The fully nonlocal, finite-temperature, adaptive 3D quasicontinuum method for bridging across scales

Thesis by
Ishan Tembhekar

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

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2018
Defended December 14, 2017

Ishan Tembhekar
ORCID: 0000-0001-5123-1958

ACKNOWLEDGMENTS

First and foremost, I would like to start by thanking my advisor Dennis. It is obvious that this entire exercise would have been impossible without him, but I mean it in the most sincere way. During this entire process, I have been stuck more times than I can count, and the amount of patience he has shown is something I will forever be grateful for. His hands-on approach, open-door policy for meetings, and encouragement to ask even the most ridiculous of questions, are some of the many reasons why I consider him to be the best research mentor one can ever hope to get. My best day at Caltech was the day I decided to attend his class on Computational Mechanics in my first year, as that was when I realized my love for research using computers and simulations.

I would also like to sincerely thank my thesis committee members Prof Ravi, Prof Kaushik, and Prof Chiara. My interaction with Ravi and Kaushik started as a SURF student at Caltech, so in a way I have known them for even longer than Dennis. They have always managed to make me feel positive and calm whenever I interacted with them throughout my five years at Caltech. When I was unsure of what path to take on the research side, a talk with Ravi helped me make a decision and I am so glad to have had such a good mentor to guide me through.

Next, I would like to thank Jeff Amelang. I owe all my knowledge about careful programming practices and code maintenance to him, and I could not have wished for a better person to inherit a massive well-written code from. He made my transition to this difficult topic very smooth, and his continued guidance even after his graduation is one of the main reasons why I can write this thesis today.

I also want to mention the numerous collaborators who were instrumental in providing me their guidance and expertise in the field. Prof. Srikanth Patala and Arash Dehghan Banadaki, the grain boundaries would have never existed nor converged without your help. Arash, thank you for all the Skype sessions and advice on my other projects too. Mauricio Ponga and Gabriela Venturini, I cannot thank you enough for helping me step by step through in my work. I really appreciate you taking time off to read through my numerous lengthy emails and responding with the best way forward.

Additionally, I would like to thank all the members of the Kochmann group who I have had the pleasure of working with. I am extremely grateful for the feedback and

support that the group has provided me over the years. I have thoroughly enjoyed our weekly group lunches, barbeques, occasional hikes, and other fun activities. As I have been here for far too long, I am sure I will miss some, but I will take a shot naming a few group members I'd like to thank: Jeff Amelang, Stan Wojnar, Ray Chang, Gebriela Venturini, Alex Zelhofer, Neel Nadkarni, Vidyasagar Ananthan, Carlos Portela, Greg Phlipot, Wei Lin Tan, Abbas Tutcuoglu, Michael Frazier, Sid Kumar, Lukas Munk and Raphael Glaesesner, among many others.

I am deeply indebted to all the wonderful people at Caltech who made sure that I did not have to worry about anything other than research. Carol, Candace, and Stefanie, you helped me get settled in Pasadena during my first ever visit overseas. Laura and Daniel, I thank you for the guidance with all the queries I had as an international student. I never had to once worry about pending paperwork and travel signatures. Being one of the international orientation leaders was extremely fun. Next, I'd like to thank all the wonderful staff from MCE and GALCIT, especially Denise. Thank you for making sure that all my office needs were taken care of, for making sure I never had any pending paperwork, and for choosing the best lunches for group meetings.

Moving on to all the friends I formed at Caltech who were instrumental in making sure I retained my sanity throughout these years: to all my *desi* friends, you are the reason I never felt homesick and always had a little India to go hangout with. Subbu, Utkarsh, Srikanth, Siddharth, Prachi, GG ,and Kavya, I thoroughly enjoyed spending time with you. Thank you for taking care of me during my times of need, I'll forever be indebted to that. My roommates —Nick and Kevin, living with you has been the best experience of my life. I would have never guessed that an informal session of tennis would give rise to a lifelong friendship, but I am so glad that I got a chance to share an apartment with you guys. 1055 will forever be the best place to party. To the other "four-one" members —Dingyi, Reid, Sean —eating at restaurants is never going to be the same again. To my hiking/running buddies —Chengyun, Tess and Andrew —thank you so much for opening my eyes to the glorious beauty of the wilderness. Nick, Andrew, and Chengyun, going on those backpacking trips with you and especially hiking Mt. Whitney will be one of the crowning achievements of my California adventure. And of course, thank you to everyone else at Caltech who made this experience memorable.

Lastly, I would like to thank my family, especially my mom Sadhana Rajhans, my dad Deodatta Tembhekar, and my sister Saee Tembhekar. You have been my pillar

of support and my source of positivity every time I needed it. As cheesy as it sounds, everything I do, I do it for you.

ABSTRACT

Computational modeling of metallic materials across various length and time scales has been on the rise since the advent of efficient, fast computing machines. From atomistic methods like molecular statics and dynamics at the nanoscale to continuum mechanics modeled by finite element methods at the macroscale, various techniques have been established that describe and predict the mechanics of materials. Many recent technologies, however, fall into a gap between length scales (referred to as mesoscales), with microstructural features on the order of nanometers (thereby requiring full atomistic resolution) but large representative volumes on the order of micrometers (beyond the scope of molecular dynamics). There is an urgent need to predict material behavior using scale-bridging techniques that build up from the atomic level and reach larger length and time scales. To this end, there is extensive ongoing research in building hierarchical and concurrent scale-bridging techniques to master the gap between atomistics and the continuum, but robust, adaptive schemes with finite-temperature modeling at realistic length and time scales are still missing.

In this thesis, we use the quasicontinuum (QC) method, a concurrent scale-bridging technique that extends atomistic accuracy to significantly larger length scales by reducing the full atomic ensemble to a small set of representative atoms, and using interpolation to recover the motion of all lattice sites where full atomistic resolution is not necessary. We develop automatic model adaptivity by adding mesh refinement and adaptive neighborhood updates to the new fully nonlocal energy-based 3D QC framework, which allows for automatic resolution to full atomistics around regions of interest such as nanovoids and moving lattice defects. By comparison to molecular dynamics (MD), we show that these additions allow for a successful and computationally efficient coarse-graining of atomistic ensembles while maintaining the same atomistic accuracy.

We further extend the fully nonlocal QC formulation to finite temperature (termed hotQC) using the principle of maximum entropy in statistical mechanics and averaging the thermal motion of atoms to obtain a temperature-dependent free energy using numerical quadrature. This hotQC formulation implements recently developed optimal summation rules and successfully captures temperature-dependent elastic constants and thermal expansion. We report for the first time the influence of temperature on force artifacts and conclude that our novel finite-temperature adap-

tive nonlocal QC shows minimal force artifacts and outperforms existing formulations. We also highlight the influence of quadrature in phase space on simulation outcomes.

We study 3D grain boundaries in the nonlocal hotQC framework (previously limited to single-crystals) by modeling coarse-grained symmetric-tilt grain boundaries in coincidence site lattice (CSL) based bicrystals. We predict relaxed energy states of various Σ-boundaries with reasonable accuracy by comparing grain boundary energies to MD simulations and outline a framework to model polycrystalline materials that surpasses both spatial and temporal limitations of traditional MD.

PUBLISHED CONTENT AND CONTRIBUTIONS

[1] I. Tembhekar et al. “Automatic adaptivity in the fully nonlocal quasicontinuum method for coarse-grained atomistic simulations”. In: *International Journal for Numerical Methods in Engineering* 110.9 (2017).
Contributions: Ishan Tembhekar performed all the numerical experiments and simulations and participated in the writing of the manuscript., pp. 878–900. ISSN: 1097-0207. DOI: 10.1002/nme.5438. URL: http://dx.doi.org/10.1002/nme.5438.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

Number *Page*

Chapter 1

INTRODUCTION

When it comes to the behavior of materials and structures, there is always more than meets the eye. All the macrostructural phenomena that we observe in our everyday lives stem from various mechanisms occurring across multiple length and time scales. There are examples of structural hierarchy of scales all around us, as seen in both natural and artificial structures, an example of each is shown in Figure 1.1. The Eiffel tower is about 324m in height and 135m in width, but consists of thousands of individual truss links that have dimensions on the order of centimeters, spanning two orders of magnitude in length. The bamboo structure, at about 5mm thick, is made up of various structural links that span down to individual cellulose molecules that are about 3nm in size. Any structural response observed at the largest length scale in these structures depends on how each of the structural links at the lower scales behave in response to external events. Thus, understanding, predicting, and ultimately designing macroscopic behavior requires looking deep into the structure across many scales.

Similarly, there is a hierarchy in material response. In simple terms, the behavior of a material can be categorized differently when observed across various length and time scales. If we consider metallic materials, we can broadly classify the hierarchy as follows.

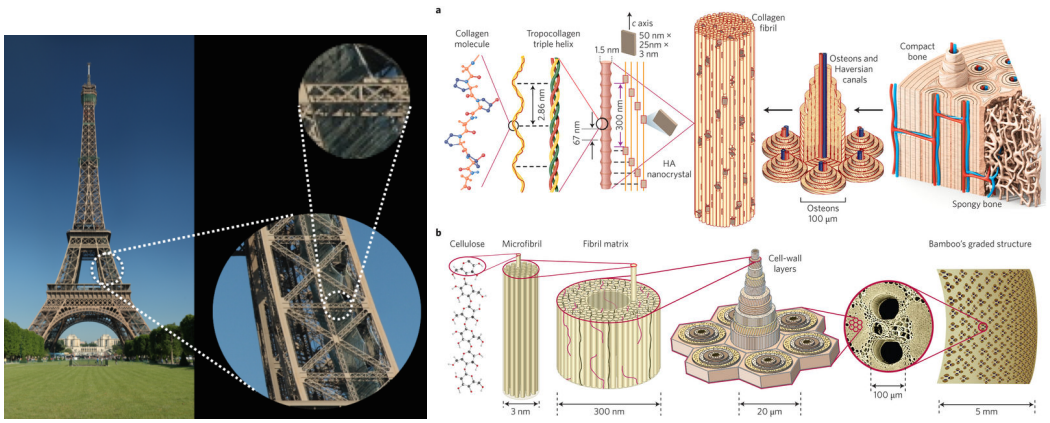

Figure 1.1: Examples of structural hierarchy in man-made structures (Eiffel tower on the left) and in nature [126] (bone and bamboo structure on the right).

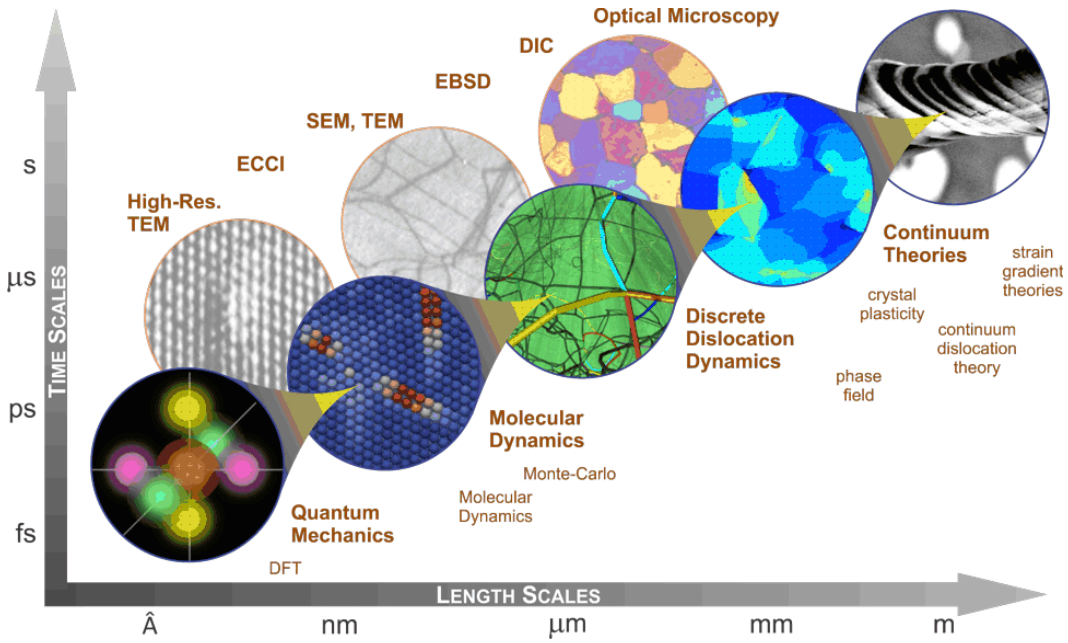

Figure 1.2: A broad categorization of different kinds of material behavior and the computational and experimental techniques used to simulate and observe the material for a variety of length and time scales. From left to right, abbreviations stand for Transmission Electron Microscopy (TEM), Density Functional Theory (DFT), Electron Channeling Contrast Imaging (ECCI), Scanning Electron Microscopy (SEM), Electron Back-Scatter Diffraction (EBSD), Digital Image Correlation (DIC).

The atomic scale: At the atomic scale, metals are composed of individual atoms that arrange into crystals defined by a unit cell (e.g. bcc, fcc, hcp). Material behavior is at the length scale of $\mathring{A}$ and atomic phenomenon, e.g. thermal vibrations, take place on orders of femtoseconds.

The microscopic scale: At this scale, the perfect crystalline nature of metals is broken up by lattice defects, giving rise to grain boundaries and polycrystals. The interaction of these defects with each other and with the grain boundaries significantly defines the material behavior at the microscopic scale, and subsequently macroscopic inelastic behavior at the macroscopic scales. Typical grain sizes range from a few nanometers up to micrometers, and the duration of plastic deformation at the atomic scale typically spans nanoseconds to a few microseconds.

The macroscopic scale: Here, the metallic material is treated as a continuous medium and the physical behavior is modeled by continuum theories. We no longer identify the crystal structure and individual grains of the metal and instead observe

and model the continuous macroscopic behavior. Systems having sizes above millimeters fall under this category, and the time scale used to observe various phenomena in continuous media starts around milliseconds.

Figure 1.2 shows the different categories of material behavior across the scales, including the commonly used experimental and computational techniques used to observe, simulate, and predict this behavior.

1.1 Need for multiscale modeling

Recently, the area of nanotechnology and nanoelectromechanical systems (NEMS) has progressed rapidly, resulting in manufacturing capabilities that allow for efforts to design structures and materials spanning multiple length scales. Figure 1.3 shows a system of nanotrusses spanning length scales from a few nanometers (thickness of individual truss members) to micrometers (outer dimensions of the entire truss network). Such structures fall into a gap between length scales (referred to as mesoscales) and simply picking one characterization of material behavior from the ones described above is not sufficient. There is an urgent need of predicting material behavior of such multiscale systems and requires in-situ experimental observation or scale-bridging computational techniques. Experimental studies, however, are often limited by resolution in both space and time as observations are generally carried out in time scales of seconds and in-situ observation of nanoscale mechanisms is usually a rare find. We therefore seek a computational technique that builds up from the atomic level, and reaches larger length and time scales.

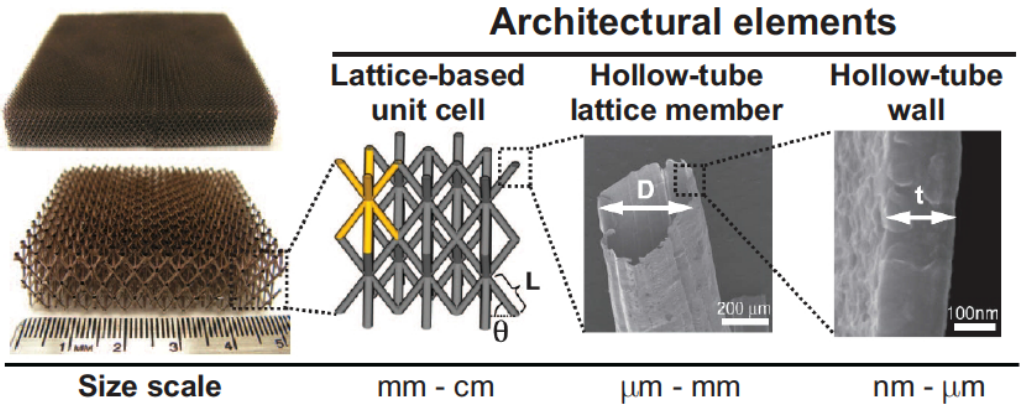

Figure 1.3: Hierarchical structures at the nanometer length scale: the structure is on the order of micrometers while the wall thickness of an individual truss link is on the order of a few nanometers[41].

1.2 Computational Modeling

1.2.1 Evolution of atomistics

Although the real beginning of atomistics can be traced back to the hard-sphere models of the 1950s, one can consider Richard Feynman's evaluation of the quantum mechanical forces on molecules as the beginning of the justification of the use of classical mechanics and electrostatics to approximate systems at a molecular state (using interatomic potentials, with classical equations of motion [37]). Since then, equation of state calculations using modified Monte Carlo (MC) integration using computers was proposed by Metropolis et al. [70], who studied interactions of a few hundreds of 2D spheres. The first known use of the term "Molecular Dynamics" (MD) was by Alder and Wainwright [2], who studied the interactions of hard spheres using many-body interactions. However, the first major advancement of using an interatomic potential to model atomic interactions was in 1964, when Rahman used the Lennard-Jones potential [79] to study a system of 864 liquid argon particles [90], and Stillinger followed with a study on liquid water [107]. Most of these approaches involved periodic boundaries, and it was Ryckaert, Ciccotti, and Berendsen [99] who first introduced a molecular dynamics system with constraints. Since then, the field has evolved exponentially with notable advancements being that of MD at constant temperature and/or pressure by Andersen [7] and unified MD and Density-Functional Theory(DFT) approaches by Car and Parrinello [19].

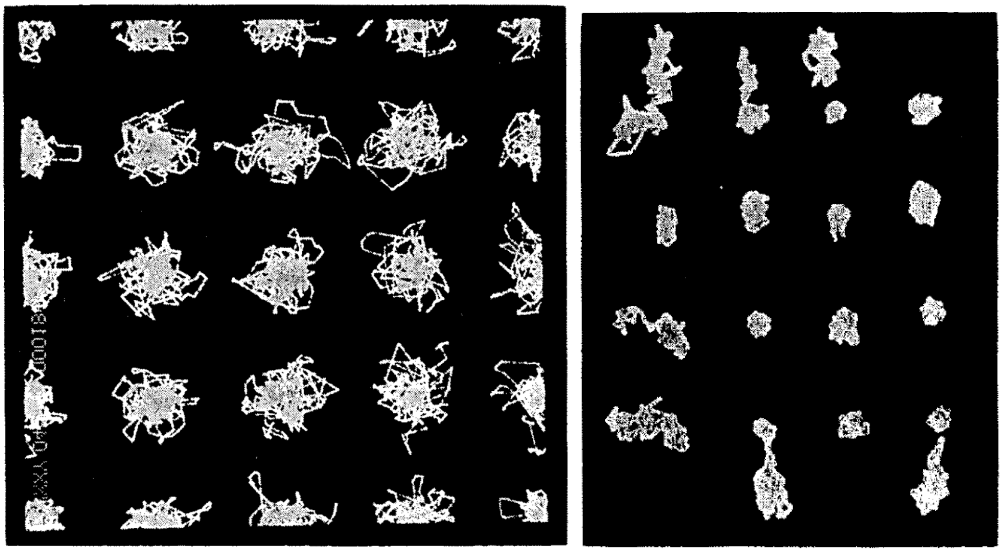

Figure 1.4: Time evolution of the positions of 32 particles using the hard-sphere model. On the left, using periodic boundary conditions. On the right, using free boundaries [2].

The atomic interactions used in molecular dynamics are nowadays modeled for a wide variety of materials, including but not limited to gasses, metals, metallic alloys, ceramics, organic substances, etc., and fast computing has increased the ability to model from trillions of atoms today. Atomistic simulations have paved the way for an unprecedented understanding of molecular and material behavior at the atomic scale [100]. LAMMPS [101] is one of the most widely used resources to run efficient MD simulations.

1.2.2 Limitations

One of the major limitations of MD methods is their inherent requirement to model the degrees of freedom of each and every atom in the system. Computational resources have come a long way since the inception of MD in the 1950s and the community has upgraded from 32 hard spheres in 1959 to billions, and more recently, a trillion atoms [40]. To model a magnitude of atoms, Germann and Kadau [40] used a simple Lennard-Jones (LJ) potential [79], which only models gases, and stored values to a single precision. However, to model a metallic cube of side $1\mu m$ using a fcc lattice, one needs to consider approximately 10^{20} atoms and use atomic interactions that are more complex than the simple pairwise interactions modeled by the LJ potential. This makes it impossible to use MD to model structures that span multiple length scales like those in Fig. 1.3. Thus, there is an urgent need for scale-bridging techniques that use accuracy of the atomistic simulations but scale up in length and time, by developing with efficient, smart models paired with a smart use of computational resources.

1.2.3 Review of concurrent scale-bridging models

There have been various approaches in the past and newer methods continue to be proposed that attempt to bridge the gap of computational modeling across various length scales. To study structures traversing multiple scales such as the ones described in Fig. 1.3, one needs the fully-atomistic resolution at the nanometer scale, while still being able to traverse to the actual length scales of these structures (μm). We briefly describe prominent methods used in computational solid mechanics to bridge this gap and state their salient features to provide a brief introduction to the state of the art in multiscale modeling of materials. Scale-bridging methodologies can be broadly classified into two types : *hierarchical* and *concurrent scale-bridging* methods. Hierarchical schemes refer to a clear separation of scales in which representative simulations at lower scales are used to extract effective consti-

tutive response at the macroscale. Some examples of these techniques are multiple-level finite-element analysis (also known as FE^n) [36, 71] and homogenization of atomistic ensembles used for material-point calculations in macroscale finite element simulations [24, 25]. Concurrent scale-bridging methods, on the other hand, integrate different constitutive descriptions into a single model by spatially separating domains and a handshake region for these descriptions to interact with each other. For this review and for the remainder of this thesis, we will focus on concurrent scale-bridging techniques.

This review follows the work of R. E. Miller and E. B. Tadmor [89] and Wernik and Meguid [128] with recent additions in scale-bridging methods included for an up-to-date state of the art. R. E. Miller and E. B. Tadmor [89] define a "multiscale model" as any model that unambiguously divides the system under consideration into two regions, the atomistic region and the continuum region. One of the main distinctions between these methods is that some are "energy-based" while others are "force-based". The energy-based methods develop a well-defined energy functional and minimize this energy as a solution to the problem. The force-based methods develop a set of physically meaningful forces on every degree of freedom, and drive these forces to zero as a solution to the problem.

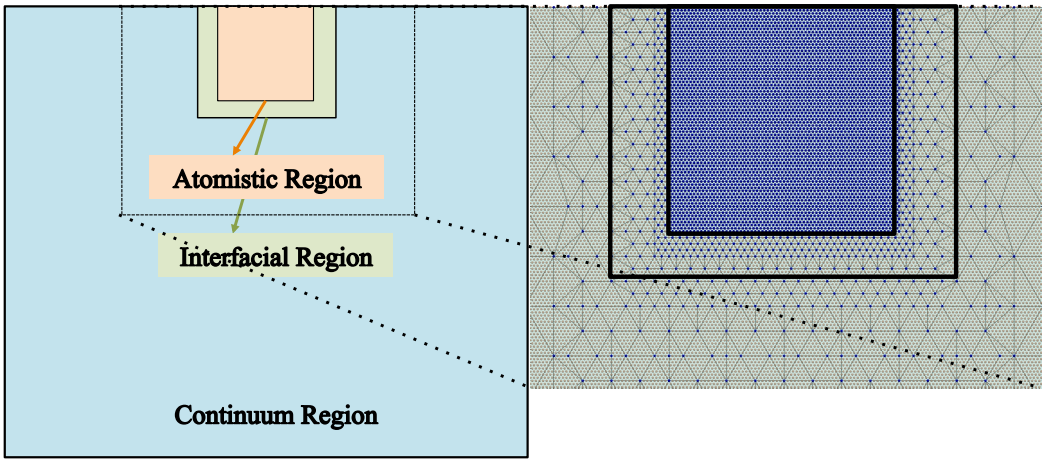

Figure 1.5: The general partitions of a concurrent multiscale model. The interface between the atomistic and continuum regions can vary based on different models. The general design is such that a bulk of the material is modeled using a continuum, whereas atomistics is localized to the relevant parts of the domain, thereby saving computational cost. The right pane shows the transition from full atomistics to coarse-grained region for a nonlocal quasicontinuum method.

The Coupling of Length Scales (CLS) method

The CLS method [97], [98] has three concurrent regimes ,as it also includes a coupling between atomistics and quantum mechanics in addition to the atomistic-continuum coupling. This is an energy-based method, with the continuum region modeled using finite elements with a continuum material model. This necessitates a form of "handshaking" between the two regions, which takes place in the interfacial region. This is done by writing the energy functional as a sum of two components, the atomistic part and the continuum part and the weights of the interfacial region are suitably modified to account for double-counting.

The Bridging Domain (BD) method

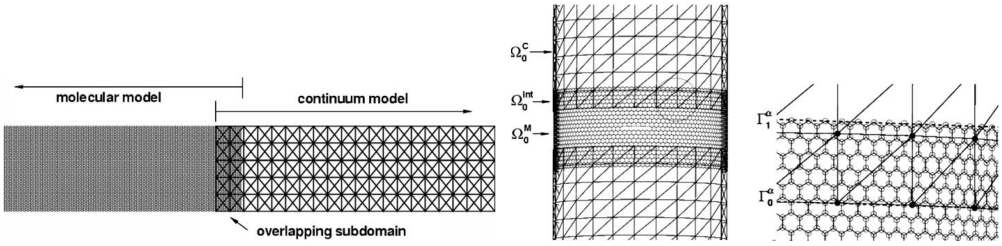

Figure 1.6: On the left: A 2D sample showing the coupling of various domains in the BD method. On the right: A domain model for a nanotube using the BD method. Atoms are represented by white spheres, solid lines represent the finite elements, and the black spheres represent the nodes. The atomistic, interfacial, and continuum regions are marked by $\Omega_0^M, \Omega_0^{int}$, and Ω_0^C respectively[130].

Xiao and Belytschko [130] introduced the BD method with the idea of an *overlap* of the atomistic and continuum regions in the interfacial region. This is an energy-based method which handles the energy in the overlapping region by taking a linear combination of the continuum and molecular Hamiltonians, and enforces compatibility of the two regions by using Lagrange multipliers. The main difference of this Lagrange multiplier approach from the CLS method is that the strong compatibility needed at the interface is reduced to weak compatibility, which makes handling of mesh generation at the interface easier but reduces the accuracy. Similar approaches have also been discussed within the quasicontinuum (QC) framework (see e.g. Blended-QC [63]).

The Bridging Scale Method (BSM)

Like the CLS and BD methods, the BSM, first introduced in [123], models the continuum region using finite elements. However, unlike the others, the finite element region extends to the *entire domain*, and there is no interfacial region. Instead, the method posits that atoms exist throughout the entire domain, but they are only explicitly tracked in the atomistic region, while they exist in the continuum region as the manifestation of heat. This requires the subtraction of the energy contribution by finite elements in the atomistic region to avoid double-counting, which is often accomplished by introducing "impedance forces" [128]. A major advantage of this method is that the atomistic and Finite Element (FE) equations of motion can be modeled at different time steps, as the continuum method exists everywhere.

Finite Element-Atomistic method (FEAt)

One of the earliest coupling methods is the FEAt which, as suggested, uses a FE description to model the continuum region, and enforces strong compatibility at the boundary of the atomistic and continuum regions, like the CLS method. However, the main difference lies in the formulation, which is force-based. FEAt uses a non-local continuum formulation in the interfacial region to provide the coupling between the atomistic and continuum domains. It was first introduced in [56].

Coupled Atomistic and Discrete Dislocations (CADD)

The CADD, introduced by Shilkrot, Miller, and Curtin [106], is similar to FEAt in that the continuum region modeled by FE methods is coupled to an atomistic region, modeled by MD. Like FEAt, it is a force-based model with no well defined energy functional. In addition, the model has the ability to accommodate discrete dislocations as they move in from the atomistic to the continuum region, where they are modeled using discrete dislocation (DD) techniques [58]. This model has strong compatibility, which means there is a displacement boundary condition on the finite element nodes at the interface of the atomistic and continuum region. One of the main limitations is that CADD has been formulated so far for 2D problems, and 3D extensions are a current area of research (see e.g. [64]).

Atomistic-to-Continuum (AtC) method

The AtC method [38] is essentially a force-based counterpart of the BD method. Like the BD method, the energy in the interfacial region is considered as a linear

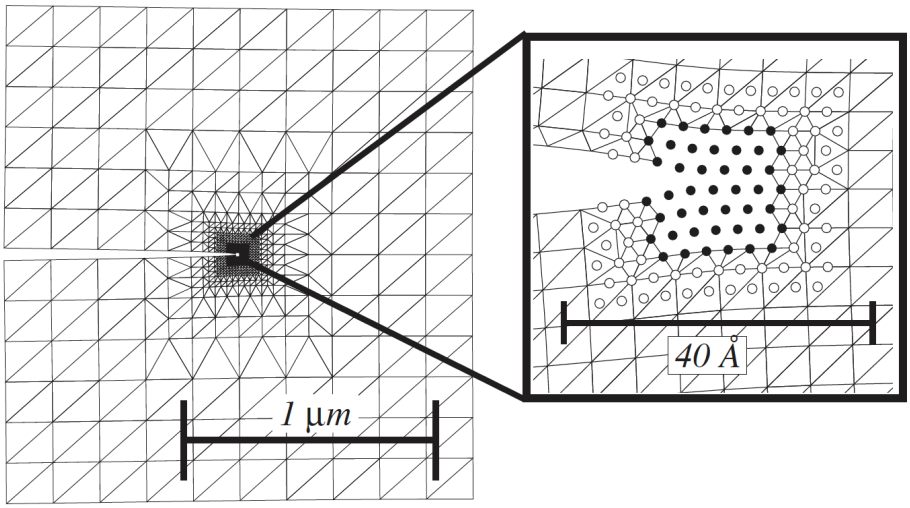

Figure 1.7: A model used to study model crack growth showing the various regions of a CADD model. The atomistic region is shown using the black atoms and the white atoms form the interfacial region that pads the atomistic and continuum regions[106].

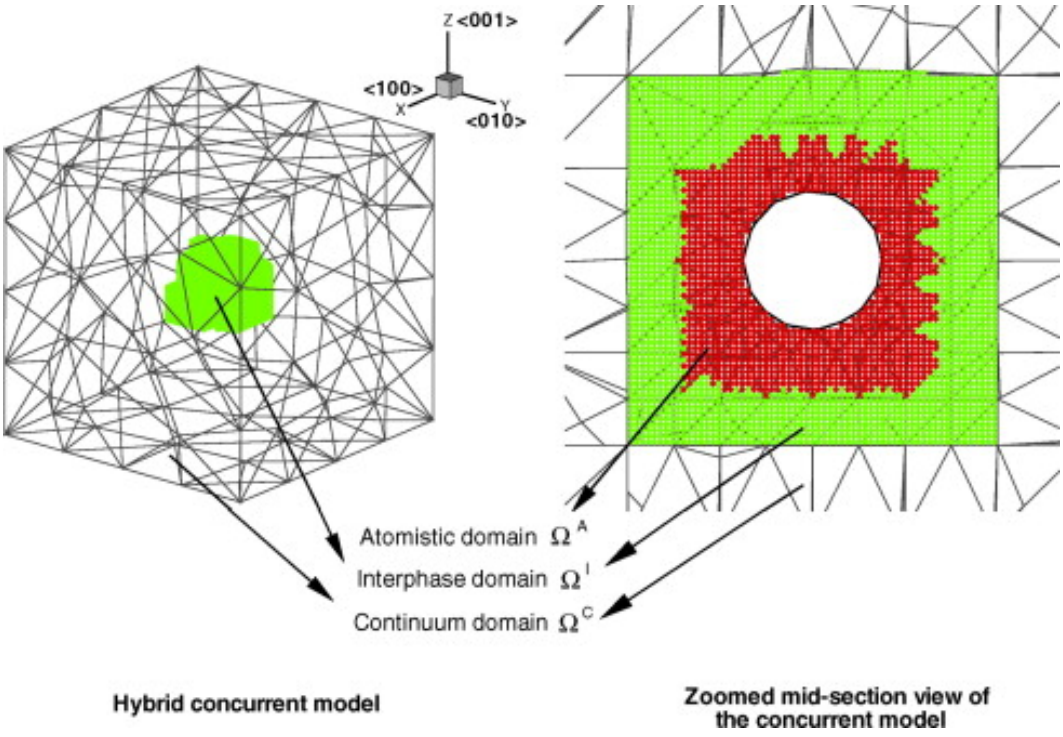

Figure 1.8: An example of the different domains in the AtC method for a void growth simulation [38].

combination of the atomistic and continuum regions. The AtC method achieves its atomistic-continuum coupling by blending at the level of forces. The motivation

comes from relating the stress in finite elements with the *virial stress* [135] at the atomistic level. The atomistic forces are gradually weakened along the interfacial region by using a weight function that goes from 1 to 0, while the finite element nodal forces are weakened similarly in the opposite direction.

Course Grained Molecular Dynamics (CGMD)

CGMD, outlined in multiple papers including [96], is another energy-based method that models the continuum region using the same MD principles that are used in the atomistic region like in the QC method. The difference from QC arises from the fact that the energy functional for the coarse-grained region (which is the continuum region) is defined as a constrained ensemble average of the atomistic energy under fixed thermodynamic relations. Also, the mesh formulation differs from all the other methods in the requirement that the finite element mesh must slowly refine down to the size of the atomistic simulation near the interface of the atomistic and continuum region. This helps it provide a seamless coupling between the two domains.

Concurrent Atomistic-Continuum (CAC) Method

The CAC method, developed by Chen et al. [21] reduces the degrees of freedom in the complete atomistic ensemble by reformulating a unified atomistic-continuum formulation of the conversation laws using classical statistical mechanics. It describes a crystalline material as a continuous medium containing lattice cells at the material-point level and embeds a group of discrete atoms (in case lattice unit cells containing more than one atom) within each lattice cell. CAC uses a single set of equations of motion that govern both the atomic and the continuum domains thereby eliminating the need of any interfacial region. As the complete field (atomistic and continuum) is represented in the governing equations, the CAC formulation can be solved efficiently using the finite element method and simulates nucleation and propagation of cracks via element separation and sliding.

The Quasicontinuum (QC) method

The QC method was first introduced in [111] can be either energy-based [10] and force-based [54]. It reduces the full atomic ensemble to a set of representative atoms and introduces kinematic constraints to described the motion of the entire atomic ensemble. This allows the use of full atomistic resolution only at places

with rapidly changing neighborhoods thereby allowing the modeling of significantly larger length scales than traditional MD. The main advantage of this method is that the entire domain in Fig. 1.5 is modeled using *coarse-grained* atomistics. There are two main flavors: local and nonlocal QC. Local-QC needs the interfacial region to connect the coarse-grained part with the fully atomistic part, whereas this coupling is completely seamless in nonlocal QC. A detailed formulation of a fully nonlocal energy-based QC method is provided in Chapter 2.

1.2.4 Finite temperature atomistic-to-continuum coupling

Most of the methods discussed so far (except CGMD) were developed for modeling materials at $0K$, neglecting the energy associated with the thermal fluctuations of atoms. Extension to finite-temperature is essential to model real-life materials as those shown in Fig. 1.3. A variety of such extensions have been postulated so far, e.g. [30, 33, 53, 60, 67, 87, 103, 113, 114, 121, 125], and are clubbed into three general approaches outlined below. As the multiscale modeling method of choice for this thesis is the QC method, the finite-temperature extensions presented below are designed specifically with the QC formulation as their basis.

