Accelerated computational micromechanics

Thesis by
Hao Zhou

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
Defended Jan 28, 2022

ACKNOWLEDGEMENTS

support.

ABSTRACT

The development of new materials is an important component of many cutting edge technologies such as space technology, electronics and medical devices. The properties of advanced materials involve phenomena across multiple scales. The material may be heterogeneous on a scale that is small compared to that of applications, or may spontaneously develop fine-scale structure. Numerical simulation of such phenomena can be an effective tool in understanding the complex physics underlying these materials, thereby assisting the development and refinement of such materials, but can also be challenging.

This thesis develops a new method to exploit the use of graphical processing units and other accelerators for the computational study of complex phenomena in heterogeneous materials. The governing equations are nonlinear partial differential equations, typically second order in space and first order in time. We propose an operator-splitting scheme to solve these equations by observing that these equations come about by a composition of linear differential constraints like kinematic compatibility and balance laws, and nonlinear but local constitutive equations. We formulate the governing equation as an incremental variational principle. We treat both the deformation and the deformation gradient as independent variables, but enforce kinematic compatibility between them as a constraint using an augmented Lagrangian. The resulting local-global problem is solved using the alternating direction method of multipliers. This enables efficient implementation on massively parallel graphical processing units and other accelerators. We use the study of elastic composites in finite elasticity to verify the method, and to demonstrate its numerical performance. We also compare the performance of the proposed method with that of other emerging approaches.

We apply the method to understand the mechanisms responsible for a remarkable in-plane liquid-like property of liquid crystal elastomers (LCEs). LCEs are rubber-like solids where rod-like nematic molecules are incorporated into the main or a side polymer chain. They undergo isotropic to nematic phase transition accompanied by spontaneous deformation which can be exploited for actuation. Further, they display a soft behavior at low temperatures due to the reorientation of the nematic directors. Recent experiments show that LCEs exhibit an in-plane liquid-like behavior under multiaxial loading, where there is shear strain with no shear stress. Our numerical studies of LCEs provides insights into the director distribution and reorientation

in polydomain specimens, and how these lead to the observed liquid-like behavior. The results show good agreement with experimental observations. In addition to providing insight, this demonstrates the ability of our computational approach to study multiple coupled fields.

The core ideas behind the method developed in this thesis are then applied elsewhere. First, we use it to study multi-stable deployable engineering structures motivated by origami. The approach uses two descriptions of origami kinematics, angle/face based approach and vertex/truss based approach independently, and enforces the relationship between them as a constraint. This is analogous to the treatment of kinematic compatibility above where both the deformation and deformation gradient are used as independent variables. The constraint is treated using a penalty. Stable and rigid-foldable configurations are identified by minimizing the penalty using alternate directions, and pathways between stable states are found using the nudged elastic band method. The approach is demonstrated using various examples.

Second, we use a balance law or equilibrium to the problem of determining the stress field from high resolution x-ray diffraction. This experimental approach determines the stress field locally, and errors lead to non-equilibriated fields. It is hypothesized that imposing equilibrium leads to a more accurate stress reconstruction. We use Hodge decomposition to project a non-equilibriated stress field onto the divergence-free (equilibriated) subspace. This projection is numerically implemented using fast Fourier transforms. This method is first verified using synthetic data, and then applied to experimental data obtained from a beta-Ti alloy. It results in large corrections near grain boundaries.

PUBLISHED CONTENT AND CONTRIBUTIONS

Kovachki, N. et al. (2022). "Multiscale modeling of materials: Computing, data science, uncertainty and goal-oriented optimization". In: *Mechanics of Materials* 165, p. 104156. DOI: `10.1016/j.mechmat.2021.104156`.
H. Zhou derived the results presented in one section of this review article.

Zhou, H. and K. Bhattacharya (2022). "Comparison of the performance of various FFT-based methods on accelerators". In: *To be submitted*.
H. Zhou derived the main results.

Zhou, H., M. Grasinger, et al. (2022). "Path finding algorithm of multistable structure". In: *To be submitted*.
H. Zhou and M Grasinger derived the main results of this paper.

Zhou, H., R. A. Lebensohn, et al. (2022). "Imposing equilibrium on experimental 3-D stress fields using Hodge decomposition and FFT-based optimization". In: *Mechanics of Materials* 164, p. 104109. DOI: `10.1016/j.mechmat.2021.104109`.
H. Zhou derived the main results of this paper.

Tokumoto, H. et al. (2021). "Probing the in-plane liquid-like behavior of liquid crystal elastomers". In: *Science Advances* 7.25, eabe9495. DOI: `10.1126/sciadv.abe9495`.
H. Zhou derived the computational results of this paper.

Zhou, H. and K. Bhattacharya (2021). "Accelerated computational micromechanics and its application to polydomain liquid crystal elastomers". In: *Journal of the Mechanics and Physics of Solids* 153, p. 104470. DOI: `10.1016/j.jmps.2021.104470`.
H. Zhou derived the main results of this paper.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

C h a p t e r   1

INTRODUCTION

Advances in many cutting-edge technologies in industries such as aviation, space technology, electronics, and medical devices call for the development of new materials. Applications in aviation and space require materials to be lighter, stronger, and thermally stable. A recent example is the development of magnesium and its alloys (Śliwa et al., 2017) due to the fact that magnesium is the lightest metal, and its alloys possess amongst the highest strength to weight ratio. Applications in medical devices and soft robotic demand materials to be functional and responsive. A recent example is the development of liquid crystal elastomers (Warner and Terentjev, 2003) that are rubber-like solids where the underlying polymer chains contain nematic mesogens leading to mechanical deformation in response to stimuli like heat, electric field, and light. There are two overwhelming challenges in the development of new materials. First, the universe of potential materials is vast and purely empirical development by trial and error is prohibitively time-consuming. Therefore theoretical and computational methods have come to play a crucial role in the search for new materials. Second, the behavior of materials at the scale of the application is the result of complex phenomena at multiple length and time scales. The brute force computation of all these phenomena at all these scales remains prohibitively expensive.

The paradigm of multiscale modeling seeks to address this complexity using a 'divide and conquer' approach shown in Figure 1.1 (Phillips, 2001; Ortiz, Cuitiño, et al., 2001; Fish, 2009; Borst and Ramm, 2011; Van Der Giessen et al., 2020). The complex range of material behavior is first divided into an ordered hierarchy of scales, the relevant mechanisms at each scale are identified and analyzed using theories/tools based on an individual scale, and the hierarchy is put back together by passing information between scales. Importantly, the passage of information between scales is pair-wise, with the larger-scale model both regulating (through average kinematic constraints like the boundary conditions) and averaging (the dynamic response like the stress) the smaller-scale model. The mathematical theory of homogenization (Bensoussan, Lions, and Papanicolaou, 2011; Pavliotis and Stuart, 2008) provides a concrete basis in specialized situations, but the underlying conceptual framework is widely used.
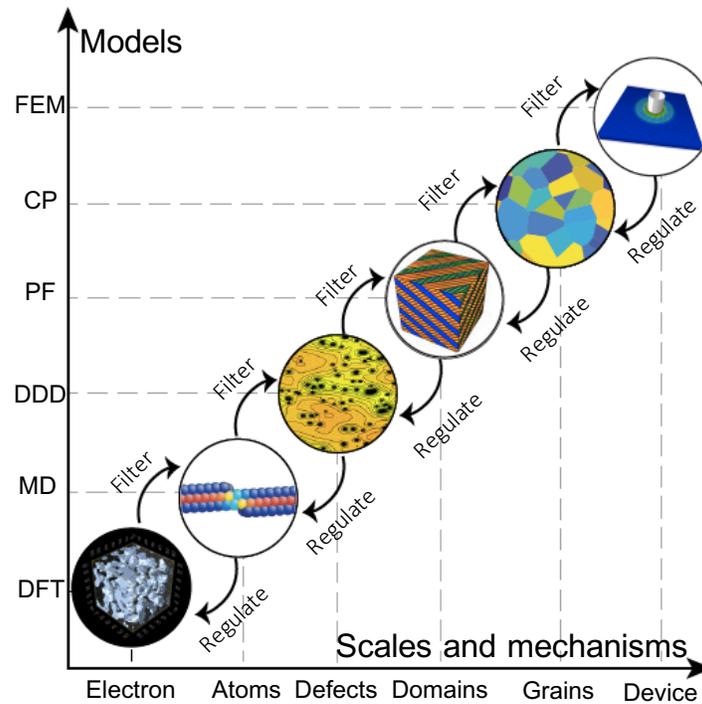
Figure 1.1: Multiscale modeling of materials is a 'divide and conquer' approach to describe the complexity of material behavior.

This thesis addresses the problem of studying multiscale phenomena where both the small and large scales are described by continuum mechanics, a subject that has come to be known as *micromechanics*. Examples include heterogeneous materials including composite materials, polycrystals and polydomain materials in phenomena including elasticity, crystal plasticity, microstructure evolution, and phase transformations (Zohdi and Wriggers, 2005). In all these situations, the state of the solid is described by the deformation gradient and a (set of) internal variables, which evolve to satisfy mechanical equilibrium and kinetic relations subject to initial and boundary conditions. These lead to nonlinear partial differential equations that are second order in space and first order in time, and the solution to such problems can be computationally demanding especially since they have to be performed repeatedly for various initial and boundary conditions to understand the overall behavior. In particular, the scale of the heterogeneity can be small compared to the scale of the application, or the internal variables may oscillate on a fine scale. In such situations, it is common to invoke two-scale analysis where the effective constitutive behavior at the coarse scale is inferred by solving a unit-cell problem with all fine details. This in turn requires the solution of the unit cell problem repeatedly under various

initial and boundary conditions. This has come to be known as *computational micromechanics*.

In pioneering work, Moulinec and P. Suquet (1994) recognized that the periodic boundary conditions that are of interest in computational micromechanics make it natural to use the fast Fourier transform (FFT) in its solution. They rewrite the problem of equilibrium in a heterogeneous linear elastic medium to a Lippmann-Schwinger type equation which they solve iteratively using FFT. Since then FFT-based methods have been widely used in a variety of applications (e.g.thermoelasticity (Anglin, Lebensohn, and Rollett, 2014), elasto-viscoplasticity (Lebensohn and Needleman, 2016), dislocations (Berbenni, Taupin, and Lebensohn, 2020), piezoelectric materials (Vidyasagar, Tan, and Kochmann, 2017), shape-memory polycrystals (Bhattacharya and P.M. Suquet, 2005), and crack prediction of brittle materials (Schneider, 2020)). The method has been understood as a Neumann series approximation (Monchiet and Bonnet, 2012; Milton, 2020). Various approaches to accelerate the convergence have been introduced (Eyre and Milton, 1999; Michel, Moulinec, and P. Suquet, 2000; Monchiet and Bonnet, 2012; Milton, 2020). We refer the reader to Moulinec and Silva (2014) and Moulinec, P. Suquet, and Milton (2018) for a discussion. More recently, the Fourier-Galerkin method of Vondřejc, Zeman, and Marek (2014) has gained in popularity. It formulates the governing equations using a variational approach, uses a Fourier basis for approximation and performs well compared to other FFT-based methods (Mishra, Vondřejc, and Zeman, 2016). Despite the development of various methods and advancing computational capability, computational micromechanics remains computationally challenging.

A recent trend in computing platforms is to complement the central processing units (CPUs) with massively parallel accelerators like graphical processing units (GPUs) (Kirk and Hwu, 2016; Kothe, Lee, and Qualters, 2019). Such accelerators contain thousands of processors, but these are not independent. Instead, they are grouped together in 'warps' that share a memory and execute the same instructions but on possible different data (SIMD). Consequently, they can provide enormous computational power if the calculations are carefully arranged to meet the limitations of the architecture. The computational implementation of computational micromechanics have largely focussed on CPUs, though there is recent work on the use of GPUs. Bertin and Capolungo (2018) used GPUs for the FFT while other researchers have used GPUs for the constitutive update (Mihaila, Knezevic, and Cardenas, 2014; Knezevic and Savage, 2014; Eghtesad et al., 2018). However, these works limit the

use of GPUs to particular aspects of the overall algorithms.

The primary goal of this thesis is to develop methods and algorithms to exploit the enormous computational power of GPUs and similar accelerators for computational micromechanics. We develop a method that exploits the parallel efficiency of GPUs in all aspects of the solution, and call this *accelerated computational micromechanics*. The key observation is to notice that the nonlinear partial differential equations come about through a composition of kinematic compatibility, balance laws (mass or compatibility, momenta, energy) and material behavior. The kinematic compatibility and balance laws are nonlocal but universal and linear if properly formulated. The material behavior is nonlinear and may involve time derivatives, but are local spatially. Thus, the core difficulty, the combination of nonlinearity and nonlocality, comes about because of the composition. Therefore the main idea is to avoid this combination. It is achieved following an operator splitting spirit. The operator splitting method in mechanics has a long history (e.g. Glowinski and Le Tallec, 1989) and a variety of applications. For example, recently, Kirchdoerfer and Ortiz (2016) applied this idea to data-driven computational mechanics. They introduced a data based constitutive cost function and separates it from the compatibility constraint through Lagrange multiplier. Further, it was extended to finite elasticity (Conti, Müller, and Ortiz, 2020) and inelastic problems (Eggersmann et al., 2019; Karapiperis et al., 2021).

The algorithm in this thesis starts with a widely used incremental variational principle, but reformulates the kinematic compatibility as a constraint using the augmented Lagrangian method (Boyd, Parikh, and Chu, 2011). Thus it treats the displacement and displacement gradient as independent variables, but adds a term linear with the constraint and a quadratic penalty to the objective functional. The functional is minimized using the alternating direction method (Boyd, Parikh, and Chu, 2011; Glowinski and Le Tallec, 1989) that first minimizes over the deformation gradient (with no compatibility requirement) and internal variables, and then over the deformation. The first step is local and thus directly amenable to parallel computation. The second step is global, but requires only the solution of a Poisson's equation for which there are fast parallel solvers (Knibbe, Oosterlee, and Vuik, 2011). The final step is to update the Lagrange multiplier, but this is again local and trivially parallel. Additional steps are taken to manage the memory efficiently.

The algorithm is applied to liquid crystal elastomers (LCEs). LCEs are obtained by loosely cross-linking the polymers with mesogens incorporated into the main or side

chains (Brömmel, Kramer, and Finkelmann, 2012; Ula et al., 2018). LCEs are of interest in soft robotics and medical applications because they exhibit a fascinating stimulus-response property owing to the coupling of LC order and rubber elasticity: the LCEs deform in response to a variation in the LC order as the LC order induces the orientation of polymer network strands along the director (principal direction of LC anisotropy). This property enables the actuation of the elastomers by various types of stimulus such as temperature change and electric, magnetic, and optical fields, each of which can drive a variation in the LC order. Various complex deformations can also be programmed by spatially controlling the director configuration in LCEs (White and Broer, 2015). LCEs are therefore attracting considerable attention as a promising material for soft actuator and bio-inspired mechanical devices.

The coupling of LC order and rubber elasticity also results in an unusual mechanical property that is called "soft elasticity." In this well-studied phenomenon, a LCE subjected to uniaxial extension deforms without (or with minimal) additional stress up to a certain stretch (Warner and Terentjev, 2003; Golubović and Lubensky, 1989; Bladon, Terentjev, and Warner, 1993; Conti, DeSimone, and Dolzmann, 2002a). Ideally, the coupling between the nematic order and elasticity implies that the elastomer deforms spontaneously (i.e., in the absence of stress) when the director rotates in the material frame. Briefly, a LCE elongates along the director as it undergoes the isotropic to nematic transition. Therefore, changing the director changes the direction of elongation, manifesting itself as a spontaneous deformation. Consequently, when a LCE is subjected to a uniaxial extension perpendicular to the director, it accommodates the imposed extension to the extent it can by rotating the director, resulting in soft elasticity.

Previous studies have used uniaxial stretching for the characterization of the nonlinear elasticity of LCEs. Uniaxial stretching, however, is only a particular one among all admissible deformations of elastomers. General biaxial strain, in which the two orthogonal strains are varied independently, covers the whole range of accessible homogeneous deformation of incompressible, isotropic materials such as elastomers (Treloar, 1975). The general biaxial stress-strain data provide a definite basis for a full understanding of the large deformation behavior of elastomeric materials (Urayama, 2017). Recent experimental observations by Kenji Urayama and his students reveal a remarkable in-plane liquid-like mechanical behavior of polydomain LCEs under biaxial stretch: The true stresses in the two directions are always identical and governed only by the area change in the loading plane, independently

of the combination of imposed strains in the two directions. This feature defines a previously unidentified state of matter that can vary its shape freely with no extra mechanical energy like liquids when deformed in the plane.

In this thesis, we apply our algorithm to the well-known Bladon-Terentjev-Warner model (Bladon, Terentjev, and Warner, 1993) of LCEs augmented with terms for non-ideality (Biggins, Terentjev, and Warner, 2008) and Franck elasticity (Warner and Terentjev, 2003). We subject a unit cell consisting of a representative volume to various deformation and study the evolution of the director. Our simulations reveal the underlying mechanism for the unusual phenomena. The model with five parameters is not only able to describe the macroscopic observations but also the evolution of directors observed with wide angle X-ray scattering.

This thesis also studies two other closely related problems. First, we show that the use of augmented Lagrangian to describe kinematic constrains are also extremely effective at the larger scale in the study of multi-stable structures such as origami. Origami is the ancient Japanese art of paper-folding, but the ideas underlying origami has proven to be extremely powerful in deployable structure (Lang, 2007). Of particular interest is rigid origami that assumes the facets are rigid and can only rotate along creases. While it is possible to develop a closed form expression (Gattas, W. Wu, and You, 2013) in regular geometries or single degree of freedom systems, applications require irregular and multi degrees of freedom (Y. Li, 2020; Tachi, 2009b; Song and Amato, 2004). Two approaches are common in numerical simulation. The first is angle or plate-based method describes the origami using the dihedral angle between neighboring facets. The dihedral angles around the same vertex are subject to a looping condition (Hull, 2002). The second is a vertex or truss model that treats all creases as trusses, with the constraint that the angle between neighboring trusses of a same facet should be a constant (Y. Li, 2020). Both approaches are subject to non-linear constraints leading to a non-convex energy landscape and have a large number of degrees of freedom, and it is common for numerical searches to get stuck in a local minima. We overcome the nonlinearity using both vertices and facets (just like we use both deformation and deformation gradient) subject to the constraint that the vertex coordinates from these two descriptions should coincide. Most importantly, this constraint is linear. We reformulate this linear constraint using augmented Lagrangian and minimize the energy using alternating direction method as before. We further develop a path finding algorithm to connect any two given states by exploiting the nudged elastic

band method commonly used in materials science (Herbol, Stevenson, and Clancy, 2017). We demonstrate these with various examples.

The second concerns the problem of experimentally identifying stress fields using non-destructive X-ray diffraction (XRD). By combining high-energy X-ray sources (synchrotrons) with tomographic techniques, micro-XRD provides unprecedented information on polycrystalline material at grain and sub-grain scales (Poulsen, 2004; Ludwig et al., 2009; Lienert et al., 2011). Rotating a sample provides multiple scattering patterns which are then compared with simulated scattering patterns to determine the grain orientation and lattice spacing. Elastic strain is obtained by comparing the inferred lattice spacing with those in the undeformed lattice, and elastic stress is then calculated using linear elasticity. However, this technique focuses only on matching the scattering pattern locally at each voxel independently. Thus, errors in the technique give rise to unequilibriated stress, and it is hypothesized that imposing equilibrium can lead to a more accurate stress distribution. We impose equilibrium, known in the field as a micromechanical constraint, at the grain and sub-grain scale, to an arbitrary (non-equilibrated) voxelized stress field obtained by means of synchrotron XED. The method finds the equilibrated stress field closest (in L^2-norm sense) to the measured non-equilibrated stress field, via the solution of an optimization problem. This leads to a projection of an unbalanced stress field onto the divergence-free (equilibrated) space. The extraction of the divergence-free (equilibrated) part of a general (non-equilibrated) field is performed using the Hodge decomposition of a symmetric matrix field. We develop an algorithm based on FFT and apply it to a 3-D XRD experiment of a beta-Ti alloy.

The thesis is organized as follows. We present our accelerated computational micromechanics algorithm in Chapter 2. We discuss the details of the algorithm, several acceleration techniques, and GPU implementation. We also study the convergence rate and convergence with mesh size. We show that the algorithm has good parallel efficiency through the study of scaling. Finally, we compare the performance of our method to other FFT-based computational methods including the Lippmann-Schwinger type method and the Fourier-Galerkin method. In Chapter 3, we show two applications of the algorithm, with a special emphasis on the physics of LCE. We introduce the fundamentals of LCE and study the problem of biaxial stretch of a polydomain specimen. We compare our results to that of experimental observations and explain how the rotation of mesogens give rise to the exotic in-plane liquid-like behavior. Chapter 4 discusses the use of augmented Lagrangian and al-

ternating direction method of minimizers to large scale mechanisms like origami, while Chapter 5 describes the problem of obtaining a self-equilibriated stress field from an arbitrary voxelized symmetric tensor field. In Chapter 6, we recall the main results of this thesis and discuss the future directions.

C h a p t e r   2

ACCELERATED COMPUTATIONAL MICROMECHANICS: METHOD

2.1   Introduction

We present the method of accelerated computational micromechanics in this chapter. We present our computational approach in Section 2.2, and the implementation on GPUs in Section 2.3. We verify our framework and study its numerical performance, using the example of a two-phase composite in finite elasticity in Section 2.4. We also use this example to compare the performance of our method with those of other FFT-based methods in Section 2.5.

2.2   Method

Formulation

It is common in a number of phenomena in solids including plasticity (e.g. Rice, 1971), phase transitions (e.g. Artemev, Jin, and Khachaturyan, 2001) and fracture (e.g. Bourdin, Francfort, and Marigo, 2008) to describe the state of the solid by a deformation gradient F and a set of internal variables η (phase fraction, plastic activity, director field, fracture field, etc.). In the absence of inertia, these are governed by a pair of coupled nonlinear partial differential equations:

$$\nabla \cdot (W_F(\nabla u, \eta, x)) = 0, \tag{2.1}$$

$$W_\eta(\nabla u, \eta, x) + D_v(\eta_t, x) = 0 \tag{2.2}$$

where $u : \Omega \times [0, T] \to \mathbb{R}^3$ is the deformation, $F : \Omega \times [0, T] \to \mathbb{R}^{3\times3}$ is the deformation gradient, $\eta : \Omega \times [0, T] \to \mathbb{R}^d$ is an internal variable or order parameter, $v = \eta_t$ is the rate of change of the internal variable, $W : \mathbb{R}^{3\times3} \times \mathbb{R}^d \times \Omega \to \mathbb{R}$ is the stored (elastic) energy density, $D : \mathbb{R}^d \times \Omega \to \mathbb{R}$ is the dissipation potential (rate of dissipation as a function of the rate of change of internal variable) that governs the evolution of the internal variables, $\Omega \subset \mathbb{R}^3$ is the reference domain assumed to be simply connected, T is the final time of interest, and the subscripts denote partial differentiation. Here and in what follows, all operators and identities are defined on the reference configuration. The first of the two equations describes the mechanical equilibrium, and the second describes the kinetic relation or configurational equilibrium that governs the evolution of the internal variables (we refer the reader to Ortiz

and Stainier (1999) and Mielke (2006) for details of rewriting the kinetic relation in this gradient-flow form).

The first equation 2.1 is a second-order nonlinear elliptic partial differential equation in space with time as a parameter. The second is typically first order in time with space as a parameter, and also nonlinear. Further, in rate independent phenomena, the dissipation potential may not be continuously differentiable and the second equation is interpreted as a differential inclusion (Mielke, 2006). These make the system (2.1, 2.2) challenging to solve. In particular, traditional finite-element approaches require significant amounts of communication to solve in parallel.

However, these equations arise from the agglomeration of a number of simpler equations:

$$\text{Compatibility:} \qquad F = \nabla u \iff \text{curl } F = 0, \qquad (2.3)$$

$$\text{Equilibrium:} \qquad \nabla \cdot S = 0, \qquad (2.4)$$

$$\text{Stress-Strain Relation:} \qquad S = W_F(F, \eta, x), \qquad (2.5)$$

$$\text{Kinetic relation:} \qquad . \qquad (2.6)$$

Note that the field equations–the first two–are linear and universal, while the constitutive updates–the last two–are local (albeit nonlinear). We want to exploit this in the GPU implementation.

In order to do so, consider an implicit time discretization

$$\text{Compatibility:} \qquad F^{n+1} = \nabla u^{n+1} \iff \text{curl } F^{n+1} = 0, \qquad (2.7)$$

$$\text{Equilibrium:} \qquad \nabla \cdot S^{n+1} = 0, \qquad (2.8)$$

$$\text{Stress-Strain Relation:} \qquad S^{n+1} = W_F(F^{n+1}, \eta^{n+1}, x), \qquad (2.9)$$

$$\text{Kinetic relation:} \qquad W_\eta(F^{n+1}, \eta^{n+1}, x) + D_v\left(\frac{\eta^{n+1} - \eta^n}{\Delta t}, x\right) = 0. \qquad (2.10)$$

This can be written as the following variational problem (e.g., Ortiz and Stainier, 1999):

$$u^{n+1}, \eta^{n+1} = \text{argmin} \int_\Omega \left(W(\nabla u, \eta, x) + \Delta t D\left(\frac{\eta - \eta^n}{\Delta t}, x\right)\right) dx \qquad (2.11)$$

or as the following constrained variational problem

$$F^{n+1}, \eta^{n+1} = \text{argmin}_{\text{curl } F=0} \int_\Omega \left(W(F, \eta, x) + \Delta t D\left(\frac{\eta - \eta^n}{\Delta t}, x\right)\right) dx. \qquad (2.12)$$

We rewrite this constrained variational problems using the *augmented Lagrangian* method or method of multipliers (e.g. Glowinski, 2015). Given $\rho > 0$, we seek to find the saddle point:

$$\int_{\Omega} \left(W(F, \eta, x) + \Delta t D \left(\frac{\eta - \eta^n}{\Delta t}, x \right) + \Lambda \cdot (\nabla u - F) + \frac{\rho}{2} |\nabla u - F|^2 \right) dx \qquad (2.13)$$

for u, F and the Lagrange multiplier $\Lambda : \Omega \to \mathbb{R}^{3 \times 3}$. We solve this problem using the *alternating direction method of multipliers (ADMM)* (Glowinski, 2015; Glowinski, 2016; Glowinski and Le Tallec, 1989) which is an iterative method.

At the $(n+1)^{th}$ time step, given $F^n, \eta^n, u^n, \Lambda^n$, set $F_0 = F^n, \eta_0 = \eta^n, u_0 = u^n, \Lambda_0 = \Lambda^n$ and iterate over i

- *Step 1: Local problem.* Update F, η by solving at each x

$$W_F(F_{i+1}, \eta_{i+1}, x) - \Lambda_i - \rho(\nabla u_i - F_{i+1}) = 0, \qquad (2.14)$$

$$W_\eta(F_{i+1}, \eta_{i+1}, x) + D_v \left(\frac{\eta_{i+1} - \eta^n}{\Delta t}, x \right) = 0. \qquad (2.15)$$

- *Step 2: Helmholtz projection.* Update u by solving the partial differential equation

$$-\Delta u_{i+1} = \nabla \cdot \left(-F_{i+1} + \frac{1}{\rho} \Lambda_i \right). \qquad (2.16)$$

- *Step 3: Update Lagrange multiplier.* Update Λ as

$$\Lambda_{i+1} = \Lambda_i + \rho(\nabla u_{i+1} - F_{i+1}). \qquad (2.17)$$

- *Step 4: Check for convergence.* Check both primal and dual feasibility:

$$r_p := ||\nabla u_{i+1} - F_{i+1}||_{L^2} \leq r_p^{\text{tolerance}}, \quad r_d := \rho/\mu ||\nabla u_{i+1} - \nabla u_i||_{L^2} \leq r_d^{\text{tolerance}} \qquad (2.18)$$

for given $r_p^{\text{tolerance}}, r_d^{\text{tolerance}}$ and representative elastic modulus μ

until convergence and update $F^{n+1} = F_i, \eta^{n+1} = \eta_i, u^{n+1} = u_i, \Lambda^{n+1} = \Lambda_i$.

We now make a series of comments about the proposed approach.

Connection to Augmented-Lagrangian method. The approach above is a special case of the Augmented-Lagrangian method of Michel, Moulinec, and P. Suquet (2000). In our language, they study the functional

$$\int_\Omega \left(W(F, \eta, x) + \Delta t D \left(\frac{\eta - \eta^n}{\Delta t}, x \right) + \Lambda \cdot (\nabla u - F) + \frac{\rho}{2} (\nabla u - F) \cdot \mathbb{C}_0 (\nabla u - F) \right) dx$$
(2.19)

where $\rho \mathbb{C}_0$ is the modulus of a comparison medium. They study convergence with ρ at low spatial resolution and use this value at high spatial resolution. Our approach is a special case of their method with $\mathbb{C}_0 = \mathrm{Id}$. The advantage of doing so is that it leads to a Helmholtz projection in Step 2 as opposed to an operator that depends on the comparison medium. Further, we tune ρ during iteration (cf. (2.23)) and explore the approximate solution of the local problem. This algorithm was also used for finite elasticity (no internal variables) in the context of general boundary value problems and finite elements by Glowinski and Le Tallec (1982) and Glowinski and Le Tallec (1984).

We comment that we could have chosen a different constant ρ' instead of ρ in equation (2.17) that updates the Lagrange multiplier (Step 3). An analysis of Eckstein and Bertsekas (1992) suggests that using $\rho' > \rho$ may improve convergence. We have not explored this in our work. We also note that the original formulation of Michel, Moulinec, and P. Suquet (2000) used a modulus $\mathbb{D}_0$ distinct from the comparison medium $\mathbb{C}_0$ in this step, though the common implementations use $\mathbb{C}_0 = \mathbb{D}_0$.

Connection to Hu-Washizu variational principle. The Hu-Washizu and other mixed methods have been used widely in mechanics (e.g. Washizu, 1968). In the Hu-Washizu variational principle of finite elasticity, one seeks the stationary point of

$$\int_\Omega (W(F) - S \cdot F - u \cdot (\nabla \cdot S)) \, dx$$
(2.20)

(up to boundary terms) over F, S, u. Integrating it by parts, we see that the functional is equivalent to

$$\int_\Omega (W(F) - S \cdot (F - \nabla u)) \, dx$$
(2.21)

(up to boundary terms). This is the same as (2.13) with $\rho = 0$ (and $S = \Lambda$). Thus, our approach with $\rho = 0$ reduces to the Hu-Washizu variational principle. It is well-known that such variational principles are delicate with regard to convergence which we avoid with $\rho > 0$.

