

Université de Montréal

**Computational Modeling and Design of  
Nonlinear Mechanical Systems and Materials**

par

**Pengbin Tang**

Département d'informatique et de recherche opérationnelle  
Faculté des arts et des sciences

Thèse présentée en vue de l'obtention du grade de  
Philosophiæ Doctor (Ph.D.)  
en Informatique

Mars 2023



# Université de Montréal

Faculté des arts et des sciences

---

Cette thèse intitulée

## Computational Modeling and Design of Nonlinear Mechanical Systems and Materials

présentée par

**Pengbin Tang**

a été évaluée par un jury composé des personnes suivantes :

*Mikhail Bessmeltsev*

---

(président-rapporteur)

*Bernhard Thomaszewski*

---

(directeur de recherche)

*Pierre Poulin*

---

(membre du jury)

*Eitan Grinspun*

---

(examineur externe)

*Andrea Bianchi*

---

(représentant du doyen de la FESP)



# Résumé

---

Les systèmes et matériaux mécaniques non linéaires sont largement utilisés dans divers domaines. Cependant, leur modélisation et leur conception ne sont pas triviales car elles nécessitent une compréhension complète de leurs non-linéarités internes et d'autres phénomènes. Pour permettre une conception efficace, nous devons d'abord introduire des modèles de calcul afin de caractériser avec précision leur comportement complexe. En outre, de nouvelles techniques de conception inverse sont également nécessaires pour comprendre comment le comportement change lorsque nous modifions les paramètres de conception des systèmes mécaniques non linéaires et des matériaux. Par conséquent, dans cette thèse, nous présentons trois nouvelles méthodes pour la modélisation informatique et la conception de systèmes mécaniques non linéaires et de matériaux.

Dans le premier article, nous abordons le problème de la conception de systèmes mécaniques non linéaires présentant des mouvements périodiques stables en réponse à une force périodique. Nous présentons une méthode de calcul qui utilise une approche du domaine fréquentiel pour la simulation dynamique et la puissante analyse de sensibilité pour l'optimisation de la conception afin de concevoir des systèmes mécaniques conformes avec des oscillations de grande amplitude. Notre méthode est polyvalente et peut être appliquée à divers types de systèmes mécaniques souples. Nous validons son efficacité en fabriquant et en évaluant plusieurs prototypes physiques.

Ensuite, nous nous concentrons sur la modélisation informatique et la caractérisation mécanique des matériaux non linéaires dominés par le contact, en particulier les matériaux à emboîtement discret (DIM), qui sont des tissus de cote de mailles généralisés constitués d'éléments d'emboîtement quasi-rigides. Contrairement aux matériaux élastiques conventionnels pour lesquels la déformation et la force de rappel sont directement couplées, la mécanique des DIM est régie par des contacts entre des éléments individuels qui donnent lieu à des contraintes de déformation cinématique anisotrope. Pour reproduire le comportement biphasique du DIM sans simuler des structures à micro-échelle coûteuses, nous introduisons une méthode efficace de limitation de la déformation anisotrope basée sur la programmation conique du second ordre (SOCP). En outre, pour caractériser de manière exhaustive la forte anisotropie, le couplage complexe et d'autres phénomènes non linéaires du DIM,

nous introduisons une nouvelle approche d'homogénéisation pour distiller des limites de déformation à grande échelle à partir de simulations à micro-échelle et nous développons un modèle macromécanique basé sur des données pour simuler le DIM avec des contraintes de déformation homogénéisées.

**Mots-clés:** Vibration non linéaire, systèmes mécaniques non linéaires, méthode d'équilibre harmonique, mouvements périodiques, cote de mailles, matériaux discrets à emboîtement, limitation de déformation, caractérisation mécanique, homogénéisation, analyse de sensibilité, Gauss-Newton, conception computationnelle.

# Abstract

---

Nonlinear mechanical systems and materials are broadly used in diverse fields. However, their modeling and design are nontrivial as they require a complete understanding of their internal nonlinearities and other phenomena. To enable their efficient design, we must first introduce computational models to accurately characterize their complex behavior. Furthermore, new inverse design techniques are also required to capture how the behavior changes when we change the design parameters of nonlinear mechanical systems and materials. Therefore, in this thesis, we introduce three novel methods for computational modeling and design of nonlinear mechanical systems and materials.

In the first article, we address the design problem of nonlinear mechanical systems exhibiting stable periodic motions in response to a periodic force. We present a computational method that utilizes a frequency-domain approach for dynamical simulation and the powerful sensitivity analysis for design optimization to design compliant mechanical systems with large-amplitude oscillations. Our method is versatile and can be applied to various types of compliant mechanical systems. We validate its effectiveness by fabricating and evaluating several physical prototypes.

Next, we focus on the computation modeling and mechanical characterization of contact-dominated nonlinear materials, particularly Discrete Interlocking Materials (DIM), which are generalized chainmail fabrics made of quasi-rigid interlocking elements. Unlike conventional elastic materials for which deformation and restoring forces are directly coupled, the mechanics of DIM are governed by contacts between individual elements that give rise to anisotropic kinematic deformation constraints. To replicate the biphasic behavior of DIM without simulating expensive microscale structures, we introduce an efficient anisotropic strain-limiting method based on second-order cone programming (SOCP). Additionally, to comprehensively characterize strong anisotropy, complex coupling, and other nonlinear phenomena of DIM, we introduce a novel homogenization approach for distilling macroscale deformation limits from microscale simulations and develop a data-driven macromechanical model for simulating DIM with homogenized deformation constraints.

**Keywords:** Nonlinear Vibration, Nonlinear Mechanical Systems, Harmonic Balance Method, Periodic Motions, Chainmail, Discrete Interlocking Materials, Strain Limiting, Mechanical Characterization, Homogenization, Sensitivity Analysis, Computational Design.



# Contents

---

<b>Résumé</b> .....	5
<b>Abstract</b> .....	7
<b>List of Tables</b> .....	13
<b>List of Figures</b> .....	15
<b>List of Symbols and Abbreviations</b> .....	21
<b>Acknowledgments</b> .....	23
<b>Chapter 1. Introduction</b> .....	25
1.1. Designing Compliant Mechanical Systems with Nonlinear Periodic Motions...	26
1.2. Modeling and Characterizing Discrete Interlocking Materials .....	27
<b>Chapter 2. A Harmonic Balance Approach for Designing Compliant     Mechanical Systems with Nonlinear Periodic Motions</b> .....	29
2.1. Introduction .....	30
2.2. Related Work .....	32
2.3. Theory .....	34
2.3.1. Equations of Motion in Frequency Space .....	34
2.3.2. Evaluation of Nonlinear Forces and Derivatives .....	35
2.3.3. Frequency Response Curves and Continuation .....	36
2.4. Computational Design .....	38
2.4.1. Dynamical Equilibrium and Sensitivity .....	38
2.4.2. Forward Sensitivity Exploration .....	39
2.4.3. Inverse Design .....	40
2.5. Results .....	41
2.5.1. Analysis & Validation .....	41

2.5.2.	Measuring Damping Parameters .....	44
2.5.3.	Forward Design with Sensitivity Exploration .....	44
2.5.4.	Optimization-Based Inverse Design .....	46
2.5.5.	Statistics & Additional Validation .....	50
2.6.	Limitations & Future Work .....	52
2.A.	Appendix .....	54
2.A.1.	Frequency-Space Equilibrium Equations .....	54
2.A.2.	DFT Operators .....	56
2.A.3.	Newmark Time Integration .....	56
2.A.4.	Accuracy of HBM vs. Newmark .....	57
2.A.5.	Mechanical Models .....	58
<b>Chapter 3.</b>	<b>A Second Order Cone Programming Approach for Simulating</b>	
	<b>Biphasic Materials .....</b>	<b>59</b>
3.1.	Introduction .....	60
3.2.	Related Work .....	61
3.3.	Theory .....	63
3.3.1.	Deformation Limits as Quadratic Constraints .....	64
3.3.2.	Constraint Projection as QCQP .....	64
3.3.3.	Conic Reformulation .....	65
3.4.	Time Stepping .....	66
3.5.	Results .....	69
3.5.1.	Performance .....	70
3.6.	Conclusion .....	71
	Acknowledgments .....	72
3.A.	Appendix A .....	72
<b>Chapter 4.</b>	<b>Beyond Chainmail: Computational Modeling of Discrete</b>	
	<b>Interlocking Materials .....</b>	<b>73</b>
4.1.	Introduction .....	74
4.2.	Related Work .....	76

4.3. Computational Model .....	78
4.3.1. Native-Scale Model .....	78
4.3.2. Macro-Scale Deformations .....	79
4.3.3. Strain-Space Representation .....	81
4.3.4. Macro-Scale Model .....	82
4.4. Results .....	83
4.4.1. Experimental Validation .....	83
4.4.2. Analysis of Discrete Interlocking Materials .....	84
4.4.3. Macro-Scale Simulation .....	89
4.5. Conclusions .....	94
4.5.1. Limitations & Future Work .....	94
Acknowledgments .....	95
<b>Chapter 5. Conclusion .....</b>	<b>97</b>
5.1. Future Work .....	98
<b>Bibliography .....</b>	<b>101</b>



# List of Tables

---

2.1	<b>Statistics for inverse design examples.</b> The columns list numbers of degrees of freedom ( $N_{DoF}$ ), harmonics ( $N_H$ ), sampling points ( $N_{AFT}$ ), parameters ( $N_p$ ), iterations required for convergence ( $N_{it}$ ), as well as the total time spent on optimization. ....	51
3.1	Performance of our method compared to unconstrained implicit Euler (IE), general nonlinear programming (NLP), and Gauss-Seidel (GS) as a function of mesh resolution under isotropic strain limiting. Timings are averaged across 100 animation steps. The symbol * indicates that some steps did not converge whereas – indicates solver failure. ....	70
4.1	Performance comparison between native- and macro-scale simulations for the examples shown in Fig. 4.9. ....	91



# List of Figures

---

1.1	In this thesis, we model and design nonlinearities at dynamical and kinematic levels. Specifically, we design nonlinear mechanical systems with nonlinear periodic motions. Furthermore, we model and characterize nonlinear kinematic deformation limits for a new class of contact-dominated nonlinear materials by considering micro-, meso-, and macro-levels. ....	26
1.2	<b>Physical prototypes of DIM deform under different states.</b> (a) The structure exhibits extreme flexibility. (b) A DIM shows stiff bending response to resist self-weight under gravity. (c) The chainmail structure presents a strong in-plane stiffness to the external deformation. ....	28
2.1	Our method enables optimization-driven design of compliant mechanical systems with periodic large-amplitude motions. For this pair of dragon wings, the initial design ( <i>a, b</i> ) exhibits only small oscillation response when driven by harmonic forcing at a frequency of $2.5Hz$ . Our approach automatically finds optimized design parameters (extra masses at the trailing edge of the wing) that lead to substantially increased amplitude ( <i>c, d</i> ). ....	30
2.2	Frequency response curve for a thin shell model. Amplitude is measured using the trajectory of a selected vertex shown in <i>green</i> . The inset figures illustrate maximum-deflection configurations at 9 rad/s, 11 rad/s and 13 rad/s, respectively.	37
2.3	Trajectory error for HBM compared to ground truth time-domain simulations. The plot shows error as a function of the number of harmonics using damping coefficients as indicated. ....	42
2.4	Trajectory comparison between fabricated L-wing and HBM simulations at driving frequency of $2.0Hz$ with different numbers of sampling points $N_{AFT}$ : physical prototype (a) and HBM simulation using (b) $N_{AFT} = 128$ , (c) $N_{AFT} = 64$ , and (d) $N_{AFT} = 32$ . ....	43
2.5	Frequency responses computed with LMA and HBM for different driving frequencies with damping coefficients $D_\alpha = 0.5$ and $D_\beta = 0.00005$ . Trajectory of a selected tip point (green) and color-coded maximum in-plane stretch. For	

	LMA, the maximum strains over one period for the three driving frequencies are 0.655, 0.582, and 0.575. The corresponding values for HBM ( $8.25e^{-5}$ , $1.00e^{-4}$ , and $8.50e^{-5}$ ) are 4-5 orders of magnitude smaller. ....	44
2.6	Forward Design with Sensitivity Exploration illustrated on an animatronic character. The initial design ( <i>left</i> ) exhibits only small oscillations at the hands (indicated in <i>green</i> ). After several steps of forward exploration, the final design exhibits an expressive large-amplitude motion ( <i>right</i> 33). ....	45
2.7	Amplitude of the wing tip before ( <i>blue</i> ) and after ( <i>orange</i> ) optimization for the dragon example. The dashed line indicates the driving frequency ( $2.5Hz$ ) used during optimization. ....	45
2.8	<b>Trajectory matching for the three-link compliant mechanism.</b> Initial design ( <i>left</i> ) and optimized design ( <i>right</i> ) with simulated ( <i>top</i> ) and real-world ( <i>bottom</i> ) end-effector trajectories shown in green. The target trajectory is shown in <i>blue</i> . ....	46
2.9	<b>Running Ostrich.</b> Four images from a running sequence of our ostrich model with two legs driven at $1.2Hz$ with a phase offset of half a period. ....	47
2.10	<b>Ostrich leg.</b> Performance of initial ( <i>left</i> ) design and optimized design ( <i>right</i> ) in simulation ( <i>top</i> ) and on the physical prototype ( <i>bottom</i> ). Target and actual trajectories are shown in <i>blue</i> and <i>green</i> , respectively. ....	47
2.11	For this animatronic wire character, we optimize the weights of three additional masses such as to achieve large-amplitude oscillation of its tail. ....	50
2.12	We optimize per-layer material stiffness for this solid such as to maximize the amplitude of the selected vertex when driving the top face with harmonic rotational excitation at $2.0Hz$ . ....	51
2.13	Material optimization on an Eiffel Tower model. For the initial design with homogeneous material ( <i>left</i> ), the amplitude at the tip is almost the same as for the driving signal. After optimizing for per-element stiffness values, the tip amplitude is substantially increased ( <i>right</i> ). ....	52
2.14	<b>Maximum error of constraints violation.</b> We use a penalty stiffness of $1e^7$ and $N_{AFT} = 64$ time-domain samples for these three compliant mechanism examples. Each curve shows the maximum error (change in angle/length divided by corresponding original value for angular and distance constraints) over all	



	constraints for each example within a period. It can be seen that constraint violations remain below $2e^{-3}\%$ at all times for all cases. ....	53
2.15	Trajectory difference between HBM and Newmark as a function of the number of time-domain samples $N_{AFT}$ for $N_H = 5$ harmonics. The step size for Newmark is set as $\Delta t = T/N_{AFT}$ , where $T$ is the period of the forcing.....	57
3.1	Two examples of 3D-printed chainmail with experimentally measured deformation limits. Orange areas in the plots indicate the variation of experimental data. ....	61
3.2	Our method can enforce stretching limits in arbitrary material directions including isotropic, orthotropic, and fully anisotropic constraints. Top row: polar plots showing input curves for deformation limits and sample locations for which constraints are enforced. Bottom row: simulation results for a static drape test using a mesh resolution of $60 \times 60$ .....	62
3.3	A schematic visualization of our quasi-symmetric step-and-reflect method. We first step to the middle of the time interval, where we apply twice the correction necessary to satisfy deformation constraints. We then integrate from mid-step to the end of the interval and apply a standard projection step. ....	67
3.4	Evolution of total energy over time for different methods on a swinging cloth example with 1% isotropic strain limits. Unconstrained implicit Euler (IE) loses energy more rapidly than our SOCP approach with end-of-step projection (SAP). Our step-and-reflect method (SAR) further improves energy conservation. For BDF-2 versions of projection (BAP) and reflection (BAR) methods, we obtain even lower numerical damping, but larger energy fluctuations.....	67
3.5	Evolution of total energy over time for different step sizes on a swinging cloth example with 1% isotropic strain limits. ....	68
3.6	Our SOCP-based strain limiting method can be integrated into standard collision resolution approaches. Here we use simple penalty functions that prevent intersection during time integration and enforce isotropic strain limits of 0.1%. ...	68
4.1	Discrete Interlocking Materials are governed by strongly coupled, highly anisotropic, and asymmetric deformation limits. Our method is able to capture and reproduce these effects for many types of interlocking materials ( <i>a</i> ). Using native-scale simulations as a basis ( <i>b</i> ), we construct macromechanical deformation limits on bending and stretching ( <i>c</i> ) which we use to develop an efficient macro-scale simulation model ( <i>d</i> ). ....	74

4.2	Element cross-section determines coupling for bending. Two threefold symmetric materials are bent using the same uniaxial load case. The elements of design (a) have circular cross-section, while cross-sections for design (b) are slightly ellipsoidal. This small asymmetry in geometry leads to much larger coupling between principal and orthogonal curvature. ....	75
4.3	Deformation limits for a three-fold symmetric material are probed in simulation using in-plane periodic boundary conditions (a) and paraboloid bending tests (b). ....	79
4.4	Strain-space boundaries for the threefold symmetric chainmail material for in-plane (left) and out-of-plane (right) deformations. Projections of the 3D shapes onto strain-space planes are shown to improve readability. ....	81
4.5	Dimensions of elements (mm) and topology of interlocking for the materials considered in this work. ....	85
4.6	<b>Deformations limits for threefold symmetric chain mail, torus knot, and scale mail materials.</b> Polar plots show directional deformation limits for uniaxial stretching (b) as well as biaxial curvature in the principal (c) and orthogonal (d) directions under uniaxial loading. For the uniaxial tests, we bend structures in a given principal direction until reaching their curvature limits. In the orthogonal direction, the structure is left free to deform under gravity and thus reaches a secondary curvature limit. We show the corresponding curvature limits in the principal (c) and orthogonal (d) directions. Solid curves (blue) correspond to simulation results which are sampled from $[0,\pi)$ and rotated to fill the range $[\pi, 2\pi)$ , experimental data is indicated using error bars (orange). These polar plots are partial views onto strain space boundaries for in-plane (e) and out-of-plane deformations (f). ....	86
4.7	4-in-1 chainmail consisting of a staggered arrangement of rings, each of which connects to their four immediate neighbors (a). Deformation limits for uniaxial stretching (b) and strain space boundary for in-plane deformations (c). Bending limits are determined by contact between non-connected elements (d). ....	87
4.8	Limits for double curvature depend on sample size. Absolute curvature values in both the principal and orthogonal directions decrease as patch size increases from 4 (a) to 6 (b) and 8 (c) elements along the diameter, corresponding to patch radii of 2.8cm, 4.8cm, and 5.8cm, respectively. ....	88
4.9	Comparison between native- and macro-scale simulations for different materials. Static equilibrium states for applied point loads of 0.5N (red arrows) with different	

patch sizes. **Threefold symmetric chainmail:** we use disk patches of size  $2.8cm$  (*a*),  $4.8cm$  (*b*), and  $5.8cm$  (*c*). Our macro-scale model uses the same resolution— but different physical sizes— of 384 faces and 217 nodes for all three cases. **Torus knot material:** we use patches with  $4 \times 4$  (*a*),  $5 \times 5$  (*b*), and  $6 \times 6$  (*c*) elements. Our macroscale model use again the same resolution ( $\#face = 200$ ,  $\#nodes = 121$ ) but different physical sizes with  $4.4cm$  (*a*),  $5.5cm$  (*b*), and  $6.4cm$  (*a*) side length, respectively. **Scale mail:** we use disk-shaped patches with 19 (*a*), 37 (*b*), and 61 (*c*) elements and corresponding radii of  $3.4cm$ ,  $4.6cm$ , and  $5.9cm$ , respectively. Our macro-scale simulations use the same mesh resolution (384 faces and 217 nodes) for all patch sizes. .... 90

4.10 Comparison of macro-scale simulation with increasing mesh resolution. We start from with coarse meshes (*a*) and increase resolution by a factor of 4 (*b*) and 16 (*c*)..... 91

4.11 Comparison of macro-scale simulations with increasing Young’s Modulus. .... 92

4.12 **Asymmetric Bending Limits.** Due to the through-the-thickness asymmetry of its elements, this scale mail prototype (*left*) exhibits very different bending limits for positive and negative curvature. Our macro-scale model (*right*) accurately captures this behavior and closely tracks the native-scale simulation results (*middle*). .... 92

4.13 Constraint violations for the examples shown in Fig. 4.9. Maximum constraint per-element constraint violations are plotted separately for in- and out-of-plane limits as indicated. It can be seen that constraint violations stay below 0.01% for all cases..... 93



## List of Symbols and Abbreviations

---

LMA	Linear Modal Analysis
HBM	Harmonic Balance Method
AFT	Alternating Frequency/Time-domain
DFT	Discrete Fourier Transforms
FDM	Fused Deposition Modeling
PLA	Polylactic Acid
SOCP	Second Order Cone Programming
QCQP	Quadratically Constrained Quadratic Program
PBD	Position-based Dynamics
ADMM	Alternating Direction of Multipliers Method
CP	Conic Programs
SAP	Step-and-project
SAR	Step-and-reflect
BAP	BDF-2 Step-and-project
BAR	BDF-2 Step-and-reflect
CST	Constant Strain Triangle
IE	Implicit Euler
NLP	General Nonlinear Programming
GS	Gauss-Seidel
DIM	Discrete Interlocking Materials



## Acknowledgments

---

First, I would like to express my deepest gratitude to my supervisor, Prof. Dr. Bernhard Thomaszewski, for his trust, patience, support, and discussions throughout my doctorate study.

I am also extremely grateful to Prof. Dr. Stelian Coros for many insightful discussions and suggestions and for inviting and helping with the visit to the CRL lab in ETH Zürich.

I would like to thank Prof. Dr. Pierre Poulin and Prof. Dr. Mikhail Bessmeltsev for helping with all issues at UdeM during my doctorate study. I am also thankful to all members at LIGUM and special thanks to Dr. Jonas Zehnder for his kind help at the beginning of my doctorate.

I would like to thank members at CRL lab, ETH Zürich and I would especially thank Yue Li, Dr. Juan Sebastián Montes Maestre, and Dr. Ronan Hinchet for valuable discussions and help. Furthermore, I would like to thank Barbara Gleich for taking care of my affairs during my stay at ETH Zürich.

I also would like to thank the Thesis Jury for spending their valuable time on my thesis review and providing insightful feedback and suggestions for possible future research.

Finally, I would like to thank my family and my friends. I could not make this without their support. I also would like to thank my wife for supporting and encouraging me with love and patience throughout my doctorate study.





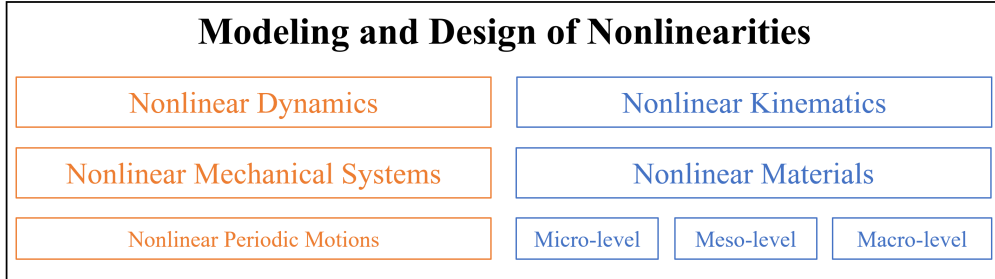
# Chapter 1

---

## Introduction

Nonlinear mechanical systems and materials are widely used in a diverse range of applications and tasks due to their fascinating mechanical properties. However, predicting and designing them is challenging due to their complex nonlinear nature. The recent advancements in 3D printing technology and computation have enabled the development of efficient digital modeling methods for designing nonlinear mechanical systems and materials. This approach allows us to iterate and validate designs in a virtual environment, and with the aid of physical simulation, we can design the desired functioning and mechanical properties of systems and materials with complex structures and other nonlinear phenomena. Nevertheless, designing nonlinear mechanical systems and materials is still a nontrivial problem: 1) nonlinear mechanical systems and materials are characterized by large deformations, strong anisotropy, internal contacts, and other inherently nonlinear phenomena that are already hard to predict and characterize; 2) local design parameters such as mass distribution and structural geometry have a complex, nonlinear effect on the performance of mechanical systems and materials. Together, these parameters combine into a high-dimensional design space that is all but expensive to navigate with trial-and-error exploration. To unlock the design of nonlinear mechanical systems and materials, we must develop accurate and efficient models for their simulation so that we can manage and leverage the internal nonlinearities for their design.

In this thesis, we aim to develop novel techniques for modeling and designing nonlinear mechanical systems and materials at various levels, as depicted in Fig. 1.1. The first part of the thesis focuses on designing nonlinear dynamical motions. Specifically, we present a new computational tool that allows for efficient and accurate design of nonlinear periodic motions for compliant mechanical systems. In the remaining chapters of the thesis, we explore a new class of contact-dominated nonlinear materials that are subject to deformation constraints imposed by contact and display kinematic deformation limits in both in-plane and out-of-plane deformation. To enable efficient simulation and exploration of potential applications



**Figure 1.1** – In this thesis, we model and design nonlinearities at dynamical and kinematic levels. Specifically, we design nonlinear mechanical systems with nonlinear periodic motions. Furthermore, we model and characterize nonlinear kinematic deformation limits for a new class of contact-dominated nonlinear materials by considering micro-, meso-, and macro-levels.

of these nonlinear materials, we develop new methods for modeling and characterizing their nonlinear kinematic deformation limits across micro-, meso-, and macro-levels.

## 1.1. Designing Compliant Mechanical Systems with Nonlinear Periodic Motions


Compliant mechanical systems are essential for a wide range of applications, and many researchers have used physics-based simulations to analyze and improve the stability of 3D-printed flexible structures [Dumas et al., 2015; Zehnder et al., 2016a; Ulu et al., 2017] while minimizing material use [Lu et al., 2014a; Martínez et al., 2015]. Various methods have also been proposed to design compliant structures such as thin shells, elastic rods, and deformable solids under static equilibrium for different purposes. Designs have varied from inflatable membranes [Skouras et al., 2014a] to flexible rod meshes [Pérez et al., 2015], Kirchhoff plateau surfaces [Pérez et al., 2017], and flat, flexible panels for assembling desired 3D shapes [Malomo et al., 2018]. However, to accomplish specified tasks or functions for compliant mechanical systems, they are always designed to perform specified motions. For example, Miyashita et al. [2020] designed a compliant mechanism to mimic the walking motion of horses via an external periodic driving force. Most structures driven by periodic forces will eventually reach stable steady-state motions, which is also referred to as limit cycle or periodic motions. Though lots of works design nonlinear mechanical systems with desired periodic motions, most of them only consider kinematics [Ceylan et al., 2013a; Megaro et al., 2017] and ignore the dynamics. To consider nonlinear dynamics with periodic motions, a naive way is using nonlinear time-domain methods, e.g., Newmark time integration. However, this approach needs to simulate an unknown number of time steps to reach the periodic motion, which depends on the internal elastic and viscous properties of the structure, external periodic stimuli, and external damping. This makes the inverse design intractable to

handle because inverting all steps of time integration would become extremely expensive to compute. Another alternative method for predicting periodic motions is linear modal analysis (LMA), which is a well-known frequency-space method. However, this method is limited to predicting structures under small-amplitude oscillations. Therefore, in the first project of the thesis, we aim to develop a new computational tool to design nonlinear mechanical systems with desired nonlinear periodic motions, as described in Chapter 2.

## 1.2. Modeling and Characterizing Discrete Interlocking Materials

Metamaterials are carefully-engineered materials that exhibit unique mechanical properties and functionalities that cannot be found in nature, e.g., negative Poisson’s ratio [Li et al., 2020a; Babaei et al., 2013; Qin and Qin, 2020], tunable stiffness [Wang et al., 2021; Chen et al., 2021; Wang and Sigmund, 2021], and shape morphing in response to mechanical forces [Overvelde et al., 2016; Li et al., 2021a]. These unusual properties of metamaterials have led to the development of new applications and products in diverse fields, including medical devices, aerospace applications, and smart clothing, which makes the design of metamaterials a popular field. Consequently, a plethora of dedicated structures and materials has been introduced in material science and engineering. Lots of works devoted to the design of dedicated structures, such as mechanism-based [Grima and Evans, 2000; Cui and Ju, 2020], origami-inspired [Filipov et al., 2015; Tang and Yin, 2017; Zhai et al., 2018], and kirigami-inspired [Isobe and Okumura, 2016; Hwang and Bartlett, 2018; Zhang and Paik, 2022] metamaterials. Researchers in the graphics community have also explored several methods to design complex fine-scale geometric structures [Schumacher et al., 2015; Panetta et al., 2015] with isotropic [Martínez et al., 2016], orthotropic [Martínez et al., 2017], and anisotropic [Zehnder et al., 2017; Schumacher et al., 2018; Martínez et al., 2019] properties.

Here, we consider contact-dominated nonlinear materials, particularly Discrete Interlocking Materials (DIM), which are generalized chainmail fabrics made of quasi-rigid interlocking elements as shown in Figure 1.2. Unlike traditional materials, DIMs are regular arrangements of elements connected through interlocking. The dimensions of the elements are chosen in such a way that they can move freely within a range defined through contact with neighboring elements. This results in nonlinear kinematic deformation limits for both in-plane and out-of-plane deformations. This fabric-like structure gives DIMs extreme flexibility to allow materials to adapt to necessary shapes; see Figure 1.2(a). After deforming beyond a certain threshold, rigidity and elasticity start to appear and the structure will show strong stiffness to the environmental stimuli; see Figure 1.2(b) and (c). The macroscopic behaviors indicated above show DIMs are very particular, biphasic materials that can be applied



**Figure 1.2 – Physical prototypes of DIM deform under different states.** (a) The structure exhibits extreme flexibility. (b) A DIM shows stiff bending response to resist self-weight under gravity. (c) The chainmail structure presents a strong in-plane stiffness to the external deformation.

to diverse potential applications, e.g., robotics, orthotics, sportswear, and many other areas of application. To support such applications, however, we must be able to predict the macro-mechanical behavior of DIM from their native-scale structure, i.e., element shape and connectivity. This relation is highly non-trivial as small changes in element shape can change the macromechanical behavior substantially. To address this problem, we first developed an efficient anisotropic strain-limiting method to replicate the mesoscopic biphasic behavior of DIM instead of simulating expensive microscale structures in Chapter 3. Additionally, to comprehensively characterize strong anisotropy, complex coupling, nonlinear kinematic deformation limits, and other nonlinear phenomena of DIM, we introduce a novel homogenization approach for distilling macroscale deformation limits from microscale simulations and developed a data-driven macromechanical model for simulating DIM with homogenized deformation constraints in Chapter 4.