Langevin dynamics

This method uses the Lagrange equations for dissipative systems [67, 88] with an additional noise term that dissipates the thermal energy in the atomistic-continuum coupling through a suitable force at the nodal level. This helps avoid the local heating of the atomistic domain due to reflection of phonons at the interface by damping them in accordance with the imposed force (which is a form of thermostat). This is a fully anhormonic method that can be used to study non-equilibrium, thermally-activated processes, but underestimates thermal properties (e.g. specific heat) which are quantified using criteria like reflection coefficients [17, 96].

Local harmonic approximation

Equilibrium Molecular Dynamics (EMD) is the study of *equilibrium properties* of a system at finite temperature using molecular dynamics. In this case, one assumes that the boundary conditions change infinitely slowly on the atomic scale, so the equations of motions are modified to simulate the effect of a heat bath or thermostat. HotQC[1] is a method developed by Tadmor et al. [113] which uses the *local*

[1] HotQC has been used as a name for multiple finite-temperature QC formulations, even though they differ in their finite-temperature realization. [113] is one such example.

harmonic approximation [61] to account for the loss of entropy in the continuum region of the QC domain and develops an effective Hamiltonian which is minimized to obtain the equilibrium equations. This method is designed to reproduce the equilibrium phase averages of systems that only depend on the atomistic region, or systems where equilibrium conditions are localized around the regions that can be treated atomistically (e.g. sharp cracks under nanoindenters).

1.2.5 Maximum-entropy (Max-ent) formulation

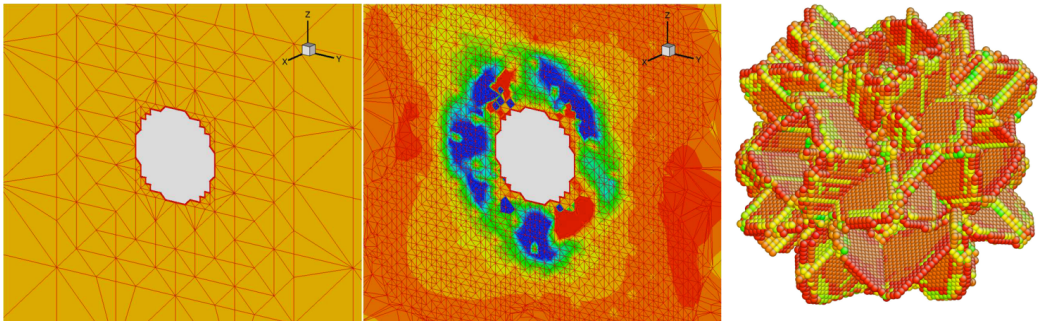

Figure 1.9: A model of nanovoid with the max-ent approach by Ariza et al. [9]. Left pane: The initial mesh before the simulation. Center pane: The mesh at the end of the simulation. Right pane: Dislocation emissions from the void at 300K.

Developed by Kulkarni, Knap, and Ortiz [60], this methods uses a variational approach and maximizes the entropy of the statistical-mechanics probability distribution in phase space to obtain the phase averages and equilibrium configuration of the system. The local nature of the probability distribution functions enables the modeling of non-equilibrium finite- temperature phenomena. More recently, this approach has been extended to non-equilibrium thermodynamics, allowing for mass and heat transport by introducing atomic-level kinetic transport relations and equipping each atom with its own temperature [87, 121]. We use this formulation to extend our fully non-local QC method to finite temperature, and the method is described in full detail in Chapter 4.

1.3 Current Shortcomings

Section 1.2 describes the current state of the art in scale-bridging techniques based on atomistics. However many of them still have multiple shortcomings that need to be addressed before they can be effectively used to model large-scale systems at finite temperature.

Most of the methods described in Section 1.2 have a static interfacial region and currently lack adaptive tracking down of full resolution wherever needed. This

necessitates the construction of a large initial atomistic region. With most of the techniques, a constitutive model for the continuum region needs to be specified (e.g. dislocation dynamics, linear elasticity, plasticity). This calls for parameter fitting at the interfacial region to match atomistics and in addition, prevents the technique from modeling the entire system using interatomic potentials. Methods like BD, BSM, CADD and CLS require special handling of the interfacial region by requiring special tricks to account for double-counting of energy (for energy-based methods) and removal of spurious force artifacts (for force-based methods). Once again, this results in little or no possibility of model adaptivity on the fly. Some of the techniques do not have finite-temperature modeling capabilities (e.g. QC by [111], BD, BSM), which severely limits the ability to predict and simulate the effect of temperature on the system. CADD and local-QC are limited to 2D scenarios. Finally, most of the work in scale-bridging models has been demonstrated for small examples and massively parallel, high-performance computing implementations are currently lacking.

These shortcomings show the continuing need to advance these techniques till a robust, adaptive, scale-bridging method that accurately models large-scale systems by combining spatial and temporal upscaling is established. The recent progress on the nonlocal QC front by Amelang, Venturini, and Kochmann [5] shows promising advances to address above limitations. Combining nonlocal QC with model adaptivity and the ability to model finite-temperature will bring us closer to efficiently model large-scale multiscale problems.

1.4 Organization of the thesis

The goal of my research is to develop a 3D fully nonlocal finite-temperature energy-based quasicontinuum framework with model adaptivity. To this end this thesis shows the development of a variety of theoretical and numerical tools to accomplish the above objectives and is structured in the following way. Chapter 2 discusses the formulation of a 3D fully nonlocal quasicontinuum method with optimal summation rules, which show improved efficiency and performance when compared to any other coarse-grained atomistic methods in the literature. In order to study large deformations using coarse-graining, an efficient and automatic way of updating and expanding the neighborhoods of sampling atoms used in the optimal summation is needed. Chapter 3 shows this need and provides an automatic adaptivity scheme which captures extremely highly deformed neighborhoods and demonstrates good agreement with fully atomistic simulations. The efficiency and

accuracy of this scheme is shown by discussion of void growth and indentation examples. In Chapter 4, the maximum-entropy principle is used to derive a finite-temperature extension for the fully nonlocal formulation from Chapter 2. Thermal expansion results for the fcc materials copper and aluminum are provided and elastic benchmark tests are conducted that validate this formulation. The chapter ends with a discussion of physical and thermal spurious force artifacts at finite temperature with a comparison of various summation rules and comments on their performance. In Chapter 5 the construction of 3D coincidence site lattice based bi-crystal symmetric-tilt grain boundaries using the new fully nonlocal finite temperature QC framework is discussed. Relaxed grain boundary energies of various symmetric tilt grain boundaries using the nonlocal scheme are compared to existing literature to show good agreement. The influence of temperature on grain boundary energy and relaxed grain boundary shapes is observed to close the chapter. As a final contribution to nonlocal QC, we discuss a proof of concept in the appendix for a second level of automatic adaptivity essential for studying plastic deformation: the local 3D adaptive mesh refinement. Appendix C discusses the theory and implementation of a constrained advancing front algorithm. Finally, we conclude in Chapter 6 by discussing the key highlights of this thesis and presenting avenues for future investigations using the developed techniques.

Chapter 2

THE FULLY-NONLOCAL QUASICONTINUUM METHOD

This chapter covers the basic derivation of the quasicontinuum (QC) method used for all the simulations in this thesis. We use a novel, fully nonlocal energy-based quasicontinuum formulation, which is described in detail in [3] and [5], but a brief derivation is carried out here to establish the formulation which will help us extend the method to finite temperature, as outlined in Chapter 4. We start from the Hamiltonian of the atomistic system, then derive the governing equations of motion for a reduced set of atoms, introduce energy-based optimal sampling rules to further improve computational efficiency, and end the chapter with a brief overview of the computational implementation of the method.

To illustrate the various aspects of the QC formulation, we show a two-void setup of fcc copper in $2D$. We study the high-deformation capabilities of this simulation later in Chapter 3, but for the purposes of this chapter, this setup works as an illustrative model. Some of the derivations in this chapter are taken from our recent publication [115].

2.1 Introduction of Representative Atoms

In atomistics, an ensemble of N atoms is described by their positions $\boldsymbol{q} = \{\boldsymbol{q}_1, \ldots, \boldsymbol{q}_N\}$ and momenta $\boldsymbol{p} = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_N\}$ with $\boldsymbol{p}_i = m_i \dot{\boldsymbol{q}}_i$ (m_i being the mass of atom i, and dots denoting material time derivatives). The system's total Hamiltonian H consists of the potential and kinetic energy, viz.

$$H(\boldsymbol{q}, \boldsymbol{p}) = V(\boldsymbol{q}) + \sum_{i=1}^{N} \frac{|\boldsymbol{p}_i|^2}{2m_i}, \qquad V(\boldsymbol{q}) = \sum_{i=1}^{N} V_i(\boldsymbol{q}), \tag{2.1}$$

with V denoting a suitable atomic interaction potential. In crystalline solids (including metals, ceramics and organic materials), the interaction potential usually allows for the above additive decomposition, where V_i is the energy contribution of atom i. The time evolution of the system is governed by Hamilton's equations, which yield Newton's equations of motion (for $i = 1, \ldots, N$), i.e.,

$$m_i \ddot{\boldsymbol{q}}_i = \boldsymbol{f}_i(\boldsymbol{q}) = -\frac{\partial V(\boldsymbol{q})}{\partial \boldsymbol{q}_i}, \tag{2.2}$$

where $\boldsymbol{f}_i(\boldsymbol{q})$ represents the total (net) force acting on atom i (including external force contributions). An example of such an ensemble is shown in Figure 2.1, in

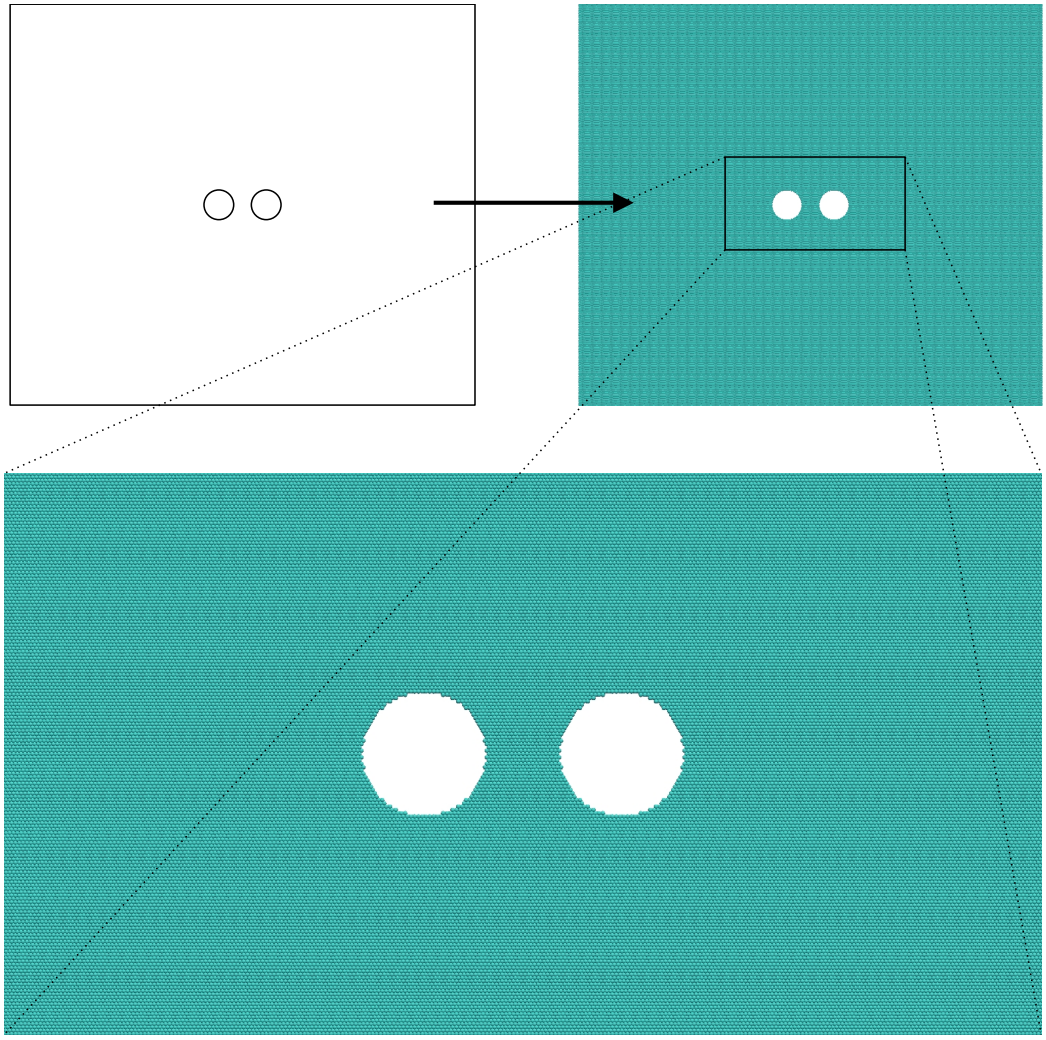

Figure 2.1: A two-void system modeled using atomistics. The material is Cu, the domain spans $0.12\mu m \times 0.12\mu m$, with void radius $4nm$ and the system is at $0K$. The top left panel shows the domain that needs to be modeled and the top right shows the atomistic ensemble (using fcc copper). The lower panel shows a zoomed in view of the individual atoms that make up part of the ensemble. All the *undeformed positions* of these atoms are on lattice sites that correspond to $[111]$ plane of fcc Cu.

which a two-void scenario is modeled using atomistics. For this specific illustrative example, we use a setup of $N = 302,003$ atoms to represent the system.

It is generally not feasible to use the above framework to large systems (consisting of over 10^{10} atoms), even for extremely short range interatomic potentials, due to the limitations imposed by computational resources. The QC method addresses this limitation by replacing the full atomic ensemble by a reduced set of $N_h (\ll N)$ representative atoms (often referred to as *repatoms* for short). The method postulates that

except in the vicinity of flaws and lattice defects such as cracks and dislocations, respectively, the local environment of each atom in a crystal lattice is almost identical up to rigid body motion. This motivates the introduction of repatoms with current positions $\boldsymbol{x}(t) = \{\boldsymbol{x}_1(t), \ldots, \boldsymbol{x}_{N_h}(t)\}$ (in contrast to their initial positions in some reference configuration, $\boldsymbol{X}(t) = \{\boldsymbol{X}_1(t), \ldots, \boldsymbol{X}_{N_h}(t)\}$). The approximate current positions $\boldsymbol{q}_i^h$ and momenta $\boldsymbol{p}_i^h$ of all atoms $i = 1, \ldots, N$ now follow from interpolation, i.e.,

$$\boldsymbol{q}_i \approx \boldsymbol{q}_i^h = \sum_{a=1}^{N_h} N_a(\boldsymbol{X}_i)\, \boldsymbol{x}_a, \qquad \boldsymbol{p}_i \approx \boldsymbol{p}_i^h = m_i \sum_{a=1}^{N_h} N_a(\boldsymbol{X}_i)\, \dot{\boldsymbol{x}}_a, \tag{2.3}$$

where $N_a(\boldsymbol{X}_i)$ denotes the shape function of repatom a evaluated at the reference position of lattice site i. The reference configuration is usually chosen as the undeformed crystal ground state. Figure 2.2 shows the reduced set of atoms for the two-void illustrative system presented in Figure 2.1. For this illustration, we choose $N_h = 3385$ repatoms, which constitute $\sim 1\%$ of the entire system. As one can expect the local environment around the voids to develop more lattice defects, or emit dislocations, or generally show significant atomic-level motion, a higher density of repatoms is chosen near the voids and this density is reduced away into the bulk, where the local environment will not be expected to change as much.

By choosing an affine interpolation on a Delaunay triangulation $\mathscr{T}$ as in the original QC method [54, 111, 118], we ensure that this coarse-graining scheme locally recovers the exact atomic ensemble when all atoms are turned into repatoms due to the Kronecker property $N_a(\boldsymbol{X}_b) = \delta_{ab}$ for all $1 \leq a, b \leq N_h$, with δ_{ij} denoting Kronecker's delta ($\delta_{ij} = 1$ if $i = j$ and 0 else). With this interpolation, the reduced (and approximate) Hamiltonian of the coarse-grained system depends on the positions and momenta of the repatoms through the approximate atomic positions $\boldsymbol{q}^h = \{\boldsymbol{q}_1^h, \ldots, \boldsymbol{q}_N^h\}$ and momenta $\boldsymbol{p}^h = \{\boldsymbol{p}_1^h, \ldots, \boldsymbol{p}_N^h\}$, and can be written as

$$H^h = H(\boldsymbol{q}^h, \boldsymbol{p}^h) = H(\boldsymbol{x}, \dot{\boldsymbol{x}}) = \sum_{i=1}^{N} \frac{|\boldsymbol{p}_i^h|^2}{2m_i} + V(\boldsymbol{q}^h). \tag{2.4}$$

Therefore, instead of solving for the positions and momenta of all the N atoms, this approximation allows us to only solve for the positions and momenta of the N_h repatoms, and obtain the approximate positions of all the remaining atoms using (2.3). This leads to a considerable and controllable reduction of the independent degrees of freedom from $d \times N$ to $d \times N_h$. The positions of the repatoms now evolve from applying Hamilton's equations to the approximate Hamiltonian, which yields

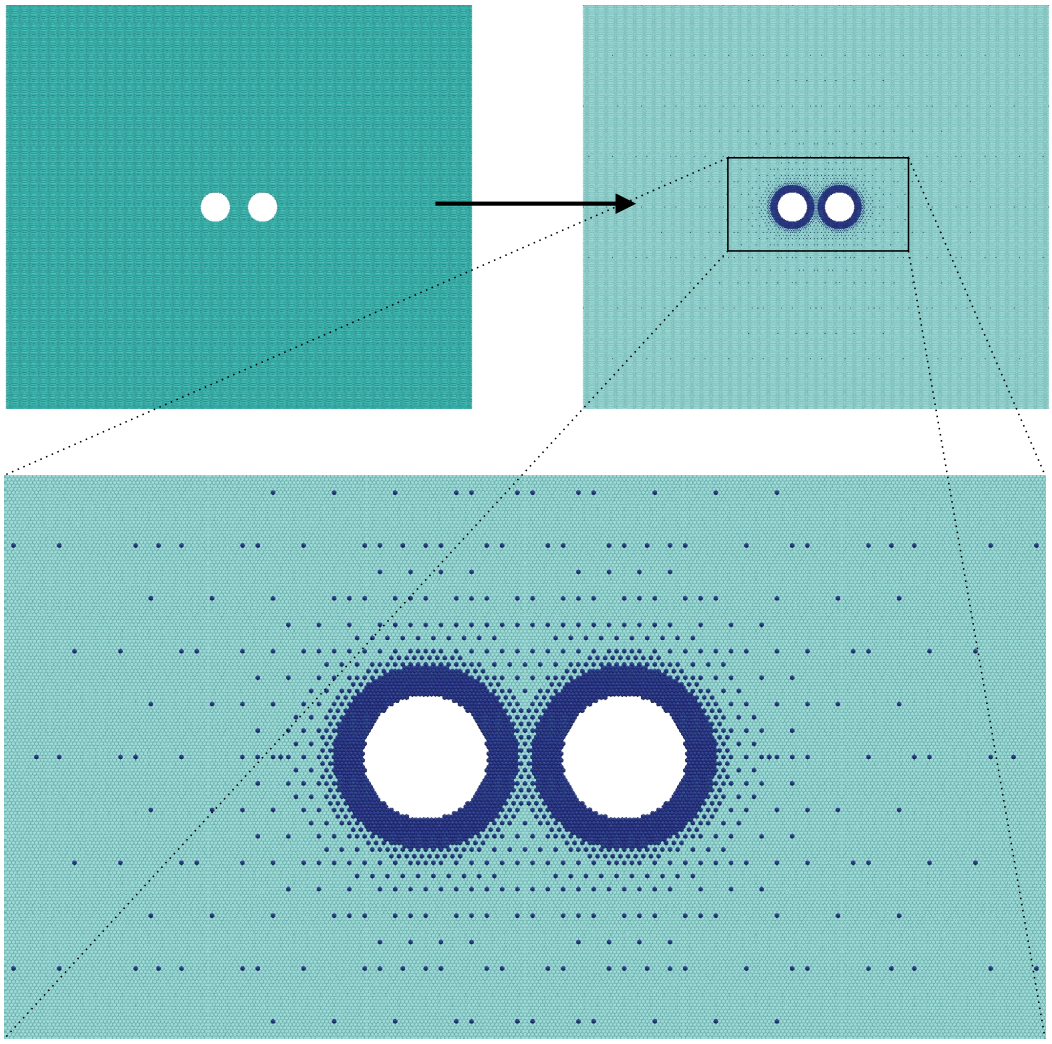

Figure 2.2: Modeling the two-void system by a *reduced set* of atoms. Top left panel shows the fully atomistic system as shown in Fig. 2.1 and the top right panel shows the representative atoms in dark blue. The bottom panel shows a zoomed version near the voids. The full set of atoms is shown in a faded color in the background.

the net force on repatom k as

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\frac{\partial H(\boldsymbol{x},\dot{\boldsymbol{x}})}{\partial \boldsymbol{x}_k} = \sum_{j=1}^{N} \boldsymbol{f}_j^h(\boldsymbol{q}^h) N_k(\boldsymbol{X}_j), \tag{2.5}$$

with

$$\boldsymbol{f}_j^h(\boldsymbol{q}^h) = -\frac{\partial V(\boldsymbol{q}^h)}{\partial \boldsymbol{q}_j^h} = -\sum_{i=1}^{N} \frac{\partial V_i(\boldsymbol{q}^h)}{\partial \boldsymbol{q}_j^h}, \tag{2.6}$$

the total force (internal and external) acting on atom j. We note that in an infinite Bravais lattice (i.e., in a defect-free single-crystal in the absence of external loading

or free surfaces), the forces on all atoms vanish in the undeformed ground state, i.e., we have $\boldsymbol{f}_i^h(\boldsymbol{q}^h) = \boldsymbol{0}$, so that the net forces on all repatoms vanish as well ($\boldsymbol{F}_k = \boldsymbol{0}$).

In any QC simulation, the initial choice of repatoms is specific to the application, and is made by the user. Our QC formulation constructs a simplicial mesh (similar to finite elements) by using the undeformed repatom positions as the set of known points. Figure 2.3 shows the formation of the mesh for our illustrative two-void system. The elements formed by this mesh are used for various purposes, including the implementation of the affine interpolation in (2.3) and determining locations and weights of sampling atoms which are introduced to calculate thermodynamic quantities like energy (see Section 2.2).

2.2 Approximating the Potential Energy using Sampling Atoms

Although the total number of degrees of freedom that are modeled has been reduced, the evaluation of (2.5) and (2.6) still requires summation over all N lattice sites, which once again makes it prohibitively expensive. Therefore, summation or sampling rules have been introduced which replace the sum over all atoms either in the computation of the Hamiltonian (so-called *energy-based QC*) [10, 112] or in the calculation of repatom forces (so-called *force-based QC*) [54] by a weighted average over carefully selected *sampling atoms* (analogous to quadrature rules employed in finite element methods). We use an energy-based formulation which approximates the total potential energy by a weighted sum over N_s sampling atoms and computes the kinetic energy based on the motion of the N_h repatoms. Thus, the Hamiltonian H^h is further approximated by $\widetilde{H}^h$ via

$$\begin{aligned} H^h(\boldsymbol{q}^h, \boldsymbol{p}^h) &= \sum_{i=1}^{N} \left(V_i(\boldsymbol{q}^h) + \frac{|\boldsymbol{p}_i^h|^2}{2m_i} \right) \\ \approx \quad \widetilde{H}^h(\boldsymbol{q}^h, \boldsymbol{p}^h) &= \sum_{\alpha=1}^{N_s} w_\alpha V_\alpha(\boldsymbol{q}^h) + \sum_{\beta=1}^{N_h} \hat{w}_\beta \frac{|\boldsymbol{p}_\beta^h|^2}{2m_\beta}, \end{aligned} \tag{2.7}$$

where w_α is the weight of sampling atom α. Physically, w_α denotes the number of lattice sites represented by sampling atom α. Out of the various existing flavors of the QC method, the main distinguishing factor is the way in which these N_s sampling atoms are chosen. This selection of sampling atoms aims for a compromise between maximum accuracy and maximum efficiency and gives rise to various summation rules. The summation rules must satisfy zeroth and first order consistency

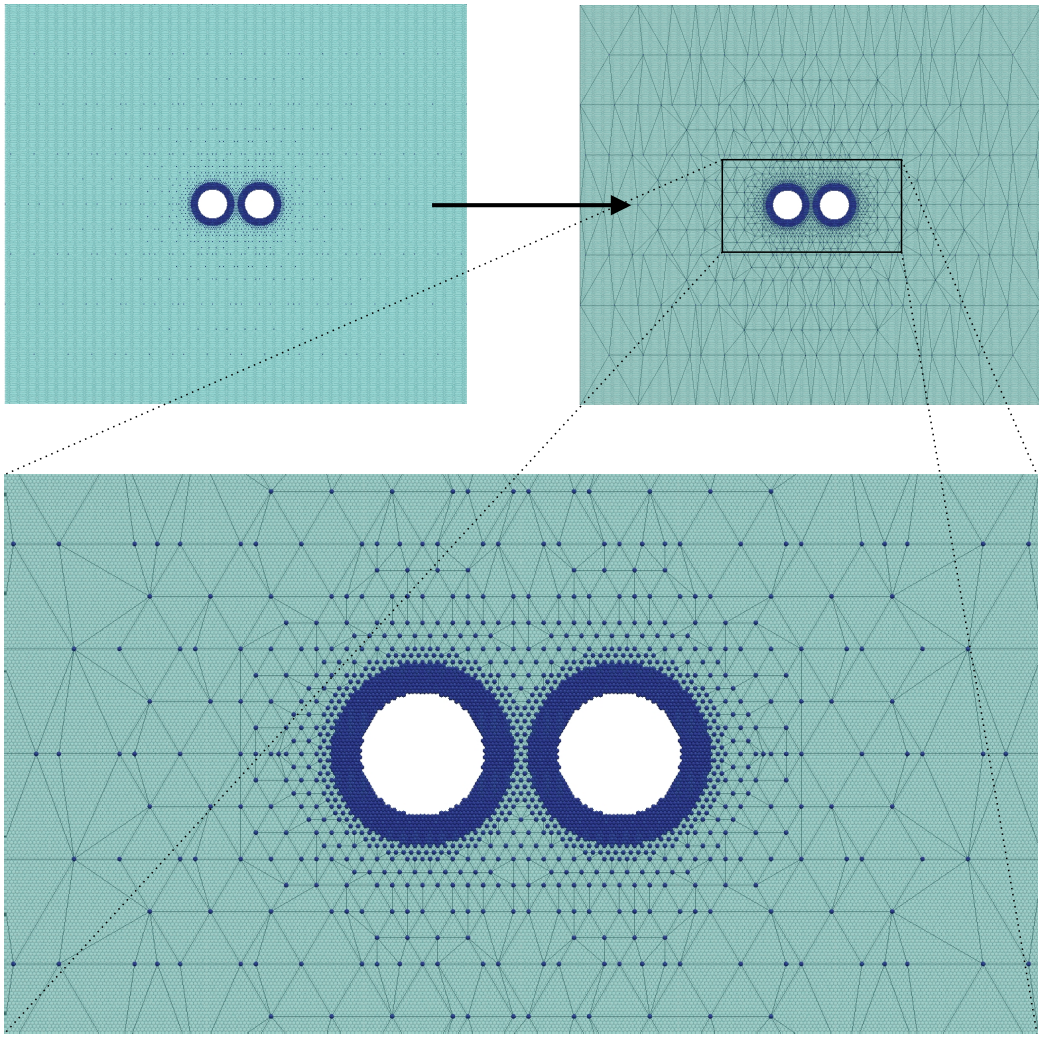

Figure 2.3: Formation of a simplicial mesh for the two-voids simulation with the repatoms as the nodes of the simplices. As before, the top panels show the meshing process for the entire simulation and the lower panel shows the zoomed in portion around the voids.

for the rule to be exact up to first order. This means that we must have

$$\sum_{\alpha=1}^{N_s} w_\alpha = N, \tag{2.8a}$$

$$\frac{1}{N}\sum_{\alpha=1}^{N_s} w_\alpha \boldsymbol{X}_\alpha = \boldsymbol{X}_s, \tag{2.8b}$$

where

$$\boldsymbol{X}_s = \frac{1}{N}\sum_{i=1}^{N} \boldsymbol{X}_i$$

is the location of the center of mass of all the N lattice sites. $\hat{w}_\beta$ denotes the weight of repatom β, here determined from the cell volumes of a Voronoi tessellation over the set of all repatoms [55], which guarantees the zeroth and first order consistency of the weights required in (2.8a) and (2.8b). m_β is the mass of repatom β (in case of a monatomic crystal with atomic mass m, we simply have $m_\beta = m$ for $\beta = 1,\ldots,N_h$).

Amelang, Venturini, and Kochmann [5] outline a comprehensive study of the state of the art of these *summation rules* which determine the choice of the locations of sampling atoms and the weights associated with them.

2.2.1 Force artifacts

Summation rules and the introduction of sampling atoms does reduce computational complexity, but also gives rise to force artifacts. This can be simply observed by writing an expression for the force on a repatom k by using the Hamiltonian from (2.7) as follows:

$$\begin{aligned}\widetilde{\boldsymbol{F}}_k(\boldsymbol{x}) &= -\frac{\partial \widetilde{H}^h(\boldsymbol{q}^h,\boldsymbol{p}^h)}{\partial \boldsymbol{x}_k},\\ &= -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j=1}^{N} \frac{\partial V_\alpha(\boldsymbol{q}^h)}{\partial \boldsymbol{q}_j^h} N_k(\boldsymbol{X}_j),\\ &= -\sum_{j=1}^{N} \left(\sum_{\alpha=1}^{N_s} w_\alpha \frac{\partial V_\alpha(\boldsymbol{q}^h)}{\partial \boldsymbol{q}_j^h} \right) N_k(\boldsymbol{X}_j).\end{aligned} \tag{2.9}$$

Comparing this form to (2.5) and (2.6), we notice that for a general choice of N_s and w_α, the term in paranthesis does not vanish, and so it is possible to have $\widetilde{\boldsymbol{F}}_k(\boldsymbol{x}) \neq \boldsymbol{0}$ even if $\boldsymbol{f}_j^h = \boldsymbol{0}$. This is what is termed a *spurious force artifact*, and needs to be addressed in order to minimize the error associated with energy-based sampling in QC.

2.2.2 Optimal summation rule

For the remainder of this thesis, unless specifically mentioned otherwise, we use the optimal sampling rules recently proposed in [5] which place sampling atoms at all repatom locations (undergoing exact local neighborhood changes during deformation) as well as Cauchy-Born-type sampling atoms within elements formed by the simplicial mesh as shown in Figure 2.3, which has two flavors: the first and the second order. The first-order rule adds one Cauchy-Born sampling atom in the interior of each element; the second-order rule also adds one Cauchy-Born sampling atom per element and additionally adds one Cauchy-Born sampling atom on

relevant element edges (one per edge), see Figure 3.1. (We note that "exact neighborhood changes" here refers to the atomic positions obtained from applying the exact interpolation to all lattice sites within elements, as opposed to approximating atomic positions by the Cauchy-Born rule.)

For both optimal sampling flavours, weights w_α are obtained such as to minimize spurious force artifacts described in Section 2.2.1. The summation rule guarantees that there are no spurious force artifacts in any affinely-deformed non-uniform QC mesh in one or two dimensions, and for three dimensions the effect is minimal. This is achieved by giving all the sampling atoms which are placed at repatom locations an identical weight w_{rep}, and choosing the weights of the inner sampling atoms such that the force in (2.9) vanishes. The actual value of w_{rep} is identified by minimizing the energy approximation error. The complete derivation and analysis can be found in [5], and is not included here. Figure 2.4 shows the formation of sampling atoms for the illustrative two-void system using the optimal second-order summation rule. For the given setup, the summation rule constructs $N_S = 7037$ sampling atoms, which means that for this example, solving for (2.1) has been reduced from $O(302,003 \times 302,003)$ computations to $O(3385 \times 7037)$ computations.

2.3 Governing Equations

Once the weights w_α and $\hat{w}_\beta$ are determined, the equilibrium configurations of the system can be solved for. Hamilton's equation of motion for a repatom β is now obtained by differentiation of (2.7) in analogy to (2.2) and (2.5):

$$\hat{w}_\beta m_\beta \ddot{\boldsymbol{x}}_\beta = \widetilde{\boldsymbol{F}}_\beta(\boldsymbol{x}) = -\frac{\partial \widetilde{H}^h}{\partial \boldsymbol{x}_\beta} = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j=1}^{N} \frac{\partial V_\alpha}{\partial \boldsymbol{q}_j^h}(\boldsymbol{x}) N_\beta(\boldsymbol{X}_j). \tag{2.10}$$

2.3.1 Example: Embedded Atom Method

A family of interatomic potentials for metals (which will be used for our entire study) is the Embedded Atom Method (EAM) potentials, introduced by [66]. For potentials of the embedded-atom type, the interatomic potential energy for a site α is given as

$$V_\alpha(\boldsymbol{q}) = \frac{1}{2} \sum_{j \in \mathscr{I}_\alpha, j \neq \alpha} \Phi(r_{\alpha j}) + \mathscr{F}(\rho_\alpha), \qquad \rho_\alpha = \sum_{j \in \mathscr{I}_\alpha} f(r_{\alpha j}). \tag{2.11}$$

The pair potential $\Phi(r_{\alpha j})$ represents the energy due to electrostatic interactions between atom α and its neighbor j, whose distance is given by $r_{\alpha j} = |\boldsymbol{r}_{\alpha j}|$ and $\boldsymbol{r}_{\alpha j} = \boldsymbol{q}_\alpha - \boldsymbol{q}_j$. ρ_α denotes the (approximated) effective electron density which atom

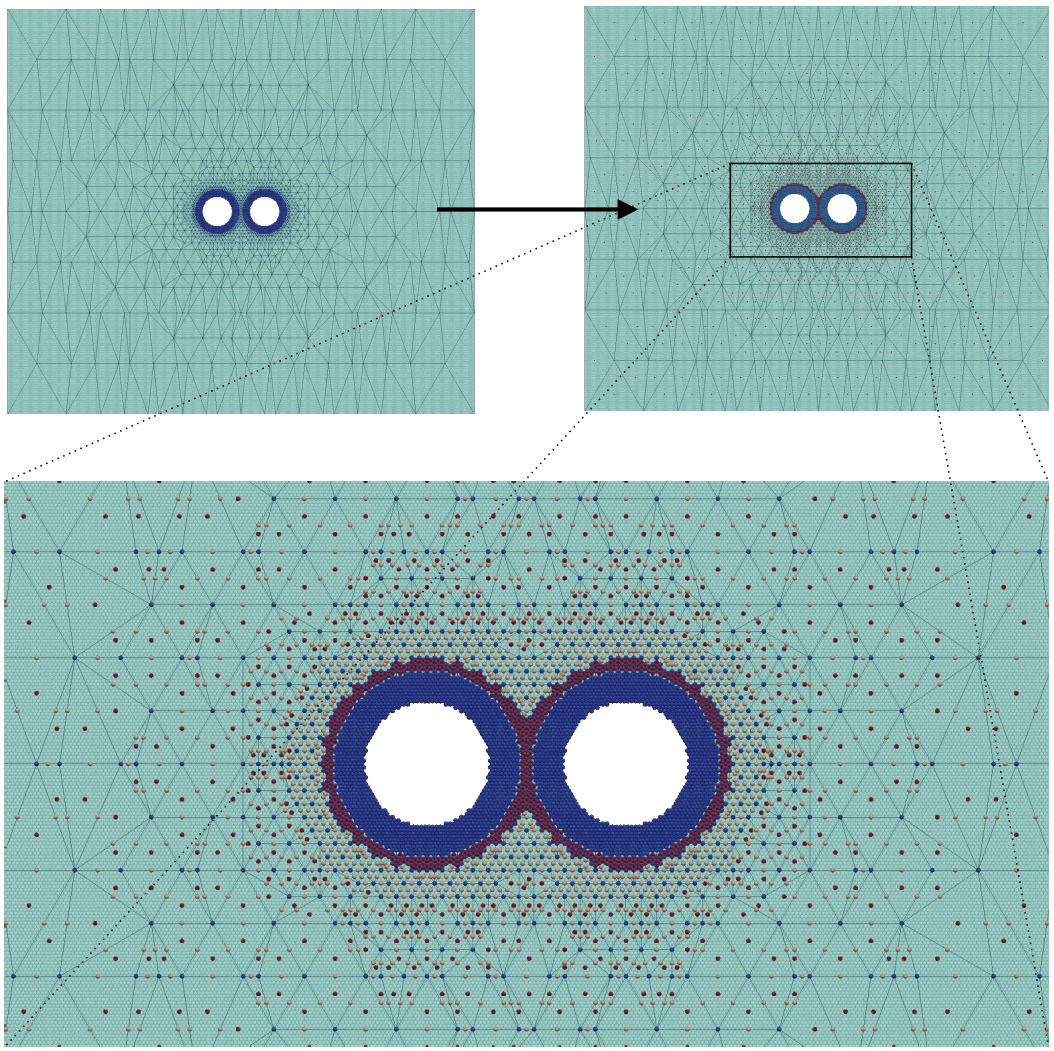

Figure 2.4: Creation of sampling atoms for the two-void system using the second-order optimal summation rule. The top left pane shows the mesh that is used in creating all the sampling atoms which are shown on the top right, and a zoomed version on the bottom pane. The colors represent the different types of sampling atoms. Blue sampling atoms are located on the nodes of elements of the mesh, red sampling atoms are internal to the element, and the cream colored ones represent sampling atoms on the edge of the elements.

