variants of Cadoux’s method on our large fiber problems, using different solvers. For instance, we have experimented solving the

dual problem with either a projected gradient method, or interior points. Solving the primal problem with interior points actually gave better results of convergence, due to the strongly convex objective in this case. However, although the latter variant turned out to perform better compared to global functional formulations, it did not allow us to handle more than a few hundreds contacting rods: beyond, the solver would fail to converge in a reasonable amount of time.

Splitting (Gauss-Seidel) formulation Instead of solving the **DFCP** globally, one may instead solve for each contact problem sequentially, and then iterate on the full set of contacts until convergence is reached. Such a sequential splitting approach, often called “Gauss-Seidel” due to its resemblance with the Gauss-Seidel algorithm to solve linear systems, is actually classically used for simulating large granular assemblies [Dubois and Jean, 2006], due to its scalability properties. Convergence to the solution at a high precision is however known to be slow in practice [Acary and Brogliato, 2008]. This was however not a problem in our case, where we only request to solver to converge to an acceptable precision, without the need to get beyond visual perception.

When adopting the Gauss-Seidel strategy, one is then left with the choice for the *local* solver. The key insight we have learnt from our multiple tests is that failing to solve one single contact properly may often cause the global loop to fail and introduce instabilities in the fiber dynamics, ultimately leading to a simulation crash. We have thus strived to build a local solver *as robust as possible*, by devising a hybrid strategy.

3.2.4 A local hybrid solving strategy

Large fiber assemblies Facing robustness and scalability issues when using state-of-the-arts approaches, we have designed a new frictional contact solver by coupling an iterative Gauss-Seidel strategy together with an extremely robust one-contact solver [Daviet et al., 2011]. The local solver combines two different methods for solving the 1-contact problem in contact force variables: on the one hand, it primarily relies upon the root finding of the modified Fischer-Burmeister function [Fukushima et al., 2002], thanks to a nonsmooth Newton method; on the other hand, in the rare case when the previous method fails to converge, the solver automatically switches to the more expensive but analytical solver which is guaranteed to find the solution to the 1-contact problem when it exists [Bonneton and Daviet, 2011]. The latter solver, called in 1% of the cases, thus acts as a *failsafe* of the method.

Our global solver proved to converge well in scenarios involving thousands fibers subject to tens thousands frictional contact points, and thus allowed us to enhance considerably the realism of hair simulations. Since then, our method has been leveraged by academia and the special effects industry for simulating hair and fur realistically [Kaufman et al., 2014]. In particular, our method has been used in production by Weta Digital, the leading company in special effects, and our accompanying software has been licensed to L’Oréal Research and AGT Digital. We have

also recently adapted our solver for coping with frictional contact in cloth modeled as a nodal system [Daviet et al., 2015], as explained below.



Figure 3.8 – Comparison of the hair collective behavior between (top) real hair motion sequences and (bottom) our corresponding simulations, based on large assemblies of (up to 2,000) individual fibers with massive self-contacts and Coulomb friction. Our model retains typical emerging effects such as transient coherent motions or stick-slip instabilities, at competitive timings (typically, a few seconds per time step, with $dt = 1\text{ms}$).

Cloth: towards a dedicated nodal solver In the fiber assembly case, the matrix \mathbf{M} is block-diagonal, so that the Delassus operator can be computed in an efficient way by leveraging sparse-block computations [Daviet et al., 2011]. This justifies solving the reduced DFCEP (3.6), where primary unknowns are forces \mathbf{r} . For cloth however, where primal variables (nodal velocities of the cloth mesh) are all interconnected via elasticity through implicit forces, the method developed above is computationally inefficient. Indeed, the matrix \mathbf{M} (only block-sparse, but not block-diagonal) is costly to invert for large systems and its inverse is dense. However, we may leverage the fact that generalized velocities of the system are 3D velocities, and slightly adapt the previous solver so as to recover efficiency.

We assume body-cloth contacts occur at cloth vertices only, i.e., each contact involves only one vertex. A key observation is that each nonzero block of \mathbf{J} related to a vertex i is simply a rotation matrix E_i , corresponding to the local contact basis for the vertex i . Our idea is then to build the square block-diagonal matrix \mathbf{G} with $\mathbf{G}_{i,i}$ equal to E_i if i is in contact, and to the 3×3 identity matrix otherwise. \mathbf{G} can thus be inverted trivially as $\mathbf{G}^{-1} = \mathbf{G}^T$. Augmenting \mathbf{u} and \mathbf{r} so that their size matches that of \mathbf{v} , we can write (3.5) as

$$\begin{cases} \mathbf{GMG}^T\mathbf{u} - \mathbf{Gf} = \mathbf{r} \\ \mathbf{v} = \mathbf{G}^T\mathbf{u} \\ (\mathbf{u}^i, \mathbf{r}^i) \in \mathcal{C}_{\mu_i} \quad \text{if } i \text{ in contact} \quad \text{and } \mathbf{r}^i = 0 \text{ otherwise.} \end{cases} \quad (3.14)$$

Denoting $\widehat{\mathbf{W}} = \mathbf{GMG}^T$ and $\widehat{\mathbf{b}} = \mathbf{Gf}$, we obtain a system that is very close to the reduced DFCEP (3.6), except that \mathbf{u} and \mathbf{r} have reversed roles. To retrieve symmetry, we apply De Saxcé's change of variable (3.8) $\tilde{\mathbf{u}}^i := \mathbf{u}^i + s_i(\mathbf{u})$ so that $\tilde{\mathbf{u}}^i$ is orthogonal to \mathbf{r}^i . For a fixed value of s , we identify (3.14) as the KKT conditions of the convex

quadratic optimization problem [Cadoux, 2009; Acary et al., 2011]

$$\min_{\tilde{\mathbf{u}}} \frac{1}{2}(\tilde{\mathbf{u}})^{\mathbf{T}}\widehat{\mathbf{W}}\tilde{\mathbf{u}} + (\widehat{\mathbf{b}} - \widehat{\mathbf{W}}\mathbf{s})^{\mathbf{T}}\tilde{\mathbf{u}}$$

$$\tilde{\mathbf{u}}^i \in K_{\frac{\perp}{\mu_i}} \quad \text{if } i \text{ in contact,}$$

which can be solved using our hybrid solver, since this time $\tilde{\mathbf{u}}$ plays the exact same role as \mathbf{r} (replacing \mathcal{K}_{μ_i} with K_{\perp}). Following [Cadoux, 2009], we compute the solution of the full problem (3.5) by iteratively updating s using the fixed point algorithm (3.13). Unlike the Delassus operator \mathbf{W} , our new operator $\widehat{\mathbf{W}}$ is easy to assemble and sparse. Solving our primal problem (3.14) thus turns out to be orders of magnitude faster compared to the reduced **DFCP**. We are currently trying to extend this simple method for dealing with cloth self-contact.



Figure 3.9 – Our body-cloth contact solver captures exact Coulomb friction both efficiently and robustly. On this dress example featuring 6000 cloth vertices and 1000 contact points on average, our solver converges at each time step ($dt = 1\text{ms}$) in a few hundred milliseconds..

3.3 Continuum modeling of granular materials

In this section we drastically change models for representing a collection of interacting objects. Instead of modeling individual objects explicitly as in Section 3.2, we adopt a *continuum* viewpoint where a fragment of this continuum represents an homogenized sample of individual elements. Such an approach is desirable when aiming at modeling hundreds thousands or even billions interacting elements, which would be intractable using a discrete element method as before¹². However, to the best of our knowledge, continuum viewpoints are not well-developed enough for taking into account individual elements with complex geometries — such as long and thin deformable fibers. To start with a continuum viewpoint, we have thus restricted ourselves to a simpler, yet widely studied material: *granular matter*.

3.3.1 Discrete vs. continuum models for granular matter

Granular materials (see, e.g., [Andreotti et al., 2011] for a comprehensive description) commonly refer to a large collection of small solid grains larger than $100\ \mu\text{m}$

¹²Note that another strategy to scale things up would be to take a look at high performance computation. However the solver we have designed in Section 3.2 is inherently sequential, and a completely new parallel algorithm would then have to be designed from scratch, requiring some specific expertise in parallel design. Motivated by physical insights instead, we have rather chosen, at least for now, to investigate the building of new, macroscopic physics-based models.

in size — which typically distinguishes them from powders, made of much smaller grains. Considering this limit size, grain-grain interactions in granulars are mainly dictated by *contact* and *dry friction*, while air-grain interactions can be neglected. *Cohesion* between grains may furthermore be considered, typically in the case of wet materials.

Due to their discrete nature at macroscopic scale, granular materials are highly dissipative systems which exhibit various physical states, ranging from “liquid” (when the material continuously flows) to “solid” (when the material rigidifies due to internal friction and undergoes plastic deformations), or even “gaseous” (when grains are dispersed in air and interact mainly through impacts). Such a rich visual behavior, very distinctive from that of Newtonian fluids (e.g., water), explains why simulating granular matter has, beyond classical applications in Physics and Mechanical Engineering, gained increased attention in Computer Graphics for a few decades [Zhu and Bridson, 2005; Alduán et al., 2009; Narain et al., 2010].

However, unlike Newtonian fluids or elastic materials for which robust physical models have been developed in the last centuries, granular matter still resists some comprehensive physical understanding: no universal equation exists yet for describing the collective behavior of such a peculiar material.

As a result, various methods have been developed for studying and simulating granular materials, following two main strategies. On the one hand, many approaches naturally rely upon a *discrete* model, where grains are represented individually as rigid bodies and grain-grain interactions are accounted for using a frictional contact model [Moreau, 1994]. Because they explicitly model the material at the grain scale, those methods are able to capture the different states of the granular matter and their transitions quite accurately. However, considering that a simple teaspoon of sugar already consists of around 10^5 grains, one may easily imagine that such approaches rapidly suffer from scalability issues.

On the other hand, *continuum*-based methods, where granular matter is modeled as a viscoplastic fluid, have started to be investigated for a few decades. Though less accurate than discrete models, such methods have nevertheless demonstrated some good qualitative agreement with real scenarios in plastic and dense flowing regimes. In particular, a popular constitutive law is the $\mu(I)$ rheology [Jop et al., 2006], relying upon a pressure-dependent yielding model — referred to as the *Drucker-Prager* yield criterion (detailed below) — with a non-constant friction coefficient. For the sake of tractability, existing numerical models however make some important simplifications, for instance assuming an incompressible flow, a uniform density, or a regularized or linearized Drucker-Prager yield criterion. Such approximations often result in a considerable loss of realism, for instance preventing a heap of sand from stabilizing after flowing, or introducing some artificial anisotropy in the flow.

Our main objective towards this work was to treat the Drucker-Prager rheology numerically with as little as possible downgrade. Typically, we addressed the non-smooth rheology directly with no regularization and no linearization, and relaxed the incompressibility condition to incorporate a more realistic, unilateral compressibility condition. Altogether, these new settings allowed us to improve substantially

the quality of simulations of large granular flows, while keeping reasonable computational timings.

Caveat Our work on the continuum modeling of granulars with a nonsmooth rheology is recent, and was the core contribution of Gilles Daviet’s PhD thesis [Daviet, 2016]. Here we only sketch the main ideas of our method, without going deeply in technical details. The interested reader may find all technical issues in our publications [Daviet and Bertails-Descoubes, 2016b,a].

Notation We consider an Eulerian fluid representing a granular medium. We use standard tensor notation for manipulating scalar and vector field associated to this continuum. Let S_d be the space of symmetric $d \times d$ tensors (represented as matrices of dimension d), where d is the space dimension ($d = 2$ or $d = 3$). For $\boldsymbol{\sigma} \in S_d$, we denote by $\text{Tr } \boldsymbol{\sigma}$ its trace (normal part) and by $\text{Dev } \boldsymbol{\sigma}$ its deviatoric (tangential) part, $\text{Dev } \boldsymbol{\sigma} = \boldsymbol{\sigma} - \frac{1}{3} \text{Tr } \boldsymbol{\sigma} \mathbb{I}$, where \mathbb{I} is the identity tensor. We use the norm $|\cdot|$ associated to the scalar product $\langle \boldsymbol{\sigma}, \boldsymbol{\tau} \rangle = \frac{\boldsymbol{\sigma} : \boldsymbol{\tau}}{2} = \frac{1}{2} \sum \tau_{ij} \sigma_{ij}$. Note that $|\cdot|$ amounts to the Frobenius norm scaled by $\frac{1}{\sqrt{2}}$. The gradient of a scalar field ϕ and of a vector field \mathbf{v} are written as $\nabla \phi$ (vector) and $\nabla \mathbf{v}$ (matrix), respectively, and the divergence of a vector \mathbf{v} field as $\nabla \cdot \mathbf{v}$ (scalar). Given a vector field \mathbf{v} , let $\mathbf{D}(\mathbf{v}) := \frac{1}{2}(\nabla \mathbf{v} + (\nabla \mathbf{v})^T)$ be the symmetric part of its gradient and $\mathbf{W}(\mathbf{v}) := \frac{1}{2}(\nabla \mathbf{v} - (\nabla \mathbf{v})^T)$ its skew-symmetric part. Finally, \mathbf{u} denotes the velocity field of the grains, and we use the shorter notation $\dot{\boldsymbol{\epsilon}} := \mathbf{D}(\mathbf{u})$ for the so-called *strain rate* tensor. Note that $\text{Tr } \dot{\boldsymbol{\epsilon}} = \nabla \cdot \mathbf{u}$.

3.3.2 Macroscopic yield criteria

Continuous models for granular materials are derived upon macroscopic *yield criteria*, i.e., inequalities that should be satisfied by the principal stresses — the eigenvalues of the stress tensor $\boldsymbol{\sigma}$ — for the material to remain stable. The most well-known is the *Mohr-Coulomb* criterion. For cohesionless 3D materials such as dry sand, and with $\sigma_3 \leq \sigma_2 \leq \sigma_1$ the principal stresses, it states

$$\sigma_1 - \sigma_3 \leq \sin \varphi (\sigma_1 + \sigma_3),$$

where φ is the so-called friction angle, which corresponds to the inclination of a stable heap of the granular material under gravity.

In the space of principal stresses, the Mohr-Coulomb criterion defines a cone with hexagonal basis, which is numerically unwieldy. It is thus often discarded in favor of the *Drucker-Prager yield criterion* on the second invariant $J_2 = \frac{1}{2} \text{Tr}(\text{Dev } \boldsymbol{\sigma})^2$ of the deviatoric stress tensor, which in 3D reads

$$\sqrt{J_2} \leq -\hat{\mu} \frac{\sigma_1 + \sigma_2 + \sigma_3}{3} \quad \text{where } J_2 = \frac{1}{6} \sum_{i \neq j} (\sigma_i - \sigma_j)^2. \quad (3.15)$$

This criterion defines a second-order cone in the space of principal stresses, and $\hat{\mu}$ is called the friction coefficient. Note that in 2D, the Mohr-Coulomb and Drucker-Prager criteria are equivalent.

3.3.3 Nonsmooth rheology $\mathcal{DP}(\mu)$

The particularity of our fluid is to be non-Newtonian: our rheology relies upon a viscoplastic constitutive law combining the Drucker-Prager yielding criterion together with an unilateral incompressibility constraint.

As done usually, we decompose the solid phase stress tensor as $\boldsymbol{\sigma}_g := \eta \dot{\boldsymbol{\epsilon}} + \boldsymbol{\sigma}_c$, where the first part corresponds to a standard Newtonian viscosity (dissipative term due to random collisions in the flowing material), and $\boldsymbol{\sigma}_c$ is the additional stress due to the Coulomb interactions between individual grains.

Drucker-Prager rheology The contact normal stress is $p = -\frac{1}{d} \text{Tr } \boldsymbol{\sigma}_c$, such that $\boldsymbol{\sigma}_c = \text{Dev } \boldsymbol{\sigma}_c - p \mathbb{1}$. The second stress invariant can be computed through the identity $J_2 = \frac{1}{2} \text{Tr}(\text{Dev } \boldsymbol{\sigma}_c^2) = |\text{Dev } \boldsymbol{\sigma}_c|^2$. The Drucker-Prager yield criterion (3.15) therefore boils down to $|\text{Dev } \boldsymbol{\sigma}_c| \leq \hat{\mu} p$.

Moreover, the maximum dissipation principle states that in the yielded regime, friction should be saturated and the frictional stress tensor should be colinear to the deviatoric part of the strain rate. The deviatoric part of $\boldsymbol{\sigma}_c$ should thus satisfy one of the two regimes,

$$\begin{cases} \text{Dev } \boldsymbol{\sigma}_c = (\hat{\mu} p) \frac{\text{Dev } \dot{\boldsymbol{\epsilon}}}{|\text{Dev } \dot{\boldsymbol{\epsilon}}|} & \text{if } \text{Dev } \dot{\boldsymbol{\epsilon}} \neq 0 \text{ (yielded)} \\ |\text{Dev } \boldsymbol{\sigma}_c| \leq \hat{\mu} p & \text{if } \text{Dev } \dot{\boldsymbol{\epsilon}} = 0 \text{ (unyielded)}. \end{cases} \quad (3.16)$$

Unilateral compressibility constraint Most continuum-based models for granulars consider the fluid to be perfectly incompressible. In contrast, we want to take into account the typically asymmetric yielding behavior of granulars by allowing the fluid to expand as much as desired, while strictly preventing compaction.