## Chapter 2

# A Harmonic Balance Approach for Designing Compliant Mechanical Systems with Nonlinear Periodic Motions

by

Pengbin Tang<sup>1</sup>, Jonas Zehnder<sup>1</sup>, Stelian Coros<sup>2</sup>, and Bernhard Thomaszewski<sup>1,2</sup>


<sup>(1)</sup> Université de Montréal

<sup>(2)</sup> ETH Zürich

### Publication

The article has been published in ACM Transactions on Graphics (TOG), 39(6), (December 2020). It was presented at the ACM SIGGRAPH Asia 2020 conference. The article was reformatted to fit the style of this thesis.

Pengbin Tang contributed the theory, numerical experiments, and physical prototypes and wrote a draft of the paper. Jonas Zehnder contributed parts of the inverse design and the experiment of the animatronic wire character. Stelian Coros and Bernhard Thomaszewski were supervising the whole process and edited the paper.



**Figure 2.1** – Our method enables optimization-driven design of compliant mechanical systems with periodic large-amplitude motions. For this pair of dragon wings, the initial design (*a*, *b*) exhibits only small oscillation response when driven by harmonic forcing at a frequency of  $2.5\text{Hz}$ . Our approach automatically finds optimized design parameters (extra masses at the trailing edge of the wing) that lead to substantially increased amplitude (*c*, *d*).

**ABSTRACT.** We present a computational method for designing compliant mechanical systems that exhibit large-amplitude oscillations. The technical core of our approach is an optimization-driven design tool that combines sensitivity analysis for optimization with the Harmonic Balance Method for simulation. By establishing dynamic force equilibrium in the frequency domain, our formulation avoids the major limitations of existing alternatives: it handles nonlinear forces, side-steps any transient process, and automatically produces periodic solutions. We introduce design objectives for amplitude optimization and trajectory matching that enable intuitive high-level authoring of large-amplitude motions. Our method can be applied to many types of mechanical systems, which we demonstrate through a set of examples involving compliant mechanisms, flexible rod networks, elastic thin shell models, and multi-material solids. We further validate our approach by manufacturing and evaluating several physical prototypes.

**Keywords:** Nonlinear Vibration, Dynamic Motion Design, Harmonic Balance Method, Mechanical Systems, Periodic Motions, Nonlinear Modes.

## 2.1. Introduction

From the seismic response of high-rise buildings, to the aeroelastic stability of turbine blades, and to the micro-vibrations of energy harvesting devices—understanding and controlling the behavior of mechanical systems subject to periodic forcing is key to many engineering applications. Designing for vibrations typically means bracing against *resonance*, i.e., the strong increase in oscillation amplitude that occurs when driving a system at its characteristic frequency. Unless explicitly prevented, resonance can induce increasingly large deformations that, ultimately, lead to failure. While avoiding resonance is therefore often a main design goal, in this paper we explore the design of flexible structures that exploit resonance to produce periodic motion in the form of large-amplitude oscillations.

**Modeling Periodic Motion.** Most real-world systems will reach a stable steady state motion—a so called limit cycle—when driven by a periodic force. However, before reaching its limit cycle, the system goes through a transient process of non-periodic motion whose

length depends on the complexity as well as the elastic and viscous properties of the system. One approach to computing limit cycles is to simulate this transient process by numerically solving the equations of motion until some periodicity condition is met. However, using such a time-domain approach as a basis for computing system parameters that yield desired steady state behavior would require inverting a transient process of *a priori* unknown length, which is computationally all but intractable.

**Frequency-Domain Approaches.** A more promising approach for designing limit cycles is to use frequency-domain methods that directly solve for the system’s steady state behavior without the need for simulating the transient process. As another advantage, periodicity is obtained by construction and does not have to be enforced explicitly. Arguably the most widely used frequency-domain method for vibration analysis is Linear Modal Analysis (LMA) [Shabana, 1990]: based on the assumption of small-amplitude oscillations, the equations of motion are linearized around the origin and transformed to frequency space using generalized Eigenvalue decomposition. The resulting modal equations decouple and can be solved efficiently to yield the approximate limit response of the system to excitation at arbitrary frequencies. LMA is particularly well suited when vibrations and resonance are to be avoided and, hence, the assumption of small oscillations is valid. For compliant mechanisms and other systems with large-amplitude oscillations and finite rotations, however, nonlinearities play a critical role and LMA is unable to predict the steady state behavior in such cases.

To embrace nonlinearities from the start, we build our approach on the basis of the Harmonic Balance Method (HBM)—a nonlinear frequency-domain method that extends to a wide range of vibration problems [Krack and Gross, 2019]. The basic idea of HBM is to project the time-continuous equations of motion to a finite-dimensional subspace spanned by a small number of trigonometric basis functions. These nonlinear basis functions allow for efficient modeling of nonlinearities in nodal positions, velocities, and forces.

**Overview & Contributions.** Using HBM as a basis, we propose an optimization-based design tool that leverages sensitivity analysis to automatically discover design parameters that best approximate target steady-state behavior. Our method is able to control steady-state motion based on high-level user input by adjusting mass distribution, stiffness, or shape parameters of the designs. By performing simulation and design optimization directly in frequency space, our method avoids long transition times, the need for periodicity constraints, and other difficulties associated with time-domain approaches. Our method is furthermore general with respect to mechanical models, and we present examples of planar mechanisms augmented with elastic elements, rod networks, thin shells, and multi-material solids. We demonstrate the capabilities of our method on a set of simulation examples and real-world prototypes that include mechanical legs, compliant mechanisms, and animatronic characters.

## 2.2. Related Work

This work aims at designing real-world mechanical systems that exhibit desired large-amplitude oscillations. To our knowledge, this exact problem has not been considered before, but there are close ties to several sub-fields of visual computing and engineering.

**Designing Mechanical Motion.** The problem of designing mechanisms that exhibit desired kinematics has received considerable attention from the visual computing community in recent years [Coros et al., 2013a; Thomaszewski et al., 2014a; Bächer et al., 2015; Zhang et al., 2017]. Perhaps closest to our setting is the work by Ceylan et al. [2013a] who optimize for periodic motion of their characters using a frequency-space formulation. However, their method is purely kinematic and does not consider dynamics. Megaro et al. [2017] and Takahashi et al. [2019] proposed optimization-based design tools for mechanisms that are augmented with or partly replaced by compliant elements. While these examples go beyond pure kinematics, they did not consider dynamics. As one particular application, our method can be seen as a continuation of these previous works toward the design of compliant mechanisms that exhibit dynamic, periodic motion.

Few works from the graphics community on fabrication-oriented design have considered dynamics so far. Notable exceptions include the work by Chen et al. [2017], who propose a dedicated coarsening approach for optimization-based design of flexible objects with dynamic motion and contact. Another example is the work by Bächer et al. [2014], who design the mass distribution of rigid bodies in order to obtain sustained rotational motion. However, neither of these works considers periodic motion, forced vibrations, or resonance.

The work by Hoshyari et al. [2019] on motion control for robotic characters is related to our effort in that it predicts oscillations induced by external forcing. However, whereas their goal is to suppress vibrations by optimizing actuation parameters, we aim at generating large-amplitude oscillations by optimizing for shape, mass, and material parameters.

In order to create real-world animations of deformable characters, Skouras et al. [2013a] optimize for multi-material distributions and the parameters of a string-based actuation system. While we also leverage heterogeneous material distributions for motion control, our actuation principle is based on resonance induced by harmonic forcing.

**Modal Subspaces for Animation.** The idea of using modal subspaces for efficient animation of flexible models in graphics goes back to the early work of Pentland and Williams [1989]. Due to the inherent inefficiency of linear modal bases to model large rotational displacements, many works have proposed more efficient alternatives using, e.g., modal derivatives [Barbič and James, 2005], higher-order derivative information [Hildebrandt et al., 2011], or rotation-strain coordinates [Pan et al., 2015]. Adaptive basis selection has also been considered [Hahn et al., 2014; Kim and James, 2009].



Compared to the problem of constructing efficient linear bases for large-deformation simulation, nonlinear bases have not received much attention so far. An exception is the work of Fulton et al. [2019], who use machine learning techniques to construct nonlinear subspaces from simulation data. But besides the fact that none of the above approaches aim at optimization, they are all time-domain methods, meaning that they must simulate the entire evolution leading up to the steady-state behavior. By building on the Harmonic Balance Method, our approach is able to sidestep the transient process and thus directly operate on the steady-state behavior.

**Audible Vibrations.** Simulating vibrations is also of importance for physics-based sound synthesis. For instance, Zheng et al. [2011] and Bonneel et al. [2008] used Linear Modal Analysis (LMA) to efficiently generate sound for rigid body impact. To increase richness and fidelity, nonlinear approaches have been investigated for vibrating thin-shell structures [Chadwick et al., 2009; Cirio et al., 2018].

In the context of fabrication-oriented design, LMA has been the predominant approach so far; see, e.g., the works of Umetani et al. [2010] and Bharaj et al. [2015] on metallophones. Another stream of work in this context has investigated the simulation [Allen and Raghuvanshi, 2015] and design [Umetani et al., 2016] of wind instruments and acoustic filters [Li et al., 2016]. While sound synthesis is concerned with high-frequency, small-amplitude vibrations, we consider the design of mechanical systems with low-frequency but large-amplitude oscillations.

**Nonlinear Vibrations in Engineering.** The study of nonlinear vibrations is fundamental for structural dynamics, fluid structure interaction, aeroelasticity, and many other fields of engineering. As an extension of LMA beyond the linear regime, nonlinear normal modes are often used to analyze nonlinear dynamical phenomena [Kerschen et al., 2009]. One approach to compute steady-state behavior is to use time-domain integration schemes for simulating the transient process of the system until stable, periodic motion is obtained. An alternative approach that avoids computing the entire transient process are shooting methods [Peeters et al., 2009] that simulate only a single cycle while optimizing initial conditions and period to obtain periodicity.

Instead of enforcing periodicity in the time domain, we leverage the Harmonic Balance Method (HBM) to directly compute nonlinear periodic motion in frequency space. HBM is a versatile approach for nonlinear vibration analysis and has been widely used for many applications including nonlinear circuits [Bandler et al., 1992], fluid dynamics [Hall et al., 2013], and nonlinear mechanical systems in general [Detroux et al., 2014]. Its basic principle is to represent motion as a truncated Fourier series composed of periodic, trigonometric functions with different frequencies and phase offsets. Unlike the shooting method which already requires optimization for computing steady-state behavior, HBM only requires the solution of a nonlinear root-finding problem. Moreover, once the steady-state behavior for

a given driving frequency has been obtained, solutions for different inputs can be computed efficiently using numerical continuation.

Besides its advantages for solving forward simulation problems, HBM holds great promise for optimization-based design automation in elasticity and aeroelasticity applications [Engels-Putzka et al., 2019]. Perhaps closest to our work, Dou and Jensen [2015] combine HBM with sensitivity analysis to suppress vibrations in a simple beam by minimizing the amplitude at resonance. While we follow a similar methodology, our goal is not to suppress but to amplify oscillations and to control large-amplitude motion. To this end, we introduce amplitude and trajectory objectives that we integrate with forward and inverse design tools based on sensitivity analysis. Together, these tools enable interactive design space exploration as well as fully-automated design optimization of compliant systems with low-frequency, large-amplitude oscillations.

## 2.3. Theory

Our optimization-based design tool builds on the Harmonic Balance Method, a frequency-domain method for simulating the steady-state behavior of nonlinear mechanical systems. We first lay out the parts of the theory as relevant to our setting and briefly describe the corresponding computational framework.

### 2.3.1. Equations of Motion in Frequency Space

To arrive at the frequency-domain formulation, we start with the canonical equations of motion for a forced dynamical system in the time domain,

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{D}\dot{\mathbf{x}} = \mathbf{f}_{\text{int}}(\mathbf{x}) + \mathbf{f}_{\text{ext}}(\mathbf{x}, \omega, t) , \quad (2.3.1)$$

where  $\mathbf{x}$ ,  $\dot{\mathbf{x}}$ ,  $\ddot{\mathbf{x}} \in \mathbb{R}^{3n}$  denote nodal displacement, velocity, and acceleration, respectively. Furthermore,  $\mathbf{f}_{\text{int}}(\mathbf{x})$  is the nonlinear internal force and  $\mathbf{f}_{\text{ext}}(\mathbf{x}, \omega, t)$  is the periodic external force with frequency  $\omega$ . Finally,  $\mathbf{M}$  denotes the mass matrix,  $\mathbf{D} = D_\alpha \mathbf{M} + D_\beta \mathbf{K}(\mathbf{x})$  is the Rayleigh damping matrix, and

$$\mathbf{K}(\mathbf{x}) = \frac{\partial \mathbf{f}_{\text{int}}(\mathbf{x})}{\partial \mathbf{x}}$$

is the tangential stiffness matrix. To simplify the subsequent derivations, we separate linear and nonlinear terms as

$$\mathbf{M}\ddot{\mathbf{x}} + \hat{\mathbf{D}}\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, \omega, t) , \quad (2.3.2)$$

where  $\hat{\mathbf{D}} = D_\alpha \mathbf{M}$  and nonlinear forces are summarized as

$$\mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, \omega, t) = \mathbf{f}_{\text{int}}(\mathbf{x}) + \mathbf{f}_{\text{ext}}(\mathbf{x}, \omega, t) - D_\beta \mathbf{K}(\mathbf{x})\dot{\mathbf{x}} . \quad (2.3.3)$$

**Discretization.** At the steady-state solution, positions and forces are periodic functions. We approximate these time-domain functions in frequency-space as finite Fourier series

$$\mathbf{x}(t) \approx \mathbf{c}_0^x + \sum_{k=1}^{N_H} \left( \mathbf{s}_k^x \sin(k\omega t) + \mathbf{c}_k^x \cos(k\omega t) \right) \quad (2.3.4)$$

$$\mathbf{f}(t) \approx \mathbf{c}_0^f + \sum_{k=1}^{N_H} \left( \mathbf{s}_k^f \sin(k\omega t) + \mathbf{c}_k^f \cos(k\omega t) \right) \quad (2.3.5)$$

truncated to the  $N_H$ -th harmonic. In the above expressions,  $\mathbf{s}_k^* \in \mathbb{R}^{3n}$  and  $\mathbf{c}_k^* \in \mathbb{R}^{3n}$  are the vectors of Fourier coefficients decorated with  $x$  and  $f$  superscripts for positions and forces, respectively. Velocities  $\mathbf{v}(t)$  follow through direct differentiation of (2.3.4) and, hence, require no additional coefficients. We gather position and force coefficients into vectors of size  $(2N_H + 1) \cdot 3n$  as

$$\mathbf{z} = [(\mathbf{c}_0^x)^T \quad (\mathbf{s}_1^x)^T \quad (\mathbf{c}_1^x)^T \quad \dots \quad (\mathbf{s}_{N_H}^x)^T \quad (\mathbf{c}_{N_H}^x)^T]^T, \quad (2.3.6)$$

$$\mathbf{b} = [(\mathbf{c}_0^f)^T \quad (\mathbf{s}_1^f)^T \quad (\mathbf{c}_1^f)^T \quad \dots \quad (\mathbf{s}_{N_H}^f)^T \quad (\mathbf{c}_{N_H}^f)^T]^T. \quad (2.3.7)$$

Since the force  $\mathbf{f}(t)$  in (2.3.3) is a nonlinear function of position and velocity, its Fourier coefficients  $\mathbf{b}$  are functions of the position coefficients  $\mathbf{z}$  and we write  $\mathbf{b} = \mathbf{b}(\mathbf{z}, \omega)$ .

Substituting (2.3.4) and (2.3.5) into (2.3.2) and balancing the harmonic terms with a Galerkin projection (see Appendix 2.A.1 for details) yields the following set of nonlinear equations in the frequency domain,

$$\mathbf{h}(\mathbf{z}, \omega) \equiv \mathbf{A}(\omega)\mathbf{z} - \mathbf{b}(\mathbf{z}, \omega) = 0, \quad (2.3.8)$$

where  $\mathbf{A} = \text{diag}(\mathbf{0}, \mathbf{A}_1, \dots, \mathbf{A}_j, \dots, \mathbf{A}_{N_H})$  is a square, block-diagonal matrix of dimension  $(2N_H + 1) \cdot 3n$  describing the linear dynamics of the system. Each of the  $6n$  blocks is defined as

$$\mathbf{A}_j = \begin{bmatrix} -(j\omega)^2 \mathbf{M} & -j\omega \hat{\mathbf{D}} \\ j\omega \hat{\mathbf{D}} & -(j\omega)^2 \mathbf{M} \end{bmatrix}. \quad (2.3.9)$$

Eq. (2.3.8) is the frequency-domain version of (2.3.1), projected to a finite-dimensional Fourier subspace. Once we find the root  $\mathbf{z}^*$  of (2.3.8), we can convert it to the corresponding time-domain motion  $\mathbf{x}^*$  using inverse Fourier transformation.

### 2.3.2. Evaluation of Nonlinear Forces and Derivatives

We solve the system of nonlinear equations (2.3.8) using Newton's method. In every iteration, the solver requires the evaluation of  $\mathbf{b}$  and the Jacobian  $\partial \mathbf{h} / \partial \mathbf{z}$  for a given driving frequency  $\omega$ . Since the elastic and viscous forces are nonlinear and often non-polynomial functions of position and velocity, expressing them directly in frequency space is difficult. An alternative approach is given by the alternating frequency/time-domain (AFT) technique [Cameron and Griffin, 1989]. Using discrete Fourier transforms (DFT), motion is first

converted from frequency space to the time domain where nonlinear forces are then evaluated and finally converted back to the frequency domain as

$$\mathbf{z} \xrightarrow{\text{DFT}^{-1}} [\mathbf{x}(t), \dot{\mathbf{x}}(t)] \rightarrow \mathbf{f}(\mathbf{x}, \dot{\mathbf{x}}, \omega, t) \xrightarrow{\text{DFT}} \mathbf{b}(\mathbf{z}) . \quad (2.3.10)$$

Let  $t_i = i\Delta t$ ,  $i = 1 \dots N$  denote uniformly distributed samples in the time domain with  $\Delta t = \frac{2\pi}{N\omega}$  and  $N \geq 2N_H + 1$ . We start by evaluating (2.3.4) at the  $N$  sample points to obtain the discrete time-domain trajectory  $\tilde{\mathbf{x}} \in \mathbb{R}^{N \cdot 3n}$  and corresponding sample velocities  $\tilde{\mathbf{v}}$ . Due to the linearity of inverse DFT, we have

$$\tilde{\mathbf{x}} = \Gamma_x \mathbf{z} , \quad \text{and} \quad \tilde{\mathbf{v}} = \Gamma_v \mathbf{z} , \quad (2.3.11)$$

where  $\Gamma_x$  and  $\Gamma_v$  are sparse linear operators of size  $3n \cdot N \times (2N_H + 1) \cdot 3n$ . We then compute nonlinear forces for each sample point according to the mechanical model (see Appendix 2.A.5) and store the result as  $\tilde{\mathbf{f}} = (\tilde{\mathbf{f}}_1^T, \dots, \tilde{\mathbf{f}}_N^T)^T$  where  $\tilde{\mathbf{f}}_i = \mathbf{f}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{v}}_i) \in \mathbb{R}^{3n}$  is the vector of nodal forces for sample  $t_i$ . The time-domain forces are transformed back to frequency space using DFT,

$$\mathbf{b} = \Gamma_f^{-1} \tilde{\mathbf{f}} , \quad (2.3.12)$$

where  $\Gamma_f^{-1}$  is again a sparse linear operator; see Appendix 2.A.2. With these transformations, the Jacobian of (2.3.8) is obtained as


$$\frac{\partial \mathbf{h}}{\partial \mathbf{z}} = \mathbf{A} - \frac{\partial \mathbf{b}}{\partial \tilde{\mathbf{f}}} \left( \frac{\partial \tilde{\mathbf{f}}}{\partial \tilde{\mathbf{x}}} \frac{\partial \tilde{\mathbf{x}}}{\partial \mathbf{z}} + \frac{\partial \tilde{\mathbf{f}}}{\partial \tilde{\mathbf{v}}} \frac{\partial \tilde{\mathbf{v}}}{\partial \mathbf{z}} \right) = \mathbf{A} - \Gamma_f^{-1} \left( \frac{\partial \tilde{\mathbf{f}}}{\partial \tilde{\mathbf{x}}} \Gamma_x + \frac{\partial \tilde{\mathbf{f}}}{\partial \tilde{\mathbf{v}}} \Gamma_v \right) \quad (2.3.13)$$

where  $\partial \tilde{\mathbf{f}} / \partial \tilde{\mathbf{x}}$  and  $\partial \tilde{\mathbf{f}} / \partial \tilde{\mathbf{v}}$  can be computed analytically from (2.3.3).

### 2.3.3. Frequency Response Curves and Continuation

With the Jacobian (2.3.13) of the governing equations in hand, we can use Newton's method to compute the solution for a given driving frequency. In order to characterize the behavior of nonlinear mechanical systems, however, it is usually necessary to compute solutions for a range of driving frequencies, leading to so-called frequency response curves that plot amplitude as a function of driving frequency; see Fig. 2.2 for an example. Frequency response curves summarize important characteristics of nonlinear mechanical oscillators in a compact form. In particular, they reveal the number, location, amplitude, and sharpness of resonance peaks—quantities that are of interest for analysis as well as design. We start by introducing our measure of amplitude, then proceed to the conditions that characterize resonance, and finally explain how to compute frequency response curves numerically using continuation.

**Amplitude.** To quantify motion magnitude, and to encourage large-amplitude oscillation during design optimization, we must first develop a general definition of amplitude. For the one-dimensional case, amplitude is defined as the maximum displacement over a period



**Figure 2.2** – Frequency response curve for a thin shell model. Amplitude is measured using the trajectory of a selected vertex shown in *green*. The inset figures illustrate maximum-deflection configurations at 9 rad/s, 11 rad/s and 13 rad/s, respectively.

of oscillation. While seemingly simple, this concept does not readily generalize to higher dimensions and it translates into computational difficulties: maximizing displacement means solving for zeroes of the velocity function, which is a trigonometric root-finding problem with potentially many local extrema. To avoid these difficulties, we instead quantify the magnitude of motion for a given vertex  $\mathbf{x}_i$  as the distance traveled within one period,

$$A_i(\mathbf{z}) = \int_0^T \|\mathbf{v}_i\| dt = \int_0^T \sqrt{(v_i^x)^2 + (v_i^y)^2 + (v_i^z)^2} dt , \quad (2.3.14)$$

where

$$v_i^j = \sum_{k=1}^{N_H} \mathbf{s}_k^{x_j} k\omega \cos(k\omega t) - \mathbf{c}_k^{x_j} k\omega \sin(k\omega t) . \quad (2.3.15)$$

While this expression could be evaluated in frequency space through quadrature, we simply compute the length of the piece-wise linear trajectory in the time domain using the AFT scheme described above.

**Resonance.** Having established a way of quantifying motion magnitude, we can now make precise the conditions for resonance as a local maximum of amplitude with respect to the driving frequency,

$$\omega_{\text{res}} = \arg \max_{\omega} A(\mathbf{z}(\omega)) , \quad (2.3.16)$$

where  $\omega_{\text{res}}$  is the corresponding resonance frequency and  $\mathbf{z}(\omega)$  are steady-state Fourier coefficients expressed as a function of the driving frequency. The map between  $\mathbf{z}$  and  $\omega$  is implicitly given through the dynamic equilibrium condition  $\mathbf{h}(\mathbf{z}, \omega) = \mathbf{0}$ . As a necessary

condition for resonance, we require that

$$\frac{dA(\mathbf{z})}{d\omega} = \frac{\partial A(\mathbf{z})}{\partial \mathbf{z}} \frac{d\mathbf{z}}{d\omega} = \mathbf{0} . \quad (2.3.17)$$

**Continuation.** To avoid solving (2.3.8) from scratch every time when computing frequency response curves, we use numerical continuation to *follow* the path of the solution while changing the input frequency. Each continuation step consists of a prediction step  $\mathbf{t}_i$  for updating the current state  $[\mathbf{z}_i^T, \omega_i]^T$ , followed by a projection step that enforces  $\mathbf{h}(\mathbf{z}, \omega) = 0$ . The prediction step requires the derivative of (2.3.8) with respect to both the state  $\mathbf{z}$  and the driving frequency  $\omega$ . The latter is obtained as

$$\mathbf{h}_\omega = \frac{\partial \mathbf{h}}{\partial \omega} = \frac{\partial \mathbf{A}}{\partial \omega} \mathbf{z} - \frac{\partial \mathbf{b}}{\partial \tilde{\mathbf{f}}} \frac{\partial \tilde{\mathbf{f}}}{\partial \omega} = \frac{\partial \mathbf{A}}{\partial \omega} \mathbf{z} - \Gamma_f^{-1} \frac{\partial \tilde{\mathbf{f}}}{\partial \omega} \quad (2.3.18)$$

where  $\frac{\partial \tilde{\mathbf{f}}}{\partial \omega} = -\mathbf{D}_\beta \mathbf{K}(\mathbf{x}) \frac{\partial \tilde{\mathbf{x}}}{\partial \omega} = -\mathbf{D}_\beta \mathbf{K}(\mathbf{x}) \frac{\Gamma_v \mathbf{z}}{\omega}$ . Note that the derivative of the external force with respect to  $\omega$  is zero; cf. the derivation of  $\Gamma$  and  $\Gamma^{-1}$  in Appendix 2.A.2. The prediction step  $\mathbf{t}_i$ , tangent to the response curve at  $[\mathbf{z}_i, \omega_i]^T$ , is then determined as

$$\begin{bmatrix} \mathbf{h}_\mathbf{z} & \mathbf{h}_\omega \\ \mathbf{t}_{i-1}^T & \end{bmatrix} \mathbf{t}_i = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix} , \quad (2.3.19)$$

where  $\mathbf{h}_\mathbf{z}$  is a shorthand for (2.3.13). The first line in the above equation asks that the step should maintain force balance to first order. The second condition requires the prediction step to have a positive dot product with the previous step, thus preventing the continuation scheme from going backwards. The resulting prediction step is then used in conjunction with an arc-length control strategy [Seydel, 2009] to compute an initial guess for Newton's method.

To reliably track down resonance peaks, we use condition (2.3.17) and monitor the sign of the gradient  $\frac{dA}{d\omega}$ . Whenever the sign changes from negative to positive between two samples, we compute the exact resonance frequency  $\omega_{\text{res}}$  by solving (2.3.17). The Jacobian matrix  $\frac{d\mathbf{z}}{d\omega}$  required for this procedure can be computed through sensitivity analysis on  $\mathbf{h}(\mathbf{z}, \omega) = \mathbf{0}$ .

## 2.4. Computational Design

Given an initial design, we would like to determine changes for parameters such as shape, mass distribution, and driving frequency such that the resulting steady-state motion best approximates a given target behavior. To this end, we consider two design approaches: user-driven forward exploration of the design space and optimization-driven inverse design.

### 2.4.1. Dynamical Equilibrium and Sensitivity

For both forward and inverse design, we must be able to predict the change in steady-state behavior induced by a given change in design parameters. To this end, we leverage

the equilibrium constraints as an implicit map between parameters and state: the Fourier coefficients  $\mathbf{z}$  must be a dynamic equilibrium configuration for the design parameters  $\mathbf{p}$ . We make this explicit by rewriting (2.3.8) as

$$\mathbf{h}(\mathbf{z}, \mathbf{p}, \omega) = \mathbf{A}(\mathbf{p}, \omega) \mathbf{z} - \mathbf{b}(\mathbf{p}, \mathbf{z}, \omega) = \mathbf{0} . \quad (2.4.1)$$

Since this relation must hold for every admissible choice of  $\mathbf{p}$ , the Fourier coefficients effectively become a function of the design parameters, i.e.,  $\mathbf{z} = \mathbf{z}(\mathbf{p})$ . Moreover, any change to the parameters will entail a corresponding state change such that the system is again at equilibrium. More formally, we have

$$\frac{d\mathbf{h}}{d\mathbf{p}} = \frac{\partial \mathbf{h}}{\partial \mathbf{z}} \frac{d\mathbf{z}}{d\mathbf{p}} + \frac{\partial \mathbf{h}}{\partial \mathbf{p}} = \mathbf{0} , \quad (2.4.2)$$

from which we obtain the so-called design sensitivity matrix as

$$\mathbf{S} = \frac{d\mathbf{z}}{d\mathbf{p}} = -\frac{\partial \mathbf{h}^{-1}}{\partial \mathbf{z}} \frac{\partial \mathbf{h}}{\partial \mathbf{p}} . \quad (2.4.3)$$

The above equations form the basis of the forward and inverse design tools that we describe next.

**Rank of Constraint Jacobian.** Eq. (2.4.1) provides exactly as many constraints as there are degrees of freedom. In order for the constraint Jacobian in (2.4.3) to be invertible, the constraint gradients must be linearly independent. Even if the Jacobian is non-singular, numerical problems can still arise if the matrix is indefinite. While the conditions on rank and definiteness cannot be guaranteed for all configurations of nonlinear oscillators, we are only interested in stable steady-state solutions, which fulfill these properties by definition. If transient rank-deficiency or indefiniteness are still encountered (manifesting through failure of the LU solver or a large residual of the linear system), we apply adaptive diagonal regularization until a valid solution is found.

## 2.4.2. Forward Sensitivity Exploration

The sensitivity matrix provides an efficient tool for interactive exploration of the design space. For a given initial design  $\mathbf{p}$  and corresponding state  $\mathbf{z}$ , we compute a first-order prediction for the new equilibrium state as

$$\mathbf{z}_p = \mathbf{z} + \frac{d\mathbf{z}}{d\mathbf{p}} \Delta \mathbf{p} . \quad (2.4.4)$$

Once the sensitivity matrix is computed, this prediction is instantaneous and thus enables interactive exploration of the design space around a given set of parameters. For larger parameter changes, however, the first-order prediction can become inaccurate, requiring a full, nonlinear update. In our interface, the user can manually issue such update commands,

which entail re-simulation with the new parameters and re-computation of the sensitivity matrix.

Besides changing parameters individually, we found it useful to provide additional compound variables that change multiple parameters simultaneously along the gradient of selected design objectives—see Sec. 2.5.3 for an example.