α senses due to its neighboring atoms. $f(r_{\alpha j})$ is the electron density at site α due to atom j as a function of their distance $r_{\alpha j}$. $\mathscr{F}(\rho_\alpha)$ accounts for the energy release upon embedding atom α into the local electron density ρ_α. $\mathscr{I}_\alpha(t)$ is the interatomic sphere of interaction around sampling α at time t, defined by the interatomic cutoff potential r_{cutoff}, i.e.,

$$\mathscr{I}_\alpha(t) = \left\{ i \in \{1,\ldots,N\} \;\middle|\; \left\| \boldsymbol{q}_i^h(t) - \boldsymbol{x}_\alpha(t) \right\| \leq r_{\text{cutoff}} \right\}. \tag{2.12}$$

Beyond this cutoff, the potential energy contributed to the site α is negligible, and can be neglected. This means that for EAM potentials, the sum over j in (2.10) can conveniently be reduced to $\mathscr{I}_\alpha(t)$, which further simplifies the force on a repatom β to be

$$\hat{w}_\beta m_\beta \ddot{\boldsymbol{x}}_\beta = \widetilde{\boldsymbol{F}}_\beta(\boldsymbol{x}) = -\frac{\partial \widetilde{H}^h}{\partial \boldsymbol{x}_\beta} = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j\in\mathscr{I}_\alpha(t)} \frac{\partial V_\alpha}{\partial \boldsymbol{q}_j^h}(\boldsymbol{x}) N_\beta(\boldsymbol{X}_j). \tag{2.13}$$

Using the form of potential given by (2.11), the force on a repatom β becomes

$$\widetilde{\boldsymbol{F}}_\beta(\boldsymbol{x}) = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j\in\mathscr{I}_\alpha(t)} \left[\frac{1}{2}\Phi'(r_{\alpha j}^h) + \mathscr{F}'(\rho_\alpha^h) f'(r_{\alpha j}^h)\right] \frac{\boldsymbol{r}_{\alpha j}^h}{r_{\alpha j}^h} \left[N_\beta(\boldsymbol{X}_\alpha) - N_\beta(\boldsymbol{X}_j)\right]. \tag{2.14}$$

2.4 Solver

We integrate (2.10) numerically by approximating the repatom accelerations by a second-order central-difference scheme in an implicit fashion for constant time steps $\Delta t = t_{n+1} - t_n$, i.e., we solve

$$\hat{w}_\beta m_\beta \frac{\boldsymbol{x}_\beta(t_{n+1}) - 2\boldsymbol{x}_\beta(t_n) + \boldsymbol{x}_\beta(t_{n-1})}{(\Delta t)^2} = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j\in\mathscr{I}_\alpha(t)} \frac{\partial E_\alpha}{\partial \boldsymbol{q}_j^h}\big(\boldsymbol{x}(t_{n+1})\big) N_\beta(\boldsymbol{X}_j). \tag{2.15}$$

for repatom positions $\boldsymbol{x}(t_{n+1})$ in a matrix-less fashion, using the Fast Inertial Relaxation Engine (FIRE) [14]. This method blends well with the QC equilibrations as they are computationally similar to molecular dynamics, which was what the method was designed to optimize. In FIRE, the velocities are reset when the global power of the system becomes negative ("uphill motion"), and the quantities like the timestep and velocity for every repatom are dynamically adjusted to accelerate progress towards the lowest energy, using the FIRE parameter α. The general steps in solving a QC problem using the formulation developed in Section 2.3 using FIRE are outlined in Algorithm 1. The implicit time discretization furnishes the scheme with unconditional stability; sufficient accuracy is confirmed by a time step convergence study (any suitable finite-difference scheme such as velocity-Verlet can be used instead).

2.5 Computational Implementation

The goal of this section is to provide a rough map of the computational implementation that goes into successfully running a QC simulation using the formulation provided in Sections 2.1, 2.2, and 2.3. Specifically we provide a brief overview of the key data structures used in our in-house QC code.

Algorithm 1 Algorithm for the QC solver

while Global forces are not equilibrated **do**
  % perform **velocity verlet** on all repatoms
  for each repatom **do**
    $\boldsymbol{a}^i = \frac{\boldsymbol{F}^i}{m}$, force obtained from (2.14)
    $\boldsymbol{v}^i = \boldsymbol{v}^{i-\frac{1}{2}} + \frac{1}{2}\mathbf{a}^i \Delta t$
  end for
  % perform **FIRE** algorithm
  $P_{total} = \sum_{i=1}^{N_h} \boldsymbol{F}_i \cdot \boldsymbol{v}_i$, summing power of all repatoms
  for each repatom **do**
    $\boldsymbol{v}^i = (1-\alpha)\boldsymbol{v}^i + \alpha \hat{\boldsymbol{F}}^i |\boldsymbol{v}^i|$, $\hat{\text{ffl}}$ denoting unit vector
  end for
  if $P_{total} > 0$ **then**
    $\Delta t \uparrow$, $\alpha \downarrow$
  else
    $\Delta t \downarrow$, $\alpha = \alpha_{start}$
    for each repatom **do**
      $\boldsymbol{v}^i = 0$, reset the repatom velocities
    end for
  end if
  % perform **verlet update** on all repatoms
  for each repatom **do**
    $\boldsymbol{a}^i = \frac{\boldsymbol{F}^i}{m}$,
    $\boldsymbol{x}^{i+1} = \boldsymbol{x}^i + \boldsymbol{v}^i \Delta t \frac{1}{2}\mathbf{a}^i \Delta t^2$
    $\boldsymbol{v}^{i+\frac{1}{2}} = \boldsymbol{v}^i + \frac{1}{2}\mathbf{a}^i \Delta t$
  end for
end while

2.5.1 Repatom Container

To store the reduced degrees of freedom in the system as shown in Section 2.1, we use a datastructure which we refer to as the **repatom container**. It contains all the information pertaining to the repatoms like positions, weights, velocities, forces, etc. In Chapter 4, we introduce a new mean-free parameter called ω, which is also defined for every repatom, and is therefore handled by this container. In any simulation, this is the first data structure that has to be populated based on the application.

2.5.2 Mesh and Elements

The QC formulation uses linear interpolation to obtain the positions of lattice sites which are not being modeled using the repatoms (if needed), and in addition, to

calculate sampling atom and neighbor positions. This necessitates the use of a simplicial **mesh** that connects the undeformed repatom positions. Once the initial set of repatom are populated, the QC formulation uses an external mesher (CGAL) and then a custom repairing process to form the final mesh that is used for subsequent calculations. The constructed "simplices" or elements are used to store important information of the system like the lattice structure (for possible multi-lattice applications), and orientation (for polycrystal applications).

2.5.3 Sampling Atoms

Once the user populates the initial entries of the repatom container, using the optimal summation rule briefly outlined in Section 2.2, the sampling atoms are automatically populated by the QC formulation. The QC code stores the collection of sampling atoms together by using an individual **sample atom** as a data structure. Sample atom population is done by looping through the elements and depends on the flavor of the summation rule (first-order or second-order), and all the different properties of the sample atom like position, weight, and centrosymmetry [16] are stored accordingly. It is important to note that while all repatoms locations are also sampling atom locations (as required by the optimal summation rule), they are stored separately, as computationally they perform different functional tasks. The sampling atom is used to perform energy calculations and derive the repatom forces derived in 2.14. Using this force and implementing it in the solver (ref. Section 2.4), the equilibrium solution of the system is determined.

2.5.4 Neighborhood Container

This is perhaps the most important (and complicated) data structure of the entire QC formulation. Its need arises from the fact that in order to calculate any property related to a sampling atom (like energy or centrosymmetry) or repatom forces, a complete representation of the neighborhood of every sampling atom is required. For further illustration, we reconsider the expression of the force on a repatom as shown in (2.14),

$$\widetilde{\boldsymbol{F}}_\beta(\boldsymbol{x}) = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j\in\mathscr{I}_\alpha(t)} \left[\frac{1}{2}\Phi'(r^h_{\alpha j}) + \mathscr{F}'(\rho^h_\alpha)\, f'(r^h_{\alpha j})\right] \frac{\boldsymbol{r}^h_{\alpha j}}{r^h_{\alpha j}} \left[N_\beta(\boldsymbol{X}_\alpha) - N_\beta(\boldsymbol{X}_j)\right].$$

The expression shows that at every sampling atom α, we need the knowledge of the locations of all atoms j that lie in its neighborhood $\mathscr{I}_\alpha$. In traditional molecular statics or dynamics, such knowledge is obtained by simply performing a distance search, but in QC that cannot be done, as we are modeling a reduced system. Our

formulation achieves it by storing the set of neighbors for every sampling atom in a data structure that we term as the **neighborhood container**. A detailed description of how our neighborhood storage achieves automatic adaptivity for high deformations is provided in Chapter 3.

2.6 Conclusions

This chapter shows a brief derivation of the general QC method and provides implementation details of the fully nonlocal energy-based QC formulation that is used throughout this thesis.

Chapter 3

AUTOMATIC ADAPTIVITY I: NEIGHBORHOOD UPDATES

Chapter 2 introduces the fully nonlocal QC method that seamlessly bridges across scales without the need to differentiate between local and nonlocal domains and does not have any notion of atomistic and continuum regions. If all atoms of the ensemble are included in the reduced set of representative atoms ($N_h \rightarrow N$), the model recovers full atomistics in an adaptive manner in a numerically elegant fashion. Therefore, in addition to the formulation presented in Chapter 2, *automatic adaptivity*, in which one does not need to require a-priori knowledge about the resolution needed at specific regions of the simulation, will in principle, enable highly efficient simulations. However, such an adaptive model refinement comes with a set of challenges which need to be addressed before we can perform adaptive simulations. (*i*) Use of the Lagrangian QC formulation with a mesh defined in the reference configuration becomes challenging in the atomistic regions where atomic neighborhoods change continuously and thus must be updated during the simulation (unlike in traditional QC, our formulation uses a mesh also in the atomistic domain in order to maintain the complete nonlocal behavior). (*ii*) Similar to MD, Verlet lists that store local atomic neighborhoods are required but the positions of all atoms are not known a-priori but must be generated on demand by interpolation (i.e., the positions of all atoms in the simulation are never explicitly known to the model). (*iii*) As in any remeshing technique, refinement criteria and geometric algorithms must be chosen to provide acceptable mesh quality. Poor mesh quality results in substantial errors as has been reported by [3]. (*iv*) In order to seamlessly refine down to full atomistics, every representative atom (i.e., every node in the mesh) is required to coincide with an existing atomic site (identified, e.g., by the aid of the discrete crystal lattice in the reference configuration). This chapter reports the adaptive model that efficiently captures high-deformation neighborhoods on the fly, and a 2D mesh refinement scheme which has enabled our QC formulation to accurately perform simulations which are verified by comparison to full atomistics.

Research presented in this chapter was published in reference [115].

3.1 Mesh Adaptivity: Refinement Criteria and Algorithms

3.1.1 Mesh Refinement in the Fully-Nonlocal QC Method

Model adaptivity enables us to efficiently deploy full atomistic resolution only where it is indeed required, such as in the immediate vicinity of lattice defects. Therefore, the initial configuration of all simulations contains as few representative atoms as possible; i.e., defect-free regions use a significantly coarsened description (assuming the crystal is in its equilibrium ground state) and full atomistic resolution is confined to the vicinity of pre-existing defects (if any). Subsequently, every simulation step going from time t_n to $t_{n+1} = t_n + \Delta t$ involves *(i)* solving for the new repatom positions $\boldsymbol{x}(t_{n+1})$ according to (2.15), *(ii)* identifying regions requiring refinement and updating the triangulation $\mathscr{T}(t_n)$, and *(iii)* updating all sampling atom neighborhoods $\mathscr{I}_\alpha$ based on the new repatom positions and the refined mesh. Step *(ii)* will be discussed in the following, whereas step *(iii)* follows in Section 3.2.1. We note that we restrict our attention to model *refinement*, i.e., the addition of new repatoms to enrich the repatom set $\mathscr{N}_h$. We do not consider model *coarsening* (i.e., the elimination of repatoms). Of course, the latter is equally important for efficiency but also conceptually more challenging (and not available in any existing QC codes to the best of our knowledge). A discussion of mesh coarsening within the fully nonlocal QC method can be found in Section 3.3.

3.1.2 Refinement Criterion

Since the interpolation is based on a mesh, it makes sense to adopt an element-wise refinement criterion [54, 111]. Specifically, the chosen affine interpolation results in constant deformation gradients within the n_e elements. Therefore, we compute

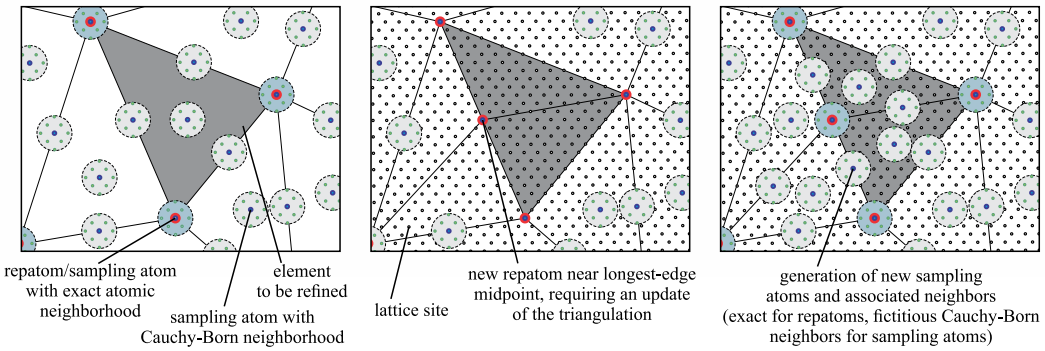

Figure 3.1: Schematic view of the remeshing process: identification of elements to refine, insertion of new repatoms (nodes) at lattice sites, update of the triangulation $\mathscr{T}$, and insertion of new sampling atoms and their associated local atomic neighborhoods for the second-order optimal sampling rule [5].

the second invariant of the right Cauchy-Green tensor $\boldsymbol{C} = \boldsymbol{F}^{\mathsf{T}}\boldsymbol{F}$ in each element and identify an element e for refinement if

$$I_e L_e \geq b\,\varepsilon \qquad \text{with} \qquad I_e = \sqrt{\frac{1}{2}\left|\mathrm{tr}\,(\boldsymbol{C}_e)^2 - (\mathrm{tr}\,\boldsymbol{C}_e)^2\right|} \qquad \text{for} \quad e = 1,\ldots,n_e. \tag{3.1}$$

b is the magnitude of the smallest Burgers' vector of the crystal, ε represents a tolerance threshold for refinement (in our examples chosen as $\varepsilon = 0.1 - 0.6$), and L_e is the longest element edge. In contrast to [54] we choose the longest element edge L_e as the characteristic length and do not consider the element volume. We note that this is, of course, only one possible choice for the refinement criterion. One can alternatively define refinement criteria based on sampling atoms (replacing the element-wise criterion), e.g., by defining thresholds for their centrosymmetry parameter or energy per atom. Rigorous error estimators (but not for the fully-nonlocal QC version) have also been proposed from a mathematical standpoint; see, e.g., [1, 80]. In the following examples, we will use criterion (3.1) which, after each solution step, results in a set of elements flagged for refinement.

3.1.3 Refinement Algorithm

Refinement necessitates two steps: *(i)* insertion of new nodes (i.e., the introduction of new repatoms) and *(ii)* their integration into the existing mesh while retaining all existing nodes. As a special difficulty of the nonlocal QC method, newly-added nodes must lie on exact lattice sites to ensure the correct, seamless refinement down to exact atomistics. We assume that the atomic crystal defines a discrete set $\mathscr{B}$ of Bravais lattice sites so that in d dimensions

$$\mathscr{B} = \left\{ \boldsymbol{X} \in \mathbb{R}^d \;\middle|\; \exists\, n_1,\ldots,n_d \in Z \text{ s.t. } \boldsymbol{X} = \sum_{i=1}^{d} n_i \boldsymbol{A}_i \right\} \tag{3.2}$$

with lattice unit vectors $\boldsymbol{A}_i$ $(i = 1,\ldots,d)$. In case of polycrystals, several such lattices coexist spatially and each element is assigned to one underlying lattice (or a non-crystalline region that requires full resolution). All remeshing is performed in the reference configuration (actual atomic positions $\boldsymbol{q}^h$ are obtained via the interpolation (2.3). This guarantees that new repatoms are inserted at physically reasonable locations and that the number of repatoms N_h cannot exceed the number of discrete lattice sites N (i.e., full refinement implies $\mathscr{N}_h = \mathscr{B}$). We thus require that each new node $N_h + 1$ be placed on a valid lattice site in the respective lattice, i.e., $\boldsymbol{X}_{N_h+1} \in \mathscr{B}$. Meshing on a discrete point set creates a challenge because – unlike in FE codes – standard techniques like mesh smoothing are not applicable and new nodes may

not be inserted anywhere in space for mesh improvement (see, e.g., [13, 52] for FE examples). Especially when approaching atomistic resolution, a proper mapping of ideal new nodal locations onto members of $\mathscr{B}$ is essential to avoid inverted or badly-shaped elements.

Here, we pursue a simple refinement strategy based on longest-edge bisection. Elements requiring refinement are bisected by the median of the longest edge, which is a common strategy used in refinement algorithms. If two or more edges are similar in length, all are bisected. For each considered edge, we identify (in the reference configuration) the lattice site $\boldsymbol{X} \in \mathscr{B}$ nearest to the edge's midpoint that is not already a node, which will be added as a new repatom. In case of several equidistant sites, we make a random selection. This procedure is applied to all elements requiring refinement in arbitrary order, resulting in a list $\mathscr{N}_h^*(t_n)$ of repatom positions in the reference configuration that must be added to the model, so that $\mathscr{N}_h(t_{n+1}) = \mathscr{N}_h(t_n) \cup \mathscr{N}_h^*(t_n)$ at the end of each completed time step. Next, the mesh $\mathscr{T}$ is updated based on the new set of repatoms $\mathscr{N}_h(t_{n+1})$ using the open-source Computational Geometry Algorithms Library (CGAL) [116]; the same library is used for generating the initial mesh. Finally, new sampling atoms (required for all newly-created elements) and their local atomic neighborhoods $\mathscr{I}_\alpha$ (required for all new sampling atoms and those whose element connectivity has changed) are identified in the reference configuration, as shown in Fig. 3.1, and their current positions and moments are obtained from (2.3).

3.1.4 Example

The above refinement algorithm for the nonlocal QC method is illustrated in Fig. 3.2 for the example of nano-void growth and coalescence. Two voids of initial radius 4 nm are placed at a distance of 4 nm in a copper single-crystal modeled by the Finnis-Sinclair potential of [29] at zero temperature. The body (constrained to undergo in-plane atomic motion) is loaded uniaxially in the vertical direction at a loading rate of $2.5 \cdot 10^{11}$ s^{-1} in tension. The snapshots show the gradual mesh refinement with increasing deformation. Dislocations nucleate at the void surfaces and propagate into the crystal on the three primary slip systems (in 2D), so that remeshing accumulates atomistic resolution in the vicinity of these moving lattice defects. This leads to localized regions of full atomistic resolution, leaving the remainder of the crystal coarse-grained.

We note that the shown mesh loses its physical meaning in highly-deformed re-

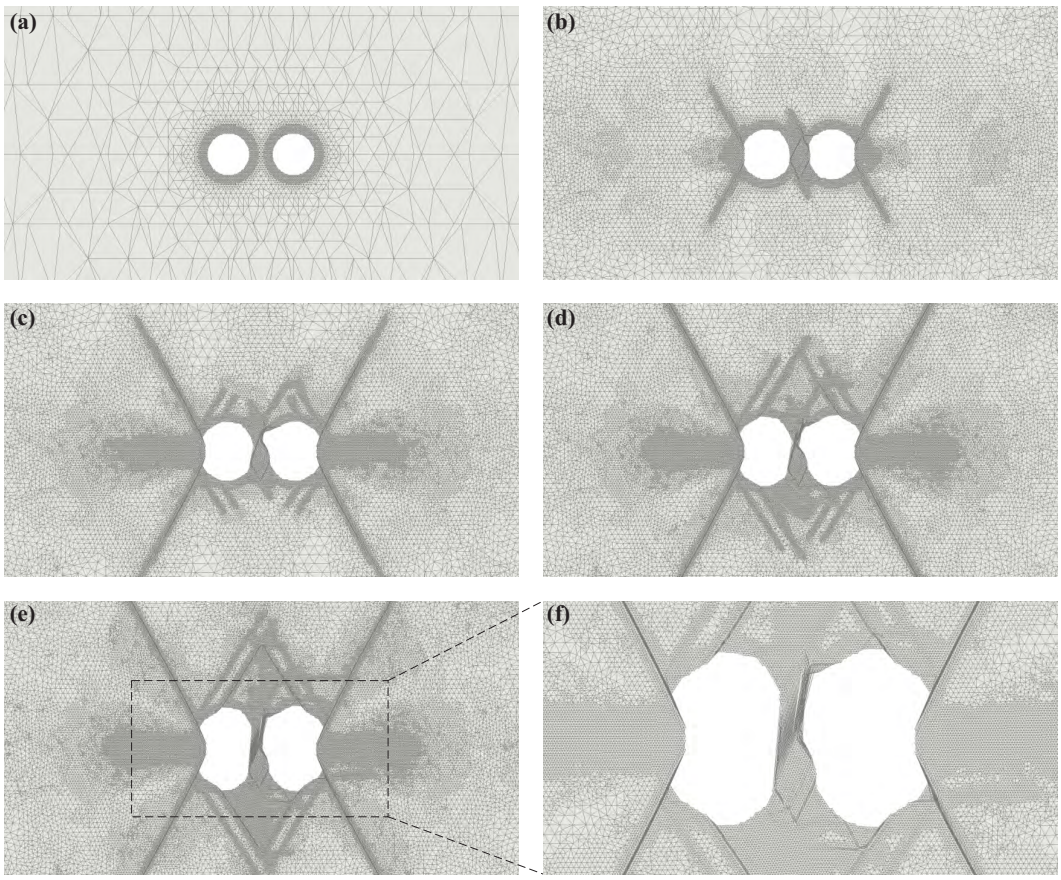

Figure 3.2: Application of the remeshing scheme to a Cu single-crystal containing two nano-voids under uniaxial extension. Shown is the reference mesh mapped into the deformed configuration at applied vertical strains of (a) 0%, (b) 4%, (c) 5%, (d) 6%, and (e) 7% with a magnified view shown in (f)

gions (shown are the element connectivities of the reference mesh mapped into the deformed final configuration, i.e., elements link atoms that were initially close but may have drifted apart). For example, the two voids in Fig. 3.2(c-f) have long coalesced, but drawing the connectivities of the reference mesh seemingly keeps a spurious ligament of elements between the voids. This becomes obvious when inspecting the actual repatom locations (to be discussed in Fig. 3.7). The mesh's only purpose is thus to enable the recovery of all underlying lattice sites, as required, e.g., for new repatoms, sampling atoms, and lattice neighbors. All equations of motion are solved in the current configuration with current neighbors.

3.2 Atomic Neighborhoods and Update Requirements

Fig. 3.2 showed an example of how the element connectivities (and thus the atomic neighborhoods) established in the reference mesh lose their physical meaning dur-

ing large atomic rearrangements, particularly because locality in the reference configuration may deviate significantly from locality in the current configuration, like in any particle method. However, eq. (2.15) requires an accurate account of all lattice sites within $\mathscr{I}_\alpha$, i.e., within the radius of interaction of each sampling atom, and the sets $\mathscr{I}_\alpha$ ($\alpha = 1, \ldots, N_s$) can change considerably during deformation. Therefore, it is imperative to keep track of all neighbors entering any given sampling atom's radius of interaction (especially those that are not close to the sampling atom in the reference configuration). An alternative would be an updated-Lagrangian description, see, e.g., [62]; that, however, may lead to incorrect atomistic ensembles and break the basis of the QC approximation (viz., that atomic positions within elements are uniquely defined by a deformed lattice); see Section 3.3. Therefore, we choose the Lagrangian description with a reference mesh (as in all prior QC versions except for meshless max-ent QC [55]) and introduce a scheme to identify atomic neighborhoods based on locality in the current configuration.

3.2.1 Neighborhood Update Scheme and Frequency

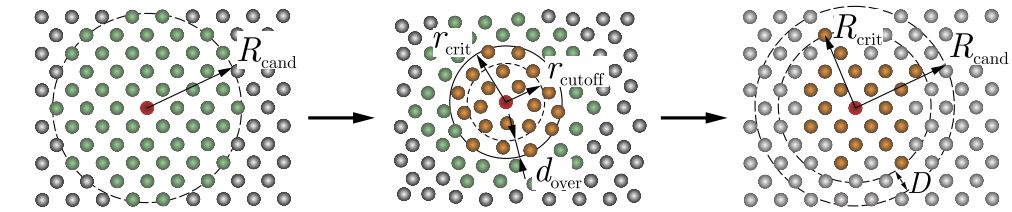

Figure 3.3: Schematic illustration of neighborhood updates: identification of candidate neighbors (green) around a sampling atom α (red) within distance $R_{\text{cand},\alpha}$ in the reference configuration (set $\mathscr{C}'_\alpha$); identification of actual neighbors (orange) within radius r_{crit} in the current configuration (set $\mathscr{C}_\alpha$); and the complete actual neighborhood mapped back into the reference configuration (orange) with the maximum distance $R_{\text{crit},\alpha}$ used to update $R_{\text{cand},\alpha}$.

In MD, up-to-date atomic neighborhoods are obtained by performing frequent neighbor searches among all atoms in the simulation. In the QC formulation, no such neighbor searches can be performed directly because we do not keep track of the positions of all atoms in the simulation (this would defeat the purpose of QC). We therefore generate and store only the local atomic neighborhoods $\mathscr{C}_\alpha(t)$ for all sampling atoms (both in atomistic and coarse-grained regions), analogous to Verlet lists in MD [122]. Specifically, each sampling atom stores only those lattice sites whose positions at time t are located within a critical radius $r_{\text{crit}} \geq r_{\text{cutoff}}$ in the *current* configuration (with r_{cutoff} the cutoff radius of the atomic potential V, see Fig. 3.3),

i.e.,

$$\mathscr{C}_\alpha(t) = \left\{ \boldsymbol{X} \in \mathscr{B} \;\middle|\; \left\| \boldsymbol{q}^h(\boldsymbol{X},t) - \boldsymbol{x}_\alpha(t) \right\| \leq r_{\text{crit}} \right\} \quad \text{with} \quad \boldsymbol{q}^h(\boldsymbol{X},t) = \sum_{\beta \in \mathscr{N}_h} N_\beta(\boldsymbol{X}) \boldsymbol{x}_\beta(t). \tag{3.3}$$

Choosing $r_{\text{crit}} = r_{\text{cutoff}}$ yields the minimally required neighborhood sets $\mathscr{I}_\alpha(t)$. However, as repatom positions $\boldsymbol{x}(t)$ change continuously, that choice would necessitate continuous updates of the neighborhood sets. To reduce the frequency of neighborhood updates, r_{crit} is chosen sufficiently large so that $d_{\text{over}} = r_{\text{crit}} - r_{\text{cutoff}} > n\, v_{\max}\, \Delta t$, where n is the number of time increments Δt after which a neighborhood update is anticipated, and $v_{\max}$ denotes the maximum atomic velocity within $\mathscr{C}_\alpha(t)$. Numerical experiments have shown that $d_{\text{over}} = 0.1\text{Å}$ is sufficient for numerical convergence for all subsequent examples. Note that we will not fix n but instead define a motion-based automatic trigger mechanism for neighborhood rebuilds.

In order to identify the neighborhood sets $\mathscr{C}_\alpha(t) \in \mathscr{B}$ at time t without knowing the current positions of all lattice sites, we generate candidate neighborhoods

$$\mathscr{C}'_\alpha(t) = \left\{ \boldsymbol{X} \in \mathscr{B} \;\middle|\; \|\boldsymbol{X} - \boldsymbol{X}_\alpha\| \leq R_{\text{cand},\alpha}(t) \right\} \tag{3.4}$$

and then find $\mathscr{C}_\alpha(t) \subset \mathscr{C}'_\alpha(t)$ for $\alpha = 1, \ldots, N_s$. To find the complete neighborhood we must choose $R_{\text{cand},\alpha}(t)$ in the reference configuration such that

$$R_{\text{cand},\alpha}(t) \geq R_{\text{crit},\alpha}(t) \qquad \text{with}$$
$$R_{\text{crit},\alpha}(t) = \max_{\boldsymbol{X} \in \mathscr{B}} \left\{ \|\boldsymbol{X} - \boldsymbol{X}_\alpha\| \;\middle|\; \left\| \boldsymbol{q}^h(\boldsymbol{X},t) - \boldsymbol{x}_\alpha(t) \right\| \leq r_{\text{crit}} \right\}, \tag{3.5}$$

i.e., $R_{\text{crit},\alpha}(t)$ is the radius that includes all neighbors in the *reference* configuration which at time t have entered radius r_{crit} of sampling atom α in the *current* configuration; see Fig. 3.3 for a schematic view. For efficiency, we assign a unique $R_{\text{cand},\alpha}(t)$ to each sampling atom. Specifically, we choose $R_{\text{cand},\alpha}(t) = R_{\text{crit},\alpha}(t) + D$, where $D \geq 0$ is an overshoot parameter. Our numerical examples use $D \approx 2.3 r_{\text{cutoff}}$ as an inexpensive, conservative choice. Note that such a choice hardly affects the computational costs because it does not effect the expenses of time stepping; it only affects the neighborhood generation, which is rare compared to force calculations and position updates. Since r_{crit} is defined in the current configuration, $R_{\text{crit},\alpha}(t)$ and $R_{\text{cand},\alpha}(t)$ in the reference configuration change with time (shown in the examples below).

Efficiency demands that neighborhoods are not updated at each time step but only as required. To this end, we do not define a fixed update frequency but assume that a new lattice site can move into the neighborhood $\mathscr{C}_\alpha(t)$ of sampling atom α only if the already existing atomic neighbors in $\mathscr{C}_\alpha(t)$ have experienced sufficient movement since the last update at time $t^* < t$, i.e., if

$$\exists\, \boldsymbol{X}_i \in \mathscr{C}_\alpha(t)\ s.t.\ \left\| \boldsymbol{q}^h(\boldsymbol{X}_i,t) - \boldsymbol{q}^h(\boldsymbol{X}_i,t^*) \right\| > q_{\max}. \tag{3.6}$$

Our choice $d_{\text{over}} = r_{\text{def}} - r_{\text{cutoff}} > 0$ makes sure that we keep track of more neighbors than strictly required; we monitor their relative movements to estimate when an update is required, and perform the neighborhood update when any neighbor's movement exceeds the tolerance q_{max}. Numerical experiments indicate that efficient convergence is achieved if $q_{max} \approx 3d_{\text{over}}$ for time steps in the range of $\Delta t = 1-5$ fs. We note that because neighborhood updates are triggered based on relative motion of a sampling atom's current neighbors, in the highly unlikely event that a distant atom approaches without sufficiently displacing the current neighbors, the present algorithm would fail [3]. However, after each neighborhood rebuild the code checks whether or not new neighbors had to be added without moving the existing ones, and for all of the following examples the check never failed. This neighborhood update scheme applies equally to 2D and 3D.

Algorithm 2 Neighborhood update algorithm

for each sampling atom $\alpha = 1,\ldots,N_s$ **do**
  if update criterion (3.6) is violated for sampling atom α **then**
    $R_{\text{cand},\alpha}(t) = 1.01 \times \left[R_{\text{crit},\alpha}(t) + D\right]$
    identify candidate neighborhood $\mathscr{C}'_\alpha(t)$ according to (3.4) in the *reference configuration*
    for each candidate neighbor $\boldsymbol{X}_\beta \in \mathscr{C}'_\alpha(t)$ **do**
      map candidate to the current configuration: $\boldsymbol{q}^h_\beta = \sum_i \boldsymbol{x}_i N_i(\boldsymbol{X}_\beta)$
      if $\left\| \boldsymbol{q}^h_\beta(t) - \boldsymbol{x}_\alpha(t) \right\| \le r_{\text{crit}} = r_{\text{cutoff}} + d_{\text{over}}$ in the
      deformed configuration **then**
        add candidate $\boldsymbol{X}_\beta$ to the neighborhood $\mathscr{C}_\alpha(t)$
      end if
    end for
    update $R_{\text{crit},\alpha}(t) = \max_\beta \left\| \boldsymbol{X}_\beta - \boldsymbol{X}_\alpha \right\|$
  end if
end for

3.2.2 Examples

We present two examples of the neighborhood updates which demonstrate significant neighborhood changes due to inelastic deformation and defect movement in the atomistic region. First, we revisit the void growth example already discussed in Section 3.1 (two identical nano-voids in a Cu single-crystal are loaded in uniaxial extension up to coalescence). Second, we simulate nanoindentation by a circular indenter of radius 5 nm penetrating into a Cu single-crystal [29] in 2D (at $2 \cdot 10^4$ m/s), using an indenter potential [16]. The potential's cutoff radius r_{cutoff} is such that each atom interacts with only six nearest neighbors in a defect-free 2D single-crystal. For each of the two scenarios, we select a few illustrative sampling atoms and follow their atomic neighborhoods during the course of the simulation.