Connection to split Bregman method. The split Bregman method introduced by Goldstein and Osher (2008) (based on an earlier method proposed by Bregman (1966) to regularize convex problems) seeks to minimize a functional $J(F)$ amongst all minimizers of another functional $H(F)$. In our case, we take these functionals to be

$$J(F) = \int_\Omega \left(W(F, \lambda, x) + \Delta t D \left(\frac{\lambda - \lambda^n}{\Delta t}, x \right) \right) dx, \quad H(F) = \int_\Omega |\nabla u - F|^2 dx. \quad (2.22)$$

It has been proven to be equivalent to the augmented Lagrangian approach (Yin et al., 2008) and used in a variety of problems including image sensing (Goldstein and Osher, 2008; Yin et al., 2008), free boundary problems (Giga and Ueda, 2020) and microstructure formation (Jaramillo and Venkataramani, 2019).

Parallel implementation. Step 1 is local, and can be solved trivially in parallel. However, it is (generally) nonlinear and therefore requires an iterative approach. In this work, we solve it using a steepest descent method with backtracking line search. A potential difficulty is that different spatial points may require different number of iterations to converge, and we address it below. Step 2 leads to a universal Poisson's equation for which there are a number of effective parallel solvers. In this work, we consider problems with periodic boundary conditions and therefore use the fast Fourier transform. Step 3 is a trivial local update, and step 4 a simple check. Thus, this iterative algorithm can be implemented effectively using accelerators like Graphical Processing Units (GPUs) as we demonstrate in the subsequent sections.

Convergence. Boyd, Parikh, and Chu (2011) prove the convergence of the spatially discretized version of the algorithm for ρ large enough under the hypotheses that W is convex and the unaugmented functional with $\rho = 0$ has a saddle point. They also review improvements of this result in the literature. However, it is not natural to expect W to be convex in F in finite deformation. Here are results in the case of elasticity where there is no internal variable. Glowinski and Le Tallec (1982) show that the weak form of equilibrium equation of incompressible elasticity is equivalent to the weak form of the first variation of the augmented Lagrangian functional. Further, Glowinski and Le Tallec (1984) show in the case of Mooney-Rivlin materials that the finite element discretization of this iterative approach converges for sufficiently large ρ. Furthermore, they show that the finite element solutions converge to the solution of the continuous problem. Their arguments can be generalized to a larger class of incompressible, isotropic, polyconvex, hyperelastic

materials (Glowinski and Le Tallec, 1989). However, the general case with internal variables remains open.

Connection to physics. This iterative method also has a close connection to the physics as the various steps and convergence criteria can be identified with the governing equations (2.3) through (2.6). Indeed, we can see that (2.15) is a time-discretized version of the kinetic relation (2.6). As the method converges, the primal convergence $(2.18)_1$ guarantees that $F \approx \nabla u$ or satisfaction of the compatibility equation (2.3). We now turn to the stress-strain relation (2.5). As a result of the the primal convergence, we observe that (2.14) becomes $W_F \approx \Lambda$. This tells us that the Lagrange multiplier converges to the stress and (2.14) approximates the stress-strain relation (2.5). Finally, we show in the appendix that the dual convergence is equivalent to the equilibrium condition (2.4).

Penalty parameter. We have noted above that the method is known to converge under suitable hypothesis on W, D for all ρ sufficiently large (Boyd, Parikh, and Chu, 2011; Glowinski and Le Tallec, 1984). However, the rate of convergence depends critically on the choice of ρ. We adapt an idea proposed by He, H. Yang, and S. Wang (2000) to adaptively change ρ with iteration guided by r_p and r_d. A large ρ better enforces compatibility (cf. 2.16) and leads to a faster drop of the primal error r_p. However, it leads to a poor enforcement of the constitutive equation (cf. 2.14) and slower drop of the dual error r_d. Conversely, a small ρ leads to a faster drop of the dual error r_d, but a slower drop of the primal error r_p. After checking for convergence, we adjust ρ as

$$
\rho_{i+1} = \begin{cases} \gamma \rho_i, & \text{if} \quad r_p/r_d > \tau \\ \max\{\rho_i/\gamma, \rho_{\min}\}, & \text{if} \quad r_d/r_p > \tau \\ \rho_i, & \text{if} \quad else \end{cases} \tag{2.23}
$$

for given γ, τ $\rho_{\min}$. We take γ and τ to be 1.3 and 10 respectively, to avoid too frequent updates. $\rho_{\min}$ enforces the requirement that ρ remains large enough for the method to converge and the choice of $\rho_{\min}$ depends on W. An important observation is that this is enabled by the fact that the Laplace operator in (2.16) is independent of ρ. We report a numerical study in Section 2.4 and specifically in Figure 2.4.

Approximate solution of the local problem. While the local problem (2.14), (2.15) is trivially parallelizable, a potential problem is that different points may re-

quire a different number of iterations to converge to a given point-wise residual error. Indeed, in practice (as we shall see in Section 2.4), a few isolated points require a very large number of iterations to converge compared to the rest. Unfortunately, the calculation can not advance to the next step until the last point has converged, and consequently, these slowly-converging points can add significantly to the computational cost. However, Eckstein and Bertsekas (1992) have proved that for convex W, the algorithm converges when the the local error is summable. In other words, it is not necessary to solve the local equations (2.14), (2.15) at every point (i.e., in L^∞ norm), but it suffices to solve them in some L^p norm for appropriate $2 \le p < \infty$. Working in $p = 2$, set the local residual to be

$$r_l = \frac{1}{\mu}||W_F(F_{i+1}, \eta_{i+1}, x) - \Lambda^n + \rho(\nabla u_i - F_{i+1})||_{L^2}^2. \tag{2.24}$$

We see from (A.4) in the Appendix that equilibrium is still satisfied if the both the dual and local residual go to zero. Thus, satisfaction of the equilibrium equation does not require pointwise convergence. Further savings can be achieved by keeping the local residual large in the initial (global) iterations, and gradually reducing it as (global) convergence is achieved (Boyd, Parikh, and Chu, 2011). We explore two strategies in Section 2.4. In one, we maintain a balance between r_l and r_d while in the other we require a fixed fraction of local points to converge.

While these approaches reduce the number of iterations of the local problem, computing either the local residual r_l or the fraction of converged points requires a communication between the various points which can be expensive in an accelerator. Therefore it is necessary to balance the cost of iteration and the cost of checking convergence. We study this balance in Section 2.4.

Gradient internal variables. In certain problems like phase transitions and microstructure evolution (like the one we shall study in Chapter 3), the state of the material is described not only by an internal variable, but also the gradient of the internal variable, i.e,. $W = W(F, \eta, \nabla \eta, x)$. These can be incorporated into the method in two ways.

The first approach is to treat the gradient of the internal variable in much the same way that we treat the deformation gradient. But this requires some care to maintain the linear structure of step 2. We introduce two internal variables, ℓ and H and enforce the constraints $\ell = \eta$, $H = \nabla \eta$ using the augmented Lagrangian, i.e.,

consider the Lagrangian density

$$W(F, \ell, H, x) + \Delta t D + \Lambda \cdot (\nabla u - F) + \rho/2 |\nabla u - F|^2$$
$$+ \mu \cdot (\nabla \eta - H) + \xi/2 |\nabla \eta - H|^2 + \kappa \cdot (\ell - \eta) + \zeta/2 |\ell - \eta|^2$$

with additional Lagrange multipliers μ, κ and penalty parameters ξ, ζ. We minimize the Lagrangian with respect to F, ℓ and H in step 1 and with respect to u and η in step 2. We then update all the Lagrange multipliers in step 3 and check convergence in step 4. Note that the equation in step 2 describing η is not the Poisson's equation but includes a linear term in η; still, it can be treated as before, similar to Section 2.2.

The second approach is to use the value of the gradient from the previous iteration. In other words, we rewrite (2.15) as

$$W_\eta(F_{i+1}, \eta_{i+1}, \nabla \eta_i, x) + D_v \left(\frac{\eta_{i+1} - \eta^n}{\Delta t}, x \right) = 0. \tag{2.25}$$

While this is approximate, it is effective. In most models, the gradient of the internal variable is introduced as a way of introducing a length scale, and this approximation does so effectively. We use this second approach in Chapter 3.

Implementation in the periodic setting

It is common in computational micromechanics to consider a material that is periodic and a representative volume element that is a unit cell. In other words, $\Omega = (-L, L)^3$ and $W(F, \eta, x), D(\eta_t, x)$ are periodic in x. Further, the average deformation gradient $\langle F \rangle$ is prescribed in strain-controlled simulations, the average stress $\langle S \rangle$ is prescribed in stress-controlled simulations and some combination with some components of $\langle F \rangle$ and complementary components of $\langle S \rangle$ are prescribed in mixed simulations (e.g., plane stress where the average planar stretch is prescribed – see for example Lucarini and Segurado, 2019 for a comprehensive discussion of mixed boundary conditions). Above, $\langle \cdot \rangle$ denotes spatial average.

We assume that the resulting deformation gradient F, internal variable η and stress S are also periodic in x. This implies that the deformation u is periodic up to a linear function; i.e., $u - \langle F \rangle x$ is periodic. It is then natural to solve Step 2 (2.16) using fast Fourier transforms (FFT). In Fourier space, (2.16) becomes local (i.e., can be solved at each ξ) as

$$\hat{u}(\xi) = -\frac{(\hat{F} - \rho^{-1} \hat{\Lambda}) i \xi}{|\xi|^2}. \tag{2.26}$$

where $\hat{f}(\xi)$ denotes the Fourier transform of $f(x)$ and i is the imaginary unit. Thus, we solve Steps 1, 3, and 4 in real space and Step 3 in Fourier space using FFT and iFFT (inverse fast Fourier transform) to go back and forth between them. Specifically, we consider a regular $N \times N \times N$ grid (for N even) on Ω in real space, and the corresponding $N \times N \times N$ on the domain $(-\frac{\pi}{L}, \frac{\pi}{L})^3$ in Fourier space. We define $\hat{f}(\xi)$ as the discrete Fourier transform on this space and use FFT to evaluate it.

A couple of comments are in order. First, it is convenient to work with

$$\tilde{u} = u - \langle \nabla u \rangle x, \quad \tilde{F} = F - \langle F \rangle, \quad \tilde{\Lambda} = \Lambda - \langle \Lambda \rangle \tag{2.27}$$

which are all periodic (recall that u is not necessarily periodic). This is also convenient since the boundary conditions are prescribed in terms of $\langle F \rangle$ or $\langle S \rangle$ (recall Λ converges to S).

Second, the material may be heterogeneous and W, D may be discontinuous functions of x in many problems of interest. In such situations, F, η may be discontinuous and thus the use of Fourier transforms to solve for u may lead to spurious oscillations. An approach around this that has proved very effective in various problems (Berbenni, Taupin, Djaka, et al., 2014; Lebensohn, Kanjarla, and Eisenlohr, 2012; Vidyasagar, Tan, and Kochmann, 2017) using the closely related Lippmann-Schwinger approach is to replace the discrete Fourier transform of the derivatives with the discrete Fourier transform of the central differences:

$$\widehat{u_{i,j}}(\xi) = i\hat{u}_i(\xi)\xi_j \quad \text{with} \quad \frac{\widehat{u_i(x+he_j)} - \widehat{u_i(x-he_j)}}{2h}(\xi) = i\hat{u}_i(\xi)\frac{\sin(h\xi_j)}{h},$$
$$\tag{2.28}$$

$$\widehat{u_{i,jj}}(\xi) = -\hat{u}_i(\xi)|\xi_j|^2 \quad \text{with} \quad \frac{\widehat{u_i(x+he_j)} + \widehat{u_i(x-he_j)} - 2\widehat{u_i(x)}}{h^2}(\xi) = -\hat{u}_i(\xi)\frac{4\sin^2(\frac{h\xi_j}{2})}{h^2}$$
$$\tag{2.29}$$

where $\widehat{u_i(x+he_j)}$ is the Fourier transform of $u_i(x+he_j)$, etc. Since $\sin\alpha/\alpha < 1$, this is equivalent to a high frequency filter and suppresses the spurious oscillations. We refer the reader to Zhou et al. (2022) for further discussion.

Quadratic functionals and Bloch waves

We conclude this section with the discussion of a closely-related linear problem that arises in the study of stability of periodic solutions to nonlinear problems. Let

$\mathbb{L}_{ijkl}(x)$ be a periodic fourth order tensor field of period 1 with $\mathbb{L}_{ijkl} = \mathbb{L}_{klij}$. We are interested in evaluating

$$\beta_k = \min_{v \in \text{Re } \mathcal{A}_k} \int_{\Omega_0} \frac{1}{2} \overline{v}_{i,j} \mathbb{L}_{ijkl} v_{k,l} dx = \min_{v \in \text{Re } \mathcal{A}_k} \int_{\Omega_0} \nabla \overline{v} \cdot \mathbb{L} \nabla v \, dx \qquad (2.30)$$

over an admissible class of functions that are unit Bloch waves:

$$\mathcal{A}_k = \left\{ v \in H^1(\Omega_0) : ||v||_{L^2(\Omega_0)} = 1, \ v(x) = p(x) \exp(i\omega_k \cdot x), \ \omega_k = \left\{ \frac{2\pi}{k_1}, \frac{2\pi}{k_2} \right\}, \ p \ 1 - \text{periodic} \right\} \qquad (2.31)$$

for $k = \{k_1, k_2\}$ with k_i integers. Above, $\Omega_0 = (0, 1)^2$ is the unit square, and $\overline{v}$ denotes the complex conjugate of v. As before, we use the augmented Lagrangian formulation to write

$$\beta_k = \min_{v \in \mathcal{A}_k, F \in L^2(\Omega_0)} \max_{\Lambda \in L^2(\Omega_0)} \int_{\Omega_0} \left(\frac{1}{2} F \cdot \mathbb{L} F + \Lambda \cdot (\nabla v - F) + \frac{\rho}{2} |\nabla v - F|^2 \right) dx. \qquad (2.32)$$

Recalling that v is a Bloch wave, and setting $F(x) = G(x) exp(i\omega_k \cdot x)$, $\Lambda(x) = g(x) exp(i\omega_k \cdot x)$, where $G, g \in L^2(\Omega_0)$ extended periodically, it follows that

$$\beta_k = \min_{p \in \mathcal{P}, G \in L^(\Omega_0)} \max_{g \in L^(\Omega_0)} \int_{\Omega_0} \left(\frac{1}{2} G \cdot \mathbb{L} G + g \cdot ((i\omega_k + \nabla)p - G) + \frac{\rho}{2} |(i\omega_k + \nabla)p - G|^2 \right) dx \qquad (2.33)$$

where $\mathcal{P} = \{p \in H^1(\Omega_0) : ||p||_{L^2(\Omega_0)} = 1, p \ 1 - \text{periodic}\}$. This saddle point problem can be solved as before using ADMM. Given G_i, p_i, g_i,

- *Step 1': Local problem.* Update G: $G_i = (\mathbb{L} + \rho\mathbb{I})^{-1}(g_i + \rho(i\omega_k + \nabla)p^n)$;

- *Step 2': Global update.* Update p: $(i\omega_k + \nabla)^2 p_{i+1} = (i\omega_k + \nabla) \cdot (G_{i+1} - \rho^{-1}g_i)$;

- *Step 3': Update Lagrange multiplier.* Update g: $g_{i+1} = g_i + \rho((i\omega_k + \nabla)p_{i+1} - G_{i+1})$;

- *Step 4': Check for convergence.*

Note that the global problem can be solved trivially in Fourier space.

Thus, a quadratic functional can be minimized over Bloch waves in the original unit cell with a slight modification of our algorithm.
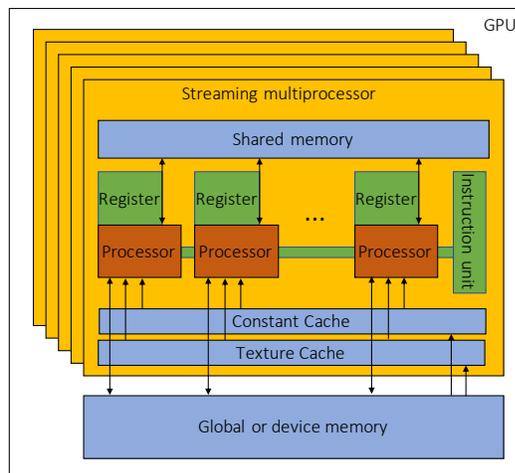
Figure 2.1: A schematic representation of the architecture of a general purpose graphical processing unit (adapted from Preis et al. (2009))

2.3 GPU implementation

We begin with a brief introduction to general purpose GPUs and their use in computing, referring the reader to Kirk and Hwu (2016) for details. A compute node typically consists of a single CPU and one or more general purpose GPUs. While the clock speed of a CPU is faster than that of the GPU, the presence of thousands of cores and the architecture enables faster overall performance if properly organized. A schematic representation of a general purpose GPU is shown in Figure 2.1. It consists of a global or device memory and a number of streaming multiprocessors (SMs). Each SM in turn contains a number of cores or processors that have access to a shared memory, various registers and an instruction unit. All SMs also have access to the constant cache and the texture cache. The calculation is organized in threads with each processor typically executing a single thread. The threads are organized into warps. All threads in the warp work following the "single instruction multiple data (SIMD)" organization, i.e. each processor executes the same instruction concurrently with possibly different data. If there are many conditional instructions and different threads fall into different conditions, then it leads to a situation described as 'warp divergence' where each of the conditional instructions is executed in a serial manner. It is important to avoid warp divergence.

The exchange of data between a CPU and a GPU is slow, and therefore has to be minimized. Even though significantly faster, the exchange of data within a GPU between the global memory and the shared memory of a SM is also slow. However, this can be made faster using a parallel transfer strategy called 'coalesced memory

---

Algorithm 1: Implementation on GPUs

Given an initial value of u^0, F^0, η^0 and macroscopic strain path $\overline{F}(t)$;

Step 0: Initialize:

   Place $\rho, \epsilon_p, \epsilon_d, \gamma_p, \gamma_d, f$ in constant cache;

   Place $u^0, F^0, \eta^0, \Lambda^0$ in global memory,

while $t < t_{max}$ **do**

    Initialize $F_0 = F^n, \eta_0 = \eta^n, u_0 = u^n, \Lambda_0 = \Lambda^n$

    while $r_p > \epsilon_p$ *or* $r_d > \epsilon_d$ **do**

       Step 1: Local problem.

       while $r_l > f r_d$ **do**

          Using kernel function:

             · Move $F_i, u_i, \eta_i, \Lambda_i$ to shared memory;

             · Fixed number of iterations to solve (2.14), (2.15) for F_{i+1}, η_{i+1};

             · Move F_{i+1}, η_{i+1} to global memory

          Compute r_l ;   `/* cuBlas */`

       end

       Step 2: Helmholtz projection.

         · FFT $\widehat{F}_{i+1}, \widehat{\eta}_{i+1}$;   `/* cuFFT */`

         · Using kernel function:

            · Move $\widehat{F}_{i+1}, \widehat{\eta}_{i+1}$ to shared memory;

            · Find $\widehat{u}_{i+1}$ from (2.26);

            · Move $\widehat{u}_{i+1}$ to global memory;

         · iFFT $\widehat{u}_{i+1}$;   `/* cuFFT */`

       Step 3: Update Lagrange multiplier. Find Λ_{i+1} from (2.17) ;

        `/* cuBlas */`

       Step 4: Check for convergence Compute r_p, r_d from (2.18) ;

        `/* cuBlas */`

    end

    Update $t, F^{n+1}, u^{n+1}, \eta^{n+1}, \Lambda^{n+1}$

end

`cuBlas` and `cuFFT` are built-in CUDA kernal functions used in the indicated
 steps.

---

access' (B. Wu et al., 2013) when a one–to–one mapping can be created between a thread and a segment of the global memory.

The implementation of the algorithm described earlier is presented in Algorithm 1. We work on a compute node consisting of a 14-core Intel Broadwell CPU and four Nvidia Tesla P100 GPUs. Each GPU contains 16GB of (global or device) memory and 56 SMs with 64 cores each for a total of 3584 cores, and has a double precision performance of 4.7 teraFLOPS. The P100 GPUs enable the compute unified de-

vice architecture (CUDA) platform with the standard programming language C++. CUDA uses warps of 32 threads (Kirk and Hwu, 2016). The implementation is general and can be expanded to other platforms.

The algorithm takes advantage of the GPU architecture in various ways:

- All calculations are performed on the GPU. The CPU is used only for kernel function calls (i.e., to provide instructions to the GPU), initialization and output of results.

- The exchange of data between CPU and GPU is limited to the first initialization and to write results.

- Following first initialization, all data is kept on the GPU global memory during the entire calculation. Since the proposed algorithm uses the result of the previous time step to initialize the current time step, it is not necessary to perform any GPU/CPU transfer between time steps when there is no need to write the result.

- Global constants like the penalty and tolerance parameters are kept in constant cache.

- The local step 1 of the proposed algorithm is well-suited for SIMD since the same equations are solved independently at each point. Further, the structure enables optimization of the exchange between global and shared memory within the GPU in two ways. The data can be kept in shared memory within the SM during the local iterations. Therefore, the exchange between global and shared memory is limited to the initialization and to the final output of the local iteration. Even these transfers can exploit the coalesced memory access since each thread (point) only requires data stored at a particular location in global memory.

- The Helmholtz projection is local in Fourier space (cf. (2.26)). Therefore is it well-suited for SIMD, and can take advantage of coalesced memory access.

- The local nature of the local step 1 in real space and the Helmholtz projection in Fourier space avoid warp divergence.

- The computation of the L^2 norms in the approximate solution of step 1, the Lagrange multiplier update (step 3), the convergence check (step 4) are executed using basic linear algebra operations.
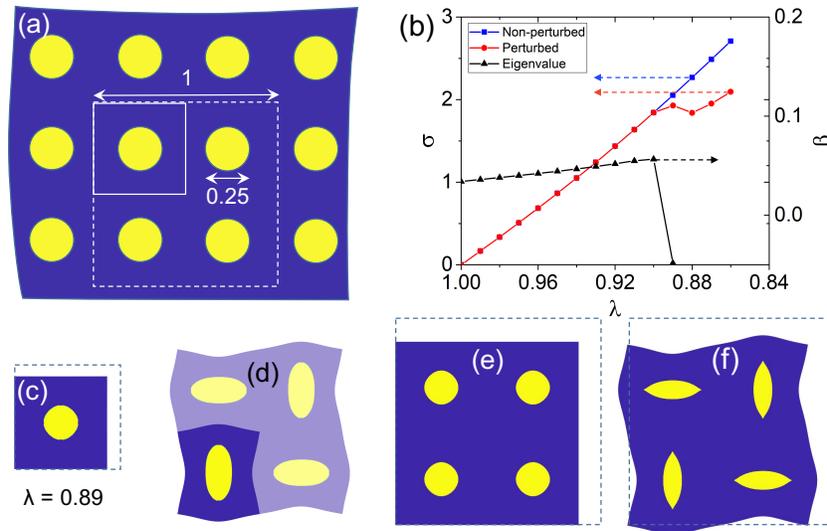
Figure 2.2: Bifurcation of a periodic composite. (a) A periodic composite with the unit cell (bold lines) and super cell (dashed line). (b) The stress-stretch curve without bifurcation (blue, square symbols) and with bifurcation (red, round symbols) along with the modulus of stability $\beta_{(2,2)}$ (black, triangle symbols). (c) Deformed shape of the unit cell (dashed line showing the undeformed size) at $\lambda = 0.89$). (d) The mode shape of the unstable mode in the period-doubling instability at $\lambda = 0.89$. (e) The deformed shape of the supercell without bifurcation at $\lambda = 0.89$. (f) The deformed shape of the supercell post bifurcation at $\lambda = 0.89$.

- Library functions that are optimized for GPUs are available for fast Fourier transform and basic linear algebra operations.

2.4 Numerical performance

We discuss the computational performance in detail using a bifurcation problem in finite elasticity. We first introduce the problem in Sec. 2.4. The problem has been previously studied using both computation in two dimensions (Triantafyllidis, Nestorović, and Schraad, 2006; Bertoldi et al., 2008) and experiment (Mullin et al., 2007; Bertoldi et al., 2008). It serves as verification of the proposed method against previous results of Triantafyllidis, Nestorović, and Schraad (2006). In Sec. 2.4, we use this example to discuss convergence and scaling. We also discuss several acceleration techniques.

Periodic medium in finite elasticity

Consider a periodic arrangement of compliant circular inclusions in a stiff matrix in two dimensions as shown in Figure 2.2(a). Both materials are modeled as

compressible Mooney-Rivlin materials with stored energy density

$$W(F, x) = \frac{\mu(x)}{2}(I_1 - \ln I_2 - 2) + \frac{\kappa(x)}{2}(\sqrt{I_2} - 1)^2 \tag{2.34}$$

where μ is the shear modulus and κ is the bulk modulus, and I_1 and I_2 are the first and second invariants of right Cauchy-Green tensor C. The moduli take the values μ_i, κ_i and μ_m, κ_m in the inclusion and the matrix, respectively, with $\mu_i < \mu_m, \kappa_i < \kappa_m$. In our numerical examples, $\kappa_i/\mu_i = \kappa_i/\mu_m = 9.8$ while $\mu_m/\mu_i = 20$.

We start with a 1×1 unit cell simulation. The periodic medium is subjected to equi-biaxial compression, $\langle F \rangle = \lambda I$ where λ decreases monotonically from an initial value of 1. At each given value of λ, the equilibrium solution is computed using Algorithm 1 on a 1024×1024 grid starting with the solution of the previous λ as an initial guess. We obtain the stress-stretch curve shown by square symbols (blue) in Figure 2.2(b) and a periodic solution with the deformed unit cell shown in Figure 2.2(c).

It is known that this example develops a long-wavelength instability. Note that a 1-periodic function is also $k-$periodic for any $k = (k_1, k_2), k_i$ integers. Thus, we may have equilibrium solutions that are periodic on a $k_1 \times k_2$ super-cell. However, it is known from Geymonat, Müller, and Triantafyllidis (1993) that the 1-periodic solution is the stable solution near $\lambda = 1$. However, this solution may become unstable as λ changes. By the second variation condition, the 1-periodic solution u remains stable as long as

$$\int_{\Omega_k} \nabla v \cdot \frac{\partial W^2}{\partial F \partial F}(\nabla u(x), x) \nabla v dx \geq 0 \tag{2.35}$$

for all non-zero $k-$periodic functions v. Using Bloch waves, this is equivalent to requiring $\beta_k \geq 0$ where β_k is as defined in (2.30).

Therefore, we compute the modulus of stability β_k for various k, and this is also shown with triangular symbols (black) in Figure 2.2(b). We see that $\beta_{(2,2)} \to 0$ as $\lambda \to 0.9$. The corresponding mode v_k is shown in Figure 2.2(d) as the darkened region. This suggests that the periodic solution will bifurcate to a solution that is periodic on a 2×2 super-cell.

We therefore repeat the finite deformation equilibrium computation on a 2×2 super-cell and a 2048×2048 mesh, once without a perturbation (i.e., with the solution to the previous λ as an initial guess), and once with the linearly unstable mode added as a perturbation (i.e., with the sum of the solution to the previous λ and a scaled

(a)

(b)
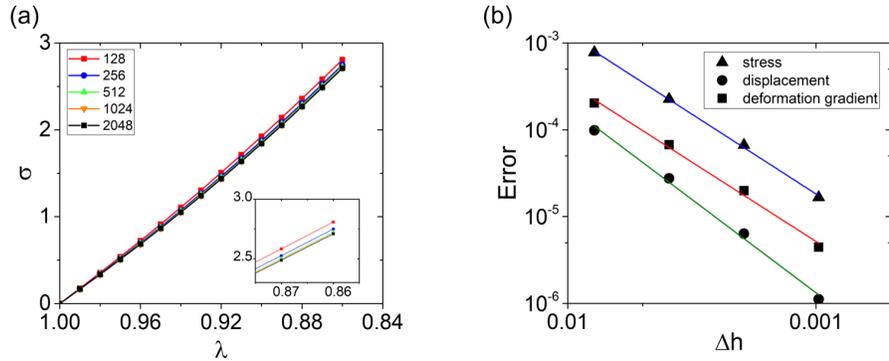
Figure 2.3: Convergence with mesh size. (a) Stress-stretch curve for various computational grid resolution and (b) Relative error in the deformation gradient and stress vs. grid size.

(a)

(b)
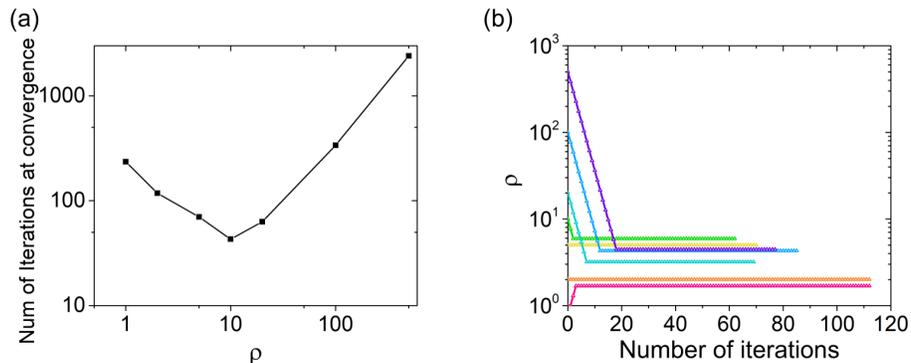
Figure 2.4: Effect of penalization parameter ρ on convergence. (a) Total number of iterations with different fixed ρ and (b) Variation of ρ starting from different values.

eigenmode v_k associated with β_k as the initial guess). The simulation without a perturbation leads to a periodic solution as before Figure 2.2(e) and with the same stress-stretch curve shown in black in Figure 2.2(b). The perturbed solution also agrees with it until $\lambda \approx 0.9$, but then bifurcates into a solution with period 2×2 shown in Figure 2.2(f) with a stress-stretch curve shown with round symbols (red).

All results agree with those of Triantafyllidis, Nestorović, and Schraad (2006), thereby verifying the method.

Convergence and performance

We now use this example to demonstrate convergence and scaling of the proposed algorithm. In all the tests, we compress the composite until $\lambda = 0.95$.