### 2.4.3. Inverse Design

Even when aided by sensitivity information, forward design based on manual exploration soon becomes unattractive as the number of parameters increases. If design goals can be quantified, automatic parameter optimization can be a very efficient and effective alternative. With the implicit relation between  $\mathbf{z}$  and  $\mathbf{p}$  defined as above, we can cast design optimization as an unconstrained minimization problem, where we aim to minimize an objective function  $f(\mathbf{z}(\mathbf{p}), \mathbf{p})$  that encodes various design goals as described below. Using Eq. (2.4.2), we obtain the objective gradient as

$$\frac{df}{d\mathbf{p}} = \frac{d\mathbf{z}^T}{d\mathbf{p}} \frac{\partial f}{\partial \mathbf{z}} + \frac{\partial f}{\partial \mathbf{p}} = -\frac{\partial \mathbf{h}^T}{\partial \mathbf{p}} \frac{\partial \mathbf{h}^{-T}}{\partial \mathbf{z}} \frac{\partial f}{\partial \mathbf{z}} + \frac{\partial f}{\partial \mathbf{p}}. \quad (2.4.5)$$

It is worth noting that this expression can be rearranged such that only a single linear system needs to be solved. We use this gradient in combination with L-BFGS-B [Byrd et al., 1995] to find parameters that minimize the objective function while enforcing bound constraints where applicable, e.g., positivity for mass values.

As we show in Sec. 2.5, HBM-based simulation and frequency-space sensitivity analysis for optimization combine into an efficient tool for inverse design. In order for this tool to be effective, however, we must define objective functions that allow users to express their design intents.

**Trajectory Matching.** An intuitive and direct way of specifying motion goals is by prescribing target trajectories that selected nodes should track. Given a time-domain target trajectory for a given vertex  $k$  as input, we first compute the corresponding Fourier coefficients  $\hat{\mathbf{z}} = (\hat{\mathbf{s}}, \hat{\mathbf{c}})$  using DFT. We then measure the distance between current and target trajectories in frequency space as

$$f_{\text{Dist}}(\mathbf{z}) = \sum_{i=1}^{N_H} \sum_{j=1}^3 \left( (\mathbf{s}_i^{3k+j} - \hat{\mathbf{s}}_i^j)^2 + (\mathbf{c}_i^{3k+j} - \hat{\mathbf{c}}_i^j)^2 \right), \quad (2.4.6)$$

where the superscript selects individual components from the coefficient vectors. The objective in its above form assumes that both target and actual trajectory have the same phase offset, but this is generally not the case. To eliminate phase dependence, we evaluate (2.4.6) for different phase offsets and define the final objective value as the smooth minimum over



the individual distances, i.e.,

$$f_{\text{Track}}(\mathbf{z}) = \frac{\sum_{i=0}^{N-1} f_{\text{Dist}}(\mathbf{z}, \phi_i) e^{-\alpha f_{\text{Dist}}(\mathbf{z}, \phi_i)}}{\sum_{i=0}^{N-1} e^{-\alpha f_{\text{Dist}}(\mathbf{z}, \phi_i)}} , \quad (2.4.7)$$

where  $\alpha$  is a parameter controlling the sharpness of the smooth minimum function,  $\phi_i = i \frac{2\pi}{N}$  are phase offsets, and  $f_{\text{Dist}}(\mathbf{z}, \phi)$  is the same as in Eq. (2.4.6) but with the Fourier coefficients of the target trajectory shifted by a phase offset  $\phi$ .

**Amplitude.** Although the trajectory matching objective affords some degree of amplitude control, we have seen in our experiments that it can be difficult to obtain large amplitudes in this way—arguably because the tracking objective is biased towards a specific motion, rather than more general, large-amplitude oscillations. We therefore introduce an objective that explicitly aims to maximize the magnitude of motion for user-selected nodes. Based on the amplitude definition in Sec. 2.3.3, we define an amplitude objective for vertex  $k$  as

$$f_{\text{Ampl}}(\mathbf{z}) = (A_k(\mathbf{z}) - \bar{A}_k)^2 \quad (2.4.8)$$

where  $\bar{A}_k$  is the target amplitude. Note that, besides optimizing for specific, reachable values for the amplitude, we can also simply encourage amplitude maximization by setting  $\bar{A}$  to an arbitrarily large value. Likewise, with a straightforward modification of (2.3.14), it is possible to optimize amplitude only along selected dimensions.


## 2.5. Results

We evaluate our method on a set of examples that highlight the impact of our design objectives both in simulation and on actual physical prototypes. We start by analyzing the accuracy of HBM compared to ground truth time-domain simulations as well as a linear frequency-space approach.

### 2.5.1. Analysis & Validation

**Comparison with Newmark.** One of the central advantages of HBM is that it directly yields periodic steady-state solutions without having to simulate the transient process of the system. In order to analyze the accuracy of HBM in our setting, we compare to ground truth simulations obtained using the Newmark integration scheme, a time-domain method that is widely used for structural dynamics and vibration analysis in general; see section 2.A.3 for details.

To this end, we consider a fork-shaped thin shell oscillator (188 elements, 144 nodes) with one end driven by a periodic force while the other two extremities are vibrating freely as shown in Fig. 2.2. To measure the difference in steady state solution for Newmark and HBM, we must first define an appropriate convergence criterion for the time-domain method.



**Figure 2.3** – Trajectory error for HBM compared to ground truth time-domain simulations. The plot shows error as a function of the number of harmonics using damping coefficients as indicated.


We measure the difference in trajectories between two successive periods as

$$e_p = \sum_{i=0}^{N_{NM}-1} \|\mathbf{x}(t_p N_{NM} + i) - \mathbf{x}(t_{(p+1)} N_{NM} + i)\|, \quad (2.5.1)$$

where the subscripts of  $t_*$  refer to time-step indices,  $p$  denotes the index of the period, and  $N_{NM} = \frac{2\pi}{\omega\Delta t}$  is the number of time steps used to simulate a single period.

Using this error metric, Fig. 2.3 shows accuracy plots for different driving frequencies and different numbers of harmonics. We use  $N_{NM} = 4096$  for computing the step size and run the Newmark simulation until the trajectory difference between two successive periods satisfies  $e_p < 1e^{-7}$ . It can be seen that the error is larger for driving frequencies near resonance, which is at  $\omega_{\text{res}} = 11.0$  rad/s for this example (see Fig. 2.2). This observation is explained by the fact that near-resonance frequencies lead to larger amplitude motion. Nevertheless, using  $N_H = 5$  harmonics leads to sufficiently good accuracy, and more terms yield virtually no improvements. Additional analysis is given in Appendix 2.A.4.

We furthermore conducted an experiment to analyze the impact of the number of samples used to evaluate the nonlinear forces in the time-domain using Eq. (2.3.3) for HBM. While a lower bound is given by the Nyquist limit, as can be seen from Fig. 2.4, the effect of using more samples is almost imperceptible. This visual impression is confirmed by an additional quantitative analysis, showing that the difference in Fourier coefficients is less than  $1e^{-10}$  in this case.




**Figure 2.4** – Trajectory comparison between fabricated L-wing and HBM simulations at driving frequency of  $2.0Hz$  with different numbers of sampling points  $N_{AFT}$ : physical prototype (a) and HBM simulation using (b)  $N_{AFT} = 128$ , (c)  $N_{AFT} = 64$ , and (d)  $N_{AFT} = 32$ .

To summarize our analysis of HBM and comparison with Newmark, we can conclude that, already with a small number of harmonics and time-domain samples, HBM accurately captures the nonlinear large-amplitude oscillation behavior that is the focus of this work. We note that, when using  $N_H = 5$  and  $N_{AFT} = 64$  for this example, HBM computes steady-state solutions one to two orders of magnitude faster than time-domain methods. For the general case, determining the *optimal* number of harmonics a priori is difficult. However, we found the following strategy to work well in practice: we first simulate using a small number of harmonics (e.g.  $N_H = 3$ ), which we increase until the trajectory error between two successive runs falls below a given threshold value. While this threshold needs to be set by the user, setting it to a small fraction of the trajectory length (e.g.  $1e^{-3}$ ) simplifies this task.

Before we present examples obtained using our optimization-based design method, we briefly comment on an alternative frequency-domain method.

**Comparison with Linear Modal Analysis.** In order to illustrate the importance of incorporating nonlinearities in the simulation model, we compare to Linear Modal Analysis (LMA), a frequency-space approach based on a linearization of modal dynamics around the rest state. As can be seen in Fig. 2.5, this first-order approximation leads to significant in-plane distortions. This is not surprising, as methods based on linearized deformation measures are known to introduce artifacts for rotational displacements. Perhaps more severe, however, is the fact that neither the resonance frequency nor the motion at resonance are captured with acceptable accuracy. These shortcomings effectively disqualify linear modal analysis as a basis for design and optimization in the large-amplitude setting.




**Figure 2.5** – Frequency responses computed with LMA and HBM for different driving frequencies with damping coefficients  $D_\alpha = 0.5$  and  $D_\beta = 0.00005$ . Trajectory of a selected tip point (green) and color-coded maximum in-plane stretch. For LMA, the maximum strains over one period for the three driving frequencies are 0.655, 0.582, and 0.575. The corresponding values for HBM ( $8.25e^{-5}$ ,  $1.00e^{-4}$ , and  $8.50e^{-5}$ ) are 4-5 orders of magnitude smaller.

### 2.5.2. Measuring Damping Parameters


Damping parameters play an important role for the dynamics of a mechanical system. To experimentally estimate damping coefficients for a given design task, we first choose a simple real-world specimen, e.g., the L-wing (see Fig. 2.4) for the thin shell model and the three-link mechanism (see Fig. 2.8) for compliant mechanisms. We then determine damping parameters such that the simulated steady-state motion is as close as possible to the corresponding real-world motion. Real-world trajectories are captured using off-the-shelf cameras for side and back views. We then extract trajectories for selected key points (such as wing tips) and use them to fit damping coefficients for simulation.

### 2.5.3. Forward Design with Sensitivity Exploration

We demonstrate the sensitivity-based forward design approach described in Sec. 2.4.2 on the mechanical character shown in Fig. 2.6. This character consists of 10 links connected through 6 joints and 8 springs. During interactive design exploration, the goal for the user is to obtain an understanding of the design space and to discover a large-amplitude motion that makes for an appealing animation. We drive the character by applying harmonic forcing




**Figure 2.6** – Forward Design with Sensitivity Exploration illustrated on an animatronic character. The initial design (*left*) exhibits only small oscillations at the hands (indicated in *green*). After several steps of forward exploration, the final design exhibits an expressive large-amplitude motion (*right*).



**Figure 2.7** – Amplitude of the wing tip before (*blue*) and after (*orange*) optimization for the dragon example. The dashed line indicates the driving frequency ( $2.5Hz$ ) used during optimization.

to its feet, which we initially choose to act in the vertical direction and in-phase. The design parameters for this exploration are extra masses for each node as well as the phase offset and amplitude for the forcing. In addition to changing each parameter individually, we add controls to the interface that change parameters simultaneously along the gradient of the




**Figure 2.8 – Trajectory matching for the three-link compliant mechanism.** Initial design (*left*) and optimized design (*right*) with simulated (*top*) and real-world (*bottom*) end-effector trajectories shown in green. The target trajectory is shown in *blue*.

amplitude objective for selected nodes. This enables convenient exploration along parameter-space directions that lead to large motion amplification.


Having computed the initial frequency-response curve and the sensitivity matrix, the user starts exploring motion variations by changing design parameters. Our system then provides instant feedback on the predicted change in motion. As best seen in the accompanying video, this approach allows the user to quickly converge towards a large-amplitude motion. Re-simulation with the new parameters exhibits only little deviation from the first-order prediction shown during sensitivity exploration.

#### 2.5.4. Optimization-Based Inverse Design

We demonstrate our optimization-driven approach on a set of examples that illustrate applications to a diverse range of mechanical models and highlight the impact of our design objectives. For validation, we manufacture physical prototypes for several examples and evaluate their performance.



**Figure 2.9 – Running Ostrich.** Four images from a running sequence of our ostrich model with two legs driven at  $1.2Hz$  with a phase offset of half a period.



**Figure 2.10 – Ostrich leg.** Performance of initial (*left*) design and optimized design (*right*) in simulation (*top*) and on the physical prototype (*bottom*). Target and actual trajectories are shown in *blue* and *green*, respectively.

**Dragon Wings.** Our first example is a dragon model obtained from a shape repository<sup>1</sup> for which we aim to create large amplitude oscillations for the wings such as to suggest flapping flight. The body of the dragon is kept static while its wings are driven with a servomotor that creates rotational motion with programmable frequency and an amplitude of 40 degrees; see Fig. 2.1. We model the wings using discrete shells [Grinspun et al., 2003a] and fit elastic material parameters as described in [Pabst et al., 2008] to match the behavior of an FDM-printed PLA cantilever, from which we obtain a Young’s modulus of  $3.3Gpa$  and a Poisson’s ratio of 0.36. With the elasticity coefficients determined, we experimentally set viscous parameters such as to minimize discrepancy between simulated and real-world motion using the setup shown in Fig. 2.4. We use as design parameters the masses of three points distributed along the trailing edge of the wing.

We start by computing the frequency response curve for the initial design, which shows only small-amplitude oscillations throughout the range of  $0.5 - 5.0Hz$  (blue curve in Fig. 2.7). We then employ our amplitude objective to determine design parameters that lead to motion amplification at  $2.5Hz$ . With an almost four-fold increase in amplitude, the optimized design exhibits greatly improved performance in simulation (orange curve in Fig. 2.7). This prediction is confirmed by the manufactured prototypes for the two designs, both of which show very good agreement with the corresponding HBM simulations. It should be noted that, rather than just amplifying the initial motion, the optimized trajectory at the wing tip is quite different from the original one. Whereas manually finding a sufficiently close target trajectory is a difficult task for the user in this case, our amplitude objective provides the freedom needed to automatically discover this large-amplitude motion.

**Compliant Three-Link Mechanism.** The ability to control large-amplitude oscillations for nonlinear mechanical systems enables new, efficient designs for robotics applications. We investigate the potential of this approach on two compliant mechanisms. The first design, shown in Fig. 2.8, is a simple three-segment chain augmented with two elastic springs. We use custom-made springs with stiffness coefficients of  $70.2N/m$  and  $16.0N/m$ , respectively, and corresponding rest lengths of  $0.146m$  and  $0.088m$ . We drive the hip joint with harmonic forcing in the vertical direction with an amplitude of  $1mm$  and a frequency of  $3Hz$  such as to mimic the footfall frequency of a fast quadrupedal walking gait [Moro et al., 2013]. To encourage walking-like motion, we prescribe a corresponding target trajectory for the bottom joint and optimize over joint masses, spring attachment points, and link lengths. As can be seen from Fig. 2.8 (*left*), the initial design shows little response. After optimization, however, the mechanism closely tracks the desired trajectory, thus converting a simple vertical input signal into a complex two-dimensional output motion. It is worth noting that the design changes found by our method are quite significant and leverage all available parameters.

---

1. <https://www.thingiverse.com/thing:2714125>




**Compliant Ostrich Leg.** In our second mechanism example we consider a complex ostrich leg inspired by the work of Cotton et al. [2012]. Our leg design consists of 5 bars connected through 5 revolute joints and 3 custom-made springs with stiffnesses of  $41.86N/m$ ,  $70.2N/m$ , and  $16.0N/m$  and corresponding rest lengths of  $0.108m$ ,  $0.146m$ , and  $0.088m$ . The leg is driven by a servomotor located at the hip joint that induces harmonic rotational oscillation in the attached link with a frequency of  $1.2Hz$  and amplitude of 40 degrees. To generate large-amplitude oscillations that approximate the characteristic running motion observed in ostriches, we prescribe target trajectories for the toe and ankle joints that we manually extracted from real-world video footage.

As can be seen from Fig. 2.10 (*left*), the initial design produces trajectories that correspond to simple, reciprocating motion along a one-dimensional curve. After optimization, however, the motion closely tracks the target trajectories for both ankle and toe (see Fig. 2.10 (*right*)). We optimize the mass for every joint to match the target trajectory during optimization. As best seen in the accompanying video, the optimized design also successfully reproduces the rotation of the toe between swing and stance phases that characterizes the real-world gait. We use the optimized design to build the physical ostrich model shown in Fig. 2.9, whose two legs are driven at a phase offset such as to mimic running motion.

**Animatronic Wire Character.** The HBM formulation extends to a large range of mechanical models and materials. Our optimization-driven approach is able to leverage this flexibility, which we demonstrate on two additional examples that use multi-material solids and elastic rods. Taking a result from Xu et al. [2018] as inspiration, we design a wire character in the form of a fish as illustrated in Fig. 2.11. We model this character using discrete elastic rods [Bergou et al., 2008] and the extension to networks described by Zehnder et al. [2016b]. In order to control the frequency response of the character, we add three extra weights to the model whose mass we optimize such as to maximize the amplitude at the tail for a driving frequency of  $2.0Hz$ . For the physical prototype, we use standard aluminum wire with a diameter of  $1.1mm$ , a Young’s modulus of  $69GPa$ , and density of  $2.7g/cm^3$ . We use brass weights customized according to the solution returned by the optimizer. Interestingly, the optimization completely removed the weight attached to the tip of the tail. To verify this somewhat counter-intuitive result, we experimented with manually-designed mass distributions with roughly the same total weight. As can be seen in the accompanying video, none of these alternatives is able to amplify the input motion, whereas the optimized design produces large-amplitude oscillation as predicted in simulation. We conclude that, even for seemingly simple cases, manually finding parameters that lead to large-amplitude oscillations at the designated driving frequency can be very challenging. Our optimization-based approach removes this burden from the user.

**Multi-material Solid.** We investigate applications of our approach to material optimization for viscoelastic solids undergoing nonlinear vibrations. Our setup consists of an inverted



**Figure 2.11** – For this animatronic wire character, we optimize the weights of three additional masses such as to achieve large-amplitude oscillation of its tail.

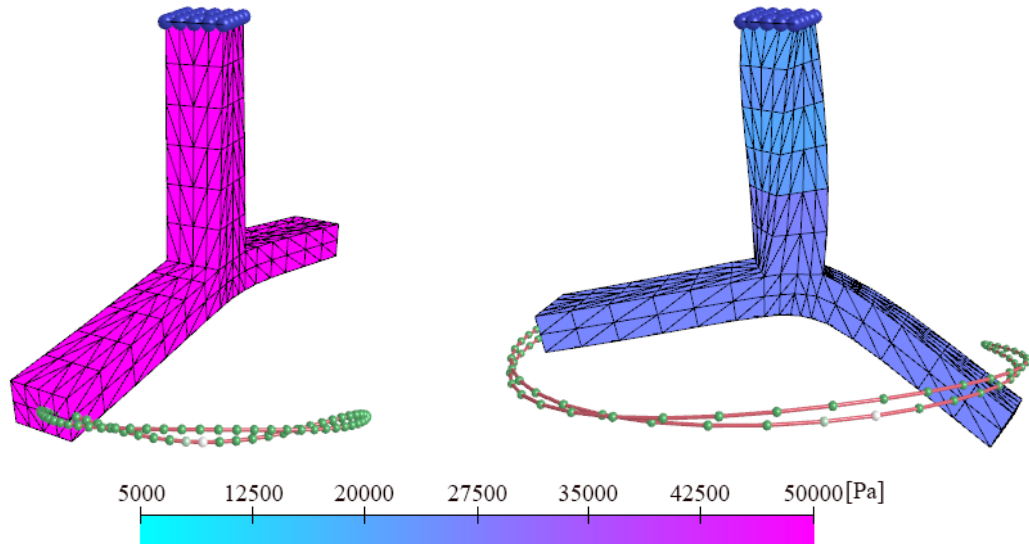
T-shaped model shown in Fig. 2.12, whose upper face we drive through harmonic rotational excitation. The solid is structured vertically into 7 layers of homogeneous material, each of which can have different viscoelastic properties. For the inverse design problem, we aim to optimize the Young’s modulus for each layer such that the amplitude of a selected node on the bottom extremity is maximized under a given driving frequency of  $2.0Hz$ . The per-layer material assignment found by the optimization successfully amplifies the motion by a factor of more than 4.

**Eiffel Tower.** Inspired by the work of [Skouras et al., 2013a], we consider a material optimization problem for the Eiffel tower model show in Fig. 2.13. We actuate the base of the model using harmonic driving in horizontal direction. We use 2D constant strain triangle elements with an St.Venant-Kirchhoff material for simulation, and optimize for per-element stiffness coefficients such as to maximize the amplitude of the top of the tower at  $2.0Hz$ . The optimized design exhibits an increase in amplitude by a factor of more than 50.

### 2.5.5. Statistics & Additional Validation

**Performance & Statistics.** All examples run on a machine with an Intel Core i9-7900X 3.3GHz processor and 32 GB of RAM. Statistics are given in Tab. 2.1.

**Constraint Violations.** For the simulation of compliant mechanisms, we use stiff penalty terms to enforce angular and distance constraints for rigid joints and links. To analyze the



**Figure 2.12** – We optimize per-layer material stiffness for this solid such as to maximize the amplitude of the selected vertex when driving the top face with harmonic rotational excitation at  $2.0Hz$ .

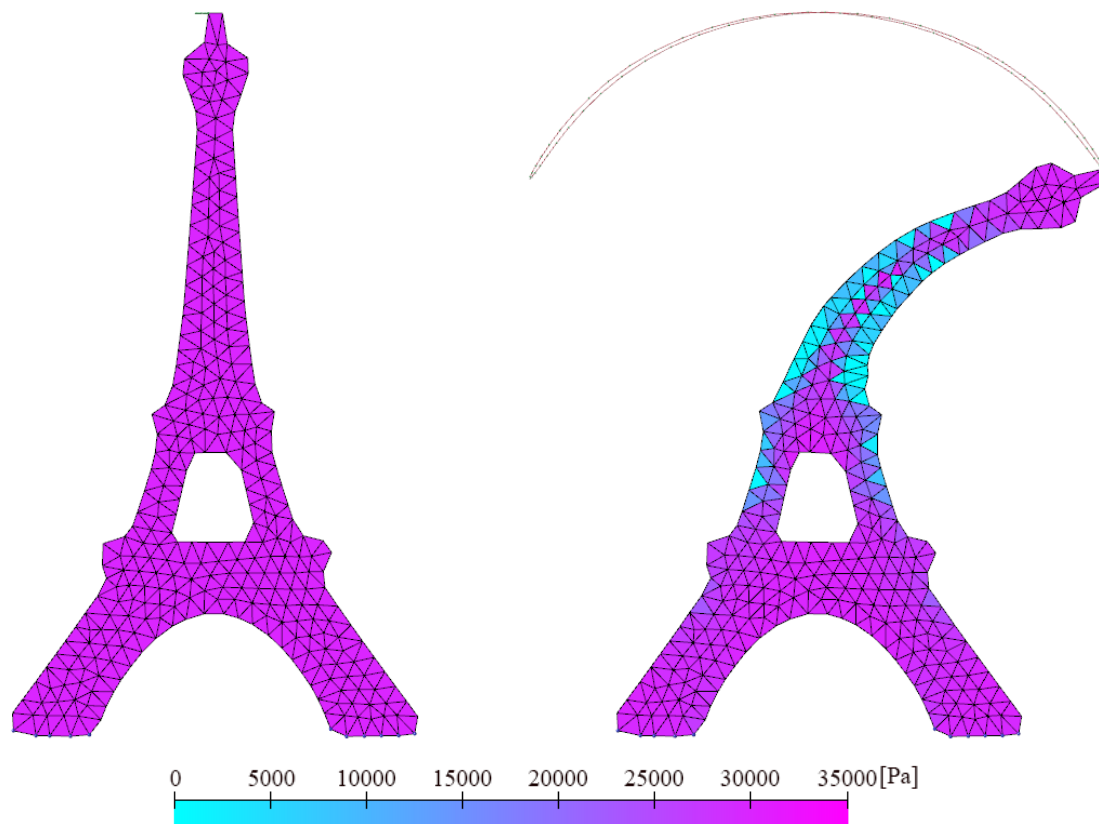
**Table 2.1 – Statistics for inverse design examples.** The columns list numbers of degrees of freedom ( $N_{DoF}$ ), harmonics ( $N_H$ ), sampling points ( $N_{AFT}$ ), parameters ( $N_p$ ), iterations required for convergence ( $N_{it}$ ), as well as the total time spent on optimization.

Example	$N_{DoF}$	$N_H$	$N_{AFT}$	$N_p$	$N_{it}$	time [s]
Dragon	780	8	64	3	12	1098.09
Three-Link Leg	8	8	64	16	146	34.89
Ostrich Leg	14	14	64	7	128	163.98
Fish	187	8	64	3	46	456.96
3D Solid	918	8	64	7	10	7926.43
Eiffel Tower	658	8	64	525	150	5534.82

validity of this approach, we monitored constraint violations, i.e., the change in angles of rigid joints and link lengths. We plot the corresponding maximum values as a function of time in Fig. 2.14, from which it can be seen that constraint violations are small for all three mechanism examples.

As another potential concern known from time-domain methods, stiff penalty terms can give rise to numerical damping when using lower-order implicit integration methods. While we cannot guarantee that our HBM simulations are free from this effect, they closely track the Newmark solutions which are known to exhibit very little numerical dissipation.

**Feasible Regions of Design Space.** Understanding and navigating the space of feasible dynamics for a given input model is crucial for successful design. Our inverse design tool can be used to answer the question whether a given motion is achievable, and one positive

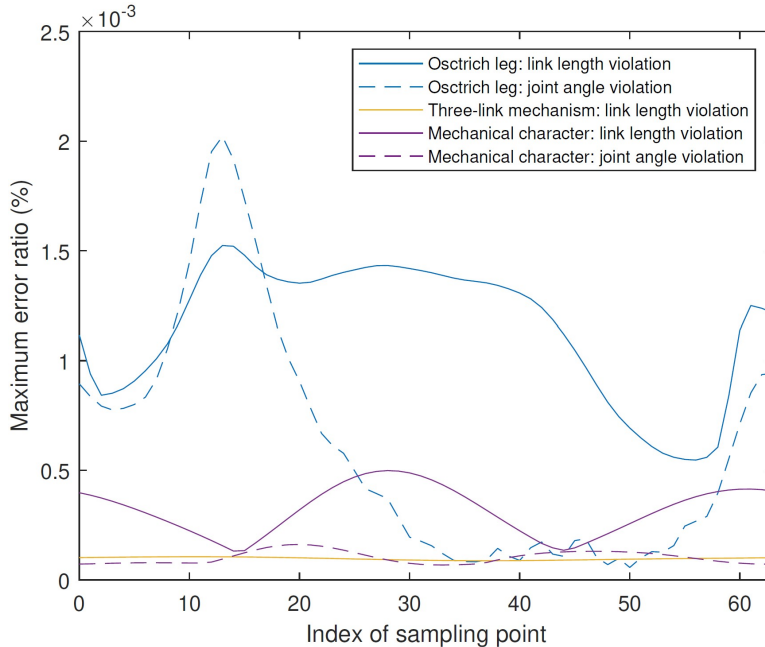


**Figure 2.13** – Material optimization on an Eiffel Tower model. For the initial design with homogeneous material (*left*), the amplitude at the tip is almost the same as for the driving signal. After optimizing for per-element stiffness values, the tip amplitude is substantially increased (*right*).

answer is often enough to fulfill the user’s intent. However, if the desired motion is infeasible, our method will return a design whose motion is, at least locally, as close as possible to the target behavior. If deviations are substantial, however, the solution computed in this way might not be subjectively optimal or even acceptable for the user. In such cases, finding good compromises requires further design space exploration. While our method does not offer an explicit representation of the space of feasible motions, if model complexity allows for it, our interactive tool provides an efficient way to explore the possibilities and limitations of a given input model.

## 2.6. Limitations & Future Work

We presented an optimization-driven frequency-space approach for designing mechanical systems that exhibit desired nonlinear oscillations. Our results indicate that Harmonic Balance paired with Sensitivity Analysis is indeed an efficient and effective combination that enables the construction of powerful forward and inverse design tools.



**Figure 2.14 – Maximum error of constraints violation.** We use a penalty stiffness of  $1e^7$  and  $N_{AFT} = 64$  time-domain samples for these three compliant mechanism examples. Each curve shows the maximum error (change in angle/length divided by corresponding original value for angular and distance constraints) over all constraints for each example within a period. It can be seen that constraint violations remain below  $2e^{-3}\%$  at all times for all cases.

There are several limitations of our method that we briefly discuss below along with other potential directions for future research. If constraints are present in the mechanical system, the truncated Fourier series will generally not satisfy them exactly. This would be the case, e.g., when modeling mechanisms as articulated multi-body systems. Nevertheless, to satisfy the requirements of a specific application, constraint violations can be made arbitrarily small by using a sufficiently large number of harmonics.

Using our amplitude objective, we found it unnecessary to explicitly enforce resonance in order to generate large-amplitude motion—and the designs that our method discovered nevertheless proved to be at or close to resonance peaks. Other applications, however, require explicit control over resonance peaks and it would be interesting to extend our formulation in this direction.

We use relatively simple constitutive models for both elastic and viscous material behavior. While our damping model does not explicitly account for air drag, the mass contribution in the Rayleigh damping model emulates this effect to some extent. This choice is justified for problems in which air drag is insignificant, but other applications might require more accurate models.

In our examples, we have not tried to generate periodic motions that include contact or friction. However, HBM can be extended to handle these effects [Krack et al., 2017] and it would be worthwhile exploring their integration in our optimization-based design approach. There are many other aspects of nonlinear vibrations that we deliberately chose to ignore in this work, including bifurcations, period doubling, and internal resonance. Nevertheless, incorporating these phenomena would increase the range of mechanical systems that can be designed with our approach.

The computational burden of our method depends on the complexity of the model, since each mesh vertex is endowed with  $3 \cdot (2N_H + 1)$  Fourier coefficients. While model complexity might necessitate a large number of vertices, the low-frequency oscillations that we aim at are typically confined to a low-dimensional, albeit nonlinear, subspace. Extending our method towards nonlinear subspaces of frequency-space is a promising direction for future research.

We have shown applications of our approach to mechanical leg designs. To be useful for robotics applications, however, the weight carried by the legs (robot body and additional payload) must be accounted for during design. Another interesting direction would be to simultaneously optimize for different driving frequencies such as to adapt leg motion according to running speed.