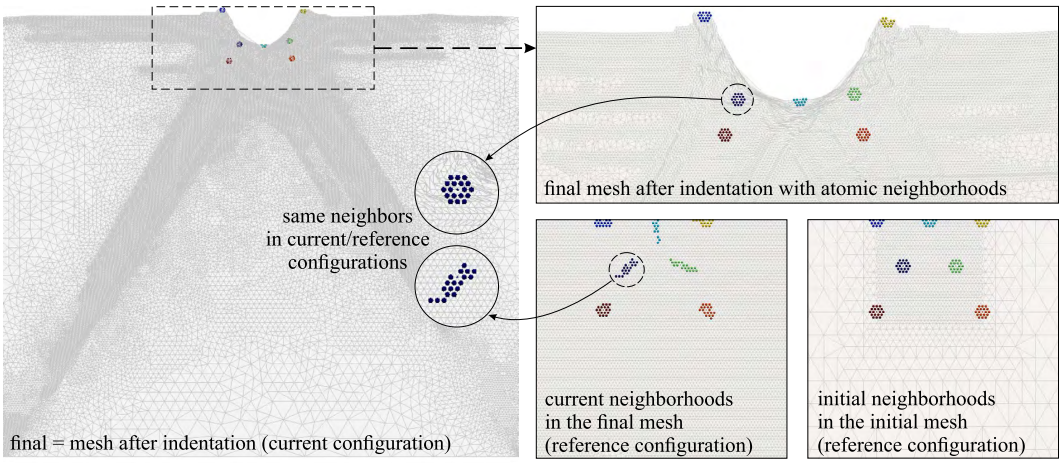

Figure 3.4: Illustration of the atomic neighborhoods of seven selected sampling atoms during nanoindentation shown in the current configuration as well as mapped back into the reference configuration; for comparison the atomic neighborhoods in the initial state (before indentation) are also shown. The magnification illustrates the same final atomic neighborhood shown both in the current and in the reference configurations.

For the indentation scenario, Fig. 3.4 shows seven sampling atoms with their respective atomic neighborhoods in the current configuration and the same atoms in the reference configuration. It becomes apparent that, with increasing plastic deformation and associated dislocation activity, neighborhoods change significantly and the distances $R_{\text{crit},\alpha}(t)$ in the reference configuration increase for all shown sampling atoms. Similarly, Fig. 3.5 shows the changes of seven sampling atom neighborhoods for the void growth example of Fig. 3.2. While some of the atomic neighborhoods remain fairly localized, others (e.g., those in the ligament between

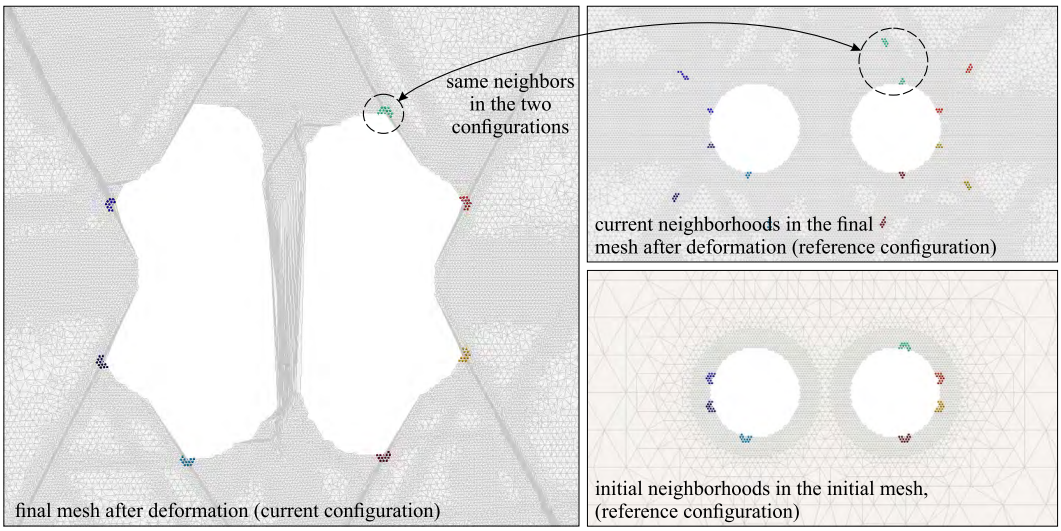

Figure 3.5: Illustration of the atomic neighborhoods of seven selected sampling atoms during void growth and coalescence shown in the current configuration as well as mapped back into the reference configuration; for comparison the atomic neighborhoods in the initial state (before extension) are also shown.

the two voids) experience considerable neighborhood changes, resulting in large distances between current neighbors when viewed in the reference mesh. The chosen update scheme accounts for this increase by gradually increasing $R_{\text{crit},\alpha}(t)$ and updating the associated neighbors of all sampling atoms.

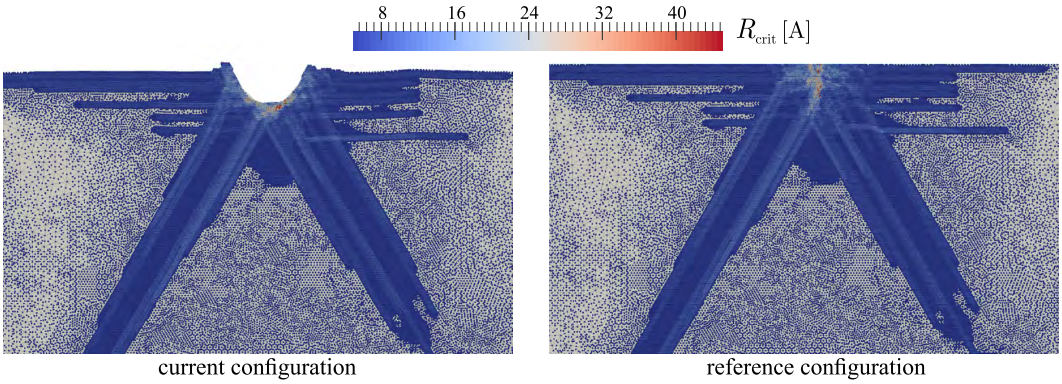

Figure 3.6: Distribution of $R_{\text{crit},\alpha}$ (in angstroms) across all sampling atoms at an indentation depth of 8 nm, shown in the deformed configuration (*left*) and mapped into the reference configuration (*right*). Note that sampling atoms are color-coded and elements left color-less (because $R_{\text{crit},\alpha}$ is defined on sampling atoms).

Fig. 3.6 summarizes the distribution of R_{crit} for all sampling atoms at an indentation depth of 8 nm, shown both in the current configuration and in the reference configuration. Obviously, large fractions of atoms right underneath the indenter un-

dergo significant plastic deformation and participate in defect nucleation and motion. Consequently, atomic neighborhoods change and R_{crit} increases considerably. As can be expected, atomic neighborhoods away from the indentation region remain intact and require only minor increases of R_{crit}. Fig. 3.7 shows the analogous map of R_{crit} for all sampling atoms in the void growth simulation. Although one might be concerned about the extra neighbors incurred by the overshoot distance d_{over}, in this simulation with the potential of [29] the average number of neighbors within interaction radius r_{cutoff} is 14, whereas the average number of neighbors within radius $r_{\text{crit}} = r_{\text{cutoff}} + d_{\text{over}}$ is 18; therefore, the extra atoms considered within each neighborhood amount to about 29%.

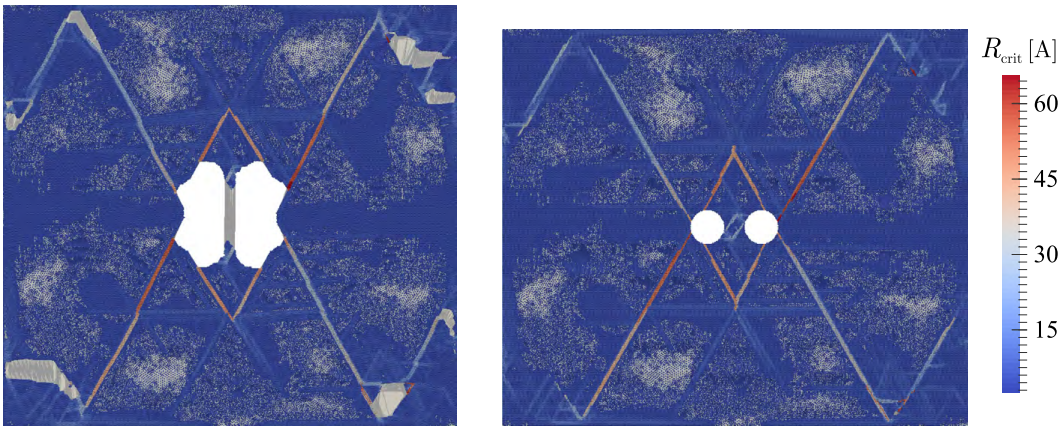

Figure 3.7: Distribution of $R_{\text{crit},\alpha}$ (in angstroms) for all sampling atoms in a void growth simulation at an applied strain of $2.5 \cdot 10^9$ s^{-1} shown in the current configuration (*left*) and mapped into the reference configuration (*right*).

Note that since each sampling atom is assigned an individual radius R_{crit}, we only enlarge the search radius of those sampling atoms that indeed participate in large neighborhood distortions. As shown, e.g., in Fig. 3.6, most of the sampling atoms retain the initial value of $R_{crit} = 4.95\mathring{A}$ throughout the simulation, as they are sufficiently far away from defects. The number of those sampling atoms whose R_{crit} grows significantly (e.g., up to $50\mathring{A}$) is shown to be fairly small. Out of the total of $600,680$ sampling atoms in this simulation, 74% retain $R_{crit} = 4.95\mathring{A}$ until the end of the simulation, 17.7% have $4.96 \leq R_{crit} \leq 7.55$, another 7.3% have $7.55 \leq R_{crit} \leq 12$, and only 1% of all sampling atoms reach a value of $R_{crit} > 12$. While neighborhood updates do add computational expenses, we deliberately chose to vary R_{crit} from atom to atom to keep the added costs as low as possible. As a matter of fact, if defects remain localized as shown here, then costs are comparable

to standard Verlet algorithms used in atomistics where the neighbors of all atoms are frequently updated.

This example illustrates not only how dislocation motion results in large local neighborhood changes but it also demonstrates how the reference mesh loses physical meaning, especially in the ligament between the two voids (showing the mesh in the current configuration completely misses the fact that the voids have long coalesced). R_{crit} serves as a metric for the local deformation history. The centrosymmetry parameter [16], e.g., returns to 0 after a dislocation has passed in 2D (in 3D a stacking fault can be left behind), although atomic neighborhoods have changed significantly, which is reflected by the R_{crit} map. This is also why the neighborhood update scheme described above is essential for the accuracy of fully-nonlocal QC simulations.

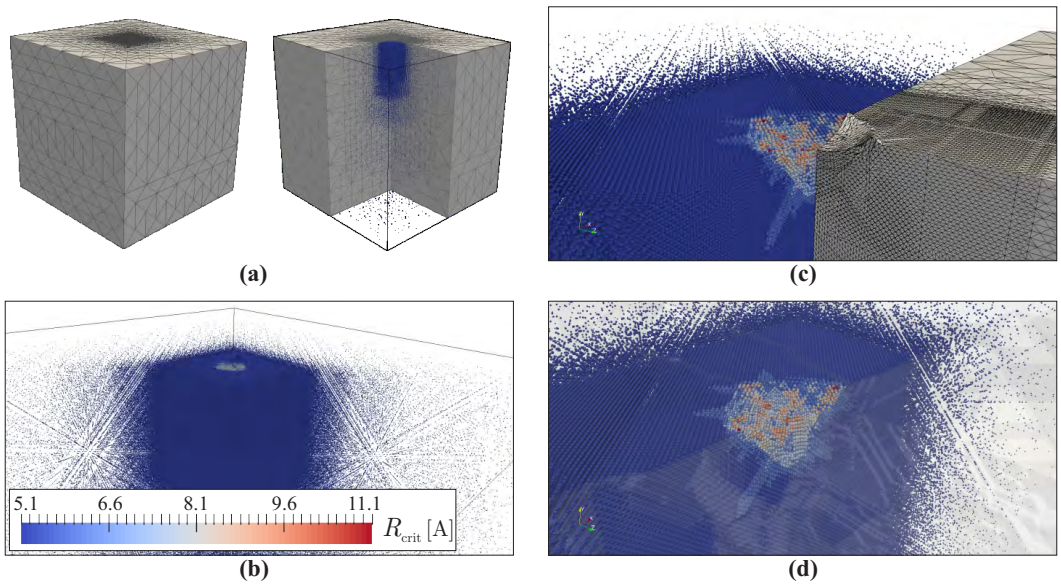

Figure 3.8: 3D indentation example using a spherical nanoindenter (diameter 10 nm, shown at a depth of 1.35 nm) into single-crystalline copper (using the neighborhood update scheme but avoiding mesh refinement). Shown are (a) the undeformed QC mesh and sampling atom distribution, (b) the deformed distribution of sampling atoms color-coded by R_{crit}, and (c) and (d) magnified views of the deformed mesh and sampling atoms with R_{crit} in the indentation region.

The last example, shown in Fig. 3.8, presents results of a 3D nanoindentation simulation. Due to the aforementioned challenges with geometric mesh refinement in 3D, we chose a configuration with a sufficiently large initial atomistic region so no refinement is necessary (the domain is coarse-grained in five steps by coarsening factors of 2). The described neighborhood update scheme applies equally in 3D, re-

sulting in the shown distribution of R_{crit}. As before, notice that the number of sampling atoms with significant neighborhood increases (large values of R_{crit}) is fairly small. The simulation contains ca. $1.17 \cdot 10^6$ repatoms and 1.75 million sampling atoms, compared to $2.26 \cdot 10^8$ atoms in the fully-atomistic case. As another benefit of the QC coarse-graining, boundary conditions (here, allowing only in-plane motion on the vertical faces with the bottom fixed) can be applied at significant distance from the indentation side without any artificial effects on the simulation outcomes in the atomistic region (and without the need for, e.g., periodic boundary conditions).

3.2.3 Benchmark Tests

The examples of nanoindentation and void growth are ideal candidates for the adaptive, fully-nonlocal QC method, since atomistic resolution can be restricted to small regions of interest. We note that the optimal second-order summation rule used here captures surface relaxation effects in an approximate sense and therefore does not require full resolution on free surfaces [5]. Let us revisit the examples of indentation and void growth, using the adaptive remeshing and neighborhood update schemes outlined in previous sections, and investigate results with respect to qualitative and quantitative errors.

Figs. 3.9 and 3.10 demonstrate the defect distribution in the two deformed Cu single-crystals undergoing nanoindentation and void growth during uniaxial extension, respectively. Edge dislocations spreading from the indentation site and from the voids, respectively, into the uniform crystals on the three primary slip systems (in the 2D close-packed crystal) are visible. Moving defects carry fully-resolved mesh fronts into initially coarse-grained regions, resulting in the shown QC representations. Dislocation slip systems, core sizes, and distributions are comparable to MD results using the same parameters. Note that white regions simply indicate coarse-grained regions that are by definition defect-free and undergo affine deformation according to the chosen interpolation scheme. The overall number of representative atoms in the final deformed state shown is considerably lower than the total number of atoms that would have to be simulated in a comparable fully-atomistic simulation (we present a quantitative comparison later).

Of course, these results encourage the development of a mesh coarsening algorithm, so that fully-refined meshes can be coarsened once a defect has passed through (viz., all blue regions in Figs. 3.9 and 3.10). However, as discussed in Section 3.3.3, mesh

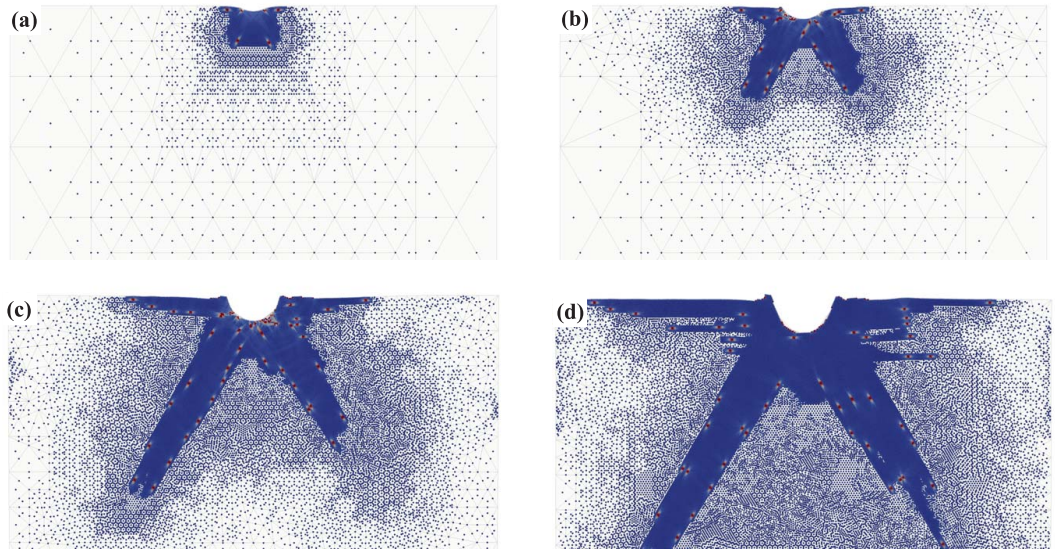

Figure 3.9: Visualization of the defect distribution during nanoindentation (at a loading rate of $2 \cdot 10^4$ m/s); shown is the centrosymmetry parameter on an arbitrary scale, highlighting dislocations on the three primary slip systems underneath the indenter at indentation depths of (a) 0.8 nm, (b) 2 nm, (c) 4 nm, and (d) 6.8 nm. The shown centrosymmetry parameter visualizes dislocations in red, whereas a pristine single crystal appears in dark blue.

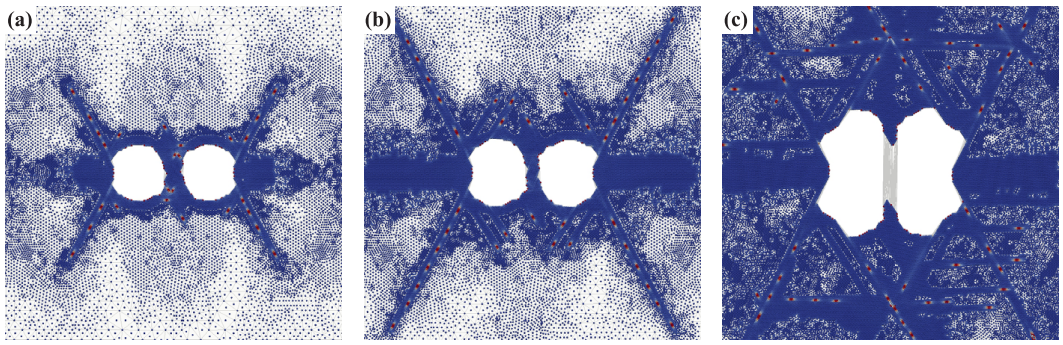

Figure 3.10: Visualization of the defect distribution during void growth and coalescence (at a loading rate of $2.5 \cdot 10^9$ s^{-1}); shown is the centrosymmetry parameter on an arbitrary scale, highlighting dislocations around the voids at vertical strains of (a) 4%, (b) 5%, and (c) 9%. The shown centrosymmetry parameter visualizes dislocations in red, whereas a pristine single crystal appears in dark blue.

coarsening is a challenge in itself due to, among other factors, the need to identify a unique Bravais lattice for each element (including elements that are created by merging existing smaller elements with different underlying deformed lattices).

Fig. 3.11 offers a quantitative error estimate based on a one-to-one comparison with an MD calculation. We perform the above void growth simulation using the fully-

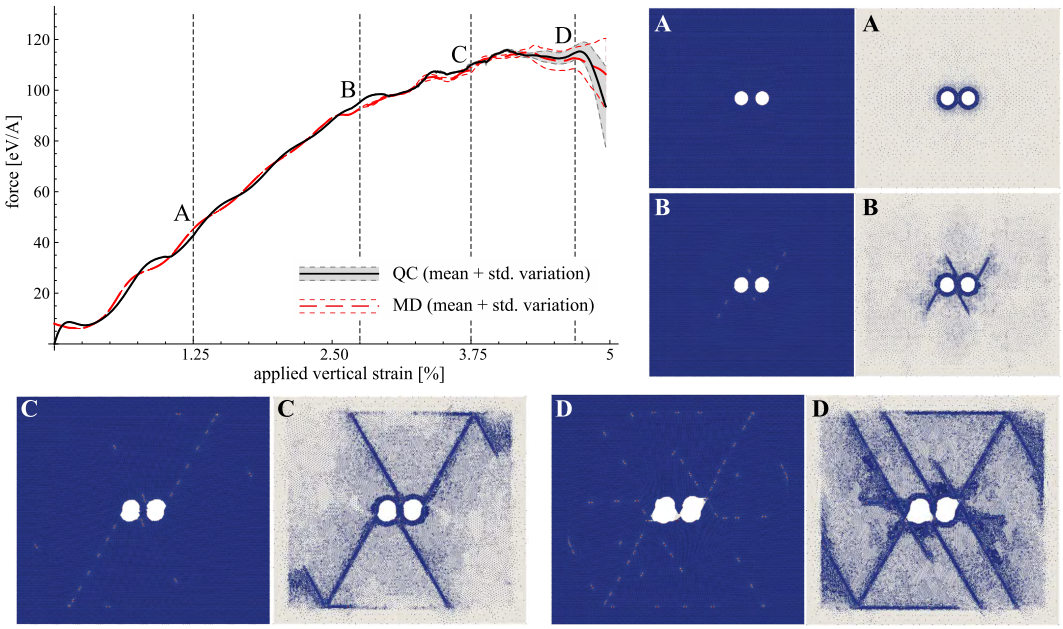

Figure 3.11: Void growth and coalescence simulated by both nonlocal QC (*right images*) and MD (*left images*) at the four strain levels **A**–**D** marked in the total force vs. strain plots for a vertical extension test (atoms or sampling atoms are color-coded by centrosymmetry on an arbitrary scale to visualize lattice defects, primarily dislocations). Curves indicate the average and standard deviations for a total of 20 QC and 20 MD simulations (strain rate is $2.5 \cdot 10^{10}$ s^{-1}).

nonlocal QC method and full atomistics (using the same time step size of 5 fs, the same boundary and initial conditions, and the same dynamic solver). Fig. 3.11 compares both the defect distributions and the total force vs. vertical extension curves. Due to floating point errors, exact results are generally not reproducible necause of the high symmetry of the problem and resulting equivalent energy minimizers. For a reasonable comparison, the shown curves were thus obtained by repeating each simulation 20 times and averaging the results (shown are both the mean values and standard deviations of QC and MD); see also the discussion in [5]. The differences up to about 3.25% strain can be attributed mainly to dynamic effects: unlike in MD, elastic waves in the QC sample cannot be represented exactly by the non-uniform QC mesh. However, deviations are small and the mean force–strain curves agree well into the inelastic regimes exhibiting dislocation plasticity and void growth and coalescence. The centrosymmetry plots confirm that defect distributions and deformed void shapes are also in qualitative agreement (we note that the unrefined boundary layers are due to the application of boundary conditions).

The remeshing and neighborhood update scheme introduces a few parameters that

are expected to affect the accuracy and efficiency of simulations. First, the remeshing criterion (3.1) introduces a tolerance ε. Second, neighborhood updates introduce the overshoot radius $d_{\text{over}} = r_{\text{crit}} - r_{\text{cutoff}}$ for the current Verlet neighborhoods, and the overshoot radius $D = R_{\text{cand}} - R_{\text{crit}}$ for the generation of candidate neighborhoods in the reference configuration, see Fig. 3.3. Third, neighborhoods are updated if atomic movements exceed the threshold q_{max}. Finally, time step Δt affects the simulated results. Of those, the chosen overshoot radii ($d_{\text{over}} = 0.1$Å, $D \approx 6 - 10$Å) and update threshold ($q_{max} = 0.3$Å) were verified to be sufficient by – on the fly – considering larger neighborhoods and checking whether or not neighboring atoms were missed at any time during the simulation, which was never the case. Once a time step Δt is selected for a particular simulation, the only adjustable parameter is then the remeshing tolerance ε, which affects the frequency of remeshing (obviously, the frequency of remeshing increases with decreasing ε).

The influence of ε is demonstrated in Fig. 3.12, which repeats the void growth simulation of Fig. 3.11 using six different values of ε (for a time step of $\Delta t = 4$ fs). Although no clear trend emerges, the low thresholds $\varepsilon = 0.1 - 0.2$ indicate lower errors than larger values of ε, as can be expected. Further, notice the spikes appearing in the QC simulation at increasing strains with increasing ε, indicating a remeshing event that causes instantaneous relaxation (those spikes have only marginal influence on the overall response). For comparison, we have included in Fig. 3.12 the time evolution of the total number of repatoms in the QC simulation (normalized by the number of atoms in the MD simulation). It becomes apparent that the spikes correspond to early remeshing events, and that the number of repatoms decreases with increasing remeshing tolerance ε (as could have been expected).

It is difficult to derive quantitative rules for the selection of ,ε but for applications in metal plasticity – one can state quantitative guidelines. We followed the strategy of Knap and Ortiz [54] and Tadmor, Ortiz, and Phillips [111], who chose the second invariant of the deformation gradient as remeshing indicator, as in (3.1). Then, $\sqrt{I_e} \cdot L_e$ is compared to the length b of the smallest Burgers' vector of the crystal lattice (i.e., we compare the maximum shear displacement in the element to the slip distance required to form a dislocation). In order to trigger dislocation motion appropriately such that dislocation motion is not inhibited in a coarse mesh, remeshing must be triggered before the maximum shear distance reaches b, and thus we choose fractions of b as the remeshing tolerance and therefore $\varepsilon < 1$. The specific values chosen here for ε were the result of numerical experiments. Results in Fig. 3.12

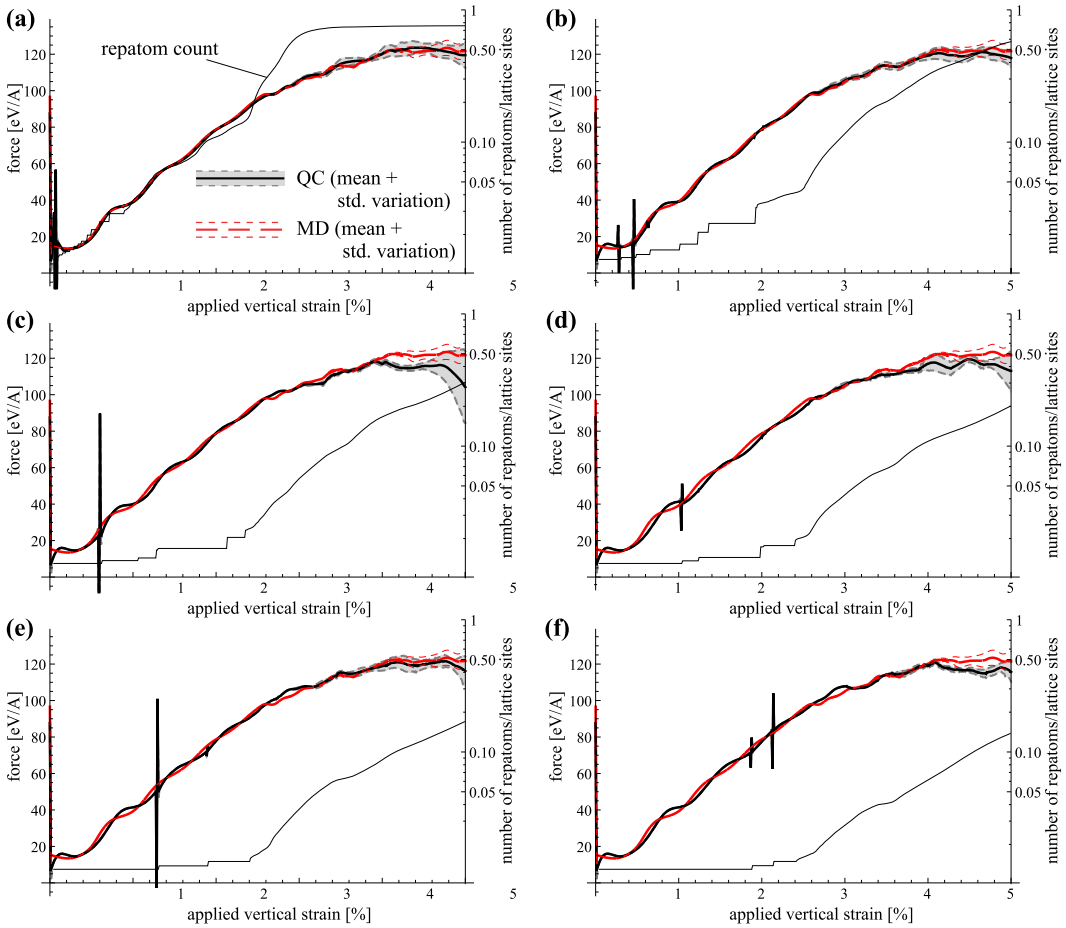

Figure 3.12: Variation of the mechanical response for the void growth example for different values of the remeshing tolerance ε: (a) $\varepsilon = 0.1$, (b) $\varepsilon = 0.2$, (c) $\varepsilon = 0.3$, (d) $\varepsilon = 0.4$, (e) $\varepsilon = 0.5$, and (f) $\varepsilon = 0.6$. The thin solid lines indicate the fraction of all lattice sites used as repatoms (on the right axes).

show that the chosen values yield results comparable to full atomistics, which reinforces that the chosen values provide a good balance of accuracy and efficiency (decreasing ε below the chosen values severely increases costs but hardly affects results as shown in Fig. 3.12; increasing ε reduces costs but leads to large errors). For comparison, the force-based QC method of [54] used significantly smaller values for ε (at least a factor of 10 smaller); however, we chose the longest element edge L_e as the characteristic element length and do not consider the simplex volume as in [54]. Ultimately, the choice of ε depends on the application and is a compromise between accuracy and efficiency.

To define a simple error metric, Fig. 3.13(a) reports the average relative errors of

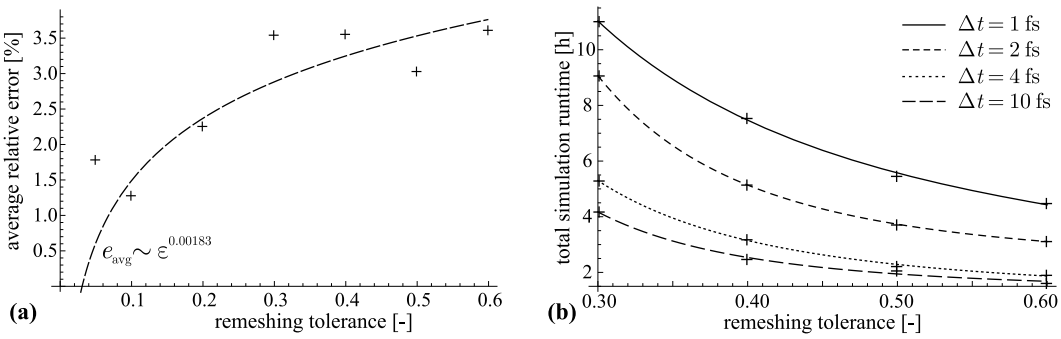

Figure 3.13: (a) Total relative error e_{avg} of the load–strain curve of the void growth example vs. the remeshing tolerance ε (with a logarithmic regression fit); (b) total simulation run time vs. remeshing tolerance ε for various time steps Δt with fitted curves.

each load–strain curve comprising n_{step} load steps, i.e.,

$$e_{\text{avg}} = \frac{1}{n_{\text{step}}} \left\| \frac{F_{\text{QC}} - F_{\text{MD}}}{F_{\text{MD}}} \right\|, \tag{3.7}$$

where F is the total vertical force applied to the sample in either QC or MD simulation. The results confirm that low values of $\varepsilon \leq 0.2$ result in low average errors (below 2.5%), but even larger values of ε keep the total average error below 3.8%.

It is worth pointing out that the agreement between QC and MD results is remarkable for the following reasons. Not only does the QC simulation arrive at qualitatively and quantitatively comparable results in terms of the microstructural defect distribution as well as in terms of the macroscopic stress-strain response. QC also performs well in capturing the elastodynamic response of the sample loaded dynamically at a high rate (comparable to typical MD simulations) despite the significantly coarsened and non-uniform mesh. The chosen remeshing and neighborhood update protocol along with the fully-nonlocal QC formulation enables us to tie full atomistic resolution to defects and thereby reduce computational expenses. In regions that are fully-atomistically resolved, the presented adaptive nonlocal QC method recovers molecular dynamics exactly, i.e., one can expect MD and QC results to agree locally in regions of full resolution.

Increasing ε implies less remeshing and thus is expected to decrease the simulation run time (as seen in the number of repatoms visualized in Fig. 3.12). As a further confirmation, Fig. 3.13(b) compares the average run times of QC simulations for the shown values of ε and Δt. Obviously, the simulation time increases with decreasing time steps for the same applied strain rate. Run times scale as $t_{\text{run}} \propto \varepsilon^c$ (with c

ε \ Δt	neighborhood update computational cost (% of total simulation time)				mesh refinement computational cost (% of total simulation time)				QC time / fully-refined time
	10 fs	**5 fs**	**4 fs**	**2 fs**	**10 fs**	**5 fs**	**4 fs**	**2 fs**	**4 fs**
0.2	35.14	37.18	36.51	35.38	5.48	4.57	4.70	4.20	3.21
0.3	33.35	33.53	34.69	33.04	7.86	6.44	6.24	5.34	1.96
0.5	32.59	32.37	31.58	29.51	11.02	9.55	9.83	8.08	0.80
0.6	33.39	32.10	31.78	28.64	12.29	10.34	10.60	8.96	0.47
0.7	32.72	32.66	31.61	28.67	15.00	11.43	11.51	10.20	0.39

Table 3.1: Overview of computational costs for the void growth simulation: percentages of total simulation time spent on neighborhood updates and mesh refinement, and a comparison of the QC simulation time and the analogous time spent on a fully-refined model for different values of the refinement tolerance ε and the time step Δt.

monotonically decreasing from $c = -1.47$ for $\Delta t = 10$ fs to $c = -2.98$ for $\Delta t = 1$ fs). Table 3.1 (right column) shows the efficiency of simulations by computing the ratio of the QC run times and the time spent on a comparable fully-refined simulation for the example case of $\Delta t = 4$ fs. For low refinement tolerance (e.g., $\varepsilon = 0.2$), the QC simulation is more than three times as expensive than the fully-refined calculation due to adaptivity overhead. By contrast, higher refinement tolerances (e.g., $\varepsilon = 0.7$) lead to more efficient simulations that take less than 40% of the fully-refined time (note that significantly higher efficiency can be gained in more aggressively coarse-grained systems). Table 3.1 also includes the relative costs of neighborhood updates and mesh refinement (as fractions of the total simulation time). With increasing refinement tolerance ε, the relative refinement costs increase (mainly because the overall simulation time decreases at steady refinement costs). By contrast, the relative cost associated with neighborhood updates shows little variation (since these are performed independently of refinement, and their cost is proportional to the number of sampling atoms in the simulation as is the overall cost).