With a dense flow hypothesis, the unilateral compressibility constraint can be expressed simply as $\nabla \cdot \mathbf{u} \geq 0$. We set the pressure p to enforce this inequality, i.e.,

$$\begin{cases} p \geq 0 & \text{if } \nabla \cdot \mathbf{u} = 0 \\ p = 0 & \text{if } \nabla \cdot \mathbf{u} > 0 \end{cases}$$

or, using an equivalent complementarity notation,

$$0 \leq p \perp \nabla \cdot \mathbf{u} \geq 0. \quad (3.17)$$

In our results, we have shown that relaxing the common incompressibility assumption $\nabla \cdot \mathbf{u} = 0$ prevents the arising of an ill-defined rheology in some typical scenarios such as the flow in the wake of an obstacle. For instance, in [Chauchat and Médale, 2014], where incompressibility is enforced, a negative pressure is observed behind the obstacle, which is outside the domain of validity of the well-accepted $\mu(I)$ rheology [Jop et al., 2006]. In contrast, our method guarantees that the pressure remains non-negative (see Figure 3.10). Moreover our complementarity constraint (3.17) naturally fits in with our numerical framework, without adding any computational cumbersomeness.

Solution set $\mathcal{DP}(\mu)$ Our full rheology (3.16 — 3.17) can be compactly rewritten as

$$\left\{ \begin{array}{ll} \text{Dev } \boldsymbol{\sigma}_c = -\frac{\hat{\mu}}{d} \text{Tr}(\boldsymbol{\sigma}_c) \frac{\text{Dev}(\dot{\boldsymbol{\epsilon}})}{|\text{Dev}(\dot{\boldsymbol{\epsilon}})|} & \text{if } \text{Dev } \dot{\boldsymbol{\epsilon}} \neq 0 \\ |\text{Dev}(\boldsymbol{\sigma}_c)| \leq \frac{\hat{\mu}}{d} \text{Tr}(\boldsymbol{\sigma}_c) & \text{if } \text{Dev } \dot{\boldsymbol{\epsilon}} = 0 \\ 0 \leq \frac{1}{d} \text{Tr}(\boldsymbol{\sigma}_c) \perp \text{Tr}(\dot{\boldsymbol{\epsilon}}) \geq 0, & \end{array} \right. \quad (3.18)$$

which is reminiscent of the Signorini-Coulomb law derived in Section 3.2.1, with $\dot{\boldsymbol{\epsilon}}$ playing the role of the relative velocity \mathbf{u} , and $\boldsymbol{\sigma}_c$ playing the role of the contact force \mathbf{r} . More precisely, by leveraging the following isomorphism converting the vector space to the space of symmetric tensors S_d ,

$$\begin{aligned} \chi : \mathbb{R} \times \mathbb{R}^{\mathbf{T}'} &\rightarrow S_d \\ (a; b, c) &\mapsto \begin{pmatrix} b & c \\ c & -b \end{pmatrix} + a\mathbb{I} && \text{if } d = 2 \\ (a; b, c, d, e, f) &\mapsto \begin{pmatrix} b - \frac{c}{\sqrt{3}} & d & e \\ d & -b - \frac{c}{\sqrt{3}} & f \\ e & f & \frac{2c}{\sqrt{3}} \end{pmatrix} \\ &+ \frac{\sqrt{2}}{\sqrt{3}} a\mathbb{I} && \text{if } d = 3, \end{aligned}$$

we can prove that our rheology is similar to the Signorini-Coulomb law, albeit expressed in higher dimension (in dimension 3 for the 2D case, and in dimension 6 for the 3D case) [Daviet and Bertails-Descoubes, 2016b].

Discretization After discretizing the full dynamic equations of our fluid using finite elements, together with an implicit discretization of the rheology (3.18), we build a discrete system which shares the exact same structure as our **DFCP** introduced in Section 3.2.2, though with higher dimension. We can thus leverage exactly the same tools as before: in particular, we use our efficient Gauss-Seidel solver featuring a modified Fischer-Burmeister (extended to higher dimensions) zero-finding solver as the local solver. Note however that the original enumerative solver was specifically designed for 3D, and cannot be trivially extended to higher dimensions.

Some results We consider a gravity-induced 2D granular medium inside a narrow channel (of width L), flowing around a cylinder of diameter $D = \frac{L}{4}$, with no-slip boundary conditions on the sides of the channel and on the cylinder. Figure 3.10 collects plots of the velocity and stress fields across the domain (with gravity directed from left to right). The pressure field is of special interest as it features a few notable phenomena:




Figure 3.10 – Simulation of a granular flow with our nonsmooth numerical method, with $\hat{\mu} = 0.3$. The velocity and stress fields are shown in the case where the granular flows from left to right in a narrow channel, around a cylindrical obstacle.

- First, it validates the benefit of allowing dilation of the flow, as the pressure in the wake (right part) of the obstacle is indeed zero, and does not become strictly negative as in [Chauchat and Médale, 2014];
- The zone of highest pressure is not located at the very front of the obstacle. Instead, we observe the formation of a high-pressure arch above this point;
- Above this arch, there exists a region where the pressure is lower than further upstream; moreover the transition between these two zones is very abrupt, as shown by the absence of two isosurfaces on the figure. Note that this

phenomenon disappears when considering free-slip boundary conditions for the channel walls.

3.3.4 Extension to a varying volume fraction

Our previous model does not consider a volume fraction field. Instead, the fluid is considered to be dense — that is, already at its maximal volume fraction ϕ_{max} — everywhere inside the simulation domain. The material is theoretically allowed to dilate, but not to compress, through the complementarity condition $0 \leq \nabla \cdot \mathbf{u} \perp p \geq 0$. However, in this model, the use of a fixed domain precludes an actual dilatation of the material through time. More fundamentally, even if the volume fraction field were somehow tracked through time (e.g., with particles), the $\nabla \cdot \mathbf{u} \geq 0$ condition would still prevent the material from recompacting after dilatating. We thus follow Narain and colleagues [2010] to account for a non-constant volume fraction ϕ , and instead enforce the complementarity condition

$$0 \leq \phi_{max} - \phi \perp p \geq 0. \quad (3.19)$$

where $\phi_{max} \leq 1$ is the maximum packing fraction for the material, beyond which grains cannot be compacted anymore. For monodisperse spherical grains, ϕ_{max} usually lies in the 0.55–0.65 range. We used $\phi_{max} = 0.6$ in all our simulations.

Considering this non-constant volume fraction ϕ , we rederived in [Daviet and Bertails-Descoubes, 2016a] the dynamics of our fluid together with the new Drucker-Prager condition. Using a semi-implicit time integration scheme together with a Material Point Method for spatial discretization, our final problem boiled down, again, to a discrete problem with a similar structure compared to our initial **DFCP**. This nice property allowed us to simulate free-flowing sand in several minutes per frame only (see Figure 3.11).

3.4 Conclusion: Towards a continuum-based model for fibrous materials

In this chapter we have motivated our choice for nonsmooth frictional contact laws, in order to capture the typical emerging effects due to contact and friction between solids. We have also presented a range of different formulations of the Signorini-Coulomb law, and presented two examples (hair and cloth) for which we adapt these formulations and design specific solvers, both robust and efficient. Finally, we have provided a few insights to our recent work on the continuum modeling of granular materials, which again leverage the nonsmooth solver we have developed in the context of discrete element modeling, albeit in higher dimension.

Towards a continuum-based model for fibrous materials Though we have recently made progress on the continuum formulation and solving of granular materials, we are still far from a continuum description of a macroscopic dry fibrous




Figure 3.11 – A steel ball impacts a granular bed with varying volume fraction. Unlike previous approaches, our fully resolved (non-linearized) Drucker-Prager rheology allows us to retrieve a perfectly round crater. Observe also the solid, liquid and gaseous phases of the granular flow that are retrieved by our model.

medium such as hair. One key ingredient that we have not been considering in our previous models is the influence of air inside divided materials. Typically, air plays a considerable role in hair motion. To advance in that direction, we have started to look at a diphasic fluid representation of granular matter, where a Newtonian fluid and the solid phase are fully coupled, while the nonsmooth Drucker-Prager rheology for the solid phase is enforced implicitly [Daviet and Bertails-Descoubes, 2017]. This first approach could be a starting point for modeling immersed granulars in a liquid, or ash clouds, for instance.

There still remains a long way to go for taking into account long fibers instead of isotropic grains in the solid phase. Coupling fiber elasticity with our current formulation is clearly an open and challenging problem that we would like to deal with in the future.

From geometry to mechanics: Inverse design of fibers

Contents

4.1	Inverse design problem: motivation and goal	58
4.2	Case of an isolated fiber	61
4.2.1	Necessary and sufficient condition for equilibrium	62
4.2.2	Sufficient condition for stability	63
4.2.3	Inverse design algorithm	64
4.2.4	From an arbitrary curve to a discrete material configuration	65
4.3	Case of a fiber assembly subject to frictional contact	66
4.3.1	Underdetermined conditions for equilibrium	68
4.3.2	A simple heuristics for a well-posed problem	69
4.4	From a geometric curve to a material curve	71
4.4.1	From a geometric curve to a material curve	72
4.4.2	Curve-angle parameterization	73
4.4.3	Direct problem: framing a Kirchhoff rod	74
4.4.4	Inverse static design problem	76
4.4.5	Discrete picture	78
4.5	Conclusion: Towards non-invasive identification	79

In a third part of our work, we have considered some new challenges regarding *inverse static design*. Indeed, while current simulators may succeed in reaching a good level of realism, they remain difficult to control in order to achieve a precise artistic goal or, more generally, to match a target observation. More precisely, to generate some desirable shapes and motions, one should be able to feed a simulator with the “right” parameters. Finding such parameters remains a very difficult task, which is often performed through a tedious trial and error process. To make this task fully automatic, we have started looking at *inverse* solutions in the case where a static shape is provided as input: the inverse model should be able to interpret automatically this shape as a stable equilibrium¹ of the simulator, under gravity and other external forces such as contact and friction.

¹In all this document, *equilibrium* refers to the *static equilibrium* of the structure, meaning that the sum of forces and the sum of torques applying onto the object vanish.

Summary of contributions In the case of an isolated fiber, we have shown that inverting any of our super-model [Bertails et al., 2006; Casati and Bertails-Descoubes, 2013] boils down to two decoupled problems that are both easy to solve [Derouet-Jourdan et al., 2010, 2011; Bertails-Descoubes, 2012]: first, an equilibrium condition which appears to be linear in the intrinsic shape of the fiber, thanks to the curvature-based parameterization of our fiber models; second, a sufficient stability condition that can be simply set by fixing a lower-bound for the ratio of stiffness over mass. Actually, the only remaining difficulty is to solve a merely geometric fitting problem — converting a curve as a piecewise helix or clothoid. In the case of helical fitting, we have already brought some efficient and robust solutions [Derouet-Jourdan et al., 2010, 2013b].

In the presence of contact and friction, Coulomb sticking constraints have to be considered, which makes the overall inverse problem nonsmooth and ill-posed. We have shown that assuming known mass and stiffness and a simplified inverse model, it is possible to recover a reasonable intrinsic shape as well as frictional contact forces at play, both of them satisfying exactly static equilibrium [Derouet-Jourdan et al., 2013a]. This work allowed us, for the first time, to animate a few hair geometries stemming from recent hair captures, such as the one depicted in Figure 4.3(a).

Finally, we are currently looking at the inverse problem in the continuous case, that is on the strong form of the Kirchhoff static equations [Bertails-Descoubes, 2017]. This ongoing work allows us to better characterize the space of solutions of the inverse problem in the case of an isolated fiber, through a generic curve-angle parameterization of the rod. This also allows us to understand how discrete models may automatically select a subspace of solutions, by imposing a certain form for the material frame of the rod. Such insights give us the hope to reach in the near future some actual *identification* process, i.e., a clear selection of *a* solution among the set of potential solutions.

Organization of the chapter Section 4.1 specifies our motivation for tackling the inverse design problem, and discards the generic nonlinear optimization approach in favor of a more effective approach dedicated to Kirchhoff rods. Sections 4.2 tackles the discrete inverse design problem relying upon super-models, in the case of an isolated fiber under gravity only. Then Section 4.3 takes a look at the more complex case of a fiber assembly subject to gravity and frictional contact. Finally, Section 4.4 elaborates on the problem of retrieving material curvatures from a merely geometric curve, first in the continuous case, then in the discrete case.

4.1 Inverse design problem: motivation and goal

Originally, our motivation stemmed out from a collaboration in 2007 with a French animation studio, *Néomis Animation*, who had a 3D movie project involving a number of characters with fancy hairstyles to be animated. An original sketch of such a hairstyle is depicted in Figure 4.1. Beyond their aesthetic appeal, these drawings




Figure 4.1 – Stylized hand-drawn hairstyle. The personality of the character is largely determined by the shape of her hair, hence such geometric features should be preserved during animation. Image courtesy of Néomis Animation Studio — Jeroen Dejonckheere.

actually suggested a very interesting scientific problem to me: how to let the artist be completely free to design the hair shape she/he desires, and to have the hair simulator *comply* with this input shape and predict its motion?

Obviously, the artist would draw hair as perceived within a physical world, that is, under external forces such as gravity and contacts. Moreover, at the character design stage, it is very likely that the drawn hairstyles intend to correspond to a hair configuration at *stable equilibrium*. However, such a shape is not readily usable by a physics-based simulator, which in contrast requires, among its parameters, the *intrinsic* shape² of the object to be simulated, that is *the shape it would have in the absence of external forces* (see Figure 4.2).

More generally, the intended hair design may come from various processes, ranging from artistic hairstyling to automatic hair capture. In all cases, one is left with, on the one hand, an *observed*³ (or *target*) shape \mathbf{q}^{obs} which integrates the effect of all surrounding forces, and on the other hand, a simulator which has to be fed with the right parameters, in particular $\bar{\mathbf{q}}$, so as to match the input shape \mathbf{q}^{obs} at stable equilibrium.

To simulate an elastic object from such a target shape \mathbf{q}^{obs} , a naïve though common practice so far in Computer Graphics was to use \mathbf{q}^{obs} as the intrinsic shape of the object, that is, to set $\bar{\mathbf{q}} = \mathbf{q}^{\text{obs}}$. However, when launching the simulator, the object would *sag* under gravity, thus losing the initial shape and ruining the design

²Recall that the intrinsic shape of a Kirchhoff rod was given in Chapter 2 by the intrinsic curvature function $\bar{\kappa}(s)$; in our finite-dimension Lagrangian settings, it corresponds to the vector $\bar{\mathbf{q}}$.

³In practice, our input data take the form of geometric curves in the case of rods, and surfaces in the case of plates/shells. Such primitives are represented with a finite number of degrees of freedom (for instance an ordered sequence of points, a spline, a triangulated mesh, etc.).




Figure 4.2 – Classical initialization of a physics-based simulator. The initial configuration $\mathbf{q}(t = 0)$ is set to a trivial intrinsic shape $\bar{\mathbf{q}}$ — here, a set of straight rods (a). Then, after feeding the simulator with other parameters (such as the mass and stiffness of each rod), a static shape \mathbf{q}^{eq} can be computed under gravity and frictional contacts (b).

work made by the artist, or the 3D geometry patiently reconstructed from real hair capture (see Figure 4.3).

Inverse design problem To address this problem effectively, an *inverse* modeling process⁴ thus has to be investigated. The inverse model should be able to interpret automatically the target shape \mathbf{q}^{obs} as a stable equilibrium of the simulator, under gravity and other external forces such as contact and friction.

As described in Section 2.2, the Kirchhoff rod model possesses two kinds of parameters. On the one hand, its intrinsic curvature $\bar{\kappa}$, discretized as vector $\bar{\mathbf{q}}$, and which may vary spatially along the centerline. On the other hand, its material parameters, composed of the linear mass density ρS and the bending/twisting stiffnesses of the rod $\mathbb{K}_3 = \text{diag}_3(\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2)$; material parameters are assumed to be constant with respect to space and time.

A generic and standard way to model our inverse problem is through least squares minimization [Kern, 2002]. That is, given the input configuration \mathbf{q}^{obs} , one seeks to solve the following constrained minimization problem,

$$\min_{\substack{(\mathbf{q}, \bar{\mathbf{q}}) \\ F(\mathbf{q}, \bar{\mathbf{q}}) \in \mathcal{K}}} \frac{1}{2} \|\mathbf{q} - \mathbf{q}^{\text{obs}}\|^2, \quad (4.1)$$

⁴In Mechanical Engineering, inverse problems are usually split in two categories [Beck and Woodbury, 1998]: (a) *inverse design* problems, where material parameters are known and one is looking for the intrinsic shape of the object, and (b) *inverse measurement* problems, where material properties are searched for. Our goal here is clearly to solve for an inverse design problem, but in the medium term, we would also like to extract some information about the material properties, from the input shape.



Figure 4.3 – When setting the intrinsic shape of the hair to the observed shape \mathbf{q}^{obs} (a), the simulator yields a new equilibrium configuration (b) which can be far from the observed shape. Acquired geometry (a) is data courtesy of Hao Li and colleagues.

where the constraint $F(\mathbf{q}, \bar{\mathbf{q}}) \in \mathcal{K}$ expresses a generalized (stable) equilibrium condition, possibly incorporating frictional contact forces when the set K is not reduced to $\{0\}$. The above formulation is nonlinear, and even nonsmooth in the presence of frictional forces: it is thus challenging to solve. We have recently attempted to solve this problem in the case of nodal models for cloth using the adjoint method, and have reported our results in the research report [Casati et al., 2016]. In particular, we have shown how to extend the adjoint method — classically used in the context of bilateral constraints — to deal with frictional contact constraints (see details in the PhD thesis of R. Casati [Casati, 2015]). Though promising, such a method is, even in the contactless case, prone to convergence issues inherent to nonlinear optimization, and much work still remains to be done to improve the robustness of our inverse process.