## 2.A. Appendix

### 2.A.1. Frequency-Space Equilibrium Equations

To derive the dynamic equilibrium equations in frequency space, we start by rewriting the time-dependent nodal positions (2.3.4) and forces (2.3.5) in matrix as

$$\mathbf{x}(t) = (\mathbf{Q}(t) \otimes \mathbb{I}_{3n})\mathbf{z} , \quad (2.A.1)$$

$$\mathbf{f}(t) = (\mathbf{Q}(t) \otimes \mathbb{I}_{3n})\mathbf{b} \quad (2.A.2)$$

where  $\mathbf{z}$  and  $\mathbf{b}$  are Fourier coefficients for positions and forces, and  $\mathbf{Q}(t)$  holds the individual terms of the sine and cosine series

$$Q(t) = [1 \ \sin(\omega t) \ \cos(\omega t) \ \dots \ \sin(N_H \omega t) \ \cos(N_H \omega t)] . \quad (2.A.3)$$

Similarly, velocities and accelerations are expressed as

$$\dot{\mathbf{x}}(t) = (\dot{Q}(t) \otimes \mathbb{I}_{3n})\mathbf{z} = ((\mathbf{Q}(t)\nabla) \otimes \mathbb{I}_{3n})\mathbf{z} , \quad (2.A.4)$$

$$\ddot{\mathbf{x}}(t) = (\ddot{Q}(t) \otimes \mathbb{I}_{3n})\mathbf{z} = ((\mathbf{Q}(t)\nabla^2) \otimes \mathbb{I}_{3n})\mathbf{z} \quad (2.A.5)$$

where the derivative operators are defined as

$$\nabla = \text{diag}(\mathbf{0}, \nabla_1, \dots, \nabla_j, \dots, \nabla_{N_H}) , \quad (2.A.6)$$

$$\nabla^2 = \text{diag}(\mathbf{0}, \nabla_1^2, \dots, \nabla_j^2, \dots, \nabla_{N_H}^2) \quad (2.A.7)$$

with

$$\nabla_j = \begin{bmatrix} 0 & -j\omega \\ j\omega & 0 \end{bmatrix} \quad \text{and} \quad \nabla_j^2 = \begin{bmatrix} -(j\omega)^2 & 0 \\ 0 & -(j\omega)^2 \end{bmatrix} . \quad (2.A.8)$$

Substituting expressions (2.A.1, 2.A.2) and (2.A.4, 2.A.5) into the equations of motion (2.3.2) yields

$$\mathbf{M}((\mathbf{Q}(t)\nabla^2) \otimes \mathbb{I}_{3n})\mathbf{z} + \hat{\mathbf{D}}((\mathbf{Q}(t)\nabla) \otimes \mathbb{I}_{3n})\mathbf{z} = (\mathbf{Q}(t) \otimes \mathbb{I}_{3n})\mathbf{b} . \quad (2.A.9)$$

The mixed-product property of the Kronecker tensor product

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD})$$

is then applied to the left-hand side of the equation as

$$\begin{aligned} \mathbf{M}((\mathbf{Q}(t)\nabla^2) \otimes \mathbb{I}_{3n})\mathbf{z} &= (1 \otimes \mathbf{M})((\mathbf{Q}(t)\nabla^2) \otimes \mathbb{I}_{3n})\mathbf{z} \\ &= ((\mathbf{Q}(t)\nabla^2) \otimes \mathbf{M})\mathbf{z} , \\ \hat{\mathbf{D}}((\mathbf{Q}(t)\nabla) \otimes \mathbb{I}_{3n})\mathbf{z} &= (1 \otimes \hat{\mathbf{D}})((\mathbf{Q}(t)\nabla) \otimes \mathbb{I}_{3n})\mathbf{z} \\ &= ((\mathbf{Q}(t)\nabla) \otimes \hat{\mathbf{D}})\mathbf{z} . \end{aligned}$$

Using this reformulation, Eq. (2.A.9) is rewritten as

$$((\mathbf{Q}(t)\nabla^2) \otimes \mathbf{M})\mathbf{z} + ((\mathbf{Q}(t)\nabla) \otimes \hat{\mathbf{D}})\mathbf{z} = (\mathbf{Q}(t) \otimes \mathbb{I}_{3n})\mathbf{b} . \quad (2.A.10)$$

The time dependency can be removed by a Galerkin procedure, projecting (2.A.10) onto the orthogonal trigonometric basis  $\mathbf{Q}(t)$  and integrating over the period  $T$  of the external force, to obtain

$$\begin{aligned} \left( \left( \frac{2}{T} \int_0^T \mathbf{Q}^T(t)\mathbf{Q}(t)dt \nabla^2 \right) \otimes \mathbf{M} \right) \mathbf{z} + \left( \left( \frac{2}{T} \int_0^T \mathbf{Q}^T(t)\mathbf{Q}(t)dt \nabla \right) \otimes \hat{\mathbf{D}} \right) \mathbf{z} = \\ \left( \left( \frac{2}{T} \int_0^T \mathbf{Q}^T(t)\mathbf{Q}(t)dt \right) \otimes \mathbb{I}_{3n} \right) \mathbf{b} \end{aligned} \quad (2.A.11)$$

Since orthogonality of  $\mathbf{Q}(t)$  implies

$$\frac{2}{T} \int_0^T \mathbf{Q}^T(t)\mathbf{Q}(t)dt = \mathbb{I}_{2N_H+1} , \quad (2.A.12)$$

Eq. (2.A.9) finally transforms into a set of algebraic equations

$$(\nabla^2 \otimes \mathbf{M})\mathbf{z} + (\nabla \otimes \hat{\mathbf{D}})\mathbf{z} = (\mathbb{I}_{2N_H+1} \otimes \mathbb{I}_{3n})\mathbf{b} \quad (2.A.13)$$

which can also be written in a more compact form to yield (2.3.8).

## 2.A.2. DFT Operators

Using Eq. (2.A.1), we can define the inverse DFT operator that transforms from Fourier coefficients to time domain positions as

$$\tilde{\mathbf{x}} = [\mathbf{Q}(t_1) \otimes \mathbb{I}_{3n}, \dots, \mathbf{Q}(t_N) \otimes \mathbb{I}_{3n}]^T \mathbf{z} = (\Gamma \otimes \mathbb{I}_{3n}) \mathbf{z} \equiv \Gamma_x \mathbf{z}. \quad (2.A.14)$$

Similarly, using Eq. (2.A.4), the inverse DFT operator for velocity is obtained as

$$\tilde{\mathbf{v}} = ((\Gamma \Delta) \otimes \mathbb{I}_{3n}) \mathbf{z} \equiv \Gamma_v \mathbf{z}. \quad (2.A.15)$$

Having transformed positions and velocities from frequency space to the time domain, we can evaluate the nonlinear forces for all sampling points and then transform them back to the frequency domain as

$$\mathbf{b} = (\Gamma \otimes \mathbb{I}_{3n})^{-1} \tilde{\mathbf{f}} = (\Gamma^{-1} \otimes \mathbb{I}_{3n}) \tilde{\mathbf{f}} \equiv \Gamma_f^{-1} \tilde{\mathbf{f}}. \quad (2.A.16)$$

The operators  $\Gamma$  and  $\Gamma^{-1}$  are given as

$$\Gamma = \begin{bmatrix} 1 & \sin(\theta_1) & \cos(\theta_1) & \dots & \sin(N_H \theta_1) & \cos(N_H \theta_1) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \sin(\theta_N) & \cos(\theta_N) & \dots & \sin(N_H \theta_N) & \cos(N_H \theta_N) \end{bmatrix}, \quad (2.A.17)$$

$$\Gamma^{-1} = \frac{1}{N} \begin{bmatrix} 1 & \dots & 1 \\ 2 \sin(\theta_1) & \dots & 2 \sin(\theta_N) \\ 2 \cos(\theta_1) & \dots & 2 \cos(\theta_N) \\ \vdots & \vdots & \vdots \\ 2 \sin(N_H \theta_1) & \dots & 2 \sin(N_H \theta_N) \\ 2 \cos(N_H \theta_1) & \dots & 2 \cos(N_H \theta_N) \end{bmatrix}. \quad (2.A.18)$$

It is worth noting that  $\theta_i = \omega t_i = \omega(i\Delta t) = 2\pi i/N$  such that  $\Gamma$  and  $\Gamma^{-1}$  do not depend on  $\omega$ .

## 2.A.3. Newmark Time Integration


The Newmark integration scheme determines end-of-step positions  $\mathbf{x}_{t+\Delta t}$ , velocities  $\dot{\mathbf{x}}_{t+\Delta t}$  and accelerations  $\ddot{\mathbf{x}}_{t+\Delta t}$  such as to satisfy

$$\dot{\mathbf{x}}_{t+\Delta t} = \dot{\mathbf{x}}_t + \Delta t(1 - \gamma)\ddot{\mathbf{x}}_t + \Delta t\gamma\ddot{\mathbf{x}}_{t+\Delta t}, \quad (2.A.19)$$

$$\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \Delta t\dot{\mathbf{x}}_t + \frac{1}{2}\Delta t^2((1 - 2\beta)\ddot{\mathbf{x}}_t + 2\beta\ddot{\mathbf{x}}_{t+\Delta t}) \quad (2.A.20)$$

where  $0 \leq \beta \leq \frac{1}{2}$  and  $0 \leq \gamma \leq 1$  are parameters that we set to  $\beta = \frac{1}{4}$  and  $\gamma = \frac{1}{2}$  to obtain second-order accuracy. We express  $\dot{\mathbf{x}}_{t+\Delta t}$  and  $\ddot{\mathbf{x}}_{t+\Delta t}$  in terms of  $\mathbf{x}_{t+\Delta t}$  and known quantities





**Figure 2.15** – Trajectory difference between HBM and Newmark as a function of the number of time-domain samples  $N_{AFT}$  for  $N_H = 5$  harmonics. The step size for Newmark is set as  $\Delta t = T/N_{AFT}$ , where  $T$  is the period of the forcing.

at time  $t$  as

$$\ddot{\mathbf{x}}_{t+\Delta t} = \frac{1}{\beta\Delta t^2}(\mathbf{x}_{t+\Delta t} - \mathbf{x}_t) - \frac{1}{\beta\Delta t}\dot{\mathbf{x}}_t - \frac{1-2\beta}{2\beta}\ddot{\mathbf{x}}_t, \quad (2.A.21)$$

$$\dot{\mathbf{x}}_{t+\Delta t} = \frac{\gamma}{\beta\Delta t}(\mathbf{x}_{t+\Delta t} - \mathbf{x}_t) + \left(1 - \frac{\gamma}{\beta}\right)\dot{\mathbf{x}}_t + \Delta t \left(1 - \frac{\gamma}{2\beta}\right)\ddot{\mathbf{x}}_t. \quad (2.A.22)$$

Substituting the above equations into the time-domain equations of motion (2.3.1) yields

$$\begin{aligned} \mathbf{R}(\mathbf{x}_{t+\Delta t}) = & \mathbf{M} \left( \frac{1}{\beta\Delta t^2}(\mathbf{x}_{t+\Delta t} - \mathbf{x}_t) - \frac{1}{\beta\Delta t}\dot{\mathbf{x}}_t - \frac{1-2\beta}{2\beta}\ddot{\mathbf{x}}_t \right) + \\ & [D_\alpha\mathbf{M} + D_\beta\mathbf{K}(\mathbf{x}_{t+\Delta t})] \left( \frac{\gamma}{\beta\Delta t}(\mathbf{x}_{t+\Delta t} - \mathbf{x}_t) + \left(1 - \frac{\gamma}{\beta}\right)\dot{\mathbf{x}}_t + \right. \\ & \left. \Delta t \left(1 - \frac{\gamma}{2\beta}\right)\ddot{\mathbf{x}}_t \right) - \mathbf{f}_{\text{int}}(\mathbf{x}_{t+\Delta t}) - \mathbf{f}_{\text{ext}}(\mathbf{x}_{t+\Delta t}) = \mathbf{0}. \end{aligned} \quad (2.A.23)$$

We solve the above system of nonlinear equations using Newton's method. Note that the tangential stiffness matrix depends on the end-of-step positions  $\mathbf{x}_{t+\Delta t}$ .

#### 2.A.4. Accuracy of HBM vs. Newmark

To further investigate the accuracy of HBM compared to Newmark, we additionally analyze the impact of the number of sample points used to evaluate the nonlinear forces in the time domain.

It can be seen from Fig. 2.3 that, with increasing number of harmonics, the error between HBM and Newmark initially decreases very rapidly. For  $N_H > 5$ , however, the error stays almost constant. This seemingly odd behavior is explained by the fact that the reference solution, computed using Newmark, is itself subject to discretization error. We can reduce this time-domain error by using smaller step sizes and, as we increase both the number of harmonics  $N_H$  and time-domain samples  $N_{AFT}$ , the error should vanish. This expectation is supported by the results shown in Fig. 2.15, indicating a sustained decrease in error for increasing number of time-domain samples and, accordingly, the step size used for Newmark.

### 2.A.5. Mechanical Models

Our approach extends to a large range of mechanical models and we show examples using compliant mechanisms, rod networks, thin shells, and volumetric solids. All computational models are implementations of standard approaches. We represent compliant mechanisms through their joint positions, enforcing angular and distance constraints for rigid joints and links with stiff penalty terms. Compliant elements are modeled using standard linear springs whose endpoints are expressed in local coordinates of the corresponding link. For the rod network example, we use discrete elastic rods [Bergou et al., 2008] together with the formulation by Zehnder et al. [2016b] for network connections. For our thin shell examples, we use the discrete shell model by Grinspun et al. [2003a] for bending and linear triangle finite elements with St. Venant-Kirchhoff material [Skouras et al., 2014b] for in-plane deformations. Our solid example uses standard linear tetrahedron elements with a St. Venant-Kirchhoff material.

Each model gives rise to a discrete elastic energy, whose negative gradients with respect to time-domain positions and velocities determine internal forces which enter the frequency-space equations of motion (2.3.8) through the nonlinear force term (2.3.3). For sensitivity analysis, we require the derivatives of the forces with respect to material parameters, masses, and rest shape. All of these derivatives are computed symbolically using compile-time automatic differentiation. Additionally, we use the same penalty term for all mechanical model to enforce harmonic motion for the boundary nodes attached to the motor or driving element.

## Chapter 3

# A Second Order Cone Programming Approach for Simulating Biphasic Materials

by

Pengbin Tang<sup>1</sup>, Stelian Coros<sup>2</sup>, and Bernhard Thomaszewski<sup>2</sup>

<sup>(1)</sup> Université de Montréal

<sup>(2)</sup> ETH Zürich

### Publication

The article has been published in *Computer Graphics Forum*, 41(8), 2022. It was presented at the SIGGRAPH / Eurographics Symposium on Computer Animation 2022 conference. The article was reformatted to fit the style of this thesis.

Pengbin Tang contributed to the method, numerical experiments, and the draft of the paper. Pengbin Tang presented the method at the conference. Stelian Coros and Bernhard Thomaszewski contributed ideas, advice, and editing of the paper.


**ABSTRACT.** Strain limiting is a widely used approach for simulating biphasic materials such as woven textiles and biological tissue that exhibit a soft elastic regime followed by a hard deformation limit. However, existing methods are either based on slowly converging local iterations, or offer no guarantees on convergence. In this work, we propose a new approach to strain limiting based on second order cone programming (SOCP). Our work is based on the key insight that upper bounds on per-triangle deformations lead to convex quadratic inequality constraints. Though nonlinear, these constraints can be reformulated as inclusion conditions on convex sets, leading to a second order cone programming problem—a convex optimization problem that a) is guaranteed to have a unique solution and b) allows us to leverage efficient conic programming solvers. We first cast strain limiting with anisotropic bounds on stretching as a quadratically constrained quadratic program (QCQP), then show how this QCQP can be mapped to a second order cone programming problem. We further propose a constraint reflection scheme and empirically show that it exhibits superior energy-preservation properties compared to conventional end-of-step projection methods. Finally, we demonstrate our prototype implementation on a set of examples and illustrate how different deformation limits can be used to model a wide range of material behaviors.

**Keywords:** Biphasic Materials, Second Order Cone Programming, Strain Limiting, convex optimization, deformation limits, energy conservation.

### 3.1. Introduction

From biological tissues in plants and animals to woven textiles and 3D-printed chainmail—many natural and human-made materials exhibit complex stress-strain responses characterized by strong nonlinearities and direction-dependent behavior. For example, due to their yarn-level structure, many textiles initially exhibit only weak resistance to deformation but stiffen rapidly beyond a direction-dependent threshold. Another extreme example is given by the 3D-printed chainmail shown in Fig. 3.1. One way of modeling this *biphasic* behavior is to use concepts from nonlinear elasticity in combination with experimentally acquired data. An alternative approach that enjoys widespread use in computer animation is to use *strain limiting*, which combines a soft elastic material with hard constraints on the maximum allowed deformation.

Many strain limiting methods enforce deformation limits per spring or element, using, e.g., Gauss-Seidel like iterations [Provot et al., 1995; Bridson et al., 2002; Thomaszewski et al., 2009]. While simplicity is an advantage of this local approach, convergence is often exceedingly slow, in particular if tight deformation bounds are used. Another class of methods combines individual constraints into a globally-coupled problem. While this strategy can greatly accelerate convergence, existing methods either solve only a linearized version of the problem [Goldenthal et al., 2007; Perez et al., 2013] or resort to nonlinear programming techniques to deal with potentially nonconvex constraints [Narain et al., 2012; Jin et al., 2017a]. In either case, there are no formal guarantees on convergence.



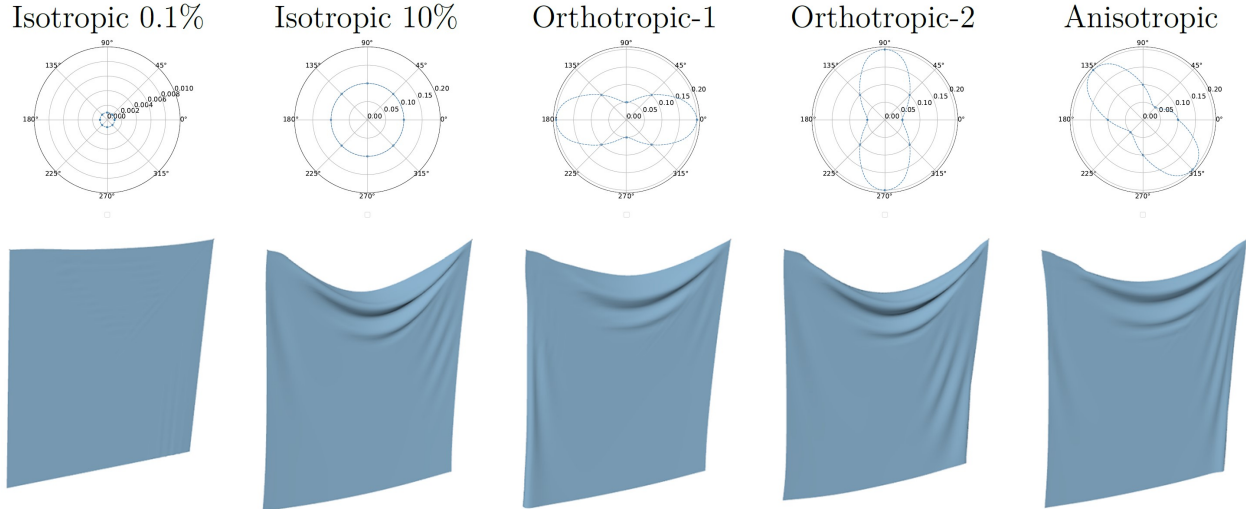
**Figure 3.1** – Two examples of 3D-printed chainmail with experimentally measured deformation limits. Orange areas in the plots indicate the variation of experimental data.

In this work, we propose a new approach to strain limiting based on second order cone programming (SOCP). Using a finite element discretization, we model bounds on triangle deformations in given material directions using quadratic inequality constraints. Our work is based on the key insight that bounds on compression lead to non-convex constraints, whereas bounds on stretching lead to convex constraints. Though nonlinear, stretching constraints can be reformulated as inclusion conditions on convex sets, leading to a second order cone programming problem—a convex optimization problem that a) is guaranteed to have a unique solution and b) can be solved much more efficiently than a general nonlinear program [Andersen et al., 2003].

To develop our formulation, we first cast strain limiting with anisotropic bounds on stretching as a quadratically constrained quadratic program (QCQP), then show how this QCQP can be mapped to a second order cone programming problem. We further propose a constraint reflection scheme and empirically show that it exhibits superior energy-preservation properties compared to conventional end-of-step projection methods. Finally, we demonstrate our prototype implementation on a set of examples and illustrate how different anisotropic deformation limits can be used to model a wide range of material behaviors.

## 3.2. Related Work

Modeling the mechanical behavior of materials is a problem that received much attention in computer animation; see, e.g., [Xu et al., 2015; Li and Barbič, 2015] for recent examples. For textiles, the spectrum ranges from simple mass spring systems with manually specified



**Figure 3.2** – Our method can enforce stretching limits in arbitrary material directions including isotropic, orthotropic, and fully anisotropic constraints. Top row: polar plots showing input curves for deformation limits and sample locations for which constraints are enforced. Bottom row: simulation results for a static drape test using a mesh resolution of  $60 \times 60$ .

stiffness coefficients [Provot et al., 1995; Bridson et al., 2002] to continuum mechanics models with material parameters extracted from real-world measurements [Wang et al., 2011; Miguel et al., 2012]. Instead of modeling fabrics using general nonlinear elasticity, another approach is to approximate them as biphasic materials that exhibit low stiffness for small deformations followed by a quasi-inextensible limit. Strain limiting is a way of modeling these biphasic materials.

The idea of strain limiting goes back to the work of Provot [1995] who experimented with mass-spring simulations using explicit time integration. Using lower stiffness coefficients allowed for larger time steps, but led to undesirable over-elongation effects. Provot corrected for these artifacts by iteratively adjusting vertex positions after every integration step. This relatively simple idea has been widely successful and was used and extended in numerous works. For example, Bridson et al. [2002] integrate strain limiting and collision handling as layers in a velocity filtering stack. Thomaszewski et al. [2009] extended strain limiting to the continuum-mechanics setting by enforcing bounds on the entries of the co-rotated Cauchy strain tensor. Subsequent work introduced limits on the principal strain to achieve isotropic [Wang et al., 2010] and anisotropic [Hernandez et al., 2013] behaviors, and proposed extensions to quadratic strain measures [Müller et al., 2015a].

While all the above approaches enforce strain limits locally by iterating over individual springs and elements, another line of work aimed at globally-coupled enforcement based on constrained optimization techniques. For example, Goldenthal et al. [2007] propose a fast projection scheme in which constraints are linearized and combined into a global system.

The corresponding system is solved repeatedly until the constraint violation is sufficiently small. While this approach is computationally more efficient than simpler per-element iterations, there are no formal convergence guarantees. Nevertheless, fast projection has been extended to triangle meshes [English and Bridson, 2008], and it has been used for augmenting conventional simulators with energy-momentum conservation [Dinev et al., 2018].

In the context of haptic simulation, Perez et al. [2013] propose a globally-coupled treatment of strain limiting and frictional contact, combining linearized constraints into a linear complementarity problem that is solved with Projected Gauss-Seidel. The resulting method is fast enough to enable haptic interaction, although tight constraint enforcement cannot be guaranteed due to the linearized formulation.

In contrast, Jin et al. [2017a] propose a fully nonlinear formulation in which inequality constraints on edge lengths are enforced in an iterative fashion using an active set-like algorithm. While Jin et al. stress the importance of their constraint relaxation and tightening scheme to avoid undesirable local minima, our formulation is convex and is therefore guaranteed to converge to a single global optimum.

Generalizing the idea of strain limiting, Position-based Dynamics (PBD) is a constraint-based animation technique that has been widely successful in computer animation; see [Bender et al., 2015] for an overview. In analogy to developments in strain limiting, Bouaziz et al. [2014] introduced a method that combines the local constraint projections of PBD into a global solve. Overby et al. [2017] showed that the formulation by Bouaziz et al. is equivalent to solving a corresponding constraint optimization problem using the Alternating Direction of Multipliers Method (ADMM) and can thus benefit from progress in ADMM solvers [Zhang et al., 2019]. While ADMM is often the method of choice when it comes to solving non-convex optimization problems, our formulation is convex from the ground up. Nevertheless, some SOCP solvers are implemented using ADMM on their inside.

### 3.3. Theory

We consider discrete elastic surfaces represented by triangle meshes with two sets of vertices,  $\mathbf{X}$  and  $\mathbf{x}$ , holding undeformed and deformed positions, respectively. Let  $\mathbf{x}_k^i, \mathbf{X}_i^k$  denote deformed and undeformed vertex positions for a given triangle  $\mathcal{T}_i$ . Using a finite element discretization over the space of piece-wise linear functions, the deformation gradient  $\mathbf{F} \in \mathbb{R}^{3 \times 2}$  is the unique matrix that maps undeformed edge vectors to their deformed counterparts as

$$(\mathbf{x}_k^i - \mathbf{x}_l^i) = \mathbf{F}(\mathbf{X}_k^i - \mathbf{X}_l^i) . \quad (3.3.1)$$

It is evident from this expression that  $\mathbf{F}$  is a linear function of the deformed positions  $\mathbf{x}$ —a fact that allows us to formulate bounds on stretching as quadratic inequality constraints.

### 3.3.1. Deformation Limits as Quadratic Constraints

Orthotropic and anisotropic materials show variations in their stress-strain response as a function of the direction of deformation. Likewise, deformation limits are generally direction-dependent. In order to properly enforce these limits, we must first establish how to measure deformation in arbitrary directions. Using the deformation gradient from (3.3.1) for a given deformed triangle, we can quantify the directional deformation for any unit-length material-space direction  $\bar{\mathbf{d}} \in \mathbb{R}^2$  as the squared length of the transformed vector, i.e.,

$$s(\bar{\mathbf{d}}) = \mathbf{d}^T \mathbf{d} = \bar{\mathbf{d}}^T \mathbf{F}^T \mathbf{F} \bar{\mathbf{d}} . \quad (3.3.2)$$

Values  $s(\bar{\mathbf{d}}) \geq 1$  indicate stretching, whereas  $s(\bar{\mathbf{d}}) \leq 1$  implies compression. By parameterizing the direction  $\bar{\mathbf{d}}$  through a single scalar,  $\bar{\mathbf{d}} = \bar{\mathbf{d}}(\alpha_i)$ , we obtain a continuous measure of deformation in polar space. We note that, for any direction  $\alpha$ , this measure is a quadratic function of the element's deformed positions  $\mathbf{x}_k^i$ . Next, we discuss how to impose limits on deformations.

We assume that deformation limits are provided through  $m$  samples,  $(\alpha_i, b_i)$ , consisting of an angle  $\alpha_i \in [-\pi/2, \pi/2]$  and a stretch limit  $b_i$ . We then impose a set of quadratic inequality constraints, i.e.,  $s(\alpha_i, \mathbf{x}) \leq b_i^2 \forall i$ . Before we can proceed to a conic reformulation, we must first transform these constraints into a different form. To this end, we observe that for any material-space direction  $\bar{\mathbf{d}}_i$ , the corresponding deformed vector  $\mathbf{d}_i(\mathbf{x})$  is a linear function of  $\mathbf{x}$  and we write

$$\mathbf{d}_i(\mathbf{x}) = \mathbf{F}(\mathbf{x}) \bar{\mathbf{d}}_i = \mathbf{A}_i \mathbf{x} , \quad (3.3.3)$$

where  $\mathbf{A}_i = \nabla_{\mathbf{x}}(\mathbf{F} \bar{\mathbf{d}}_i) \in \mathbb{R}^{3 \times 9}$  holds products between the partial derivatives of  $\mathbf{F}$  and components of  $\bar{\mathbf{d}}_i$  (see Appendix). We can now rewrite the quadratic constraints as

$$s_i(\mathbf{x}) = s(\alpha_i, \mathbf{x}) = \|\mathbf{A}_i \mathbf{x}\|^2 = \mathbf{x}^T \mathbf{A}_i^T \mathbf{A}_i \mathbf{x} \leq b_i^2 . \quad (3.3.4)$$

### 3.3.2. Constraint Projection as QCQP

We simulate biphasic materials using a step-and-project approach [Goldenthal et al., 2007]: we first compute candidate positions  $\hat{\mathbf{x}}$  using a standard implicit time stepping scheme without deformation constraints. We then project these candidate positions onto the constraint manifold by computing the closest configuration that satisfies all constraints, i.e.,

$$\min \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}\|^2 \quad \text{such that} \quad (3.3.5)$$

$$\mathbf{x}^T \mathbf{A}_i^T \mathbf{A}_i \mathbf{x} \leq b_i^2 \forall i .$$

The above expression describes a quadratically-constrained quadratic program (QCQP). Since the matrices  $\mathbf{A}^T \mathbf{A}$  are symmetric positive semi-definite, this QCQP is convex and thus has a single global optimum. Conversely, imposing lower bounds on deformation corresponds



to multiplying the quadratic constraint by  $-1$ . The matrix  $-\mathbf{A}_i^T \mathbf{A}_i$  is evidently negative semi-definite and the constraint is no longer convex. Interestingly, the bilateral version of this constraint is likewise nonconvex. This can be understood geometrically since the feasible set for the quadratic equality constraint corresponds to an ellipsoidal surface which is not a convex set: for any two points on the surface, the line connecting the two points does not belong to the feasible set. These observations are valuable as they assert that our projection problem—which enforces bounds on stretching but not compression—has a unique solution. This is a crucial property that previous formulations did not possess or expose. However, in order to leverage the full computational advantages that this setting offers, we must transform (3.3.5) into a conic formulation as follows.