3.3 Alternative routes to mesh refinement

3.3.1 Discrete lattice site constraint

The nonlocal QC method outlined in Chapter 2 places all repatoms onto atomic sites of the underlying lattice, which presents a strong constraint. Some competing schemes relax this constraint in the coarse-grained region to reduce complexity (e.g., to avoid the geometric refinement difficulties mentioned above), see Fig. 3.14(a) and (b) for a schematic comparison. Unfortunately, relaxing the constraint necessitates a small, well-defined transition region between atomistic and coarse-grained domains, in which repatoms lie on lattice sites (so that atoms in the

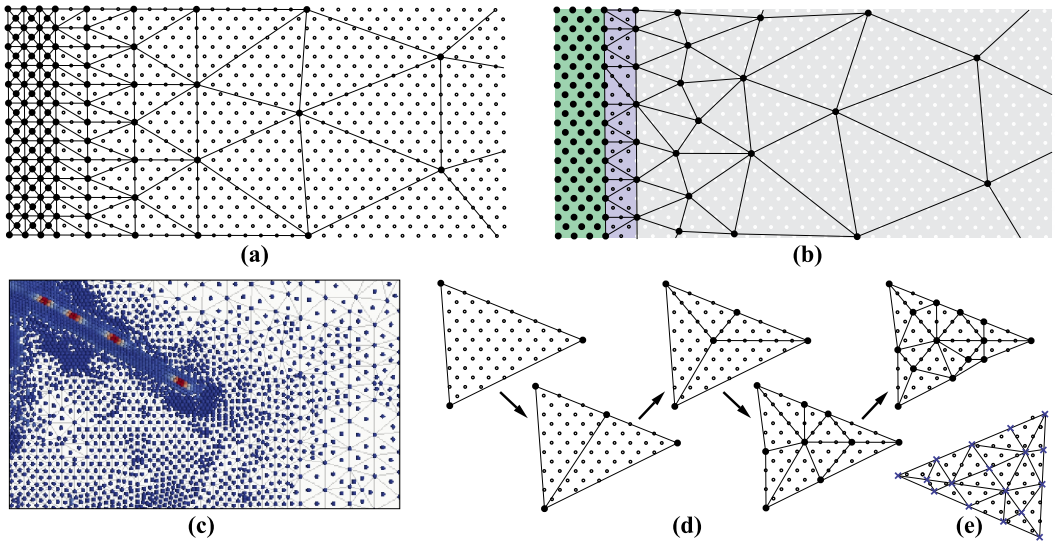

Figure 3.14: Examples of QC representations using (a) our fully nonlocal QC method with repatoms coinciding with lattice sites, and (b) an alternative formulation whose repatoms do not fall onto lattice sites in the coarse-grained region (gray), requiring a transition layer (purple) between both. (c) Shown is an actual mesh from our void growth simulations with adaptive remeshing (blue dots denote sampling atoms; the centrosymmetry color code highlights four edge dislocations). The transition from atomistic to coarse-grained is seamless, i.e., there is no sharp, well-defined interface between those. (d) A schematic illustration of mesh refinement with repatoms placed on lattice sites, starting with an element a few refinement steps away from atomistics. Placing nodes not on lattice sites can cause serious problems when passing to full atomistics, as shown in (e) where refinement is performed for the same element using non-lattice sites (crosses) and some ad-hoc method is required to transition to a feasible atomistic representation.

atomistic domain possess physically-sensible atomic neighborhoods). As shown in Fig. 3.14(c), the fully-nonlocal QC formulation does not require strict interfaces nor transition regions, so the automatic remeshing scheme can arbitrarily create atomistic and coarse-grained regions throughout the entire simulation domain (i.e., there is no clear interface between the two but the transition is seamless). This is possible only if nodes fall onto lattice sites, as shown in the schematic remeshing sequence in Fig. 3.14(d). If instead the lattice site constraint is relaxed, (see Fig. 3.14(e)), passing to the atomistic limit becomes impossible or requires ad-hoc assumptions. Thus, one may reformulate the method without the discrete lattice site constraint, but this significantly reduces the adaptive flexibility that was one of the key objectives of the fully nonlocal formulation.

3.3.2 Lagrangian vs. updated-Lagrangian QC formulation

Like almost all prior QC techniques, our nonlocal QC formulation uses a Lagrangian reference mesh, which necessitates the neighborhood updates and geometrically complicated remeshing algorithms introduced in Sections 3.1 and 3.2. To circumvent the associated difficulties, one could reformulate the technique using an updated-Lagrangian formulation, as described, e.g., in [15, 55]. Then, the reference configuration of each isoparametric element and the corresponding shape function values for all atomic locations are updated after each step. To retain all information required to reproduce the underlying atomic lattice (required, e.g., for atomic force calculations but also for the insertion of new repatoms), each element must store and frequently update its local deformation gradient $\boldsymbol{F}$ and/or the deformed Bravais lattice basis. This scenario, however, runs into the following severe conceptual problem near the atomistic limit.

The QC method's main paradigm is that atomic positions within elements do not have to be tracked individually but can be computed uniquely from the underlying lattice (e.g., if the deformation gradient or deformed Bravais vectors are known). In addition, as discussed in Section 3.3.1, we require repatoms locations to coincide with lattice sites. However, if adjacent elements carry different deformation histories, their underlying lattices can vastly differ, as shown in the example in Fig. 3.15.

In this example, mesh refinement cannot place the new repatom to be inserted on a valid lattice site, so that either choice of the new repatom (see (d) and (e)) results in elements whose nodes do not lie on the same lattice anymore. Further refinement leads to even more severe cases (see (f) and (g)), where small elements after refinement would not have all contained atoms on the same lattice. Thus, it becomes impossible to recover all underlying lattice sites and, ultimately, each atomic positions would have to be stored instead of only the Bravais vectors or the deformation gradient. This defeats the purpose of the QC method and becomes prohibitively expensive in large simulations as those shown in Section 3.2.3. One may argue that this problem primarily arises near atomistics; however, adaptive refinement can lead to atomistic resolution anywhere in the simulation domain, and the elegance of the nonlocal formulation lies in not having to specify where such atomistic regions emerge nor having to make a conceptual difference between atomistics and coarse regions. Besides, the above example is highly idealized; in reality repatoms move continuously so that changes in deformation gradients across element boundaries amplify significantly over time.

By contrast, if all atomic sites are defined in a reference configuration and refinement is performed on a reference mesh, as discussed in Sections 3.1 and 3.2, atomic positions are uniquely defined and can always be recovered (see Fig. 3.15).

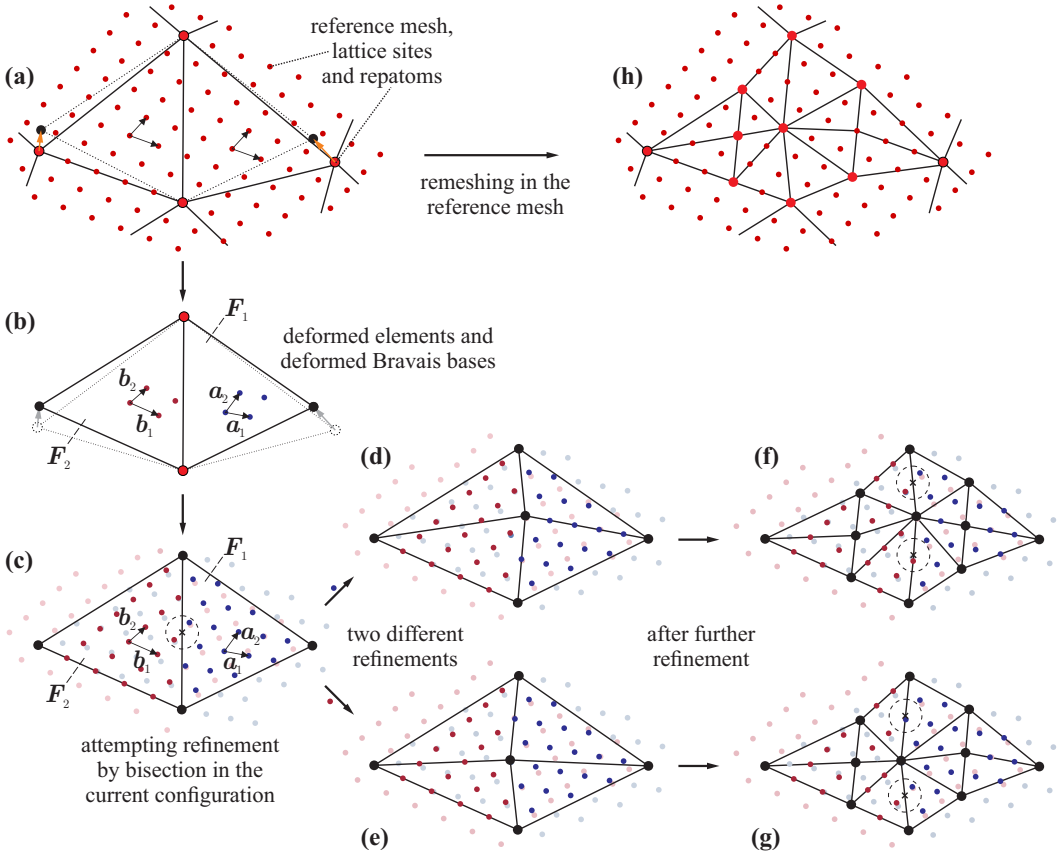

Figure 3.15: Mesh refinement with an *updated-Lagrangian* formulation: (a) starting from a uniform lattice in the initial/reference configuration, repatom motion leads to each element carrying its deformation history via deformation gradient $\boldsymbol{F}_i$ and the associated deformed Bravais bases, shown in (b). (c) The two deformed adjacent elements shall refine by bisection but no unique, common repatom location exists (the ideal location of the new node is marked by a cross, and all nearby lattice sites fall onto only one of the two adjacent lattices). (d) and (e) show the result of refinement if one of the two nearby lattice sites is chosen as the new repatom. Further refinement leads to (f) and (g) where the same problem arises due to a lattice mismatch; (h) shows remeshing in the reference configuration for comparison.

As an alternative, Biyikli and To [15] recently proposed an updated-Lagrangian MMM formulation that circumvents some of the above problems by tracking the exact positions of all lattice sites within an element when required (e.g., when an atomistic region is adaptively coarsened). Then, the QC approximation interpolates atomic positions, but to compute repatom forces the position of every atom

contained within the element is computed based on their exact deformed positions (those no longer lie on a lattice). This comes with increased computational costs and memory requirements since, effectively, no approximating sampling/summation rule is applied in those elements.

3.3.3 Mesh coarsening

Although the adaptivity techniques for rebuilding atomic neighborhoods apply generally to all (atomistic and coarse-grained) sampling atoms as well as to mesh refinement and coarsening, adaptivity of the mesh is limited to refinement here. Mesh coarsening is – in principle – a logical extension to increase efficiency (e.g., to revert back to a coarse description after defects have passed through an atomistic region). However, mesh coarsening comes with a crucial problem, and one may individually consider mesh coarsening in the coarse-grained regions as opposed to starting from full atomistic resolution.

First, in the coarse-grained domain, the QC description resembles the finite element method with a special quadrature rule (viz., a specific sampling rule) and material model (e.g., Cauchy-Born continuum). Here, if elements are not severely distorted, coarsening can be performed even against a Lagrangian mesh, which was described, e.g., in [104]. However, caution is required if anything but a local Cauchy-Born description is used (as here) since – upon coarsening – the merging of adjacent elements with slightly different deformed lattices can lead to sampling atoms with incorrect local neighborhoods that may produce spurious force artifacts. Second, coarsening from full atomistics generally requires an updated-Lagrangian description as discussed in [15, 55]. In our fully nonlocal QC formulation we refrained from the latter for reasons of efficiency (this avoids storing element-wise Bravais vectors, recomputing shape function values, etc.)

Adaptive mesh coarsening only in the coarse-grained regions, as mentioned above, can be readily applied in our method. However, coarsening only in the coarse regions is hardly sufficient to provide superior efficiency: if the atomistic region is not coarsened and mesh quality must be preserved, coarse regions cannot be effectively coarsened. Mesh coarsening from atomistics can, in principle, be applied analogous to [15] but is not considered in the present work since it challenges the fully-nonlocal paradigm: it requires having special elements in which sampling rules are replaced by summing over all atomic sites, which will have to be investigated separately.

3.4 Conclusions

We have discussed strategies for automatic model adaptivity (specifically, neighborhood updates and mesh refinement) in the fully-nonlocal QC method, which leads to full atomistic resolution being tied to evolving regions of interest such as moving lattice defects. Unlike all previous QC formulations, the chosen framework does not fundamentally differentiate between atomistic and coarse-grained regions and therefore enables a truly seamless bridging across scales, from atomistics to the continuum, solely based on interatomic potentials. Drawbacks are complicated remeshing and neighborhood update protocols that must be performed against a Lagrangian mesh. We have shown that an element-based constrained bisection algorithm (with mesh nodes restricted to the underlying crystal lattice) in combination with a neighborhood update algorithm for restructuring the Verlet neighborhood lists (that involves mapping atomic positions and distances between the current and the reference configuration) allows for a successful coarse-graining of atomistic ensembles that shows controllable, acceptable errors as compared to MD for the examples of nanoindentation and nanovoid growth. The neighborhood update scheme has been successfully deployed in both 2D and 3D simulations, while the mesh adaptation is currently limited to two dimensions (the 3D extension is, in principle, straight-forward but requires a constrained 3D mesh generator that is the subject of ongoing research, ref. Appendix C).

Chapter 4

FINITE TEMPERATURE EXTENSION USING THE *MAX-ENT* APPROACH

Chapters 2 and 3 develop the fully nonlocal QC method for $0K$ as the formulation assumes that at any point of time the exact positions and momenta of all the repatoms in the system can be calculated without considering thermal vibrations. However, the modeling of thermal and temperature-varying properties of materials using QC necessitates the modeling of the energy associated with the thermal vibration of atoms, which brings the requirement of a finite temperature extension. Section 1.2.4 reviews the current state of the art modeling techniques used to extend a coarse-grained formulation at finite temperature. For this thesis, we use the maximum entropy (abbreviated as *max-ent* approach) to extend our QC formulation to finite-temperature, and this chapter details our implementation as follows. First, we devlop the governing equations for finite temperature QC using the max-ent approach previously shown in [9, 60, 87, 120]. We then validate our model by comparing thermal expansion results with both experiments from [76] and data presented in [59], and by predicting the elastic constants of pure copper at finite temperatures. Finally, we calculate spurious force artifacts for physical and thermal forces at finite temperatures and show that the optimal summation rules, used here for the first time in connection with finite-temperature QC, minimize these ghost forces more effectively than all others.

Note that the mex-ent finite-temperature QC concept has been derived previously [60]. Here, we show derivations to the extent necessary for subsequent discussions, and we show details of importance for the implementation that are not shown in previous publications.

4.1 Phase Space Average and the Principle of Maximum Entropy

Like in Chapter 2, we start the formulation with an ensemble of N atoms with positions $\boldsymbol{q} = \{\boldsymbol{q}_1, \ldots, \boldsymbol{q}_N\}, \boldsymbol{q} \in \mathscr{R}^{3N}$ and momenta $\boldsymbol{p} = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_N\}, \boldsymbol{p} \in \mathscr{R}^{3N}$, and the Hamiltonian $H(\boldsymbol{q}, \boldsymbol{p})$ given by (2.1). Considering the phase space of all positions and momenta $\left(\mathscr{R}^{3N} \times \mathscr{R}^{3N}\right)$, a point in this phase space will be denoted by $(\boldsymbol{q}, \boldsymbol{p})$, and any function that can be determined as a function of positions and momenta at any instant of time (e.g. (2.1)) will be termed a *phase function*. This

motivates the definition of *phase space average* for any phase function $A(\boldsymbol{q},\boldsymbol{p})$ as:

$$\langle A\rangle = \frac{1}{h^{3n}} \underbrace{\int\int}_{\mathscr{R}^{3N}\times\mathscr{R}^{3N}} A(\boldsymbol{q},\boldsymbol{p})\,\rho(\boldsymbol{q},\boldsymbol{p})\,\mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{p}, \quad \text{where} \quad \mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{q} = \prod_{i=1}^{N}\prod_{a=1}^{3} dq_{ia}dp_{ia}. \tag{4.1}$$

In (4.1), h is Planck's constant, and the phase function $\rho(\boldsymbol{q},\boldsymbol{p})$ is the probability density function, which provides the probability that the system is at a point $(\boldsymbol{q},\boldsymbol{p})$ in phase space. Thus, by definition, $\rho(\boldsymbol{q},\boldsymbol{p}) \geq 0$. The pre-factor $\left(h^{3N}\right)^{-1}$ comes from the the requirement that entropy is extensive in classical statistical thermodynamics. Following the convention presented in [59], we define the entire phase space as $\Gamma = \left(\mathscr{R}^{3N}\times\mathscr{R}^{3N}\right)$. To determine this probability density function, we will use the principle of maximum entropy from statistical mechanics in the next section. To that end, as shown in [59] for an atom a, we assume that it moves in the vicinity of a position $\overline{\boldsymbol{q}_a}$ with standard deviation $\sqrt{3}\tau_a$ and it has a momentum of $\overline{\boldsymbol{p}_a}$ with standard deviation $\sqrt{3\sigma_a}$. From the definition of the probability density function, these mean positions and momenta for every atom a are

$$\overline{\boldsymbol{q}_a} = \langle \boldsymbol{q}_a\rangle, \qquad \overline{\boldsymbol{p}_a} = \langle \boldsymbol{p}_a\rangle \qquad \forall a = 1,2,...N. \tag{4.2}$$

We can physically interpret $\overline{\boldsymbol{q}_a}$ and $\overline{\boldsymbol{p}_a}$ as variables that are used to solve the dynamics of the system at the macroscopic time scale. In other words, at any instant of time (on the macroscopic scale), we solve for the phase average of positions and momenta instead of the actual values[1]. From these, we obtain N new constraints, which are

$$\begin{aligned} \langle|\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2\rangle &= 3\tau_a^2, \\ \langle|\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2\rangle &= 3\sigma_a^2, \qquad \forall a = 1,2,...N. \end{aligned} \tag{4.3}$$

As seen in [60], in order to obtain a physical interpretation and for simplifying expressions, we define another quantity[2]

$$\omega_a = \frac{\sigma_a}{\tau_a}. \tag{4.4}$$

Rearranging (4.3) to eliminate τ_a we rewrite our constraints as

$$\langle|\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2\rangle + \omega_a^2\langle|\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2\rangle = 6\sigma_a^2, \qquad \forall a = 1,2,...N. \tag{4.5}$$

[1] It is easy to observe that at $0K$ the QC formulation from Chapter 2 is a special case of (4.2) where $\rho(\boldsymbol{q},\boldsymbol{p})$ is a Direc δ function.

[2] The ratio of the standard deviations of position and momentum, ω_a, has a physical meaning and interpretation only at finite temperatures. At $0K$, $\boldsymbol{q}_a = \overline{\boldsymbol{q}_a}$ and $\boldsymbol{p}_a = \overline{\boldsymbol{p}_a}$, resulting in both τ_a and σ_a being 0. In that case, ω_a is undefined.

To derive the governing equations of motion and solve for $\boldsymbol{q}_a$ and $\boldsymbol{p}_a$ at the macroscopic time scale, the expression for the probability density function $\rho(\boldsymbol{q},\boldsymbol{p})$ is needed, which is obtained by using the *principle of maximum entropy* [60] by Jaynes [49], which states that the probability distribution that best represents the current state of knowledge is the one with the largest entropy. Therefore, in the set of all trial probability distributions, the one with the maximum entropy is the proper distribution. The expression for the global entropy is given by Boltzmann as

$$S[\rho] = -k_B \langle \log \rho(\boldsymbol{q},\boldsymbol{p}) \rangle, \tag{4.6}$$

where k_B is the Boltzmann constant. We solve for the probability distribution function which maximizes this entropy with the additional constraints

$$\begin{aligned} \langle 1 \rangle &= 1, \\ \langle |\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2 \rangle + \omega_a^2 \langle |\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2 \rangle &= 6\sigma_a^2 \qquad \forall a = 1,2,...N, \end{aligned}$$

where the second set of constraints comes from (4.5), and the first constraint comes from the requirement that the system always has a solution in phase space at any instant of time. Using Lagrange multipliers λ and β_a, the expression to maximize becomes[3]

$$\mathscr{L}[\rho] = S[\rho] - \lambda \langle 1 \rangle - \sum_a \beta_a \left[\langle |p_a - \overline{p_a}|^2 \rangle + \omega_a^2 \langle |q_a - \overline{q_a}|^2 \rangle \right], \tag{4.7}$$

and by taking the variation of (4.7) with respect to ρ, we obtain

$$\int_\Gamma \left[-k_B (1 + \log \rho) - \lambda - \sum_a \beta_a \left[|\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2 + \omega_a^2 |\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2 \right] \right] \delta\rho \, \mathrm{d}\boldsymbol{q} \, \mathrm{d}\boldsymbol{p} = 0. \tag{4.8}$$

For the above equation to be satisfied for every variation $\delta\rho$, we conclude

$$-k_B (1 + \log \rho) - \lambda - \sum_a \rho \beta_a \left[|\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2 + \omega_a^2 |\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2 \right] = 0, \tag{4.9}$$

which gives the probability density function as

$$\rho(\boldsymbol{q},\boldsymbol{p}) = \frac{1}{e^{1+\lambda}} e^{-\left(\sum_a \beta_a \left[|\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2 + \omega_a^2 |\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2 \right] \right)}. \tag{4.10}$$

Once an expression for ρ is obtained, the quantity $Z = \exp(1+\lambda)^{-1}$ (also known as the *partition function* of the system) is obtained as

$$Z = \frac{1}{h^{3N}} \prod_i^N \left[\left(\frac{\pi}{\beta_i \omega_i} \right)^3 \right]. \tag{4.11}$$

[3] The constants λ and $6\beta_a^2\sigma_a^2$ are not included because they vanish upon differentiation.

With this expression for the partition function, we obtain the remaining N Lagrange multipliers by substituting the probability density expression in (4.2), which simplifies to

$$\beta_a = \frac{1}{2\sigma_a^2} \qquad \forall a = 1, 2, ...N. \tag{4.12}$$

Once all Lagrange multiplies have been obtained, the final expression for the *max-ent* probability distribution function is

$$\begin{aligned} \rho\left(\boldsymbol{q}, \boldsymbol{p}\right) &= \frac{1}{Z} \exp\left(- \sum_a \frac{|\boldsymbol{p}_a - \overline{\boldsymbol{p}_a}|^2 + \omega_a^2 |\boldsymbol{q}_a - \overline{\boldsymbol{q}_a}|^2}{2\sigma_a^2} \right), \qquad \text{where} \\ Z &= \frac{1}{h^{3N}} \prod_i^N \left[\left(\frac{2\pi\sigma_i^2}{\omega_i} \right)^3 \right]. \end{aligned} \tag{4.13}$$

4.2 Thermodynamic Potentials

It is worthwhile to note that because of the form of (4.13), the latter can be decomposed into a product of exponentials, each depending on the terms associated with each atom. This form is obtained because of the nature of the imposed constraints, which are local in nature. Such a form implicitly assumes the *local-equilibrium hypothesis*, which states that as long as the system can be divided into different subsystems, each of which are close to equilibrium, we can solve a non-equilibrium macroscopic problem by looking at equilibrium relations at the microscopic level. For this formulation, one considers every atom as a subsystem. In other words, there will be two kinds of "relaxation times", one for establishing local equilibrium at every subsystem, and one for the entire system. The local equilibrium hypothesis allows the definition of thermodynamic properties like entropy and temperature for every atom locally. Using the probability function derived in (4.13), we proceed to calculate different thermodynamic potentials, and in particular, their local forms.

4.2.1 Entropy

Recalling the expression of entropy from (4.6) and using (4.13), we obtain the entropy of the system as

$$S = k_B \left(3N + 3 \sum_a^N \log \left[\frac{\sigma_a^2}{\hbar\omega_a} \right] \right). \tag{4.14}$$

Using the local equilibrium hypothesis outlined above, we decompose this expression into a summation of entropies of the local subsystems having the form $S = \sum_a^N S_a$, which gives us an expression of the local entropy for every atom to be

$$S_a = 3k_B \log \frac{\sigma_a^2}{\hbar\omega_a} + 3k_B. \tag{4.15}$$

For the purposes of evaluating the expressions for different thermodynamic quantities, we have made the assumption that the standard deviations of atomic positions and momenta σ_a and ω_a are known quantities. Subsequent sections elaborate on ways to determine these parameters. For that purpose, it is useful to rewrite (4.15) to obtain an explicit expression for σ_a as

$$\sigma_a = \sqrt{\hbar\omega_a}\exp\left[\frac{1}{3}\left(\frac{S_a}{k_B} - 3\right)\right]. \tag{4.16}$$

4.2.2 Internal energy

To obtain the expression for the internal energy, we start with the Hamiltonian of the system (2.1). Expressing the Hamiltonian of the system as a summation of local Hamiltonians,

$$H = \sum_a^N H_a, \tag{4.17}$$

gives the Hamiltonian for every atom as

$$H_a = \frac{1}{2}\frac{|p_a|^2}{m_a} + V_a(\boldsymbol{q}). \tag{4.18}$$

We observe that the kinetic energy in (4.18) is obtained locally; however, the potential energy requires the positions of all atoms in the system, and is hence nonlocal. In statistical mechanics, internal energy is defined as the phase average of the total Hamiltonian of the system:

$$E = \frac{1}{h^{3N}}\int_\Gamma H(\boldsymbol{q},\boldsymbol{p})\rho(\boldsymbol{q},\boldsymbol{p})\,\mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{p}. \tag{4.19}$$

Using (4.13), the additive nature of the Hamiltonian (4.17) and (4.18), the integral to evaluate simplifies to

$$\begin{aligned} E &= \frac{1}{h^{3N}}\sum_a\int_\Gamma\left[\frac{1}{2}\frac{|p_a|^2}{m_a}\right]\rho(\boldsymbol{q},\boldsymbol{p})\,\mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{p} + \frac{1}{h^{3N}}\sum_a\int_\Gamma\left[V_a(\boldsymbol{q})\right]\rho(\boldsymbol{q},\boldsymbol{p})\,\mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{p}. \\ &= \sum_a\left[\frac{1}{2m_a}\langle|p_a|^2\rangle + \langle V_a(\boldsymbol{q})\rangle\right]. \end{aligned}$$

The phase average of the kinetic energy is simplified by manipulating (4.2), which reduces to

$$\langle|\boldsymbol{p}_a|^2\rangle = \left(3\sigma_a^2 + |\overline{\boldsymbol{p}_a}|^2\right). \tag{4.20}$$

Using (4.20) and (4.16), the final expression of the internal energy is obtained as

$$E = \sum_a\left[\frac{3}{2m_a}\hbar\omega_a\exp\left(\frac{S_a}{3k_B} - 1\right) + \frac{1}{2}\frac{|\overline{\boldsymbol{p}_a}|^2}{m_a} + \langle V_a(\boldsymbol{q})\rangle\right]. \tag{4.21}$$

4.2.3 Helmholtz free energy

The Helmholtz free energy is defined as a Legendre transform of the internal energy with respect to entropy:

$$F\left(\overline{\boldsymbol{q}}, \boldsymbol{\omega}, \boldsymbol{\theta}\right) = \inf_{\boldsymbol{S}} \left\{ E\left(\overline{\boldsymbol{q}}, \boldsymbol{\omega}, \boldsymbol{S}\right) - \sum_a \theta_a S_a \right\},$$

and the local equilibrium relation for every particle is:

$$\theta_a = \frac{\partial E}{\partial S_a}, \tag{4.22}$$

where θ_a is the local temperature of atom a. The dependent variables for the Helmholtz free energy will be the variables which we solve for in the macroscopic equilibrium equations. We restrict our analysis to quasistatics and so $\overline{\mathbf{p}}_\mathbf{a} = \mathbf{0}$ for every atom. Due to the nonlocal nature of the potential energy $V_a(\boldsymbol{q})$, we invoke another fundamental result of statistical mechanics to obtain the equilibrium relation (4.22) without solving for the phase average, called *equipartition of energy*. It states that the quadratic term in the Hamiltonian (which is the kinetic energy for our system) contributes $\frac{k_B\theta}{2}$ to the internal energy of the system. In addition, because of the local equilibrium hypothesis, we enforce this locally as

$$\left\langle \frac{|\boldsymbol{p}_a|^2}{2m_a} \right\rangle = \frac{3}{2} k_B \theta_a. \tag{4.23}$$

Comparing (4.23) to (4.20) lets us relate σ_a to the local temperature as

$$\sigma_a^2 = k_B \theta_a m_a. \tag{4.24}$$

Finally, using (4.14) and (4.24), the equilibrium relation between local entropy and local temperature is

$$S_a = 3k_B \log\left(\frac{k_B m_a \theta_a}{\hbar \omega_a}\right) + 3k_b, \tag{4.25}$$

which along with (4.25), gives

$$F(\overline{\boldsymbol{q}}, \boldsymbol{\omega}, \boldsymbol{\theta}) = \sum_{a=1}^{N} \left[3k_B\theta_a \log \frac{\hbar\omega_a}{k_b m_a \theta_a} + \langle V_a(\boldsymbol{q}, \boldsymbol{p}, \boldsymbol{\omega}) \rangle \right]. \tag{4.26}$$

Once we have calculated the internal energy of the system, the problem of finding equilibrium configurations becomes a minimization problem with respect to the mean positions. In the following sections and for the rest of this work, we restrict our studies to isothermal conditions, for which the Helmholtz free energy is a

suitable energy functional. Using variational mean field theory [59], for an *isothermal, quasistatic* process, the equilibrium configurations of a system are obtained by solving

$$\inf_{\overline{\boldsymbol{q}}} \inf_{\boldsymbol{\omega}} F\left(\overline{\boldsymbol{q}}, \boldsymbol{\omega}, \boldsymbol{\theta}\right). \tag{4.27}$$

However, before we proceed to obtain equilibrium expressions by performing this minimization, we simplify the phase average of the potential energy in the next section.

4.3 Potential Phase Average

In (4.26), the phase average for the potential cannot be evaluated analytically. This arises due to the nonlocal nature of the potential for atomistics. In order to be able to get equilibrium configurations of the system, the phase average has to be approximated numerically. The Gaussian nature of the probability distribution function (4.13) motivates the use of numerical Gaussian quadrature.

Consider any general potential at an atomic site i. Its phase average is given (from (4.1)) as:

$$\begin{aligned}\langle V_i\left(\boldsymbol{q}\right)\rangle &= \frac{1}{h^{3N}} \int_{\Gamma} V_i(\boldsymbol{q})\, \rho\left(\boldsymbol{q}, \boldsymbol{p}, \boldsymbol{\omega}\right) \mathrm{d}\boldsymbol{q}\, \mathrm{d}\boldsymbol{p}, \\ &= \frac{1}{Zh^{3N}} \int_{\Gamma} V_i(\boldsymbol{q}) \exp\left(-\sum_a \frac{|p_a - \overline{p_a}|^2 + \omega_a^2 |q_a - \overline{q_a}|^2}{2\sigma_a^2}\right) \mathrm{d}\boldsymbol{q}\, \mathrm{d}\boldsymbol{p},\end{aligned}$$

using (4.13). Noting that the potential only depends on positions, performing the integral over the $\boldsymbol{p}$ space independently gives

$$\langle V_i\rangle = \frac{1}{Zh^{3N}} \prod_{j=1}^{N} \left[\sqrt{2\sigma_j^2\pi}\right]^3 \int_{\boldsymbol{q}} V_i(\boldsymbol{q}) \exp\left(-\sum_a \frac{\omega_a^2 |q_a - \overline{q_a}|^2}{2\sigma_a^2}\right) \mathrm{d}\boldsymbol{q}.$$

Recognizing that $Z = \frac{1}{h^{3N}} \prod_i^N \left[\left(\frac{2\sigma_i^2\pi}{\omega_i}\right)^3\right]$, and expressing the potential as a function of positions of all the atoms in the system, we are left with

$$\langle V_i\rangle = \frac{1}{(2\pi)^{\frac{3N}{2}}} \prod_{j=1}^{N} \frac{\omega_j^3}{\sigma_j^3} \int_{\boldsymbol{q}} V_i(\boldsymbol{q}_1, \boldsymbol{q}_2, \ldots \boldsymbol{q}_N) \exp\left(-\sum_a \frac{\omega_a^2 |q_a - \overline{q_a}|^2}{2\sigma_a^2}\right) \mathrm{d}\boldsymbol{q}_1\, \mathrm{d}\boldsymbol{q}_2 \ldots \mathrm{d}\boldsymbol{q}_N.$$

We only consider potentials of the embedded-atom type so to obtain the potential at an atomic site i, only those atoms that lie in the interaction distance ($\boldsymbol{q}_J \in \mathscr{S}_i$) of atomic site i are considered. The phase average now is re-expressed as

$$\langle V_i\rangle = \frac{1}{(2\pi)^{\frac{3N}{2}}} \left(\prod_{k=1}^{N} \frac{\omega_k^3}{\sigma_k^3}\right) \int_{\boldsymbol{q}} V_i(\boldsymbol{q}_j : \boldsymbol{q}_j \in \mathscr{S}_i) \exp\left(-\sum_a \frac{\omega_a^2 |q_a - \overline{q_a}|^2}{2\sigma_a^2}\right) \mathrm{d}\boldsymbol{q}_1 \ldots \mathrm{d}\boldsymbol{q}_N.$$

Noting that the atomic positions that lie outside the interaction distance can be integrated independently, we are left with

$$\langle V_i \rangle = \frac{1}{(2\pi)^{\frac{3n_N}{2}}} \left(\prod_{k \in \mathscr{S}_i} \frac{\omega_k^3}{\sigma_k^3} \right) \prod_{j \in \mathscr{S}_i} \int V_i(\boldsymbol{q}_j : \boldsymbol{q}_j \in \mathscr{S}_i) \exp\left(-\frac{\omega_j^2 |q_j - \overline{q_j}|^2}{2\sigma_j^2} \right) \mathrm{d}\boldsymbol{q}_j, \tag{4.28}$$

where n_N is the number of neighbors in the interaction distance for atomic site i. Before we numerically solve (4.28), we define a change of variables for each of the coordinates, viz. for $a = 1,2,3$

$$\frac{\omega_j^2}{2\sigma_j^2}(q_{j,a} - \overline{q}_{j,a})^2 = x_{j,a}^2 \qquad \Leftrightarrow \qquad q_{j,a}(x_{j,a}, \overline{q}_{j,a}) = \sqrt{\frac{2\sigma_j^2}{\omega_j^2}} x_{j,a} + \overline{q}_{j,a}, \tag{4.29}$$

so that

$$\mathrm{d}q_{j,a} = \sqrt{\frac{2\sigma_j^2}{\omega_j^2}}\, \mathrm{d}x_{j,a}.$$

The simplified expression of the phase average of the potential energy becomes

$$\langle V_i \rangle = \left(\frac{1}{\sqrt{\pi}} \right)^{3n_N} \prod_{j \in \mathscr{S}_i} \int_{-\infty}^{\infty} V_i(\boldsymbol{q}_j\left(\boldsymbol{x}_j, \overline{\boldsymbol{q}_j}\right) : \boldsymbol{q}_j \in \mathscr{S}_i) \prod_{a=1}^{3} \exp\left(-x_{j,a}^2\right) \mathrm{d}x_{j,a}. \tag{4.30}$$

4.3.1 Numerical quadrature

At this point, numerical quadrature is required to solve (4.30). Hermite-Gauss quadratures are integral approximations of the type

$$\int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} f(x_1, \dots, x_n) \exp(-x_1^2) \cdot \dots \cdot \exp(-x_n^2)\, \mathrm{d}x_1 \cdot \dots \cdot \mathrm{d}x_n \approx \sum_{p=1}^{n_{QP}} f(\boldsymbol{x}_p) W_p,$$

for n scalar independent variables. n_{QP} represent the number of quadrature points chosen, W_p is the weight at every point and $\boldsymbol{x}_p = \{x_{1,p}, x_{2,p}, \dots x_{n_P}\}$ is the set of quadrature values at sampling p. Equation (4.30) fits into this description, and so our exact phase average expression for the potential energy of an atom at a site i turns into the approximation

$$\langle V_i \rangle \approx \left(\frac{1}{\sqrt{\pi}} \right)^{3n_N} \sum_{p=1}^{n_{QP}} W_p \widetilde{V}_i\left[\boldsymbol{x}_{1,p}, \dots, \boldsymbol{x}_{n_N,p}\right]. \tag{4.31}$$

The actual sampling values and weights are chosen based on quadrature rules given by A.H. Stroud [108]. Based on different degrees of quadrature, $\boldsymbol{x}_P$ and W_P have to be determined. To avoid computational overhead, our emphasis is to seek quadrature rules that use the least number of quadrature points for a given degree.