We begin by investigating the convergence with resolution. The simulation is performed with 128×128, 256×256, 512×512, 1024×1024, and 2048×2048

grids, and the stress-stretch curves are shown in Figure 2.3(a). Further, taking the 2048×2048 grid as the reference, Figure 2.3(b) shows how the relative error (L_2 norm) of deformation gradient and stress depend on resolution. We observe polynomial convergence with rates of 1.83 and 1.84 for the error in deformation gradient and displacement, respectively. These are very close to the expected rate of 2 for the discrete differential operator of FFT. We believe that these are due to the change in residual spurious oscillation at the interface as well as the change in pixellated geometric representation with resolution.

The effect of the penalization parameter ρ on the number of iterations required for a given error is shown in Figure 2.4. Figure 2.4(a) shows the number of iterations to achieve a given convergence (in primal and dual error) when ρ is held fixed at different values. We observe that the number of iterations increase for both small ρ and large ρ with the optimal at about $\rho = 10$. As noted earlier in Section 2.2, the primal error is large for small ρ and the dual error is large for large ρ. This is the reason that we adjust ρ following (2.23). We show the evolution of ρ for various initial values of ρ in Figure 2.4(b) with γ and τ in (2.23) set to 5 and 1.3, respectively. We note that in all cases, ρ converges exponentially to values from 3 to 10. Further, in contrast to the case with fixed ρ, the simulation converges well before 100 iterations. Thus, (2.23) ensures a robust convergence of ρ and significantly speeds up the algorithm. We have also observed in our numerical experiments that $\rho < 1$ leads to divergence early in the iteration.

We study the approximate solution of the local step 1 in Figure 2.5 using two strategies. The first strategy is to require convergence of the local iterations of Step 1 only on a fixed fraction of spatial points, and these results are shown in Figure 2.5(a,b). We check how many spatial points have converged to a given (pointwise) residual after a given number of local iterations and proceed to step 2 if a given percentage of spatial points have converged. This check requires a communication between the shared and device memory which adds time, but it can be expedited using coalesced memory (the time required for the memory transfer and check is comparable to the time required for a single local iteration in our examples). Figure 2.5(a,b,c) shows the results when the local residual is held to 10^{-11}, a check is performed every two local iterations and the percentage of converged points varied from 10% to 90%. We see from Figure 2.5(a) that the wall clock time for the global iteration to converge decreases monotonically with the percentage of converged points. Further, the number of global iterations necessary for global
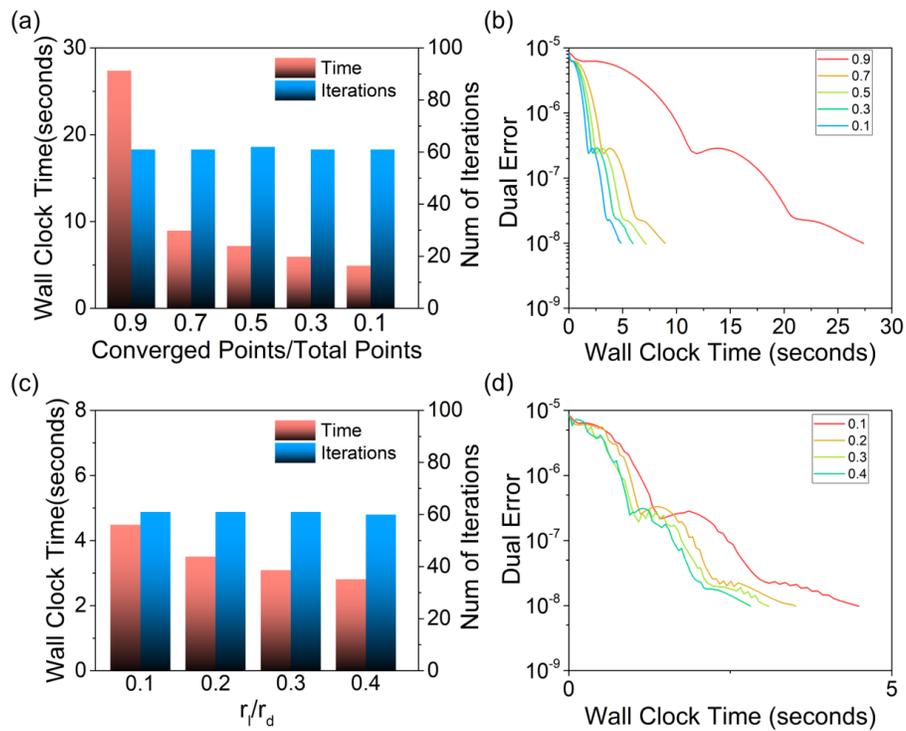
Figure 2.5: Performance with approximate solution of the local problem. (a,b) Local convergence on a fixed fraction of spatial points: (a) Wall clock time and number of global iterations for global convergence for various fractions. (b) The global dual error versus wall clock time for various fractions. (c,d) Fixed ratio of local (r_l) to global dual (r_d) residual: (c) Wall clock time and number of global iterations for global convergence for various ratios and (d) The global dual error versus wall clock time for various ratios.

convergence is largely independent of the number of converged points. In other words, the approximate solution of the local step 1 has relatively little adverse effect on the global convergence. Indeed, we see from Figure 2.5(b) that the global dual residual decreases in the same manner as the calculations proceed, except each global iteration is faster thereby reducing the overall clock time.

The second strategy is to require that the local residual r_l be a fixed fraction of the global dual residual r_d, and the results are shown in Figure 2.5(c,d). We observe from Figure 2.5(c) that the wall clock time for global convergence decreases while the total number of global iterations necessary for global convergence remains unchanged as we increase the ratio r_l/r_d. In fact, at around a ratio of 0.3, we only require one or two iterations in the local step after a few global steps. Again, we see from Figure 2.5(d) that the global dual residual decreases in the same manner independent of the ratio r_l/r_d as the calculations proceed, except each global iteration is faster with
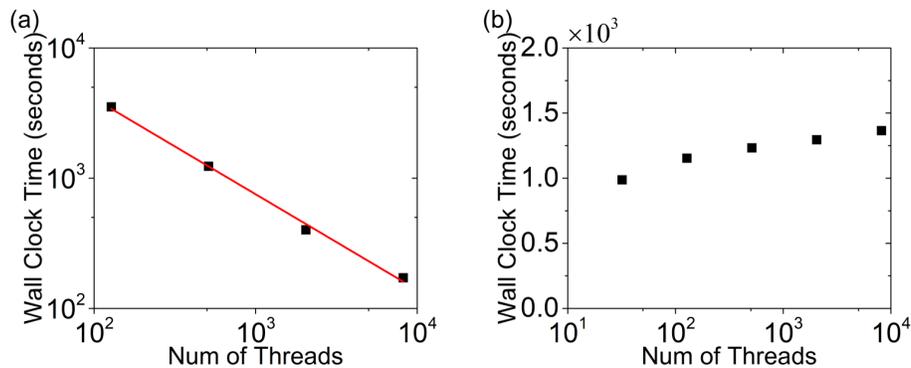
Figure 2.6: Parallel performance of the algorithm. (a) Strong scaling and (b) Weak scaling.

increasing r_l/r_d thereby reducing the overall clock time.

Finally, we have verified that the error in the final solution obtained by either of these approximate approaches to local step 1 is negligible when compared to the solution obtained by the exact solution of step 1 (difference in the deformation gradient is comparable to machine error).

These results show that approximate solution of the local step 1 is an effective strategy to improving performance of our method. In effect, we allow for larger tolerance in the local step when the global residual is large and exploit these in future iterations. Further, recall from Section 2.2 that the error in satisfying the (physical) equilibrium equation is bounded by the local and residual global error. In particular, the vanishing dual error implies vanishing local error as required by the Eckstein-Bertsekas condition (Eckstein and Bertsekas, 1992). Therefore, requiring the local residual to be a fraction of the global residual ensures physically meaningful solutions. Therefore, we adopt this strategy. Finally, we remark that this strategy is especially useful in highly nonlinear problems. The analogous results for our example in liquid crystal elastomers are shown in Figure A.1 of the Appendix.

Finally, the parallel performance of the algorithm is examined in Figure 2.6. The problem with a 1024×1024 mesh is simulated with 128, 512, 2048, 8192 threads. We observe a steady decrease of wall time with increased threads of GPU. The slope is -0.73, suggesting a good scalability of the algorithm. We attribute the deviation from the perfect slope of -1 to two reasons. First, we use FFT which scales as $O(n \log n)$ with system size, and second, operations such as sum and norm also takes communication and do not scale linearly with the number of nodes. We note that the scaling improves for nonlinear problems as the local step 1 takes a larger

fraction of the time. The analogous result for liquid crystal elastomers is shown in Figure A.2 of the Appendix and the slope is -0.80. The scaling efficiency is confirmed by weak scaling. The same configuration is studied with a 128×128 grid using 32 threads, a 256×256 grid using 128 threads, a 512×512 grid using 512 threads, a 1024×1024 grid using 2048 threads, and a 2048×2048 grid using 8192 threads. Overall, the algorithm and GPU implementation show good parallel efficiency with system size.

2.5 Comparison to other FFT-based methods

We conclude this chapter by comparing the performance proposed method to that of other FFT-based methods, the original Lippman-Schwinger-based method and the emerging Fourier-Galerkin method. We first present these methods in the context of finite elasticity on a periodic domain Ω:

$$\nabla \cdot (W_F(\nabla u, x)) = 0 \tag{2.36}$$

where $u : \Omega \to \mathbb{R}^3$ is the deformation, $F : \Omega \to \mathbb{R}^{3\times3}$ is the deformation gradient, $W(F, x) : \mathbb{R}^{3\times3} \times \mathbb{R}^d \times \Omega \to \mathbb{R}$ is the elastic energy density.

Fourier-Galerkin

From the virtual work principle, we reformulate (2.36) as,

$$\int_\Omega \delta F : P dx = 0 \tag{2.37}$$

where $P = \frac{\partial W}{\partial F}(\nabla u, \eta, x)$ is the stress, and δF is all the deformation gradient induced by displacement perturbation. Note that δF is curl-free. The Fourier-Galerkin method ensures the compatibility constraint using a projection operator G that projects an arbitrary second order tensor field F into its curl-free component. Thus for arbitrary test function δF (not necessarily curl-free), the weak formulation reads,

$$\int_\Omega (G * \delta F) : P d\Omega = 0, \quad \forall \, \delta F. \tag{2.38}$$

Notice that G is self-conjugate (Vondřejc, Zeman, and Marek, 2014), so we obtain

$$\int_\Omega \delta F : (G * P) d\Omega = 0, \quad \forall \, \delta F. \tag{2.39}$$

Given the arbitrariness of the test function δF, it leads to the final equation,

$$G * P = 0 \tag{2.40}$$

Algorithm 2: Fourier-Galerkin method

Given an initial value of u^0, F^0, η^0 and macroscopic strain path $\overline{F}(t)$;
Step 0: Initialize:
  Initialize $r_0, \mu, \epsilon, F^0, A^0$
while $t < t_{max}$ **do**
  | Initialize $\Delta F_0, P = P^n, A = A^n, r_0 = -G * P - G * (A : \Delta F_0), q_0 = r_0, i = 0$
  | **while** $||r_i||_{L^2} > \epsilon$ **do**
  | | $\alpha_i := \frac{r_i^T r_i}{q_i^T G * (A : q_i)}$
  | | $\Delta F_{i+1} := \Delta F_i + \alpha_i q_i$
  | | $r_{i+1} := r_i + \alpha_i G * (A : q_i)$
  | | $\beta_i := \frac{r_{i+1}^T r_{i+1}}{r_i^T r_i}$
  | | $q_{i+1} := r_{i+1} + \beta_i q_i$
  | | $i := i + 1$
  | **end**
  | Update $t, F^{n+1} = F^n + \Delta F_i, P^{n+1}, A^{n+1}$
end

This nonlinear equation can be solved by Newton iteration,

$$G * (P^i + A^i : \Delta F^{i+1}) = 0 \qquad \Longleftrightarrow \qquad G * (A^i : \Delta F^{i+1}) = -G * P^i \qquad (2.41)$$

where $A^i = \partial^2 W / \partial F^2 \big|_{F^i}$ is the incremental modulus. Since G is a projection operator, it is linear. We solve this linear system using the conjugate gradient method. The algorithm is summarized in Alg. 2

Lippmann-Schwinger

The solution of (2.36) using Newton iteration leads to the linear problem

$$\nabla \cdot \left(A^i \nabla \Delta u^i + P^i \right) = 0. \qquad (2.42)$$

Introduce a fixed modulus A^0 and rewrite the above as

$$\nabla \cdot A^0 \nabla \Delta u^i + \nabla \cdot (A^i - A^0) \nabla \Delta u^i = -\nabla \cdot P^i. \qquad (2.43)$$

Let Γ^0 be the operator

$$\Gamma^0 = \nabla (\nabla \cdot A^0 \nabla)^{-1} \nabla \cdot . \qquad (2.44)$$

Applying the operatore $\nabla (\nabla \cdot A^0 \nabla)^{-1}$ to (2.43), we can rewrite it as

$$\left(I + \Gamma^0 (A^i - A^0) \right) \Delta F^i = -\Gamma^0 P^i \qquad (2.45)$$

where $\Delta F^i = \nabla \Delta u^i$. Setting $A^0 = I$ reduces this to

$$(I + G * (A^i - I)) \Delta F^{i+1} = -G * P^i \qquad (2.46)$$

Algorithm 3: Lippmann-Schwinger method

Given an initial value of u^0, F^0, η^0 and macroscopic strain path $\overline{F}(t)$;
Step 0: Initialize:
   Initialize $r_0, \mu, \epsilon, F^0, A^0$
while $t < t_{max}$ **do**
    Initialize ΔF_0 randomly, $P = P^n, A = A^n, i = 0$
    while $||r_i||_{L^2} > \epsilon$ **do**
       $\Delta F_{i+1} := -(G * (A - I) : \Delta F_i) - G * P$
       $r_i := F_{i+1} - F_i$
       $i := i + 1$
    end
    Update $t, F^{n+1} = F^n + \Delta F_i, P^{n+1}, A^{n+1}$
end

---

where G is as defined before. We solve this by fixed point iteration following Alg. 3.

Implementation and verification

These two methods are implemented on GPUs. Note that the calculation of P at each step is local, and therefore can be done in parallel. As before, the FFT is implemented using FFT the CUDA cuFFT library while all basic linear algebraic operations are implemented using the cuBLAS library. Again as before, we work on a compute node consisting of a 14-core Intel Broadwell CPU and four Nvidia Tesla P100 GPUs. It features 16GB of (global or device) memory and has 56 SMs with 64 cores each for a total of 3584 cores, and has a double precision performance of 4.7 teraFLOPS.

We study the same finite elastic composite problem as in Section 2.4. The stress-stretch curve obtained from the various methods are in Figure 2.7. The agreement verifies our implementation.

Step-size

Note in Figure 2.7 that the step size used for our method is large compared to those used for Fourier-Galerkin and Lippmann-Schwinger methods. This is because the latter two methods would not converge for large step-size. Further, we encounter convergences issues at large strains even for very small step-size. We believe that this is for two reasons. First, our approach uses gradient descent with back-tracking for the nonlinear problem. In contrast, the Fourier-Galerkin and Lippmann-Schwinger use Newton's method. While the latter provides faster convergence, it has a smaller
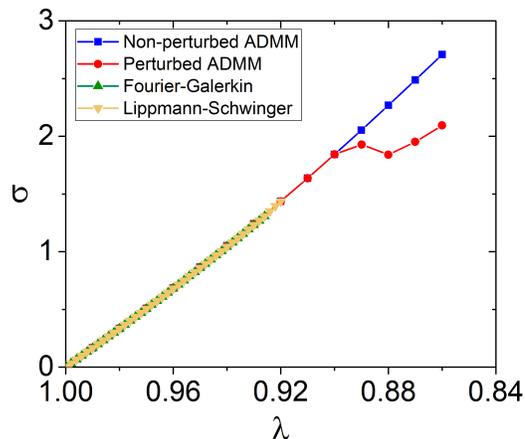
Figure 2.7: Stress-stretch curve of the elastic two phase composite computed using our approach (ADMM) without perturbation (blue squares); our approach (ADMM) with perturbation induced bifurcation (red circle), Fourier-Galerkin method (green triangle), and Lippmann-Schwinger method (yellow triangle)

radius of convergence. Second, the nonlinear optimization problem is local in our method. In contrast, in our methods, the energy minimization step is global, and this potentially has more local minima.

Scaling

The parallel performance of three algorithms are examined in Figure 2.8. We consider five increments (timestep) of the finite elasticity problem. The problem with a 1024×1024 mesh is computed with 32 (omitted for our method), 128, 512, 2048, 8192 threads. We observe a steady decrease of wall time with increased threads of GPU. The slope of our method, our method with approximate local solution, Fourier-Galerkin method and Lippmann-Schwinger are -0.73, -0.85, -0.79, and -0.80, respectively, suggesting a good scalability of all algorithms. We attribute the deviation from the perfect slope of -1 to two reasons. First, we use FFT which scales as $O(n \log n)$ with system size; and second, operations such as sum and norm also takes communication and do not scale linearly with the number of nodes. The scaling efficiency is confirmed by weak scaling. The same configuration is studied with a 128×128 grid using 32 threads, a 256×256 grid using 128 threads, a 512×512 grid using 512 threads, a 1024×1024 grid using 2048 threads, and a 2048×2048 grid using 8192 threads. We note that the last data points of weak scaling of our method with approximate local and Fourier-Galerkin has a sudden increase. This is because the first four cases uses only one GPU while the last case
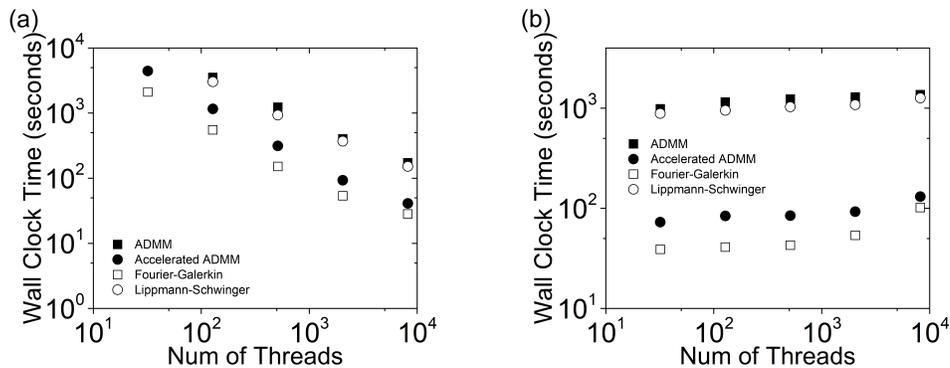
Figure 2.8: Scaling of various methods: our method (ADMM), our method with approximate local solution (accelerated ADMM), Fourier-Galerkin and Lippmann-Schwinger:(a) strong scaling; (b) weak scaling

uses four GPUs.

We also note the difference in efficiency. Note that the Fourier-Galerking method is the fastest, almost twice as fast as the next method, our method with approximate local solution. The other two methods are about ten-times slower.

Convergence rate of different methods

The convergence rates of the various methods for a single increment of the finite elasticity problem is shown in Figure 2.9. We note that Fourier-Galerkin method has a faster convergence rate compared with other methods. Fourier-Galerkin uses onjugate gradient method which has a convergence rate of $\frac{\sqrt{\rho}-1}{\sqrt{\rho}+1}$, where ρ is the radius of the linear operator. In contrast, our method and the Lippmann-Schwinger have a convergence rate of $\frac{\rho-1}{\rho+1}$. We also note that our approach with the approximate local solution is able to achieve a comparable convergence rate with Fourier-Galerkin.

ADMM and accelerated ADMM method. Since Galerkin method adopts conjugate gradient method while ADMM here adopts gradient-based method. Conjugate gradient method is converging at rate $\frac{\sqrt{\rho}-1}{\sqrt{\rho}+1}$, where ρ is the radius of the linear operator. ADMM and Lippmann-Schwinger method are converging with rate $\frac{\rho-1}{\rho+1}$. However, using ADMM with partial solve, we could achieve comparable convergence rate compared with Fourier-Galerkin method.

Summary

The Fourier-Galerkin method offers the fastest rate of convergence and fastest numerical performance on a single time step compared to the Lippmann-Schwinger
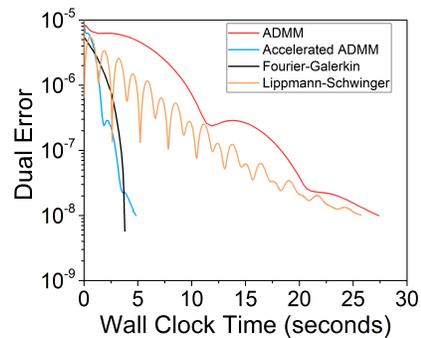
Figure 2.9: Convergence rates of various methods: our method (ADMM), our method with approximate local solution (accelerated ADMM), Fourier-Galerkin and Lippmann-Schwinger

and our methods. Our approach with the approximate local solution provides comparable performance to the Fourier-Galerkin method. However, we are able to take large step size in our approach (including with approximate local solution) compared to both the Fourier-Galerkin and Lippmann-Schwinger approaches. Further the latter approaches have convergence issues near bifurcations. Therefore, we conclude that our approach with the approximate local solution provides the best performance for nonlinear problems.

Chapter 3

ACCELERATED COMPUTATIONAL MICROMECHANICS: APPLICATION TO LIQUID CRYSTAL ELASTOMERS

3.1 Introduction

Liquid crystal elastomers (LCEs) are synthetic materials made by incorporating rod-like nematic mesogens into the polymer chains of an elastomer (Warner and Terentjev, 2003). The combination of polymer elasticity and nematic ordering gives rise to exotic mechanical behavior. At high temperatures, the mesogens are randomly oriented and the LCE is an isotropic rubbery solid. However, on cooling, they undergo a phase transition where the steric interactions between the mesogens cause them to align in a particular direction. The isotropic to nematic phase transition is accompanied by a spontaneous deformation. Thus, they have been proposed for applications as actuators and for shape-morphing. Further, below the transition temperature, the nematic director can reorient resulting in a soft behavior (Kundler and Finkelmann, 1995). We refer the reader to Warner and Terentjev (2003) for a comprehensive introduction.

Our focus is *isotropic-genesis* LCEs, those that are synthesized in the isotropic phase and then cooled. There is typically some small non-uniformity of cross-linking during synthesis, and this leads to a local preference for the nematic orientation that is not uniform ('random'). This results in a *polydomain* state where the nematic director is highly non-uniform. The application of stress can reorient the director, but there is an energetic cost (Biggins, Terentjev, and Warner, 2008). This leads to semi-soft behavior in such isotropic-genesis polydomain materials (Urayama et al., 2009), and this is attractive for a number of damping applications. Indeed, the viscoelastic response of such materials has also been a subject of recent interest (Azoug et al., 2016). Finally, recent experiments have shown rather unusual behavior in biaxial loading where unequal stretch leads to equal true stress (Tokumoto et al., 2021). The goal of this chapter is to study the underlying mechanism of microstructure evolution in isotropic-genesis polydomain materials subjected to multiaxial stress.

3.2 In-plane liquid-like behavior of isotropic-genesis polydomain LCEs

We summarize recent experimental observations of Kenji Urayama and his students (reported in a joint publication, Tokumoto et al. (2021)) on isotropic-genesis main-
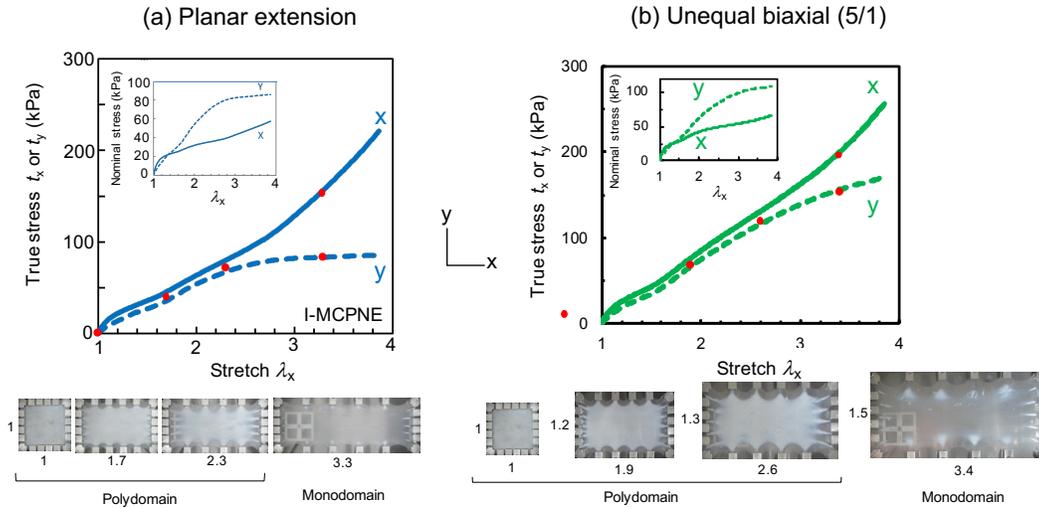
Figure 3.1: In-plane liquid like behavior of polydomain isotropic genesis LCE. Stress vs. stretch for (a) Planar extension where $\lambda_x = \lambda, \lambda_y = 1$ and (b) Unequal biaxial where $(\lambda_x - 1)/(\lambda_y - 1) = 5/1$.

chain polydomain LCE synthesized using a thiol-acrylate Michael addition reaction between di-acrylate mesogen, di-thiol chain extender, and tetra-thiol cross-linker.

Figure 3.1(a) shows the stress vs. stretch behavior of the LCE subjected to planar extension (where one edge is held fixed while the other is extended: $\lambda_x = \lambda, \lambda_y = 1$. Quite unexpectedly, the nominal stress in the stretching direction $(x-)$ is in fact smaller than that in the fixed direction $(y-)$ (inset). Further, the true stresses in both directions are equal till a stretch of about 2.3 despite the fact that the two directions are extended unequally. The specimen is initially cloudy indicating a polydomain state and becomes clear indicating a monodomain state at around a stretch of 2.3. Figure 3.1(b) shows the analogous behavior an unequal biaxial stretch where the nominal strain ratio is 5:1 $(\lambda_x - 1)/(\lambda_y - 1) = 5/1)$. The nominal stress in the stretching direction is again smaller, while the true stresses are equal till a stretch of about 2.6.

Various biaxial loading protocols where tested, and the equality of the true stress $(t_x = t_y)$ was always found to always hold. Further, it was found that the value of the (equal) true stress is independent of the individual stretches and depends only on the areal stretch or the product of the two individual stretches. These results are collected in Figure 3.2(a). Note that we use the incompressibility of LCEs $(\lambda_x \lambda_y \lambda_z = 1)$ to plot the result as a function of λ_z. In short, the in-plane true stresses are always equal, independent of the individual stresses, and depend only
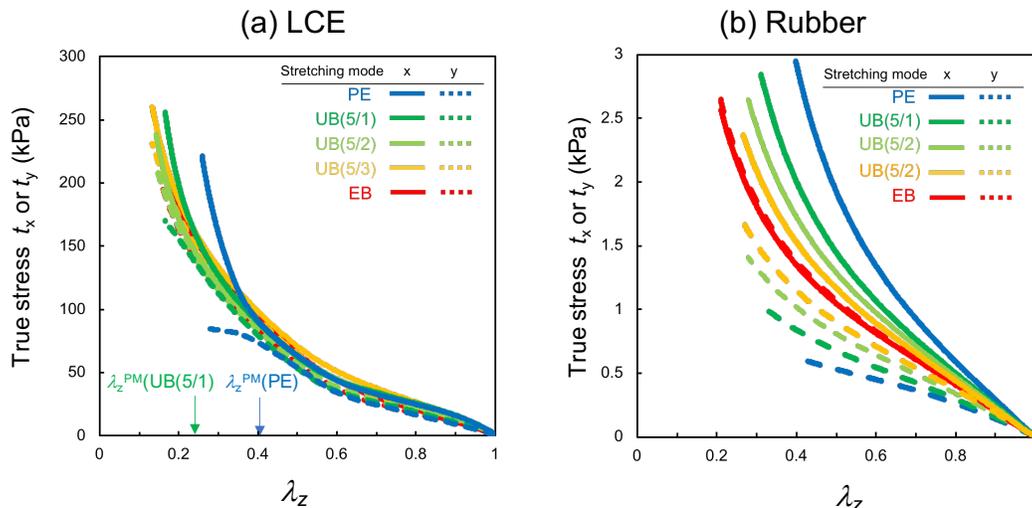
Figure 3.2: In-plane liquid like behavior of polydomain isotropic genesis LCE. True stress vs. thickness stretch for various loading protocols. (a) LCE: the in-plane true stresses are always equal, independent of the individual stresses, and depend only on the areal stretch. (B) Rubber: the in-plane true stresses are not necessarily equal and depend on the individual stretches.

on the areal stretch. This is the planar analog of a liquid where the true stresses are equal and depend only on the volume change. Therefore, we call this the *in-plane liquid-like behavior*. Figure 3.2(b) is a control that repeats the experiment in styrene butadiene rubber and shows that this remarkable property is not present in ordinary rubber.

This in-plane liquid-like behavior was anticipated using an ideal theory of LCE by Cesana, Plucinsky, and Bhattacharya (2015). This theory also predicts perfectly soft behavior in uniaxial extension, but such perfectly soft behavior is not observed experimentally due to non-ideal corrections necessary to model real materials. Remarkably, the in-plane liquid-like behavior is experimentally observed and we seek to understand it within a theory that includes the non-ideal corrections.

3.3 Theoretical framework for isotropic-genesis LCEs

We consider a continuum model that treats the director as an internal variable, has a free energy density that includes contributions from the entropy of the polymer chains (Bladon, Terentjev, and Warner, 1993), the energetic cost of reorientation (Biggins, Terentjev, and Warner, 2008), nematic or Frank elasticity that penalizes gradients in the director (Warner and Terentjev, 2003), and dissipation that accounts

for director reorientation and viscosity. The 'ideal' model, where the typically small reorientation and Frank contributions are neglected, has been widely studied in the static setting. This theory is not convex and leads to stripe domains and other fine-scale microstructure, resulting in soft behavior (Warner and Terentjev, 2003). Indeed, DeSimone and Dolzmann (2002) explicitly calculated the relaxed energy or the effective energy that implicitly accounts for the microstructure in this ideal situation, and this energy has been the basis of macroscopic finite element simulations (Conti, DeSimone, and Dolzmann, 2002b; Plucinsky and Bhattacharya, 2017). There are also some detailed study of microstructure evolution (Mbanga et al., 2010; Y. Zhang et al., 2019), but they focus on monodomain specimens. In polydomain specimens, Biggins, Warner, and Bhattacharya (2009) and Biggins, Warner, and Bhattacharya (2012) used bounds to establish the semi-soft behavior in uniaxial tension while Uchida (1999) and Uchida (2000) studied semi-soft behavior with a related model based on random networks. The focus of this work is microstructure evolution in polydomain materials, and its macroscopic consequence under multiaxial loading.