### 3.3.3. Conic Reformulation

Conic Programs (CP) are optimization problems in which inequality constraints are expressed as set inclusion conditions on convex cones [Lobo et al., 1998; Alizadeh and Goldfarb, 2003]. A set  $\mathcal{K} \subset \mathbb{R}^n$  is called a convex cone if for each  $\mathbf{y}, \mathbf{z} \in \mathcal{K}$  we have  $\mathbf{y} + \mathbf{z} \in \mathcal{K}$  and for each  $\mathbf{y} \in \mathcal{K}$  we have  $\alpha \mathbf{y} \in \mathcal{K}$  for all  $\alpha \geq 0$ . We focus on quadratic (i.e., second-order) cones, which are defined as

$$\mathcal{Q}^{n+1} = \{(t, \mathbf{y}) \in \mathbb{R}^{n+1} \mid \|\mathbf{y}\|_2 \leq t\} . \quad (3.3.6)$$

While quadratic cones directly allow for Euclidean norm constraints, it is often more convenient to impose constraints on squared norms. To this end, we use so-called rotated second-order cones, which are defined as

$$\mathcal{Q}_r^{n+1} = \{(u, v, \mathbf{y}) \in \mathbb{R}^{n+2}, u, v \geq 0 \mid \|\mathbf{y}\|_2^2 \leq 2uv\} . \quad (3.3.7)$$

To convert the quadratic inequality constraints (3.3.4) into equivalent second-order cone constraints, we first introduce two auxiliary variables  $t_i$  and  $\mathbf{y}_i$  such that

$$\begin{aligned} t_i &= b_i^2 , \\ \mathbf{y}_i &= \mathbf{A}_i \mathbf{x} . \end{aligned} \quad (3.3.8)$$

The quadratic constraint is now readily expressed in a second-order cone form as

$$\mathbf{x}^T \mathbf{A}_i^T \mathbf{A}_i \mathbf{x} \leq b_i^2 \iff \left(\frac{1}{2}, t_i, \mathbf{y}_i\right) \in \mathcal{Q}_r^{k+2} . \quad (3.3.9)$$

With all constraints reduced to conic form, we must now transform the quadratic objective function in (3.3.5) into a linear form. This can be accomplished by first writing out the quadratic objective as

$$\min \frac{1}{2} \mathbf{x}^T \mathbf{x} - \mathbf{x}^T \hat{\mathbf{x}} + \frac{1}{2} \hat{\mathbf{x}}^T \hat{\mathbf{x}} . \quad (3.3.10)$$

---

**Algorithm 1** Step-and-Reflect
 

---

```

1: procedure STEPANDREFLECT( $\mathbf{x}^0, \mathbf{v}^0$ )
2:    $\tilde{\mathbf{x}}^{1/2}, \tilde{\mathbf{v}}^{1/2} \leftarrow \text{ImplicitEulerStep}(\mathbf{x}^0, \mathbf{v}^0, \frac{\Delta t}{2})$ 
3:    $\mathbf{x}^{1/2} \leftarrow \text{StrainLimitingSOCP}(\tilde{\mathbf{x}}^{1/2})$ 
4:    $\mathbf{v}^{1/2} = \frac{\mathbf{x}^{1/2} - \mathbf{x}^0}{\Delta t/2}$ 
5:    $\hat{\mathbf{x}}^{1/2} = 2\mathbf{x}^{1/2} - \tilde{\mathbf{x}}^{1/2}, \hat{\mathbf{v}}^{1/2} = \mathbf{v}^{1/2}$ 
6:    $\tilde{\mathbf{x}}^1, \tilde{\mathbf{v}}^1 \leftarrow \text{ImplicitEulerStep}(\hat{\mathbf{x}}^{1/2}, \hat{\mathbf{v}}^{1/2})$ 
7:    $\mathbf{x}^1 \leftarrow \text{StrainLimitingSOCP}(\tilde{\mathbf{x}}^1)$ 
8:    $\mathbf{v}^1 = \frac{\mathbf{x}^1 - \hat{\mathbf{x}}^{1/2}}{\Delta t/2}$ 
9: end procedure

```

---

Omitting the constant term and introducing an additional auxiliary variable  $t_0$ , this quadratic objective is equivalent to

$$\min_{\mathbf{x}, t_0} \quad t_0 - \hat{\mathbf{x}}^T \mathbf{x} \quad \text{s.t.} \quad \mathbf{x}^T \mathbf{x} \leq 2t_0, \quad (3.3.11)$$


where the quadratic inequality constraint can be directly translated into conic form as outlined above. Combining the conic reformulations for the objective and all constraints, we arrive at

$$\begin{aligned} \min_{\mathbf{x}, t_0, t_i, \mathbf{y}_i} \quad & t_0 - \hat{\mathbf{x}}^T \mathbf{x} \quad \text{such that} & (3.3.12) \\ & (1, t_0, \mathbf{x}) \in \mathcal{Q}_r^{n+2} \\ & t_i = b_i^2 \quad \forall i, \\ & \mathbf{y}_i = \mathbf{A}_i \mathbf{x} \quad \forall i, \\ & (\frac{1}{2}, t_i, \mathbf{y}_i) \in \mathcal{Q}_r^{k+2} \quad \forall i. \end{aligned}$$


With the above formulation, we have converted our constraint projection problem into SOCP form. It is therefore guaranteed to have a unique solution, and that solution can be computed with any SOCP solver (we use Mosek 9.3).

### 3.4. Time Stepping

When implemented as a post-step velocity filter in the spirit of [Bridson et al., 2002], our SOCP formulation integrates seamlessly with standard time-stepping schemes and collision response methods. This step-and-project (SAP) approach generally leads to stable and visually pleasing motion. We notice, however, that it largely inherits the numerical dissipation of lower-order implicit integration schemes, which manifests as a loss of kinetic energy and dynamic detail over time; see the accompanying video.




**Figure 3.3** – A schematic visualization of our quasi-symmetric step-and-reflect method. We first step to the middle of the time interval, where we apply twice the correction necessary to satisfy deformation constraints. We then integrate from mid-step to the end of the interval and apply a standard projection step.




**Figure 3.4** – Evolution of total energy over time for different methods on a swinging cloth example with 1% isotropic strain limits. Unconstrained implicit Euler (IE) loses energy more rapidly than our SOCP approach with end-of-step projection (SAP). Our step-and-reflect method (SAR) further improves energy conservation. For BDF-2 versions of projection (BAP) and reflection (BAR) methods, we obtain even lower numerical damping, but larger energy fluctuations.

To improve upon this behavior, we follow the idea by Zehnder et al. [2018] and experiment with a quasi-symmetric constraint reflection scheme that we refer to as *step-and-reflect* (SAR).

To this end, we first integrate to the middle of the time interval, where we enforce constraints by solving (3.3.12). We then apply twice the correction required to satisfy all



**Figure 3.5** – Evolution of total energy over time for different step sizes on a swinging cloth example with 1% isotropic strain limits.



**Figure 3.6** – Our SOCP-based strain limiting method can be integrated into standard collision resolution approaches. Here we use simple penalty functions that prevent intersection during time integration and enforce isotropic strain limits of 0.1%.

constraints which, visually, reflects the current configuration to the other side of the constraint manifold (i.e., the boundary between feasible and infeasible configurations). We step from this reflected configuration to the end of the interval, where we apply a standard constraint projection. See Fig. 3.3 for a visual summary and Algorithm 1 for a pseudo-code description of this scheme.

The rationale of this step-and-reflect approach is that, in analogy to symmetric (i.e., time-reversible) integration schemes, the quasi-symmetric arrangement of constraint correction and integration steps should lead to better energy conservation [Hairer et al., 2006]. Similar to Zehnder et al. [2018], we indeed observe that the reflection scheme leads to better energy conservation, as can be seen in Fig. 3.4. The plot also shows that, unlike the implicit Euler and step-and-project approach, the energy evolution for the step-and-reflect scheme shows slight oscillations around a smoothly evolving curve. We conjecture that this behavior

might be akin to the fluctuations observed in symplectic integration methods which, although non-monotonic, exhibit excellent long-term energy conservation [Hairer et al., 2006]. Using a second-order reflection scheme [Narain et al., 2019] might further improve energy conservation.

We additionally experimented with BDF-2 variants of our end-of-step projection and reflection schemes for comparison. By directly replacing implicit Euler with BDF-2 in SAP, we recover the projection method by English and Bridson [2008] to which we refer as BAP. In addition, we also implemented a reflection version of BAP, termed BAR, but noticed instability problems when using the BDF-2 formula for the mid-step velocity  $\mathbf{v}^{1/2}$ . As shown in Fig. 3.4, the energy curves of BDF-2 methods exhibit larger fluctuations, but show overall better energy conservation as expected. Finally, we also investigate the impact of step size on our SAR method. As can be seen in Fig. 3.5, SAR is robust even for large step sizes even though energy conservation deteriorates.

### 3.5. Results

We tested our method on a number of static and dynamic examples that we discuss below. In all cases, we use a square sheet of fabric with  $0.5m$  side length initialized in the  $xz$ -plane and fixed at two corners. We use constant strain triangle (CST) elements [Thomaszewski et al., 2009] and hinge-based elements [Grinspun et al., 2003b] for modeling the in- and out-of-plane behavior of the material, respectively. We set the material thickness to  $0.5mm$ , its density to  $200kg/m^3$ , and use a Young’s modulus  $1800Pa$  and a Poisson’s ratio of 0.41. Using these parameters, we evaluate our method for different deformation limits and different mesh resolutions.

Our method allows us to freely impose arbitrary bounds on stretching in any material direction. This, in turn, enables simulation of a large variety of materials, including quasi-inextensible isotropic materials, orthotropic fabrics with tight and loose bounds in different directions, as well as fully anisotropic weaves. As illustrated in Fig. 3.2, different combinations of deformation limits lead to different static drapes—and a similarly rich variety in dynamic behavior, as shown in the accompanying video. Another dynamic example illustrates the feasibility of combining our method with standard collision resolution algorithms (see Fig. 3.6).

During all these experiments, we observed that our SOCP-based method is extremely robust, yielding stable results for both static and dynamic examples that precisely satisfy all constraints, regardless of the deformation limits. We attribute this favorable behavior to the strict convexity of our formulation.

**Table 3.1** – Performance of our method compared to unconstrained implicit Euler (IE), general nonlinear programming (NLP), and Gauss-Seidel (GS) as a function of mesh resolution under isotropic strain limiting. Timings are averaged across 100 animation steps. The symbol \* indicates that some steps did not converge whereas – indicates solver failure.

Mesh resolution		$10 \times 10$	$20 \times 20$	$40 \times 40$	$60 \times 60$
	IE	0.017s	0.064s	0.361s	1.318s
1% strain	Ours	0.108s	0.285s	1.289s	4.060s
	NLP	0.436s	*21.009s	–	–
	GS	0.422s	*7.184s	*54.725s	*100.033s
10% strain	Ours	0.107s	0.278s	1.242s	3.790s
	NLP	0.405s	*27.001s	–	–
	GS	0.016s	0.238s	6.311s	*48.427s

### 3.5.1. Performance

We analyze the performance of our method by measuring the average time spent inside the SOCP solver across 100 simulation steps of our swinging cloth example using different mesh resolutions. Note that our SAR method involves two projections, so we average across 200 SOCP solves. For reference, we also present the average time spent inside the Newton solver for implicit integration. As can be seen from Tab. 3.1, the timings for the SOCP solver and implicit Euler are comparable for small and moderate problem sizes. For larger problems, however, computation times increase much more rapidly for SOCP, taking roughly 8s per time step for 10,000 degrees of freedom and approximately 30,000 inequality constraints.

We further evaluated the impact of constraint tightness on computation times, but observed almost no effect when decreasing bounds from 20% to 0.1%. Similarly, we noticed only slight changes (approximately 20%) in computation time when increasing the number of constraints per element from 4 to 18.

To analyze the performance of our SOCP formulation, we compare to two common strain limiting methods: 1) solving a general nonlinear programming (NLP) problem (similar to [Jin et al., 2017a]) and 2) using Gauss-Seidel (GS) to iteratively enforce constraints per element (similar to [Thomaszewski et al., 2009]). We solve the SOCP using MOSEK 9.3 with default parameters, whereas the NLP is solved using Ipopt 3.14 with default parameters. For GS, flag convergence once constraint violations drop below  $1e-6$  and set the maximum number of iterations to  $1e5$ . As shown in Tab. 3.1, our method stands out over these two alternatives, especially for high-resolution meshes and tight deformation bounds. In addition, we can see that both NLP and GS may not converge or even fail to handle the problem for high-resolution meshes. In contrast, our SOCP method handles all of these cases robustly.

All examples were run on a machine with an Intel Core i9-7900X 3.3GHz processor and 32 GB of RAM.

### 3.6. Conclusion

We presented a new formulation for strain limiting based on second order cone programming. Unlike previous methods for globally-coupled strain limiting, our formulation is convex from the ground up and is guaranteed to have a single optimum. We have translated our quadratic inequality constraints into the conic form, allowing us to leverage powerful conic programming solvers.

While our SOCP excels at enforcing bounds on stretching, it cannot handle bounds on compression since the corresponding constraints are not convex. We argue that, for thin sheet materials such as textiles, this is a worthwhile compromise as compressions generally lead to immediate out-of-plane buckling.

Our initial results demonstrate the feasibility of our approach and indicate that, by setting arbitrary bounds on stretching in arbitrary directions, a large range of interesting material behaviors can be modeled. Our experiments further indicate that our SOCP-based approach is unconditionally robust, yielding stable and temporally smooth animations even for tight deformation bounds.

Computation times seem to be largely unaffected by constraint tightness and the number of constraints per element, but they grow rather rapidly with respect to mesh resolution. However, our method still outperforms conventional approaches by large margins, especially for high-resolution meshes and tight deformation bounds. While our formulation will automatically benefit from progress in SOCP solvers, multi-resolution approaches that enforce constraints on coarser levels than the simulation mesh might be a worthwhile direction for future exploration.

We have demonstrated that our method can be combined with a simple collision response approach. However, future work should investigate ways of integrating our method with advanced collision handling frameworks such as [Narain et al., 2012] or [Li et al., 2021b].

While we have focused on in-plane behavior, it would also be interesting to develop quadratic formulations for bending deformation limits. Enforcing limits on bending deformation would enable simulating materials such as quasi-rigid chainmail.

Our results, both qualitative and quantitative, focus on simple geometries and simple test cases. While they constitute a proof-of-concept, we would like to confirm the promising behavior of our method on clothing examples with production-level complexity.

Finally, we believe that many other animation problems can benefit from conic programming reformulations. One particularly interesting case would be frictional contact modeling, for which SOCP methods have already shown their promise [Boyd and Wegbreit, 2007].

## Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. This work was supported by the Discovery Grants Program and the Discovery Accelerator Awards program of the Natural Sciences and Engineering Research Council of Canada (NSERC). Computing equipment has been funded through an infrastructure grant from the Canada Foundation for Innovation (CFI). This work was also supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research, innovation program (grant agreement No. 866480), and the Swiss National Science Foundation through SNF project grant 200021\_200644.

### 3.A. Appendix A

For the deformation gradient  $\mathbf{F} \in \mathbb{R}^{3 \times 2}$ , we compute it as

$$\mathbf{F} = \frac{d\mathbf{x}}{d\bar{\mathbf{x}}} = \frac{d\mathbf{x}\mathbf{N}(\bar{\mathbf{x}})}{d\bar{\mathbf{x}}} = \mathbf{x} \frac{d\mathbf{N}(\bar{\mathbf{x}})}{d\bar{\mathbf{x}}} = \mathbf{x}\mathbf{B} , \quad (3.A.1)$$

where  $\mathbf{N}$  is the shape functions for the triangle. Therefore, we get the matrix  $\mathbf{A}$  in Equation (3.3.3) as

$$\mathbf{A} = \begin{bmatrix} \bar{\mathbf{d}} & & \\ & \bar{\mathbf{d}} & \\ & & \bar{\mathbf{d}} \end{bmatrix}^T \begin{bmatrix} \mathbf{B}_1 & & \\ & \mathbf{B}_1 & \\ & & \mathbf{B}_1 \\ \mathbf{B}_2 & & \\ & \mathbf{B}_2 & \\ & & \mathbf{B}_2 \\ \mathbf{B}_3 & & \\ & \mathbf{B}_3 & \\ & & \mathbf{B}_3 \end{bmatrix}^T , \quad (3.A.2)$$

where  $\mathbf{B}_i$  indicates the  $i$ -th row of matrix  $\mathbf{B}$ .



## Chapter 4

# Beyond Chainmail: Computational Modeling of Discrete Interlocking Materials

by

Pengbin Tang<sup>1,2</sup>, Stelian Coros<sup>2</sup>, and Bernhard Thomaszewski<sup>2</sup>

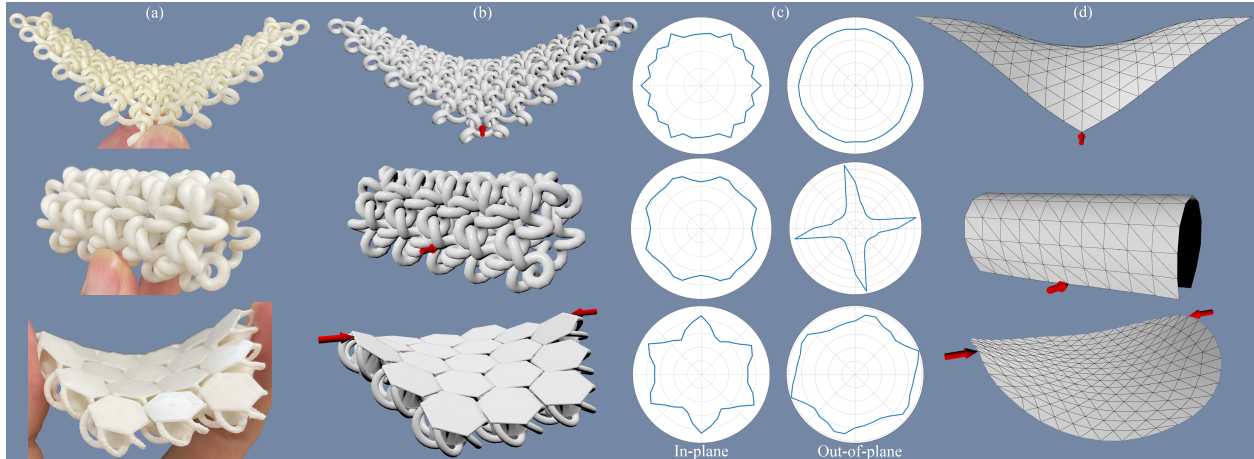
<sup>(1)</sup> Université de Montréal

<sup>(2)</sup> ETH Zürich

### Publication

The article has been published in ACM Transactions on Graphics (TOG), 42(4), (July 2023). It will be presented at the ACM SIGGRAPH 2023. The article was reformatted to fit the style of this thesis.

Pengbin Tang contributed the method, numerical experiments, the physical prototypes, and wrote a draft of the paper. Stelian Coros and Bernhard Thomaszewski contributed to the whole process and edited the paper.




**Figure 4.1** – Discrete Interlocking Materials are governed by strongly coupled, highly anisotropic, and asymmetric deformation limits. Our method is able to capture and reproduce these effects for many types of interlocking materials (a). Using native-scale simulations as a basis (b), we construct macromechanical deformation limits on bending and stretching (c) which we use to develop an efficient macro-scale simulation model (d).

**ABSTRACT.** We present a method for computational modeling, mechanical characterization, and macro-scale simulation of discrete interlocking materials (DIM)—3D-printed chainmail fabrics made of quasi-rigid interlocking elements. Unlike conventional elastic materials for which deformation and restoring force are directly coupled, the mechanics of DIM are governed by contacts between individual elements that give rise to anisotropic deformation constraints. To model the mechanical behavior of these materials, we propose a computational approach that builds on three key components. (a): We explore the space of feasible deformations using native-scale simulations at the per-element level. (b): Based on this simulation data, we introduce the concept of strain-space boundaries to represent deformation limits for in- and out-of-plane deformations, and (c): We use the strain-space boundaries to drive an efficient macro-scale simulation model based on homogenized deformation constraints. We evaluate our method on a set of representative discrete interlocking materials and validate our findings against measurements on physical prototypes.

**Keywords:** Bi-Phasic Materials, Homogenization, Mechanical Characterization, Strain Limiting, Data-Driven Macromechanical Model.

## 4.1. Introduction

Designing materials with desired mechanical properties is an important problem in science and engineering. Often, designers must create materials whose flexibility allows for proper functioning while providing the strength needed to prevent excessive deformation and support external loads. For conventional elastic materials, reconciling these soft-stiff requirements means designing for strong nonlinearities—and with increasing stiffness contrast, this task becomes ever more challenging. Here we consider a new class of Discrete Interlocking Materials (DIM): generalized chainmail fabrics made of quasi-rigid interlocking elements.



**Figure 4.2** – Element cross-section determines coupling for bending. Two threefold symmetric materials are bent using the same uniaxial load case. The elements of design (a) have circular cross-section, while cross-sections for design (b) are slightly ellipsoidal. This small asymmetry in geometry leads to much larger coupling between principal and orthogonal curvature.

Unlike conventional elastic materials for which restoring force increases in lockstep with deformation, the mechanics of DIM is governed by contacts between individual elements. Their particular structure leads to extremely high contrast in deformation resistance: when contacts are sparse or unstructured, deformation is possible without any restoring force (the *slack* regime). Once a chain of contacts is formed along a given loading direction, further deformation is resisted with the full strength of the base material (the *taut* regime). This bi-phasic behavior lends DIM both strength and flexibility, making them attractive for use in robotics, orthotics, sportswear, and many other areas of application. To support such applications, however, we must predict the macro-mechanical behavior of DIM from their native-scale structure, i.e., element shape and connectivity. This relation is highly non-trivial as small changes in element shape can change the macromechanical behavior substantially (see Fig. 4.2).

In this work, we propose a computational approach for modeling and characterizing discrete interlocking materials composed of quasi-rigid elements. Rather than through the properties of their base material, we describe DIM through their strain-space boundaries that delineate the space of kinematically feasible configurations. These boundaries describe multi-dimensional, highly anisotropic limits for in- and out-of-plane deformations. In order to capture this complex mechanical behavior, we propose a novel homogenization approach for distilling macroscale deformation limits from native scale simulations. For each material, we perform thousands of virtual deformation tests for uni- and bi-axial stretching and bending in different directions. We thus establish a sample-based representation of a material’s feasible set of deformations, from which we extract its strain-space boundary. We use this boundary to derive various metrics for characterization and to develop a homogenized model

for DIM that combines thin-shell simulation with anisotropic deformation constraints for stretching and bending. Using this computational framework, we explore the space of discrete interlocking materials and characterize the macro-mechanical behavior of a diverse set of samples. Our simulation-based results are validated through real-world experiments, showing good agreement on both qualitative and quantitative levels. In summary, the key technical contributions of our work are

- the first formal framework for computational modeling and mechanical characterization of discrete interlocking materials,
- a method for homogenizing strain limits for in- and out-of-plane deformations,
- a representation of multi-dimensional deformation limits as feasible sets in strain space,
- a data-driven macromechanical model for discrete interlocking materials that combines thin shell simulation with homogenized deformation constraints.

## 4.2. Related Work

**Metamaterial Design.** Through precisely architected microstructures, flexible metamaterials can achieve a broad range of macromechanical properties [Bertoldi et al., 2017]. Fueled by the increasing availability of 3D printing technology, the graphics community has started to embrace the problem of generating 3D-printable content such as models optimized for stability [Stava et al., 2012; Lu et al., 2014b; Zhou et al., 2013], mechanical assemblies [Zhu et al., 2012; Ceylan et al., 2013b; Coros et al., 2013b; Thomaszewski et al., 2014b], or characters that can be posed and deformed in desired ways [Skouras et al., 2013b; Bächer et al., 2012]. One particular line of research in this context investigates the creation of 3D-printable metamaterials. The spectrum includes layered materials fabricated with multi-material printers [Bickel et al., 2010], materials with lattice- [Panetta et al., 2015, 2017], voxel- [Schumacher et al., 2015; Zhu et al., 2017], and foam-like [Martínez et al., 2016, 2017] structures, as well as two-dimensional, sheet-like materials [Schumacher et al., 2018; Martínez et al., 2019; Tozoni et al., 2020; Leimer and Musialski, 2020]. While these previous works have explored many aspects of mechanical metamaterials, they all focus on elastic behavior. In contrast, we investigate a new class of metamaterials whose macro-scale behavior is regulated by internal contact, not elastic deformation.

**Homogenization.** The macromechanical properties of metamaterials can be determined using the toolset of homogenization theory [Bensoussan et al., 1978]. The central idea of homogenization is to infer macroscopic descriptions for the mechanical behavior of structured materials, typically by analyzing a representative patch of material subject to periodic boundary conditions. Most approaches use the same discretization (e.g., volumetric finite elements) for both native and macro scale simulation [Kharevych et al., 2009; Schumacher

et al., 2015; Panetta et al., 2015]. Nevertheless, there are also cross-discretization approaches that map between very different computational models, e.g., from elastic rods to thin plates [Schumacher et al., 2018] or thin shells [Sperl et al., 2020]. Our method is likewise a cross-discretization approach, mapping from rigid body simulation to a 3D shell model. Unlike previous methods, however, we use homogenization to infer macroscopic deformation limits, not elasticity constants.

**Bi-Phasic Materials.** Capturing and modeling the behavior of elastic materials has a long history in graphics [James and Pai, 1999; Bickel et al., 2009; Hahn et al., 2019]. An alternative to using highly nonlinear material models is to combine a weak elastic base material with deformation constraints [Provot et al., 1995; Goldenthal et al., 2007; Jin et al., 2017b]. Whereas previous methods focused on edge-based strain limiting, Thomaszewski et al. [2009] suggested to enforce limits on per-triangle strain tensors, thus taking into account the full state of deformation. Wang et al. [2010] extended this concept to strain limiting based on singular value decomposition, which subsequent work applied to fully anisotropic limits on nonlinear deformation measures [Perez et al., 2013; Hernandez et al., 2013; Müller et al., 2015b; Tang et al., 2022]. Our discrete interlocking materials can be considered an extreme case of bi-phasic materials: they initially show no resistance to deformation, but exhibit a hard stop once a limit is reached. However, instead of selecting deformation limits to approximate the behavior of highly nonlinear materials, we characterize materials by inferring deformation limits from native-scale simulation data.

**Simulating Contact.** Robust simulation of mechanical systems with frictional contact has been a core focus of graphics research for many years [Hahn, 1988; Baraff, 1989; Kaufman et al., 2008; Bertails-Descoubes et al., 2011; Erleben, 2018; Peiret et al., 2019; Geilinger et al., 2020; Li et al., 2022]. We focus our survey on methods most relevant to interlocking assemblies with tight contacts. In this context, Qu and James [2021] propose a method that computes certificates for topological validity between arrangements of closed curves found, e.g., when simulating knitwear [Kaldor et al., 2008; Cirio et al., 2014] or chainmail. Robust time stepping with tight contacts is also the focus of Li et al. [2020b], who propose smoothly clamped logarithmic barriers to guarantee intersection-free states at all times. Our native-scale simulation model builds on an extension by Ferguson et al. [2021] for intersection-free rigid body simulation. Building on data from this native-scale model, we propose a macro-scale simulation model that captures the high-level deformation behavior without the need to handle inter-element contacts.

**Interlocking Materials & Structures.** Using interlocking as a mechanism for creating stable assemblies is a concept that is used across architecture, robotics, and material sciences. For example, interlocking can be used to build functional furniture without nails or adhesives [Yao et al., 2017; Song et al., 2017], rigid assembly puzzles [Xin et al., 2011; Song et al., 2012; Wang et al., 2018], and stable surfaces made from flexible [Skouras et al., 2015] or rigid [Wang

et al., 2019] components. Interestingly, imperfections during manufacturing can lead to loose joints such that assemblies, despite the rigidity of their components, can produce macroscopic deformations that can be harnessed, e.g., for robotics applications [Lensgraf et al., 2020].


Interlocking can also be used to create strong but flexible materials such as chainmail—a technique that has been known since ancient times. There are comparatively few works in material science and engineering that have revisited this technique. One example is the work by Engel and Liu [2007] who proposed a micro-machining process for creating chainmail-like fabrics. Ransley et al. [2017] combined a chainmail material with temperature controlled actuation for soft robotics applications. More recently, Wang et al. [2021] described a friction-modulated stiffening effect when subjecting sealed chainmail fabrics to external pressure. Our work uses the same material principle, i.e., periodic fabric-like arrangements of interlocking quasi-rigid elements. However, whereas Wang et al. focus on elasticity as an emerging property, we address the fundamental question of quantifying the kinematic behavior of discrete interlocking materials.

### 4.3. Computational Model

Our goal is to construct macroscopic descriptions for discrete interlocking materials (DIM) with a wide range of element shapes and topologies. To this end, we must first be able to predict the behavior of given DIM. We start by describing our native-scale model that simulates DIM on the level of individual elements (Sec. 4.3.1). We then use native-scale simulations to create macroscopic states of in-plane and out-of-plane deformations (Sec. 4.3.2). Using these macroscopic deformation tests, we construct a strain-space representation—the boundary between slack and taut regimes—that describes the space of feasible deformations for a given material. Finally, we propose a macro-scale simulation model that leverages this strain-space representation to efficiently predict the macromechanical behavior of DIM based on homogenized deformation constraints.

#### 4.3.1. Native-Scale Model

The mechanical behavior of Discrete Interlocking Materials is determined by the *shape* of their elements and the way in which they interconnect, i.e., their *topology*. By varying element shape and topology, a large range of macromechanical behavior can be obtained. Here we focus on quasi-rigid elements such that the mechanical behavior is entirely determined by kinematics—deformation limits—not elasticity. A natural choice to model this class of DIM is to represent individual elements as rigid bodies and model interactions between neighboring elements through contact forces. To predict deformation limits, we must compute static equilibrium configurations for given boundary conditions. We follow Ferguson et al. [2021] and cast rigid-body simulation with contact as an unconstrained minimization problem with



**Figure 4.3** – Deformation limits for a three-fold symmetric material are probed in simulation using in-plane periodic boundary conditions (a) and paraboloid bending tests (b).

logarithmic penalty functions that guarantee intersection-free states,

$$\min_{\mathbf{q}} E_{\text{Ext}}(\mathbf{q}) + E_{\text{Coll}}(\mathbf{q}) , \quad (4.3.1)$$

where  $\mathbf{q} = (\mathbf{q}_1, \dots, \mathbf{q}_n)^T \in \mathbb{R}^{6n}$  is a vector holding six degrees of freedom  $\mathbf{q}_i = (\mathbf{x}_i, \boldsymbol{\omega}_i)$  for each of the  $n$  rigid elements, with  $\mathbf{x}_i$  and  $\boldsymbol{\omega}_i$  denoting center-of-mass positions and rotations, respectively. Furthermore,  $E_{\text{Ext}}(\mathbf{q})$  and  $E_{\text{Coll}}(\mathbf{q})$  are potentials due to gravity and contacts. We refer to [Ferguson et al., 2021] for details.