Third-Degree Quadrature

For n scalar variables, this third degree-quadrature contains $n_{QP} = 2n$ samples. These are obtained by requiring that the summation integrates all monomials of degree ≤ 3 exactly. The n_{QP} samples are shown in Table 4.1. For the implementa-

p	W_P	x_1	x_2	...	x_n
1	$\frac{V}{2n}$	r	0	...	0
2	$\frac{V}{2n}$	0	r	...	0
...	...				
n	...	0	0	...	r
$n+1$	...	$-r$	0	...	0
$n+2$		0	$-r$	...	0
...					
$2n$	$\frac{V}{2n}$	0	0	...	$-r$

$$\text{where} \qquad V = I(1) = \pi^{\frac{n}{2}}, \qquad r = \sqrt{\frac{n}{2}}.$$

Table 4.1: Quadrature weights and values for each of the $2n$ samples

tion of this quadrature rule, we note that if a site i has n_N nearest neighbors which interact for the potential, then $n = \textit{Spatial Dimension} \times (n_N + 1)$ as the potential also depends on the position of the site in question, and for every involved position, the number of scalar variables is equal to the spatial dimension of the simulation. Thus, for 3D simulations, $n = 3\,(n_N + 1)$. This makes the total number of quadrature points to be $n_{QP} = 2n = 6n_N + 6$. This also means that for every quadrature sample p, as shown in Table 4.1, we update one coordinate of one of the positions by the quadrature shift $\pm r$.

Fifth-Degree Quadrature

For n scalar variables, the general fifth-degree quadrature contains $n_{QP} = 2n^2 + 1$ samples. The general form of the approximation of the integral is written as

$$
\begin{aligned}
I(f(\boldsymbol{x})) \approx Q(f) = & B_0 f(0) + \\
& B_1 \left[\sum_{\text{permutations of } r} (r, 0, \ldots, 0) + (-r, 0, \ldots, 0) \right] \\
& + B_2 \left[\sum_{\text{permutations of } r} (r, r, \ldots, 0) + \sum_{\text{permutations of } r} (r, -r, \ldots, 0) \right] + \\
& B_2 \left[\sum_{\text{permutations of } r} (-r, r, \ldots, 0) + \sum_{\text{permutations of } r} (-r, -r, \ldots, 0) \right],
\end{aligned}
$$

where

$$
\begin{aligned}
r &= \sqrt{\frac{3}{5}}, \qquad V = I(1) = \pi^{\frac{n}{2}}, \\
B_0 &= \frac{25n^2 - 115n + 162}{162} V, \\
B_1 &= \frac{70 - 25n}{162} V, \\
B_2 &= \frac{25}{324} V.
\end{aligned}
$$

Similar to our observation for the third-degree case, for a 3D simulation and for n_N number of neighbors, here we obtain $n = 3(n_N + 1)$, which gives $n_{QP} = 18(n_N + 1)^2 + 1$ samples. For the EAM Johnson potential [51] for Cu, we have 18 neighbors in the interaction distance, which gives $n_{QP} = 6499$. This is the number of quadrature sampling we perform for every site where the potential is evaluated, which makes the fifth-degree quadrature prohibitively expensive.

4.4 Equilibrium Configurations

After obtaining an approximate expression for the potential energy at an atomic site using Gaussian quadrature in Section 4.3, we proceed to perform the minimization in (4.27) in order to solve for the mean positions and mean-field parameters ω. We once again note that we restrict this study to isothermal cases at known temperature, so the temperature at every atomic site θ_a is known. Thus, for the system to be in

equilibrium, we need

$$\frac{\partial}{\partial \overline{\boldsymbol{q}}_i} F(\overline{\boldsymbol{q}}, \boldsymbol{\omega}, \theta) = \mathbf{0}, \tag{4.32}$$

$$\frac{\partial}{\partial \omega_i} F(\overline{\boldsymbol{q}}, \boldsymbol{\omega}, \boldsymbol{\theta}) = 0. \qquad \forall i = 1,2,3...N. \tag{4.33}$$

4.4.1 Position minimization

We solve for the time evolution of the positions by minimizing the free energy with respect to the position. From (4.26) and (4.33),

$$\mathbf{0} = \frac{\partial}{\partial \overline{\boldsymbol{q}}_i} \sum_{a=1}^{N} \left[3k_B\theta_a \log \frac{\hbar\omega_a}{k_b m_a \theta_a} + \langle V_a(\boldsymbol{q}) \rangle \right],$$

and using the phase average approximation (4.31), we obtain

$$\begin{aligned} \mathbf{0} &= \sum_{a=1}^{N} \left[\frac{\partial}{\partial \overline{\boldsymbol{q}}_i} \left(\frac{1}{\sqrt{\pi}} \right)^{3n_N} \sum_{p=1}^{n_{QP}} W_p \widetilde{V}_a \left[\boldsymbol{x}_{1,p}, \ldots, \boldsymbol{x}_{n_N,p} \right] \right] \\ &= \left(\frac{1}{\sqrt{\pi}} \right)^{3n_N} \sum_{a=1}^{N} \sum_{p=1}^{n_{QP}} W_p \left[\frac{\partial}{\partial \overline{\boldsymbol{q}}_i} V_a \left[\boldsymbol{q}_1(\boldsymbol{x}_{1,p}, \overline{\boldsymbol{q}}_1), \ldots, \boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N,p}, \overline{\boldsymbol{q}}_{n_N}) \right] \right], \end{aligned}$$

where, for every quadrature sample,

$$\widetilde{V}_a \left[\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{n_N} \right] = V_a \left[\boldsymbol{q}_1(\boldsymbol{x}_1, \overline{\boldsymbol{q}}_1), \ldots, \boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N}, \overline{\boldsymbol{q}}_{n_N}) \right], \qquad \boldsymbol{q}_j(\boldsymbol{x}_j, \overline{\boldsymbol{q}}_j) = \sqrt{2} \frac{\sigma_j}{\omega_j} \boldsymbol{x}_j + \overline{\boldsymbol{q}}_j.$$

We also observe that $\boldsymbol{q}_i$ and $\overline{\boldsymbol{q}}_i$ are linearly related, so

$$\frac{\partial}{\partial \overline{\boldsymbol{q}}_i} V(\boldsymbol{q}_1, \ldots, \boldsymbol{q}_{n_N}) = \frac{\partial}{\partial \boldsymbol{q}_i} V(\boldsymbol{q}_1, \ldots, \boldsymbol{q}_{n_N}).$$

So the governing equation of motion becomes

$$\begin{aligned} \mathbf{0} &= \left(\frac{1}{\sqrt{\pi}} \right)^{3n_N} \sum_{a=1}^{N} \sum_{p=1}^{n_{QP}} W_p \left[\frac{\partial}{\partial \boldsymbol{q}_i} V_a \left[\boldsymbol{q}_1(\boldsymbol{x}_1, \overline{\boldsymbol{q}}_1), \ldots, \boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N}, \overline{\boldsymbol{q}}_{n_N}) \right] \right], \\ &= -\left(\frac{1}{\sqrt{\pi}} \right)^{3n_N} \sum_{a=1}^{N} \sum_{p=1}^{n_{QP}} W_p \boldsymbol{f}_{ai} \left[\boldsymbol{q}_1(\boldsymbol{x}_1, \overline{\boldsymbol{q}}_1), \ldots, \boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N}, \overline{\boldsymbol{q}}_{n_N}) \right], \end{aligned} \tag{4.34}$$

where $\boldsymbol{f}_{ai} = -\frac{\partial V_a}{\partial \boldsymbol{q}_i}$ is the atomic force between sites a and i as defined in (2.2).

4.4.2 ω minimization

Solving for the equilibrium values of the free parameter ω follows similar steps as in Section 4.4.1. Minimizing with respect to ω gives

$$0 = \frac{\partial}{\partial \omega_i} \sum_{a=1}^{N} \left[3k_B\theta_a \log \frac{\hbar\omega_a}{k_b m_a \theta_a} + \langle V_a(\boldsymbol{q}) \rangle \right].$$

Once again, using (4.26) and (4.30), the expression reduces to

$$-\frac{3k_B\theta_i}{\omega_i} = \left(\frac{1}{\sqrt{\pi}}\right)^{3n_N} \sum_{a=1}^{N}\sum_{p=1}^{n_{QP}} W_p \frac{\partial}{\partial \omega_i} V_a\left[\boldsymbol{q}_1(\boldsymbol{x}_{1,p},\overline{\boldsymbol{q}}_1),\ldots,\boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N,p},\overline{\boldsymbol{q}}_{n_N})\right].$$

To obtain the partial derivative of the potential with respect to ω, we use

$$\frac{\partial}{\partial \omega_i} = \frac{\partial}{\partial \boldsymbol{q}_i}\cdot\frac{\partial \boldsymbol{q}_i}{\partial \omega_i} = \frac{\partial}{\partial \boldsymbol{q}_i}\cdot\frac{\partial}{\partial \omega_i}\left(\sqrt{2}\frac{\sigma_i}{\omega_i}\boldsymbol{x}_i + \overline{\boldsymbol{q}}_i\right) = -\sqrt{2}\frac{\sigma_i}{\omega_i^2}\boldsymbol{x}_i\cdot\frac{\partial}{\partial \boldsymbol{q}_i} = -\sqrt{2}\frac{\sigma_i}{\omega_i^2}\boldsymbol{x}_i\cdot\frac{\partial}{\partial \overline{\boldsymbol{q}}_i},$$

which turns the second set of governing equations into

$$\frac{3k_B\theta_i}{\omega_i} = \left(\frac{1}{\sqrt{\pi}}\right)^{3n_N} \sum_{a=1}^{N}\sum_{p=1}^{n_{QP}} W_p \sqrt{2}\frac{\sigma_i}{\omega_i^2}\boldsymbol{x}_{i,p}\cdot\frac{\partial}{\partial \overline{\boldsymbol{q}}_i} V_a\left[\boldsymbol{q}_1(\boldsymbol{x}_{1,p},\overline{\boldsymbol{q}}_1),\ldots,\boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N,p},\overline{\boldsymbol{q}}_{n_N})\right]. \tag{4.35}$$

Using (4.24) and following steps from Section 4.4.1, the governing equation for the time evolution of ω is given as

$$0 = \frac{3k_B\theta_i}{\omega_i} + \left(\frac{1}{\sqrt{\pi}}\right)^{3n_N} \sum_{a=1}^{N}\sum_{p=1}^{n_{QP}} W_p \frac{\sqrt{2k_B\theta_i m_i}}{\omega_i^2}\boldsymbol{x}_{i,p}\cdot\boldsymbol{f}_{ai}\left[\boldsymbol{q}_1(\boldsymbol{x}_1,\overline{\boldsymbol{q}}_1),\ldots,\boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N},\overline{\boldsymbol{q}}_{n_N})\right]. \tag{4.36}$$

Thus (4.34) and (4.36) are the equations that we solve for every atomic site i in order to obtain the equilibrium configurations of the system. These undergo modifications once coarse-graining is introduced similarly to Chapter 2 for the finite temperature case. The coarse-grained equations are derived in Section 4.6.

4.5 Recovering 0K behavior

An essential check in any finite-temperature formulation is the recovery of the governing equation (2.6) if we set the temperature to zero. As noted before, the mean-field parameter ω is meaningless for this case and therefore should not evolve with time. From (4.29), the position shifts for Gaussian quadrature ($\boldsymbol{x}_j$) are defined for every quadrature sample and they relate $\boldsymbol{q}_j$ and $\overline{\boldsymbol{q}}_j$ as

$$\boldsymbol{q}_{j,a}(\boldsymbol{x}_{j,a},\overline{\boldsymbol{q}}_{j,a}) = \sqrt{\frac{2\sigma_j^2}{\omega_j^2}}\boldsymbol{x}_{j,a} + \overline{\boldsymbol{q}}_{j,a},$$

but for $\theta_j = 0$, we have $\sigma_j = 0$ from (4.24). So each of the quadrature samples are identical and we have

$$\langle V_a\rangle = V_a\left[\boldsymbol{q}_1,\boldsymbol{q}_2\ldots,\boldsymbol{q}_{n_N}\right]\left(\frac{1}{\sqrt{\pi}}\right)^{3n_N}\sum_{p=1}^{n_{QP}} W_p.$$

For both the third- and fifth-degree quadratures from Section 4.3.1, we have

$$\sum_{p=1}^{n_{QP}} W_p = \sqrt{\pi}^{3n_N} = I(1),$$

which simplifies the phase average of the potential to

$$\langle V_a \rangle = V_a\left[\boldsymbol{q}_1, \boldsymbol{q}_2 \ldots, \boldsymbol{q}_{n_N}\right].$$

Substituting this in (4.34), we indeed recover (2.6).

4.6 Quasicontinuum Formulation

As shown in Chapter 2, in QC we coarse-grain the system into repatoms and sampling atoms. After obtaining (4.34) and (4.36), we proceed to apply this coarse-graining method for the finite-temperature extension of our QC formulation. To this end, we recall that the approximate current (mean) position $\overline{\boldsymbol{q}}_i^h$ and (mean) momentum $\overline{\boldsymbol{p}}_i^h$ of atom i are obtained by interpolation:

$$\begin{aligned} \overline{\boldsymbol{q}}_i \approx \overline{\boldsymbol{q}}_i^h &= \sum_{a=1}^{N_h} N_a(\boldsymbol{X}_i)\, \boldsymbol{x}_a, \\ \overline{\boldsymbol{p}}_i \approx \overline{\boldsymbol{p}}_i^h = m_i \dot{\boldsymbol{x}}_i^h &= m_i \sum_{a=1}^{N_h} N_a(\boldsymbol{X}_i)\, \dot{\boldsymbol{x}}_a \qquad \text{for} \quad i = 1, \ldots, N. \end{aligned} \tag{4.37}$$

$N_a(\boldsymbol{X}_i)$ is the shape function of repatom a evaluated at the position $\boldsymbol{X}_i$ of lattice site i in the reference configuration. Using the same affine interpolation as a natural extension for the new mean-free parameter ω_i allows us write the approximate value for an atom i as

$$\omega_i \approx \omega_i^h = \sum_{a=1}^{N_h} N_a(\boldsymbol{X}_i)\, s_a, \qquad \text{for} \quad i = 1, \ldots, N. \tag{4.38}$$

Here, we have introduced the notation of $\{\boldsymbol{x}_1, \boldsymbol{x}_2 \ldots \boldsymbol{x}_{N_H}\}$ as the repatom positions and $\{s_1, s_2 \ldots s_{N_H}\}$ as the repatom mean-free parameters corresponding to ω. Analogous to (2.5), we re-evaluate the expressions for the net force on a repatom k to obtain the evolution of the positions and ω of the repatoms. For ease of notation, we refer to the "force" exerted on the mean-free parameter ω as the *thermal force*[4].

[4] This is a misnomer, as we do not calculate a force that evolves the temperature, but rather one that evolves ω. However, noting that our study is restricted to isothermal cases, and that this is a convenient notation that distinguishes the two forces, we stick to this terminology.

4.6.1 Physical force on a repatom

We combine our approaches from Chapter 2 and Section 4.4.1 to obtain the physical force on a repatom k. For the isothermal, quasistatic finite temperature problem, the expression is given by

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\frac{\partial}{\partial \boldsymbol{x}_k} F\left(\overline{\boldsymbol{q}^h}, \boldsymbol{\omega}^h, \boldsymbol{\theta}^h\right).$$

Using (4.26) and noting that only the potential remains after differentiation, the force expression simplifies to

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\sum_{a=1}^{N} \frac{\partial}{\partial \boldsymbol{x}_k} \langle V_a(\boldsymbol{q}^h)\rangle.$$

As only the positions $\boldsymbol{q}^h$ depend on the repatom positions $\boldsymbol{x}_k$, using the chain rule gives

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\sum_{a=1}^{N} \left[\sum_{j=1}^{N} \frac{\partial}{\partial \boldsymbol{q}_j^h} \langle V_a(\boldsymbol{q}^h)\rangle\right] \frac{\partial}{\partial \boldsymbol{x}_k}\left(\overline{\boldsymbol{q}_j^h} + \frac{\sqrt{2}\sigma_j}{\omega_j}\boldsymbol{x}_j\right),$$

where (4.29) is used to write $\boldsymbol{q}^h$ in terms of the mean position $\overline{\boldsymbol{q}^h}$. Now applying the interpolation from (2.3) results in

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\sum_{a=1}^{N} \left[\sum_{j=1}^{N} \frac{\partial}{\partial \boldsymbol{q}_j^h} \langle V_a(\boldsymbol{q}^h)\rangle\right] N_k(\boldsymbol{X}_j).$$

Introducing the energy-based formulation from (2.7) and restricting the potentials to be of the EAM type, we further approximate the expression by performing the potential calculation only at a set of N_s *sampling atoms*. Thus, the summation of all lattice sites over index a, is replaced with a reduced summation of sampling atoms over index α, as shown in Section 2.2. Rewriting the force expression with this reduced summation gives

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j \in \mathscr{S}_\alpha} \frac{\partial}{\partial \boldsymbol{q}_j^h} \langle V_\alpha(\boldsymbol{q}^h)\rangle N_k(\boldsymbol{X}_j).$$

Using numerical quadrature from Section 4.3, the final expression for the force on a repatom k becomes

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\sum_{\alpha=1}^{N_s} w_\alpha \sum_{j \in \mathscr{S}_\alpha} \sum_{p=1}^{n_{QP}} W_p \frac{\partial}{\partial \boldsymbol{q}_j} V_\alpha\left[\boldsymbol{q}_1(\boldsymbol{x}_1, \overline{\boldsymbol{q}}_1), \ldots, \boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N}, \overline{\boldsymbol{q}}_{n_N})\right] N_k(\boldsymbol{X}_j), \tag{4.39}$$

where we have absorbed the term $\left(\frac{1}{\sqrt{\pi}}\right)^{3n_N}$ into the weight $\tilde{W}_p$ by defining the new quadrature weight as $\tilde{W}_p = W_p / \left(\frac{1}{\sqrt{\pi}}\right)^{3n_N}$. In the following, we assume that for the quasicontinuum formulation, we deal with a reduced set of atoms, and will drop the superscript h from atomic positions for conciseness. As all of our studies are performed by EAM-based potentials, computational implementations require us to evaluate

$$\frac{\partial}{\partial \boldsymbol{q}_l} V_i = \frac{\partial}{\partial \boldsymbol{q}_l} \left(\frac{1}{2} \sum_{j \in \mathscr{S}_i} \Phi(r_{ij}) + \mathscr{F}(\rho_i) \right),$$

which gives the force on a site l for the EAM potential as

$$\frac{\partial}{\partial \boldsymbol{q}_l} V_i = \sum_{j \in \mathscr{S}_i} \left[\frac{1}{2} \Phi^{'}(r_{ij}) + \mathscr{F}^{'}(\rho_i) \rho^{'}(r_{ij}) \right] \frac{1}{r_{ij}} \left(\delta_{il} - \delta_{jl} \right) \cdot \left(\boldsymbol{q}_i - \boldsymbol{q}_j \right). \tag{4.40}$$

Substituting (4.40) back into (4.39), the force on a repatom for an EAM potential becomes

$$\boldsymbol{F}_k(\boldsymbol{x}) = -\sum_{p=1}^{n_{QP}} W_p \sum_{\alpha=1}^{N_s} w_\alpha \sum_{m \in \mathscr{S}_\alpha} \left[\frac{1}{2} \Phi^{'}(r_{\alpha m}) + \mathscr{F}^{'}(\rho_\alpha) \rho^{'}(r_{\alpha m}) \right] \frac{(\boldsymbol{q}_\alpha - \boldsymbol{q}_m)}{r_{\alpha m}} [N_k(\boldsymbol{X}_\alpha) - N_k(\boldsymbol{X}_m)]. \tag{4.41}$$

4.6.2 Thermal force on a repatom

Analogous to the derivation of the physical force on the repatom, we evaluate the expression for the thermal force on a repatom k as

$$\begin{aligned} T_k(s) &= -\frac{\partial}{\partial s_k} F(\overline{\boldsymbol{q}^h}, \boldsymbol{\omega}^h, \boldsymbol{\theta}^h), \\ &= -\sum_{a=1}^{N} \left[3k_B \theta_a \frac{\partial}{\partial s_k} \log \frac{\hbar \omega_a}{k_b m_a \theta_a} + \frac{\partial}{\partial s_k} \langle V_a(\boldsymbol{q}) \rangle \right]. \end{aligned}$$

Considering the summation rules as before and using (4.38), the expression simplifies to

$$T_k(s) = -\sum_{\alpha=1}^{N_s} w_\alpha \left[\frac{3k_B \theta_\alpha}{\omega_\alpha} N_k(\boldsymbol{X}_\alpha) + \frac{\partial}{\partial s_k} \langle V_\alpha(\boldsymbol{q}) \rangle \right].$$

To express the partial derivative with respect to s in a form that allows us to use (2.6), we consider the term

$$\frac{\partial}{\partial s_k} \langle V_\alpha(\boldsymbol{q}) \rangle = \sum_{j=1}^{N} \frac{\partial}{\partial \boldsymbol{q}_j} \langle V_\alpha(\boldsymbol{q}) \rangle \cdot \frac{\partial \boldsymbol{q}_j}{\partial s_k},$$

where we use the fact that the dependence on s_k comes implicitly from the positions q_j. Further evaluation gives

$$\frac{\partial}{\partial s_k}\langle V_\alpha(\boldsymbol{q})\rangle = -\sum_{j=1}^{N}\sqrt{2}\frac{\sigma_j}{\omega_j^2}\frac{\partial}{\partial \boldsymbol{q}_j}\langle V_\alpha(\boldsymbol{q})\rangle\cdot\mathbf{x_j}\mathbf{N_k}(\boldsymbol{X_j}).$$

Invoking the numerical quadrature from Section 4.3, the expression for the thermal force for EAM potentials on a repatom k becomes

$$\begin{aligned}T_k(s) = &-\sum_{\alpha=1}^{N_s} w_\alpha\left[\frac{3k_B\theta_\alpha}{\omega_\alpha}N_k(\boldsymbol{X}_\alpha)\right]+\\ &\sum_{p=1}^{n_{QP}}W_p\sum_{\alpha=1}^{N_s}w_\alpha\sum_{j\in\mathscr{S}_\alpha}\sqrt{2}\frac{\sigma_j}{\omega_j^2}\frac{\partial}{\partial \boldsymbol{q}_j}V_\alpha\left[\boldsymbol{q}_1(\boldsymbol{x}_1,\overline{\boldsymbol{q}}_1),\ldots,\boldsymbol{q}_{n_N}(\boldsymbol{x}_{n_N},\overline{\boldsymbol{q}}_{n_N})\right]\cdot\mathbf{x_j}\mathbf{N_k}(\boldsymbol{X_j}).\end{aligned} \tag{4.42}$$

Like in the case of the physical force, we evaluate (4.42) further for the case of EAM potentials. Recalling the general form of EAM potentials given by (2.11), a substitution into (4.42) gives the thermal force on a repatom k as

$$\begin{aligned}T_k(s) = &-\sum_{\alpha=1}^{N_s} w_\alpha\left[\frac{3k_B\theta_\alpha}{\omega_\alpha}N_k(\boldsymbol{X}_\alpha)\right]+\\ &\sum_{p=1}^{n_{QP}}W_p\sum_{\alpha=1}^{N_s}w_\alpha\left(\sqrt{2}\frac{\sigma_\alpha}{\omega_\alpha^2}\sum_{m\in\mathscr{S}_\alpha}\left[\frac{1}{2}\Phi'(r_{\alpha m})+\right.\right.\\ &\qquad\left.\mathscr{F}'(\rho_\alpha)\rho'(r_{\alpha m})\right]\frac{(\boldsymbol{q}_\alpha-\boldsymbol{q}_m)}{r_{\alpha m}}\cdot\mathbf{x}_\alpha\mathbf{N_k}(\boldsymbol{X}_\alpha)+\\ &\qquad\sum_{m\in\mathscr{S}_\alpha}\sqrt{2}\frac{\sigma_m}{\omega_m^2}\left[\frac{1}{2}\Phi'(r_{\alpha m})+\right.\\ &\qquad\left.\left.\mathscr{F}'(\rho_\alpha)\rho'(r_{\alpha m})\right]\frac{-(\boldsymbol{q}_\alpha-\boldsymbol{q}_m)}{r_{\alpha m}}\cdot\mathbf{x_m}\mathbf{N_k}(\boldsymbol{X_m}),\right).\end{aligned} \tag{4.43}$$

We note that while (4.43) appears more involved, its form is similar to (4.41).

4.7 Uniform Thermal Expansion

We now present the results from numerical simulations that were performed with the finite-temperature QC framework derived in the previous sections. We begin with a fundamental yet important test of the finite-temperature implementation, the uniform thermal expansion test. Here, we model the change in the relaxed configuration of a system, when it is subjected to uniform temperature. Such a study has

already been performed by Kulkarni [59] and Venturini et al. [120] for the force-based max-ent hot-QC formulation, but we perform it for the first time for a fully non-local energy-based formulation at finite temperature. In addition, we use summation rules introduced by Amelang, Venturini, and Kochmann [5], and also model our hot-QC formulation using both third- and fifth-degree quadrature, whereas previous studies have been restricted to the third-degree quadrature. This study is also performed using a variety of interatomic potentials listed in Table 4.2, providing an extensive comparison of the performance of these potentials that has previously not been reported.

We perform thermal expansion tests on FCC copper and aluminum, and we observe the *linear thermal expansion*, which we will do by plotting the relative lattice parameter as a function of temperature. We note that the thermal expansion coefficient α is the instantaneous slope of this curve, but we restrict our observations to the curve itself. Thus, we plot ε as a function of temperature T, where ε is given as

$$\varepsilon = \frac{a}{a_0} - 1, \tag{4.44}$$

where a_0 is the lattice parameter at a given temperature T_0 and a is the lattice parameter at temperature T. This is consistent with [60] and allows us to compare our results with previous results. As noted before, the thermal expansion coefficient is

$$\alpha(T) = \frac{\partial \varepsilon}{\partial T}.$$

Here, we show the effects of uniform thermal expansion using two types of simulations.

4.7.1 Infinite crystal simulation

In this simulation, we assume that the atoms lie in an infinite FCC single crystal, which means every atom sees the exact same neighborhood. This setting is identical to the full-periodic simulations shown in [60], [120]. However, we formulate our system differently from those by modeling the neighborhood of exactly one atom. This changes the governing equations, as the only mechanical degree of freedom in the system is the lattice parameter for an FCC system. A detailed derivation of the new set of equilibrium equations that we obtain is shown in Appendix A.

4.7.2 Finite large-scale simulation

In addition to the approach in Section 4.7.1, we observe thermal expansion without changing the governing equations as done in Appendix A.1. Another way to

Name	r_{cutoff}	a_0	Reference
Copper			
Johnson	3.249 A	3.609 A	Johnson [50]
EFS	4.320 A	3.609 A	Dai et al. [29]
Sutton Chen	3.249 A	3.610 A	Sutton and Chen [109]
Mishin	5.507 A	3.615 A	Mishin et al. [73]
Mendelev	6.000 A	3.636 A	Mendelev and King [68]
Aluminum			
Winey Kubota Gupta	6.365 A	4.025 A	Winey, Kubota, and Gupta [129]
Zhou	6.404 A	4.050 A	Zhou, Johnson, and Wadley [136]
Mishin	6.287 A	4.050 A	Zope and Mishin [137]
Mendelev	6.500 A	4.045 A	Mendelev et al. [69]

Table 4.2: List of the interatomic potentials for Cu and Al used for all simulations with their cut-off radius and the lattice parameter provided to construct an FCC lattice.

approximate the same response as the infinite case but by modeling a finite set of atoms is to model a large system of atoms and measure the response away from the surfaces, thereby mitigating the errors introduced by surface effects. For this set of simulations, we model an $8 \times 8 \times 8$ cube of an FCC material which consists of $2,547$ atoms. After setting a uniform temperature, we let the system relax and measure the volumetric expansion of a small subset of this cube, and take a cube root of this volume to obtain the linear expansion. We use this measure instead of a in (4.44) and plot the results.

4.7.3 Interatomic potentials

To test the performance of the interatomic EAM potentials with finite temperature, we chose a variety of potentials from the Interatomic Potentials Repository Project (IPR) [75]. Our studies are performed with FCC metals Al and Cu, and Table 4.2 lists the various potentials used for the simulations discussed in this chapter.

4.7.4 Results

We show the results from simulations described in Sections 4.7.1 and 4.7.2. The plots show ε as a function of temperature (4.44) by comparing a variety of cases : (*i*) Experimental data for single-crystal Al and Cu from Nix and MacNair [76], (*ii*) infinite crystal simulation using both third- and fifth-degree Gaussian quadrature, (*iii*) finite large-scale simulations with volume calculations using a $1 \times 1 \times 1$ cell, a $2 \times 2 \times 2$ cell and a $4 \times 4 \times 4$ cell extracted from the center of the simulation

cube, (*iv*) results from [59] for the Johnson potential [50], (*v*) results from MD simulations performed using the LAMMPS [101] with the OpenKim database [34]. We use a $T_0 = 273K$ as the reference temperature to match the experimental data.

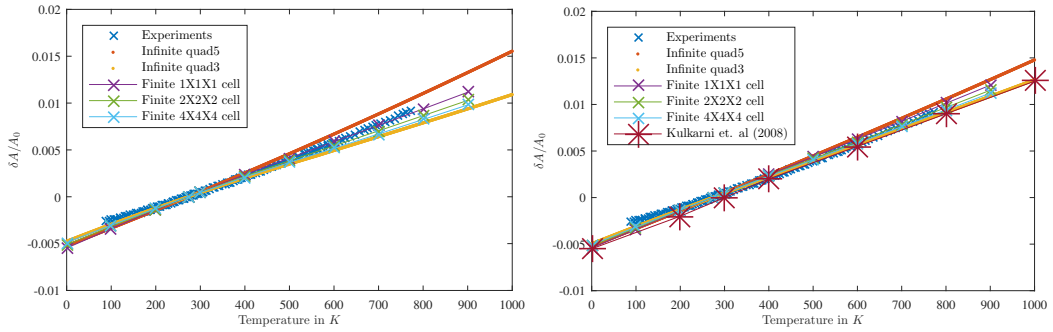

Figure 4.1: Uniform thermal expansion of single crystal FCC copper using the EFS potential [29] to the left and the Johnson potential [50] to the right. Experimental data from [76].

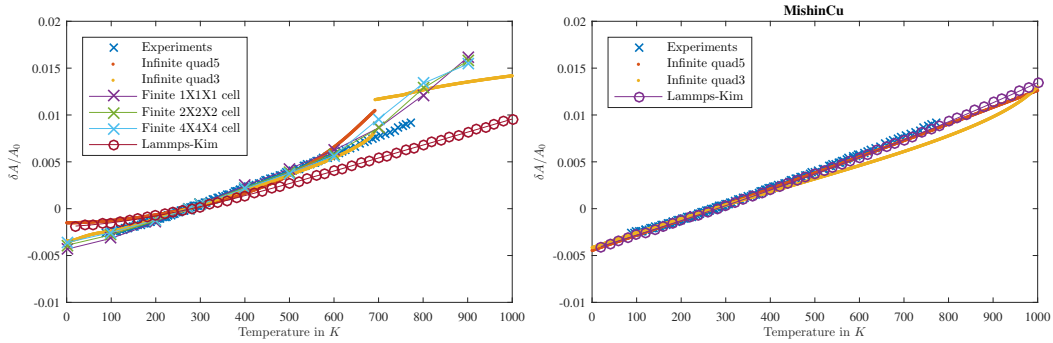

Figure 4.2: Uniform thermal expansion of single crystal FCC copper using the Mendelev potential [68] to the left and the Mishin potential [73] to the right. Experimental data from [76].

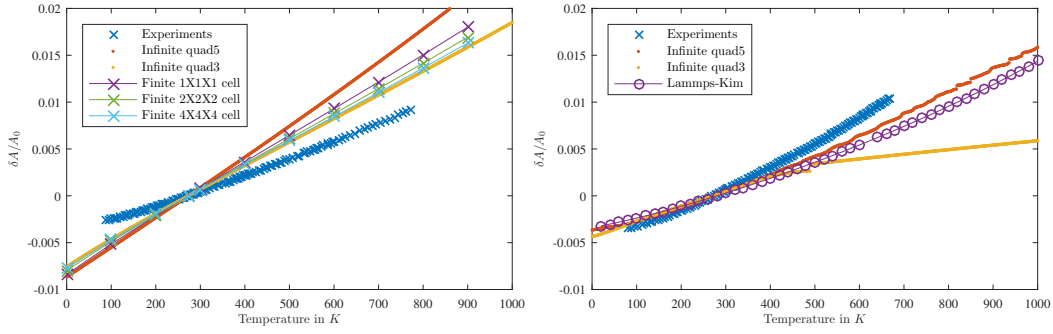

Figure 4.3: Uniform thermal expansion for a single crystal FCC lattice. Left panel: Cu using the Sutton Chen potential [109]. Right panel: Al using the Mishin potential [137]. Experimental data from [76].

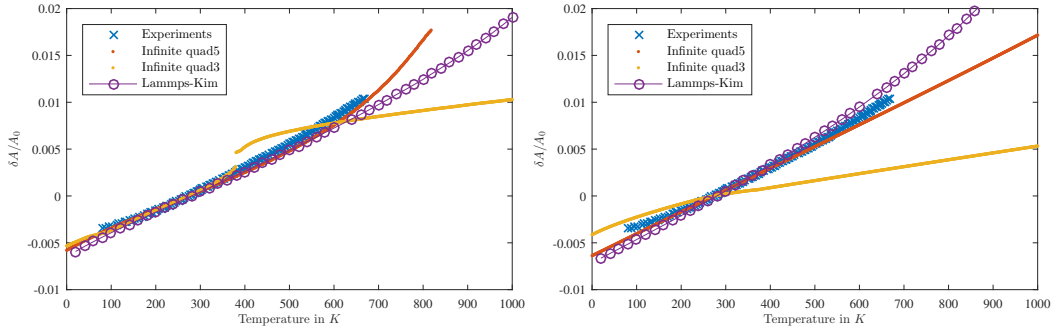

Figure 4.4: Uniform thermal expansion of single crystal FCC aluminum using the Mendelev potential [69] on the left and Winey Kubota Gupta[129] on the right.Experimental data from [76].