Inspecting Kirchhoff equations as well as our discrete curvature-based models (described in Chapter 2), we have fortunately noted that a much more effective inverse model could be built, which avoids nonlinear optimization. Result is not only the design of a simple, fast and robust inversion process, but also a better characterization of the whole space of solutions, which opens the way for a reliable parameter identification process in the short term.

4.2 Case of an isolated fiber

Let us consider a discrete curvature-based model (“super-model”) clamped at one end, and subject to gravity only (see Figure 4.4(b)). We assume the configuration \mathbf{q}

to be given as input, and we are looking for the parameters $\bar{\mathbf{q}}$ (intrinsic shape), ρS (linear mass density and $\mathbb{K}_3 = \text{diag}_3(\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2)$ (bending and twisting stiffnesses) so that \mathbf{q} corresponds to a stable equilibrium of the super-model.

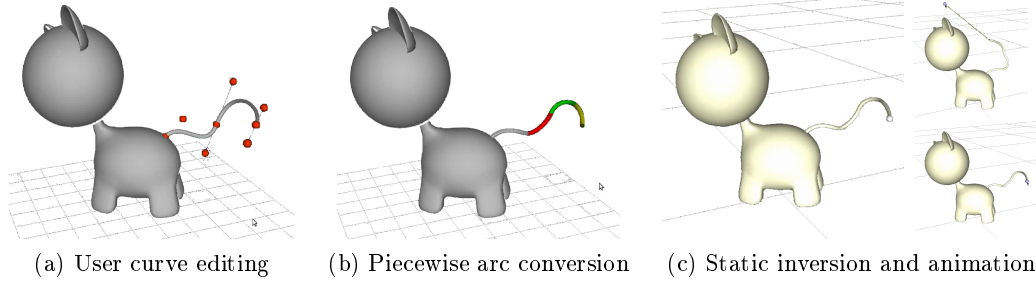


Figure 4.4 – Example of fiber design session. The user first carves the tail of the cat geometrically as she desires (a), then converts this curve into a piecewise helical arc curve (b) using our floating tangent algorithm (see Section 4.2.4), and finally automatically computes the parameters of a super-helix so that its centerline matches the input curve at stable equilibrium under gravity (c).

4.2.1 Necessary and sufficient condition for equilibrium

Let us first write down the equilibrium equation. From Equation (2.8), dropping all time-dependent terms and making parameter dependence explicit with brackets yields

$$\mathbb{K}_{[\mathbb{K}_3]}(\mathbf{q} - \bar{\mathbf{q}}) + \mathbf{G}_{[\rho S]}(\mathbf{q}) = 0, \quad (4.2)$$

where

$$\mathbb{K}_{[\mathbb{K}_3]} = \mathbb{K}_3 \otimes \mathbb{L}_N \quad \text{and} \quad \mathbf{G}_{[\rho S]}(\mathbf{q}) = -\rho S \underbrace{\int_0^L \left(\frac{\partial \mathbf{r}}{\partial \mathbf{q}} \right)^\top \mathbf{g} \, ds}_{\mathcal{G}(\mathbf{q})},$$

where \mathbb{L}_N is a symmetric positive-definite matrix of size N , which only depends on the length $\boldsymbol{\ell} = \{\ell_0, \dots, \ell_{N-1}\}$ of the rod elements. For a super-helix, \mathbb{L}_N is simply diagonal (see Expression (2.11)), whereas for a super-clothoid, \mathbb{L}_N is tridiagonal (see Expression (2.12)).

Note that if $E_p(\mathbf{q})$ is the potential energy of the rod, Equation (4.2) is equivalent to writing

$$\nabla E_p[\bar{\mathbf{q}}, \mathbb{K}_3, \rho S](\mathbf{q}) = 0,$$

meaning that we are looking for the parameters $\bar{\mathbf{q}}$, \mathbb{K}_3 and ρS so that \mathbf{q} is a critical point of energy of the rod.

Assuming the material parameters ρS and \mathbb{K}_3 to be constant, Equilibrium (4.2)

is exactly satisfied for a unique family of solutions,

$$\begin{aligned}\bar{\mathbf{q}} &= \mathbf{q} + \mathbb{K}_{[\mathbb{K}_3]}^{-1} \mathbf{G}_{[\rho S]}(\mathbf{q}) \\ &= \mathbf{q} - (\rho S \operatorname{diag}_3(\mathcal{K}_0^{-1}, \mathcal{K}_1^{-1}, \mathcal{K}_2^{-1}) \otimes \mathbb{L}_N^{-1}) \mathcal{G}(\mathbf{q})\end{aligned}\quad (4.3)$$

using inverse of Kronecker product [Prasolov, 1994, Section 27.4],

where $\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2$ and ρS can span the entire positive space.

In practice, solving (4.3) amounts to choosing a positive set of parameters $\{\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2, \rho S\}$, and then computing $\bar{\mathbf{q}}$ by solving a full rank linear system of $\sim 3N$ equations, which turns out to be diagonal in the case of super-helices, and tridiagonal in the case of super-clothoids. The number of elements N being generally low (≤ 20), the latter problem is thus extremely fast to solve.

4.2.2 Sufficient condition for stability

Evaluating the stability of the equilibrium requires the computation of the Hessian matrix $\nabla^2 E_p$ of potential energy. The equilibrium will be stable if $\nabla^2 E_p$ is a positive-definite matrix, i.e., if all its eigenvalues are positive.

Differentiating the left-hand member of Equation (4.2) gives the expression for the Hessian of potential energy,

$$\nabla^2 E_{p[\mathbb{K}_3, \rho S]} = \mathbb{K}_3 \otimes \mathbb{L}_N + \rho S \mathbb{S}(\mathbf{q}), \quad (4.4)$$

where \mathbb{S} is a real symmetric matrix of size $3N$ which (nonlinearly) depends on the input configuration \mathbf{q} , but remains independent of our unknown parameters. Note that $\nabla^2 E_p(\mathbf{q})$ thus only depends on the material parameters \mathbb{K}_3 and ρS , in a linearly way, while dependence w.r.t. the intrinsic shape $\bar{\mathbf{q}}$ has vanished.

Using Expression (4.4), we can find a sufficient condition for the Hessian $\nabla^2 E_p$ to be positive-definite. Indeed, Horn's theorem [Fulton, 2000] implies that the minimum eigen value of a sum of real symmetric matrices is greater than or equal to the sum of the eigen values that are minimum for each matrix. Let $\lambda^m, \mu^m(\mathbf{q}), \tau^m$ be the minimum eigen values of $\nabla^2 E_p, \mathbb{S}$, and \mathbb{L}_N , respectively. Let $\mathcal{K} = \min\{\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2\}$. From [Prasolov, 1994, Section 27.4], we deduce that the minimum eigen value of the Kronecker product $\mathbb{K}_3 \otimes \mathbb{L}_N$ is the product of minimum eigen values $\mathcal{K} \tau^m$. Applying Horn's theorem then gives

$$\lambda^m \geq \mathcal{K} \tau^m + \rho S \mu^m(\mathbf{q}) \quad (4.5)$$

with $\mathcal{K} > 0$ and $\tau^m > 0$. Inequality (4.5) provides a lower bound for the smallest eigenvalue of $\nabla^2 E_p$. It thus yields a sufficient condition on \mathbb{K}_3 and ρS for guaranteeing that all the eigenvalues of $\nabla^2 E_p$ are positive,

$$\frac{\mathcal{K}}{\rho S} > -\frac{\mu^m(\mathbf{q})}{\tau^m}. \quad (4.6)$$

Recall that the two eigen values $\mu^m(\mathbf{q})$ and τ^m can be directly computed from the input configuration \mathbf{q} and the element lengths ℓ , using standard eigen value extraction algorithms⁵.

According to Inequality (4.6), by either increasing \mathcal{K} or decreasing ρS , the stability of the rest shape under gravity can thus *always* be guaranteed, whatever the input configuration \mathbf{q} . This is actually in line with intuition. Imagine an input fiber oriented downwards. To be stable, this configuration will not require a high stiffness since it is already pointing towards a direction that is preferred by gravity. Now, if the fiber is oriented upwards, a large stiffness (or equivalently, a small mass) will be required to guarantee that the fiber can stably holds in this position and will not bend downwards. This is quite similar to applying gel on hair fibers so that they comply to some arbitrary desired shapes, even if those are in strong contradiction with gravity forces.

4.2.3 Inverse design algorithm

Computing the parameters of the super-model so as to match an arbitrary configuration \mathbf{q} at stable equilibrium is thus elementary in the case where only gravity is involved. The user may first set the stiffness \mathcal{K} and the linear mass density ρS so that their ratio lies in the authorized halfspace given by (4.6), then compute the intrinsic curvature $\bar{\mathbf{q}}$ satisfying equilibrium by solving the system of equations (4.3). If unhappy with the motion of the fiber which looks to stiff or too soft, she can modify the stiffness, check that stability is still guaranteed (if not, automatically retrieve the minimum stiffness satisfying stability), and finally recompute the intrinsic curvature yielding equilibrium. This simple design process is summarized in Figure 4.5, and an example of interactive design session is illustrated in Figure 4.4.

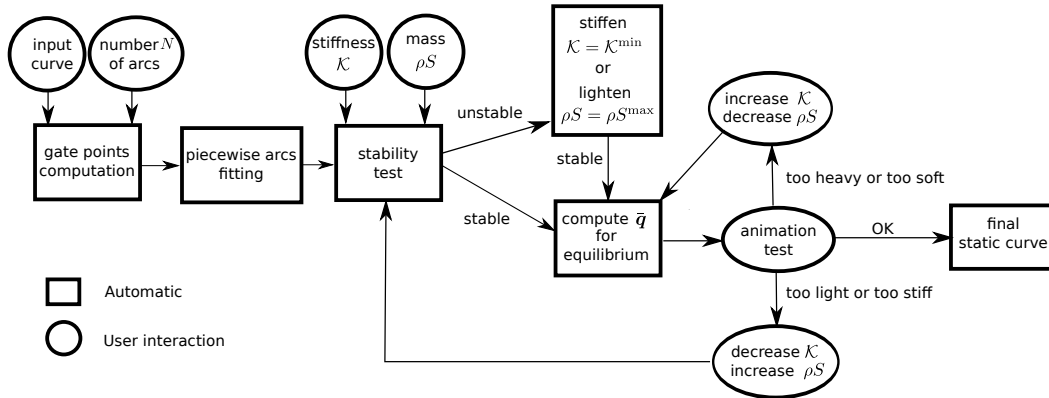


Figure 4.5 – Our inverse design process for an isolated fiber.

Note that our stability criterion (4.6) only provides a *sufficient* condition for sta-

⁵In the case of a super-helix, the minimum eigen value τ^m boils down to the minimum length of elements, $\min_i\{\ell_i\}$.

bility, meaning that equilibrium may still be stable when decreasing stiffness beyond the suggested lower bound. In practice however, we have noticed that violating this condition always causes equilibrium to become unstable. This gives us the hope to establish in the short term a *necessary* condition for stability, useful this time not just for designing stable rods matching an input curve, but for characterizing precisely the range of potential stiffness and mass parameters from a real observed rod under gravity (see Conclusion of this chapter).

4.2.4 From an arbitrary curve to a discrete material configuration

So far we have considered that a discrete material configuration \mathbf{q} was provided as input. However, this is not exact. The true input data is an arbitrary curve $\Gamma(s)$, represented with a finite number of degrees of freedom, for instance as a spline (see Figure 4.4(a)). Two steps then need to be completed for reconstructing an admissible material configuration \mathbf{q} , before performing inverse statics.

The first step is to convert the arbitrary input curve into a *kinematically admissible* centerline $\mathbf{r}(s)$ for a super-model, that is, into a C^1 -smooth piecewise helical curve if dealing with the super-helix model, or a C^1 -smooth piecewise clothoidal⁶ curve if dealing with the super-clothoid model.

Once a piecewise arc curve $\mathbf{r}(s)$ has been reconstructed, a second step consists in reconstructing admissible material curvatures and twist \mathbf{q} for the discrete rod, which is equivalent to reconstructing a kinematically admissible *material frame* $\mathcal{R}(s)$ for the rod. To anticipate on the following, we have shown that the only admissible material frame for a super-helix consists in building a continuous frame which shares, on each element, the same Darboux vector as the Frenet frame. We defer the corresponding proof to Section 4.4 where the general question of retrieving a material frame from a curve will be deeply examined, from the continuous (Kirchhoff) case to the discrete case. Here, we focus only on the first step, that is on the geometric problem aiming at converting an arbitrary curve into a piecewise arc curve.

Piecewise helical fitting In our work we have tackled the first step mentioned above in 3D in the case where the target super-model is a super-helix⁷. That is, we have addressed the geometric problem of converting an arbitrary input curve into a C^1 -smooth piecewise helix.

⁶In 3D, the notion of a “clothoidal” curve is not clear given our definition of a space clothoid element in Chapter 2. Indeed, unlike what happens in the piecewise uniform curvature case (super-helix), assuming that the material curvatures and twist are piecewise linear does *not* imply that the geometric (or Frenet) curvature and torsion of the centerline are piecewise linear (the reader may take a look at Equations (4.16) to get convinced about it). As noted already in Chapter 2, the shape of the centerline of the super-clothoid model is thus more general than the so-called 3D Euler spiral [Harary and Tal, 2012].

⁷In 2D, approximation algorithms are much simpler to design due to the absence of torsion. Our 3D floating tangent algorithm directly translates to the 2D problem consisting in fitting a curve to a smooth piecewise circular curve [Derouet-Jourdan et al., 2010]. Moreover, efficient methods for fitting a 2D curve to a smooth piecewise clothoid curve have been designed, for instance [McCrae and Singh, 2008] that we leverage in [Bertails-Descoubes, 2012].

In a nutshell, our algorithm consists in segmenting the input rod into N pieces of equal length, and approximating each curve segment with a helix under the constraint that two neighboring helical segments should have a C^1 -smooth connection, i.e., tangents should remain continuous at the junction. To achieve this, we enforce tangents at junction points to match initial tangents, but since having both tangents and positions fixed may not be compatible with an interpolating helix, we relax the initial positions of the junction points, and minimize the error between admissible and initial positions. The key of our approach relies upon the following theorem [Ghosh, 2010; Derouet-Jourdan et al., 2013b],

Theorem 4.1. *Given two points \mathbf{p}_0 and \mathbf{p}_1 such that $\mathbf{p}_0 \neq \mathbf{p}_1$ and two tangents \mathbf{t}_0 and \mathbf{t}_1 such that $\mathbf{t}_1 \neq \pm \mathbf{t}_0$, there exists a unique short helix starting at \mathbf{p}_0 with tangent \mathbf{t}_0 and ending at \mathbf{p}_1 with tangent \mathbf{t}_1 if and only if*

$$\langle \mathbf{p}_1 - \mathbf{p}_0, \mathbf{t}_1 - \mathbf{t}_0 \rangle = 0, \quad (4.7)$$

which allows us to formulate admissible positions recursively as a simple linear function of the input tangents. Our algorithm, coined “floating tangents algorithm”, then simply amounts to solving a full-rank linear system of size N ; this obviously proves order of magnitude faster compared to a nonlinear least-squares optimization method (see Figure 4.7)

Theorem 4.1, illustrated in Figure 4.6, was first stated in Ghosh’s PhD thesis [Ghosh, 2010], but proof was incomplete. We have completed the missing parts in [Derouet-Jourdan et al., 2013b], especially thanks to the PhD work of Alexandre Derouet-Jourdan [Derouet-Jourdan, 2013]. Note that the proof of existence is constructive, and provides an efficient algorithm for building the unique short helix once Condition (4.7) is satisfied.

4.3 Case of a fiber assembly subject to frictional contact

While inverse design of an isolated fiber is interesting from a theoretical point of view, and may even apply to a few examples such as the tail of an animal or a plant stem, in practice most useful scenarios involve *interactions* with other objects, and typically *contact* and *friction* forces.

We have thus started to investigate the inverse design problem for a fiber subject not only to gravity but also to frictional contact forces. More specifically, we have considered a fiber assembly subject to external contacts (with a body) as well as fiber-fiber contacts, in the presence of Coulomb friction. A typical example of such a system is a head of hair, usually modelled in Computer Graphics with a few hundreds to thousands thin elastic rods. With the recent advance in 3D hair reconstruction from real hair photographs, many geometric datasets, composed of 3D curves which resemble real hair shapes, are nowadays available (see Figure 4.8(a)).




Figure 4.6 – In red, the unique short helix that interpolates \mathbf{p}_0 and \mathbf{p}_1 with tangents \mathbf{t}_0 and \mathbf{t}_1 when Equation (4.7) is satisfied. The helix is said to be short if $\varphi^* \leq 2\pi$.