### 4.3.2. Macro-Scale Deformations

The native-scale model allows us to study the response of DIM to arbitrary external forces. However, in order to describe these materials at a level of abstraction that is independent of the size and shape of a given sample, we must choose forces that lead to uniform macroscopic states of deformation. While we draw inspiration from homogenization theory for elastic media, the discrete nature of DIM requires a different approach as described next. **In-plane Deformations.** To characterize the in-plane behavior of DIM, we first define a tileable unit cell as shown in Fig. 4.3(a). In order for this cell to be tileable, the positions of corresponding elements on opposite boundaries must be related by the same translation and their orientations must match. For each boundary pair, these requirements lead to a set of *periodicity conditions* of the form  $\mathbf{x}_j = \mathbf{x}_i + \mathbf{t}_{ij}$  for translation and  $\boldsymbol{\omega}_j = \boldsymbol{\omega}_i$  for rotation. Analogous conditions are imposed for the second boundary pair, i.e.,  $\mathbf{x}_k = \mathbf{x}_i + \mathbf{t}_{ik}$  and  $\boldsymbol{\omega}_k = \boldsymbol{\omega}_i$ . Apart from enforcing tileability, these periodic boundary conditions can also serve for prescribing in-plane strains. For a given angle  $\theta$ , any such macroscopic strain can be defined through a principal direction  $\mathbf{d}_p = (\cos(\theta), \sin(\theta))^T$ , its orthogonal direction

$\mathbf{d}_o = (\cos(\theta + \frac{\pi}{2}), \sin(\theta + \frac{\pi}{2}))^T$ , and corresponding directional strains  $\boldsymbol{\varepsilon}_p$  and  $\boldsymbol{\varepsilon}_o$ . Given these quantities, we can define periodic translations for biaxial states of deformation as

$$\mathbf{t}_{ij} = \boldsymbol{\varepsilon}_p(\bar{\mathbf{t}}_{ij}^T \mathbf{d}_p) \mathbf{d}_p + \boldsymbol{\varepsilon}_o(\bar{\mathbf{t}}_{ij}^T \mathbf{d}_o) \mathbf{d}_o, \quad (4.3.2)$$

where the translations  $\mathbf{t}_{ij}$  and  $\bar{\mathbf{t}}_{ij}$  map vertices from one boundary to their counterparts on the opposite boundary in the current and reference configurations, respectively. The translation  $\mathbf{t}_{ik}$  for the second boundary pair is determined analogously, see Fig. 4.3. To draw a native-scale simulation towards a given macroscopic target strain  $\hat{\boldsymbol{\varepsilon}}$ , we solve the constrained minimization problem

$$\min_{\mathbf{q}, \boldsymbol{\varepsilon}} \quad \frac{1}{2}(\boldsymbol{\varepsilon} - \hat{\boldsymbol{\varepsilon}})^2 + E_{\text{coll}}(\mathbf{q}) \quad \text{s.t.} \quad \mathbf{C}_{\text{BC}}(\mathbf{q}, \boldsymbol{\varepsilon}) = \mathbf{0}, \quad (4.3.3)$$

where  $\mathbf{C}_{\text{BC}}(\mathbf{q}, \boldsymbol{\varepsilon})$  encode the periodic boundary conditions from (4.3.2) and  $\boldsymbol{\varepsilon} = (\boldsymbol{\varepsilon}_x, \boldsymbol{\varepsilon}_y, \boldsymbol{\varepsilon}_{xy})^T$ . Thanks to the coupling between strain variables  $\boldsymbol{\varepsilon}$  and element transformations  $\mathbf{q}$  through boundary conditions and collision penalties, this formulation ensures that strains remain valid even if target values are outside the feasible set.

**Out-of-plane Deformations.** To analyze deformation limits for bending, we initially consider using periodic boundary conditions analogous to the planar setting. While single curvature states—i.e., cylindrical bending—can be conveniently modeled in this way [Sperl et al., 2020], the concept of periodic boundary conditions does not readily extend to the doubly-curved setting: one cannot, in Euclidean geometry, define a finite-sized, double-curvature patch that tiles with itself [Sausset and Tarjus, 2007]. Nevertheless, finite patches of DIM can generally assume states of double curvature. Some materials will even strongly resist uniaxial bending and prefer negative Gaussian curvature instead, see Fig. 4.2(b). For this reason, we lay aside periodic boundary conditions for out-of-plane behavior and instead turn towards a penalty approach that *encourages* states of double curvature without strictly enforcing them through constraints. To this end, we define target surfaces with prescribed curvatures and ask that a given patch of interlocking elements should match this target as closely as possible. We use circular patches of paraboloid surfaces for this purpose, which are defined as  $z = Ax^2 + By^2 + Cxy$ . To achieve a given target curvature  $\hat{\boldsymbol{\kappa}}$ , we reparameterize the positions of all elements as  $\mathbf{x}_i(\boldsymbol{\kappa}) = (x_i, y_i, Ax_i^2 + By_i^2 + Cx_i y_i)^T$ , where  $\boldsymbol{\kappa} = (A, B, C)^T$ . We then find the feasible curvature closest to the target value by solving the minimization problem

$$\min_{\mathbf{q}, \boldsymbol{\kappa}} \quad \frac{1}{2}(\boldsymbol{\kappa} - \hat{\boldsymbol{\kappa}})^2 + E_{\text{Coll}}(\mathbf{q}(\boldsymbol{\kappa})). \quad (4.3.4)$$

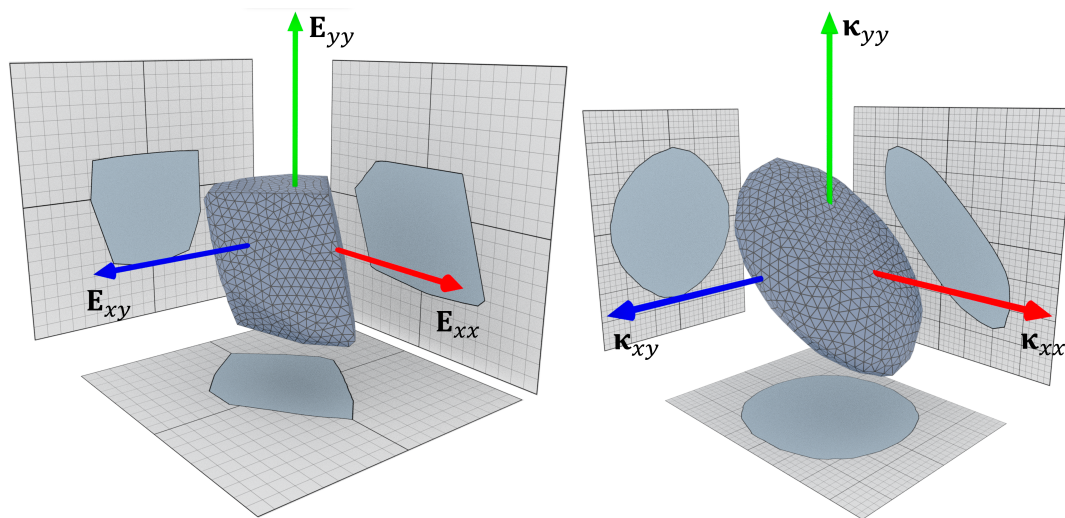
As in the planar setting, simultaneously optimizing for paraboloid parameters  $\boldsymbol{\kappa}$  and element transformations  $\mathbf{q}$  ensures that the corresponding curvature remains in the feasible set: any deviation would lead to intersections between elements, which are strongly penalized by the logarithmic penalty term  $E_{\text{Coll}}$ .



### 4.3.3. Strain-Space Representation

Having established the means to enforce macro-scale deformations for native-scale simulations, we can now start to build macromechanical representations of discrete interlocking materials. Unlike elastic materials, DIM can freely deform in a given direction until a hard stop due to tight contact between elements prevents further deformation. These deformation limits can vary strongly with direction and often depend on orthogonal deformations in highly nonlinear ways. To derive these limits, we observe that DIM can be fully described by enumerating all feasible macroscopic deformations. We assume that the set of all feasible in- and out-of-plane deformations forms a closed, simply-connected region in a six-dimensional strain space. While it is conceivable that materials violating this assumption can be constructed using, e.g., bi-stable/snap-through connections, we do not consider this case here. Even though bending and stretching deformations are coupled for the case of double curvature, our experiments and analysis indicate that treating in- and out-of-plane limits separately is a reasonable approximation (see Sec. 4.4.2). We, therefore, use distinct three-dimensional strain spaces for stretching and bending. See Fig. 4.4 for an example.

**Stretching.** Planar states of strain form a three-dimensional space that can be parameterized in different ways. Using a canonical coordinate system with two axial stretches  $\epsilon_x$ ,  $\epsilon_y$ , and one shear strain  $\epsilon_{xy}$  is one possible choice. However, simple sampling strategies in this space will not lead to uniform distributions of samples with respect to the direction of deformation. For this reason, we parameterize planar strain through an angle  $\theta$  that transforms the canonical coordinate axes into the orthogonal directions of principal stretch. This system



**Figure 4.4** – Strain-space boundaries for the threefold symmetric chainmail material for in-plane (*left*) and out-of-plane (*right*) deformations. Projections of the 3D shapes onto strain-space planes are shown to improve readability.

has zero shear strain such that two auxiliary variables  $\epsilon'_p$  and  $\epsilon'_o$  for stretch along the principal axes fully describe the planar state of strain as in Eq. (4.3.2). To explore the strain space boundary, we start by uniformly sampling directions  $\theta \in [0, \pi)$ . For each sampled direction  $\theta_i$ , we gradually increase (and decrease) the principal stretch  $\epsilon'_p$ . For each principal stretch value we then explore the limits of stretching and compression in the orthogonal direction. We always enforce zero shear to ensure that our principal strain coordinate system remains valid. Whenever we find a state for which deformation in either the principal or orthogonal direction cannot be further extremized, we have found a point on the boundary of the feasible set. In preparation for our macro-scale model, we then convert this point into its equivalent Green strain representation  $\mathbf{E} = (\mathbf{E}_x, \mathbf{E}_y, \mathbf{E}_{xy})^T$  in material coordinates and add it to the list of points used for constructing the strain space boundary.

**Bending.** To enable directional sampling of curvatures in analogy to the planar setting, we observe that any paraboloid with nonzero torsion  $C \neq 0$  can be reparameterized in terms of an angle  $\theta$  indicating the rotation of the canonical axes onto the directions of principal curvature. In this rotated system, we have nonzero curvatures  $A' \neq 0$ ,  $B' \neq 0$  but vanishing torsion  $C' = 0$  such that we can use the same boundary exploration strategy as for in-plane deformations.

**Strain-Space Representation.** The data generation strategy outlined above yields sets of points on the boundary between slack and taut regimes. We triangulate these point sets to obtain closed and watertight surface meshes that serve as explicit boundary representations  $\mathcal{B}_s$  and  $\mathcal{B}_b$  for stretching and bending, respectively. See also Fig. 4.4. Based on these strain-space boundaries, we define high-level directional deformation limits for characterizing discrete interlocking materials. We furthermore use these representations to construct macromechanical simulation models as explained next.

#### 4.3.4. Macro-Scale Model

Native-scale simulation allows for accurate modeling and, consequently, analysis of the mechanical behavior of discrete interlocking materials. While the costs of native-scale simulation are acceptable for distilling macroscopic deformation limits from small samples, they become an impediment when exploring the behavior of larger assemblies. To enable efficient preview of DIM with larger numbers of elements, we develop a macro-scale model that combines thin-shell simulation with homogenized deformation constraints. To this end, we use the bending model by Gingold et al. [2004] as a basis, which provides direct access to curvature tensors using simple four-triangle elements. We complement this bending model using linear triangle finite elements for in-plane deformations, which we quantify using the rotationally-invariant Green strain.

A key insight in this context is that we can enforce per-element stretching and bending limits as set inclusion constraints in strain space. Since the set boundaries are represented with triangle meshes, we can leverage established collision detection algorithms for modeling and enforcing strain-space constraints. To this end, we interpret per-triangle strain and curvature tensors as points in three-dimensional strain space. We enforce deformation limits through soft constraint functions that penalize the motion of points in strain space outside their boundaries  $\mathcal{B}_s$  and  $\mathcal{B}_b$ , respectively. Let  $\mathbf{s}_i \in \mathbb{R}^3$  denote a point in strain space corresponding to the in- or out-of-plane deformation of triangle  $i$ . We first determine the distance  $d$  of  $\mathbf{s}_i$  to the closest primitive—point, edge, or triangle—on  $\mathcal{B}$  using standard geometry tests and bounding volume hierarchies for acceleration. We then set up smooth, unilateral penalty functions as

$$E_{ss}(d) = \begin{cases} k_{ss}d^3 & d \geq 0 \\ 0 & \text{otherwise} \end{cases}, \quad (4.3.5)$$

where positive distance values  $d \geq 0$  indicate points outside the strain-space boundary and  $k_{ss}$  is a penalty coefficient. We use the same strategy for defining penalty functions for stretching and bending limits, which are complemented by weak elastic potentials that regularize behavior inside the feasible region. For these regularizing materials, we aim to select stiffness coefficients that are large enough to prevent numerical problems, but small enough to not interfere with the deformation limits. To this end, we use a soft elastic base material with a Young’s modulus of  $1250Pa$ , a Poisson’s ratio of 0.4, and a mass density of  $250kg/m^3$ . To enable both implicit time stepping and static equilibrium simulation for our macro-scale model, we combine elastic and penalty potentials with an inertia term as described by Martin et al. [2011].


By reducing the number of variables and avoiding contact resolution at the element level, our macro-scale model is able to reduce computation times by one to two orders of magnitude compared to native-scale simulations. We present a detailed analysis of accuracy and computation times in Sec. 4.4.

## 4.4. Results

We evaluate our method on a representative set of discrete interlocking materials and compare results obtained through simulation to measurements on physical prototypes. We start with a brief description of our experimental setup, after which we present our analysis of individual discrete interlocking materials.

### 4.4.1. Experimental Validation

For experimental validation, we fabricate sets of physical prototypes using 3D-printing. To obtain controlled deformation for both in-plane and out-of-plane motion, we use a custom experimental setup as shown in the inset figure, consisting of a linear stage for precision actuation and a load cell for force measurement.



**In-plane Deformation Limits.** To experimentally identify the limits of the slack region, we stretch and compress samples along given directions until the load cell signals that a force threshold of  $2N$  has been reached. This process is repeated 5 times in each direction, and we report lower and upper bounds for the measurements in our plots.

These experiments result in a set of deformation limits for stretching and compression measured in absolute units of length. To remove dependencies on sample size and increase readability, we convert these limits into a Cauchy strain-like relative measure,

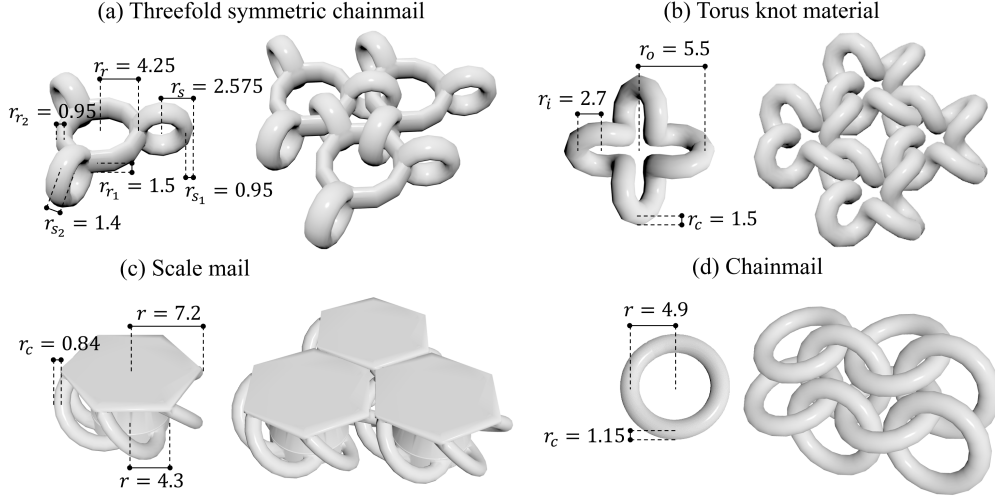
$$c = \frac{l_s - l_c}{l_c} = \frac{(1 + \varepsilon_s)L - (1 + \varepsilon_c)L}{(1 + \varepsilon_c)L} = \frac{\varepsilon_s - \varepsilon_c}{1 + \varepsilon_c}, \quad (4.4.1)$$

where  $l_s$  and  $l_c$  are the measured lengths for stretching and compression experiments. Visually, this directional Cauchy strain indicates the free range of travel relative to the corresponding minimal length under compression. Eq. (4.4.1) also shows that we can easily compute this metric from simulation data, with  $L$  denoting a reference length while  $\varepsilon_s$  and  $\varepsilon_c$  are corresponding stretch and compression limits. Summarizing strain values for different directions in a polar plot then enables direct comparison between experimental measurements and simulation data.

**Out-of-plane Deformation Limits.** We measure bending limits on physical prototypes using the same basic setup. To capture the state of curvature of a given sample, we additionally use an optical tracking system. To this end, we place markers on selected elements such as to sample both curvatures in principal and orthogonal directions. We then recover the closest state of constant curvature from the data by fitting circular arcs to the corresponding marker positions. We use the same experimental protocol as for the in-plane measurements. Due to the finite accuracy of our setup, however, we restrict bending tests to a smaller set of discrete directions that can be easily measured.

#### 4.4.2. Analysis of Discrete Interlocking Materials

Our approach allows us to analyze a wide range of discrete interlocking materials. Here we show results for a set of four representative examples, illustrated in Fig. 4.5, that cover

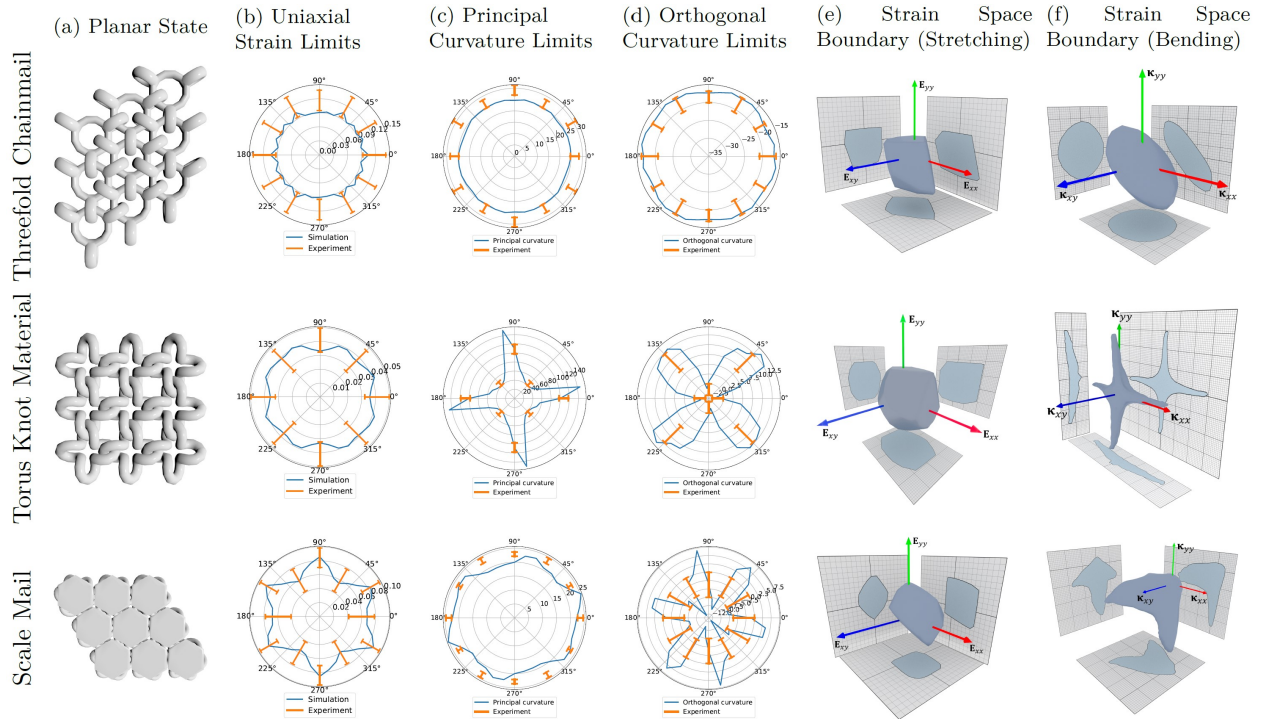


**Figure 4.5** – Dimensions of elements (mm) and topology of interlocking for the materials considered in this work.

some of the major effects to be observed in these materials. For each example, we validate our method by comparing simulation results to measurements on physical prototypes.

**Threefold Symmetric Chainmail.** Our first example is the three-fold symmetric structure shown in Figs 4.5(a) and 4.6(a). Each element of this material connects to its six immediate neighbors, leading to an overall threefold symmetric structure. As can be seen from Fig. 4.6(b), the material shows an almost isotropic limit profile for in-plane deformations with only small higher-frequency fluctuations. The simulation data is furthermore within the error bars (*orange*) of the experimental data. Here, the large discrepancy in the measurements is arguably due to displacements being small compared to the finite accuracy of our experimental setup. While the bending profiles under uniaxial bending shown in Fig. 4.6(c) and 4.6(d) are also almost isotropic, the difference in sign between principal and orthogonal curvature limits indicates a rather interesting coupling behavior: when applying uniaxial bending in any given direction, the material will bend with opposite curvature in the orthogonal direction, leading to characteristic saddle shapes as shown in Fig. 4.1(a). Interestingly, the amount of coupling seems to be determined by the aspect ratio of the elements’ cross-sections. As shown in Fig. 4.2, ellipsoidal shapes induce hyperbolic curvature, whereas circular shapes lead to almost cylindrical bending. The limit plots are just a partial view onto the full set of kinematically-feasible deformations that is captured by our strain-space boundaries, shown in Fig. 4.6(e, f). These strain-space boundaries, which are best appreciated in the accompanying video, feed directly into our macro-scale simulation model, allowing us to predict the complex coupling behavior of this material in an accurate and efficient way. See Fig. 4.1(d).

**Torus Knot Material.** Our second example is a fourfold symmetric material consisting of torus knot elements arranged in a regular quadrilateral grid as shown in Fig. 4.5(b). As can

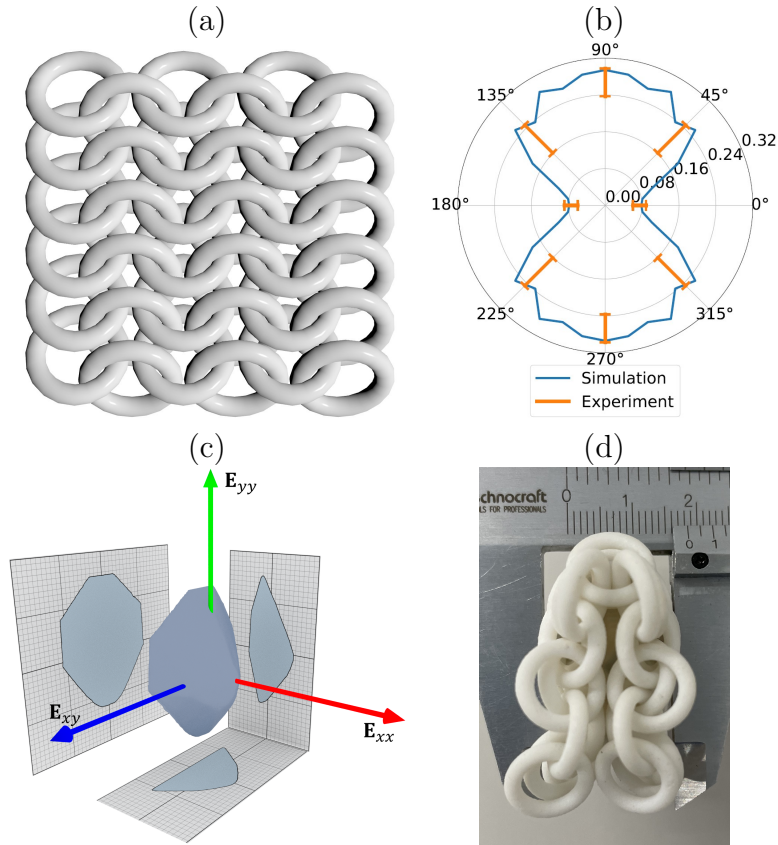


**Figure 4.6 – Deformations limits for threefold symmetric chain mail, torus knot, and scale mail materials.** Polar plots show directional deformation limits for uniaxial stretching (b) as well as biaxial curvature in the principal (c) and orthogonal (d) directions under uniaxial loading. For the uniaxial tests, we bend structures in a given principal direction until reaching their curvature limits. In the orthogonal direction, the structure is left free to deform under gravity and thus reaches a secondary curvature limit. We show the corresponding curvature limits in the principal (c) and orthogonal (d) directions. Solid curves (*blue*) correspond to simulation results which are sampled from  $[0, \pi)$  and rotated to fill the range  $[\pi, 2\pi)$ , experimental data is indicated using error bars (*orange*). These polar plots are partial views onto strain space boundaries for in-plane (e) and out-of-plane deformations (f).

be seen from the uniaxial strain plot in Fig. 4.6(b), this material exhibits almost isotropic in-plane limits, albeit with somewhat larger margins in 30 and 60 degrees. Compared to the threefold symmetric chainmail, however, these limits are much tighter. This material nevertheless admits significantly larger curvature, especially in 0 and 90 degrees. Interestingly, whereas this torus knot material only allows cylindrical curvature in 0 and 90 degrees, it admits doubly-curved states in 45 degrees. All of these observations can again be derived from the corresponding strain-space boundaries shown in Fig. 4.6(e, f).

**Scale Mail.** Our method can also be applied to elements with more complex geometry. In our third example, we consider a material made from elements with a sixfold symmetric base and a hexagonal plate on top of one side in Fig. 4.5(c)<sup>1</sup>. As can be seen from the

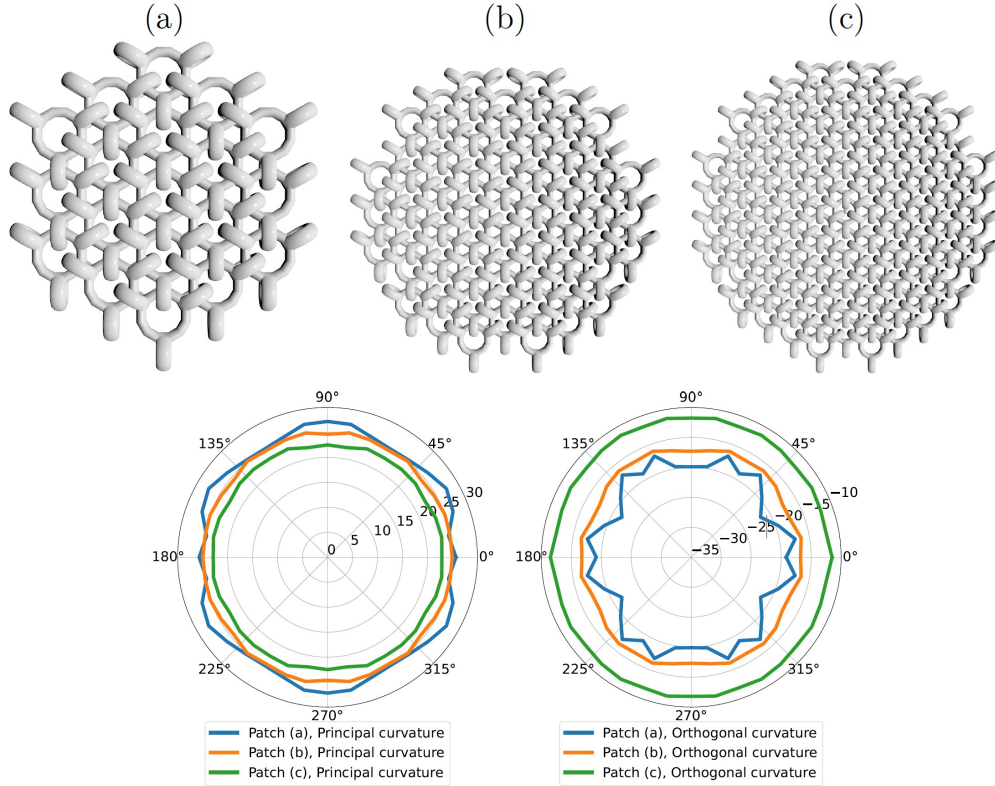
1. This model is a modified version of the NASA chain mail material available at <https://www.thingiverse.com/thing:3095799>



**Figure 4.7** – 4-in-1 chainmail consisting of a staggered arrangement of rings, each of which connects to their four immediate neighbors (a). Deformation limits for uniaxial stretching (b) and strain space boundary for in-plane deformations (c). Bending limits are determined by contact between non-connected elements (d).

uniaxial strain plot in Fig. 4.6(b), our simulation results are in good agreement with the measurements, showing a sixfold symmetric behavior with local maxima and minima for range of motion separated by 15 degrees. The curvature limits computed for this structure, shown in Fig. 4.6(c, d), show slightly larger deviations from the measurements but are still within acceptable margins.

Unlike all other materials considered here, this scale material is not symmetric through the thickness. This asymmetry in shape translates into an asymmetric bending behavior as shown in Fig. 4.12. For positive curvature, contact between the scales puts an early stop to bending. In the negative curvature direction, the bending limit is determined through contact between the interlocking rings and the base, which allows for much larger deformations. This peculiar asymmetric bending behavior can also be observed from the strain space boundaries shown in Fig. 4.6(f), which—unlike for the other materials—shows no symmetry with respect to curvature sign changes.



**Figure 4.8** – Limits for double curvature depend on sample size. Absolute curvature values in both the principal and orthogonal directions decrease as patch size increases from 4 (a) to 6 (b) and 8 (c) elements along the diameter, corresponding to patch radii of  $2.8\text{cm}$ ,  $4.8\text{cm}$ , and  $5.8\text{cm}$ , respectively.

**4-in-1 Chainmail.** As our fourth example, we characterize the behavior of classic chainmail. This material is made of a staggered arrangement of ring-shaped elements that connect to their four immediate neighbors as shown in Fig. 4.5(d). As can be seen from Fig. 4.7(b), our simulation results for in-plane limits again show good agreement with experimental measurements. Although the range of travel for in-plane deformation is not substantially different from other materials, this chainmail admits very large bending deformations. As shown in Fig. 4.7(d), the bending limits are indeed not determined by contact between directly connected elements but occur when remote elements come into contact upon folding. Since these folded configurations are far away from the constant-curvature case that we use for determining bending limits, our method cannot extract meaningful bending limits for this type of material. Nevertheless, accurate macro-mechanical behavior can still be obtained by combining the in-plane limits extracted from simulation data with a penalty-based collision handling approach such as IPC [Li et al., 2020b].