We take the configuration in the isotropic state to be the reference configuration, but consider the material at a temperature below the phase transition temperature so that it is in the nematic phase. The state of the material at a material point x is then described by the deformation gradient $F(x)$, and a nematic director $n(x)$ that describes the orientation of the nematic mesogens in an infinitesimal volume around x. The material is typically incompressible and so $\det F = 1$, and n is a unit vector ($|n| = 1$) since it describes an orientation.

The free energy per unit volume of the material is given by,

$$W(F, n, \nabla n, x) = W_{el}(F, n) + W_{ni}(F, n, x) + W_F(\nabla n) \qquad (3.1)$$

where the three terms describe three different physics. The first term,

$$W_{el}(F, n) = \frac{1}{2}\mu Tr(F^T \ell^{-1} F) \quad \text{where} \quad \ell = r^{-1/3}(\delta - (r - 1)n \otimes n) \qquad (3.2)$$

describes the entropic elasticity of the polymer network (Bladon, Terentjev, and Warner, 1993). μ is the shear modulus, $\delta = Id$ is identity tensor, $r > 1$ represents the degree of nematic order that depends on temperature, and $\otimes$ represents the tensor product. We take r to be fixed since we fix temperature. Note that if $r = 1$, then $\ell = \delta$ and W_{el} reduces to the neo-Hookean energy (Atkin and Fox, 1980). For $r > 1$, the set of ground states ($W_{el} = 0$) corresponds to $F = R\ell_0^{1/2}Q, n = Re$ for

rotations R, Q and fixed unit vector e. In other words, the material elongates along the director by factor $r^{1/3}$ and contracts perpendicular to it by factor $r^{-1/6}$, and the director is free to take any orientation. The second term,

$$W_{ni}(F, n, x) = \frac{1}{2}\mu\alpha Tr((\delta - n_0(x) \otimes n_0(x))F^T(n \otimes n)F) \tag{3.3}$$

describes the 'non-ideality' originating from non-uniformity in the cross-link density (Biggins, Terentjev, and Warner, 2008). α is the strength of the non-ideality and $n_0(x)$ is a fixed random unit vector field. The non-uniformity in the cross-link density seeks to orient the director n parallel to Fn_0 at x, but this is a weak preference since α is typically small. Finally, the third term

$$W_F(\nabla n) = K|\nabla n|^2 = \frac{1}{2}K|\nabla(n \otimes n)|^2 \tag{3.4}$$

is (a constant coefficient approximation of) Frank elasticity (Warner and Terentjev, 2003; Wojtowicz, Sheng, and Priestley, 1975). It reflects the preference of the directors to align spatially. It is easy to verify that both forms of the expression shown are equivalent using the fact that $|n| = 1$. We note for later use that $\sqrt{K/\mu}$ determines a length-scale of the domain wall and is typically $O(10\text{nm})$ (Warner and Terentjev, 2003).

The evolution is controlled by a dissipation potential which we take to be

$$D(\dot{F}, \dot{n}) = \frac{1}{2}\nu_F|\dot{F}|^2 + \frac{1}{2}\nu_n|\dot{n}|^2 \tag{3.5}$$

Note that this expression is not frame-indifferent and there are a number of generalizations (Mielke, Ortner, and Şengül, 2014). However, this is commonly used when the boundary conditions do not involve large rigid body rotations. The implicit time discretization of the evolution equation gives rise to the following variation problem (cf. 2.11)

$$u^{k+1}, n^{k+1} = \text{argmin}_{|n|=1, \det \nabla u=1} \int_\Omega \left(W(\nabla u, n, \nabla n, x) + \Delta t D\left(\frac{\nabla u - \nabla u^k}{\Delta t}, \frac{n - n^k}{\Delta t}\right)\right) dx. \tag{3.6}$$

Since ∇u and n satisfy constraints, one should consider the non-Euclidian metrics along the constraint manifold instead of the Euclidian metric in the embedding space. However, the approximate expressions are accurate to first order since the manifolds are smooth with curvature bounded from below.

We discretize space using finite differences and solve the resulting equations according to the massively parallel approach described in Sections 2.2 and 2.3 with
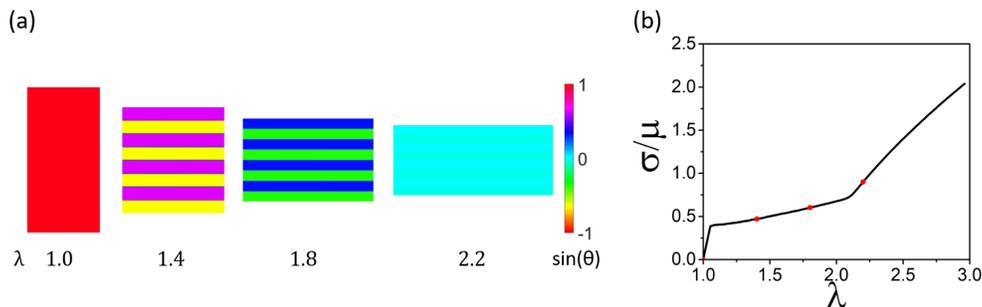
Figure 3.3: Two-dimensional study of a monodomain LCE subjected to uniaxial stress. (a) Evolution of the director (θ is the angle between the director and the horizontal loading direction). (b) Stress-stretch curve (stress is normalized by the modulus μ).

two modifications. First, the constraint of incompressibility, $\det F = 1$, is enforced in the local Step 1 using a Lagrange multiplier; while the constraint on the director, $|n| = 1$ is enforced by introducing Euler angles. Second, we also have a gradient of $(n \otimes n)$ in our functional. We could proceed by introducing an auxiliary variable for $\nabla(n \otimes n)$ and using a constraint for it. However, we have found that we obtain satisfactory results by treating this term explicitly. All simulations are performed under periodic boundary conditions on the deformation gradient and the director.

3.4 Monodomain LCE

We begin by studying a simple problem in two dimensions inspired by the experiments of Küpfer and Finkelmann (1991). We start with a monodomain specimen where $n_0 = e_2$ is uniformly in the vertical direction as shown in Figure 3.3. We normalize the energy density with μ and take the rest of the parameters to be $r = 4$, $\alpha = 0.1$, $K/\mu = 7.63 \times 10^{-6}$ (in non-dimensional length units). We set dissipation to zero taking $\nu_n = \nu_F = 0$ so that we solve for equilibrium at each time step. Finally, we subject the specimen to an average deformation gradient $\langle F_{11} \rangle = \lambda, \langle F_{12} \rangle = \langle F_{21} \rangle = 0$ and $\langle F_{22} \rangle$ free and solve it with a resolution of 256×256. We use the previous configuration along with a small periodic perturbation of the order 10^{-4} in F as the initial guess at each time step. The resulting domain patterns are shown in Figure 3.3(a) while the stress-stretch curve is shown in Figure 3.3(b). We observe the semi-soft behavior and stripe domains consistent with the experimental observations. The director is initially aligned with n_0 and there is no stress. At small applied stretch, the LCE reacts elastically as the non-ideal term keeps the director n aligned with n_0. At a critical stretch, the director

can align, but doing so leads to a shear inconsistent with the imposed condition on the deformation gradient. Therefore, it forms stripe domains where the director rotates in opposite directions in alternating stripes; the two regions have the same stretch but opposing shear so that they can satisfy the imposed average deformation gradient condition. The spacing depends on the perturbation and κ. The formation of stripe domains is accompanied by a softening in the stress-stretch curve. The director continues to rotate as the stretching continues until it is fully rotated to the horizontal when both domains merge (since the sign of the director has no meaning). The stress-stretch curve then stiffens as the material responds elastically. All of this is consistent with the observations of Küpfer and Finkelmann (1991) and prior theoretical considerations (Warner and Terentjev, 2003).

3.5 Polydomain LCE

In this section, the parameters are $r = 7.71$, $\mu = 23.63$ kPa, $\alpha = 0.06$, $K = 3.61 \times 10^{-9}$N, $\nu_F = 2.65$ kPa.s and $\nu_n = 0.005$ kPa.s unless otherwise specified. These parameters are chosen to match the experimental results of Tokumoto et al. (2021) as will be described later. We conduct our simulations on a $(1 \ \mu\text{m})^3$ cubic unit cell with a 128^3 resolution at a strain rate of $1 \ \text{s}^{-1}$ with time steps of 0.02s.

Polydomain material   We begin by exploring the effect of the random director field n_0 on the initial configuration. Figure 3.4(a) shows three 'random' director fields n_0 with fluctuations on a different length-scale. All three of them are generated by starting with the same Gaussian random field of angles and then filtering to different length scales. With an initial guess of $n = n_0$, $F = I$, we let the system relax under zero average stress and we obtain the director field n shown in Figure 3.4(b). We see that the relaxed n does not follow n_0, and the system is internally stressed as shown in Figure 3.4(c). Interestingly, the length-scale on which n fluctuates is similar in each of the three cases. It is larger than the length-scale of n_0 fluctuation in the first two cases, but similar in the third.

Importantly, n remains largely equi-distributed. We recall the definition of the orientation tensor (Warner and Terentjev, 2003)

$$S = Q\left(\langle n \otimes n - \frac{1}{3}I\rangle\right) \tag{3.7}$$

where $\langle \cdot \rangle$ denotes the average over the computational domain and $Q = (r-1)/(r+2)$ is a material parameter ($r = 7.71$ and so that $Q = 0.69$ in our computations). Note that S is a trace-free matrix whose eigenvalues are bounded by $2Q/3$ and $-Q/3$.
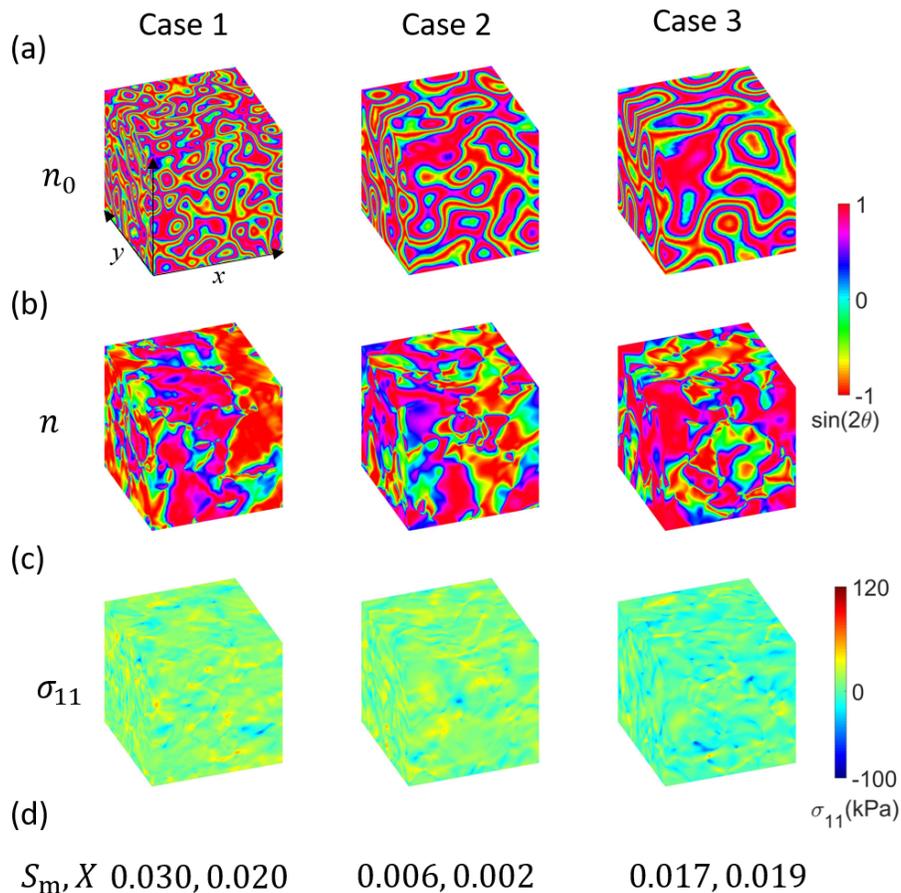
Figure 3.4: Creating the initial polydomain material. (a) Distribution of the preferred director n_0 with fluctuations on three length-scales, (b) Distribution of director n after relaxation, (c) Internal stress distribution and (d) Orientation parameters (below (3.7)) after relaxation. θ is the in-plane angle of mesogens w.r.t. x−axis.

The *mean orientational order parameter* S_m is the largest eigenvalue of S while the *biaxial order parameter* X is the difference between the two smallest eigenvalues of S. It follows that $Sm = 0, X = 0$ for equidistributed domains, $S_m = 2Q/3$ (= 0.46 in our calculations), $X = 0$ for perfectly aligned domains where n is uniform, $S_m = Q/6$ (= 0.115 in our calculations), $X = Q/2$ (= 0.345 in our calculations) for equidistributed domains in the plane. Figure 3.4(d) shows $S_m \approx 0, X \approx 0$ and n is essentially equi-distributed.

We may understand this initial relaxation as follows. Recall that the non-ideal term prefers that the director n follow the prescribed n_0 and the elastic energy prefers an elongation along director. However, this resulting distortion field may not be compatible leading to elastic energy. Further, the Frank energy penalizes the
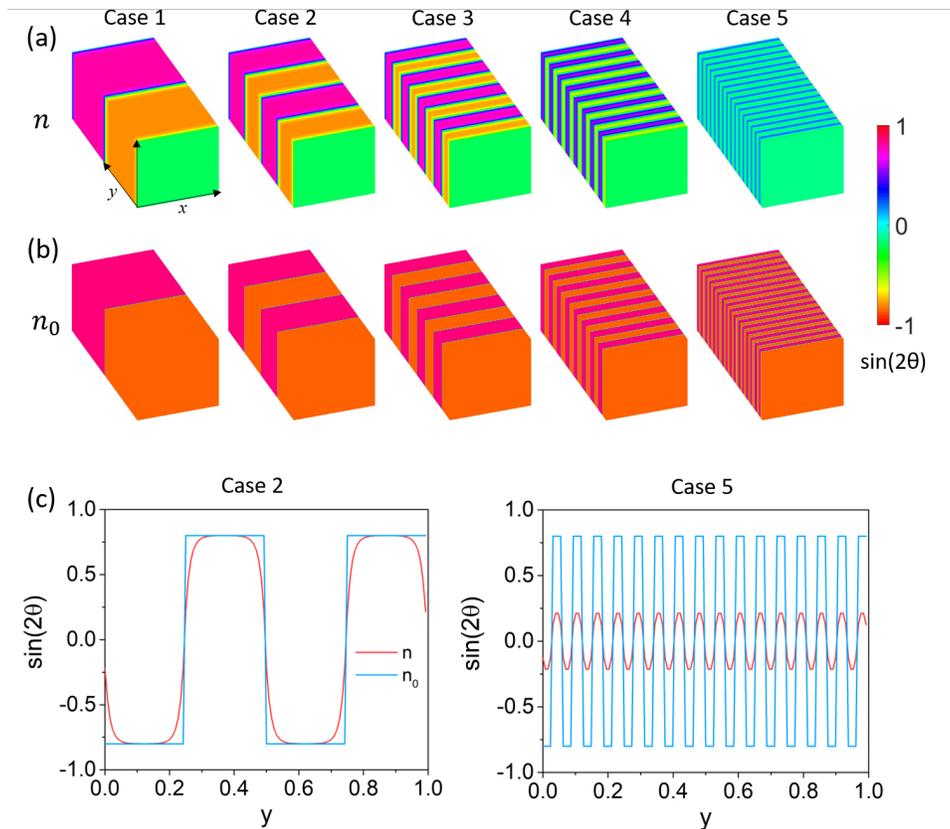
Figure 3.5: Relaxation of a stripe polydomain. (a) Distribution of the compatible preferred director n_0 with multiple length-scales, (b) Distribution of director n after relaxation (θ is the in-plane angle of mesogens w.r.t. x−axis and $\sin 2\theta$ is the product of the x and y components of the director.). (c,d) Details of the director field: the y-component of the director versus the y-coordinate for the second (c) and last cases (d) of (a,b).

fluctuations in the director field. Thus, the competition among these three terms drives the relaxation, and the resulting director pattern is a compromise among them.

To verify this, we start with a 'compatible' initial director field n_0 that takes two distinct values $n_0^{\pm} = \{\pm 0.42, 0.91, 0\}$ in alternating stripes as shown in Figure 3.5(a,c). It is easy to verify that the two corresponding spontaneous stretches $(\ell(n_0^{\pm}))^{1/2}$ are kinematically compatible across an interface with normal e_2, i.e., we can find a rotation Q and a vector a such that $Q(\ell(n_0^{+}))^{1/2} - Q(\ell(n_0^{-}))^{1/2} = a \otimes e_2$. Thus, the director field $n = n_0$ and deformation with gradient $F = (\ell(n))^{1/2}$ is admissible and minimizes the sum of the first two terms in the energy (3.1). Thus,
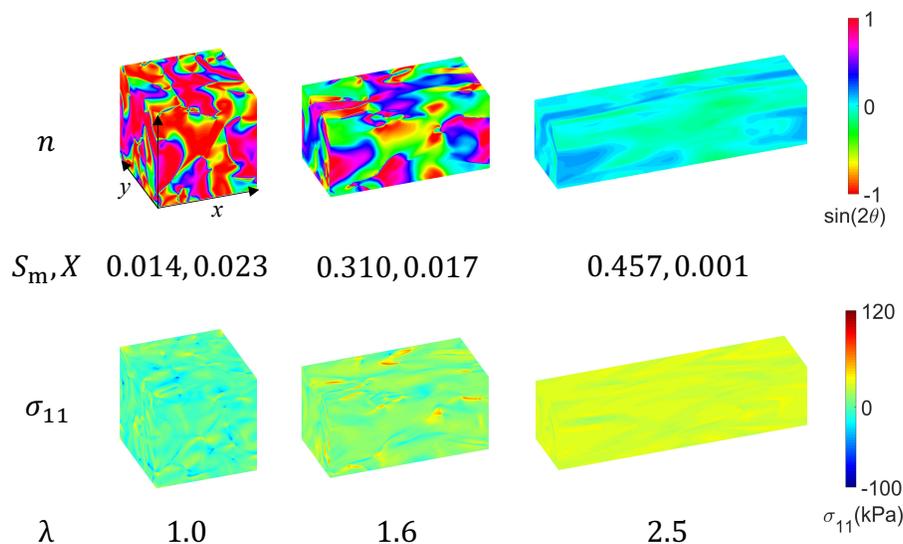
Figure 3.6: Uniaxial stress (UNI). Evolution of the director, the orientation parameters (below (3.7)) and the stress with stretch.

we expect the solution to follow $n = n_0$, $F = (\ell(n))^{1/2}$ except close to the interface where we expect a transition layer with thickness of the order $\sqrt{K/\mu}$. We study five cases with distinct length-scales. As before, we start with $n = n_0$, $F = I$, and let the system relax under zero average stress. We obtain the director field n shown in Figure 3.5(b,c). If the length-scale is is sufficiently large (the first four cases), then n follows n_0 except near the interface where we see a transition layer as we expect. At smaller length-scales the Frank elasticity prevents n from completely relaxing to n_0; in other words the interfaces dominate. This calculation shows that kinematic compatibility drives the relaxation with the Frank elasticity setting the length-scale.

Uniaxial and biaxial deformation   These simulations are motivated by the experiments reported in Tokumoto et al. (2021). They took $65 \times 65 \times 0.7$ mm sheets and subjected them to uniaxial stress and biaxial stretch protocols. To replicate these experiments, we start with a relaxed polydomain specimen prepared as described above and subject it to three loading protocols.

- UNI: uniaxial stress ($\langle F_{11} \rangle$ is prescribed while all other components are free);

- PE: uniaxial stretch of a sheet in plane stress ($\langle F_{11} \rangle = \lambda_x$, $\langle F_{12} \rangle = 0$, $\langle F_{22} \rangle = 1$ while the other components are free);
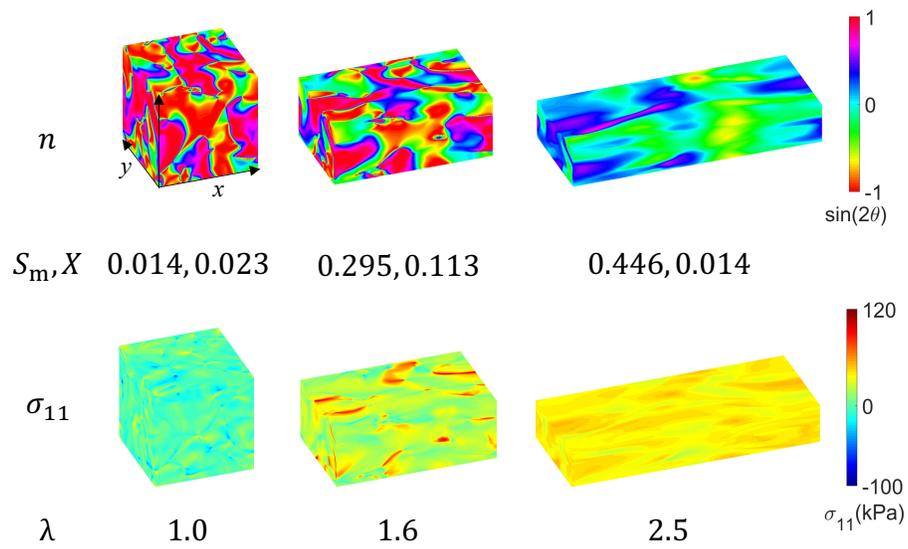
Figure 3.7: Uniaxial stretch of a sheet in plane stress (PE). Evolution of the director, the orientation parameters (below (3.7)) and the stress with stretch.
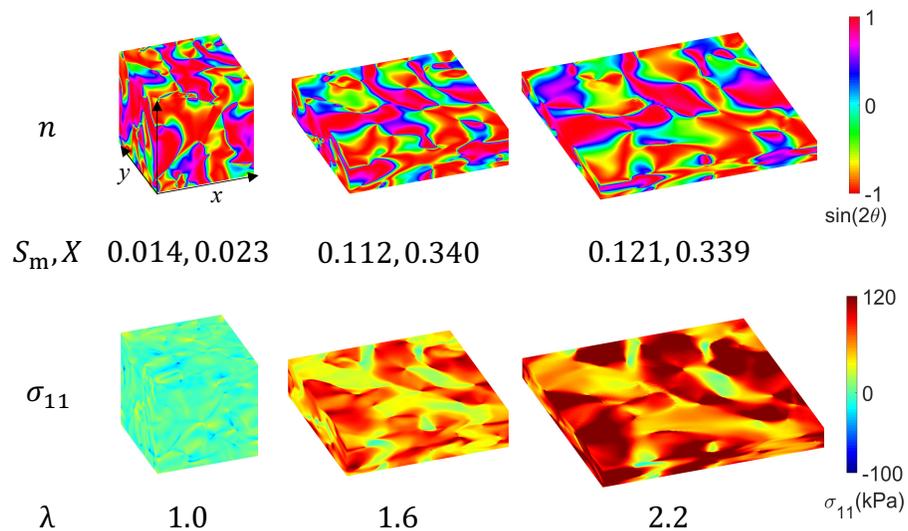
Figure 3.8: Equi-biaxial stretch of a sheet in plane stress (EB). Evolution of the director, the orientation parameters (below (3.7)) and the stress with stretch.
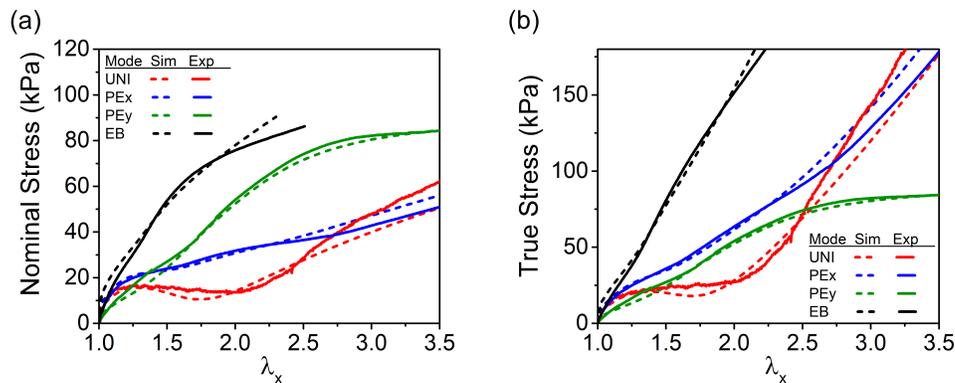
Figure 3.9: Stress vs. stretch for various loading protocols (UNI, PE and EB) obtained by both simulation (dashed line) and experiment (solid line). (a) Nominal stress vs. stretch and (b) True stress vs. stretch.

- EB: equi-biaxial extension of a sheet in plane stress ($\langle F_{11} \rangle = \langle F_{22} \rangle = \lambda_x$, $\langle F_{12} \rangle = 0$ while the other components are free).

Note that PE and EB are mixed boundary conditions on the unit cell. Since $\langle F_{i3} \rangle$ and $\langle F_{3i} \rangle$ are left free and part of the minimization, the average tractions on the faces with normal e_3 are zero, and the average tractions on the other faces are planar. This corresponds to plane stress. At the same time, the average planar stretch is prescribed. The evolution of the director and the stress are shown in Figures 3.6, 3.7 and 3.8. The macroscopic stress-strain curve is shown in Figure 3.9 and compared to experimental observations.

We observe that the director pattern, residual stress and overall stress-strain curve are very different in the different loading scenarios. In uniaxial stress (UNI, Figure 3.6), the directors rotate until they all eventually align. This is similar to the situation in the ideal material (Figure 3.3) though the actual patterns are more complicated. Any residual fluctuations are small and this is also reflected in the small stress heterogeneity. This microstructure evolution leads to a soft plateau in the macroscopic stress-strain response which eventually stiffens when all the directors are aligned.

In the case of uniaxial stretch in plane stress (PE, Figure 3.7), the directors again try to rotate to the direction of elongation, but are prevented from doing so by the lateral constraint. Therefore, significant amounts of residual microstructure and some residual stress persist. Further, the macroscopic stress-strain response shows only a small plateau. The macroscopic stress-strain response also shows another

rather interesting feature. The nominal stress and thus the applied force in the stretching (x) direction is smaller than those in the constrained (y) direction. This is counter-intuitive, and different from the behavior of ordinary elastomers. The reason for it is evident by examining the true or Cauchy stress: we observe the true stresses are (almost) equal in the two directions despite the fact that the stretches in the two directions. In other words, we are in a state of equi-biaxial stress with no shear stress. Cesana, Plucinsky, and Bhattacharya (2015) predicted a region of equi-biaxial stress in ideal materials ($\alpha = K = 0$). This behavior remains in non-ideal materials.

In the case of equi-biaxial stretch (EB, Figure 3.8), the directors orient gradually to become planar, but there is little, if any, evolution beyond that. There is also consequently significant heterogeneity in the state of stress. Together, the three results show that shear of unequal stretch drives microstructure evolution.

Comparison with experiments We compare the results of our simulations with the experimental observations of Tokumoto et al. (2021). To do so, we pick values for ν_n and K consistent with prior work. We then fit the remaining four parameters (r, μ, α, ν_F) to the experimental (nominal) stress-stretch relationship for UNI and PE using least squares (i.e., minimizing the L_2 norm). These lead to the values reported at the beginning of the section. The comparison between the experimental and simulated stress-stretch relations is shown in Figure 3.9. It also serves as a verification of our simulation. We see the fit as well as the agreement in EB (which is not used in the fit) is excellent. Figure 3.10 compares the results of all the mechanical tests. Again, the agreement is excellent.

Finally, we compare the statistics of the mesogen distribution. This can be measured using wide-angle x-ray scattering (WAXS) measurements. The patterns (Figure.3.11(a)) shows the statistics of mesogens orientation in the plane of the LCE sheet: the brightness indicates a component the density of mesogens with component aligned in that direction. The patterns start with uniform distribution, which shows mesogens are randomly distributed. In EB stretching, the pattern remains almost isotropic regardless of strain, exhibiting that the local nematic director stays nearly random in the x-y plane. In UB stretching including PE, as strain is applied, a diffuse peak arises in the patterns. The azimuthal redistribution of intensity proves the reorientation of mesogens of LCEs. Mesogens are reorienting towards the main stretching direction. The azimuthal distribution of mesogen orientation
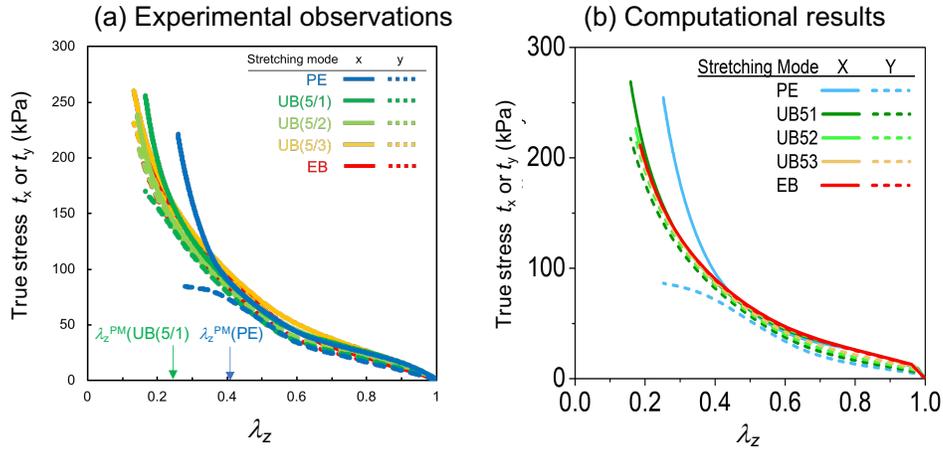
Figure 3.10: Comparison of the true stress vs. thickness stretch. (a) Experimental observations. (b) Computational results.
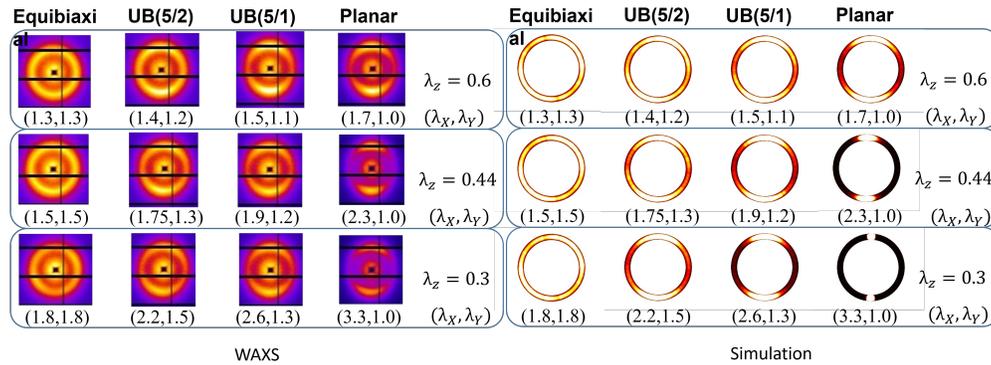
Figure 3.11: 2D wide-angle diffraction patterns in various types of biaxial strain for LCEs obtained by (a)experiments; (b)the full-field 3D simulations with the free energy including non-ideal and Frank elasticity and a time-dependent evolution for the director.

is also calculated from the simulations(Figure.3.11(b)). Again the agreement is excellent for various stretches indicating that the theoretical model in fact captures the microscopic physics accurately.