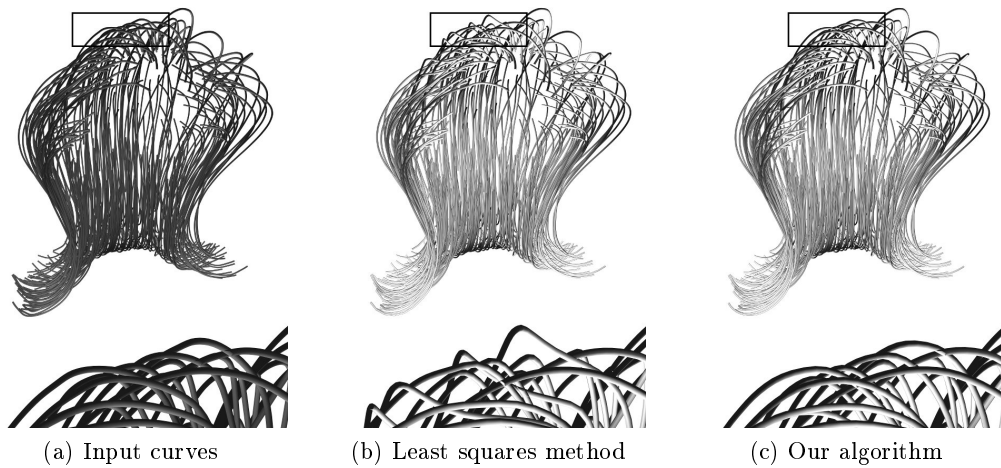


Figure 4.7 – Comparison of helical approximations (with $N = 15$ elements) on a input curve dataset (a), between least-squares optimization (b) and our floating tangents algorithm (c). Our method proves not only more accurate, but also faster ($10\times$ on this example) than least-squares optimization.




Figure 4.8 – Real curly wig (a) reconstructed in 3D by Luo et al. [2013], inverted by our method in [Derouet-Jourdan et al., 2013a] and physically animated (b) and trimmed (c).

4.3.1 Underdetermined conditions for equilibrium

Let us consider a dynamic system composed of a set of $N_{\mathcal{F}}$ fibers discretized as super-models. Let \mathbf{q} be the vector containing the m generalized coordinates of the full system. We assume the system to be subject to n frictional contacts, which may occur between a fiber and an external rigid object, or between two different fibers, or within a single fiber.

Using notations of Chapter 3, our previous equilibrium equation (4.2) gets transformed into the constrained problem

$$\begin{cases} \mathbb{K}_{[\mathbb{K}_3^1, \mathbb{K}_3^2, \dots, \mathbb{K}_3^{N_{\mathcal{F}}}] } (\mathbf{q} - \bar{\mathbf{q}}) + \mathbf{G}_{[\rho S^1, \rho S^2, \dots, \rho S^{N_{\mathcal{F}}}] } (\mathbf{q}) = \mathcal{H}(\mathbf{q})^\top \mathbf{r} \\ \forall i = 1 \dots n, \quad \mathbf{r}^i \in \mathcal{K}_{\mu_i}, \end{cases} \quad (4.8)$$

meaning that elastic and gravitational forces should balance contact forces, and all contacts should be in sticking mode (no sliding velocity), and thus each local contact force \mathbf{r}^i should belong to the interior of the Coulomb friction cone \mathcal{K}_{μ_i} .

Compared to Section 4.2, new parameters enter the game: the n friction coefficients μ_i corresponding to the n contacts. Material parameters (stiffness and linear mass density) are also augmented with the number of fibers, and likewise the intrinsic curvature vector $\bar{\mathbf{q}}$, which is of same size as \mathbf{q} . Finally, as in the direct problem, the n contact forces \mathbf{r}^i collected in vector \mathbf{r} are part of the unknowns.

To simplify the problem, we have so far only considered $\bar{\mathbf{q}}$ and \mathbf{r} as unknowns. All the $N_{\mathcal{F}}$ material parameters \mathbb{K}_3^j and ρS^j as well as the n friction coefficients μ_i

are thus assumed to be fixed in advance by the user (and explicit dependence to those denoted by brackets will be omitted from now on). Our problem amounts to searching for intrinsic shapes $\bar{\mathbf{q}}$ and contact forces \mathbf{r} which are compatible with Equilibrium (4.8),

$$\begin{cases} \bar{\mathbf{q}} = \mathbf{q} + \mathbb{K}^{-1} \left(\mathbf{G}(\mathbf{q}) - \mathcal{H}(\mathbf{q})^\top \mathbf{r} \right) \\ \mathbf{r} \in \mathcal{K}_\mu \quad \text{with } \mathcal{K}_\mu = \mathcal{K}_{\mu_1} \times \mathcal{K}_{\mu_2} \times \dots \times \mathcal{K}_{\mu_n}. \end{cases} \quad (4.9)$$

One immediate observation is that Problem (4.9) is underdetermined, as any valid choice for \mathbf{r} leads to a mathematically valid solution for $\bar{\mathbf{q}}$ (but not the reverse). In particular, it is possible to choose $\mathbf{r} = \mathbf{0}$, meaning that contact forces are inactive; one is then left with the same solution for $\bar{\mathbf{q}}$ as in the isolated case treated in Section 4.2. However, this choice may be fairly unrealistic in many situations, as illustrated in Figure 4.9, where it is perceptible that contact definitely plays a role in the observed configuration.




Figure 4.9 – Two examples of input configurations \mathbf{q} that should, intuitively, be accounted for by contact forces and not just by the intrinsic shape $\bar{\mathbf{q}}$ of the rods.

4.3.2 A simple heuristics for a well-posed problem

Our idea to better pose the inverse problem (4.9) is to prevent the intrinsic shape $\bar{\mathbf{q}}$ from taking any fancy value, and instead to restrict its value to be close to a good “guess” $\bar{\mathbf{q}}^0$. For instance, in Figure 4.9, it is unrealistic to think that the sudden curvature deviation of the curves is due to a brutal change in the intrinsic shape $\bar{\mathbf{q}}$, especially if we have in mind representing natural fibers (hair, plant stem) which are rather homogeneous; instead, a more realistic assumption is to consider that $\bar{\mathbf{q}}$ is probably vanishing, and thus, that the resulting curved shape is mainly explained by contact itself.

More generally, we may formalize our problem as finding the contact forces \mathbf{r} which minimize the drift between an admissible intrinsic shape $\bar{\mathbf{q}}$, which satisfies

Equilibrium (4.9), and an initial guess $\bar{\mathbf{q}}^0$. More precisely, our new problem reads

$$\begin{aligned} \min_{\mathbf{r}} \quad & \frac{1}{2} \left\| \overbrace{\mathbf{q} - \mathbb{K}^{-1} \left(\mathcal{H}(\mathbf{q})^\top \mathbf{r} - \mathbf{G}(\mathbf{q}) \right)}^{\bar{\mathbf{q}}} - \bar{\mathbf{q}} \right\|^2 + \gamma \|\mathbf{r}\|^2 \\ \text{s.t.} \quad & \mathbf{r} \in \mathcal{K}_\mu, \end{aligned} \quad (4.10)$$

where we have introduced a regularization term $\gamma \|\mathbf{r}\|^2$ for preventing contact forces to reach overly big values.

Problem (4.10) may be put under the canonical form

$$\begin{aligned} \min_{\mathbf{r}} \quad & \frac{1}{2} \mathbf{r}^\top \mathbf{W} \mathbf{r} + \mathbf{b}^\top \mathbf{r} \\ \text{s.t.} \quad & \mathbf{r} \in \text{int}(\mathcal{K}_\mu), \end{aligned} \quad (4.11)$$

with $\mathbf{W} = \mathcal{H}(\mathbf{q}) \mathbb{K}^{-2} \mathcal{H}(\mathbf{q})^\top + \gamma \mathbb{I}$ and $\mathbf{b} = -\mathcal{H}(\mathbf{q}) \mathbb{K}^{-2} (\mathbf{G}(\mathbf{q})(\mathbf{q}) + \mathbb{K}(\mathbf{q} - \bar{\mathbf{q}}^0))$, where \mathbb{I} is the identity matrix. Note that the objective function $f(\mathbf{r}) = \frac{1}{2} \mathbf{r}^\top \mathbf{W} \mathbf{r} + \mathbf{b}^\top \mathbf{r}$ is quadratic, and has to be minimized under second-order conic constraints: our problem is a second-order cone quadratic program (SOCQP), and belongs to the family of the so-called quadratically constrained quadratic programs (QCQP) [Boyd and Vandenberghe, 2004]. Since \mathbf{W} is symmetric positive-definite (definiteness comes from the regularization term), our problem is strictly convex and admits a unique solution \mathbf{r} . Refining our inverse model by taking a priori information on the parameters we wish to retrieve has thus allowed us, unsurprisingly, to transform the initial underdetermined problem (4.9) into the well-posed convex problem (4.11).

Choosing a good guess For now we have assumed the intrinsic curliness of a given hairstyle could be guessed easily. If one adopts the simplistic assumption that the fiber grows regularly with a uniform intrinsic shape, it is indeed possible to retrieve this value approximately. Boundary conditions for thin elastic rods subject to gravity tell us that at the free end of hair fibers, the actual curvature equals the intrinsic curvature (see Equations (2.6d) and (2.5)). Provided no contact substantially deforms the tip of the fiber, one may thus simply measure the actual curvature at the tip to get a good estimation of the intrinsic curvature.

However, real fibers seldom feature a purely uniform intrinsic shape, which is in reality modulated by small defects due to a non-perfect growth process or history involving irreversible (plastic) changes. When dealing with input data coming from real hair captures, we found out that taking as an estimation the full actual configuration \mathbf{q} yielded better results than taking a uniform intrinsic curvature matching \mathbf{q} at the tip only. Looking back at our minimization problem (4.10), this choice can be interpreted mechanically. Taking $\bar{\mathbf{q}}^0 = \mathbf{q}$ means that we are searching for contact forces that guarantee an exact equilibrium state for hair while minimizing hair internal elastic energy. This implies that, *as far as possible*, we rely on the contact forces to compensate for gravity. Of course, as the number of contact points is arbitrarily sparse and as the contact forces are bounded and constrained to belong to the

friction cone, contact is unlikely to support any arbitrary input shape on its own. In the (common) case where contacts are not sufficient to compensate for gravity, then the elastic energy becomes active and also contributes to the equilibrium.

Efficient solving The optimality conditions of our SOCQP (4.11) are equivalent to the following complementarity condition,

$$\forall i = 1 \dots n, \quad \mathcal{K}_{\perp}^i \ni (\nabla f(\mathbf{r}))^i \perp \mathbf{r}^i \in \mathcal{K}_{\mu_i}. \quad (4.12)$$

Interestingly, this condition is mathematically similar to the complementarity condition (3.9) of the dynamic case, except that the dynamic variable $\tilde{\mathbf{u}}^i$ has been replaced with $(\nabla f(\mathbf{r}))^i$, i.e., the i^{th} 3×1 block of the gradient. Now, our solver derived in [Daviet et al., 2011] was designed to solve exactly this kind of conic complementarity, i.e., problems which can be formulated as

$$\forall i = 1 \dots n, \quad K_{\perp}^i \ni (\mathbf{W}\mathbf{r} + \mathbf{c})^i \perp \mathbf{r}^i \in K_{\mu_i},$$

where \mathbf{W} is a symmetric positive (semi-)definite matrix of size $(3n, 3n)$, and \mathbf{c} a vector of size $3n$. In the direct (dynamic) case, \mathbf{W} was identified to the sparse (discrete) Delassus operator $\mathbf{H}\mathbf{M}^{-1}\mathbf{H}^{\top}$ [Daviet et al., 2011]. In our static inversion problem, \mathbf{W} is identified to our \mathbf{W} matrix, which is sparse, symmetric and positive-definite.

We can thus apply our **DFCP**'s Gauss-Seidel solver straightforwardly to solve our inverse problem (4.12), and thus our initial SOCQP (4.11). Unlike the dynamic problem (3.5), our inverse problem (4.12) is convex. It is thus likely to be easier to solve. In practice, simply using the primary Fischer-Burmeister solver (with no fail-safe) proved to converge well for all the problems we have tested.

4.4 From a geometric curve to a material curve

One important problem that we have put under the carpet so far is the following: Given a piecewise helical (or clothoidal) curve, how is it possible to reconstruct the material frame $\mathcal{R}(s)$ — or equivalently, the material curvatures and twist \mathbf{q} — of the corresponding super-model? This question is indeed fundamental, because it raises the issue of transforming a merely *geometric* curve, i.e., our input curve, into a *material* curve, i.e., the geometry of a (discrete) Kirchhoff rod, composed of a centerline $\mathbf{r}(s)$ and of a material frame $\mathcal{R}(s)$. While the input curve only carries geometric information, the material rod carries some mechanical information, such as how matter deforms along the centerline.

To investigate this question, we shall first have a look at the continuous picture, taking as input a general curve $\Gamma(s)$ with no prescribed shape. We shall establish the necessary and sufficient conditions enforced by the (strong) Kirchhoff equations on the framing of $\Gamma(s)$, both in the *direct* case (i.e., when the intrinsic curvature $\bar{\kappa}(s)$ is known), and in the *inverse* case (i.e., when the $\bar{\kappa}(s)$ is not known and looked

for to guarantee the equilibrium of the rod). Then we shall examine the special case of our super-models, as well as that of other discrete models of the literature, in particular the discrete elastic rod model [Bergou et al., 2008, 2010].

Note that this section partly covers unpublished material [Bertails-Descoubes, 2017], and we present here only the main conclusions of our study. Since the defense of this habilitation, more results on the inverse design problem have been obtained and published in [Bertails-Descoubes et al., 2018].

4.4.1 From a geometric curve to a material curve

Let $\Gamma(s)$ be a C^2 -smooth curve⁸ of length L parameterized by its arclength $s \in [0, L]$. The idea is to equip $\Gamma(s)$, which is a merely geometric curve, with a material frame $\mathcal{R}(s) = (\mathbf{n}_0(s), \mathbf{n}_1(s), \mathbf{n}_2(s))$, supposed to be *orthonormal* and *adapted* to the curve $\Gamma(s)$, i.e., such that $\mathbf{n}_0(s) = \Gamma'(s)$ and the two other axes \mathbf{n}_1 and \mathbf{n}_2 lie in the normal plane $\mathcal{P}^\perp(s)$ of the curve (see Figure 4.10(a)).




Figure 4.10 – Framing an arbitrary curve $\Gamma(s)$ with an adapted frame $\mathcal{R}(s)$, parameterized by the angle θ with respect to a reference frame $\mathcal{R}^{\text{ref}}(s)$.

Then, we shall further assume that $\mathcal{R}(s)$ is C^1 -smooth and that the frame $\mathcal{R}(s)$ is transported along the curve $\Gamma(s)$ through infinitesimal rotation about a Darboux vector $\mathbf{\Omega}(s)$, just similarly as in Section 2.2 of Chapter 2. This mathematically reads

$$\begin{aligned} \mathcal{R}'(s) &= [\mathbf{\Omega}(s)]_\times \mathcal{R}(s) \\ &= \mathcal{R}(s) [\mathbf{\kappa}(s)]_\times, \end{aligned}$$

⁸At this stage, required degree of smoothness varies upon the reference frame that will be chosen in the curve-angle parameterization (see Section 4.4.2). If the Frenet frame is chosen, C^3 -continuity is required to define torsion as a continuous function of s (provided curvature does not vanish), allowing us to interpret the Frenet frame as a C^1 -smooth material frame.

where $\boldsymbol{\kappa}(s) = \{\kappa_0(s), \kappa_1(s), \kappa_2(s)\}$ is the *material curvature* vector, which collects the local coordinates (the so-called material twist and curvatures) of the Darboux vector in the material frame $\mathcal{R}(s)$.

Note that at this stage, we have already converted our input curve $\Gamma(s)$ into a *material* curve $\{\mathbf{r}(s); \mathcal{R}(s)\}$ with only material twist and curvatures as degrees of freedom. This however does not mean it is an acceptable frame for a Kirchhoff rod, as Kirchhoff equations impose further assumptions about the material frame. We shall now examine such conditions, in two different settings:

Direct case: We shall study the law $\mathcal{R}(s)$ should obey when $\Gamma(s)$ corresponds to the centerline of a Kirchhoff rod, with all parameters fixed;

Inverse static case: We shall study the law the intrinsic curvatures $\bar{\kappa}(s)$ should obey when $\Gamma(s)$ corresponds to the centerline of a Kirchhoff rod at equilibrium.

4.4.2 Curve-angle parameterization

Let us first reduce the dimension of our problem, by observing that $\mathcal{R}(s)$ can only rotate about the tangent vector $\mathbf{n}_0(s)$, since it should remain orthonormal and adapted to $\Gamma(s)$. This degree of freedom can be parameterized by an angle $\theta(s)$ representing the rotation angle about $\mathbf{n}_0(s)$ between a *reference* material frame $\mathcal{R}^{\text{ref}}(s)$, also orthonormal and adapted to $\Gamma(s)$, and our material frame $\mathcal{R}(s)$. The relationship between $\mathcal{R}(s)$ and $\mathcal{R}^{\text{ref}}(s)$ is illustrated in Figure 4.10(b), and mathematically reads

$$\mathcal{R}(s) = \mathcal{R}^{\text{ref}}(s) \mathcal{R}_\theta(s), \quad (4.13)$$

where $\mathcal{R}_\theta(s) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta(s) & -\sin \theta(s) \\ 0 & \sin \theta(s) & \cos \theta(s) \end{bmatrix}$ is the rotation matrix of axis \mathbf{e}_x and angle θ , which represents the change of basis matrix from $\mathcal{R}^{\text{ref}}(s)$ to $\mathcal{R}(s)$.