**Discussion.** In summary, our method is able to capture the deformation limits for a wide range of discrete interlocking materials. In our analysis of different materials, we have identified several interesting phenomena, including strong coupling between bending modes

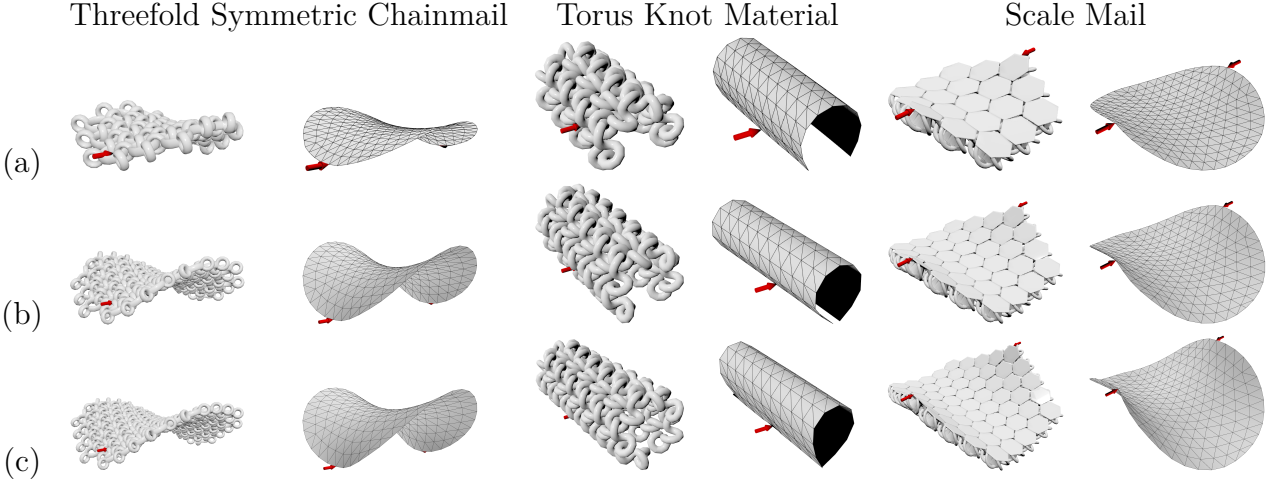


as well as asymmetric bending limits. Although some of the results obtained through native-scale simulation may not align perfectly with experimental measurements, all of them can be deemed within acceptable margins. While many of the plots obtained from simulation data are smooth and symmetric, others exhibit imperfections and high-frequency fluctuations that are not in obvious ways connected to the symmetry of the pattern. One explanation for such perturbations is the fact that in-plane strain is a function of the minimum length obtained for compressive loading. Compared to stretching, these compression tests are inherently more sensitive to noise, both in simulation and reality. Discrepancies between simulated and measured limits can also partly be attributed to the finite precision of our measurement setup. Bending tests generally showed much larger variations, which is arguably due to inaccuracies in marker placement, elastic deformations, as well as friction and gravity.

While our characterization of in-plane limits applies to arbitrary sample sizes, the same is not true for bending limits. As explained in Sec. 4.3.2, periodic tilings of doubly-curved patches do not exist in Euclidean geometry. Experimentally, however, we observe that finite sample sizes do admit double curvature. As shown in Fig. 4.8, our analysis of the relation between double curvature limits and sample size suggests that both principal and orthogonal curvature limits tighten with increasing sample size. We conjecture that this behavior is a direct consequence of Gauss’s *Theorema Egregium*, i.e., the fact that nonzero Gaussian curvature inevitably leads to in-plane deformation. To see this, consider a given patch size with constant nonzero Gaussian curvature. Relative to its flat configuration, this patch must exhibit differential in-plane deformation, e.g., stretching at its center and compression along its boundaries. Increasing the patch size while maintaining the same Gaussian curvature, we must either increase stretch at the center, or compression at the boundary—but both deformations are ultimately bounded by in-plane limits. Our conclusion from this observation is that in-plane limits, not bending limits, prevent double curvature states as sample size increases. We argue that, by measuring double curvature on smaller samples, our method nevertheless provides meaningful information on the intrinsic curvature preferences of discrete interlocking materials. Furthermore, our macro-scale simulations will automatically produce the effect of diminishing double curvature as sample size is increased even though we do not directly model coupling between in- and out-of-plane limits.

### 4.4.3. Macro-Scale Simulation

Having validated the deformation limits obtained through native-scale simulation and our strain-space boundary representation, we now evaluate the behavior of our macro-scale simulation model. All examples were run on a machine with an Intel Core i9-12900HK 2.5GHz processor and 32 GB of RAM.




**Figure 4.9** – Comparison between native- and macro-scale simulations for different materials. Static equilibrium states for applied point loads of  $0.5\text{N}$  (*red arrows*) with different patch sizes. **Threefold symmetric chainmail**: we use disk patches of size  $2.8\text{cm}$  (*a*),  $4.8\text{cm}$  (*b*), and  $5.8\text{cm}$  (*c*). Our macro-scale model uses the same resolution—but different physical sizes—of 384 faces and 217 nodes for all three cases. **Torus knot material**: we use patches with  $4 \times 4$  (*a*),  $5 \times 5$  (*b*), and  $6 \times 6$  (*c*) elements. Our macroscale model use again the same resolution ( $\#face = 200$ ,  $\#nodes = 121$ ) but different physical sizes with  $4.4\text{cm}$  (*a*),  $5.5\text{cm}$  (*b*), and  $6.4\text{cm}$  (*a*) side length, respectively. **Scale mail**: we use disk-shaped patches with 19 (*a*), 37 (*b*), and 61 (*c*) elements and corresponding radii of  $3.4\text{cm}$ ,  $4.6\text{cm}$ , and  $5.9\text{cm}$ , respectively. Our macro-scale simulations use the same mesh resolution (384 faces and 217 nodes) for all patch sizes.

**Comparison between Native- and Macro-Scale Simulations.** Using the materials analyzed in Sec. 4.4.1, we apply the strain space representations captured from native-scale simulation data to drive the deformation constraints in our macro-scale model. Fig. 4.9 shows a comparison between macro- and native-scale simulation results for the threefold symmetric chainmail, torus knot, and scale mail materials. For each material, we further compare results for different physical patch sizes. For native-scale simulations, we increase the number of elements according to patch size. For macro-scale simulations, we simply change the size of the mesh while keeping its resolution fixed. As can be seen from these comparisons, our macro-scale simulation results are in good agreement with their native-scale counterparts.

As can be seen from Tab. 4.1, computation times for our macro-scale model are between one and two orders of magnitude smaller than for the native-scale model. It can further be noted that computation times for the native-scale model grow very rapidly with increasing problem size. Our macroscale model, however, shows almost constant cost since the same mesh resolution is used for all physical sample sizes. Small differences between macro-scale timings are due to the fact that more solver iterations are required to reach equilibrium states with larger deformations.


**Table 4.1** – Performance comparison between native- and macro-scale simulations for the examples shown in Fig. 4.9.

Example	Native-scale [s]	Macro-scale [s]	Speedup
Threefold (a)	133.526	68.394	2.0x
Threefold (b)	883.701	61.104	20.5x
Threefold (c)	1883.044	61.734	14.5x
Torus knot (a)	216.75	78.618	2.8x
Torus knot (b)	722.203	97.602	7.4x
Torus knot (c)	1647.41	141.261	11.7x
Scale mail (a)	179.057	43.581	4.1x
Scale mail (b)	962.517	36.821	26.1x
Scale mail (c)	4929.636	40.755	121.0x




**Figure 4.10** – Comparison of macro-scale simulation with increasing mesh resolution. We start from with coarse meshes (*a*) and increase resolution by a factor of 4 (*b*) and 16 (*c*).

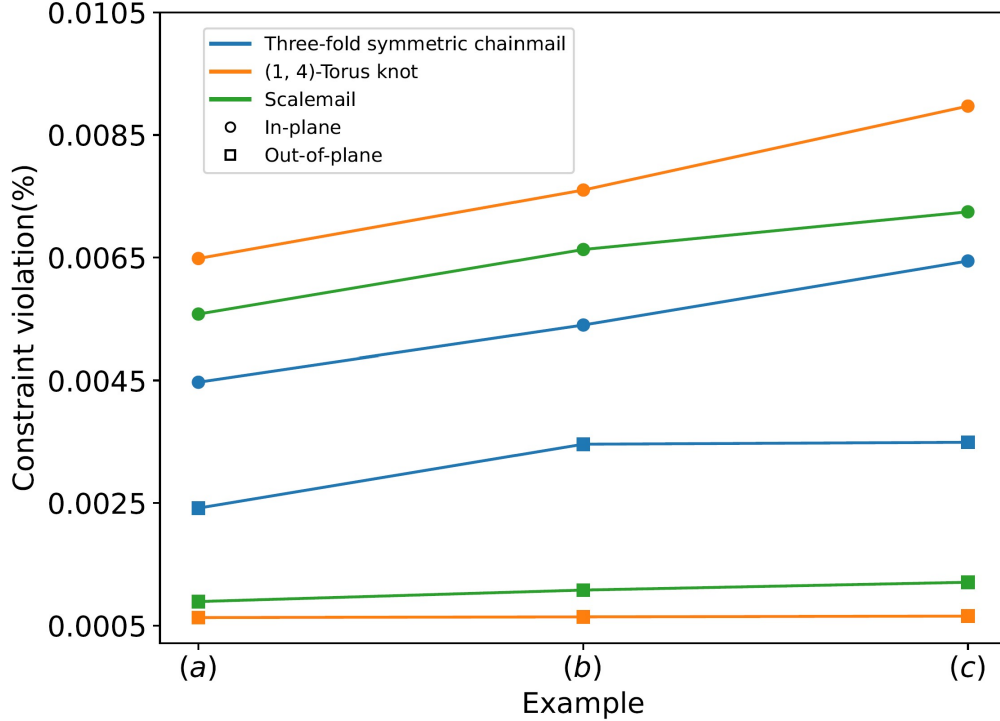
**Increasing Mesh Resolution and Young’s Modulus.** To investigate the mesh dependence of our method, we compare macro-scale simulation results with increasing mesh resolution in Fig. 4.10. It can be seen that our macro-scale simulations behave consistently under refinement. We furthermore analyze the influence of the regularizing elastic material by comparing macro-scale simulation results for increasing Young’s moduli. From the results shown in Fig. 4.11, it can be seen that using too large a stiffness value (5000Pa) prevents the materials from reaching their deformation limits. It should be noted that the lower bound on the regularizing stiffness is determined by purely numerical considerations—below a certain stiffness, the linear solver will need additional diagonal regularization which slows down convergence, but the solution will be largely unaffected.



**Figure 4.11** – Comparison of macro-scale simulations with increasing Young's Modulus.



**Figure 4.12** – **Asymmetric Bending Limits.** Due to the through-the-thickness asymmetry of its elements, this scale mail prototype (*left*) exhibits very different bending limits for positive and negative curvature. Our macro-scale model (*right*) accurately captures this behavior and closely tracks the native-scale simulation results (*middle*).



**Figure 4.13** – Constraint violations for the examples shown in Fig. 4.9. Maximum constraint per-element constraint violations are plotted separately for in- and out-of-plane limits as indicated. It can be seen that constraint violations stay below 0.01% for all cases.

**Bending under Gravity.** As an additional test, we further analyze the ability of our macro-scale model to reproduce the asymmetric bending behavior of the scale mail material. To this end, we clamp one side of a hexagonal patch with 19 elements and let it deform into its equilibrium state under gravity. We then turn the sample upside down and repeat the experiment. As shown in Fig. 4.12, our macro-scale model captures the salient asymmetry in bending and closely tracks the behavior of both native-scale simulation and the physical prototype.

**Constraint Violations.** We use penalty terms to enforce both in- and out-of-plane deformation limits. For all our macro-scale simulations, we use a stiffness coefficient of  $1e^7$  to ensure that constraint violation remains below 0.01%. To further analyze the accuracy of our macro-scale model, we investigate maximum constraint violations for the examples shown in Fig. 4.9. For each element, we compute the smallest positive distance between its corresponding point in strain space and the strain-space boundary. We normalize these values with respect to the diagonal length of the bounding box of the strain-space boundary and plot the maximum value as a function of sample size. The results shown in Fig. 4.13 indicate that constraint violations are small for all cases.

## 4.5. Conclusions

We presented a method for computational modeling, mechanical characterization, and efficient macro-scale simulation of Discrete Interlocking Materials—generalized chainmail fabrics made from quasi-rigid elements. Unlike elastic materials, the mechanical behavior of DIM is governed by inter-element contacts that give rise to anisotropic deformation constraints. We proposed to explore these limits using thousands of native-scale simulations for in- and out-of-plane deformations. Using this simulation data, we have introduced the concept of strain-space boundaries as an explicit representation of the set of feasible deformations. These strain-space boundaries provide insight into the mechanics of DIM, and they also give rise to an efficient macro-scale simulation model based on homogenized deformation constraints. We have applied our method to a set of representative discrete interlocking materials and validated our findings against measurements on physical prototypes. We conclude that our methodology is well-suited for characterizing and modeling a broad range of discrete interlocking materials.

### 4.5.1. Limitations & Future Work

There are nevertheless several limitations of our work that indicate avenues for future work. Our analysis has focused on kinematic motion and, consequently, we have neither considered friction nor elastic deformations in the structure. While this approximation is justified for many types of discrete interlocking materials, friction forces can become dominant for tight materials in which elements have little freedom to move. Moreover, as shown by Wang et al. [2021], resistance to deformation can be controlled through normal pressure—a mechanism which crucially relies on friction. To apply our method to such friction-dominated scenarios, we would like to extend our macro-scale model to account for internal friction. The work by Miguel et al. [2013] for modeling hysteresis in textiles through internal friction might be a good starting point.

While we modeled elements as rigid bodies, this is clearly an idealized assumption. None of our physical prototypes showed noticeable deformations for in-plane loading. For bending, however, an elastic regime can be *felt* for some of the materials. Modeling and characterizing elasticity of discrete interlocking materials outside the slack region is an interesting direction for future research.

Our method successfully captured the strain spaces for four different types of DIM using an increasing-and-decreasing biaxial loading scheme. Nevertheless, this simple scheme might not be able to fully explore strain spaces that are not star-shaped, i.e., for which rays from the origin can intersect the boundary more than once. We expect DIM with elements admitting twist or screw-like motions to fall into this category.

Our method offers an intuitive, systematic way for macromechanical characterization which can pave the way to using DIM for garment design. An interesting direction for future work would be to extend our formulation to inverse design problems such as finding element shapes that lead to desired deformation limits.

While strain space boundaries could perhaps be derived analytically for simple elements, obtaining closed-form deformation limits for complex-shaped elements seems elusive. Nevertheless, further analysis into the relation between structural symmetry and deformation limits is another exciting direction for future work.

Our macro-scale model offers a computationally efficient alternative that closely tracks native-scale behavior. However, our method does not provide geometric detail at the element level, which may be important for applications. Exploring mechanically-aware geometry synthesis and rendering methods in the spirit of Sperl et al. [2021a] would be a worthwhile direction for future work.

Finally, while we have focused on homogeneous patterns with planar rest states, the space of possible discrete interlocking materials is much larger. In the future, we would like to explore the design and simulation of heterogeneous materials with curved rest shapes and spatially varying element shapes and connectivities.

## Acknowledgments

We would like to thank Charles Gingras for the valuable discussion, and Ronan Hinchet for helping with the measuring setup and fabricating some of the physical prototypes. We are also grateful to the anonymous reviewers for their valuable comments. This work was supported by the Discovery Grants Program and the Discovery Accelerator Awards program of the Natural Sciences and Engineering Research Council of Canada (NSERC). Computing equipment has been funded through an infrastructure grant from the Canada Foundation for Innovation (CFI). This work was also supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866480), and the Swiss National Science Foundation through SNF project grant 200021\_200644.





# Chapter 5

---

## Conclusion

In this thesis, we present three methods for the modeling and design of compliant mechanical systems and materials.

In Chapter 2, we introduce a harmonic balance approach for designing compliant mechanical systems with desired nonlinear periodic motions. Combined with sensitivity analysis, our method allows us to interactively explore the design space. However, as the number of parameters increases, manual exploration soon becomes unattractive. Therefore, we further introduce amplitude and trajectory objectives to enable efficient, automatic inverse design of steady-state motion for a given target behavior. We validate our method by fabricating a set of physical prototypes. Our results indicate that harmonic balance paired with sensitivity analysis is indeed an efficient and effective combination that enables the construction of powerful forward and inverse design tools.

In Chapter 3, we develop a strain-limiting method based on second-order cone programming (SOCP) for the modeling and simulating of biphasic materials. Unlike previous methods, our method is convex from the ground up and is guaranteed to have a unique solution. Our SOCP method can seamlessly combine with existing implicit Euler time integration, however, the step-and-projection approach presents a severe loss of kinetic energy and dynamic detail over time. Thus, we further introduce a quasi-symmetric constraint reflection scheme, named as step-and-reflection scheme, with SOCP for simulating biphasic materials. Our results show our method can simulate a large range of materials with different deformation limits and is much faster than previous methods.

In Chapter 4, we introduce a new method for computational modeling, mechanical characterization, and efficient macro-scale simulation of quasi-rigid Discrete Interlocking Materials (DIM). To characterize all nonlinear mechanical properties of DIM, we explore both in- and out-of-plane deformation limits using thousands of native-scale simulations. Based on all simulation data in strain space, we build both in- and out-of-plane strain-space boundaries as an explicit representation of the set of feasible deformations. These strain-space

boundaries provide insight into the mechanics of DIM, and they also give rise to an efficient macro-scale simulation model based on homogenized deformation constraints. Our results reveal that our method can characterize a broad range of discrete interlocking materials.

In summary, this thesis contributes new knowledge for modeling and designing of mechanical nonlinearities at different levels. Specifically, we propose three methods for modeling and designing nonlinear mechanical systems and materials, each advancing the state-of-the-art. Our first approach is the design of nonlinear dynamics with periodic motions for nonlinear mechanical systems. This method is general with respect to mechanical models and can be applied to design a wide range of complex nonlinear mechanical systems. We then explore the modeling and characterization of nonlinear materials dominated by contact, particularly Discrete Interlocking Materials. The micro-level nonlinearities of DIM, such as complex structure and discontinuity of contact, give rise to complex mechanical properties at the meso- and macro-level. We develop two methods to efficiently simulate and explore these materials. Although these two approaches are inspired by DIM, our methods are not limited to modeling and simulating DIM and can be applied to other materials with similar properties.

## 5.1. Future Work

For the design of nonlinear periodic motions for compliant mechanical systems, we have developed a leg design that mimics the running motion of the ostrich. However, to further explore promising mechanical systems, it would be interesting to consider nonsmooth contact. Rather than developing advanced machine learning control algorithms, it would be also interesting to use our design tool to explore stable periodic motions of walking and running for compliant mechanical systems [Miyashita et al., 2020]. This direction is also appealing because animals are capable of efficient walking and running motions without conscious thought, and our design tool can help to replicate these motions in mechanical systems.

For individuals without a background in mechanical system design, it can be challenging to propose an initial design that aligns with their goals. Often, they may only know their target goal or desired functionality, without any knowledge of how to create an initial design for subsequent optimization. Moreover, without a clear understanding of nonlinear mechanical systems, they may not know whether their initial design will ultimately achieve their final goals. Additionally, a potentially feasible initial design may behave quite differently from the final optimized version. To address this, we recommend introducing new methods, such as machine learning-based techniques, for the automatic generation of mechanical systems. This can serve as a helpful starting point for our design tool, particularly for individuals without a professional background in design.

Simulating discrete interlocking materials with certain shapes, e.g., clothing and orthoses, requires mapping rigid discrete interlocking material onto the specified surfaces without introducing intersection. In addition, designing shapes with discrete interlocking materials that are stable and structural sound under static equilibrium is also a potential direction. Achieving these goals requires careful consideration of the geometry and mechanical properties of the interlocking materials and the surfaces onto which they are mapped. Therefore, future research can focus on developing new design approaches that can ensure the successful integration of discrete interlocking materials into complex shapes while maintaining stability and structural integrity.

We can use our macro-scale model to simulate a wide range of discrete interlocking materials. It provides a much more efficient way of simulating interlocking materials than native-scale simulation. However, we did not provide the rendering of native-scale geometry along with the macro-scale simulation. Several works [Montazeri et al., 2019; Sperl et al., 2021b] have been proposed to render mechanics-aware yarn-level cloth before. However, unlike super elastic yarn-level rods, ours are almost rigid so their methods cannot directly apply to our setting. Therefore, developing new mechanical-aware rendering techniques for quasi-rigid interlocking materials is also a good direction for future research. Such techniques would allow us to generate more realistic visualizations of the interlocking materials and their behavior for more potential applications.

Our results demonstrate that discrete interlocking materials exhibit a range of interesting mechanical properties that differ from those of traditional continuous materials. However, interlocking materials are currently passive materials that only respond to external stimuli. To enable more diverse functionality for this type of material, future research can explore the use of discrete interlocking materials in the design of active metamaterials. By incorporating active elements into the interlocking structure, such as sensors or actuators, we could create materials that can actively respond to changes in their environment, opening up new possibilities for their use in various applications.

Nowadays, machine learning algorithms have also shown good performance for material characterization. Due to complex geometry and topology, designing interlocking materials via machine learning will be also a potential direction.



## Bibliography

---

- J eremie Dumas, An Lu, Sylvain Lefebvre, Jun Wu, and Christian Dick. By-example synthesis of structurally sound patterns. *ACM Transactions on Graphics (TOG)*, 34(4):1–12, 2015.
- Jonas Zehnder, Stelian Coros, and Bernhard Thomaszewski. Designing structurally-sound ornamental curve networks. *ACM Transactions on Graphics (TOG)*, 35(4):99, 2016a.
- Erva Ulu, James Mccann, and Levent Burak Kara. Lightweight structure design under force location uncertainty. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017.
- Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. Build-to-last: strength to weight 3d printed objects. *ACM Transactions on Graphics (TOG)*, 33(4):97, 2014a.
- Jon s Mart nez, J eremie Dumas, Sylvain Lefebvre, and Li-Yi Wei. Structure and appearance optimization for controllable shape design. *ACM Transactions on Graphics (TOG)*, 34(6):1–11, 2015.
- M lina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. Designing inflatable structures. *ACM Transactions on Graphics (TOG)*, 33(4):1–10, 2014a.
- Jes s P rez, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, Jos  A. Canabal, Robert Sumner, and Miguel A. Otaduy. Design and fabrication of flexible rod meshes. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 34(4), 2015.
- Jes s P rez, Miguel A Otaduy, and Bernhard Thomaszewski. Computational design and automated fabrication of kirchhoff-plateau surfaces. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017.
- Luigi Malomo, Jes s P rez, Emmanuel Iarussi, Nico Pietroni, Eder Miguel, Paolo Cignoni, and Bernd Bickel. Flexmaps: computational design of flat flexible shells for shaping 3d objects. *ACM Transactions on Graphics (TOG)*, 37(6):1–14, 2018.
- Kazuhiro Miyashita, Yoichi Masuda, Megu Gunji, Akira Fukuhara, Kenjiro Tadakuma, and Masato Ishikawa. Emergence of swing-to-stance transition from interlocking mechanism in horse hindlimb. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7860–7865. IEEE, 2020.

- Duygu Ceylan, Wilmot Li, Niloy J. Mitra, Maneesh Agrawala, and Mark Pauly. Designing and fabricating mechanical automata from mocap sequences. *ACM Trans. Graph.*, 32(6), November 2013a. ISSN 0730-0301. doi: 10.1145/2508363.2508400. URL <https://doi.org/10.1145/2508363.2508400>.
- Vittorio Megaro, Jonas Zehnder, Moritz Bächer, Stelian Coros, Markus H Gross, and Bernhard Thomaszewski. A computational design tool for compliant mechanisms. *ACM Trans. Graph.*, 36(4):82–1, 2017.
- Dewen Li, Xiaochen Bu, Zongpu Xu, Yingwu Luo, and Hao Bai. Bioinspired multifunctional cellular plastics with a negative poisson’s ratio for high energy dissipation. *Advanced Materials*, 32(33):2001222, 2020a.
- Sahab Babae, Jongmin Shim, James C Weaver, Elizabeth R Chen, Nikita Patel, and Katia Bertoldi. 3d soft metamaterials with negative poisson’s ratio. *Advanced Materials*, 25(36):5044–5049, 2013.
- Guangzhao Qin and Zhenzhen Qin. Negative poisson’s ratio in two-dimensional honeycomb structures. *npj Computational Materials*, 6(1):1–6, 2020.
- Yifan Wang, Liuchi Li, Douglas Hofmann, José E Andrade, and Chiara Daraio. Structured fabrics with tunable mechanical properties. *Nature*, 596(7871):238–243, 2021.
- Tian Chen, Mark Pauly, and Pedro M Reis. A reprogrammable mechanical metamaterial with stable memory. *Nature*, 589(7842):386–390, 2021.
- Fengwen Wang and Ole Sigmund. 3d architected isotropic materials with tunable stiffness and buckling strength. *Journal of the Mechanics and Physics of Solids*, 152:104415, 2021.
- Johannes TB Overvelde, Twan A De Jong, Yanina Shevchenko, Sergio A Becerra, George M Whitesides, James C Weaver, Chuck Hoberman, and Katia Bertoldi. A three-dimensional actuated origami-inspired transformable metamaterial with multiple degrees of freedom. *Nature communications*, 7(1):1–8, 2016.
- Yanbin Li, Qiuting Zhang, Yaoye Hong, and Jie Yin. 3d transformable modular kirigami based programmable metamaterials. *Advanced Functional Materials*, 31(43):2105641, 2021a.
- JN Grima and KE Evans. Auxetic behavior from rotating squares. *Journal of Materials Science Letters*, 19(17):1563–1565, 2000.
- Zhiming Cui and Jaehyung Ju. Topologically reconfigurable mechanical metamaterials with motion structures. *Mechanics of Materials*, 143:103317, 2020.
- Evgueni T Filipov, Tomohiro Tachi, and Glaucio H Paulino. Origami tubes assembled into stiff, yet reconfigurable structures and metamaterials. *Proceedings of the National Academy of Sciences*, 112(40):12321–12326, 2015.
- Yichao Tang and Jie Yin. Design of cut unit geometry in hierarchical kirigami-based auxetic metamaterials for high stretchability and compressibility. *Extreme Mechanics Letters*, 12:77–85, 2017.

- Zirui Zhai, Yong Wang, and Hanqing Jiang. Origami-inspired, on-demand deployable and collapsible mechanical metamaterials with tunable stiffness. *Proceedings of the National Academy of Sciences*, 115(9):2032–2037, 2018.
- Midori Isobe and Ko Okumura. Initial rigid response and softening transition of highly stretchable kirigami sheet materials. *Scientific reports*, 6(1):1–6, 2016.
- Doh-Gyu Hwang and Michael D Bartlett. Tunable mechanical metamaterials through hybrid kirigami structures. *Scientific reports*, 8(1):1–8, 2018.
- Hongying Zhang and Jamie Paik. Kirigami design and modeling for strong, lightweight metamaterials. *Advanced Functional Materials*, page 2107401, 2022.
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. Microstructures to control elasticity in 3d printing. *ACM Transactions on Graphics (Tog)*, 34(4):1–13, 2015.
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. Elastic textures for additive fabrication. *ACM Transactions on Graphics (TOG)*, 34(4):1–12, 2015.
- Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. Procedural voronoi foams for additive manufacturing. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016.
- Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. Orthotropic k-nearest foams for additive manufacturing. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017.
- Jonas Zehnder, Espen Knoop, Moritz Bächer, and Bernhard Thomaszewski. Metasilicone: design and fabrication of composite silicone with desired mechanical properties. *ACM Transactions on Graphics (TOG)*, 36(6):1–13, 2017.
- Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. Mechanical characterization of structured sheet materials. *ACM Transactions on Graphics (TOG)*, 37(4):1–15, 2018.
- Jonàs Martínez, Mélina Skouras, Christian Schumacher, Samuel Hornus, Sylvain Lefebvre, and Bernhard Thomaszewski. Star-shaped metrics for mechanical metamaterial design. *ACM Transactions on Graphics (TOG)*, 38(4):1–13, 2019.
- Ahmed A. Shabana. *Theory of Vibration, Volume II: Discrete and Continuous Systems*. Springer, New York, NY, USA, 1990.
- Malte Krack and Johann Gross. *Harmonic Balance for Nonlinear Vibration Problems*. Springer, 2019.
- Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W. Sumner, Wojciech Matusik, and Bernd Bickel. Computational design of mechanical characters. *ACM Trans. Graph.*, 32(4), July 2013a. ISSN 0730-0301. doi: 10.1145/2461912.2461953. URL <https://doi.org/10.1145/2461912.2461953>.