This level of agreement is remarkable because all simulations are carried out with only six parameters.

C h a p t e r 4

MULTISTABLE STRUCTURES

4.1   Introduction

In this chapter, we apply one of the key ideas of accelerated computational micromechanics–the use of augmented Lagrangian instead of a restricted basis set to impose kinematic compatibility–to macroscopic mechanical structures. The focus is on multistable deployable structures inspired by origami. Origami is the ancient Japanese art of paper folding, where a flat sheet is folded into complex shapes with (relatively) rigid and flat facets and (relatively) flexible creases. The complexity of shapes that such structures can achieve, and the ability of some such structures to deform freely have inspired applications in science and engineering including depolyable space structures (Schenk, Viquerat, et al., 2014), robotics (Rus and Tolley, 2018), biomedical devices (M. Johnson et al., 2017), architecture (Reis, Jiménez, and Marthelot, 2015), and metamaterials (Silverberg et al., 2014; Y. Li, 2020; B. Liu et al., 2018). The systematic design of origami fold-patterns and structures have therefore become an area of active research.

A natural question is whether a structure is "rigid foldable": i.e., can the structure be deformed with rigid facets. Some fold patterns are not rigid foldable. Other patterns lead to an achievable shape, but there is no rigid foldable path from the flat to the shape. There are other structures that are extremely flexible with multiple families of rigid foldable shapes and paths. And then there are those that are multi-stable: they have multiple rigid foldable shapes but no rigid foldable paths between them. Another question is to understand the energetic cost of deformation and low energy pathways between rigid shapes when a structure is not rigid foldable.

For rigid-foldable patterns, once can use the geometry to parametrize the origami, and one can obtain explicit closed-form expressions for periodic and regular patterns (Gattas, W. Wu, and You, 2013; K. Wang et al., 2011; Zakirov and Alekseev, 2010; Hanna et al., 2014). This has been used extensively to study patterns such as Miura-Ori pattern (Gattas, W. Wu, and You, 2013) and generalized quadrilateral patterns (Tachi, 2009a; Lang, Magleby, and Howell, 2016) with one degree of freedom (DOF). However, it is difficult to use such an approach for irregular patterns.

A variety of numerical methods have been developed to study general patterns,

and they can generally be categorized into two sets: angle/plate based model and vertex/truss based model. The angle-based model assumes the facet to be rigid and describes the structure using dihedral angles between facets. For facets around a single vertex, the dihedral angles satisfy the loop closure constraint (Hull, 2002; Tachi, 2009b; Tachi, 2012). Hull (2002) described the constraint by enforcing the product of rotation matrices around a single vertex to be identity matrix. Tachi (Tachi, 2009a; Tachi, 2012) projected the infinitesimal increments of folding angles into the linearized constraint space and further extended for origami with holes. W. Wu and You (2010) proposed a rotating vector model based on quaternion and dual-quaternion following the same loop closure constraint. Hu and Liang (2020) enforced the constraint with a Lagrange multiplier.

Vertex-based model describes the origami structure using vertex coordinates. A typical example is the bar-hinge model (Schenk and Guest, 2011). The in-plane deformation is restricted by bars. The out-of-plane rotation and folding are described by hinges, where the dihedral angle is represented by inner and cross product of vectors. If the bars are restricted with no extension and hinges with no rotation, then the model describes a rigid origami. The bar-hinge model can naturally be extended to non-rigidly-foldable configuration (Wei et al., 2013; Filipov, K. Liu, et al., 2017) and include complex material behaviors(K. Liu and Paulino, 2017; Cehula and Prša, 2020). Filipov, Tachi, and Paulino (2015) and Filipov, K. Liu, et al. (2017) introduced extra bars and nodes to incorporate detailed in-plane deformation and out-of-plane folding. K. Liu and Paulino (2017) includes nonlinear constitutive relations of bars and hinges to model non-rigid behaviors. Y. Li (2020) proposed a mixed method where both vertex coordinates and angles are constrained.

A more recent approach is to consider deformable structures, and use the finite element method (Pinson et al., 2017; Ma and You, 2014; Q. Zhang et al., 2017). However, this suffers from significant computational expense.

Much of this work above seeks to identify the equilibrium shapes of a given fold pattern. However, it is often necessary to understand transitions or trajectories between stable shapes in multi-stable structures. This is an emerging area of research. Song and Amato (2004) proposes a tree structure model using probabilistic roadmap methods. Each folding step is independently stored in the tree structure and thus may not satisfy global compatibility constraint. Tachi (2009b) proposes a method of projecting the trajectory onto the linearized constraint space. Xi and Lien (2014) and Xi and Lien (2015) randomly samples between the initial and target configurations.

The compatible configuration along the path is then found by this random search, which is computationally expensive. Y. Li (2020) uses a shooting method. The increment is calculated by projecting the difference between current configuration and target to the linearized constraint space. This method can be slow if the path is "wavy" since the increment vector could be perpendicular or even opposite to the tangential of the path.

An important difficulty in all these methods is the enforcement of the kinematic compatibility condition that is geometrically non-linear. In this chapter, we propose a new approach that uses both vertices and facets, and then formulates the kinematic compatibility condition as a constraint.

4.2   Method

We use the following notation convention. Lowercase letters describe scalars or vectors, while uppercase letters represent constants, matrices, or subsets. Bold uppercase letters represent sets. Subscript i enumerate vertices, subscript j enumerate facets, subscript k enumerate states and n timestep.

Kinematics

We consider a structure with I vertices, whose coordinates forms the set $\mathbf{Y} = \{y_i | i = 1, 2, \ldots, I\}$. The structure is composed of J facets or rigid components. The jth rigid component includes a subset of vertices, $F_j \subset \{1, 2, \ldots, I\}, j = 1, 2, ..., J$. Together they form the set $\mathbf{F} = \{F_j\}, j = 1, 2, \ldots, J$. Similarly, the ith vertex shares multiple rigid components, $G_i \subset \{j = 1, 2, \ldots, J\}$. Together they form the set $\mathbf{G} = \{G_i\}$.

Within each rigid component, the center of rigid component is defined as the average of all vertex coordinates in the set F_j,

$$c_j = \frac{1}{|F_j|} \sum_{y_i \in F_j} y_i \tag{4.1}$$

where $|\cdot|$ represents the number of elements in a set.

In the initial configuration, represented by subscript 0, the coordinate of each vertex in the rigid component is as

$$y_{0i} = c_{0j} + r_{0ji}, \quad \forall i \in F_j \tag{4.2}$$

where r_{0ji} is the position vector of the ith vertex relative to the center of the jth rigid. Note that we do not require a structure that is flat in the reference configuration.
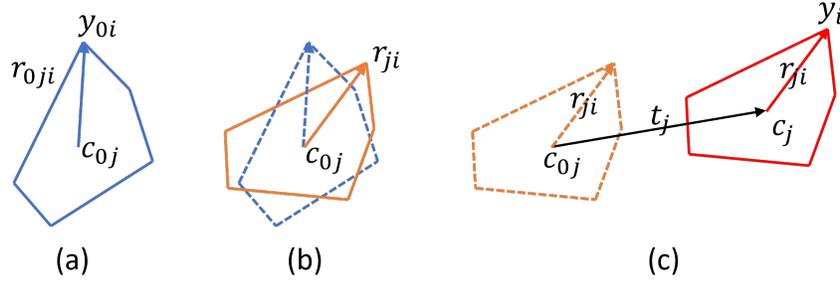
Figure 4.1: Kinematics of a rigid component: (a) initial configuration; (b) rotation; (c) translation.

When the structure deforms, each rigid component undergoes a translation t_j and rotation R_j. In the new configuration, the ith vertex has the position

$$y_i = c_j + r_{ji} = c_{0j} + t_j + R_j r_{0ji}, \quad \forall\, i \in F_j \tag{4.3}$$

where c_j is the center, and r_{ji} is the relative position. Therefore, $c_j = c_{0j} + t_j$, $r_{ji} = R_j r_{0ji}$. However, the rigid components do not deform independently; thus c_j, R_j are not independent, but satisfy

$$c_{0j} + t_j + R_j r_{0ji} = c_{0j'} + t_{j'} + R_j r_{0j'i} \quad \forall\, j, j' \in G_i. \tag{4.4}$$

We say that a configuration is rigid foldable if the constraint (4.4) is satisfied.

Finding stable and rigid foldable configurations

We are interested in structures that are rigid foldable as well as structures that are not. Further, if a structure is not rigid foldable, we seek to understand how close the structure is to being rigid foldable. This enables us to identify almost rigid foldable paths between two rigid foldable states. Therefore, we do not require the constraint (4.4) exactly, but use a penalty away from it. Specifically, we regard $\mathbf{Y} = \{y_1, y_2, \dots y_I\}$, $\mathbf{C} = \{c_1, c_2, \dots c_J\}$ and $\mathbf{R} = \{R_1, R_2, \dots, R_J\}$ as independent variables and define an energy

$$W(\mathbf{Y}, \mathbf{C}, \mathbf{R}) = \sum_{j \in J} \sum_{i \in F_j} (y_i - c_j - R_j r_{ji})^2. \tag{4.5}$$

Note that this is non-negative and zero if and only if the configuration is rigid foldable.

We define stable configurations to be those that minimize the energy (4.5) over $(\mathbf{Y}, \mathbf{C}, \mathbf{R})$ using alternating directions.

We first minimize with respect to vector coordinates y_i. Since W is quadratic in y, it is straight forward,

$$y_i = \frac{1}{|T_i|} \sum_{j \in T_i} (c_j + R_j r_{ji}) \tag{4.6}$$

where $T_i = \{j | y_i \in F_j\}$ is the index set of all rigid components that includes y_i. Then, we minimize with respect to rigid body center c_j and this is again easy since W is quadratic in c_j. We obtain

$$c_j = \frac{1}{|F_j|} \sum_{i \in F_j} (y_i - R_j r_{ji}) = \frac{1}{|F_j|} \sum_{i \in F_j} y_i, \tag{4.7}$$

The second equality follows from the fact that $\sum_{i \in F_j} r_{ji} = 0$. Note that c_j is independent of R_j.

Finally, we minimize with respect to R_j. W is quadratic in R_j, but R_j is a rotation matrix. So we introduce quaternions which allows us to solve the problem (Kuehnel, 2003). A quaternion p is a combination of a scalar q_0 and a vector $q = \{q_1, q_2, q_3$ written as $p = (q_0, q)$. It follows the multiplication rules

$$pp' = (q_0 q'_0 - q \cdot q', q_0 q + q'_0 q - q \times q'). \tag{4.8}$$

Further, a unit quaternion (i..e., one that satisfies $q_0^2 + |q|^2 = q_0^2 + +q_1^2 + q_2^2 + q_3^2 = 1$) can be identified with the rotation R following

$$R = \begin{pmatrix} 1 - 2q_2^2 - 2q_3^2 & 2(q_1 q_2 - q_0 q_3) & 2(q_1 q_3 + q_0 q_2) \\ 2(q_1 q_2 + q_0 q_3) & 1 - 2q_1^2 - 2q_3^2 & 2(q_2 q_3 - q_0 q_1) \\ 2(q_1 q_3 - q_0 q_2) & 2(q_2 q_3 - q_0 q_1) & 1 - 2q_2^2 - 2q_2^2 \end{pmatrix}. \tag{4.9}$$

We can rewrite the energy with respect to the quaternion. p_j minimizes the following functional,

$$p_j = \text{argmin}_{p_j} \sum_{i \in F_j} (y_i - c_j - R_j(p_j) r_{ji})^2 = \text{argmin}_{p_j} \sum_{i \in F_j} p_j B_{ji}^T B_{ji} p_j^T \tag{4.10}$$

where B_{ji} is a 4×4 matrix,

$$B_{ji} = \begin{pmatrix} 0 & (r_{ji} + c_j - y_i)^T \\ (c_j - y_i - r_{ji}) & [c_j - y_i + r_{ji}]_\times \end{pmatrix} \tag{4.11}$$

and the cross product operator is defined as

$$[(r1, r2, r3)]_\times = c \begin{pmatrix} 0 & -r3 & r2 \\ r3 & 0 & -r1 \\ -r2 & r1 & 0 \end{pmatrix},$$

Algorithm 4: Configuration finding

Given an initial value of $\mathbf{Y}^0, \mathbf{C}^0, \mathbf{R}^0$, tolerance ϵ and position vectors r_{ji},
$\quad i = 1, 2, 3, ..., I, \ j = 1, 2, 3, ..., J;$
while $|W^{n+1} - W^n| > \epsilon$ **do**

$\quad$ **for** $i = 1, 2, 3, ..., I$ **do**

$\quad\quad \bigg| \quad y_i^{n+1} = \frac{1}{|T_i|} \sum_{j \in T_i} (C_j^n + R_j^n r_{ji})$

$\quad$ **end**

$\quad$ **for** $j = 1, 2, 3, ..., J$ **do**

$\quad\quad \bigg| \quad C_j^{n+1} = \frac{1}{|F_j|} \sum_{i \in F_j} y_i^{n+1}$

$\quad$ **end**

$\quad$ **for** $j = 1, 2, 3, ..., J$ **do**

$\quad\quad$ **for** $i = 1, 2, 3, ..., I$ **do**

$\quad\quad\quad \bigg| \quad B_{ji} = \begin{pmatrix} 0 & (r_{ji} + C_j - y_i)^T \\ (C_j - y_i - r_{ji}) & [C_j - y_i + r_{ji}]_\times \end{pmatrix}$

$\quad\quad$ **end**

$\quad\quad p_j^{n+1} = \text{argmin}_{p_j} \sum_{i \in F_j} p_j B_{ji}^T B_{ji} p_j^T$

$\quad$ **end**

$\quad$ Update $W^{n+1} = \sum_{j \in J} \sum_{i \in F_j} (y_i^{n+1} - C_j^{n+1} - R_j^{n+1} r_{ji})^2$

end

The minimizer of (4.10) is given by the eigenvector (with the least eigenvalue) of the corresponding matrix.

Finally, we need boundary conditions to prevent overall rigid body deformation or to specify some overall shape. Typical boundary conditions combine freezing some vertices and restricting others to a plane. It is also common to prescribe a time-dependent boundary condition to define a path. For example, to find the path of expanding a Miura-Ori origami, we could prescribe the coordinates of two vertices.

The algorithm is summarized in Algorithm 4.

Path finding

Origami-inspired structures are often used as deployable structures. A common strategy here is to seek structures that are multi-stable, i.e., have more than one discrete rigid foldable shapes. This simplifies the control strategy since one has to push the structure from one rigid foldable state to another. In such situations, one seeks the low energy path between the two stable states.

Consider a structure with two stable (or rigid foldable states) $\mathbf{Y}_0$ and $\mathbf{Y}_K$. We seek the lowest energy path between these states. To do so, we introduce the reduced
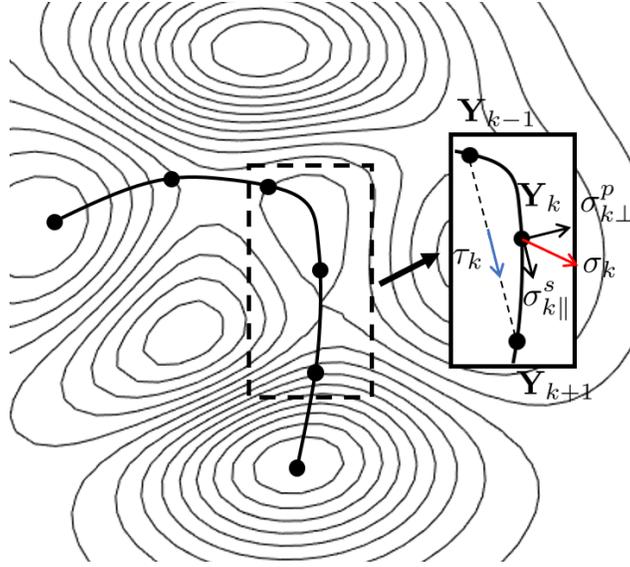
Figure 4.2: Illustration of nudged elastic band method.

energy which treats $\mathbf{Y}$ as the variable:

$$W_c(\mathbf{Y}) = \min_{\mathbf{C},\mathbf{R}} \sum_{j \in J} \sum_{i \in F_j} (y_i - C_j - R_j r_{ji})^2.$$

This energy is also zero if and only if $\mathbf{Y}$ is rigid foldable.

We now use the *nudged elastic band* (NEB) method that is widely used in materials science (Herbol, Stevenson, and Clancy, 2017) to find the path between $\mathbf{Y}_0$ and $\mathbf{Y}_K$ according to the energy W_c. We specifically seek to find a a chain of states $\mathbf{Y}_0, \mathbf{Y}_1, ..., \mathbf{Y}_K$ connecting $\mathbf{Y}_0$ and $\mathbf{Y}_K$ that follows the minimum energy path–see Figure 4.2. Since we want this path to have low energy, we could minimize the energy with respect to $\mathbf{Y_k}$, but that would lead all the states to fall either at $\mathbf{Y}_0$ or $\mathbf{Y}_K$. We try to make the intermediate states roughly equidistant to each other by penalizing the distance between adjacent states with linear springs. Therefore we consider the energy

$$W_t(\mathbf{Y}_0, \mathbf{Y}_1, ..., \mathbf{Y}_K) = \sum_{k=1}^{K-1} W_c(\mathbf{Y}_k) + \sum_{k=1}^{K} c_k (\mathbf{Y}_k - \mathbf{Y}_{k-1})^2 \qquad (4.12)$$

for some constants $c_k > 0$.

This energy is usually updated by a gradient/force-based method. Note that there are two contributions to the gradient, also called the force on the point:

$$\sigma_k^p = \nabla W(\mathbf{Y}_k), \quad \sigma_k^s = c_k(\mathbf{Y}_{k+1} - \mathbf{Y}_k) - c_k(\mathbf{Y}_k - \mathbf{Y}_{k-1}). \qquad (4.13)$$

An important idea in NEB is to project the first component σ_k^p coming from the energy tangential to the path and the second component σ_k^s coming from the springs; i.e., define the force as

$$\sigma_k = \sigma_{k\perp}^p + \sigma_{k\parallel}^s \tag{4.14}$$

where the subscripts $\perp$, $\parallel$ represent, respectively, the parallel and perpendicular component with respect to the tangent τ_k The idea is that the energy landscape rearranges the points along the path, and the distance constraints rearranges the path to maintain the distance. This avoids certain pathologies. The discrete tangent τ_k to the path at the state Y_k is approximated by the neighboring states as shown in Figure 4.2. Finally, to further accelerate the computation by avoiding jagged minimization, we use a conjugate gradient approach where the descent direction s^{n+1} is always conjugate to the previous steps:

$$s^{n+1} = \sigma^{n+1} + \beta s^n$$

where

$$\beta = \frac{< \sigma^{n+1}, \sigma^{n+1} >}{< \sigma^n, \sigma^n >} \tag{4.15}$$

$< \cdot, \cdot >$ is the inner product, S and σ without subscript stand for the entire set $\{\sigma_k | k = 0, 1, 2, .., K\}$ and $\{s_k | k = 0, 1, 2, .., K\}$

4.3   Examples

Miura-Ori

We start with the simple example of the unit cell of the Miura-Ori pattern with four patterns. The results of folding when certain points are prescribed is shown in Figure 4.3. The points A, B are fixed and point D is constrained in the $x - y$ plane to prevent rigid body deformation of the entire structure. The deformation path is designed by lifting point C (controlling the z-coordinate of C). We find the minimum for each prescribed height of the point C. Figure 4.4 shows another result for a Miura-Ori pattern when the folds are perpendicular to each other. We prescribe the two end states and seek the low (zero) energy path between them. This is non-trivial since the operations have to be sequenced properly first unfolding it and then folding it in the other direction. The NEB does this naturally as shown in the figure.

Multiple degree of freedom patterns

We now study more complex patterns, in particular those that have multiple families of rigid foldable states. The first is the Ron Resch pattern shown in Figure 4.5.

Algorithm 5: Path finding

Given the initial state $\mathbf{Y}_0$ and final state $\mathbf{Y}_K$, tolerance ϵ and position vectors r_{ji}
  $,i = 1, 2, 3, ..., I,\ j = 1, 2, 3, ..., J$;
Initialize $\mathbf{Y}_k^0$, $k = 1, 2, 3, ..., K - 1$ by linear interpolation between $\mathbf{Y}_0$ and $\mathbf{Y}_K$.;
Calculate τ^0, σ^0;
while $|W_t^{n+1} - W_t^n| > \epsilon$ **do**
    if $n = 0$ **then**
        $\sigma^n = s^n$;
    else
        if $< \sigma_{n-1}, \sigma_{n-1} >= 0$ **then**
            $\beta = 1$;
        else
            $\beta = \frac{<\sigma^n, \sigma^n>}{<\sigma^{n-1}, \sigma^{n-1}>}$;
        end
    end
    $s^n = \alpha s^n$;
    $\mathbf{Y}^{n+1} = \mathbf{Y}^n + s^n$;
    Update τ^{n+1}, σ^{n+1};
    Update α using backtrack line search;
end
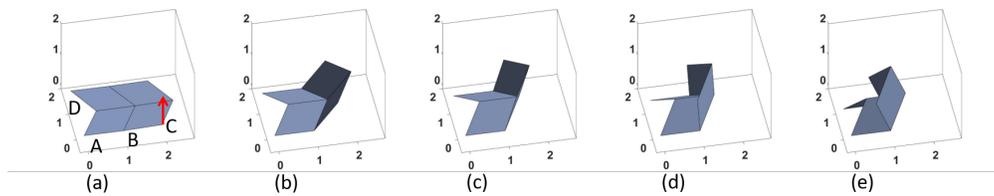
Figure 4.3: Snapshots of intermediate state of folding a 4–piece Miura-Ori pattern by prescribing points A, B, C, D.

This pattern has two flat states shown in Figure 4.5 (a) and (f). Figure 4.5(a) also shows the fold lines. We prescribe the two end states and use the NEB algorithm to find the intermediate states, and these are shown in the figure. Figure 4.6 shows the analogous results for the Resch-waterbomb pattern invented by Robert Lang. In both cases, the algorithm finds the appropriate sequence of operations through the various families of rigid foldable structures.

Generalized folding structures

Our final example consists of an Origami-inspired structure invented by Robert Lang and studied in detail by L. Yang, Robert, and Pelligrino, 2020.
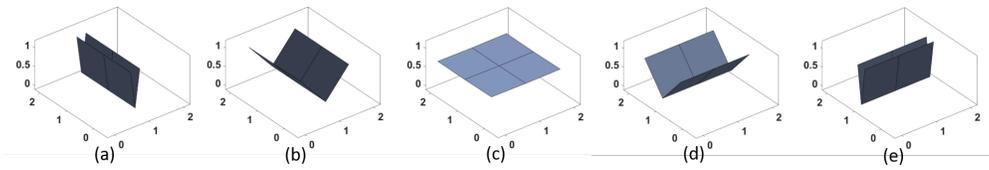
Figure 4.4: Snapshots of intermediate state of folding a 4–piece Miura-Ori pattern with 90° cut. The path connecting (a) and (f) is found by NEB method.
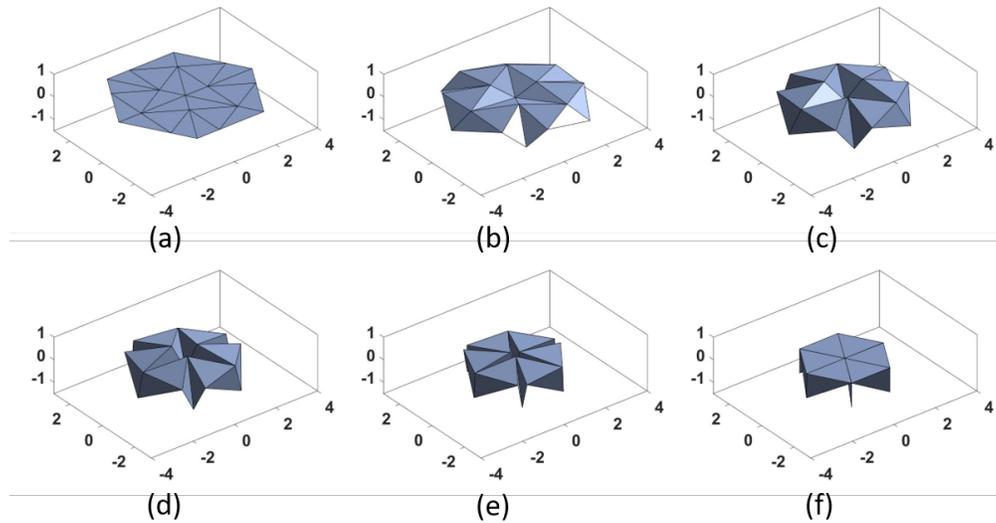
The

Figure 4.5: Snapshots of intermediate state of folding a Ron Resch pattern with 90° cut. The path connecting (a) and (f) is found by NEB method.

The 3×3 version of the structure is shown in Figure 4.7. This consists of a 3×3 array of (non-planar) rigid components, each having a square face and perpendicular protrusions. The elements connect to each other via folds or hinges between the protrusions of neighboring elements. Note that there are no "diagonal" connections. This structure only has one degree of freedom. We use our energy minimization algorithm to study the structure. We apply boundary conditions to fix the rigid body motion to the central component, and then pull two corners towards each other. The results are shown in Figure 4.8. We find a sequence of rigid foldable configurations that takes it from one shape to another.

The 5×5 version of the structure is shown in Figure 4.9. We apply a boundary condition to the central component to eliminate the overall rigid body rotation. We find two stable configurations shown in Figure 4.9 (a) and (f). The energy
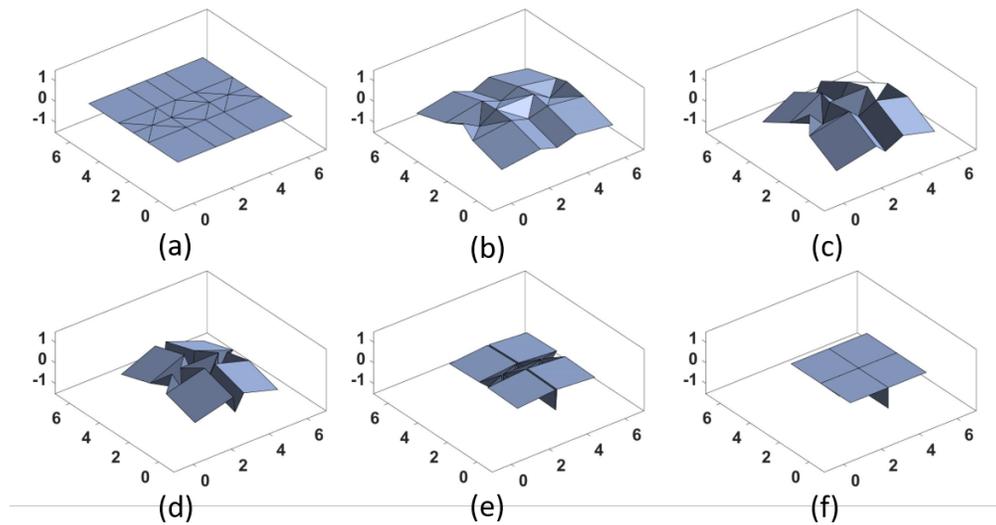
Figure 4.6: Snapshots of intermediate state of folding a waterbomb pattern with 90°
cut. The path connecting (a) and (e) is found by NEB method.
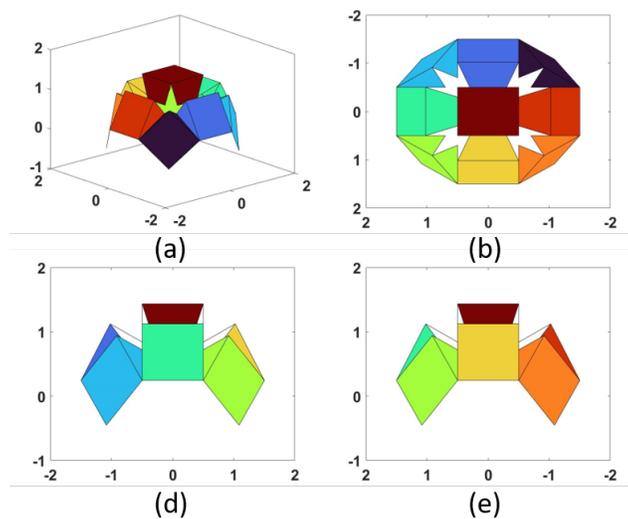
Figure 4.7: Illustration of a 3 by 3 origami inspired structure: (a) overview; (b) top
view; (c) side view; (d) front view.
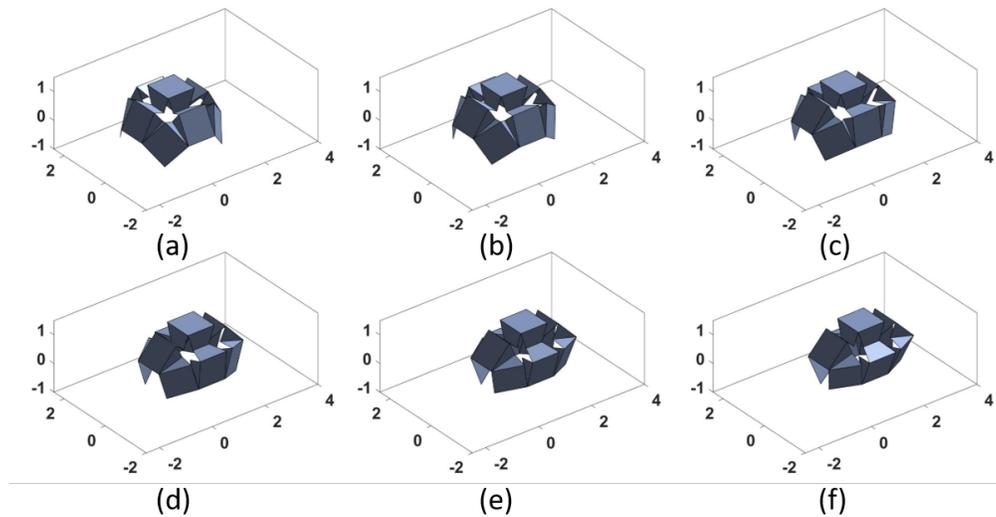
Figure 4.8: Snapshots of intermediate state of a 3×3 structure; the path is found by pulling the bottom corner two points.

minimization algorithm is unable to find any rigid foldable path between these two stable states. Therefore, we use the NEB algorithm to study the transition between the two states, and this confirms the lack of any rigid foldable path between these states. Figure 4.11(a) shows the energy barrier over various iterations. The low energy path corresponds to the states shown in Figure 4.9, and we can recover these using the energy minimization algorithm with the height of the structure prescribed. Figure 4.11(b)) shows the energies of the intermediate states obtained using energy minimization and this corresponds to the results of NEB.