We further suppose that the reference frame $\mathcal{R}^{\text{ref}}(s)$ can be fully determined by the curve $\Gamma(s)$: this is typically the case for the Frenet frame and the Bishop frame⁹, which are explicitly taken as examples of reference frames hereafter.

Finally, we assume the reference frame to be sufficiently smooth for being interpreted as a material frame transported along the curve $\Gamma(s)$ through infinitesimal rotation about a Darboux vector $\boldsymbol{\Omega}^{\text{ref}}(s)$, similarly to $\mathcal{R}(s)$. This allows us to define a reference material curvature $\boldsymbol{\kappa}^{\text{ref}}(s)$, which gives the coordinates of $\boldsymbol{\Omega}^{\text{ref}}(s)$ in the local frame $\mathcal{R}^{\text{ref}}(s)$.

After a few derivations, we can express our (unknown) curvature vector $\boldsymbol{\kappa}$ function of the (known) curvature $\boldsymbol{\kappa}^{\text{ref}}$, as

$$\boldsymbol{\kappa}(s) = \mathcal{R}_{-\theta}(s) \left(\boldsymbol{\kappa}^{\text{ref}}(s) + \theta'(s) \mathbf{e}_x \right), \quad (4.14)$$

⁹Note that unlike the Frenet frame which is uniquely determined by the curve (provided some extra regularity assumptions for $\Gamma(s)$ are met, as mentioned earlier), the Bishop frame is not unique but determined by the curve up to a rigid rotation: the initial frame at $s = 0$ may be arbitrarily chosen.

which can be expanded into

$$\begin{cases} \kappa_0(s) = \kappa_0^{\text{ref}}(s) + \theta'(s) & (4.15a) \\ \kappa_1(s) = \cos \theta(s) \kappa_1^{\text{ref}}(s) + \sin \theta(s) \kappa_2^{\text{ref}}(s) & (4.15b) \\ \kappa_2(s) = -\sin \theta(s) \kappa_1^{\text{ref}}(s) + \cos \theta(s) \kappa_2^{\text{ref}}(s). & (4.15c) \end{cases}$$

Examples of reference frames The *Frenet frame* is a possible reference frame, which admits the Darboux vector $\mathbf{\Omega}^{\text{ref}}(s) = \tau_g(s) \mathbf{n}_0^f(s) + \kappa_g(s) \mathbf{n}_2^f(s)$ and the material curvature vector $\mathbf{\kappa}^{\text{ref}}(s) = (\tau_g(s), 0, \kappa_g(s))$, with $\mathbf{n}_2^f(s)$ the binormal of the curve $\Gamma(s)$, and $\kappa_g(s)$ and $\tau_g(s)$ the geometric curvature and torsion, respectively. From (4.15) we retrieve the equations

$$\begin{cases} \kappa_0(s) = \tau_g(s) + \theta'(s) & (4.16a) \\ \kappa_1(s) = \sin \theta(s) \kappa_g(s) & (4.16b) \\ \kappa_2(s) = \cos \theta(s) \kappa_g(s). & (4.16c) \end{cases}$$

which were already derived directly in [Love, 1927, Section 253].

Another popular reference frame is the *Bishop frame*, which has the particularity to have a vanishing material twist. It has for Darboux vector $\mathbf{\Omega}^{\text{ref}}(s) = \kappa_g(s) \mathbf{n}_2^b(s)$ and for material curvature vector $\mathbf{\kappa}^{\text{ref}}(s) = (0, \kappa_1^b(s), \kappa_2^b(s))$ with $\kappa_1^b(s)$ and $\kappa_2^b(s)$ two functions satisfying $(\kappa_1^b(s))^2 + (\kappa_2^b(s))^2 = \kappa_g^2(s)$.

When $\mathcal{R}^{\text{ref}}(s)$ is the Bishop frame, System (4.15) boils down to

$$\begin{cases} \kappa_0(s) = \theta'(s) & (4.17a) \\ \kappa_1(s) = \cos \theta(s) \kappa_1^b(s) + \sin \theta(s) \kappa_2^b(s) & (4.17b) \\ \kappa_2(s) = -\sin \theta(s) \kappa_1^b(s) + \cos \theta(s) \kappa_2^b(s). & (4.17c) \end{cases}$$

In particular, we find that the material twist $\kappa_0(s)$ of the material frame $\mathcal{R}(s)$ is given by the derivative of the angle $\theta(s)$. This nice property made the curve-angle (Γ, θ) parametrization of a material rod from the Bishop frame particularly appealing [Langer and Singer, 1996] and is now preferred by many authors to the classical Euler angles parametrization. This reduced curve-angle parametrization was in particular leveraged in Computer Graphics for designing the so-called *discrete elastic rod* model, for which primary variables are discrete node positions of the centerline [Bergou et al., 2008].

4.4.3 Direct problem: framing a Kirchhoff rod

Let us now state the condition on $\mathcal{R}(s)$ guaranteeing it is a suitable frame for a Kirchhoff rod whose centerline exactly coincides with $\Gamma(s)$. As before, we assume the Kirchhoff rod to be clamped at $s = 0$, and free at $s = L$. We also suppose that all its parameters $\bar{\mathbf{\kappa}}(s)$, \mathbb{K} , and ρS are known, and that the rod is only subject to a known external force such as gravity. Then from the curve $\Gamma(s)$, it is possible to recover an admissible material frame (and thus admissible material twist and curvatures) for the corresponding Kirchhoff rod, according to the following property,

Property 4.1. $\{\Gamma(s), \mathcal{R}(s)\}$ is an admissible Kirchhoff rod configuration, i.e., it is compatible with dynamic Kirchhoff equations, if and only if one of the two equivalent conditions is satisfied,

(i) $\theta(s)$ is solution of the axial projection of the angular momentum equation,

$$\mathbf{M}'(s) \cdot \mathbf{n}_0(s) = 0 \quad (4.18)$$

subject to the boundary conditions (2.6b–2.6d), meaning that $\theta(s)$ is solution of the second-order differential equation

$$\theta''(s) + A(s) \cos^2 \theta(s) + B(s) \cos \theta(s) \sin \theta(s) + C(s) \cos \theta(s) + D(s) \sin \theta(s) + E(s) = 0$$

with

$$\begin{aligned} A(s) &= 2(\mathcal{K}_2 - \mathcal{K}_1) \kappa_1^{ref}(s) \kappa_2^{ref}(s) \\ B(s) &= (\mathcal{K}_2 - \mathcal{K}_1) ((\kappa_2^{ref}(s))^2 - (\kappa_1^{ref}(s))^2) \\ C(s) &= \mathcal{K}_1 \kappa_2^{ref}(s) \bar{\kappa}_1(s) - \mathcal{K}_2 \kappa_1^{ref}(s) \bar{\kappa}_2(s) \\ D(s) &= - \left(\mathcal{K}_1 \kappa_1^{ref}(s) \bar{\kappa}_1(s) + \mathcal{K}_2 \kappa_2^{ref}(s) \bar{\kappa}_2(s) \right) \\ E(s) &= - (\mathcal{K}_2 - \mathcal{K}_1) \kappa_1^{ref}(s) \kappa_2^{ref}(s) + \mathcal{K}_0 \left((\kappa_0^{ref})'(s) - \bar{\kappa}_0'(s) \right), \end{aligned}$$

subject to the boundary conditions (2.6).

(ii) $\theta(s)$ is a stationary point of the potential elastic energy of the rod,

$$\mathcal{E}_{el} = \int_0^L (\boldsymbol{\kappa}(s) - \bar{\boldsymbol{\kappa}}(s))^\top \mathbb{K}_3 (\boldsymbol{\kappa}(s) - \bar{\boldsymbol{\kappa}}(s)) \, ds. \quad (4.19)$$

Moreover, in the particular case of an isotropic rod ($\mathcal{K}_1 = \mathcal{K}_2$) with vanishing intrinsic curvatures and twist, then the material frame is necessarily the Bishop frame.

Proof. We only sketch here the main lines of the proof. Conditions on the material frame are given by the Kirchhoff angular momentum equation (2.4b), which can be expressed in the local frame $\mathcal{R}(s)$ as

$$[\boldsymbol{\kappa}(s)]_\times \mathbb{K}_3 (\boldsymbol{\kappa}(s) - \bar{\boldsymbol{\kappa}}(s)) + \mathbb{K}_3 (\boldsymbol{\kappa}'(s) - \bar{\boldsymbol{\kappa}}'(s)) + \mathbf{e}_x \times \mathbf{T}_{loc}(s) = 0 \quad (4.20)$$

Projecting (4.20) onto each vector of $\mathcal{R}(s)$ leads to

$$\begin{cases} \kappa_1 \kappa_2 (\mathcal{K}_2 - \mathcal{K}_1) + \mathcal{K}_1 \kappa_2 \bar{\kappa}_1 - \mathcal{K}_2 \kappa_1 \bar{\kappa}_2 + \mathcal{K}_0 (\kappa_0' - \bar{\kappa}_0') & = 0 & (4.21a) \\ \kappa_0 \kappa_2 (\mathcal{K}_0 - \mathcal{K}_2) + \mathcal{K}_2 \kappa_0 \bar{\kappa}_2 - \mathcal{K}_0 \kappa_2 \bar{\kappa}_0 + \mathcal{K}_1 (\kappa_1' - \bar{\kappa}_1') - T_{loc,2} & = 0 & (4.21b) \\ \kappa_0 \kappa_1 (\mathcal{K}_1 - \mathcal{K}_0) + \mathcal{K}_0 \kappa_1 \bar{\kappa}_0 - \mathcal{K}_1 \kappa_0 \bar{\kappa}_1 + \mathcal{K}_2 (\kappa_2' - \bar{\kappa}_2') + T_{loc,1} & = 0 & (4.21c) \end{cases}$$

Equation (4.21a) represents the axial projection of the Kirchhoff angular momentum conservation (4.18) expressed in the material frame. Note that this equation is

independent of the tension $\mathbf{T}(s)$, and thus of the acceleration of the centerline. It can actually be reformulated function of the single angular variable $\theta(s)$, all other quantities depending only of the centerline $\mathbf{r}(s)$ (thus being known). This can be simply achieved by replacing the curvatures with their angle parameterization given by formulas (4.15), leading to the second-order nonlinear differential equation (4.1) with boundary conditions (2.6) for $\theta(s)$.

Then, we consider the total potential energy of the rod and examine how it gets perturbed as the material frame is perturbed through an infinitesimal rotation $\delta\varphi(s)$ about the tangent vector to the centerline $\Gamma(s)$. The difference with the general analysis in [Audoly and Pomeau, 2010, Section 3.6] is that we restrict here $\delta\varphi(s)$ to be a cross-sectional perturbation, as the tangent vector of the material frame should not be perturbed since the centerline is fixed. Thus, $\delta\varphi(s)$ may be written as $\delta\varphi(s) = \delta\theta(s) \mathbf{n}_0(s)$, which represents an infinitesimal rotation of angle $\delta\theta(s)$ about the tangent vector $\mathbf{n}_0(s)$. One may then show that for a rod with free end, the perturbed potential energy boils down to the perturbed potential elastic energy, and also that it vanishes for any small perturbation $\delta\varphi(s)$ if and only if Equation (4.18) subject to the boundary conditions (2.6b—2.6d) is satisfied.

Finally, by choosing the reference frame to be the Bishop frame, one may easily show that when $\mathcal{K}_1 = \mathcal{K}_2$ and $\bar{\boldsymbol{\kappa}} = 0$, then Equation (4.1) boils down to $\theta'(s) = 0$, meaning that the unique set of admissible material frames is given by the Bishop frame. \square

This *quasistatic frame* property, and in particular the interpretation of the material frame as a minimizer of potential elastic energy, has already been leveraged in [Bergou et al., 2008] for building the discrete elastic rod model (see Section 4.4.5). To our knowledge, the optimality condition (4.1) has however not been explicitly derived before, furthermore with respect to a generic angular function $\theta(s)$.

Note that a material frame retrieved by Property 4.1 may be that of a Kirchhoff rod under motion, i.e., with a non-vanishing linear acceleration. Moreover, if the potential elastic energy of the rod is non convex, there may exist several admissible material frames, yielding different dynamical states for the rod. However, the constant feature is that whatever the dynamic state of the rod is, the material frame always obeys a quasistatic equation, due to the Kirchhoff assumption neglecting torsional inertia. Finally, it is interesting to note that Equation (4.1) takes some familiar forms in special cases. For instance, when the rod is isotropic with a piecewise helical centerline and uniform intrinsic curvatures and twist $\bar{\boldsymbol{\kappa}}$, then $\theta(s)$ satisfies the equation of the *Euler elastica*. We shall further study these special cases in future work.

4.4.4 Inverse static design problem

We now assume that the intrinsic curvatures and twist $\bar{\boldsymbol{\kappa}}(s)$ are unknown too (other parameters \mathbb{K} and ρS being still fixed). From the input curve $\Gamma(s)$, we aim at finding conditions on $\bar{\boldsymbol{\kappa}}(s)$ and the material frame $\mathcal{R}(s)$ so that $\Gamma(s)$ coincides with

the centerline of a Kirchhoff rod *at equilibrium* under gravity, framed by $\mathcal{R}(s)$. We have the following property, which, to the best of our knowledge, has not been formulated before,

Property 4.2. *For any choice of function $\theta(s)$ (i.e., any choice of the material frame $\mathcal{R}(s)$), Kirchhoff's equilibrium equations are satisfied if and only if $\bar{\boldsymbol{\kappa}}(s)$ satisfies the explicit linear ODE of first order,*

$$\bar{\boldsymbol{\kappa}}'(s) + \mathbb{K}_3^{-1} [\boldsymbol{\kappa}(s)]_{\times} \mathbb{K}_3 \bar{\boldsymbol{\kappa}}(s) = \boldsymbol{\kappa}'(s) + \mathbb{K}_3^{-1} b(s), \quad (4.22)$$

with

$$b(s) = \begin{bmatrix} \kappa_1(s) \kappa_2(s) (\mathcal{K}_2 - \mathcal{K}_1) \\ \kappa_0(s) \kappa_2(s) (\mathcal{K}_0 - \mathcal{K}_2) - T_{loc,2} \\ \kappa_0(s) \kappa_1(s) (\mathcal{K}_1 - \mathcal{K}_0) + T_{loc,1} \end{bmatrix},$$

where $\boldsymbol{\kappa}(s)$ is a function of $\theta(s)$ given by (4.14), $\boldsymbol{\kappa}(s) = \mathcal{R}_{-\theta}(s) (\boldsymbol{\kappa}^{ref}(s) + \theta'(s) \mathbf{e}_x)$ and the tension $\mathbf{T}(s) = \mathcal{R}(s) \mathbf{T}_{loc}(s)$ satisfies the linear momentum equation (2.4a) at equilibrium,

$$\mathbf{T}'(s) - \rho S g \mathbf{e}_z = 0, \quad \text{i.e.,} \quad T(s) = - \int_s^L \rho S g \mathbf{e}_z = \rho S g (s - L) \mathbf{e}_z.$$

Moreover, we have the boundary condition

$$\bar{\boldsymbol{\kappa}}(L) = \boldsymbol{\kappa}(L). \quad (4.23)$$

Equations (4.24) and (4.23) form a Cauchy problem, which admits a unique solution.

Proof. Proof follows from Equations (4.21), which are reorganized function of the new unknown $\bar{\boldsymbol{\kappa}}(s) = \{\bar{\kappa}_0(s), \bar{\kappa}_1(s), \bar{\kappa}_2(s)\}$ as

$$- \mathbb{K}_3 \bar{\boldsymbol{\kappa}}'(s) + A(s) \bar{\boldsymbol{\kappa}}(s) = - \mathbb{K}_3 \boldsymbol{\kappa}'(s) - b(s), \quad (4.24)$$

with

$$A(s) = \begin{bmatrix} 0 & \mathcal{K}_1 \kappa_2(s) & -\mathcal{K}_2 \kappa_1(s) \\ -\mathcal{K}_0 \kappa_2(s) & 0 & \mathcal{K}_2 \kappa_0(s) \\ \mathcal{K}_0 \kappa_1(s) & -\mathcal{K}_1 \kappa_0(s) & 0 \end{bmatrix} = -[\boldsymbol{\kappa}(s)]_{\times} \mathbb{K}_3,$$

and then multiplied by $-\mathbb{K}_3^{-1}$ and simplified. \square

The above property allows us to characterize the set of solutions $\bar{\boldsymbol{\kappa}}(s)$: it is the set of (unique) solutions to the Cauchy problem (4.24 – 4.23) parameterized by the function $\theta(s)$. At this stage however, one may feel that it is not possible to identify the intrinsic shape of the rod from the sole data of a curve $\Gamma(s)$. But actually, after the defense of this habilitation, we have proved that although the set of admissible natural curvatures $\bar{\boldsymbol{\kappa}}(s)$ is infinite, the intrinsic shape of the curve $\bar{\Gamma}(s)$ is *unique*, and can be computed easily, starting from *any* framing of the input curve $\Gamma(s)$ [Bertails-Descoubes et al., 2018].

4.4.5 Discrete picture

Having in mind the continuous picture for framing a Kirchhoff rod, it is now interesting to get back to discrete Kirchhoff models. We especially study the framing of two models built upon very different concepts: the discrete elastic rod model [Bergou et al., 2008, 2010], and the super-helix model.