Infinite vs finite: We observe in Figures 4.1, 4.2, and 4.4 that both the infinite- and finite-crystal simulations provide similar trends for all the copper potentials. This shows that despite solving a completely different formulation, the relaxed value of the lattice parameter increases at the same rate for both. In addition, for the finite-crystal simulation, the results for the $(1 \times 1 \times 1)$, $(2 \times 2 \times 2)$, and $(4 \times 4 \times 4)$ are almost identical, which shows that the interior of the cube expands uniformly and does not encounter surface effects.

Third- and fifth-degree quadrature: An important comparison to observe is the one between the two types of numerical quadrature implemented. As the fifth-degree quadrature is computational intensive, this study is performed only with the infinite-crystal simulation. Figures 4.1, 4.2, 4.3, and 4.4 all show the comparison between the third- (quad3) and fifth-degree (quad5) quadrature for various potentials. We observe that quad3 generally under-predicts the value of thermal expansion for all the cases; however, for EFS, Johnson and Mishin potentials in copper, the difference between the two quadrature rules is negligible. However, in aluminum, the quad5 performs significantly better for the Sutten-Chen, Mishin, Mendelev and Winey Kubota Gupta. It better matches the experiments and also the MD calculations. This shows that despite the computational overhead associated with fifth-degree quadrature, depending on the choice of potential, it is better to use quad5 over quad3 for correctness.

Comparison to LAMMPS: To compare the effectiveness of the max-ent formulation with actual MD results, we use LAMMPS with the OpenKim database [34].

Comparisons are only made for the exact same potential files used in both [34] and our hot-QC formulation. The MD results are performed using the tests provided by Wen [127]. Figures 4.2, 4.3, and 4.4 show the comparisons of the MD data with hot-QC. We observe that while our quad3 results do not agree well with the LAMMPS data (with the notable exception of the Mishin potential [137]), the quad5 results match well for every potential that is compared. This shows that the approximation introduced by considering phase averages and the max-ent formulation to solve the governing equations of motion, performs well only if the phase averages are computed with higher degree quadrature. Other possible sources of deviation from MD results include differences in the MD solver and QC solver and convergence criteria for the simulation.

Comparison to Experiments: For most potentials, the thermal expansion results match experimental data from [76], with the exception of Sutton Chen for copper. Some deviations are expected due to the difficulty in manufacturing pure metals for experimental purposes. In addition to experimental data, we obtain a good comparison with Kulkarni [59] for the Johnson potential for copper.

The above results show the importance of higher order quadrature and a good choice of potential in order to accurately model a system using hot-QC. Looking at the overall comparisons with both experiments and MD, we conclude that this hot-QC formulation captures lower-temperature thermal expansion better, and the third-degree quadrature rule deviates at higher temperatures. These thermal expansion results provide a basis for selection of a temperature range while choosing a potential for future applications. Despite a close match of the fifth-order quadrature hot-QC with MD, the increasing mismatch at higher temperatures for some potentials suggests limitations for this model to replicate results from atomistics.

4.8 Finite Temperature Elastic Constants

After the promising results of thermal expansion tests, we test our finite-temperature QC formulation in its ability to approximate long-range elastic fields by using coarse-graining.

4.8.1 Theory

Phillips [86] and Venturini [119] give a comprehensive overview of the theory behind elastic constants calculations using a microscopic approach, so a short paragraph is provided here.

If we restrict the deformation in the system to small deformations, then linear elasticity theory says that the strain energy stored in the system undergoing a strain ε can be expressed using the elastic energy density w as

$$w(\varepsilon) = \frac{1}{2} \sum_{i,j,k,l} \mathscr{C}_{ijkl} \varepsilon_{ij} \varepsilon_{kl}. \tag{4.45}$$

We note that (4.45) is a continuum expression, but can also be transformed into the atomistic degrees of freedom by writing the strain energy in the form

$$w(\varepsilon) = \frac{1}{V} \left(F(\overline{\boldsymbol{q}}, \varepsilon) - F(\overline{\boldsymbol{q}}, 0) \right), \tag{4.46}$$

where $F(\overline{\boldsymbol{q}}, \varepsilon)$ is the Helmholtz free energy in the deformed state of a unit cell and $F(\overline{\boldsymbol{q}}, 0)$ is the free energy of a perfect unit cell. The justification of equating (4.45) and (4.46) is that both represent the energy per unit volume. By expanding the Helmholtz free energy about the reference state ($\boldsymbol{\varepsilon} = \mathbf{0}$), only considering terms up to the second order, and noting that the linear term vanishes about the reference state, one obtains [86, 119]

$$\mathscr{C}_{ijkl} = \frac{1}{V} \frac{\partial^2 F}{\partial \varepsilon_{ij} \partial \varepsilon_{kl}} \bigg|_{\boldsymbol{\varepsilon}=\mathbf{0}} = \frac{\partial^2 w}{\partial \varepsilon_{ij} \partial \varepsilon_{kl}} \bigg|_{\boldsymbol{\varepsilon}=\mathbf{0}}. \tag{4.47}$$

Writing the strain tensor $\boldsymbol{\varepsilon} = \gamma \xi$, where γ is a scalar measure of the applied strain and ξ is a symmetric transformation matrix pertaining to a given mode of deformation, (4.47) simplifies to

$$\frac{1}{V} \frac{\partial^2 F}{\partial \gamma^2} \bigg|_{\gamma=0} = \frac{\partial^2 w}{\partial \gamma^2} \bigg|_{\gamma=0}. \tag{4.48}$$

Venturini [119] showed that by using thermodynamic relations at constant temperature, and applying appropriate transformation matrices, one can recover the three fundamental elastic constants of a fcc material: bulk modulus κ, uniaxial modulus C_{11}, and shear modulus $\mu = C_{44}$. The expressions for these moduli are given as

$$\kappa = \frac{1}{9V} \frac{\partial^2 F}{\partial \gamma^2} \bigg|_{\gamma=0} \text{ with } \xi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{4.49}$$

for the bulk modulus,

$$\mu = \frac{1}{4V} \frac{\partial^2 F}{\partial \gamma^2} \bigg|_{\gamma=0} \text{ with } \xi = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{4.50}$$

for the shear modulus, and

$$C_{11} = \frac{1}{V} \frac{\partial^2 F}{\partial \gamma^2}\bigg|_{\gamma=0} \quad \text{with } \xi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{4.51}$$

for the uniaxial modulus.

4.8.2 Scenario setup

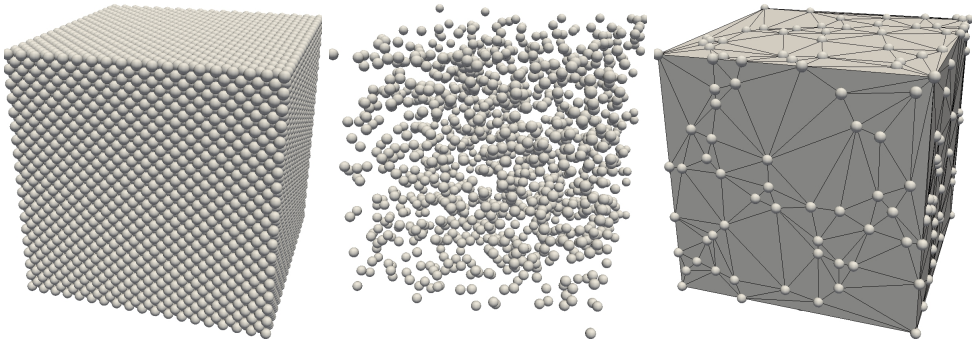

Figure 4.5: The initial setup of the ElasticConstants scenario for single-crystal fcc Cu. The left panel shows the fully atomistic case containing 25357 atoms. The center panel shows the coarse-grained scenario with 5% repatom density with a total of 1274 atoms used to simulate the entire region. The right panel shows the mesh formed by this coarse set of atoms.

This simulation achieves two main goals. We identify the errors associated with coarse-graining as a function of temperature. To this end, we will evaluate the elastic constants of single-crystal copper at various temperatures for two sets of systems: a *fully resolved* system where we model all degrees of freedom, and a *coarse-grained* system where we heavily reduce the number of degrees of freedom. Considering the fully resolved values as the "exact solution", we compare them to the coarse-grained results at various temperatures and plot the relative errors associated with each of the elastic constants. The other goal is to compare the calculated values of elastic constants to experimental observations, which aims for a validation of the finite-temperature nonlocal QC formulation. To this end, we plot the elastic constants obtained by our hot-QC model as a function of temperature and compare them to experimental observations for single-crystal Cu.

The scenario is set up is as follows. A three-dimensional QC representation of single-crystal fcc Cu, using the Finnis-Sinclair potential of [29], is set up and affine boundary conditions are applied remotely to a large crystalline ensemble. The reduced set of repatom locations is chosen *randomly*, which provides an extremely

robust test to the capabilities of the QC formulation. As a representative example, we choose a repatom density of 5%, where repatom density is defined as the ratio of the number of repatoms to the total number of lattice sites contained in the sample. This represents a rather coarse setup, as shown in Figure 4.5. The random choice of repatoms affects the results, and so 20 different realizations for the same value of repatom density are performed, and an average is taken. Finally, the same simulation is performed using a density of 100% (representing the exact solution) and the elastic constants obtained from this simulation are used to calculate the errors shown by the coarse simulations.

4.8.3 Results

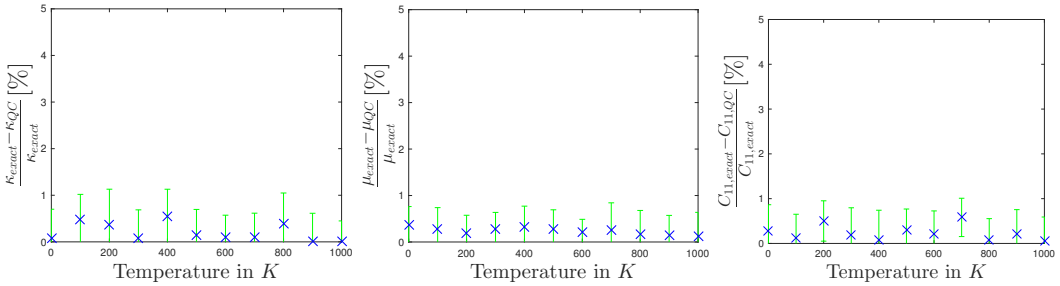

Figure 4.6: Results of the elastic constant simulations: (a) relative error in the bulk modulus (b) relative error in the shear modulus (c) relative error in the uniaxial modulus, plotted for a range of temperatures, for a repatom density of 5%. The relative error is measured with respect to the exact values obtained by full atomistics (repatom density 100%).

Figure 4.6 shows the results of the elastic constants simulations. Our coarse-grained QC model shows low errors (less than 1%) for all three elastic constants. It is also worth noting that the errors do not change as we increase the temperature in the simulations. As mentioned in Section 4.8.2, the exact values are determined by using the fully-resolved atomistic ensemble to obtain the values of elastic constants. Figure 4.7 shows the comparison of our calculated elastic constants with known experimental values, and other atomistic simulations using MD. The experimental data is obtained from Overton and Gaffney [81] and Chang and Himmel [20]. We notice a slight offset in the values compared to the experiments, but overall we capture the expected softening with increase in temperature quite well. For both the bulk modulus and the uniaxial modulus, the offset is maintained throughout the temperature range. [29] provide values for elastic constants, which we reproduce exactly for $0K$. It is interesting to note that the elastic constants of [29] also have the offset that we observe in our coarse-grained QC simulations. This can be attributed

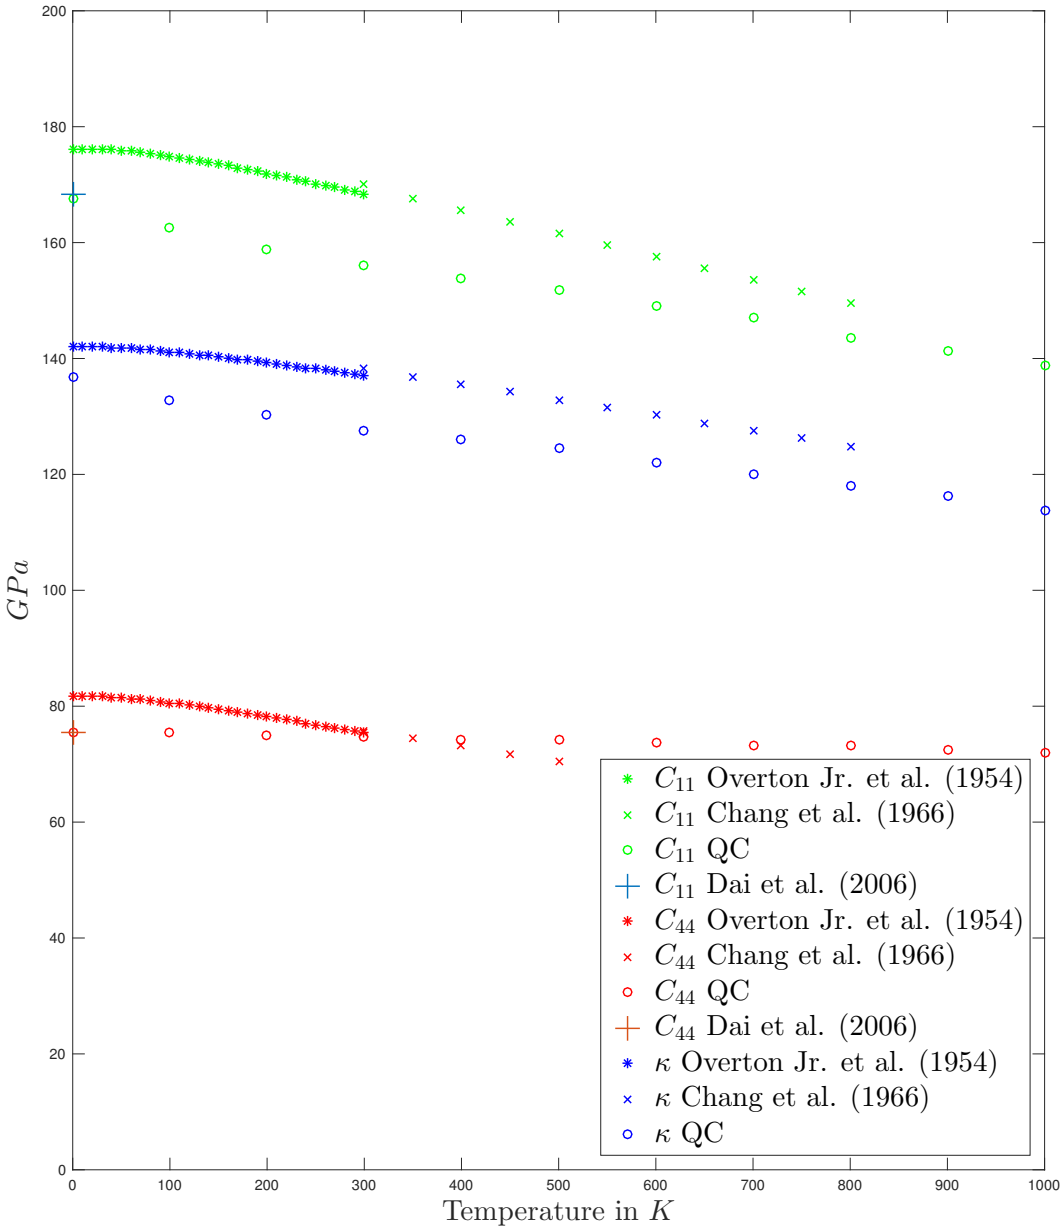

Figure 4.7: Results of the elastic constant simulations on copper: Comparison of the bulk modulus, shear modulus, and uniaxial modulus to various experimental results plotted at a range of temperatures. Experimental results from [81], [20]

to the effect of fitting parameters while designing the potential, such that quantities like elastic constants match actual experiments at room temperature. So, the reason we maintain this offset could be due to the fact that the $0K$ simulations performed by coarse-grained QC, while correctly reproducing MD, assume that the parameters provided by the potential are for $0K$.

4.9 Finite Temperature Force Artifacts

In Section 1.2 we describe the various approaches to multiscale modeling and one of the main challenges faced during scale bridging are spurious force artifacts [5], [10], [89]. Section 2.2.1 outlines this problem for the energy-based QC method when we evaluate the expression for the force on a repatom. Force artifacts are non-physical forces which change the behavior of the system and give erroneous results. Depending on the choice of sampling atoms and weights, different summation rules in QC give rise to different magnitudes of force artifacts. Amelang, Venturini, and Kochmann [5] have studied various summation rules in detail to quantify the ghost force errors at $0K$. Their results show conclusively that the optimal summation rules outperform all others in their ability to greatly minimize errors due to force artifacts. While the errors due to physical forces at $0K$ have been quantified, force artifacts due to physical and thermal forces at higher temperatures have not been studied before. We perform a study similar to [5] to investigate force artifacts at finite temperatures.

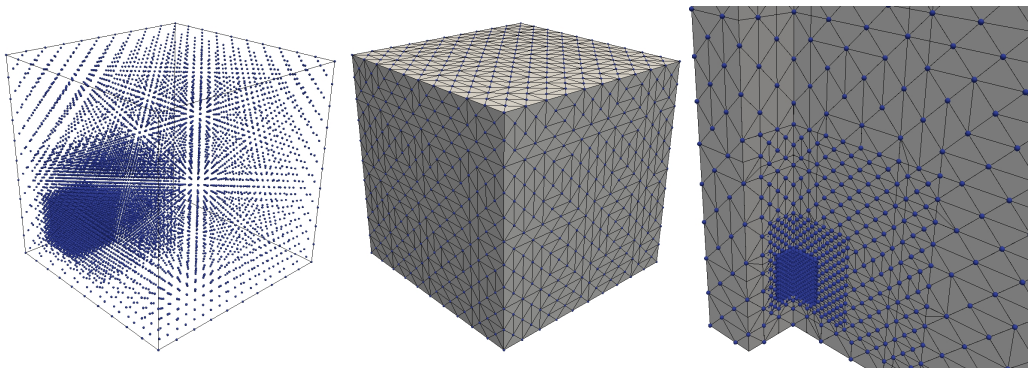

Figure 4.8: The initial setup to observe force artifacts across interfaces in fcc Cu single-crystal modeled using the EAM potential by Dai et al. [29]. The domain is coarsened in four layers and contains 24 million atoms that are represented by 16528 repatoms. The left panel shows all repatoms used in this simulation. The center panel shows the mesh constructed by these repatoms. The right panel shows a zoomed view of a clip of this mesh that shows the four interfaces formed by the four layers of coarsening.

4.9.1 Scenario setup

The main goal we achieve in this study is to quantify the errors produced by physical force artifacts and force artifacts on the mean-free parameter ω, which we have termed thermal forces, at finite temperatures, due to the effect of coarse-graining of the domain of study. In addition, we also test the effectiveness of the optimal

summation rule to minimize these force artifacts, and form comparisons with other summation rules used in other QC formulations. To this end, we construct a cube of fcc Cu having sharp and diffused interfaces due to rapid coarse-graining from full atomistic resolution to a coarse region in factors of 2 in a unified setting that tests all the summation rules with every other setting being identical. We use the EAM potential by [29], and the study is performed at $100K$ as a representative example. The summation rules used for the comparative study are as follows: (*i*) *nodal summation rules* where the sampling atom locations coincide with the repatom locations, (*ii*) *quadrature-type summation* [43, 132], (*iii*) *node based cluster summation* [54, 47], and (*iv*) *optimal summation* [5]. For each of these summation rules, we use both the third- and fifth-degree quadrature from Section 4.3.1.

To effectively calculate force artifacts, no boundary conditions are applied to the cube, and all representative atoms of the cube are constructed using the equilibrium lattice spacing and equilibrium ω for an infinite crystal obtained from results of Section 4.7. In the absence of any force artifacts, any force calculation with this setting should yield no net physical or thermal force on the repatoms, except at the surfaces. Therefore, performing a force calculation and plotting the results at repatom locations shows the force artifacts for an *undeformed, non-equilibrated* sample. As the aim of this study is to observe force artifacts arising purely due to coarse-graining, force artifacts on the surface are not considered.

4.9.2 Results

For the third-degree quadrature, Figure 4.9 shows the comparison of the physical forces and thermal forces for all the summation rules considered for sharp interfaces, and Figure 4.11 shows the same for the diffuse interfaces. We note that the thermal force is a scalar quantity, but to show a visual comparison, a vector in the y direction is shown with its magnitude equal to the magnitude of the thermal force.

For both the physical and thermal force artifacts, we define a summation rule to be effective if the magnitude of force artifacts at interfaces is small and if these artifacts are observed at fewer repatom locations. Thus, a summation rule is considered to be more effective than another if it has a lower average magnitude of spurious forces and has these artifacts at fewer locations than the other. We use the $>$ sign with the definition above for comparing effectiveness of the summation rules that are studied.

For the physical forces, our results at $100K$ show the same trends as those seen in

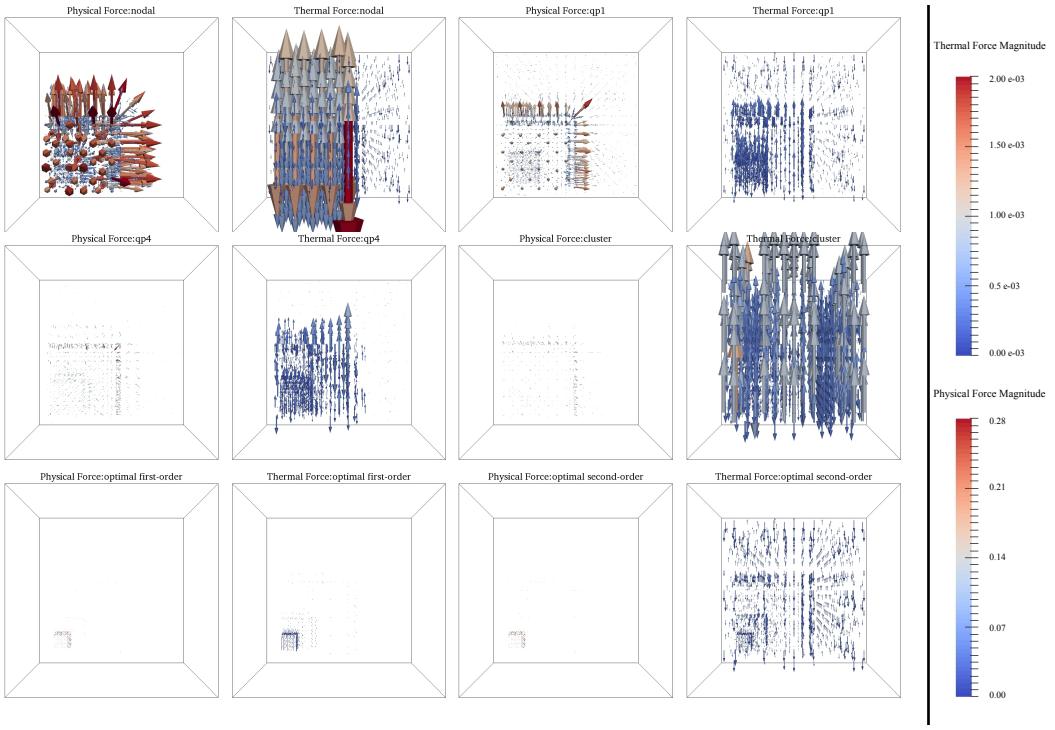

Figure 4.9: Results showing the physical and thermal force artifacts on all repatoms from the system shown in Figure 4.8. The labels above the box indicate the force that is observed and the summation rule used for that simulation. The nodal summation rules are labeled as "nodal", the cluster-based summation rules are "cluster", and the "qp1" and "qp4" labels represent the quadrature-type summation with 1 and 4 quadrature points respectively. This figure shows the results using *sharp interfaces* and by using the *third-degree gaussian quadrature* to approximate the potential energy.

[5] for $0K$. The optimal summation rules outperform all the other rules and the general order of the effectiveness of the summation rules is: *optimal second-order* $>$ *optimal first-order* $>$ *cluster summation* $>$ *quadrature summation with 4 points* $>$ *quadrature summation with 1 point* $>$ *nodal summation*. This order mirrors the results from [5] and shows that the hot-QC formulation does not contribute any additional force artifacts to the physical forces. The thermal forces follow a similar pattern, with in general the diffuse interfaces showing a higher magnitude than the sharp interfaces. This is expected as the sharp interfaces provide a certain amount of symmetry that helps mitigate the spurious artifacts. However, for the thermal forces the order of effectiveness of the summation rules is slightly different from the physical forces: *optimal first-order* $>$ *optimal second-order* $>$ *quadrature summation with 4 points* $>$ *cluster summation* $>$ *quadrature summation with 1 point* $>$ *nodal summation*.

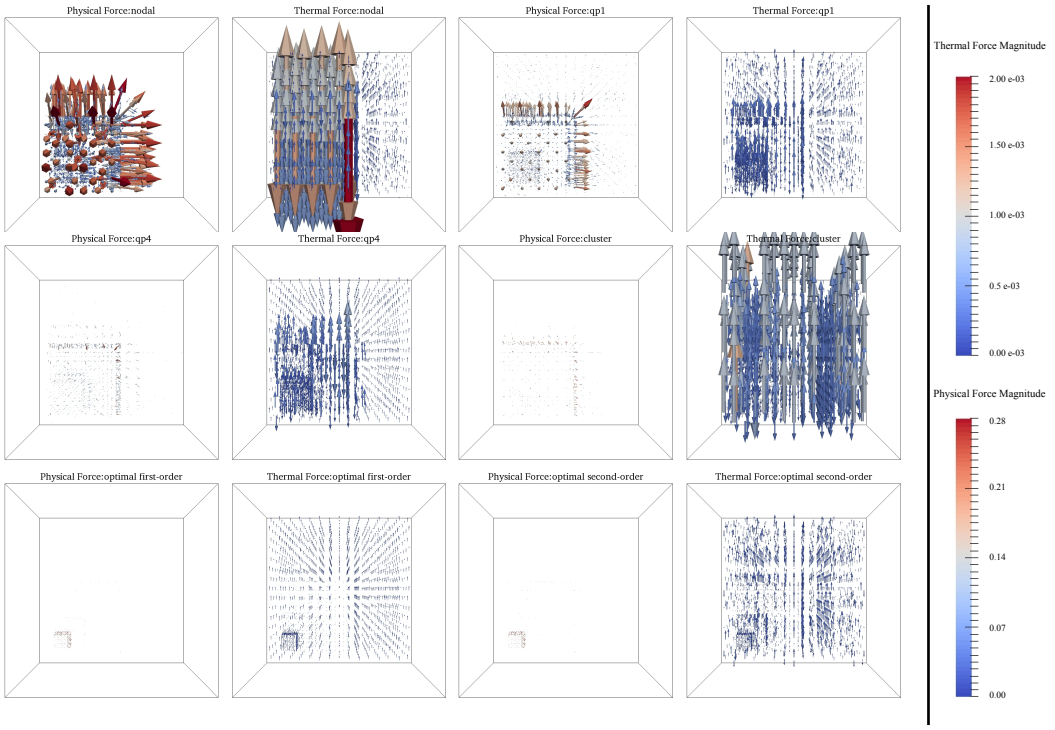

Figure 4.10: Results showing the physical and thermal force artifacts on all repatoms from the system shown in Figure 4.8. The labels above the box indicate the force that is observed and the summation rule used for that simulation. The nodal summation rules are labeled as "nodal", the cluster-based summation rules are "cluster" and the "qp1" and "qp4" labels represent the quadrature-type summation with 1 and 4 quadrature points respectively. This figure shows the results using *sharp interfaces* and by using the *fifth-degree gaussian quadrature* to approximate the potential energy.

Figures 4.10 and 4.12 show the results of the ghost force calculations for the fifth-degree quadrature. For better comparing and showing how effective the optimal summation rules are, the scale on the physical force and the thermal force is kept the same for both the third and fifth-degree quadrature plots. The results look similar to those of the third-degree quadrature, however the optimal first-order and second-order force artifacts are much closer for the fifth-degree quadrature than the third-degree. The other summation rules show nearly similar magnitudes for the different quadratures. The two non-intuitive results from this simulation are (*i*) the better performance of the optimal first-order as compared to the second-order for the thermal force artifacts, and (*ii*) the relatively poor performance of cluster-based summation for the thermal ghost forces.

One possible explanation for this behavior is the assignment of weights and the

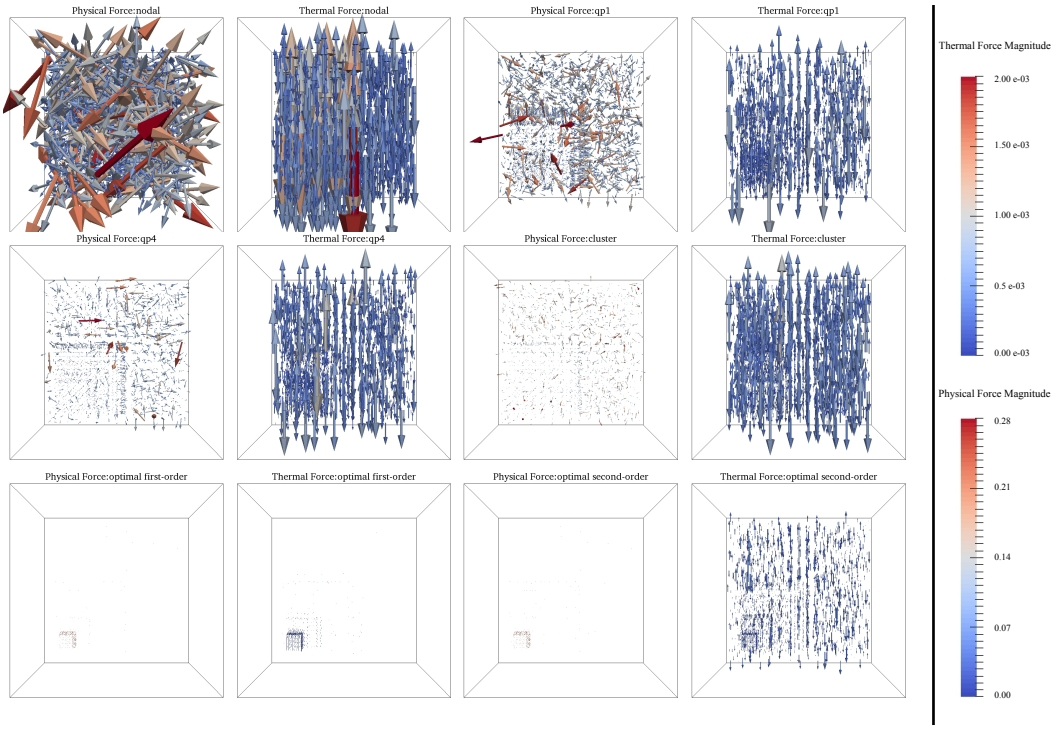

Figure 4.11: Results showing the physical and thermal force artifacts on all repatoms from the system shown in Figure 4.8. The labels above the box indicate the force that is observed and the summation rule used for that simulation. The nodal summation rules are labeled as "nodal", the cluster-based summation rules are "cluster" and the "qp1" and "qp4" labels represent the quadrature-type summation with 1 and 4 quadrature points, respectively. This figure shows the results using *diffuse interfaces* and by using the *third-degree Gaussian quadrature* to approximate the potential energy.

sensitivity of the thermal forces to the exact value of ω. We expect the thermal force to be zero only if the given value of ω of the repatom is exact. Any deviation will result in the right side of (4.36) to not be zero and give rise to a spurious force. As sampling atoms on the edge of elements for the optimal second-order summation rule have significantly higher weights as compared to the first-order (see [5] for details), a small deviation from the correct ω causes a larger force for second-order sampling atoms than the first-order ones, which is what we observe in our results.

4.10 Conclusions

In this chapter, we have formulated a finite temperature extension of the nonlocal quasicontinuum method outlined in Chapter 2 using the *max-ent* approach. We first derived the governing equations of motion for a fully atomistic ensemble, then formed the energy-based QC approximation by introducing repatoms and sam-

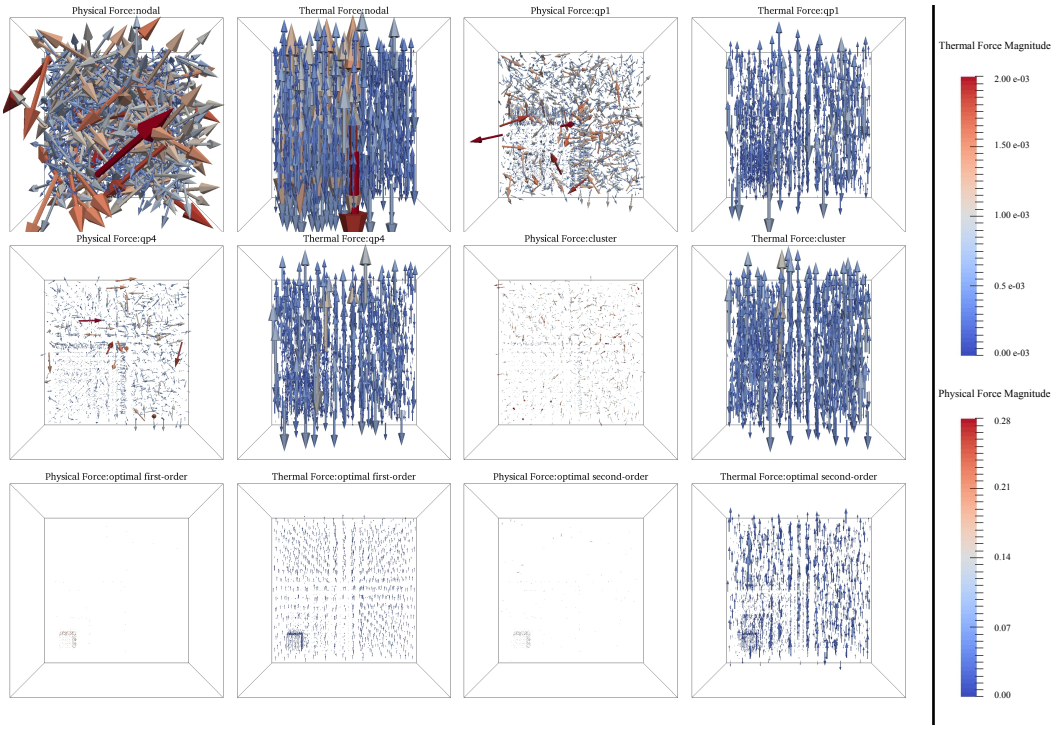

Figure 4.12: Results showing the physical and thermal force artifacts on all repatoms from the system shown in Figure 4.8. The labels above the box indicate the force that is observed and the summation rule used for that simulation. The nodal summation rules are labeled as "nodal", the cluster-based summation rules are "cluster" and the "qp1" and "qp4" labels represent the quadrature-type summation with 1 and 4 quadrature points, respectively. This figure shows the results using *diffuse interfaces* and by using the *fifth-degree Gaussian quadrature* to approximate the potential energy.

pling atoms, and approximating the potential energy by using third- and fifth-degree Gaussian quadrature. This hot-QC formulation is validated by performing uniform thermal expansion tests by constructing simulations modeling both large-finite and infinite single-crystal fcc lattices of Cu and Al. The results show good agreement with previous reports and also with MD calculations using LAMMPS. We then show that our QC formulation recovers the elastic constants of pure Cu almost exactly by using only 5% repatoms from the total atomistic ensemble. These elastic constants also match experimental trends by showing softening with temperature, and having a similar dependence on temperature. The small discrepancy in the values can be attributed to parameter fitting of the potential used [29] in which they match the values of elastic constants at room temperature for a $0K$ atomistics model. Lastly, we analyze spurious forces in various summation rules from QC literature at finite temperatures. We observe that, like the physical forces, coarse-graining

results in spurious force artifacts on the mean-free parameter ω. We quantify these forces and compare their magnitudes for various summation rules and show that the optimal summation rules introduced by Amelang, Venturini, and Kochmann [5] minimize the artifacts more effectively than the cluster-based and quadrature based summations.