- Bernhard Thomaszewski, Stelian Coros, Damien Gauge, Vittorio Megaro, Eitan Grinspun, and Markus Gross. Computational design of linkage-based characters. *ACM Trans. Graph.*, 33(4), July 2014a. ISSN 0730-0301. doi: 10.1145/2601097.2601143. URL <https://doi.org/10.1145/2601097.2601143>.
- Moritz Bächer, Stelian Coros, and Bernhard Thomaszewski. Linkedit: Interactive linkage editing using symbolic kinematics. *ACM Trans. Graph.*, 34(4), July 2015. ISSN 0730-0301. doi: 10.1145/2766985. URL <https://doi.org/10.1145/2766985>.
- Ran Zhang, Thomas Auzinger, Duygu Ceylan, Wilmot Li, and Bernd Bickel. Functionality-aware retargeting of mechanisms to 3d shapes. *ACM Trans. Graph.*, 36(4), July 2017. ISSN 0730-0301. doi: 10.1145/3072959.3073710. URL <https://doi.org/10.1145/3072959.3073710>.
- Takuto Takahashi, Jonas Zehnder, Hiroshi G. Okuno, Shigeki Sugano, Stelian Coros, and Bernhard Thomaszewski. Computational design of statically balanced planar spring mechanisms. *IEEE Robotics and Automation Letters*, 4:4438–4444, 2019.
- Desai Chen, David I. W. Levin, Wojciech Matusik, and Danny M. Kaufman. Dynamics-aware numerical coarsening for fabrication design. *ACM Trans. Graph.*, 36(4), July 2017. ISSN 0730-0301. doi: 10.1145/3072959.3073669. URL <https://doi.org/10.1145/3072959.3073669>.
- Moritz Bächer, Emily Whiting, Bernd Bickel, and Olga Sorkine-Hornung. Spin-it: Optimizing moment of inertia for spinnable objects. *ACM Trans. Graph.*, 33(4), July 2014. ISSN 0730-0301. doi: 10.1145/2601097.2601157. URL <https://doi.org/10.1145/2601097.2601157>.
- Shayan Hoshiyari, Hongyi Xu, Espen Knoop, Stelian Coros, and Moritz Bächer. Vibration-minimizing motion retargeting for robotic characters. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019.
- Mélina Skouras, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, and Markus Gross. Computational design of actuated deformable characters. *ACM Trans. Graph.*, 32(4), July 2013a. ISSN 0730-0301. doi: 10.1145/2461912.2461979. URL <https://doi.org/10.1145/2461912.2461979>.
- A. Pentland and J. Williams. Good vibrations: Modal dynamics for graphics and animation. In *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '89, page 215–222, New York, NY, USA, 1989. Association for Computing Machinery. ISBN 0897913124. doi: 10.1145/74333.74355. URL <https://doi.org/10.1145/74333.74355>.
- Jernej Barbič and Doug L. James. Real-time subspace integration for st. venant-kirchhoff deformable models. *ACM Trans. Graph.*, 24(3):982–990, July 2005. ISSN 0730-0301. doi: 10.1145/1073204.1073300. URL <https://doi.org/10.1145/1073204.1073300>.



- Klaus Hildebrandt, Christian Schulz, Christoph Von Tycowicz, and Konrad Polthier. Interactive surface modeling using modal analysis. *ACM Trans. Graph.*, 30(5), October 2011. ISSN 0730-0301. doi: 10.1145/2019627.2019638. URL <https://doi.org/10.1145/2019627.2019638>.
- Zherong Pan, Hujun Bao, and Jin Huang. Subspace dynamic simulation using rotation-strain coordinates. *ACM Trans. Graph.*, 34(6), October 2015. ISSN 0730-0301. doi: 10.1145/2816795.2818090. URL <https://doi.org/10.1145/2816795.2818090>.
- Fabian Hahn, Bernhard Thomaszewski, Stelian Coros, Robert W. Sumner, Forrester Cole, Mark Meyer, Tony DeRose, and Markus Gross. Subspace clothing simulation using adaptive bases. *ACM Trans. Graph.*, 33(4), July 2014. ISSN 0730-0301. doi: 10.1145/2601097.2601160. URL <https://doi.org/10.1145/2601097.2601160>.
- Theodore Kim and Doug L. James. Skipping steps in deformable simulation with online model reduction. *ACM Trans. Graph.*, 28(5):1–9, December 2009. ISSN 0730-0301. doi: 10.1145/1618452.1618469. URL <https://doi.org/10.1145/1618452.1618469>.
- Lawson Fulton, Vismay Modi, David Duvenaud, David I. W. Levin, and Alec Jacobson. Latent-space Dynamics for Reduced Deformable Simulation. *Computer Graphics Forum*, 2019. ISSN 1467-8659. doi: 10.1111/cgf.13645.
- Changxi Zheng and Doug L. James. Toward high-quality modal contact sound. *ACM Trans. Graph.*, 30(4), July 2011. ISSN 0730-0301. doi: 10.1145/2010324.1964933. URL <https://doi.org/10.1145/2010324.1964933>.
- Nicolas Bonneel, George Drettakis, Nicolas Tsingos, Isabelle Viaud-Delmon, and Doug James. Fast modal sounds with scalable frequency-domain synthesis. *ACM Trans. Graph.*, 27(3):1–9, August 2008. ISSN 0730-0301. doi: 10.1145/1360612.1360623. URL <https://doi.org/10.1145/1360612.1360623>.
- Jeffrey N Chadwick, Steven S An, and Doug L James. Harmonic shells: a practical nonlinear sound model for near-rigid thin shells. *ACM Trans. Graph.*, 28(5):119–1, 2009.
- Gabriel Cirio, Ante Qu, George Drettakis, Eitan Grinspun, and Changxi Zheng. Multi-scale simulation of nonlinear thin-shell sound with wave turbulence. *ACM Transactions on Graphics (TOG)*, 37(4):110, 2018.
- Nobuyuki Umetani, Jun Mitani, Takeo Igarashi, and Kenshi Takayama. Designing custom-made metallophone with concurrent eigenanalysis. In *New Interfaces for Musical Expression++ (NIME++)*, pages 26–30, 2010.
- Gaurav Bharaj, David I. W. Levin, James Tompkin, Yun Fei, Hanspeter Pfister, Wojciech Matusik, and Changxi Zheng. Computational design of metallophone contact sounds. *ACM Trans. Graph.*, 34(6), October 2015. ISSN 0730-0301. doi: 10.1145/2816795.2818108. URL <https://doi.org/10.1145/2816795.2818108>.
- Andrew Allen and Nikunj Raghuvanshi. Aerophones in flatland: Interactive wave simulation of wind instruments. *ACM Trans. Graph.*, 34(4), July 2015. ISSN 0730-0301. doi: 10.

- 1145/2767001. URL <https://doi.org/10.1145/2767001>.
- Nobuyuki Umetani, Athina Panotopoulou, Ryan Schmidt, and Emily Whiting. Printone: Interactive resonance simulation for free-form print-wind instrument design. *ACM Trans. Graph.*, 35(6), November 2016. ISSN 0730-0301. doi: 10.1145/2980179.2980250. URL <https://doi.org/10.1145/2980179.2980250>.
- Dingzeyu Li, David I. W. Levin, Wojciech Matusik, and Changxi Zheng. Acoustic voxels: Computational optimization of modular acoustic filters. *ACM Trans. Graph.*, 35(4), July 2016. ISSN 0730-0301. doi: 10.1145/2897824.2925960. URL <https://doi.org/10.1145/2897824.2925960>.
- Gaëtan Kerschen, Maxime Peeters, Jean-Claude Golinval, and Alexander F Vakakis. Nonlinear normal modes, part i: A useful framework for the structural dynamicist. *Mechanical Systems and Signal Processing*, 23(1):170–194, 2009.
- Maxime Peeters, Régis Vigué, Guillaume Sérandour, Gaëtan Kerschen, and J-C Golinval. Nonlinear normal modes, part ii: Toward a practical computation using numerical continuation techniques. *Mechanical systems and signal processing*, 23(1):195–216, 2009.
- J. W. Bandler, R. M. Biernacki, and S. H. Chen. Harmonic balance simulation and optimization of nonlinear circuits. In *[Proceedings] 1992 IEEE International Symposium on Circuits and Systems*, volume 1, pages 85–88 vol.1, May 1992. doi: 10.1109/ISCAS.1992.230008.
- Kenneth C. Hall, Kivanc Ekici, Jeffrey P. Thomas, and Earl H. Dowell. Harmonic balance methods applied to computational fluid dynamics problems. *International Journal of Computational Fluid Dynamics*, 27(2):52–67, 2013. doi: 10.1080/10618562.2012.742512. URL <https://doi.org/10.1080/10618562.2012.742512>.
- T. Detroux, L. Renson, and G. Kerschen. The harmonic balance method for advanced analysis and design of nonlinear mechanical systems. In Gaëtan Kerschen, editor, *Nonlinear Dynamics, Volume 2*, pages 19–34, Cham, 2014. Springer International Publishing. ISBN 978-3-319-04522-1.
- Anna Engels-Putzka, Jan Backhaus, and Christian Frey. Forced Response Sensitivity Analysis Using an Adjoint Harmonic Balance Solver. *Journal of Turbomachinery*, 141(3), 01 2019. ISSN 0889-504X. doi: 10.1115/1.4041700. URL <https://doi.org/10.1115/1.4041700>. 031014.
- Suguang Dou and Jakob Søndergaard Jensen. Optimization of nonlinear structural resonance using the incremental harmonic balance method. *Journal of Sound and Vibration*, 334: 239–254, 2015. ISSN 0022-460X. doi: 10.1016/j.jsv.2014.08.023.
- TM Cameron and JH Griffin. An alternating frequency/time domain method for calculating the steady-state response of nonlinear dynamic systems. *Journal of applied mechanics*, 56(1):149–154, 1989.
- Rüdiger Seydel. *Practical bifurcation and stability analysis*, volume 5. Springer Science & Business Media, 2009.

- Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on scientific computing*, 16(5):1190–1208, 1995.
- Eitan Grinspun, Anil N Hirani, Mathieu Desbrun, and Peter Schröder. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 62–67. Eurographics Association, 2003a.
- Simon Pabst, Sybille Krzywinski, Andrea Schenk, and Bernhard Thomaszewski. Seams and Bending in Cloth Simulation. In Francois Faure and Matthias Teschner, editors, *Workshop in Virtual Reality Interactions and Physical Simulation "VRIPHYS" (2008)*. The Eurographics Association, 2008. ISBN 978-3-905673-70-8. doi: 10.2312/PE/vriphys/vriphys08/031-038.
- Federico L. Moro, Alexander Spröwitz, Alexandre Tuleu, Massimo Vespignani, Nikos G. Tsagarakis, Auke J. Ijspeert, and Darwin G. Caldwell. Horse-like walking, trotting, and galloping derived from kinematic motion primitives (kmeps) and their application to walk/trot transitions in a compliant quadruped robot. *Biological Cybernetics*, 107(3), 2013.
- S. Cotton, I. M. C. Olaru, M. Bellman, T. van der Ven, J. Godowski, and J. Pratt. Fastrunner: A fast, efficient and robust bipedal robot. concept and planar simulation. In *2012 IEEE International Conference on Robotics and Automation*, pages 2358–2364, 2012.
- Hongyi Xu, Espen Knoop, Stelian Coros, and Moritz Bächer. Bend-it: Design and fabrication of kinetic wire characters. *ACM Trans. Graph.*, 37(6):239:1–239:15, December 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275089. URL <http://doi.acm.org/10.1145/3272127.3275089>.
- Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. Discrete elastic rods. In *ACM SIGGRAPH 2008 Papers*, SIGGRAPH '08, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781450301121. doi: 10.1145/1399504.1360662. URL <https://doi.org/10.1145/1399504.1360662>.
- Jonas Zehnder, Stelian Coros, and Bernhard Thomaszewski. Designing structurally-sound ornamental curve networks. *ACM Trans. Graph.*, 35(4), July 2016b. ISSN 0730-0301. doi: 10.1145/2897824.2925888. URL <https://doi.org/10.1145/2897824.2925888>.
- Malte Krack, Loic Salles, and Fabrice Thouverez. Vibration prediction of bladed disks coupled by friction joints. *Archives of Computational Methods in Engineering*, 24(3):589–636, 2017.
- Mélina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. Designing inflatable structures. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 33(4), 2014b.
- Xavier Provot et al. Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Graphics interface*, pages 147–147. Canadian Information Processing Society,

1995.

- Robert Bridson, Ronald Fedkiw, and John Anderson. Robust treatment of collisions, contact and friction for cloth animation. *ACM Trans. Graph.*, 21(3):594–603, jul 2002. ISSN 0730-0301.
- Bernhard Thomaszewski, Simon Pabst, and Wolfgang Straßer. Continuum-based strain limiting. *Computer Graphics Forum*, 28(2):569–576, 2009.
- Rony Goldenthal, David Harmon, Raanan Fattal, Michel Bercovier, and Eitan Grinspun. Efficient simulation of inextensible cloth. *ACM Trans. Graph.*, 26(3):49–es, jul 2007. ISSN 0730-0301.
- Alvaro G Perez, Gabriel Cirio, Fernando Hernandez, Carlos Garre, and Miguel A Otaduy. Strain limiting for soft finger contact simulation. In *2013 World Haptics Conference (WHC)*, pages 79–84. IEEE, 2013.
- Rahul Narain, Armin Samii, and James F. O’Brien. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph.*, 31(6), nov 2012. ISSN 0730-0301.
- Ning Jin, Wenlong Lu, Zhenglin Geng, and Ronald P. Fedkiw. Inequality cloth. In *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation, SCA ’17*, New York, NY, USA, 2017a. Association for Computing Machinery. ISBN 9781450350914.
- Erling D Andersen, Cornelis Roos, and Tamas Terlaky. On implementing a primal-dual interior-point method for conic quadratic optimization. *Mathematical Programming*, 95(2):249–277, 2003.
- Hongyi Xu, Funshing Sin, Yufeng Zhu, and Jernej Barbič. Nonlinear material design using principal stretches. *ACM Transactions on Graphics (TOG)*, 34(4):1–11, 2015.
- Yijing Li and Jernej Barbič. Stable anisotropic materials. *IEEE transactions on visualization and computer graphics*, 21(10):1129–1137, 2015.
- Huamin Wang, James F. O’Brien, and Ravi Ramamoorthi. Data-driven elastic models for cloth: Modeling and measurement. *ACM Trans. Graph.*, 30(4), jul 2011. ISSN 0730-0301.
- Eder Miguel, Derek Bradley, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Miguel A Otaduy, and Steve Marschner. Data-driven estimation of cloth simulation models. In *Computer Graphics Forum*, volume 31, pages 519–528. Wiley Online Library, 2012.
- Huamin Wang, James O’Brien, and Ravi Ramamoorthi. Multi-resolution isotropic strain limiting. *ACM Trans. Graph.*, 29(6), dec 2010. ISSN 0730-0301.
- Fernando Hernandez, Gabriel Cirio, Alvaro G Perez, and Miguel A Otaduy. Anisotropic strain limiting. In *Proc. of Congreso Español de Informática Gráfica*, volume 2, 2013.
- Matthias Müller, Nuttapong Chentanez, Tae-Yong Kim, and Miles Macklin. Strain based dynamics. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA ’14*, page 149–157, Goslar, DEU, 2015a. Eurographics Association.

- Elliot English and Robert Bridson. Animating developable surfaces using nonconforming elements. In *ACM SIGGRAPH 2008 Papers*, SIGGRAPH '08, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781450301121.
- Dimitar Dinev, Tiantian Liu, Jing Li, Bernhard Thomaszewski, and Ladislav Kavan. Fepr: Fast energy projection for real-time simulation of deformable objects. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018.
- Jan Bender, Matthias Müller, and Miles Macklin. Position-based simulation methods in computer graphics. In *Eurographics (tutorials)*, page 8, 2015.
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. Projective dynamics: Fusing constraint projections for fast simulation. *ACM transactions on graphics (TOG)*, 33(4):1–11, 2014.
- Matthew Overby, George E Brown, Jie Li, and Rahul Narain. Admm  $\supseteq$  projective dynamics: Fast simulation of hyperelastic models with dynamic constraints. *IEEE transactions on visualization and computer graphics*, 23(10):2222–2234, 2017.
- Juyong Zhang, Yue Peng, Wenqing Ouyang, and Bailin Deng. Accelerating admm for efficient simulation and optimization. *ACM Trans. Graph.*, 38(6), nov 2019. ISSN 0730-0301.
- Miguel Sousa Lobo, Lieven Vandenbergh, Stephen Boyd, and Hervé Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284(1):193–228, 1998. ISSN 0024-3795. International Linear Algebra Society (ILAS) Symposium on Fast Algorithms for Control, Signals and Image Processing.
- Farid Alizadeh and Donald Goldfarb. Second-order cone programming. *Mathematical programming*, 95(1):3–51, 2003.
- Jonas Zehnder, Rahul Narain, and Bernhard Thomaszewski. An advection-reflection solver for detail-preserving fluid simulation. *ACM Trans. Graph.*, 37(4), jul 2018. ISSN 0730-0301.
- Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric numerical integration*, volume 31 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, second edition, 2006. ISBN 3-540-30663-3; 978-3-540-30663-4. Structure-preserving algorithms for ordinary differential equations.
- Rahul Narain, Jonas Zehnder, and Bernhard Thomaszewski. A second-order advection-reflection solver. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 2(2):1–14, 2019.
- Eitan Grinspun, Anil N. Hirani, Mathieu Desbrun, and Peter Schröder. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '03, page 62–67, Goslar, DEU, 2003b. Eurographics Association. ISBN 1581136595.
- Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. Codimensional incremental potential contact. *ACM Trans. Graph.*, 40(4), jul 2021b. ISSN 0730-0301.

- Stephen P. Boyd and Ben Wegbreit. Fast computation of optimal contact forces. *IEEE Transactions on Robotics*, 23(6):1117–1132, 2007.
- Katia Bertoldi, Vincenzo Vitelli, Johan Christensen, and Martin Van Hecke. Flexible mechanical metamaterials. *Nature Reviews Materials*, 2(11):1–11, 2017.
- Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Měch. Stress relief: improving structural strength of 3d printable objects. *ACM Transactions on Graphics (TOG)*, 31(4):48, 2012.
- Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. Build-to-last: Strength to weight 3d printed objects. *ACM Trans. Graph.*, 33(4), jul 2014b. ISSN 0730-0301. doi: 10.1145/2601097.2601168. URL <https://doi.org/10.1145/2601097.2601168>.
- Qingnan Zhou, Julian Panetta, and Denis Zorin. Worst-case structural analysis. *ACM Trans. Graph.*, 32(4):137–1, 2013.
- Lifeng Zhu, Weiwei Xu, John Snyder, Yang Liu, Guoping Wang, and Baining Guo. Motion-guided mechanical toy modeling. *ACM Transactions on Graphics (TOG)*, 31(6):1–10, 2012.
- Duygu Ceylan, Wilmot Li, Niloy J Mitra, Maneesh Agrawala, and Mark Pauly. Designing and fabricating mechanical automata from mocap sequences. *ACM Transactions on Graphics (TOG)*, 32(6):1–11, 2013b.
- Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W. Sumner, Wojciech Matusik, and Bernd Bickel. Computational design of mechanical characters. *ACM Trans. Graph.*, 32(4), jul 2013b. ISSN 0730-0301. doi: 10.1145/2461912.2461953. URL <https://doi.org/10.1145/2461912.2461953>.
- Bernhard Thomaszewski, Stelian Coros, Damien Gauge, Vittorio Megaro, Eitan Grinspun, and Markus Gross. Computational design of linkage-based characters. *ACM Trans. Graph.*, 33(4), jul 2014b. ISSN 0730-0301. doi: 10.1145/2601097.2601143. URL <https://doi.org/10.1145/2601097.2601143>.
- Mélina Skouras, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, and Markus Gross. Computational design of actuated deformable characters. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013b.
- Moritz Bächer, Bernd Bickel, Doug L. James, and Hanspeter Pfister. Fabricating articulated characters from skinned meshes. *ACM Trans. Graph.*, 31(4), jul 2012. ISSN 0730-0301. doi: 10.1145/2185520.2185543. URL <https://doi.org/10.1145/2185520.2185543>.
- Bernd Bickel, Moritz Bächer, Miguel A Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. Design and fabrication of materials with desired deformation behavior. *ACM Transactions on Graphics (TOG)*, 29(4):1–10, 2010.
- Julian Panetta, Abtin Rahimian, and Denis Zorin. Worst-case stress relief for microstructures. *ACM Transactions on Graphics (TOG)*, 36(4):1–16, 2017.

- Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. Two-scale topology optimization with microstructures. *ACM Trans. Graph.*, 36(4), jul 2017. ISSN 0730-0301. doi: 10.1145/3072959.3095815. URL <https://doi.org/10.1145/3072959.3095815>.
- Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. A low-parametric rhombic microstructure family for irregular lattices. *ACM Transactions on Graphics (TOG)*, 39(4):101–1, 2020.
- Kurt Leimer and Przemyslaw Musialski. Reduced-order simulation of flexible meta-materials. Association for Computing Machinery, Inc, 11 2020. ISBN 9781450381703. doi: 10.1145/3424630.3425411.
- Alain Bensoussan, Jacques-Louis Lions, and George Papanicolaou. Studies in mathematics and its applications, 1978.
- Lily Kharevych, Patrick Mullen, Houman Owhadi, and Mathieu Desbrun. Numerical coarsening of inhomogeneous elastic materials. *ACM Transactions on graphics (TOG)*, 28(3): 1–8, 2009.
- Georg Sperl, Rahul Narain, and Chris Wojtan. Homogenized yarn-level cloth. *ACM Trans. Graph.*, 39(4):48, 2020.
- Doug L. James and Dinesh K. Pai. Artdefo: Accurate real time deformable objects. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, page 65–72, USA, 1999. ACM Press/Addison-Wesley Publishing Co. ISBN 0201485605. doi: 10.1145/311535.311542. URL <https://doi.org/10.1145/311535.311542>.
- Bernd Bickel, Moritz Bächer, Miguel A Otaduy, Wojciech Matusik, Hanspeter Pfister, and Markus Gross. Capture and modeling of non-linear heterogeneous soft tissue. *ACM transactions on graphics (TOG)*, 28(3):1–9, 2009.
- David Hahn, Pol Banzet, James M Bern, and Stelian Coros. Real2sim: Visco-elastic parameter estimation from dynamic motion. *ACM Transactions on Graphics (TOG)*, 38(6): 1–13, 2019.
- Ning Jin, Wenlong Lu, Zhenglin Geng, and Ronald P Fedkiw. Inequality cloth. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 1–10, 2017b.
- Matthias Müller, Nuttapong Chentanez, Tae-Yong Kim, and Miles Macklin. Strain based dynamics. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '14, page 149–157, Goslar, DEU, 2015b. Eurographics Association.
- Pengbin Tang, Stelian Coros, and Bernhard Thomaszewski. A Second Order Cone Programming Approach for Simulating Biphasic Materials. *Computer Graphics Forum*, 2022. ISSN 1467-8659. doi: 10.1111/cgf.14626.
- James K. Hahn. Realistic animation of rigid bodies. In *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '88, page 299–308,

- New York, NY, USA, 1988. Association for Computing Machinery. ISBN 0897912756. doi: 10.1145/54852.378530. URL <https://doi.org/10.1145/54852.378530>.
- D. Baraff. Analytical methods for dynamic simulation of non-penetrating rigid bodies. *SIGGRAPH Comput. Graph.*, 23(3):223–232, jul 1989. ISSN 0097-8930. doi: 10.1145/74334.74356. URL <https://doi.org/10.1145/74334.74356>.
- Danny M. Kaufman, Shinjiro Sueda, Doug L. James, and Dinesh K. Pai. Staggered projections for frictional contact in multibody systems. *ACM Trans. Graph.*, 27(5):164:1–164:11, December 2008. ISSN 0730-0301.
- Florence Bertails-Descoubes, Florent Cadoux, Gilles Daviet, and Vincent Acary. A non-smooth newton solver for capturing exact coulomb friction in fiber assemblies. *ACM Trans. Graph.*, 30(1), feb 2011. ISSN 0730-0301. doi: 10.1145/1899404.1899410. URL <https://doi.org/10.1145/1899404.1899410>.
- Kenny Erleben. Methodology for assessing mesh-based contact point methods. *ACM Trans. Graph.*, 37(3), jul 2018. ISSN 0730-0301. doi: 10.1145/3096239. URL <https://doi.org/10.1145/3096239>.
- Albert Peiret, Sheldon Andrews, József Kövecses, Paul G. Kry, and Marek Teichmann. Schur complement-based substructuring of stiff multibody systems with contact. *ACM Trans. Graph.*, 38(5), oct 2019. ISSN 0730-0301. doi: 10.1145/3355621. URL <https://doi.org/10.1145/3355621>.
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. Add: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph.*, 39(6), nov 2020. ISSN 0730-0301. doi: 10.1145/3414685.3417766. URL <https://doi.org/10.1145/3414685.3417766>.
- Yifei Li, Tao Du, Kui Wu, Jie Xu, and Wojciech Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Trans. Graph.*, 42(1), oct 2022. ISSN 0730-0301. doi: 10.1145/3527660. URL <https://doi.org/10.1145/3527660>.
- Ante Qu and Doug L. James. Fast linking numbers for topology verification of loopy structures. *ACM Trans. Graph.*, 40(4), jul 2021. ISSN 0730-0301. doi: 10.1145/3450626.3459778. URL <https://doi.org/10.1145/3450626.3459778>.
- Jonathan M. Kaldor, Doug L. James, and Steve Marschner. Simulating knitted cloth at the yarn level. In *ACM SIGGRAPH 2008 Papers*, SIGGRAPH '08, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781450301121. doi: 10.1145/1399504.1360664. URL <https://doi.org/10.1145/1399504.1360664>.
- Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. Yarn-level simulation of woven cloth. *ACM Trans. Graph.*, 33(6), nov 2014. ISSN 0730-0301. doi: 10.1145/2661229.2661279. URL <https://doi.org/10.1145/2661229.2661279>.
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy R Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M Kaufman. Incremental potential contact:



- intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.*, 39(4): 49, 2020b.
- Zachary Ferguson, Minchen Li, Teseo Schneider, Francisca Gil-Ureta, Timothy Langlois, Chenfanfu Jiang, Denis Zorin, Danny M Kaufman, and Daniele Panozzo. Intersection-free rigid body dynamics. *ACM Transactions on Graphics (SIGGRAPH)*, 40(4), 2021.
- Jiaxian Yao, Danny M. Kaufman, Yotam Gingold, and Maneesh Agrawala. Interactive design and stability analysis of decorative joinery for furniture. *ACM Trans. Graph.*, 36(2), mar 2017. ISSN 0730-0301. doi: 10.1145/3054740. URL <https://doi.org/10.1145/3054740>.
- Peng Song, Chi-Wing Fu, Yueming Jin, Hongfei Xu, Ligang Liu, Pheng-Ann Heng, and Daniel Cohen-Or. Reconfigurable interlocking furniture. *ACM Trans. Graph.*, 36(6), nov 2017. ISSN 0730-0301. doi: 10.1145/3130800.3130803. URL <https://doi.org/10.1145/3130800.3130803>.
- Shiqing Xin, Chi-Fu Lai, Chi-Wing Fu, Tien-Tsin Wong, Ying He, and Daniel Cohen-Or. Making burr puzzles from 3d models. *ACM Trans. Graph.*, 30(4), jul 2011. ISSN 0730-0301. doi: 10.1145/2010324.1964992. URL <https://doi.org/10.1145/2010324.1964992>.
- Peng Song, Chi-Wing Fu, and Daniel Cohen-Or. Recursive interlocking puzzles. *ACM Trans. Graph.*, 31(6), nov 2012. ISSN 0730-0301. doi: 10.1145/2366145.2366147. URL <https://doi.org/10.1145/2366145.2366147>.
- Ziqi Wang, Peng Song, and Mark Pauly. Desia: A general framework for designing interlocking assemblies. *ACM Trans. Graph.*, 37(6), dec 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275034. URL <https://doi.org/10.1145/3272127.3275034>.
- Mélina Skouras, Stelian Coros, Eitan Grinspun, and Bernhard Thomaszewski. Interactive surface design with interlocking elements. *ACM Trans. Graph.*, 34(6), nov 2015. ISSN 0730-0301. doi: 10.1145/2816795.2818128. URL <https://doi.org/10.1145/2816795.2818128>.
- Ziqi Wang, Peng Song, Florin Isvoranu, and Mark Pauly. Design and structural optimization of topological interlocking assemblies. *ACM Trans. Graph.*, 38(6), nov 2019. ISSN 0730-0301. doi: 10.1145/3355089.3356489. URL <https://doi.org/10.1145/3355089.3356489>.
- Samuel Lensgraf, Karim Itani, Yinan Zhang, Zezhou Sun, Yijia Wu, Alberto Quattrini Li, Bo Zhu, Emily Whiting, Weifu Wang, and Devin J. Balkcom. Puzzleflex: kinematic motion of chains with loose joints. In *2020 IEEE International Conference on Robotics and Automation, ICRA 2020, Paris, France, May 31 - August 31, 2020*, pages 6730–6737. IEEE, 2020. doi: 10.1109/ICRA40945.2020.9196854. URL <https://doi.org/10.1109/ICRA40945.2020.9196854>.
- Jonathan Engel and Chang Liu. Creation of a metallic micromachined chain mail fabric. *Journal of Micromechanics and Microengineering*, 17(3):551, feb 2007. doi: 10.1088/0960-1317/17/3/018. URL <https://dx.doi.org/10.1088/0960-1317/17/3/018>.

- Mark Ransley, Peter Smitham, and Mark Miodownik. Active chainmail fabrics for soft robotic applications. *Smart Materials and Structures*, 26(8):08LT02, jul 2017. doi: 10.1088/1361-665X/aa7221. URL <https://dx.doi.org/10.1088/1361-665X/aa7221>.
- F Sausset and G Tarjus. Periodic boundary conditions on the pseudosphere. *Journal of Physics A: Mathematical and Theoretical*, 40(43):12873, oct 2007. doi: 10.1088/1751-8113/40/43/004. URL <https://dx.doi.org/10.1088/1751-8113/40/43/004>.
- Yotam Gingold, Adrian Secord, Han Jefferson, Eitan Grinspun, and Dennis Zorin. A discrete model for inelastic deformation of thin shells. In *Proc. of the ACM/Eurographics Symposium on Computer Animation (SCA) 2004 (Posters track)*, 2004.
- Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. Example-based elastic materials. In *ACM SIGGRAPH 2011 Papers*, SIGGRAPH '11, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450309431.
- Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C. Schwartzman, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Steve Marschner, and Miguel A. Otaduy. Modeling and estimation of internal friction in cloth. *ACM Trans. Graph.*, 32(6), nov 2013. ISSN 0730-0301. doi: 10.1145/2508363.2508389. URL <https://doi.org/10.1145/2508363.2508389>.
- Georg Sperl, Rahul Narain, and Chris Wojtan. Mechanics-aware deformation of yarn pattern geometry. *ACM Trans. Graph.*, 40(4), jul 2021a. ISSN 0730-0301. doi: 10.1145/3450626.3459816. URL <https://doi.org/10.1145/3450626.3459816>.
- Zahra Montazeri, Chang Xiao, Yun Fei, Changxi Zheng, and Shuang Zhao. Mechanics-aware modeling of cloth appearance. *IEEE transactions on visualization and computer graphics*, 27(1):137–150, 2019.
- Georg Sperl, Rahul Narain, and Chris Wojtan. Mechanics-aware deformation of yarn pattern geometry. *ACM Trans. Graph.*, 40(4), jul 2021b. ISSN 0730-0301. doi: 10.1145/3450626.3459816. URL <https://doi.org/10.1145/3450626.3459816>.