Framing a discrete elastic rod The quasistatic frame property 4.1 was a core motivation for building the discrete elastic rod model [Bergou et al., 2008]. In this model, discrete positions on the centerline are the main degrees of freedom of the dynamics, and discrete angles θ_i , determining the material frame from the Bishop frame, are computed at each time step so as to minimize the potential elastic energy of the rod. Note that in [Bergou et al., 2010], the θ_i may be free degrees of freedom in the case where torsional inertia is considered. In that case the quasistatic frame property is satisfied at equilibrium only (in the limit of a vanishing torsional inertia, the property is fully recovered). For the discrete elastic rod model, Property 4.1 is thus satisfied by construction, at any resolution. Moreover, as predicted by Property 4.2, an input curve might be arbitrarily framed, and a set of intrinsic curvatures (that should all be equivalent, given our latest results [Bertails-Descoubes et al., 2018]), can be found to satisfy the equilibrium of the corresponding rod.

Framing a super-helix In contrast, the super-helix relies on a strong assumption regarding the shape of its elements, as it imposes a piecewise uniform shape function for the material curvatures and twist. This assumption has a strong impact on the shape of the material frame. Indeed, from (4.16) we get that on each element, $\theta(s)$ should be a *uniform* function, and as a consequence $\tau_g(s)$ and $\kappa_g(s)$ should also be uniform. We then retrieve the fact that a super-helix element represents a circular helix (since its Frenet curvature and torsion are uniform), plus the fact that the angle between the material frame and the Frenet frame is uniform. The latter condition implies that the two Darboux vectors are equal (but not the frames), that is,

$$\Omega^{\text{SH}}(s) = \Omega^f(s) \quad \forall s \in [0, L]. \quad (4.25)$$

We note that the $\theta(s) = C^{te}$ condition breaks the continuous law (4.1) written in the case of a helical centerline. The super-helix model thus does not satisfy the quasistatic frame condition exactly, but only at the limit when the number of elements tends to infinity.

When dealing with inverse design, once a piecewise helix has been reconstructed by some geometric approximation (see Section 4.2.4), then no much choice is left for the material frame: it may start arbitrarily, but then should be transported on each element like the Frenet frame, while remaining continuous at joints (thus shifting from the Frenet frame at each joint). Interestingly, the choice of an arbitrary shape function for the material strains thus acts as a selecting process for the set of admissible material frames. It would be interesting to investigate the case of super-clothoids, which in turn impose a piecewise linear shape function for material

strains. Such an assumption is often compatible with real fibers for which the growing process remains quite smooth, and would thus pre-filter admissible material frames in a more realistic way.

4.5 Conclusion: Towards non-invasive identification

In this chapter we have introduced the inverse static design problem for fibers, and have proposed a full pipeline for finding some solutions in the case of the super-helix model, possibly subject to frictional contact. We have also brought some new insights on the full space of solutions, both in the discrete and continuous settings.

Since our first work on fiber inverse design [Derouet-Jourdan et al., 2010], many other authors in Computer Graphics got interested in the problem, usually for more general deformable objects discretized with low-order finite elements [Twigg and Kačić-Alesić, 2011; Skouras et al., 2012; Zhao and Barbič, 2013; Chen et al., 2014; Pérez et al., 2015]. This growing interest has been especially boosted by the advent of additive fabrication, making it possible to fabricate the models with the predicted properties, and then check for their validity in practice.

However, the latter methods, heavily relying upon nonlinear optimization, are meant to converge to one particular solution of the inverse problem. As we have shown here, there can be many ways to explain a given static configuration. In contrast to concurrent methods, we would like to better characterize the exact space of solutions of the problem, in order to understand the precise mechanical information which can be extracted from the geometry of objects. Such a “non-blind” methodology has already proven useful for computing solutions to the inverse design problem efficiently, without having to resort to nonlinear optimization. In the longer run, it would also serve to control the exact amount of input information needed (e.g., a minimal number of poses) for dealing with an accurate *identification* process. This dream for designing a merely *non-invasive* identification pipeline, that is, recovering all physical parameters of an object from the observation of its shape only, is the very topic of my research program, which is provided in next chapter.

Research Perspectives: From Geometry to Mechanics

With the considerable advance of automatic image-based capture in Computer Vision and Computer Graphics these latest years, it becomes now affordable to acquire quickly and precisely the full 3D geometry of many mechanical objects featuring intricate shapes. Yet, while more and more geometrical data get collected and shared among the communities, there are currently very few studies about how to infer the underlying mechanical properties of the captured objects merely from their geometrical configurations.

An important challenge consists in developing a *non-invasive* method for inferring the mechanical properties of complex objects from a minimal set of geometrical poses, in order to *predict* their dynamics. In contrast to classical inverse reconstruction methods, my research project is built upon the claim that 1/ the mere geometrical shape of physical objects reveals a lot about their underlying mechanical properties and 2/ this property can be fully leveraged for a wide range of objects featuring rich geometrical configurations, such as slender structures subject to contact and friction (e.g., folded cloth or twined filaments).

To achieve this goal, we shall develop an original inverse modeling strategy based upon a/ the design of reduced and high-order discrete models for slender mechanical structures including rods, plates and shells, b/ a compact and well-posed mathematical formulation of our nonsmooth inverse problems, both in the static and dynamic cases, c/ the design of robust and efficient numerical tools for solving such complex problems, and d/ a thorough experimental validation of our methods relying on the most recent capturing tools.

In addition to significant advances in fast image-based measurement of diverse mechanical materials stemming from physics, biology, or manufacturing, this research is expected in the long run to ease considerably the design of physically realistic virtual worlds, as well as to boost the creation of dynamic human doubles.

5.1 From geometry to mechanics, a broken pipe

A profusion of shapes acquired from physical objects With the considerable advance of automatic image-based capture in Computer Vision and Computer Graphics these latest years, it becomes now affordable to acquire quickly and precisely the full 3D geometry of many mechanical objects featuring intricate shapes such as cloth and skin [Miguel et al., 2012], or even hair fibers [Luo et al., 2013].




Figure 5.1 – Recent acquisition technologies make it now affordable to reconstruct the 3D geometry of many complex objects featuring intricate shapes, such as (a) cloth [Bradley et al., 2008], (b) skin [Beeler et al., 2011], (c) hair [Luo et al., 2013], and (d) trees [Livny et al., 2010]. However, the consistent connection of such a geometrical description to the mechanical characterization of the underlying object remains an open issue.

Acquisition technologies range from expensive structured light or laser scans to new cheap devices such as depth cameras [Chang et al., 2012] which are often sufficient for capturing a static pose precisely.

Yet, while more and more geometrical data is collected and shared among the communities, there is currently very little study about how to infer the underlying mechanical properties of the captured objects merely from their geometrical configurations. One can however suspect that the pure static shape of a physical object may already give some insights about the constitutive material of the object and the interplaying contacts: from the folding patterns of a tablecloth or a curtain, the human eye may perceive whether the fabric is made of rough cotton or silk, and identify zones of contacts. One may then have the dream that feeding a well-designed physics-based simulator with such easy-available initial data could help predict the deformations or even the dynamics of the physical objects of interest.

Material tests for measuring physical parameters In parallel, contactless measurement methods, which reconstruct full-displacement fields based on camera capture and digital image correlation, have recently gained much interest in Experimental Mechanics [Avril et al., 2008]. Indeed, unlike sensor-based capture, image-based capture does not interfere with the displacement field being measured. Combined with FEM-based inverse modeling, contactless measurement methods allow for a complete parameter identification of complex materials. They however request that a number of specific material tests (e.g., tensile and shear tests) be performed, which may often require some expensive material and time-costly measurement protocols, and sometimes may even be impracticable when objects are not directly manipulable. Moreover, although some recent developments in Computer Graphics have extended the range of studies from small to moderate 3D deformations [Miguel et al., 2012] and partly lightened the necessary amount of control in the experimental setup [Wang et al., 2011; Miguel et al., 2013], such methods remain

limited to the study of *contact-free* objects.

In contrast, using geometrical acquisitions of a minimum number of *uncontrolled* static poses would release the burden of material testing and provide a breakthrough in the non-invasive and fast measurement of many mechanical features including rest configuration, material parameters, self-contacting forces, and friction coefficients at contact.

Inaccurate reconstruction of motion Besides, while geometry acquisition from a static pose has become a mature technology, a current major challenge in Computer Graphics and Vision deals with the accurate reconstruction of motion. More precisely, the *non-rigid shape registration*, which aims at finding consistent correspondences between successive acquisitions of the moving object and resolving ambiguities or occlusions, is still an active area of research [Chang et al., 2012]. Physical predictions of the motion of the physical object could be of great help to resolve such inconsistencies. Conversely, the recovery of accurate dynamic geometrical poses would enrich mechanical interpretation and greatly help refine the identification of all physical parameters at work.

Giving a physical meaning to mere geometrical data would not only serve as an innovative parameters measurement method, but also as a powerful strategy to convert a purely descriptive approach into a generative one, able to predict an infinite number of new and rich dynamic scenarios.

Some sparse connection tentatives A few works have attempted to identify the physical parameters of slender objects such as fibers, cloth or skin, from an arbitrary shape at static equilibrium [Bhat et al., 2003; Twigg and Kačić-Alesić, 2011; Bickel et al., 2012]. However, because these approaches rely upon nodal mechanical models whose rest shape has to conform to some constraints (e.g., spring rest length should be nonnegative), they cannot guarantee that a static equilibrium configuration will match the input shape exactly. Moreover, as the functions to minimize are non convex and of large size, the cost of problem solving tends to get prohibitive, precluding an accurate treatment of contact and friction. Finally, such a straight numerical treatment provides only a single solution to the (under-determined) inverse problem and does not give insights about the general structure of the subspace of solutions.

Even sparser, inverse dynamic studies greatly suffer from the limited quality of the 3D reconstruction of moving geometry. To complete the geometrical reconstruction consistently, physically-guided approaches have started to be explored. So far however, they have only proven to be successful in the case of smooth dynamic systems such as Newtonian fluids [Wang et al., 2009]. Recently, Bouman et al. [2013] have proposed to skip the 3D reconstruction step and to use instead statistics characterizing temporal textures in order to predict the material properties of fabric from gentle motion. However, their method is not tailored for making a perfect match between the real phenomenon and a (predictive) physic-based model of it,

nor for retrieving contact forces and friction coefficients at play.

To the best of my knowledge, our preliminary work on the inversion of fiber assemblies was the first to take into account the presence of dry frictional contact for (static) inversion [Derouet-Jourdan et al., 2013a]. We however made some important assumptions, such as the *a priori* knowledge of the material parameters and an *a priori* warm start estimation of the rest shape.

5.2 Interpreting geometry as a mechanical state

The key challenge of my research project is the automatic connection between the *geometrical shape* of physical objects and their underlying *mechanical properties*. More precisely, I intend to focus this study on complex deformable objects featuring detailed geometrical configurations. Typical objects of interest include slender deformable structures such as rods, plates and shells, all of them being widespread in our environment, from the macroscopic scale (e.g., tree branches and leaves, hair, cloth, skin, paper) to nanoscopic and molecular scales (e.g., carbon nanotubes, DNA). Such structures, which are prone to strongly *nonlinear* behaviors as well as to possibly prominent *self-contacting* causing knots, plectonemes, or folds, exhibit very rich geometrical configurations.

My claim in this research project, supported by the preliminary results we have gained on hair fibers after years of research, is that these complex geometrical features *reveal a lot about the underlying mechanical structures*.

5.3 Scientific problems

To be able to extract such mechanical properties from a minimum set of geometrical shapes, my goal is to develop *specific effective computational models* addressing four major scientific problems:

SP1: Design of well-suited discrete models for slender structures I believe that the quality of the upstream, reference physics-based model is essential to the effective connection between geometry and mechanics. Typically, such a model should properly account for the nonlinearities due to large displacements of the structures. It should also be parameterized and discretized in such a way that inversion gets simplified mathematically, possibly avoiding the huge cost of large and nonconvex optimization. In that sense, unlike concurrent methods which impose inverse methods to be compatible with a generic physics-based model, I instead advocate the design of specific physics-based models which are *tailored* for the inversion process.

More precisely, from our experience on fiber modeling, I believe that *reduced Lagrangian models*, based on a minimal set of coordinates and physical parameters (as opposed to maximal coordinates models such as mass-springs), are particularly well-suited for inversion and physical interpretation of geometrical data. Furthermore, choosing a *high-order* coordinate system (e.g., curvatures instead of angles)

allows for a precise handling of curved boundaries and contact geometry, as well as the simplification of constitutive laws (which are transformed into a linear equation in the case of rods).

SP2: Static inversion of physical objects from geometrical poses I believe that pure *static inversion* may by itself reveal many insights regarding a range of parameters such as the undeformed configuration of the object, some material parameters or contact forces.

The typical settings that I will consider will be composed of, on the one hand, a reference mechanical model of the object of interest (**SP1**), and on the other hand a single or a series of *complete* geometrical poses corresponding each to a *static equilibrium* under gravitational load. The core challenge will consist in analysing theoretically the amount of information that can be gained from one or several geometrical poses, and to understand how the fundamental under-determinacy of the inverse problem can be reduced, for each unknown quantity (parameter or force) at play. Both the *equilibrium* condition and the *stability* criterion of the equilibrium will be leveraged towards this goal.

SP3: Dynamic inversion of physical objects from geometrical poses To refine the solution subspaces searched for in **SP2** and estimate dynamic parameters (e.g., damping coefficients), a *dynamic inversion* process accounting for the motion of the object of interest is necessary.

In contrast to the static case **SP2** where we can afford to rely on exact geometrical poses, our analysis in the dynamic case will have to take into account the imperfect quality of input data with possible missing parts or outliers. One interesting challenge will be to combine our physics-based model from **SP1** together with the acquisition process in order to refine both the parameter estimation and the geometrical acquisition.

SP4: Experimental validation with respect to real data The goal will be to confront the theories developed in **SP2** and **SP3** to real experiments. Compared to the statics, the dynamic case will be particularly involving as it will be highly dependent on the quality of input data as well as the accuracy of the motion predicted by the physics-based model designed in **SP1**. Such experiments will not only serve to refine the models developed in **SP1**, **SP2** and **SP3**, but will also be used to improve the 3D geometrical acquisition of moving objects. Besides, once validation will be performed, we shall work on the setting up of *new non-invasive measurement protocols* to acquire physical parameters of slender structures from a minimal amount of geometrical configurations.

5.4 Expected Impact

Converting geometry into physics promises a considerable amount of scientific impact and potential applications for a number of fields ranging from Computer Graphics to Physics and Biology. Thanks to the funding of this research project from the *European Research Council* (grant started in September 2015), I intend to create and lead a group internationally active on this pluridisciplinary topic.

Scientific impact Fundamentally, this research project may considerably enhance our understanding of how geometry is tightly linked to mechanics. From the mere observation of shapes through imaging, it will allow scientists to better identify and even predict the physical behavior of many physical and biological slender structures ranging from soft tissues to DNA. Regarding Computer Science, our project may significantly contribute to pave the way for robust 3D dynamic reconstruction of geometry for slender structures, a longstanding challenge in Computer Vision.

Applications of the future In the long run, I expect this project to contribute to new *performant and non-invasive measurement systems* for capturing automatically the physical properties of slender structures, from a few samples geometry only. Beyond huge applications in the measurement of human features such as hair, skin, cloth and even internal thin organs from 3D capture, one interesting application would be the automatic retrieval of physical objects properties from 2D images such as photographs or even paintings. Furthermore, the ability to *predict* the motion of objects from a mere static geometry is highly desirable not only in the artistic design of dynamic scenes for movies or games, but also in the dynamic virtual try-on industry of garments and hairstyles. Ultimately, our contributions could serve for the fast, image-based creation of complete *dynamic* digital human doubles, desirable not only for the virtual acting in movies but also for patient-specific surgery.

Conclusion

Now listen to the rule of the last inch. The realm of the last inch. The job is almost finished, the goal almost attained, everything possible seems to have been achieved, every difficulty overcome - and yet the quality is just not there. The work needs more finish, perhaps further research. In that moment of weariness and self-satisfaction, the temptation is greatest to give up, not to strive for the peak of quality. That's the realm of the last inch - here, the work is very, very complex, but it's also particularly valuable because it's done with the most perfect means. The rule of the last inch is simply this - not to leave it undone. And not to put it off - because otherwise your mind loses touch with that realm. And not to mind how much time you spend on it, because the aim is not to finish the job quickly, but to reach perfection.

– Alexander Solzhenitsyn, *The First Circle*

In this memoir we have presented our main ideas for designing realistic, robust and computationally efficient models for thin deformable objects and divided materials prone to frictional contact. The general philosophy we are tending to adopt is a pluridisciplinary approach, driven by applications but also inspired by theoretical concepts. To build a consistent and effective model, we strive to master the full modeling pipeline, from the continuous physical model to its discretization and implementation. From this transverse viewpoint, which led us to meet different scientific communities, we retain a number of clues and preferred paths for advancing our research work; these are listed below.

Upfront modeling Throughout this long-term work on the numerical modeling of fibers and frictional contact, we have learnt that systematically concentrating the efforts on the upstream modeling and formulation of problems often pays off: even for very complex problems, the resulting numerics may be greatly simplified and thus solved more easily and robustly. Keeping in mind this key lesson, we are starting to investigate the case of 2D slender structures (plates and shells), for which many exciting challenges remain open, regarding both direct and inverse modeling.