4.4 Conclusion

In this chapter, we proposed a method to study the folding mechanism inspired by origami. The method describes such structures using both vertices and faces and then using an penalty to describe the compatibility between them. We find stable and rigid foldable configurations by minimizing the penalty following an alternating direction approach. We also find the lowest energy transition pathways using a nudged elastic band method. We demonstrate the method with a series of examples from origami and origami-inspired structures.

While this work only addresses rigid structures, the method can naturally be extended to incorporate other mechanisms. For example, springs or torsional springs could be added to the structure, or the facets could be flexible. We would then add appropriate terms to the energy and use the same minimization and NEB algorithms to find a
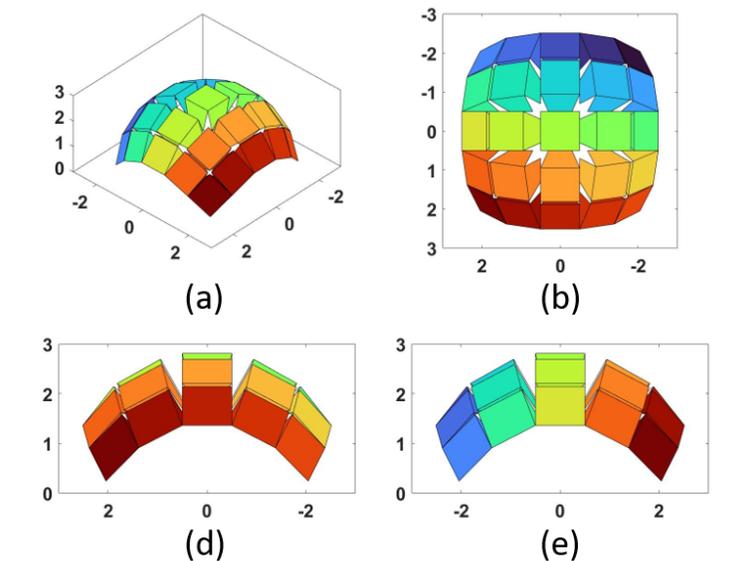
Figure 4.9: Illustration of a 5 by 5 origami inspired structure: (a) overview; (b) top view; (c) side view; (d) front view.
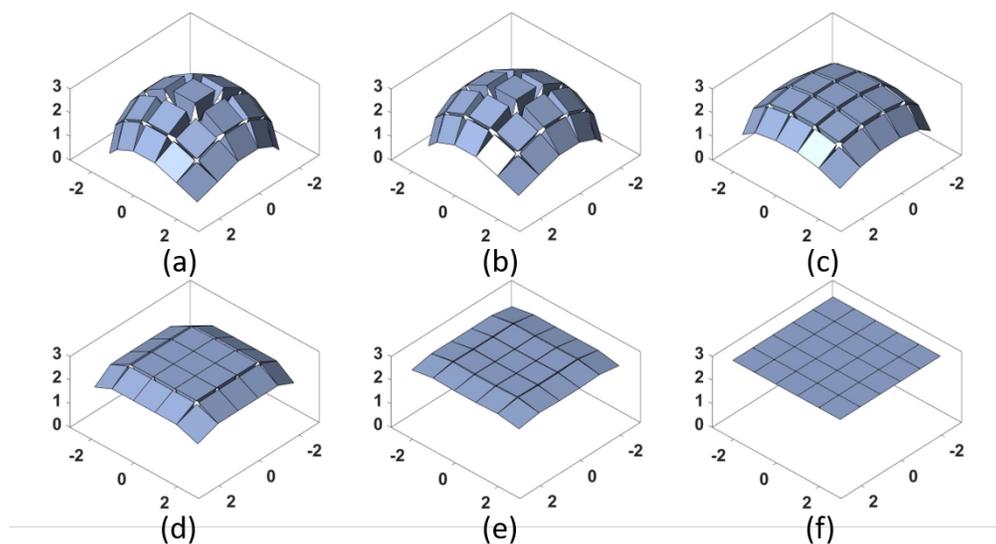
Figure 4.10: Snapshots of intermediate state of 5×5 structure; the path is found by pushing the top piece.
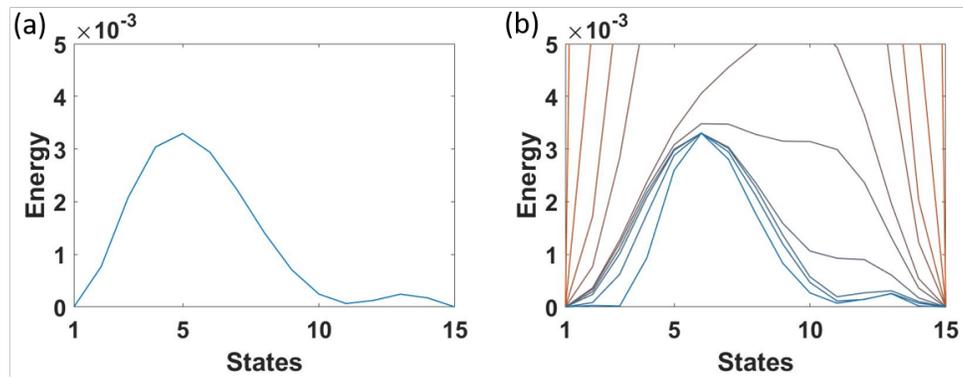
Figure 4.11: The configuration energy W_c during deformation of a 5×5 structure: (a) NEB path at different iterations(from 0 to 300 iterations); (b) prescribed path.

stable structure and low energy paths. Thus, the method is extremely versatile in studying multistable structures.

Chapter 5

EXPERIMENTAL CHARACTERIZATION OF STRESS FIELD

5.1 Introduction

Emerging non-destructive X-ray diffraction (XRD) techniques for 3-D material characterization performed at light sources, collectively known as 3DXRD (e.g. Poulsen, 2004; Ludwig et al., 2009; Lienert et al., 2011), provide previously inaccessible in-situ microstructural and micromechanical information on polycrystalline materials at grain and sub-grain scales. In-situ 3-D measurements are enabled by the collection of multiple diffraction patterns as the sample is rotated and deformed. These patterns are then used in combination with forward-modelling of the diffraction experiment to optimize the agreement between those measured detector images and simulated patterns corresponding to every possible orientation and micromechanical field configuration. Near-field (nf) techniques (Suter et al., 2006; S. Li and Suter, 2013; Viganò et al., 2016) are used to obtain crystal orientation fields, resulting in space-resolved voxelized microstructural images with intragranular resolution. Further, far-field (ff) techniques (Poulsen et al., 2001; Bernier et al., 2011; Oddershede et al., 2010) provide information on local micromechanical fields. Due to a trade-off between direct and reciprocal space resolutions, the original ff data inversion methods were able to deliver average stresses/elastic strains in single crystal grains and the volume and location of the centers of mass of those grains, but not voxelized intragranular fields.

Recently, 3DXRD data inversion methods were improved to provide not only voxelized crystal orientation but also stress fields with intragranular resolution using a variety of experimental procedures, including a microbeam (1D), line beam (2D) or box beam (3D) illumination, each demanding a tailored optimization-based data processing approach (Reischig and Ludwig, 2019; Hayashi et al., 2019; Reischig and Ludwig, 2020; Shen, H. Liu, and Suter, 2020; Henningsson et al., 2020). However, these types of adopted optimization methods have been exclusively based on imposing diffraction constraints, by minimizing the difference between experimental and simulated detector images obtained by forward-modelling of the diffraction experiment, and retaining as *measured* fields the orientation and stress fields that produced that minimum. In either case, using the original data inversion based on

grain averages, or the more advanced methods that deliver voxelized information, the resulting stress fields are in general unbalanced (not divergence-free, violating the stress equilibrium condition).

Since the micro-mechanical field that directly affects diffraction patterns—by changing the local lattice spacing—is elastic strain, accounting for stress equilibrium at the sub-grain level requires a fairly accurate knowledge of the single crystal properties, such as zero-stress lattice parameters, and single crystal elastic constants. Thus, failing to fulfill mechanical equilibrium in the solution may be unavoidable. In other instances, ignoring the latter allows for a simplified and faster solution process, where large-scale optimization is already challenging and computationally intensive. However, when the material properties are known, taking into account mechanical equilibrium clearly provides a more constrained solution, and may significantly improve the stress/strain fields determination, and consequently, the spatial resolution, reliability, and application range of 3DXRD methods. The elastic deformation solvers are typically iterative, and thus mechanical constraints could be employed in one of many ways: to correct the final stress field solution, in every iteration step by enforcing equilibrium in the latest solution, or as an additional component in an optimization target function.

Given the aforementioned problem, an approach to solve the lack of equilibrium of the experimental stress field is to combine micromechanical modelling and simulations with experimental observations. In different contexts, this combination has been recently done by Pokharel and Lebensohn (2017) using a fast Fourier transform (FFT)-based elasto-viscoplastic (EVPFFT) model, and by Chatterjee et al. (2017) adopting a field dislocation mechanics finite element model, to complement X-ray diffraction based experimental information and characterize residual and internal stresses. Also, Pagan and Beaudoin (2019) used lattice orientation and crystal plasticity kinematics to recover the geometry, and further calculated the stress through finite element simulation.

While the above authors solve physics-based field equations, McNelis, Dawson, and Miller (2013) proposed a new approach where they seek to impose equilibrium to match the lattice stresses inferred from XRD and elasticity. In particular, they were interested in large mechanical parts (at the order of meters) where the detection spots are limited. Therefore, they proposed a two-scale method, where the continuum scale stress field imposes equilibrium to match the lattice scale stresses. The approach was extended to three dimensions by Demir et al. (2013), and has found successful

applications (Park, Lienert, et al., 2013; Park, Yildizli, et al., 2018). However, this approach does not impose equilibrium at the grain and sub-grain scale, which is the focus of our work.

This chapter presents a novel methodology to impose micromechanical constraints, i.e. stress equilibrium at the subgrain and grain scale, to an non-equilibrated voxelized stress field obtained, for example, by means of synchrotron X-ray diffraction techniques. The main idea is to find the equilibrated stress field closest (in L^2-norm sense) to an experimental (and possibly) non-equilibrated stress field, via the solution of an optimization problem based on the Hodge decomposition of a symmetric 2nd rank tensorial field. The Hodge decomposition is a generalization of the classical Helmholtz decomposition that states that any sufficiently smooth vector field can be decomposed into the sum of an irrotational vector field and a solenoidal vector field. We use the version of the Hodge decomposition for symmetric tensorial fields developed by Geymonat and Krasucki (2009). We present an efficient numerical implementation of the Hodge decomposition of symmetric 2nd rank tensorial fields using FFTs, and apply it to the analysis/determination of equilibrated stress fields in polycrystalline materials, including 3-D space-resolved stress inferred from X-ray diffraction. As mentioned before, these stress fields are in general not divergence-free, due to piecewise approximation based on grain averages and/or absent the consideration of any micromechanical constraint in the data inversion procedure. Using the proposed decomposition and formulating an optimization problem, unbalanced stress fields can be filtered to extract their divergence-free part. The proposed methodology has been further modified/extended to deal with interfaces and discontinuities occurring at grain boundaries. We begin by demonstrating that the method is capable of good recovery of synthetic data obtained in the elastic regime with an FFT-based micromechanical model that provides an objective ground-truth (a validation methodology originally used by Shen, H. Liu, and Suter (2020)), and then apply it to a stress field actually obtained from diffraction data.

The outline of the chapter is as follows. In Section 2, we present a summary of the 3D orientation and strain mapping approach developed by Reischig and Ludwig (2019), and its application to Gum Metal, a beta-Ti-based alloy, deformed and measured in-situ within the elastic regime. In Section 3, we recall the fundamentals of the Hodge decomposition and formulate the optimization problem that allows us to extract the divergence-free part of a general non-equilibrated stress field. In Section 4, we present the FFT-based method that enables an efficient numerical resolution

of the previously formulated optimization problem when applied to a voxelized field, the type of data structure that naturally results from data reduction of 3DXRD experiments. Section 5 shows applications of the proposed method to synthetic piecewise constant stress fields with a known ground truth, and stress fields in Gum Metal inferred from 3DXRD. In both cases, we show that the largest corrections are obtained near grain boundaries. In Section 6, we draw conclusions and give perspectives of the adoption of the proposed method in 3DXRD data inversion packages.

5.2 Experimental characterization of elastic strain tensors

Data acquisition

The experimental dataset used in this work to demonstrate the proposed stress filtering method was recorded in a regular Diffraction Contrast Tomography (DCT) scan (Ludwig et al., 2009; Reischig, King, et al., 2013). The method uses a parallel monochromatic synchrotron X-ray beam and records diffraction spots from individual grains of the polycrystalline sample on a near-field area detector while the sample is rotated continuously over $360°$. The image stack is subject to a number of pre-processing steps, at the end of which the diffraction spots are segmented and indexed according to their grain of origin and their (hkl) Miller indices. The indexing is based on crystallographic principles and the knowledge of the crystal structure and approximate undeformed lattice cell parameters of the one (or possible several) phases present in the polycrystal. Each diffraction spot is stored as a (u, v, w) volume array representing the diffracted intensity distribution as a 3D scalar field: u and v being the horizontal and vertical image coordinates, and w is the rotation angle. The use of monochromatic X-rays and a high-resolution detector provides sensitivity to the local orientation and local unit cell parameters within the crystallites. Local variations in lattice parameters, i.e. a change of the shape and dimensions of the unit cell, in this case is interpreted as mechanical elastic strain, and any effects from a possible local change of chemical composition is neglected.

The results presented in this study were obtained from a polycrystalline Gum Metal sample with a composition of Ti-36Nb-2Ta-3Zr-0.3O $wt\%$ and mean grain diameter of 61μm, from the "low load" scan presented in (Reischig and Ludwig, 2020). Gum Metal is able to sustain elastic strains up to 2% elongation and above, thus its nickname. A 500μm tall gauge volume of a tensile specimen with an approximately 600μm wide rectangular cross-section was illuminated in its entirety with a box beam during the $360°$ rotation, using 40 keV beam energy, 1.5 sec exposure time per

image, 0.05° angular step size, resulting in 7200 images in total. The detector pixel size was 1.4μm, and the rotation axis to detector distance was 7 mm. The sample was mounted in a miniature tensile rig (Gueninchault, Proudhon, and Ludwig, 2016) and a DCT scan at a uniaxial external tensile stress of 34 MPa was acquired. On average 29 diffraction spots from the first three families of reflections could be recorded for each of the ~ 1430 grains in the illuminated sample volume, which served as the input data for the reconstruction.

Reconstruction of the sub-grain elastic deformation field

The first stage of the data analysis follows the now standard DCT processing route, with a calibration procedure based on the dataset itself (Reischig and Ludwig, 2019; Reischig and Ludwig, 2020). The centroid positions and mean orientations of the grains were determined based on their diffraction spot centroid metadata. The initial 3D shapes of the grains were reconstructed iteratively from the diffraction spots using the SIRT algorithm, and assuming a constant orientation and elastic strain distribution throughout the grain volume.

The deformation analysis applied here corresponds to an earlier version of the novel Iterative Tensor Field (ITF) Reconstruction method described in detail in Reischig and Ludwig (2020), the main difference being that here no Tikhonov regularization term was applied in the elastic deformation optimization target function (Equation (39) in Reischig and Ludwig (2020)). The ITF method was used to retrieve the complete local elastic strain tensor (6 parameters), the local lattice misorientations from a reference grain orientation (3 parameters) and refine the grain boundaries over the entire 3D voxelated gauge volume in the specimen. ITF is based on a vector representation of the 3D distribution of X-ray diffracting powers of the crystallites p (one scalar per voxel) and a 3D intragranular deformation field d (9 deformation components per voxel) of their crystal lattice, and aims to reconstruct this $\mathcal{F}(x, y, z)$: $\mathbb{R}^3 \mapsto \mathbb{R}^{1+9}$ tensor field on a grain-by-grain basis in a static deformed state, measured in a single scan. Mosaicity, i.e. a variation of crystal lattice orientations without associated stress, is allowed in the solution, hence the compatibility equations are not enforced. A diffraction spot is effectively treated as a projection of the 3D grain shape into the 3D detector domain (u, v, w), where the sub-grain deformations (local misorientations and elastic strains) determine the projection geometry that is unknown. ITF builds on the principles of kinematical diffraction and ray tracing, and the solution for the elastic deformation field is found by minimizing the differences between the measured pixel intensities of the detector q^{mes} (ordered as a column

vector) and the simulated intensities q in the diffraction spots of a grain in an iterative large scale optimization. This is a strongly non-linear and usually under-determined problem but in each iteration step a local linear problem at the latest solution can be formulated in matrix form:

$$\Delta q = q^{mes} - q = A_d \, \Delta d \tag{5.1}$$

where Δd (a column vector) contains the corrections sought to be applied to each elastic deformation component of each active grain voxel in the current iteration step, either positive or negative. The A_d matrix contains the intensity contributions to each pixel (one row per pixel) from each elastic deformation component correction of the active grain voxels. The A_d matrix is non-sparse, and it needs to be recomputed frequently due to the strong non-linearity of the model. Translated into a least-squares problem, the solver minimizes the objective function Γ which in this case is the square of the L_2 norm (denoted by $\|.\|_2$) of the residual vector:

$$\Gamma = \|\Delta q - A_d \Delta d\|_2^2 \tag{5.2}$$

$$\Delta d^* = \underset{\Delta d}{\text{argmin}}(\Gamma). \tag{5.3}$$

The elastic deformation corrections Δd^* found in the latest iteration step are applied to the current elastic deformation field of the grain. A 3D smoothing operation is then performed on the elastic deformation field in each step, which effectively acts as regularization in the optimization and results in a more realistic field.

The corrections to the diffracting powers, which determine the grain shapes, are recomputed in a separate step in a similar way but less frequently, using the latest elastic deformation field (i.e. the latest projection geometry).

Errors in the elastic deformation fields result from a number of sources in the acquisition and data processing. The strain sensitivity here is expected to be in the high single digits of 10^{-4}. The overall deformation sensitivity of the setup was limited by the angular step size, pixel size and detector distance. The relatively small number of diffraction spots meant the problem was more ill-posed. Some systematic deviations in intensity were caused by the simplifications in the model of the image formation process. Random noise from the photon counting statistics, fluorescence, and scattering processes was often significant. The grain-by-grain reconstruction approach is inherently prone to higher errors near grain boundaries. These aspects and potential future improvements of the acquisition and data processing, including

the handling of all grains simultaneously and ways to include the local mechanical equilibrium, are described in more detail in (Reischig and Ludwig, 2020). The fit quality in the ITF deformation solver can be monitored through the residual. In this study, the fitting is considered to have failed in several percent of the grains, mainly at the top and bottom of the gauge volume and at the free surface. Although the dataset does not constitute the state-of-art measurement capabilities, it is perfectly suitable to demonstrate the proposed stress filtering method.

5.3 Proposed method

Hodge decomposition

We use the following notation. Let $\Omega \subset \mathbb{R}^3$ denote a rectangular domain and $L^2_{per}(\Omega, \mathbb{R})$ be the set of all functions $u \in H^1_{loc}(\mathbb{R}^3, \mathbb{R})$ that are periodic with period Ω. The Helmholtz decomposition is a fundamental theorem in vector calculus that has applications in many fields including elasticity, incompressible flows and electromagnetism (see for example, (Cantarella, DeTurck, and Gluck, 2002)). Given any vector field $v \in L^2(\Omega, \mathbb{R}^3)$, there exists a scalar field $\varphi \in H^1_{per}(\Omega, \mathbb{R})$ and a vector field $w \in H^1_{per}(\Omega, \mathbb{R}^3)$, such that,

$$v = \nabla \times w + \nabla \varphi, \quad \nabla \cdot w = 0.$$

Further φ, w are unique up to a constant. The Hodge decomposition is a generalization to tensor fields, and we use a version due to Geymonat and Krasucki (2009) for symmetric fields. Let $\mathbb{M}^3$ denotes the linear space of all second-order matrices and $\mathbb{M}^3_{sym}$ the linear space of all second-order symmetric matrices. Given a symmetric matrix field $A \in L^2(\Omega, \mathbb{M}^3_{sym})$, there exist a vector field $y \in H^1_{per}(\Omega, \mathbb{R}^3)$, a symmetric matrix field $H \in H^2(\Omega, \mathbb{M}^3_{sym})$ and a constant matrix field c_0 such that,

$$A = \nabla \times H \times \nabla + \frac{1}{2}(\nabla y + \nabla y^T) + c_0, \quad \nabla \cdot H = 0$$

where $(\nabla \times H \times \nabla)_{kn} = \epsilon_{ijk}\epsilon_{lmn}\nabla_i\nabla_l H_{jm}$ in indicial notation. Further y, H up to a constant, and c_0 are unique.

Problem formulation

Given a symmetric matrix field $S_{exp} \in L^2(\Omega, \mathbb{M}^3_{sym})$, we seek to find the closest field $S_{equil} \in H_{div}(\Omega, \mathbb{M}^3_{sym})$ in L^2 norm, such that $\nabla \cdot S_{equil} = 0$, where

$$H_{div} = \{S \in L^2, div S \in L^2\}.$$

In other words,

$$S_{equil} = \text{argmin}\left\{||S - S_{exp}||_2 : S \in H_{div,per}(\Omega, \mathbb{M}^3_{sym}), \nabla \cdot S = 0\right\}.$$

It is easy to verify using the divergence theorem that $\langle(\nabla y + \nabla y^T), S\rangle_{H^1} = 0$ for any $y \in H^1_{per}(\Omega, \mathbb{R}^3)$ and $S \in H_{div,per}(\Omega, \mathbb{M})$ that satisfies $\nabla \cdot S = 0$. Since $\nabla \cdot S_{equil} = 0$, it follows that the Hodge decomposition of S_{equil} has $y = 0$, or

$$S_{equil} = \nabla \times H \times \nabla + c_0, \quad \nabla \cdot H = 0 \tag{5.4}$$

for some $H \in H^2_{per}(\Omega, \mathbb{M}^3_{sym})$ and $c_0 \in \mathbb{M}^3_{sym}$. It follows that

$$S_{exp} = S_{equil} + \frac{1}{2}(\nabla y + \nabla y^T) = \nabla \times H \times \nabla + c_0 + \frac{1}{2}(\nabla y + \nabla y^T). \tag{5.5}$$

Now, taking the average over Ω and again using the divergence theorem, $c_0 = \overline{S}_{exp}$ where we use the overhead bar to denote spatial average $\overline{A} = \frac{1}{|\Omega|}\int_\Omega A dX$. Putting all this together, the original minimization problem can then be written as

$$H = \text{argmin}\left\{\mathcal{L}(A) : A \in H^2_{per}(\Omega, \mathbb{M}^3_{sym}), \nabla \cdot A = 0\right\}$$
$$\text{where} \quad \mathcal{L}(A) := ||\nabla \times A \times \nabla - \widetilde{S}_{exp}||_2 \tag{5.6}$$

where tilde denotes the demeaned part $\widetilde{S}_{exp} = S_{exp} - \overline{S}_{exp}$. We use this variational problem to find S_{equil}.

The first variation

$$\delta\mathcal{L} = \int_\Omega (\nabla \times H \times \nabla - \widetilde{S}_{exp}) \cdot (\nabla \times \delta H \times \nabla) \, dX = 0 \quad \forall \, \delta H \in \mathcal{A}$$

where $\mathcal{A} = \{A \in H^2_{per}(\Omega, \mathbb{M}^3_{sym}) : \nabla \cdot A = 0\}$. Notice that for any $A, B \in H^1_{per}(\Omega, \mathbb{M}^3_{sym})$,

$$\int_\Omega (\nabla \times A) \cdot B - A \cdot (\nabla \times B) \, dX = \int_\Omega \nabla \cdot (A \times B) \, dX = 0$$

where the second equality follows from the divergence theorem. Thus,

$$\int_\Omega (\nabla \times A) \cdot B \, dX = \int_\Omega A \cdot (\nabla \times B) \, dX.$$

Applying this equality twice to the functional, we obtain

$$\delta\mathcal{L} = \int_\Omega (\nabla \times \nabla \times H \times \nabla \times \nabla - \nabla \times \widetilde{S}_{exp} \times \nabla) \cdot \delta H \, dX = 0 \quad \forall \, \delta H \in \mathcal{A}.$$

We obtain the Euler-Lagrange equation,

$$\nabla \times \nabla \times H \times \nabla \times \nabla = \nabla \times \widetilde{S}_{exp} \times \nabla.$$

Since $\nabla \times (\nabla \times H) = \nabla(\nabla \cdot H) - \nabla^2 H$ and $\nabla \cdot H = 0$, the Euler-Lagrange equation simplifies to

$$\nabla^4 H = \nabla \times \widetilde{S}_{exp} \times \nabla. \tag{5.7}$$

We solve this equation to find H and then S_{equil}. Recall that H is unique up to the constant, but the constant does not affect S_{equil}: so we choose the constant to make $\overline{H} = 0$.

We conclude the formulation with a few comments. First, while we refer to S_{equil} as a stress field, it does not mean it follows the underlying physical problem. Instead it is the self-equilibriated field (i.e., in H_{div}) that is closest in the sense of L^2 norm to a given unequilibriated field. Second, in this work, we use the L^2 norm. This has the advantage that material information is unnecessary and that it gives rise to a bi-harmonic operator in (5.7). We could, of course, use other norms. For example, we could use the energetic norm to include more material information see (Bellis and P. Suquet, 2019). This leads to an equation similar to (5.7) but with a different operator (perhaps with non-uniform coefficients in a polycrsytal). Finally, an alternative is to adapt an information theoretic point of view and use an entropic penalty, but this leads to a nonlinear problem.

Solution strategy

Given S_{exp}, we average over the volume to find c_0 and $\widetilde{S}_{exp}$. We then solve Eq. (5.7) to find H. Finally, we use Eq. (5.4) to obtain S_{equil}. The solution of Eq. (5.7) poses a challenge since it involves the fourth order biharmonic equation for $H \in H^2$, which requires continuous functions and derivatives. This in turn requires either higher order elements using finite-element discretization, or higher order differences in finite-difference methods. In either situation, it gives rise to stiff numerical problems (Argyris and Dunne, 1976; Gupta and Manohar, 1979). Another approach is to break it up into two harmonic problems and to solve them iteratively (Cheng, Han, and Huang, 2000). Here, we solve Eq. (5.7) using Fourier transforms. This poses two challenges. The first is that the experimental stress-field 5.4 is not periodic. To overcome this, a buffer is filled around the voxelized representation of the material to make the problem. The properties of the buffer near the lateral surfaces are chosen to give zero stress. This can introduce artifacts, especially near the boundaries. However, our numerical study presented in the appendix shows that while there is some difference at the boundaries, the role of the buffer is minimal in the interior. Second, the stress may suffer discontinuities across grain boundaries, interfaces and across the boundaries of the periodic unit cell. A

Fourier series does not converge to discontinuous functions in the L^∞ norm, and this in turn leads to ringing artifacts in practice. We address this in the next section.

5.4 Different operators in Fourier space

Let $\Omega = (0, L)^3$ be a cube and let us discretize it with a $N \times N \times N$ uniform cubic grid for N even. The corresponding domain in the Fourier space is $(-\pi/L, \pi/L)^3$ and again discretized uniformly with a $N \times N \times N$ grid. For any $f : \Omega \to \mathbb{R}^3$, the discrete Fourier transform (DFT) $\hat{f}$ satisfies

$$f(i, j, k) = \sum_{l=1}^{N} \sum_{m=1}^{N} \sum_{n=1}^{N} \hat{f}(l, m, n) exp(\frac{2\pi i}{L}((i-1)(l-1)+(j-1)(m-1)+(l-1)(k-1)),$$

$$(5.8)$$

$$\hat{f}(l, m, n) = \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{k=1}^{N} f(i, j, k) exp(\frac{-2\pi i}{L}((i-1)(l-1)+(j-1)(m-1)+(l-1)(k-1)),$$

$$(5.9)$$

where $i = \sqrt{-1}$ (distinct from index i).

Continuous differential operator

The continuous differential operator (CDO) of f in Fourier domain is obtained point-wise multiplication. For example, the partial derivative of f with respect to x is

$$\frac{\partial f(i, j, k)}{\partial x} = \sum_{l=1}^{N} \sum_{m=1}^{N} \sum_{n=1}^{N} i\xi_x \hat{f}(l, m, n) exp(\frac{2\pi i}{L}((i-1)(l-1)+(j-1)(m-1)+(l-1)(k-1))$$

$$(5.10)$$

where

$$\xi_x = \begin{cases} \frac{2\pi}{L}l, & \text{if } 1 \le l < \frac{N}{2} \\ 0, & \text{if } l = \frac{N}{2} \\ \frac{2\pi}{L}(l - N), & \text{if } \frac{N}{2} < l \le N. \end{cases} \qquad (5.11)$$

Higher order derivatives are a direct composition of first derivatives, with modification on the coefficients of the highest frequency term (see for example S. Johnson (2011)).