Small data for big understanding¹ My current research trajectory is clearly an attempt to build stronger connections with physics and physicists, with the aim to better understand and master the macroscopic modeling of complex natural phenomena; indeed, there is still a huge gap between our actual need for relevant macroscopic models, and the set of available models of the literature. However, in a time where big data is considerably expanding and seems to prevail over standard

¹First coined by Pascal Barla while talking together at Siggraph 2016.

modeling techniques, one may argue that such a tentative, definitely upstreams, is vain. Indeed, big data analysis spreads the word that models may not need to be designed by humans anymore, but rather spontaneously emerge from the automatic analysis and learning of terabytes of collected data, performed by supercomputers.

Yet, in this deep change of modeling strategies where statistical and computational approaches may want to take the lead, I believe that the physical modeling point of view is more than ever desirable. First of all, the human mind feels like understanding the world, and not just reproducing it blindly. Then, the idea of feeding computers with unlimited and inexpensive data is probably illusory, and unsustainable in certain fields like human biology where medical imaging cannot be overused on a given patient. In contrast, I believe that compact models, able to give some physical meaning and to provide some predictive tools from a minimal amount of input data, are still extremely valuable. They should also complement nicely statistical approaches which treat huge amounts of data, by providing them with some useful hints for interpreting classified objects.

Physics and Computer Science: a promising alliance? Finally, as pointed out in the introduction of this document, there are some striking common points between Nonlinear Physics and Computer Graphics, as some phenomena of interest are becoming more and more similar. Links between the two fields are however quite loose, with only a handful of experts in the world working at the frontier between the two fields. Such a disconnection can be explained by higher education, which usually clearly separates Physics (and Mechanics) from Computer Science, the latter being more willingly attached to Applied Mathematics. In France, such a segmentation is clear at university or in engineering schools: computer scientists will not tackle Natural Sciences during their study, and physicists will probably learn that Computer Science is just a programming tool, readily accessible through black boxes such as Mathematica or Abaqus. In research labs, this segmentation persists, although more and more efforts are made to establish connections between digital science and other fields. For instance, the last decades have seen the emergence of Computational Biology, an actual pluridisciplinary field where biologists, mathematicians and computer scientists have indeed combined their expertise to advance a modern vision of biology.

At my (much humbler) scale, I believe valuable connections could be established between Physics of Nonlinear Elasticity and Computer Animation, where scientists share the common goal to model complex and integrated mechanical phenomena at the macroscopic scale. Indeed, while computer scientists are eager to learn and understand new physical models, physicists get more and more interested in the numerical tools, in which they perceive not only a means to confirm predictions afterwards, but also a support for testing hypothesis and for getting insights into the search for analytic solutions, starting at the modeling stage. However, they may be limited by a blind usage of numerical black boxes, which may not be dedicated to their specific needs. According to me, promoting a science of modeling in *numerical physics* would thus be a promising and rich avenue for our two research fields.



Figure 6.1 – My daily activity, a constant search for the right balance. Between mechanics and graphics, between theory and applications, between paper work and programming, between supervision and learning, between publication and industrial transfer... and between work and private life.

Bibliography

- Acary, V. and Brogliato, B. (2008). *Numerical methods for nonsmooth dynamical systems*, volume 35 of *Lecture Notes in Computational and Applied Mechanics*. Springer. (Cited on pages 36, 40, 41, 42 and 46.)
- Acary, V., Cadoux, F., Lemarechal, C., and Malick, J. (2011). A formulation of the linear discrete Coulomb friction problem via convex optimization. *ZAMM / Z angew Math Mech; Zeitschrift für Angewandte Mathematik und Mechanik*, 91:155–175. (Cited on pages 45 and 48.)
- Alart, P. and Curnier, A. (1991). A mixed formulation for frictional contact problems prone to Newton like solution methods. *Comput. Methods Appl. Mech. Eng.*, 92(3):353–375. (Cited on page 43.)
- Alduán, I., Tena, A., and Otaduy, M. A. (2009). Simulation of high-resolution granular media. In *Proc. of Congreso Español de Informática Gráfica*, number 4 in 1. (Cited on page 49.)
- Andreotti, B., Forterre, Y., and Pouliquen, O. (2011). *Granular media: between fluid and solid*. Cambridge University Press. (Cited on page 48.)
- Anitescu, M. (2005). Optimization-based simulation of nonsmooth rigid multibody dynamics. *Math. Program.*, 105(1):113–143. (Cited on page 45.)
- Anjyo, K., Usami, Y., and Kurihara, T. (1992). A simple method for extracting the natural beauty of hair. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'92)*, pages 111–120. (Cited on page 31.)
- Antman, S. (1995). *Nonlinear Problems of Elasticity*. Springer Verlag. (Cited on page 13.)
- Audoly, B. and Neukirch, S. (2005). Fragmentation of rods by cascading cracks: Why spaghetti does not break in half. *Physical Review Letters*, 95(9):095505. (Cited on page 3.)
- Audoly, B. and Pomeau, Y. (2010). *Elasticity and Geometry: from hair curls to the nonlinear response of shells*. Oxford University Press. (Cited on pages 7, 13 and 76.)
- Avril, S., Bonnet, M., Bretelle, A.-S., Grédiac, M., Hild, F., Ienny, P., Latourte, F., Lemosse, D., Pagano, S., Pagnacco, E., and Pierron, F. (2008). Overview of identification methods of mechanical parameters based on full-field measurements. *Experimental Mechanics*, 48(4):381–402. (Cited on page 82.)

- Baraff, D. (1989). Analytical methods for dynamic simulation of non-penetrating rigid bodies. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'89)*, pages 223–232, New York, NY, USA. ACM. (Cited on pages 33 and 35.)
- Baraff, D. (1991). Coping with friction for non-penetrating rigid body simulation. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'91)*, pages 31–40. ACM. (Cited on pages 35 and 41.)
- Baraff, D. (1993). Issues in computing contact forces for non-penetrating rigid bodies. *Algorithmica*, 10:292–352. (Cited on page 35.)
- Baraff, D. (1994). Fast contact force computation for nonpenetrating rigid bodies. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'94)*, pages 23–34, New York, NY, USA. ACM. (Cited on page 35.)
- Baraff, D. and Witkin, A. (1998). Large steps in cloth simulation. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'98)*, pages 43–54. (Cited on pages 15 and 36.)
- Beck, J. and Woodbury, K. (1998). Inverse problems and parameter estimation: integration of measurements and analysis. *Measurement Science and Technology*, 9(6):839. (Cited on page 60.)
- Beeler, T., Hahn, F., Bradley, D., Bickel, B., Beardsley, P., Gotsman, C., Sumner, R. W., and Gross, M. (2011). High-quality passive facial performance capture using anchor frames. *ACM SIGGRAPH'11*. (Cited on page 82.)
- Ben Amar, M. and Pomeau, Y. (1997). Crumpled paper. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 453(1959):729–755. (Cited on page 3.)
- Benham, C. and Mielke, S. (2005). DNA mechanics. *Annual Review of Biomedical Engineering*, 7:21–53. (Cited on page 12.)
- Bergou, M., Audoly, B., Vouga, E., Wardetzky, M., and Grinspun, E. (2010). Discrete viscous threads. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'10)*, 29(4). (Cited on pages 14, 72 and 78.)
- Bergou, M., Wardetzky, M., Robinson, S., Audoly, B., and Grinspun, E. (2008). Discrete elastic rods. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'08)*, 27(3):1–12. (Cited on pages 14, 25, 72, 74, 76 and 78.)
- Bertails, F. (2009). Linear time super-helices. *Computer Graphics Forum (Proc. Eurographics'09)*, 28(2). (Cited on pages 9, 12 and 27.)
- Bertails, F., Audoly, B., Cani, M.-P., Querleux, B., Leroy, F., and Lévêque, J.-L. (2006). Super-helices for predicting the dynamics of natural hair. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'06)*, 25:1180–1187. (Cited on pages 9, 11, 36 and 58.)

- Bertails-Descoubes, F. (2012). Super-clothoids. *Computer Graphics Forum (Proc. Eurographics'12)*, 31(2pt2):509–518. (Cited on pages 9, 12, 58 and 65.)
- Bertails-Descoubes, F. (2017). About curves, frames, and material rods. To be submitted in 2017. (Cited on pages 58 and 72.)
- Bertails-Descoubes, F., Cadoux, F., Daviet, G., and Acary, V. (2011). A nonsmooth Newton solver for capturing exact Coulomb friction in fiber assemblies. *ACM Transactions on Graphics*, 30:6:1–6:14. (Cited on pages 32 and 44.)
- Bertails-Descoubes, F., Derouet-Jourdan, A., Romero, V., and Lazarus, A. (2018). Inverse design of an isotropic suspended kirchhoff rod: theoretical and numerical results on the uniqueness of the natural shape. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 474(2212). (Cited on pages 72, 77 and 78.)
- Beusmans, J. M. H. and Silk, W. K. (1988). Mechanical properties within the growth zone of corn roots investigated by bending experiments. ii. distributions of modulus and compliance in bending. *American Journal of Botany*, 75(7):996–1002. (Cited on page 17.)
- Bhat, K. S., Twigg, C. D., Hodgins, J. K., Khosla, P. K., Popović, Z., and Seitz, S. M. (2003). Estimating cloth simulation parameters from video. In *Proc. ACM SIGGRAPH/EG Symp. Comp. Animation, SCA '03*, pages 37–51. (Cited on page 83.)
- Bickel, B., Kaufmann, P., Skouras, M., Thomaszewski, B., Bradley, D., Beeler, T., Jackson, P., Marschner, S., Matusik, W., and Gross, M. (2012). Physical face cloning. *ACM SIGGRAPH'12*. (Cited on page 83.)
- Blumentals, A., Bertails-Descoubes, F., and Casati, R. (2016a). Dynamics of a developable shell with uniform curvatures. In *The 4th Joint International Conference on Multibody System Dynamics*, Montréal, Canada. (Cited on pages 9 and 29.)
- Blumentals, A., Brogliato, B., and Bertails-Descoubes, F. (2016b). The contact problem in Lagrangian systems subject to bilateral and unilateral constraints, with or without sliding coulomb's friction: a tutorial. *Multibody System Dynamics*, 38(1):43–76. (Cited on pages 9 and 41.)
- Bonnefon, O. and Daviet, G. (2011). Quartic formulation of Coulomb 3D frictional contact. Technical report, INRIA - Laboratoire Jean Kuntzmann. (Cited on pages 42 and 46.)
- Bouman, K. L., Xiao, B., Battaglia, P., and Freeman, W. T. (2013). Estimating the material properties of fabric from video. In *International Conference on Computer Vision (ICCV'13)*. (Cited on page 83.)

- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press. (Cited on page 70.)
- Bradley, D., Boubekur, T., and Heidrich, W. (2008). Accurate multiview reconstruction using robust binocular stereo and surface meshing. In *Computer Vision and Pattern Recognition (CVPR'08)*. (Cited on page 82.)
- Brun, P.-T., Ribe, N. M., and Audoly, B. (2012). A numerical investigation of the fluid mechanical sewing machine. *Physics of fluids*, 24(4):043102. (Cited on page 3.)
- Cadoux, F. (2009). *Méthodes d'optimisation pour la dynamique non-régulière*. PhD thesis, Université Joseph Fourier. (Cited on pages 40, 41, 45 and 48.)
- Casati, R. (2015). *Quelques contributions à la modélisation numérique de structures élancées pour l'informatique graphique*. PhD thesis, Grenoble Alpes Universités. (Cited on page 61.)
- Casati, R. and Bertails-Descoubes, F. (2013). Super space clothoids. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'13)*, 32(4):48:1–48:12. (Cited on pages 9, 12, 13, 21, 23, 25 and 58.)
- Casati, R., Daviet, G., and Bertails-Descoubes, F. (2016). Inverse elastic cloth design with contact and friction. Research report, Inria Grenoble Rhône-Alpes, Université de Grenoble. (Cited on page 61.)
- Chambon, G., Bouvarel, R., Laigle, D., and Naaim, M. (2011). Numerical simulations of granular free-surface flows using smoothed particle hydrodynamics. *Journal of Non-Newtonian Fluid Mechanics*, 166(12-13):698–712. (Cited on page 33.)
- Chang, W., Li, H., Mitra, N. J., Pauly, M., and Wand, M. (2012). Dynamic geometry processing. In *Eurographics 2012: Tutorial Notes*. (Cited on pages 82 and 83.)
- Chauchat, J. and Médale, M. (2014). A three-dimensional numerical model for dense granular flows based on the $\mu(i)$ rheology. *Journal of Computational Physics*, 256:696–712. (Cited on pages 51 and 53.)
- Chen, X., Zheng, C., Xu, W., and Zhou, K. (2014). An asymptotic numerical method for inverse elastic shape design. *ACM Trans. Graph.*, 33(4):95:1–95:11. (Cited on page 79.)
- Chentanez, N., Alterovitz, R., Ritchie, D., Cho, L., Hauser, K., Goldberg, K., Shewchuk, J., and O'Brien, J. (2009). Interactive simulation of surgical needle insertion and steering. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'09)*, pages 88:1–10. (Cited on page 12.)
- Cosserat, E. and Cosserat, F. (1909). *Théorie des corps déformables*. Hermann. (Cited on page 13.)

- Coulomb, C.-A. (1781). *Théorie des machines simples: en ayant égard au frottement de leurs parties et à la roideur des cordages*. Chapitre.com (new edition). (Cited on page 38.)
- Crisfield, M. A. and Jelenić, G. (1998). Objectivity of strain measures in the geometrically exact three-dimensional beam theory and its finite-element implementation. *Proc. Royal Society of London, Series A*, 455(1983):1125–1147. (Cited on page 14.)
- Cundall, P. (1971). A computer model for simulating progressive large scale movements of blocky rock systems. in proceedings of the symposium of the international society of rock mechanics. In *Proceedings of the Symposium of the International Society of Rock Mechanics*, volume 1, pages 132–150. (Cited on page 33.)
- Daviet, G. (2016). *Modèles et algorithmes pour la simulation du contact frottant dans les matériaux complexes : application aux milieux fibreux et granulaires*. PhD thesis, Grenoble Alpes Universités. (Cited on pages 38, 41 and 50.)
- Daviet, G. and Bertails-Descoubes, F. (2016a). A semi-implicit material point method for the continuum simulation of granular materials. *ACM Transactions on Graphics*, 35(4):13. (Cited on pages 8, 9, 33, 50 and 54.)
- Daviet, G. and Bertails-Descoubes, F. (2016b). Nonsmooth simulation of dense granular flows with pressure-dependent yield stress. *Journal of Non-Newtonian Fluid Mechanics*, 234:15–35. (Cited on pages 8, 9, 33, 50 and 52.)
- Daviet, G. and Bertails-Descoubes, F. (2017). Simulation of Drucker–Prager granular flows inside Newtonian fluids. working paper or preprint. (Cited on page 55.)
- Daviet, G., Bertails-Descoubes, F., and Boissieux, L. (2011). A hybrid iterative solver for robustly capturing Coulomb friction in hair dynamics. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia’11)*, 30:139:1–139:12. (Cited on pages 8, 9, 32, 33, 44, 46, 47 and 71.)
- Daviet, G., Bertails-Descoubes, F., and Casati, R. (2015). Fast Cloth Simulation with Implicit Contact and Exact Coulomb Friction. ACM SIGGRAPH / Eurographics Symposium on Computer Animation. Poster. (Cited on pages 9, 32 and 47.)
- De Saxcé, G. and Feng, Z.-Q. (1998). The bipotential method: a constructive approach to design the complete contact law with friction and improved numerical algorithms. *Math. Comput. Modelling*, 28(4-8):225–245. (Cited on pages 38 and 43.)
- Derouet-Jourdan, A. (2013). *Inversion statique de fibres: de la géométrie de courbes 3D à l’équilibre d’une assemblée de tiges mécaniques en contact frottant*. PhD thesis, Grenoble Alpes Universités. (Cited on page 66.)

- Derouet-Jourdan, A., Bertails-Descoubes, F., Daviet, G., and Thollot, J. (2013a). Inverse dynamic hair modeling with frictional contact. *ACM Trans. Graph.*, 32(6):159:1–159:10. (Cited on pages 8, 9, 58, 68 and 84.)
- Derouet-Jourdan, A., Bertails-Descoubes, F., and Thollot, J. (2010). Stable inverse dynamic curves. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia'10)*, 29:137:1–137:10. (Cited on pages 9, 58, 65 and 79.)
- Derouet-Jourdan, A., Bertails-Descoubes, F., and Thollot, J. (2011). 3D inverse dynamic modeling of strands. In Wexler, D., editor, *ACM SIGGRAPH 2011 Posters*, page Article No. 55, Vancouver, Canada. ACM SIGGRAPH, ACM. Poster. (Cited on pages 9 and 58.)
- Derouet-Jourdan, A., Bertails-Descoubes, F., and Thollot, J. (2013b). Floating tangents for approximating spatial curves with G^1 piecewise helices. *Computer Aided Geometric Design*, 30. (Cited on pages 9, 58 and 66.)
- Dill, E. (1992). Kirchhoff's theory of rods. *Archive for History of Exact Sciences*, 44(1):1–23. (Cited on page 13.)
- Dubois, F. and Jean, M. (2006). The non smooth contact dynamic method: recent LMGC90 software developments and application. In Wriggers, P. and Nackenhorst, U., editors, *Analysis and Simulation of Contact Problems*, volume 27 of *Lecture Notes in Applied and Computational Mechanics*, pages 375–378. Springer Berlin / Heidelberg. 10.1007/3-540-31761-9-44. (Cited on page 46.)
- Erleben, K. (2007). Velocity-based shock propagation for multibody dynamics animation. *ACM Transaction on Graphics*, 26(2). (Cited on page 42.)
- ESPCI (2016). Rencontre en l'honneur de Yves Pomeau, octobre 2016. <https://www.sfpnet.fr/rencontre-celebrant-la-medaille-boltzmann-d-yves-pomeau>. (Cited on page 3.)
- Featherstone, R. (1983). The calculation of robot dynamics using articulated-body inertias. *International Journal of Robotics Research*, 2(1):13–30. (Cited on pages 12 and 27.)
- Featherstone, R. (1999a). A divide-and-conquer articulated-body algorithm for parallel $o(\log(n))$ calculation of rigid-body dynamics. part 1: Basic algorithm. *The International Journal of Robotics Research*, 18(9):867–875. (Cited on page 28.)
- Featherstone, R. (1999b). A divide-and-conquer articulated-body algorithm for parallel $o(\log(n))$ calculation of rigid-body dynamics. part 2: Trees, loops, and accuracy. *The International Journal of Robotics Research*, 18(9):876–892. (Cited on page 28.)
- Frigaard, I. and Nouar, C. (2005). On the usage of viscosity regularisation methods for visco-plastic fluid flow computation. *Journal of Non-Newtonian Fluid Mechanics*, 127(1):1–26. (Cited on page 34.)

- Fukushima, M., Luo, Z.-Q., and Tseng, P. (2002). Smoothing functions for second-order-cone complementarity problems. *SIAM J. on Optimization*, 12:436–460. (Cited on pages 44 and 46.)
- Fulton, W. (2000). Eigenvalues, invariant factors, highest weights, and Schubert calculus. *Bull. Amer. Math. Soc.*, 37:209–249. (Cited on page 63.)
- Ghosh, S. (2010). *Geometric approximation of curves and singularities of secant maps. A differential geometric approach*. PhD thesis, University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science. (Cited on page 66.)
- Goldberg, D. (1991). What every computer scientist should know about floating-point arithmetic. *ACM Comp. Surveys*, 23:5–48. (Cited on page 22.)
- Goldstein, R. E., Warren, P. B., and Ball, R. C. (2012). Shape of a ponytail and the statistical physics of hair fiber bundles. *Phys. Rev. Lett.*, 108:078101. (Cited on page 3.)
- Goriely, A. and Neukirch, S. (2006). Mechanics of climbing and attachment in twining plants. *Physical Review Letters*, 97:184302. (Cited on pages 3 and 12.)
- Goyal, S., Perkins, N., and Lee, C. (2008). Non-linear dynamic intertwining of rods with self-contact. *International Journal of Non-Linear Mechanics*, 43(1):65 – 73. (Cited on pages 12 and 14.)
- Grinspun, E., Hirani, A., Desbrun, M., and Schröder, P. (2003). Discrete Shells. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'03)*, pages 62–67. ACM-EG SCA. (Cited on page 36.)
- Hadap, S. (2003). *Hair Simulation*. PhD thesis, Université de Genève. (Cited on page 28.)
- Hadap, S. (2006). Oriented strands - dynamics of stiff multi-body system. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'06)*, pages 91–100. ACM-EG SCA. (Cited on pages 14 and 28.)
- Hadap, S. and Magnenat-Thalmann, N. (2001). Modeling dynamic hair as a continuum. *Computer Graphics Forum (Proc. Eurographics'01)*, 20(3):329–338. (Cited on page 14.)
- Harary, G. and Tal, A. (2012). 3D Euler spirals for 3D shape completion. *Computational Geometry*, 45(3):115–126. (Cited on pages 20 and 65.)
- Haslinger, J. (1983). Approximation of the Signorini problem with friction, obeying the coulomb law. *Mathematical Methods in the Applied Sciences*, 5(1):422–437. (Cited on page 45.)

- Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). *Fundamentals of convex analysis*. Grundlehren Text Editions. Springer. (Cited on page 44.)
- Ivanova, E. (2000). On one approach to solving the Darboux problem. *Mechanics of Solids*, 35:36–43. (Cited on page 16.)
- Jean, M. (1999). The non smooth contact dynamics method. *Computer Methods in Applied Mechanics and Engineering*, 177:235–257. Special issue on computational modeling of contact and friction, J.A.C. Martins and A. Klarbring, editors. (Cited on page 31.)
- Jop, P., Forterre, Y., and Pouliquen, O. (2006). A constitutive law for dense granular flows. *Nature*, 441(7094):727–730. (Cited on pages 49 and 51.)
- Kaufman, D., Edmunds, T., and Pai, D. (2005). Fast frictional dynamics for rigid bodies. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'05)*, 24(3):946–956. (Cited on page 35.)
- Kaufman, D., Sueda, S., James, D., and Pai, D. (2008). Staggered projections for frictional contact in multibody systems. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia'08)*, 27(5):164:1–164:11. (Cited on pages 35 and 42.)
- Kaufman, D., Tamstorf, R., Smith, B., Aubry, J.-M., and Grinspun, E. (2014). Adaptive nonlinearity for collisions in complex rod assemblies. *ACM Trans. on Graphics (SIGGRAPH 2014)*. (Cited on page 46.)
- Kern, M. (2002). *Problèmes inverses*. (Cited on page 60.)
- Klarbring, A. (1987). Contact problems with friction by linear complementarity. *Unilateral Problems in Structural Analysis - 2*, pages 197–219. (Cited on page 42.)
- Langer, J. and Singer, D. A. (1996). Lagrangian aspects of the Kirchhoff elastic rod. *SIAM Review*, 38(4):605–618. (Cited on page 74.)
- Livny, Y., Yan, F., Olson, M., Chen, B., Zhang, H., and El-Sana, J. (2010). Automatic reconstruction of tree skeletal structures from point clouds. *ACM SIGGRAPH Asia'10*. (Cited on page 82.)
- Love, A. (1927). *A Treatise on the Mathematical Theory of Elasticity*. Cambridge University Press. (Cited on page 74.)
- Luo, L., Li, H., and Rusinkiewicz, S. (2013). Structure-aware hair capture. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'13)*, 32(4). (Cited on pages 68, 81 and 82.)
- Mazhar, H., Heyn, T., Negrut, D., and Tasora, A. (2015). Using Nesterov's method to accelerate multibody dynamics with friction and contact. *ACM Transactions on Graphics*, 34(3):1–14. (Cited on page 45.)

- McAdams, A., Selle, A., Ward, K., Sifakis, E., and Teran, J. (2009). Detail preserving continuum simulation of straight hair. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'09)*, 28(3):1–6. (Cited on page 31.)
- McCrae, J. and Singh, K. (2008). Sketching piecewise clothoid curves. *Computers & Graphics*, 33(4):452–461. (Cited on page 65.)
- Miguel, E., Bradley, D., Thomaszewski, B., Bickel, B., Matusik, W., Otaduy, M. A., and Marschner, S. (2012). Data-driven estimation of cloth simulation models. *Comp. Graph. Forum*, pages 519–528. (Cited on pages 81 and 82.)
- Miguel, E., Tamstorf, R., Bradley, D., Schwartzman, S., Thomaszewski, B., Bickel, B., Matusik, W., Marschner, S., and Otaduy, M. (2013). Modeling and estimation of internal friction in cloth. *ACM SIGGRAPH Asia'13*. (Cited on page 82.)
- Miller, A. and Christensen, H. (2003). Implementation of multi-rigid-body dynamics within a robotic grasping simulator. In *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, volume 2, pages 2262–2268 vol.2. (Cited on page 42.)
- Moore, M. and Wilhelms, J. (1988). Collision detection and response for computer animation. In *Computer Graphics Proceedings (Proc. ACM SIGGRAPH'88)*, pages 289–298. (Cited on page 33.)
- Moreau, J. (1994). Some numerical methods in multibody dynamics: Application to granular materials. *European Journal of Mechanics - A/Solids*, supp.(4):93–114. (Cited on page 49.)
- Moreau, J.-J. (1988). Unilateral contact and dry friction in finite freedom dynamics. *Nonsmooth mechanics and applications*, CISM Courses Lect. 302, 1-82 (1988). (Cited on pages 31, 39 and 40.)
- Moreau, J.-J. and Jean, M. (1996). Numerical treatment of contact and friction: the contact dynamics method. In *The 1996 3rd Biennial Joint Conference on Engineering Systems Design and Analysis, ESDA*, pages 201–208, Montpellier; Fr. (Cited on pages 31 and 36.)
- Narain, R., Golas, A., and Lin, M. C. (2010). Free-flowing granular materials with two-way solid coupling. *ACM Transactions on Graphics*, 29(6):1–10. (Cited on pages 33, 49 and 54.)
- Neher, M. (1999). An enclosure method for the solution of linear ODEs with polynomial coefficients. *Numerical Functional Analysis and Optimization*, 20:779–803. (Cited on pages 22 and 24.)
- Neukirch, S. (2004). Extracting DNA twist rigidity from experimental supercoiling data. *Physical Review Letters*, 93. (Cited on page 37.)

- Otaduy, M., Tamstorf, R., Steinemann, D., and Gross, M. (2009). Implicit contact handling for deformable objects. *Computer Graphics Forum (Proc. Eurographics'09)*, 28(2). (Cited on page 42.)
- Pai, D. (2002). Strands: Interactive simulation of thin solids using Cosserat models. *Computer Graphics Forum (Proc. Eurographics'02)*, 21(3):347–352. (Cited on pages 12, 13 and 14.)
- Pérez, J., Thomaszewski, B., Coros, S., Bickel, B., Canabal, J. A., Sumner, R., and Otaduy, M. A. (2015). Design and fabrication of flexible rod meshes. *ACM Trans. Graph.*, 34(4):138:1–138:12. (Cited on page 79.)
- Poole, E. (1936). *Introduction to the theory of linear differential equations*. Clarendon Press. (Cited on page 22.)
- Prasolov, V. (1994). *Problems and theorems in linear algebra*. American Mathematical Society; 1st edition (June 13, 1994). (Cited on page 63.)
- Qi, L. and Sun, J. (1993). A nonsmooth version of Newton’s method. *Mathematical Programming*, 58:353–367. (Cited on page 44.)
- Redon, S., Galoppo, N., and Lin, M. (2005). Adaptive dynamics of articulated bodies. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'05)*, 24(3):936–945. (Cited on page 28.)
- Reissner, E. (1973). One one-dimensional large-displacement finite-strain beam theory. *Studies in App. Math.*, 52(2):87–95. (Cited on pages 14 and 18.)
- Simo, J. and Vu-Quoc, L. (1986). A three-dimensional finite-strain rod model. part ii: Computational aspects. *Computer Methods in Applied Mechanics and Engineering*, 58(1):79 – 116. (Cited on page 14.)
- Skouras, M., Thomaszewski, B., Bickel, B., and Gross, M. (2012). Computational design of rubber balloons. *Comput. Graphics Forum (Proc. Eurographics)*. (Cited on page 79.)
- Spillmann, J. and Teschner, M. (2007). CoRdE: Cosserat rod elements for the dynamic simulation of one-dimensional elastic objects. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'07)*, pages 63–72. ACM-EG SCA. (Cited on page 14.)
- Spillmann, J. and Teschner, M. (2008). An adaptive contact model for the robust simulation of knots. *Computer Graphics Forum*, 27(2). Proc. Eurographics’08. (Cited on page 34.)
- Stewart, D. and Trinkle, J. (1996). An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and Coulomb friction. *International Journal for Numerical Methods in Engineering*, 39(15). (Cited on page 42.)

- Sugiyama, H., Gertsmayr, J., and Mikkola, A. (2013). Flexible multibody dynamics — Essential for accurate modeling in multibody system dynamics. *Journal of Computational Nonlinear Dynamics*, 9. (Cited on page 29.)
- Twigg, C. and Kačić-Alesić, Z. (2011). Optimization for sag-free simulations. In *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'11)*, pages 225–236. ACM-EG SCA. (Cited on pages 79 and 83.)
- Wang, H., Liao, M., Zhang, Q., Yang, R., and Turk, G. (2009). Physically guided liquid surface modeling from videos. *ACM SIGGRAPH'09*. (Cited on page 83.)
- Wang, H., Ramamoorthi, R., and O'Brien, J. (2011). Data-driven elastic models for cloth: modeling and measurement. *ACM Transactions on Graphics (SIGGRAPH 2011)*, 30(4):71:1–71:12. (Cited on page 82.)
- Ward, K., Bertails, F., Kim, T.-Y., Marschner, S., Cani, M.-P., and Lin, M. (2007). A survey on hair modeling: Styling, simulation, and rendering. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 13(2):213–34. (Cited on page 12.)
- Williams, R. (2009). *The animator's survival kit*. Faber & Faber, Inc. (Cited on page 5.)
- Zhao, Y. and Barbič, J. (2013). Interactive authoring of simulation-ready plants. *ACM Trans. Graph.*, 32(4):84:1–84:12. (Cited on page 79.)
- Zhu, Y. and Bridson, R. (2005). Animating sand as a fluid. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'05)*, 24(3):965–972. (Cited on page 49.)

Numerical Modeling of elastic slender structures subject to contact and friction:

From dynamic simulation to inverse static design

Abstract: Slender structures (filaments or surfaces) are ubiquitous around us. It is actually striking to note that such thin structures represent most of the deformable objects composing our direct environment, stemming either from natural biological phenomena — plant stems and leaves, flower petals, hair, skin — or manufacturing processes — cables, ribbons, paper, cloth. When multiple thin structures are coupled together with contact and friction, the range of emerging phenomena is even more exacerbated, giving rise to stick-slip dynamical instabilities, entangling, or spontaneous collective behavior.

This memoir presents an overview of our work on the simulation of slender elastic structures and entangled materials, such as hair, with a specific interest for virtual prototyping and computer graphics applications.

I first introduce a family of high-order, reduced models for discretizing Kirchhoff's equations for thin elastic rods in a both faithful and robust way. Such models are particularly well-suited for simulating inextensible fibers subject to bending and twisting, and featuring an arbitrary curly resting geometry. Then I show how such models can be coupled to frictional contact using the nonsmooth contact dynamics framework, and I present a hybrid iterative solver suitable for robustly handling thousands packed fibers at reasonable frame rates. A first continuum-based approach. Finally, I give some insights into the inverse modeling of fibers, consisting in taking an arbitrary curve geometry as input and inferring corresponding geometric and physical parameters of the simulator such that the input geometry corresponds to a stable configuration at equilibrium.

Keywords: Thin elastic rods, frictional contact, numerical simulation, direct and inverse modeling

Modélisation numérique de structures élancées en présence de contact frottant :

De la simulation dynamique à la conception statique inverse

Résumé : Les structures élancées (objets filiformes ou surfaciques), font partie intégrante de notre environnement quotidien : les exemples biologiques naturels — branches d’arbre, feuilles, pétales, cheveux, peau — rivalisent en nombre avec les objets issus de l’industrie — câbles, rubans, papier, vêtements. La diversité des formes et la non-linéarité des déformations de ces structures procurent une grande richesse visuelle, qui est davantage exacerbée lorsque des contacts et du frottement entrent en jeu, donnant lieu à des phénomènes d’instabilité de type glissement-adhérence, des comportements collectifs spontanés, ou de l’emmêlement. Synthétiser de manière réaliste la forme et le mouvement de ces phénomènes à l’échelle macroscopique intéresse grandement l’industrie du cinéma d’animation et des effets spéciaux, mais aussi, et de plus en plus, des secteurs industriels plus traditionnels en ingénierie mécanique, soucieux de raffiner et d’optimiser leurs procédés de conception et de fabrication grâce au prototypage virtuel.

Ce mémoire présente une synthèse de dix années de recherche visant à capturer quelques-uns de ces phénomènes mécaniques complexes de manière à la fois réaliste, robuste, et efficace en temps de calcul. J’introduis tout d’abord une famille de modèles réduits et de haut degré en espace permettant de discrétiser les équations de Kirchhoff pour les tiges élastiques minces. Ces modèles sont particulièrement adaptés à la simulation de filaments capables de se fléchir et de se tordre tout en restant parfaitement inextensibles. Après l’évocation de perspectives d’extension au cas surfacique, je montre comment ces structures minces peuvent être couplées entre elles via du contact et du frottement solide, en m’appuyant sur des principes de mécanique non-régulière. En particulier, j’introduis un solveur numérique hybride capable de résoudre efficacement le problème de contact frottant discret, pour des dizaines de milliers de contacts. Ce modèle numérique nous a notamment permis de simuler la dynamique de milliers de fibres enchevêtrées, et de commencer à envisager la simulation dynamique de vêtements. Nous avons également pu étendre la portée de notre approche à la résolution du mouvement de millions de grains en contact frottant, modélisés comme un écoulement fluide dont la loi de comportement (loi de Drucker-Prager), est non-régulière. Finalement, j’évoque nos travaux en conception statique inverse de fibres, consistant à inférer les paramètres physiques de nos modèles à partir de la seule observation de leur forme sous gravité et éventuellement en présence de contact frottant. Au-delà de la perspective d’améliorer le contrôle utilisateur d’une simulation physique dans un contexte de création artistique, ces recherches laissent espérer à long terme la conception de protocoles d’identification paramétrique non-invasifs pour les structures élancées en contact frottant.