Discrete differential operator

We are interested in studying situations with grain boundaries and interfaces across which the stress may suffer discontinuities. A Fourier series does not converge to such a function in the L^∞ norm even as $N \to \infty$. In practice, we encounter ringing artifacts or spurious oscillations at finite N. A solution to this problem is to construct

the differential operator by taking finite differences and then using DFT (Berbenni, Taupin, Djaka, et al., 2014; Vidyasagar, Tan, and Kochmann, 2017). The first order central difference approximation reads,

$$\frac{\partial f(i, j, k)}{\partial x} = \frac{f(i+1, j, k) - f(i-1, j, k)}{2\Delta h} \tag{5.12}$$

etc., and therefore the discrete differential operator is given by Eq. (5.10) with the substitution

$$i\xi_x \leftarrow \frac{i}{\Delta h} \sin\left(\frac{2\pi(m-1)}{N}\right),$$

and analogously for the other partial derivatives. For example,

$$(\nabla \times H \times \nabla)_{kn} = \epsilon_{ijk}\epsilon_{lmn}\nabla_i\nabla_l H_{jm},$$

we substitute ∇_x with $\frac{i}{\Delta h} \sin\left(\frac{2\pi(m-1)}{N}\right)$ for the discrete approximation. Note that this first order operator approximates the continuous differential operator for small ξ_x, but then decays thereby acting as a low-pass filter. This resolves the overshooting and ringing at the interfaces. However, the convergence with respect to mesh size is reduced to second order from the exponential convergence of the CDO.

The biharmonic operator adopts the following 25–point stencil that is consistent,

$$\begin{aligned}
\nabla^4 f(i, j, k) = {} & 42f(i, j, k) - 12(f(i-2, j, k) + f(i+2, j, k) + f(i, j-2, k) + f(i, j+2, k) \\
& + f(i, j, k-2) + f(i, j, k+2)) + f(i-4, j, k) + f(i+4, j, k) + f(i, j-4, k) \\
& + f(i, j+4, k) + f(i, j, k-4) + f(i, j, k+4) + 2(f(i-2, j-2, k) \\
& + f(i-2, j+2, k) + f(i+2, j-2, k) + f(i+2, j+2, k) + f(i, j-2, k-2) \\
& + f(i, j-2, k+2) + f(i, j+2, k-2) + f(i, j+2, k+2) + f(i-2, j, k-2) \\
& + f(i-2, j, k+2) + f(i+2, j, k-2) + f(i+2, j, k+2)).
\end{aligned}$$

Notice that this is based on the second neighbor to make it consistent with the first order difference , Eq. (5.12). We obtain the corresponding operator in Fourier space from the usual Fourier space biharmonic operator with the substitution

$$\begin{aligned}
\xi_x^4 + \xi_y^4 + \xi_z^4 \leftarrow {} & (42 - 24(cos(4\pi(m-1)/N) + cos(4\pi(n-1)/N)) + cos(4\pi(l-1)/N)) \\
& + 2(cos(8\pi(m-1)/N) + cos(8\pi(n-1)/N) + cos(8\pi(l-1)/N)) \\
& + 8(cos(4\pi(m-1)/N)cos(4\pi(n-1)/N) + cos(4\pi(n-1)/N)cos(4\pi(l-1)/N) \\
& + cos(4\pi(m-1)/N)cos(4\pi(l-1)/N)))\frac{1}{16\Delta h^4}.
\end{aligned}$$
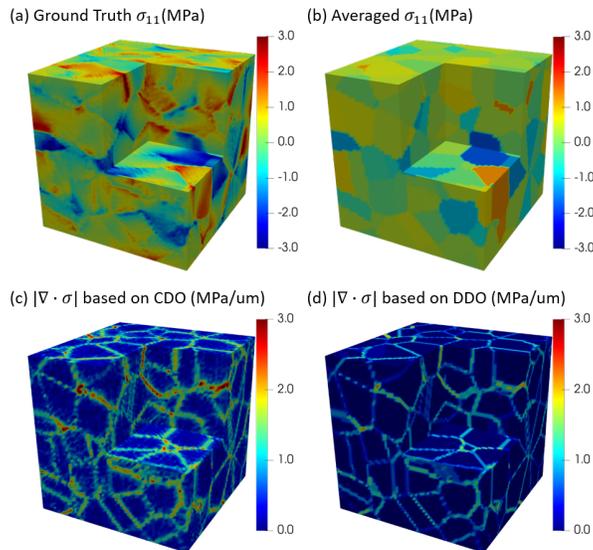
Figure 5.1: Synthetic data: (a) lateral stress field σ_{11} obtained with FFT-based model; (b) grain averaged σ_{11}; (c) divergence based on CDO (the divergence operator is using CDO approximation); (d) divergence based on DDO (the divergence operator is using DDO approximation).

5.5 Numerical examples

We study two examples, one using synthetic data with a known ground truth and another with experimental data.

Synthetic data with known ground truth

We begin with synthetic data from the results of a simulation using the FFT-based method (Lebensohn, Kanjarla, and Eisenlohr, 2012) to obtain an equilibrated stress field on a periodic unit cell representing a Cu polycrystal deformed in tension in the elastic regime. This polycrystalline unit cell is discretized by a 64^3 grid on a $[0, 64]\mu m^3$ box, and contains 100 Cu single crystal grains generated by periodic Voronoi tesselation. The Cu single crystal elastic constants reflect their cubic anisotropy, with $C_{11} = 168.4$ GPa, $C_{12} = 121.4$ GPa and $C_{44} = 75.4$ GPa, which determines heterogeneous, piece-wise constant elastic properties associated with the different crystal orientations, and, therefore, a non-uniform stress field, when the polycrystal is subjected to load in the elastic regime. For this analysis, the unit cell was loaded to a strain of 10^{-5} in uniaxial tension along axis x_3, corresponding to a longitudinal stress $\Sigma_{33} = 6.32$ MPa, and zero lateral stresses: $\Sigma_{11} = \Sigma_{22} = 0$, applied to the unit cell. Figure 5.1(a) shows the σ_{11} component of the local stress

	Diff with ground truth	Diff with average
CDO (L^2/L^∞)	0.0024/4.6041	0.0018/2.6945
DDO (L^2/L^∞)	0.0025/4.6007	0.0018/2.6222

Table 5.1: Error in the recovery from grain averaged data.

field (in MPa), which fluctuates with respect to the macroscopic value $\Sigma_{11} = 0$. Since this synthetic data was generated using a physically meaningful model that uses CDO to solve the governing equations of micro-mechanics (stress equilibrium and strain compatibility), the resulting stress field is indeed equilibrated (and thus divergence-free with respect to the CDO).

We start with this equilibrated stress field and obtain a *grain-averaged* stress field by averaging the stress over the grains – Figure 5.1(b) shows the σ_{11} component of the grain-averaged stress. This piece-wise constant field is not divergence-free, and the divergence is concentrated at the grain boundaries. We compute the divergence of the grain-averaged stress using both CDO and DDO–the results are shown in Figure 5.1(c) and Figure 5.1(d), respectively. We observe that the divergence is generally concentrated at the grain boundaries, but CDO smears it into the grain interior. This is a result of the spurious oscillation and overshooting that is characteristic of CDO on discontinuous functions.

We now study if it is possible to obtain the original stress field from the grain-averaged data using our method. We apply the proposed algorithm to the grain-averaged stress field using both CDO and DDO. The results are shown in Figure 5.2. We see that the recovery is imperfect, and this is confirmed in Table 5.1 which lists the L^2 and L^∞ norms of the error ($|f(x)|_{L^\infty} = sup\{f(x)\forall\ x\ \text{in}\ \Omega\}$). However, both CDO and DDO yield similar errors. In fact, Figure 5.2(g) shows the σ_{11} component of the stress along the line $0.5 \times 0.5 \times [0, 1]$). Both operators recover similar fields though the CDO has more oscillations within the grains and an overshoot at grain boundaries (see for example the grain boundary around $z = 0.9$).

To understand the error in recovery, recall that our algorithm only filters out the symmetrized gradient (curl-free) part and keeps all the divergence-free part of the given stress-field. So, if the difference between the given data and the ground truth deviate by both curl-free and divergence-free fields, our algorithm will filter out the symmetrized gradient portion of the deviation but retain the divergence-free part of that deviation. In our example, the difference between grain-averaged stress field and
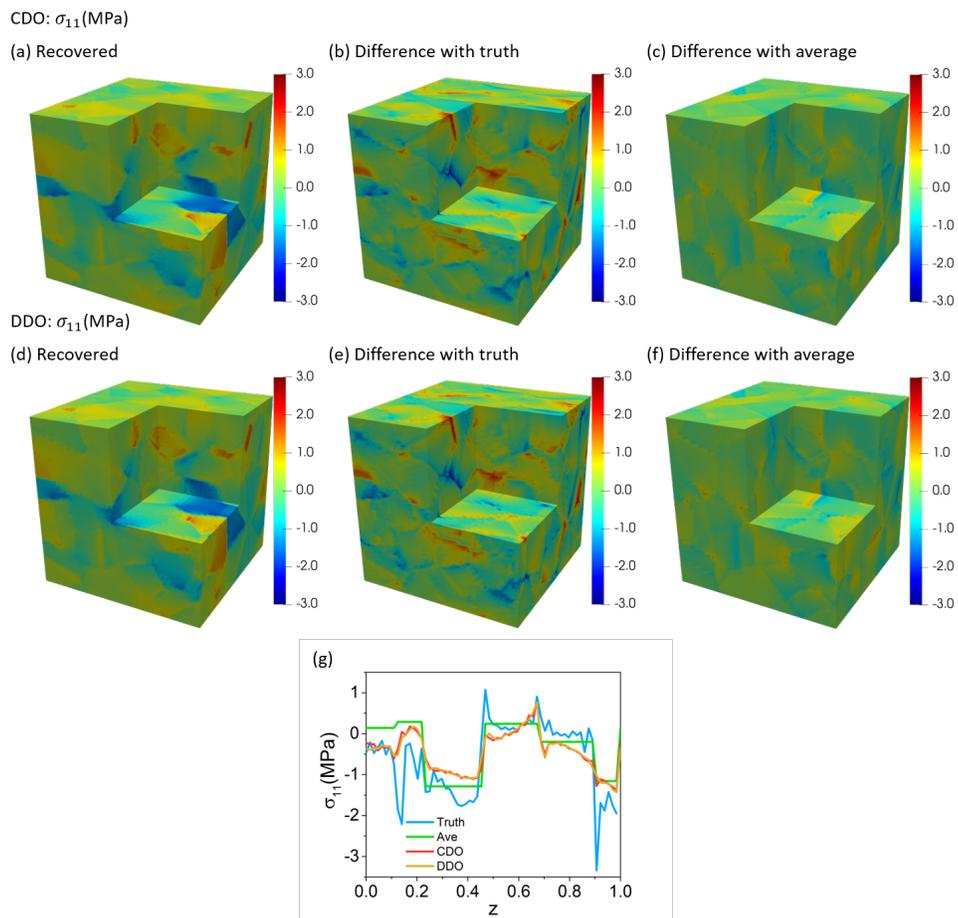
Figure 5.2: Recovery from grain averaged data, σ_{11}. (a-c) Results of recovery with CDO: (a) recovered stress field, (b) difference between ground truth an recovered stress field and (c) difference between grain averaged and recovered stress field. (d-f) Results of recovery with DDO: (a) recovered stress field, (b) difference between ground truth an recovered stress field and (c) difference between grain averaged and recovered stress field. (g) σ_{11} along the line $0.5 \times 0.5 \times [0, 1]$.
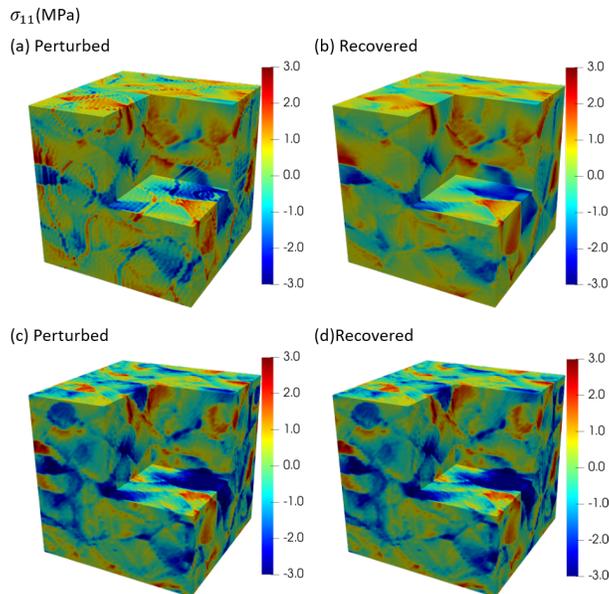
Figure 5.3: Recovery from synthetic data:(a,b) Recovery from a perturbation that is a symmetrized gradient (curl-free) (a) perturbed and (b) recovered fields. Compared to the ground truth in Figure 5.1(a), the recovery is very good – 9.016×10^{-6} in L^2 norm and 0.014 in L^{infty} norm. (c,d) Recovery from a perturbation that is divergence-free (c) perturbed and (d) recovered fields. The recovered field agrees with the perturbed field.

ground truth contains both curl-free and divergence-free components. Our algorithm filtered out the former but not the latter and this is the error in both Figure 5.2 and Table 5.1. We have verified this by taking the difference between grain-averaged stress field and ground truth, and applying our algorithm to it. The resulting residual equals (close to machine precision) to the recovery error (difference between the ground truth and recovered) in both CDO and DDO.

We further verify this in Figure 5.3. We perturb the ground truth by adding the symmetrized gradient $\nabla a + (\nabla a)^T$ of some vector field a – see Figure 5.3(a). For a comparison of the scale, $|\sigma_{11}|_{L^\infty}$ of the original field is 6.02 MPa. $|\sigma_{11}|_{L^\infty}$ of the perturbation is 1.00 MPa. Our algorithm is able to recover the ground truth from the perturbation – see Figure 5.3(b) with minimal error: 9.016×10^{-6} in L^2 norm and 0.014 in L^{infty} norm. On the other hand, when we perturb the ground truth by adding a divergence free field (Figure 5.3(c)), our algorithm returns the perturbed field and not the ground truth (Figure 5.3(d)). For a comparison of the scale, $|\sigma_{11}|_{L^\infty}$ of the original field is 6.02 MPa. $|\sigma_{11}|_{L^\infty}$ of the perturbation is 4.81 MPa. We note

that whether the perturbation remains does not depend on the scale of perturbation.

In conclusion, our approach finds the best projection to divergence-free fields. Therefore, it is able to filter out artifacts due to fields that are symmetrized gradients/curl-free, but unable to filter out artifacts that happen to be divergence-free. Indeed, it is not possible to filter out divergence-free artifacts without additional knowledge in terms of material behavior. For this reason, describing the average stress in each grain only provides limited information about the actual state of stress in the material.

Experimental Sample

In this section, we apply the proposed filtering methodology to the DCT experimental data on Gum Metal processed with the ITF method to obtain voxelized stress fields as described in Section 5.2. We focus on the experimental stress field from a sample subjected to an average normal stress of 34 MPa along the z-axis while the other two lateral surfaces are stress free. In order to use our algorithm that assumes periodicity, we introduce a buffer region around the specimen. The buffer region on the side (along x-axis and y-axis) has zero stress to accommodate the free surface. We also note that we implicitly make the top and bottom to be periodically connected. The buffer region at the bottom is an extrusion of the stress state of the sample's bottom. Each pixel of the buffering region has the same stress state as the pixel above it. This spurious connection between top and bottom in the longitudinal direction evidently deteriorates the stress balance near those regions. However, we should note that the experimental measurement is also less reliable in those locations and in grains on the free surface of the sample that have high mosaicity as a result of the spark cutting process. This treatment still preserves the information in the bulk of the sample. The resulting unit cell of $[0, 640\mu m]^3$, resolved at 256^3 voxels is shown in Figure 5.4(a,b) along with the σ_{11} and σ_{33} components of the experimental inferred stress. We note that the state of stress is quite heterogeneous. Figure 5.4(c) shows the magnitude of the divergence of the experimental inferred stress-field calculated using DDO, and we observe high values of divergence at the grain boundaries and triple junctions. This shows the the difficulty of accurate measurement of lattice strains and thus stresses where lattice structure is not preserved.

We present the filtering results with DDO (Fig. 5.5). After filtering, the peak value is reduced. Most of the profile patterns are preserved. Comparing the difference, we notice that most of the region undergoes minor modification. Significant differences

mostly occur around the grain boundaries. We also observe a significant correction of stress in the upper left grain (Fig. 5.5(a)), likely because the deformation solver failed for this grain and some other grains close to the surface. The main trends in Figure 5.5(g) are well preserved though and most of the smaller features within the grains are retained, even if the stress magnitudes changed. This is an encouraging result, and suggests that the elastic deformation solver based on diffraction constraints returns realistic trends in the stress fields. The difference in the actual magnitudes of this measurement are somewhat less concerning, as these were less reliable due to the deformation resolution of the experimental setup not being fully adequate. The deformation sensitivity of the setup will be improved in the future.

There are a number of outlier grains where the strain/stress state is extreme and significantly deviates from their neighbours. The ITF deformation fitting has potentially failed in these grains. A notable result of the filtering is that most of those grains are retained as outliers. Such high local deviations do not seem realistic, although their validity cannot be ruled out. Thus, it is not clear whether the filtered field is indeed more accurate in those regions. The measurement errors might be so high in those grains that the filtering cannot be expected to correct for them.

A detailed comparison of the two operators is conducted by looking closely at the stress distribution along line $320\mu m \times 320\mu m \times [0, 640]\mu m$) (Fig. 5.5)(g). Both methods yield similar results. In some grains, the stress is not corrected much. But in the grain near $z \approx 300\mu m$, there is a significant correction. The main difference between these two operators lie in grain boundaries. CDO and DDO both have some oscillations around the boundaries (e.g. $z \approx 250\mu m$). But CDO provides a higher estimation at the boundaries. In contrast, DDO has a smoother transition but it is underestimating the values at the interface.
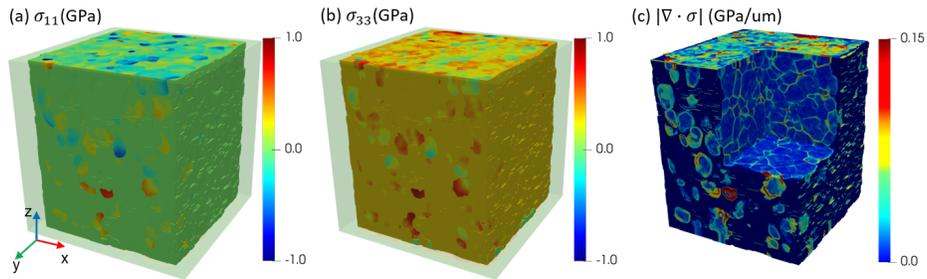
Figure 5.4: Stress field as inferred from DCT-ITF: (a) σ_{11} (transparent box is the computational domain); (b) σ_{33}. (c) Magnitude of the divergence based on DDO.
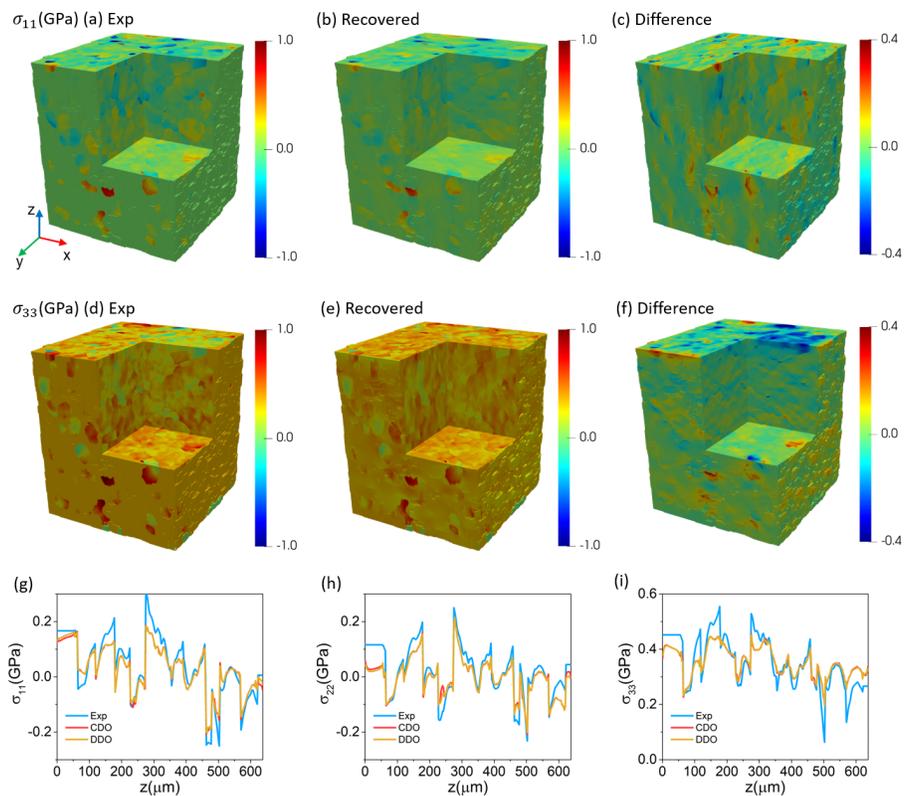
Figure 5.5: Stress distribution: (a-c) σ_{11}; (d-f) σ_{33}; (g-i) σ_{11}, σ_{22}, σ_{33} along the line $320\mu m \times 320\mu m \times [0, 640]\mu m$.

Chapter 6

CONCLUSION AND FUTURE DIRECTION

In this thesis, we present an approach to solving problems of computational micromechanics that is amenable to massively parallel calculations through the use of graphical processing units and other accelerators. The approach is based on splitting the solution operator in a manner that exploits the structure of continuum models that combine linear and universal physical laws (kinematic compatibility, balance laws), and nonlinear but local constitutive relations. We verify the approach against previous numerical simulations and study convergence and performance using finite elasticity. We also compared this method with other Fourier-based method. All algorithms have similar scaling, and our algorithm achieves comparable convergence rate with Fourier-Galerkin method. Importantly, our algorithm allows very large steps in nonlinear problems, and is able to stably handle bifurcations.

We applied our algorithm to study some unusual observations in liquid crystal elastomers. The materials incorporate rod-like nematic mesogens into the polymer chains of an elastomer. The combination of polymer elasticity and nematic ordering gives rise to exotic mechanical behavior. At high temperatures, the mesogens are randomly oriented and the LCE is an isotropic rubbery solid. However, on cooling, they undergo a phase transition where the steric interactions between the mesogens cause them to align in a particular direction. This makes it attractive for various applications. We focus on a highly unusual in-plane liquid-like behavior where sheets of these materials have zero shear stress even when the shear strain is nonzero. The results of our calculations agree with experimental observations, and provide insight into the mechanisms that give rise to this remarkable phenomenon.

We note that the method can be applied to a variety of problems. These include crystal plasticity, martensitic phase transformations, twinning, precipitation and Landau-Ginzburg models since these problems lead to equations of the form (2.1,2.2). Further, the implementation and examples presented here concern periodic boundary conditions which enabled the use of fast Fourier transforms to solve Poisson's equation. However, periodic boundary conditions are not inherent to this approach. The key issue is the solution of Poisson's equation, and there are a number of parallel iterative approaches that have been implemented with accelerators (Rees

et al., 2014). The problem of liquid crystal elastomers showed that we can incorporate point-wise constraints (incompressibility and prescribed norm on a vector internal variable) naturally in this method. It is possible to extend this approach to problems like fracture and contact where one has inequality constraints. Finally, one can extend this method to phenomena that include higher derivatives by introducing additional auxilliary variables.

An important open question is the convergence of the algorithm and error estimates. We have noted in Chapter 2 that there are partial results in the case of convexity. However, the general case where W is quasi-convex in F and convex in the internal variables remains open. Further, systematic analysis of the error remains a topic for the future.

The thesis also shows that the core idea behind the accelerated computational micromechanics may also be used to study multistable structures inspired by origami. The idea is to approach combines two kinematic descriptions of origami, vertex-based and facet-based and treat the compatibility condition as a constraint. This allows us to use the alternating directions and nudged elastic band to identify stable or rigid-foldable states, as well as transition paths between them. We demonstrate this with several examples.

Looking forward, there are two important directions. The first is inverse design, one of finding the folding pattern that would give rise to desired states and transition paths. The second is to combine such structures with actuation to autonomously morph from one state to another on the application of stimulus.

Finally, the thesis applied the idea that balance laws are projections to determining stress from high resolution micro xray diffraction tomography. experimental stress data. It uses Hodge decomposition to project a non-equilibriated stress field onto the divergence-free(equilibriated) subspace. This projection is numerically achieved using FFT-based solver, taking care to deal with discontinuity across grain boundaries. This method first verified using synthetic data from simulations. Then, it is applied to experimental data of a 3D X-ray diffraction experiment of a beta-Ti alloy, resulting in large corrections near grain boundaries.

This method is general, and it can be applied other experiment-derived stress field and serve as the post-processing step. Similarly analogous methods can be applied to enforce other constraints like kinematic compatibility constrain in micro-mechanical problems. Finally, we presented the method as a post-processing step, but it can

be combined with numerical methods that infer local stress fields from the x-ray diffraction patterns. This remains a task for the future.

BIBLIOGRAPHY

Anglin, B.S., R.A. Lebensohn, and A.D. Rollett (2014). "Validation of a numerical method based on Fast Fourier Transforms for heterogeneous thermoelastic materials by comparison with analytical solutions". In: *Computational Materials Science* 87, pp. 209–217.

Argyris, J.H. and P.C. Dunne (1976). "The finite element method applied to fluid mechanics". In: *Computational methods and problems in aeronautical fluid dynamics.(A77-12551 02-02) London and New York*, pp. 158–197.

Artemev, A., Y. Jin, and A. Khachaturyan (2001). "Three-dimensional phase field model of proper martensitic transformation". In: *Acta Materialia* 49.7, pp. 1165–1177.

Atkin, R.J. and N. Fox (1980). *An introduction to the theory of elasticity*. Longman.

Azoug, A. et al. (2016). "Viscoelasticity of the polydomain-monodomain transition in main-chain liquid crystal elastomers". In: *Polymer* 98, pp. 165–171.

Bellis, C. and P. Suquet (2019). "Geometric variational principles for computational homogenization". In: *Journal of Elasticity* 137.2, pp. 119–149.

Bensoussan, A., J.L. Lions, and G. Papanicolaou (2011). *Asymptotic analysis for periodic structures*. Vol. 374. American Mathematical Soc.

Berbenni, S., V. Taupin, K.S. Djaka, et al. (2014). "A numerical spectral approach for solving elasto-static field dislocation and g-disclination mechanics". In: *International Journal of Solids and Structures* 51, pp. 4157–4175.

Berbenni, S., V. Taupin, and R.A. Lebensohn (2020). "A fast Fourier transform-based mesoscale field dislocation mechanics study of grain size effects and reversible plasticity in polycrystals". In: *Journal of the Mechanics and Physics of Solids* 135, p. 103808.

Bernier, J.V. et al. (2011). "Far-field high-energy diffraction microscopy: a tool for intergranular orientation and strain analysis". In: *The Journal of Strain Analysis for Engineering Design* 46.7, pp. 527–547.

Bertin, N. and L. Capolungo (2018). "A FFT-based formulation for discrete dislocation dynamics in heterogeneous media". In: *Journal of Computational Physics* 355, pp. 366–384.

Bertoldi, K. et al. (2008). "Mechanics of deformation-triggered pattern transformations and superelastic behavior in periodic elastomeric structures". In: *Journal of the Mechanics and Physics of Solids* 56, pp. 2642–2668.

Bhattacharya, K. and P.M. Suquet (2005). "A model problem concerning recoverable strains of shape-memory polycrystals". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 461, pp. 2797–2816.

Biggins, J.S., E.M. Terentjev, and M. Warner (2008). "Semisoft response of nematic elastomers to complex deformations". In: *Physical Review E* 78, p. 041704.

Biggins, J.S., M. Warner, and K. Bhattacharya (2009). "Supersoft elasticity in polydomain nematic elastomers". In: *Physical Review Letters* 103.3, p. 037802.

– (2012). "Elasticity of polydomain liquid crystal elastomers". In: *Journal of the Mechanics and Physics of Solids* 60.4, pp. 573–590.

Bladon, P., E.M. Terentjev, and M. Warner (1993). "Transitions and instabilities in liquid crystal elastomers". In: *Physical Review E* 47, R3838.

Borst, R. de and E. Ramm (2011). *Multiscale Methods in Computational Mechanics*. Heidelberg: Springer.

Bourdin, B., G.A. Francfort, and J. Marigo (2008). "The variational approach to fracture". In: *Journal of Elasticity* 91, pp. 5–148.

Boyd, S., N. Parikh, and E. Chu (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc.

Bregman, L. M. (1966). "Relaxation method for finding a common point of convex sets and its application to optimization problems". In: *Doklady Akademii Nauk* 171, pp. 1019–1022.

Brömmel, F., D. Kramer, and H. Finkelmann (2012). "Preparation of liquid crystalline elastomers". In: *Liquid Crystal Elastomers: Materials and Applications*, pp. 1–48.

Cantarella, J., D. DeTurck, and H. Gluck (2002). "Vector calculus and the topology of domains in 3-space". In: *The American mathematical monthly* 109.5, pp. 409–442.

Cehula, J. and V. Prša (2020). "Computer modelling of origami-like structures made of light activated shape memory polymers". In: *International Journal of Engineering Science* 150, p. 103235.

Cesana, P., P. Plucinsky, and K. Bhattacharya (May 2015). "Effective behavior of nematic elastomer membranes". In: *Arch. Rat. Mech. Anal.* 218.2, pp. 863–905.

Chatterjee, K. et al. (2017). "Study of residual stresses in Ti-7Al using theory and experiments". In: *Journal of the Mechanics and Physics of Solids* 109, pp. 95–116.

Cheng, X., W. Han, and H. Huang (2000). "Some mixed finite element methods for biharmonic equation". In: *Journal of computational and applied mathematics* 126.1-2, pp. 91–109.

Conti, S., A. DeSimone, and G. Dolzmann (2002a). "Semisoft elasticity and director reorientation in stretched sheets of nematic elastomers". In: *Physical Review E* 66.6, p. 061710.

Conti, S., A. DeSimone, and G. Dolzmann (2002b). "Soft elastic response of stretched sheets of nematic elastomers: A numerical study". In: *Journal of the Mechanics and Physics of Solids* 50, pp. 1431–1451.

Conti, S., S. Müller, and M. Ortiz (2020). "Data-driven finite elasticity". In: *Archive for Rational Mechanics and Analysis* 237.1, pp. 1–33.

Demir, E. et al. (2013). "A computational framework for evaluating residual stress distributions from diffraction-based lattice strain data". In: *Computer Methods in Applied Mechanics and Engineering* 265, pp. 120–135.

DeSimone, A. and G. Dolzmann (2002). "Macroscopic response of nematic elastomers via relaxation of a Class of SO(3)-invariant energies". In: *Archive for Rational Mechanics and Analysis* 161.3, pp. 181–204.

Eckstein, J. and D.P. Bertsekas (1992). "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators". In: *Mathematical Programming* 55, pp. 293–318.

Eggersmann, R. et al. (2019). "Model-free data-driven inelasticity". In: *Computer Methods in Applied Mechanics and Engineering* 350, pp. 81–99.

Eghtesad, A. et al. (2018). "Spectral database constitutive representation within a spectral micromechanical solver for computationally efficient polycrystal plasticity modelling". In: *Computational Mechanics* 61, pp. 89–104.

Eyre, D.J. and G.W. Milton (1999). "A fast numerical scheme for computing the response of composites using grid refinement". In: *The European Physical Journal-Applied Physics* 6.1, pp. 41–47.

Filipov, E.T., K. Liu, et al. (2017). "Bar and hinge models for scalable analysis of origami". In: *International Journal of Solids and Structures* 124, pp. 26–45.

Filipov, E.T., T. Tachi, and G.H. Paulino (2015). "Toward optimization of stiffness and flexibility of rigid, flat-foldable origami structures". In: *The 6th International Meeting on Origami in Science, Mathematics and Education*, p. 121.

Fish, J.E. (2009). *Multiscale Methods: Bridging the Scales in Science and Engineering*. Oxford: Oxford University Press.

Gattas, J.M., W. Wu, and Z. You (2013). "Miura-base rigid origami: parameterizations of first-level derivative and piecewise geometries". In: *Journal of Mechanical design* 135.11.

Geymonat, G. and F. Krasucki (2009). "Hodge decomposition for symmetric matrix fields and the elasticity complex in Lipschitz domains". In: *Communications on Pure & Applied Analysis*, pp. 295–309.

Geymonat, G., S. Müller, and N. Triantafyllidis (1993). "Homogenization of nonlinearly elastic materials, microscopic bifurcation and macroscopic loss of rank-one convexity". In: *Archive for Rational Mechanics and Analysis* 122, pp. 231–290.

Giga, Y. and Y. Ueda (2020). "Numerical computations of split Bregman method for fourth order total variation flow". In: *Journal of Computational Physics* 405, pp. 109–114.

Glowinski, R. (2015). *Variational Methods for the Numerical Solution of Nonlinear Elliptic Problems*. Philadelphia, PA: Society for Industrial and Applied Mathematics.

– (2016). "ADMM and Non-convex Variational Problems". In: *Splitting Methods in Communication, Imaging, Science and Engineering*. Ed. by R. Glowinski, S.J. Osher, and W. Yin. Scientific Computation. Cham, Switzerland: Springer, pp. 251–299.

Glowinski, R. and P. Le Tallec (1982). "Numerical solution of problems in incompressible finite elasticity by augmented Lagrangian methods. I. Two-dimensional and axisymmetric problems". In: *SIAM Journal on Applied Mathematics* 42, pp. 400–429.

– (1984). "Numerical solution of problems in incompressible finite elasticity by augmented Lagrangian methods II. Three-dimensional problems". In: *SIAM Journal on Applied Mathematics* 44, pp. 710–733.

– (1989). *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*. Philadelphia: SIAM.

Goldstein, T. and S. Osher (2008). "The Split Bregman Method for L1-Regularized Problems". In: *SIAM Journal on Imaging Science* 2, pp. 323–343.

Golubović, L. and T.C. Lubensky (1989). "Nonlinear elasticity of amorphous solids". In: *Physical review letters* 63.10, p. 1082.

Gueninchault, N., H. Proudhon, and W. Ludwig (Nov. 2016). "Nanox: a miniature mechanical stress rig designed for near-field X-ray diffraction imaging techniques". In: *Journal of Synchrotron Radiation* 23.6, pp. 1474–1483.

Gupta, M.M. and R.P. Manohar (1979). "Direct solution of the biharmonic equation using noncoupled approach". In: *Journal of Computational Physics* 33.2, pp. 236–248.

Hanna, B.H. et al. (2014). "Waterbomb base: a symmetric single-vertex bistable origami mechanism". In: *Smart Materials and Structures* 23.9, p. 094009.

Hayashi, Y. et al. (2019). "Intragranular three-dimensional stress tensor fields in plastically deformed polycrystals." In: *Science (New York, N.Y.)* 366.6472, pp. 1492–1496.

He, B.S., H. Yang, and S.L. Wang (2000). "Alternating direction method with self-adaptive penalty parameters for monotone variational inequalities". In: *Journal of Optimization Theory and Applications* 106, pp. 337–356.

Henningsson, N.A. et al. (Apr. 2020). "Reconstructing intragranular strain fields in polycrystalline materials from scanning 3DXRD data". In: *Journal of Applied Crystallography* 53.2, pp. 314–325.

Herbol, H.C., J. Stevenson, and P. Clancy (2017). "Computational implementation of nudged elastic band, rigid rotation, and corresponding force optimization". In: *Journal of Chemical Theory and Computation* 13.7, pp. 3250–3259.

Hu, Y. and H. Liang (2020). "Folding simulation of rigid origami with Lagrange multiplier method". In: *arXiv preprint arXiv:2006.05025*.

Hull, T.C. (2002). "Modelling the folding of paper into three dimensions using affine transformations". In: *Linear Algebra and its applications* 348.1-3, pp. 273–282.

Jaramillo, G. and S. Venkataramani (2019). "A modified split Bregman algorithm for computing microstructures through Young measures". In: *arXiv preprint arXiv:1912.03360*.

Johnson, M. et al. (2017). "Fabricating biomedical origami: A state-of-the-art review". In: *International Journal of Computer Assisted Radiology and Surgery* 12.11, pp. 2023–2032.

Johnson, S.G. (2011). "Notes on FFT-based differentiation". In: *MIT Applied Mathematics*.

Karapiperis, K. et al. (2021). "Data-driven multiscale modeling in mechanics". In: *Journal of the Mechanics and Physics of Solids* 147, p. 104239.

Kirchdoerfer, T. and M. Ortiz (2016). "Data-driven computational mechanics". In: *Computer Methods in Applied Mechanics and Engineering* 304, pp. 81–101.

Kirk, D.B. and W.W. Hwu (2016). *Programming Massively Parallel Processors*. Morgan-Kaufman.

Knezevic, M. and D.J. Savage (2014). "A high-performance computational framework for fast crystal plasticity simulations". In: *Computational Materials Science* 83, pp. 101–106.

Knibbe, H., C.W. Oosterlee, and C. Vuik (2011). "GPU implementation of a Helmholtz Krylov solver preconditioned by a shifted Laplace multigrid method". In: *Journal of Computational and Applied Mathematics* 236.3, pp. 281–293.

Kothe, D., S. Lee, and I. Qualters (2019). "Exascale Computing in the United States". In: *Computing in Science and Engineering* 21, pp. 17–29.

Kuehnel, F.O. (2003). *On the minimization over SO (3) manifolds*. Tech. rep. Citeseer.

Kundler, I. and H. Finkelmann (1995). "Strain-induced director reorientation in nematic liquid single crystal elastomers". In: *Macromolecular Rapid Communications* 16.9, pp. 679–686.

Küpfer, J. and H. Finkelmann (1991). "Nematic liquid single crystal elastomers". In: *Die Makromolekulare Chemie, Rapid Communications* 12, pp. 717–726.

Lang, R.J. (2007). "The science of origami". In: *Physics World* 20.2, p. 30.

Lang, R.J., S. Magleby, and L. Howell (2016). "Single degree-of-freedom rigidly foldable cut origami flashers". In: *Journal of Mechanisms and Robotics* 8.3.

Lebensohn, R.A., A.K. Kanjarla, and P. Eisenlohr (2012). "An elasto-viscoplastic formulation based on fast Fourier transforms for the prediction of micromechanical fields in polycrystalline materials". In: *International Journal of Plasticity* 32, pp. 59–69.

Lebensohn, R.A. and A. Needleman (2016). "Numerical implementation of nonlocal polycrystal plasticity using fast Fourier transforms". In: *Journal of the Mechanics and Physics of Solids* 97, pp. 333–351.

Li, S.F. and R.M. Suter (2013). "Adaptive reconstruction method for three-dimensional orientation imaging". In: *Journal of Applied Crystallography* 46.2, pp. 512–524.

Li, Y. (2020). "Motion paths finding for multi-degree-of-freedom mechanisms". In: *International Journal of Mechanical Sciences*, p. 105709.

Lienert, U. et al. (2011). "High-energy diffraction microscopy at the advanced photon source". In: *Jom* 63.7, pp. 70–77.

Liu, B. et al. (2018). "Topological kinematics of origami metamaterials". In: *Nature Physics* 14.8, pp. 811–815.

Liu, K. and G.H. Paulino (2017). "Nonlinear mechanics of non-rigid origami: An efficient computational approach". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 473.2206, p. 20170348.

Lucarini, S and J Segurado (2019). "An algorithm for stress and mixed control in Galerkin-based FFT homogenization". In: *International Journal for Numerical Methods in Engineering* 119, pp. 797–805.

Ludwig, W. et al. (2009). "Three-dimensional grain mapping by x-ray diffraction contrast tomography and the use of Friedel pairs in diffraction data analysis". In: *Review of scientific instruments* 80.3, p. 033905.

Ma, J. and Z. You (2014). "Energy absorption of thin-walled square tubes with a prefolded origami pattern—part I: geometry and numerical simulation". In: *Journal of applied mechanics* 81.1.

Mbanga, B.L. et al. (2010). "Modeling elastic instabilities in nematic elastomers". In: *Physical Review E* 82.5, p. 051701.

McNelis, K.P., P.R. Dawson, and M.P. Miller (2013). "A two-scale methodology for determining the residual stresses in polycrystalline solids using high energy X-ray diffraction data". In: *Journal of the Mechanics and Physics of Solids* 61.2, pp. 428–449.

Michel, J.C., H. Moulinec, and P. Suquet (2000). "A computational method based on augmented Lagrangians and fast Fourier transforms for composites with high contrast". In: *CMES(Computer Modelling in Engineering & Sciences)* 1.2, pp. 79–88.

Mielke, A. (2006). "A Mathematical Framework for Generalized Standard Materials in the Rate-Independent Case". In: *Multifield Problems in Solid and Fluid Mechanics*. Ed. by R. Helmig, A. Mielke, and B.K. Wohlmuth. Springer, pp. 399–428.

Mielke, A., C. Ortner, and Y. Şengül (2014). "An approach to nonlinear viscoelasticity via metric gradient flows". In: *SIAM Journal on Mathematical Analysis* 46, pp. 1317–1347.

Mihaila, B., M. Knezevic, and A. Cardenas (2014). "Three orders of magnitude improved efficiency with high-performance spectral crystal plasticity on GPU platforms". In: *International Journal for Numerical Methods in Engineering* 97, pp. 785–798.

Milton, G.W. (2020). "Substitution of subspace collections with nonorthogonal subspaces to accelerate Fast Fourier Transform methods applied to conducting composites". arXiv Preprint: 2001.08289.

Mishra, N., J. Vondřejc, and J. Zeman (2016). "A comparative study on low-memory iterative solvers for FFT-based homogenization of periodic media". In: *Journal of Computational Physics* 321, pp. 151–168.

Monchiet, V. and G. Bonnet (2012). "A polarization-based FFT iterative scheme for computing the effective properties of elastic composites with arbitrary contrast". In: *International Journal for Numerical Methods in Engineering* 89, pp. 1419–1436.

Moulinec, H. and F. Silva (2014). "Comparison of three accelerated FFT-based schemes for computing the mechanical response of composite materials". In: *International Journal for Numerical Methods in Engineering* 97, pp. 960–985.

Moulinec, H. and P. Suquet (1994). "A fast numerical method for computing the linear and nonlinear mechanical properties of composites". In: *Comptes rendus de l'Académie des sciences. Série II. Mécanique, physique, chimie, astronomie.*

Moulinec, H., P. Suquet, and G.W. Milton (2018). "Convergence of iterative methods based on Neumann series for composite materials: Theory and practice". In: *International Journal for Numerical Methods in Engineering* 114, pp. 1103–1130.

Mullin, T. et al. (2007). "Pattern transformation triggered by deformation". In: *Physical Review Letters* 99, p. 084301.

Oddershede, J. et al. (2010). "Determining grain resolved stresses in polycrystalline materials using three-dimensional X-ray diffraction". In: *Journal of Applied Crystallography* 43.3, pp. 539–549.

Ortiz, M., A.M. Cuitiño, et al. (2001). "Mixed atomistic-continuum models of material behavior: The art of transcending atomistics and informing continua". In: *MRS Bulletin* 26.3, pp. 216–221.

Ortiz, M. and L. Stainier (1999). "The variational formulation of viscoplastic constitutive updates". In: *Computer Methods in Applied Mechanics and Engineering* 171, pp. 419–444.

Pagan, D.C. and A.J. Beaudoin (2019). "Utilizing a novel lattice orientation based stress characterization method to study stress fields of shear bands". In: *Journal of the Mechanics and Physics of Solids* 128, pp. 105–116.

Park, J.S., U. Lienert, et al. (2013). "Quantifying three-dimensional residual stress distributions using spatially-resolved diffraction measurements and finite element based data reduction". In: *Experimental Mechanics* 53.9, pp. 1491–1507.

Park, J.S., K. Yildizli, et al. (2018). "Non-destructive characterization of subsurface residual stress fields and correlation with microstructural conditions in a shot-peened Inconel component". In: *Experimental Mechanics* 58.9, pp. 1389–1406.

Pavliotis, G. and A. Stuart (2008). *Multiscale methods: Averaging and homogenization*. Springer Science & Business Media.

Phillips, R. (2001). *Crystals, defects and microstructures: Modeling across scales*. Cambridge University Press.

Pinson, M.B. et al. (2017). "Self-folding origami at any energy scale". In: *Nature Communications* 8.1, pp. 1–8.

Plucinsky, P. and K. Bhattacharya (2017). "Microstructure-enabled control of wrinkling in nematic elastomer sheets". In: *Journal of the Mechanics and Physics of Solids* 102, pp. 125–150.

Pokharel, R. and R.A. Lebensohn (2017). "Instantiation of crystal plasticity simulations for micromechanical modelling with direct input from microstructural data collected at light sources". In: *Scripta Materialia* 132, pp. 73–77.

Poulsen, H.F. (2004). *Three-dimensional X-ray diffraction microscopy: Mapping polycrystals and their dynamics*. Vol. 205. Springer Science & Business Media.

Poulsen, H.F. et al. (2001). "Three-dimensional maps of grain boundaries and the stress state of individual grains in polycrystals and powders". In: *Journal of Applied Crystallography* 34.6, pp. 751–756.

Preis, T. et al. (2009). "GPU accelerated Monte Carlo simulation of the 2D and 3D Ising model". In: *Journal of Computational Physics* 228, pp. 4468–4477.

Rees, W. van et al. (Jan. 2014). "High performance CPU/GPU multiresolution Poisson solver". In: *Advances in Parallel Computing* 25, pp. 481–490. DOI: 10.3233/978-1-61499-381-0-481.

Reis, P.M., F.L. Jiménez, and J. Marthelot (2015). "Transforming architectures inspired by origami". In: *Proceedings of the National Academy of Sciences* 112.40, pp. 12234–12235.

Reischig, P., A. King, et al. (2013). "Advances in X-ray diffraction contrast tomography: Flexibility in the setup geometry and application to multiphase materials". In: *Journal of Applied Crystallography* 46.2, pp. 297–311.

Reischig, P. and W. Ludwig (2019). "Three-dimensional sub-grain mapping of elastic strain state and orientations in the bulk of polycrystals". In: *40th Risoe International Symposium on Metal Microstructures in 2D, 3D and 4D*.

– (2020). "Three-dimensional reconstruction of intragranular strain and orientation in polycrystals by near-field X-ray diffraction". In: *Current Opinion in Solid State and Materials Science*, p. 100851.

Rice, J.R. (1971). "Inelastic constitutive relations for solids: An internal-variable theory and its application to metal plasticity". In: *Journal of the Mechanics and Physics of Solids* 19, pp. 433–455.

Rus, D. and M.T. Tolley (2018). "Design, fabrication and control of origami robots". In: *Nature Reviews Materials* 3.6, pp. 101–112.

Schenk, M. and S.D. Guest (2011). "Origami folding: A structural engineering approach". In: *Origami* 5, pp. 291–304.

Schenk, M., A.D. Viquerat, et al. (2014). "Review of inflatable booms for deployable space structures: packing and rigidization". In: *Journal of Spacecraft and Rockets* 51.3, pp. 762–778.

Schneider, M. (2020). "An FFT-based method for computing weighted minimal surfaces in microstructures with applications to the computational homogenization of brittle fracture". In: *International Journal for Numerical Methods in Engineering* 121, pp. 1367–1387.

Shen, Y.F., H. Liu, and R.M. Suter (2020). "Voxel-based strain tensors from near-field High Energy Diffraction Microscopy". In: *Current Opinion in Solid State and Materials Science* 24.4, p. 100852.

Silverberg, J.L. et al. (2014). "Using origami design principles to fold reprogrammable mechanical metamaterials". In: *science* 345.6197, pp. 647–650.

Śliwa, R.E. et al. (2017). "Metal forming of lightweight magnesium alloys for aviation applications". In: *Archives of Metallurgy and Materials* 62.

Song, G. and N.M. Amato (2004). "A motion-planning approach to folding: From paper craft to protein folding". In: *IEEE Transactions on Robotics and Automation* 20.1, pp. 60–71.

Suter, R.M. et al. (2006). "Forward modeling method for microstructure reconstruction using x-ray diffraction microscopy: Single-crystal verification". In: *Review of scientific instruments* 77.12, p. 123905.

Tachi, T. (2009a). "Generalization of rigid-foldable quadrilateral-mesh origami". In: *Journal of the International Association for Shell and Spatial Structures* 50.3, pp. 173–179.

– (2009b). "Simulation of rigid origami". In: *Origami* 4.08, pp. 175–187.

– (2012). "Design of infinitesimally and finitely flexible origami based on reciprocal figures". In: *Journal for Geometry and Graphics* 16.2, pp. 223–234.

Tokumoto, H. et al. (2021). "Probing the in-plane liquid-like behavior of liquid crystal elastomers". In: *Science Advances* 7.25, eabe9495.

Treloar, L.R.G. (1975). *The physics of rubber elasticity*. OUP Oxford.

Triantafyllidis, N., M.D. Nestorović, and M.W. Schraad (2006). "Failure surfaces for finitely strained two-phase periodic solids under general in-plane loading". In: *Journal of Applied Mechanics* 73, pp. 505–515.

Uchida, N. (1999). "Elastic effects in disordered nematic networks". In: *Physical Review E* 60.1, R13.

– (2000). "Soft and nonsoft structural transitions in disordered nematic networks". In: *Physical Review E* 62.4, p. 5119.

Ula, S.W. et al. (2018). "Liquid crystal elastomers: an introduction and review of emerging technologies". In: *Liquid Crystals Reviews* 6.1, pp. 78–107.

Urayama, K. (2017). "New aspects of nonlinear elasticity of polymer gels and elastomers revealed by stretching experiments in various geometries". In: *Polymer International* 66.2, pp. 195–206.

Urayama, K. et al. (2009). "Polydomain- monodomain transition of randomly disordered nematic elastomers with different cross-linking histories". In: *Macromolecules* 42.12, pp. 4084–4089.

Van Der Giessen, E. et al. (2020). "Roadmap on multiscale materials modeling". In: *Modelling and Simulation in Materials Science and Engineering* 28, p. 043001.

Vidyasagar, A., W.L. Tan, and D.M. Kochmann (2017). "Predicting the effective response of bulk polycrystalline ferroelectric ceramics via improved spectral phase field methods". In: *Journal of the Mechanics and Physics of Solids* 106, pp. 133–151.

Viganò, N. et al. (2016). "Three-dimensional full-field X-ray orientation microscopy". In: *Scientific reports* 6.1, pp. 1–9.

Vondřejc, J., J. Zeman, and I. Marek (2014). "An FFT-based Galerkin method for homogenization of periodic media". In: *Computers & Mathematics with Applications* 68.3, pp. 156–173.

Wang, K. et al. (2011). "Folding a patterned cylinder by rigid origami". In: *Origami* 5, pp. 265–276.

Warner, M. and E.M. Terentjev (2003). *Liquid crystal elastomers*. Oxford: Oxford University Press.

Washizu, K. (1968). *Variational methods in elasticity and plasticity*. Oxford, Pergamon Press, x, 349p.

Wei, Z.Y. et al. (2013). "Geometric mechanics of periodic pleated origami". In: *Physical review letters* 110.21, p. 215501.

White, T.J. and D.J. Broer (2015). "Programmable and adaptive mechanics with liquid crystal polymer networks and elastomers". In: *Nature Materials* 14.11, p. 1087.

Wojtowicz, P.J., P. Sheng, and E.B. Priestley (1975). *Introduction to liquid crystals*. Springer.

Wu, B. et al. (2013). "Complexity analysis and algorithm design for reorganizing data to minimize non-coalesced memory accesses on GPU". In: *ACM SIGPLAN Notices* 48, pp. 57–68.

Wu, W. and Z. You (2010). "Modelling rigid origami with quaternions and dual quaternions". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 466.2119, pp. 2155–2174.

Xi, Z. and J.M. Lien (2014). "Folding rigid origami with closure constraints". In: *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. Vol. 46377. American Society of Mechanical Engineers, V05BT08A052.

– (2015). "Plan folding motion for rigid self-folding machine via discrete domain sampling". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 2938–2943.

Yang, L., L. Robert, and S. Pelligrino (2020). In: *Unpublished work*.

Yin, W. et al. (2008). "Bregman iterative algorithms for L1-minimization with applications to compressed sensing". In: *SIAM Journal on Imaging Sciences* 1, pp. 143–168.

Zakirov, I.M. and K.A. Alekseev (2010). "Design of a wedge-shaped folded structure". In: *Journal of Machinery Manufacture and Reliability* 39.5, pp. 412–417.

Zhang, Q. et al. (2017). "Origami and kirigami inspired self-folding for programming three-dimensional shape shifting of polymer sheets with light". In: *Extreme Mechanics Letters* 11, pp. 111–120.

Zhang, Y. et al. (2019). "Continuum mechanical modeling of liquid crystal elastomers as dissipative ordered solids". In: *Journal of the Mechanics and Physics of Solids* 126, pp. 285–303.

Zhou, H. et al. (2022). "Imposing equilibrium on experimental 3-D stress fields using Hodge decomposition and FFT-based optimization". In: *Mechanics of Materials* 164, pp. 104–109.

Zohdi, T.I. and P. Wriggers (2005). *Introduction to computational micromechanics*. Springer Verlag.

ACCELERATED COMPUTATIONAL MICROMECHANICS: METHOD

Equilibrium condition We show that the dual feasibility ensures satisfaction of the equilibrium equation of mechanics. We begin with the case when the local problem (step 1) is solved exactly. Consider a smooth test function $\varphi : \Omega \to \mathbb{R}^3$ that vanishes on the boundary. Multiply (2.16) with φ, integrate over the domain, and use the divergence theorem to obtain

$$\int_\Omega \nabla\varphi \cdot \left(\rho(F^{n+1} - \nabla u^{n+1}) - \Lambda^n \right) dx = 0. \tag{A.1}$$

Now multiply (2.14) with $\nabla\varphi$ and integrate over the domain to obtain

$$\int_\Omega \nabla\varphi \cdot \left(W_F(F^{n+1}, \eta^{n+1}, x) - \Lambda^n + \rho(\nabla u^n - F^{n+1}) \right) dx = 0. \tag{A.2}$$

Subtract one from the other, and we obtain

$$\int_\Omega \nabla\varphi \cdot W_F(F^{n+1}, \eta^{n+1}, x) dx = \rho \int_\Omega \nabla\varphi \cdot (\nabla u^{n+1} - \nabla u^n) dx. \tag{A.3}$$

By the dual feasibility $(2.18)_2$ and the Cauchy-Schwarz inequality, the right hand side above goes to zero. Further, the left hand side converges to the weak form of the equilibrium equation since this holds for arbitrary φ.

When the local step is not exact, we rewrite (A.1)

$$\int_\Omega \nabla\varphi \cdot W_F(F^{n+1}, \eta^{n+1}, x) dx =$$
$$\rho \int_\Omega \nabla\varphi \cdot (\nabla u^{n+1} - \nabla u^n) dx + \int_\Omega \nabla\varphi \cdot \left(W_F(F^{n+1}, \eta^{n+1}, x) - \Lambda^n + \rho(\nabla u^n - F^{n+1}) \right) dx. \tag{A.4}$$

The first term on the right is bounded by dual feasibility $(2.18)_2$ and the Cauchy-Schwarz inequality as before, and the second term is bounded by the local error estimate and the Cauchy-Schwarz inequality. Thus the weak form of the equilibrium equation holds.
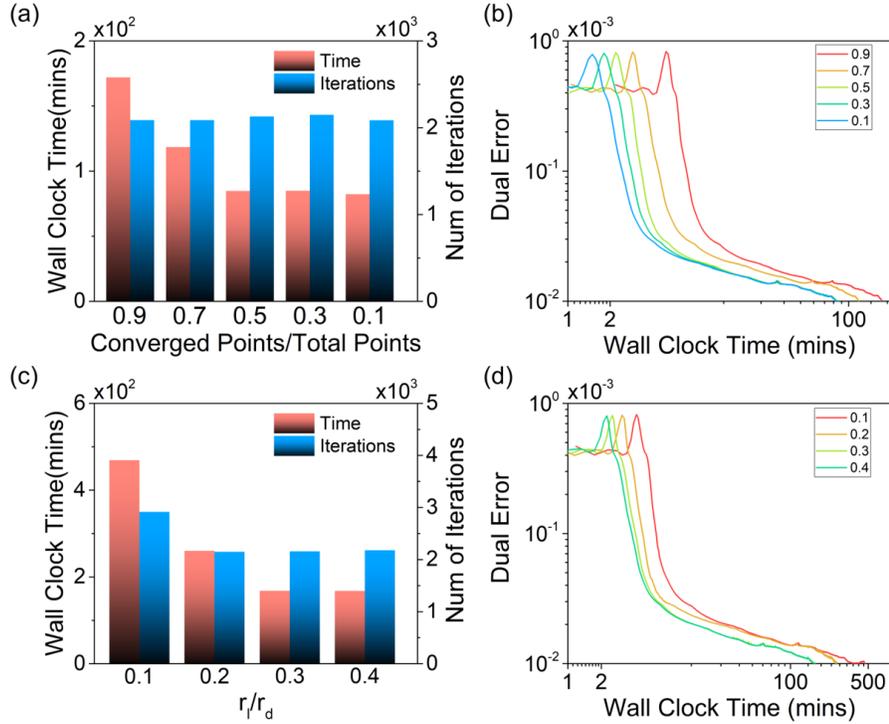
Figure A.1: Performance with an approximate solution of the local problem in the case of liquid elastomers. (a,b) Local convergence on a fixed fraction of spatial points: (a) Wall clock time and number of global iterations for global convergence for various fractions. (b) The global dual error versus wall clock time for various fractions. (c,d) Fixed ratio of local (r_l) to global dual (r_d) residual: (c) Wall clock time and number of global iterations for global convergence for various ratios and (d) The global dual error versus wall clock time for various ratios.

Convergence and performance in the case of LCEs   Figure A.1 shows the performance of the algorithm when we have an approximate solution of the local problem (step 1). As in the case of the elasticity problem discussed in Section 2.4, we find that the approximate solution provides savings in time without affecting the overall global convergence.

Figure A.2 shows the strong scaling in the case of liquid crystal elastomers. The slope is -0.80, which is in fact better than that observed in the case of elasticity. This is because the local step 1 which scales linearly takes a larger fraction of time compared to the case of elasticity. We have not performed the analysis of weak scaling since the specification of n_0 typically depends on the spatial resolution and therefore one-to-one comparison between simulations with different resolutions is not possible.
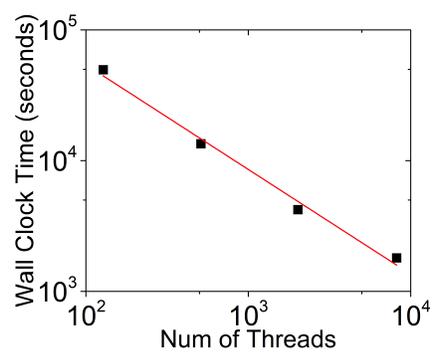
Figure A.2: Strong scaling in the case of liquid crystal elastomers.

A p p e n d i x B

EXPERIMENTAL CHARACTERIZATION OF THE STRESS FIELD

Effect of buffer on stress of the grains at surface   We note that to make the non-periodic cell periodic, a buffer zone is filled on the outside of grains. The buffer zone adds artifacts to the stress of grains at the surface. However, in practice, this artifact decays proportionally to $1/r^2$ and makes little difference on the inside grains. In section 5, we filled the buffer zone such that the cell is of size $256 \times 256 \times 256$. We perform another test by extending to buffer to make the cell of size $512 \times 512 \times 256$. Fig.B.1 shows the comparison of the two treatments in terms of the diagonal stress component σ_{11} as a function of the x-coordinate along a line of constant y- and z-coordinates going through the unit cell. The stress in the buffer zone and the outer grains slightly differs, but in the inner grains does not change much.
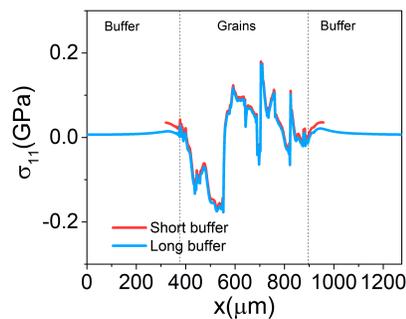
Figure B.1: Stress profile of σ_{11} along the line $(y, z) = 320\mu m \times 320\mu m$.