Chapter 5

INVESTIGATION OF SYMMETRIC TILT GRAIN BOUNDARIES

All the steps presented in Chapters 2, 3, and 4 have been taken to enable the study of severely deforming, finite-temperature polycrystalline materials. In this chapter, we focus on the most fundamental building block of polycrystals, viz. grain boundaries. Instead of modeling large networks of grains, we aim to demonstrate the applicability of the techniques introduced above for the most fundamental example: symmetric-tilt grain boundaries. First, we provide a brief background and motivation for atomistic modeling of grain boundaries. We then discuss the construction of several fully 3D symmetric-tilt grain boundaries and calculate their relaxed energies at $0K$. Next, we extend this study to finite temperature using our approach from Chapter 4 and comment on the trends of the relaxed structures and energy as a function of temperature. We finally close the chapter with a future outlook towards potential applications of our grain boundary modeling framework.

5.1 Background

As introduced in Chapter 1, some of the main influences of mechanical properties at the microscale of crystalline materials are material defects like grain boundaries, dislocations, vacancies, voids, etc. One of the most prominent relations showing the influence of grain size on material strength is the Hall-Petch relation [44, 85] given as

$$\sigma_y = \sigma_0 + \frac{k_y}{\sqrt{d}}, \tag{5.1}$$

where σ_y is the yield stress of the material, σ_0 is the starting stress for dislocation motion (material constant), k_y is the strengthening coefficient (material constant), and d is the average grain size. The Hall-Petch relation predicts that a decrease in grain size increases the yield stress (a measure of strength) of a material and generally works until the average grain size reaches about 10 nm, after which a reverse effect is observed (also termed as the "inverse Hall-Petch relation" [28, 65, 74]). The two relations are qualitatively shown in Figure 5.1.

Grain boundaries and dislocations play a very important role in characterizing the microstructure of a material, which influences the material strength, deformation

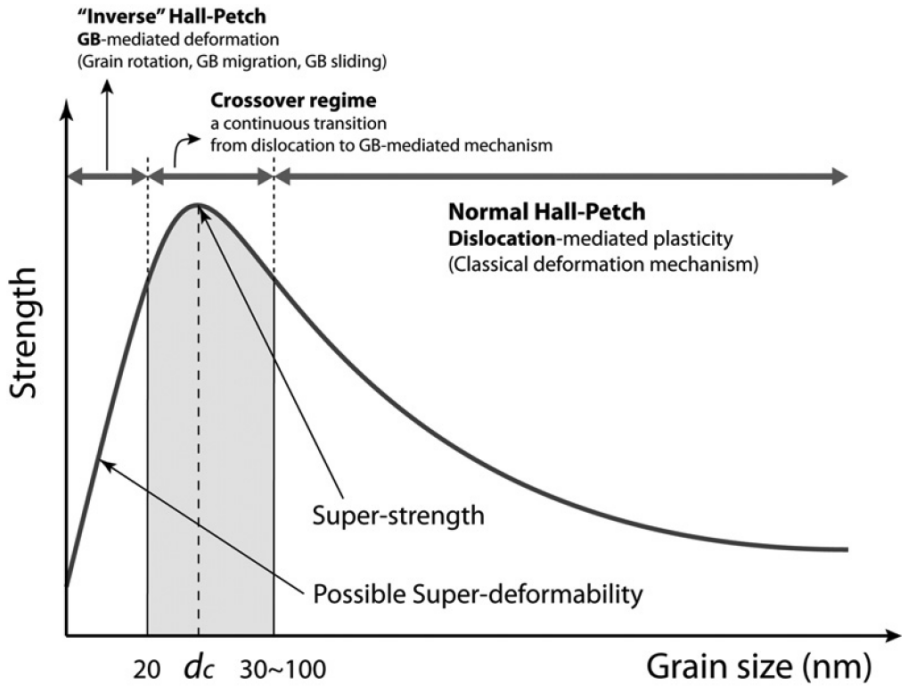

Figure 5.1: A diagram showing the Hall-Petch and inverse Hall-Petch effect in polycrystalline materials, including the commonly accepted deformation mechanisms explaining the two regimes [42].

mechanisms [8, 131], toughness, and other properties of importance. This has resulted in a lot of active studies in the modeling of grain boundaries of pure metals and metal alloys, as it is extremely difficult to synthesize them experimentally without contamination, and because of the advent of microstructure modeling using extremely accurate atomistic simulations (see e.g. [22]).

5.1.1 Crystallographic description

In basic terms, a *grain boundary* in a solid crystalline material is a region separating two crystals (grains) of the same phase where the grains differ in crystallographic orientation. The grain boundary of a *bicrystal* (a sample containing two grains with a planar separating interface) can be fully characterized by *five independent parameters* or macroscopic degrees of freedom (DOF). The first three specify the mutual misorientation of the adjoining grains, which is represented by a rotation described by *(i)* a rotation axis $\mathbf{o}$ (2DOF), and *(ii)* a rotation angle θ (1DOF). The last two DOFs are the orientation of the grain boundary inbetween these misoriented grains (e.g., described by the normal to the grain boundary plane $\mathbf{n}$). The different degrees of freedom are shown in Figure 5.2.

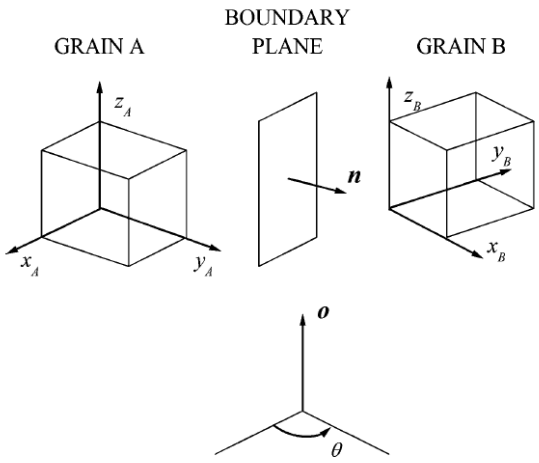

Figure 5.2: The variables that define a grain boundary [83].

Grain Boundaries (GB) are categorized into different groups according to the relationships among different DOFs: *(i) tilt GB*: $\mathbf{o} \perp \mathbf{n}$, *(ii) twist GB*: $\mathbf{o} \parallel \mathbf{n}$, *(iii) mixed GB*: $\mathbf{o}$ is neither parallel nor perpendicular to $\mathbf{n}$, *(iv) symmetric GB*: the boundary plane represents the plane of *mirror symmetry*, and *(v) asymmetrical GB*: a GB that is not symmetrical. The different types of GBs are shown in Figure 5.3. In this work, we restrict our attention to *symmetric tilt GBs*, so the grain boundary plane will form the plan of symmetry and the grain boundary normal will be perpendicular to the axis of rotation between the two grains.

5.1.2 Coincidence-site lattice model

Proposed in 1949 by Kronberg and Wilson [57], the assumption of this model is that the grain boundary energy is low when the coincidence of atomic positions in both adjoining grains is high because the number of bonds broken across the boundary is small. If we construct a bicrystal where the grains are misoriented by an angle θ around an axis $\mathbf{o}$ by superimposing them such that some atomic sites coincide, then such sites are called *coincident sites*. If these coincident sites are spread uniformly throughout the entire superposition such that they form a crystal lattice of their own, then such an arrangement is termed *Coincident-Site Lattice* (CSL). An important quantity for such sites is the density (or rather its reciprocal) of the coincident sites, denoted by Σ.

$$\Sigma = \frac{\text{number of coincident sites in an elementary cell}}{\text{total number of all lattice sites in an elementary cell}} \quad \text{or} \tag{5.2}$$

$$\Sigma = \frac{\text{volume of an elementary cell of CSL}}{\text{volume of an elementary cell of the crystal lattice}}. \tag{5.3}$$

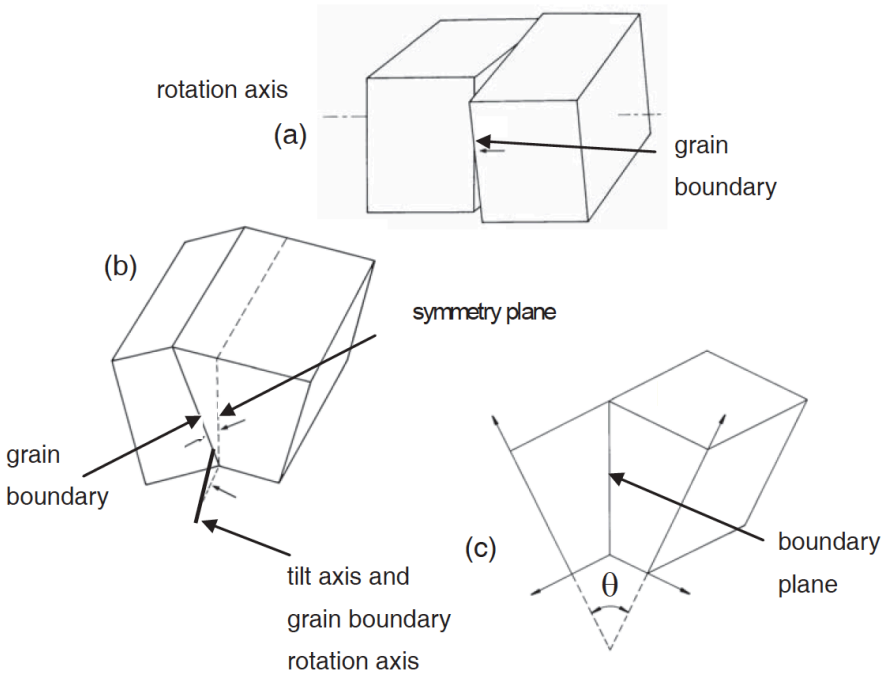

Figure 5.3: Schematic depictions of a) a twist grain boundary, b) an asymmetric tilt grain boundary, c) a symmetric tilt grain boundary [83].

If the purely geometric character of a CSL is considered, then atoms from the two grains need to be matching exactly for the grain boundary to have a coincident form as any small change of grain misorientation from the grain boundary, results in big change in coincidence. Thus, in practice, a *range* of existence of Σ CSL can be considered with small deviations from true coincidence. An example of a construction of a CSL is shown in Figure 5.4. A detailed implementation of the construction of the of a CSL grain boundary is provided in the next section.

5.2 Grain Boundary Construction for QC

As mentioned before, we restrict our analysis to symmetric tilt grain boundaries. One of the main challenges in atomistics is to construct a physically intuitive initial setup of atoms to perform simulations. For single-crystal simulations, we set our initial geometry such that the atoms lie on lattice sites corresponding to linear combinations of Bravias vectors for a given crystallographic lattice. However, while modeling grain boundaries, we model multiple grains, which necessitates keeping track of the *orientation* (the rotation of the grain with respect to the global coordi-

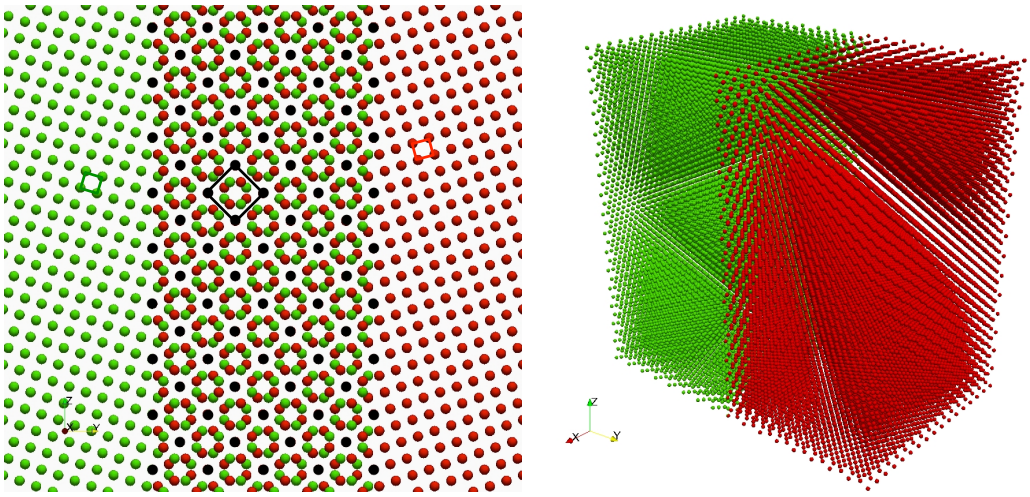

Figure 5.4: An example of a Σ3 CSL and a Σ3[111] grain boundary. The *xz*-plane is the grain boundary plane and the boundary normal is along the *y*-axis. The left pane shows a 2*D* schematic of the construction of the CSL in the *yz*-plane. The left and right lattices have different orientations and the black sites are the coincident sites. The value of Σ comes from the ratio of the volumes of the elementary cells of the CSL (shown as a black square) and the actual lattice (shown in red or green squares). The right pane shows the bicrystal formed by choosing *one* of the multiple planes of coincidence and the colors of the atoms indicate different orientations.

nate system) of every grain as its initial geometry is set up. In addition, in order to take advantage of coarse-graining, we keep the region around the grain boundary fully resolved and chose a coarse set of repatoms in the interior of each grain. While the interior of every grain might have a well-defined set of atoms located at sites represented by Bravais lattice vectors multiplied by a suitable rotation, the atoms along the grain boundary are not part of any specific lattice. One can, in principle, model atoms using suitable Bravias sites in every grain and simply remove atoms that are placed too close to each other on the grain boundary [77]. However, the meshes that are generated by such arbitrary selections of orientation and atom placements result in poorly shaped elements, which results in erroneous results. For this purpose, all bicrystals in our work are modeled using a careful construction of the simulation box as described in the following section.

5.2.1 Bicrystal using CSL

We set the following requirements to accurately model a CSL grain boundary: *(i)* a bicrystal with the grain boundary along the plane $y = 0$, *(ii)* the *y*-axis as the grain boundary normal, *(iii)* the grain boundary plane only consists of sites constructed by the CSL formed by a suitable misorientation between the two grains forming

the bicrystal, *(iv)* the grain boundary is fully resolved by choosing all lattice sites around it as repatoms and coarse-graining in both crystals away from the boundary plane, and *(v)* the repatoms that are chosen, form a well constructed mesh that does not produce poorly shaped elements.

Banadaki and Patala [11] have developed an efficient tool to construct simulation boxes for a CSL bicrystal to run an atomistic grain boundary simulation. The tool GBpy [82] helps calculate rotation matrices that construct the grain boundary simulation box with requirements mentioned above. Table B.1 shows the various grain boundaries that we use for our study, with the relevant rotation matrices. It should be noted that this tool generates boxes for atomistic simulations, so suitable modifications have to be made to construct a coarse-grained grain boundary box. The relevant steps required for this process are provided in Algorithm 3 and a brief description is given in Appendix B. The algorithm allows for generation of arbitrarily sized coarse-grained grain boundaries for *any* CSL bicrystal.

5.3 Relaxed Grain Boundary Energy at $0K$

One of the most commonly studied properties of grain boundaries is the grain boundary energy. Simply put, it is the excess free energy associated with the presence of a grain boundary with the reference being a perfect single-crystal lattice. Read and Shockley [92] provide a theoretical model that studies grain boundary energies for low-angle grain boundaries. In addition, there have been several studies performed using theory, experiments, and atomistic simulations (e.g. [26, 46, 78] among others) to observe the effect of various geometric parameters on properties like the grain boundary energy and mobility, which in turn determine grain boundary mechanics (involving GB-defect interactions, GB mibration and grain growth, e.g. during recrystallization or recovery). The GB energy is also of importance in determining a material's fracture properties, as it affects the tendency to form inter- vs intragranular cracks. This makes it important to be able to model the relaxed energy state of a grain boundary accurately. Here, we use the fully non-local QC method to determine the relaxed energy state of CSL bicrystals in order to validate the method's ability to accurately measure relaxed energy. We compare our results to the database provided by Olmsted, Foiles, and Holm [78] to quantify the errors occurring due to coarse-graining of the grain boundary.

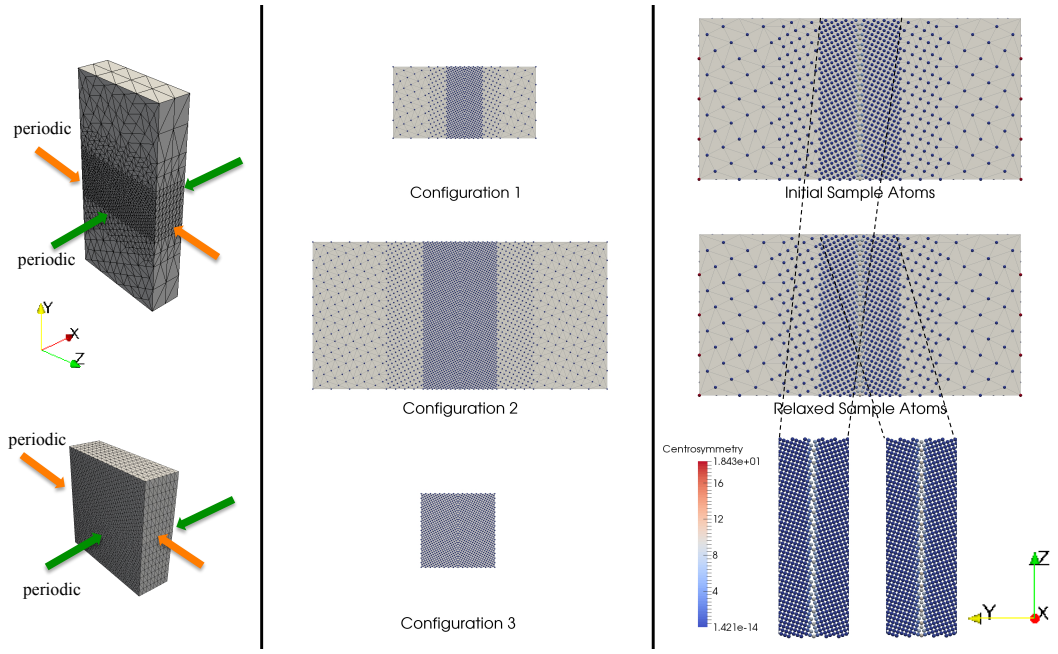

Figure 5.5: The setup of the a sample CSL bicrystal modeled with fcc Au using the Mishin potential [137] for calculating the relaxed grain boundary energy. The left pane shows meshes of coarse and fully atomistic configurations that are used for comparison along with the boundary conditions. The center pane shows a 2D perspective of the various configurations studied (sizes are to scale for comparison). Configurations 1 and 2 are coarse-grained and modeled using 12849 and 49361 repatoms respectively. Configuration 3 contains 21929 repatoms. The right pane shows the initial and relaxed state of the system for the $\Sigma 5(0\overline{1}2)$ boundary, with a zoomed view at the bottom. The atoms are colored by the centrosymmetry parameter [16].

5.3.1 Scenario setup

The scenario is set up as follows. We model 14 CSL bicrystals shown in B.1 for fcc Al using the EAM-based Mishin potential [137] at $0K$. For each of these systems, we use Algorithm 3 to generate the grain boundary box. To test the effects of coarse-graining, we model three configurations: a fully resolved configuration that acts as the exact solution, and two coarse-grained bicrystals with varied levels of coarsening. We model an *infinite grain boundary plane* by setting periodic boundary conditions along the *yz*- and the *xy*-planes. This system is allowed to relax by solving for the equilibrium configuration using quasistatics. Once equilibrated using the FIRE scheme [14], the relaxed grain boundary energy is calculated by evaluating the excess energy of the sampling atoms around the grain boundary and comparing it to the energy of a perfect single-crystal.

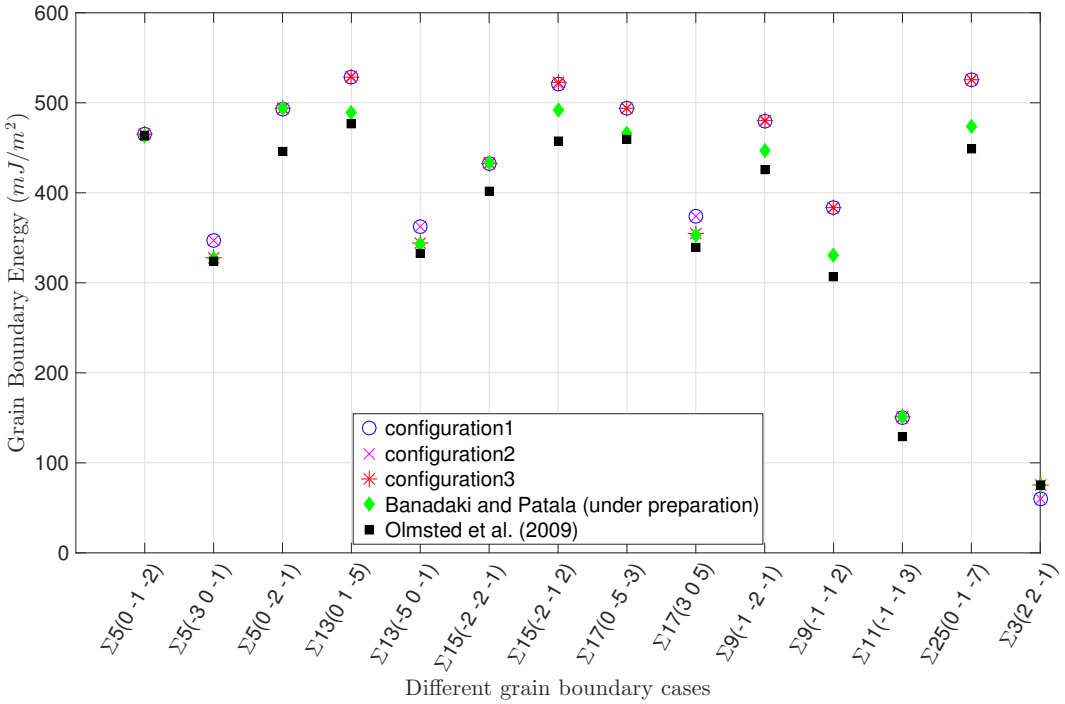

Figure 5.6: Results of the relaxed grain boundary energy for various Σ boundaries in a fcc Al bicrystal modeled with the Mishin potential [137]. The configurations are as shown in Figure 5.5. The results from Olmsted, Foiles, and Holm [78] were obtained using the potential by Ercolessi and Adams [35]. The results by Banadaki and Patala were obtained by using the Mishin potential [137].

Ten repeats are run for every configuration and results are averaged to account for various configurations possible for the equilibrium states [78]. The schematics of the scenario are shown in Figure 5.5.

5.3.2 Results

Figure 5.6 shows the results of the relaxed energy values for the different grain boundaries studied for this scenario. We observe that for most grain boundaries, all three configurations yield an identical result, so the coarse configurations have the same relaxed energy as the fully resolved one. This indicates that coarse graining does not introduce errors in the simulation. In comparison with [78], the fully non-local QC formulation overpredicts the relaxed energy value for most grain boundaries. However, out of the fourteen boundaries studied, the error introduced is less than 10% for ten of them. The high errors occur the high-Σ boundaries, e.g. $\Sigma 25(0\bar{1}7)$, with the exception of $\Sigma 9(\bar{1}\bar{1}2)$. In case of the $\Sigma 3(22\bar{1})$ boundary, the QC formulation predicts a lower value for the relaxed energy.

Some of the reasons for these discrepancies in energy values are anticipated as fol-

lows. Olmsted, Foiles, and Holm [78] used the EAM potential by Ercolessi and Adams [35] for their calculations, which result in different values for the value of the energy of a perfect single-crystal. Banadaki and Patala have rerun the simulations from [78] using MD and the EAM Mishin potential [137] and their results are also included in Figure 5.6. We see that the errors between MD and QC drop considerably when the same potential is used, which addresses one of the possible sources of discrepancies. Another difference may stem from the *offset vectors* used in [78] as a part of their initial configuration, where they laterally translate one grain with respect to other using predetermined vectors. One of the reasons for this is to try and capture the multiplicity in the equilibrium configurations [124]. In addition, [78] uses a variable cutoff distance to remove atoms near the grain boundary that are too close to each other in the initial configuration, while our QC formulation uses a fixed distance of $2\mathring{A}$. This indicates that even while modeling a fully resolved system using QC, the initial set of atoms near the grain boundary is not the same. Finally, performing an exhaustive search of a global energy minimum for a grain boundary cannot be guaranteed [78, 124], and since a reasonable match of energy was obtained after ten repeats for most cases, that was set as a cap for further simulations.

Thus we see that the hot-QC formulation acts as an effective method to reproduce GB energies for various systems. It allows for simple and computationally inexpensive studies while saving on the degrees of freedom of the system and allowing for real quasistatic relaxations that do not require resolving the thermal vibrations of every single atom in the system.

5.4 Relaxed Grain Boundaries at Finite Temperature

As mentioned earlier, the GB energy study in Section 5.3 is modeled at $0K$. This has largely been the case for most of the computational studies performed in literature, e.g. [12, 78, 117]. The aim of these studies is to get an insight into the influence of GB-energy on material properties and to characterize the energy as a continuous function of the macroscopic degrees of freedom that define the boundary. However, most real life applications and experiments that observe grain boundaries are performed at finite temperature, e.g. High Resolution Transmission Electron Microscopy (HRTEM) observations on Σ-boundaries are conducted at room temperature [6, 32, 72]. It is therefore imperative that atomistic simulations used to observe the effect of GB energy be performed at finite temperature.

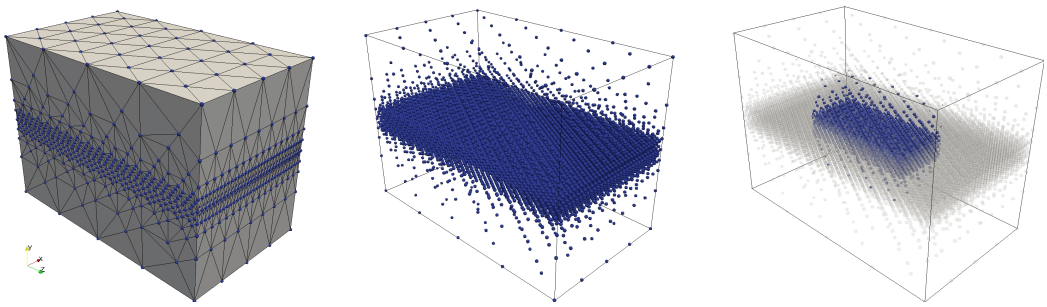

Figure 5.7: The setup of a $\Sigma 3(22\bar{1})$ CSL-based bicrystal of fcc Au modeled with the Mishin potential [137]. The left pane shows the initial mesh constructed for the bicrystal having dimensions of $9.2nm \times 11.2nm \times 15.8nm$. The center pane shows the outline of the simulation box to be modeled along containing $17,521$ repatoms. The right pane shows the subset of the atoms that are selected for energy calculations, selected using dimensions of $4.8nm \times 5.6nm \times 7.9nm$.

To this end, we use our hotQC formulation from Chapter 4 to perform finite-temperature simulations and re-evaluate GB energy for a few of the boundaries modeled in Section 5.3.

5.4.1 Scenario setup

The scenario is set up as follows. We model $\Sigma 5(0\bar{1}\bar{2})$, $\Sigma 5(0\bar{2}\bar{1})$ and $\Sigma 3(22\bar{1})$ boundaries for fcc Al using the Mishin potential [137] as before, at temperature values of $0K$, $50K$, $100K$, $150K$, $200K$, $250K$ and $300K$, using Algorithm 3. Third-degree gaussian quadrature is used from Section 4.3.1 as it is sufficient to capture thermal expansion accurately for the temperature ranges chosen (see Fig. 4.3). Isothermal conditions are maintained throughout the system for each individual temperature case. To allow for free expansion at finite temperature, no boundary conditions are applied to the outer surfaces of the simulation box. Optimal summation rules [5] are used that allow for accurate energy calculations at coarse free surfaces [4]. A small section in the interior of the grain boundary is used for energy calculations to negate surface effects. Equilibration and energy calculations are performed similar to the approach from Section 5.3.

5.4.2 Results

The evolution of relaxed energy of the boundaries with temperature is shown in Figure 5.8. We note that the $0K$ energy values for these boundaries are higher than the results from Figure 5.6 on account of the lack of periodicity in the system, thereby introducing a length scale into the observations. We observe that the energy of the $\Sigma 3(22\bar{1})$ remains more or less the same throughout the temperature range, which

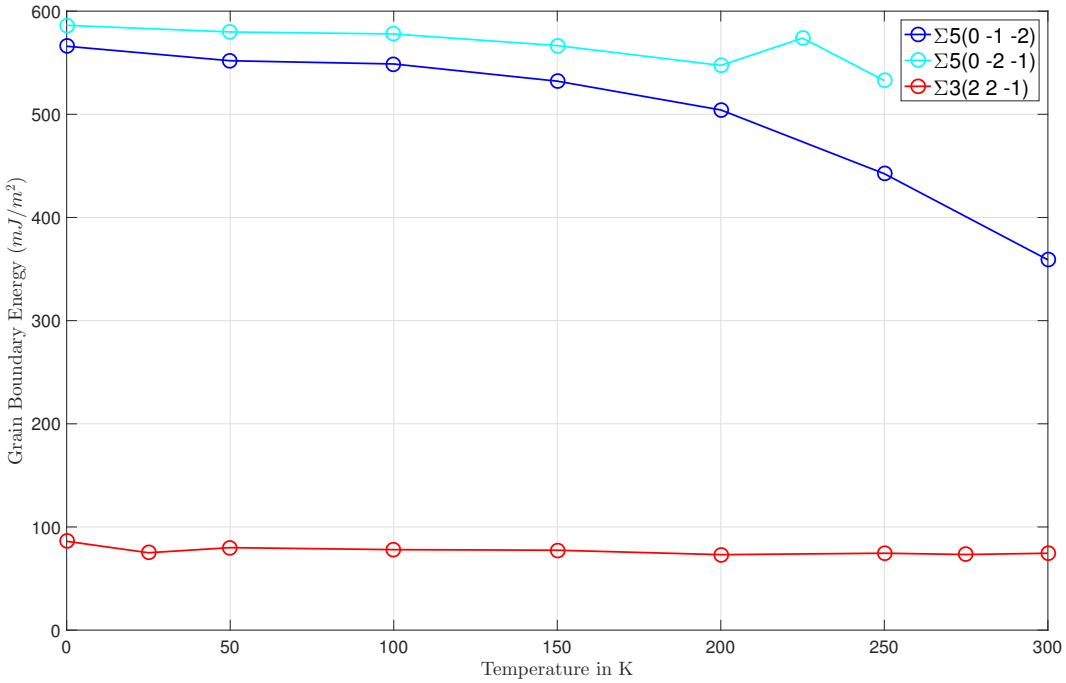

Figure 5.8: Results of the relaxed grain boundary energy for three Σ boundaries in a fcc Al bicrystal modeled with the Mishin potential [137] as a function of temperature.

can be attributed to the highly stable twin structure of the boundary. $\Sigma5(0\bar{1}\bar{2})$ and $\Sigma5(0\bar{2}\bar{1})$ show a gradual decrease in energy with temperature, with $\Sigma5(0\bar{2}\bar{1})$ showing a considerable dip near room temperature. The energy values for most of the temperature measurements lie within 20% of the predicted $0K$ values. One reason for the relatively low values of energy for some cases at $300K$ could be due to the finite nature of the grain boundary, allowing for free surface interactions. However, overall the results indicate that the data obtained from MD studies at $0K$ (e.g [78]) can be suitably used while making room temperature predictions. Figure 5.9 shows the relaxed shapes of the grain boundaries at some representative temperature values. We observe that the relaxed GB shape for the $\Sigma5(0\bar{1}\bar{2})$ and $\Sigma5(0\bar{2}\bar{1})$ boundaries is significantly different compared to the initial shape generated by Algorithm 3. As can be expected from energy calculations shown in Figure 5.8, the relaxed GB shapes do not show any appreciable geometric change over temperature. This is another indication of validity of $0K$ results at room temperature.

5.5 Conclusions

In this chapter, we extend the fully nonlocal, finite-temperature, adaptive QC framework developed throughout the coarse of this thesis to model coarse-grained 3D

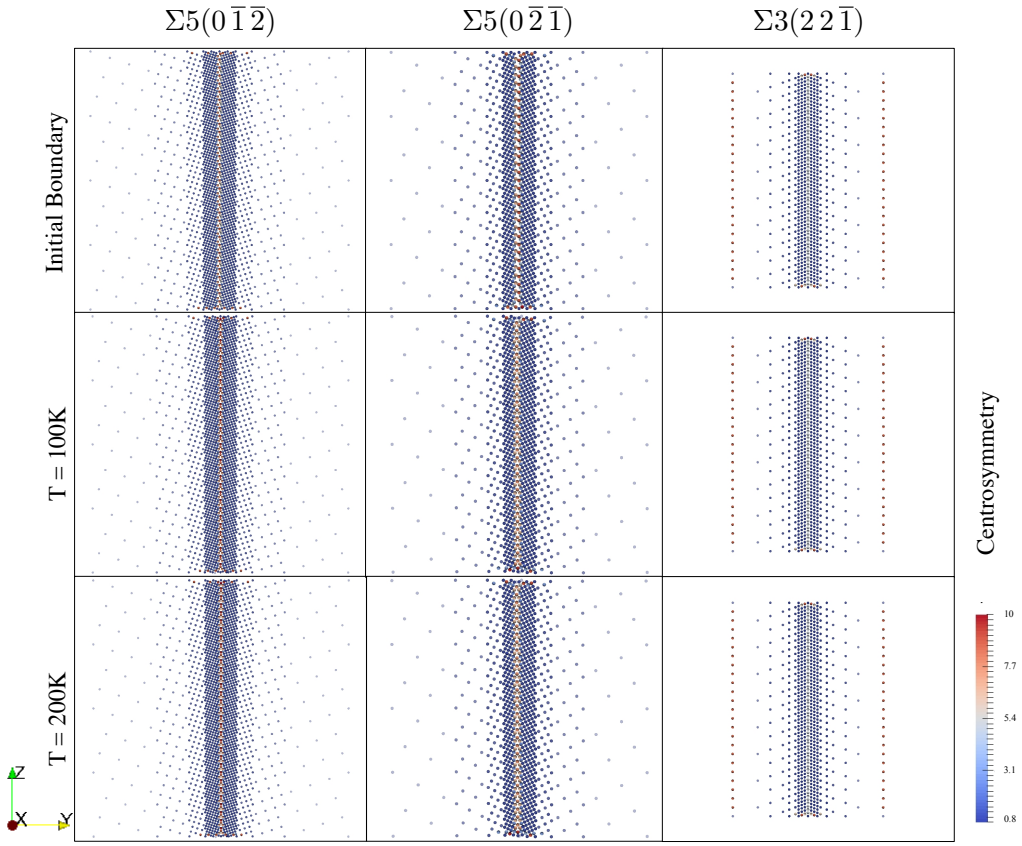

Figure 5.9: Showing the relaxed energy shapes of grain boundaries modeled in 5.4 at different temperatures. The boundaries are viewed in 2-D perspective along the *yz*-plane. The left pane shows the $\Sigma5(0\,\bar{1}\,\bar{2})$ boundary. The center pane shows the $\Sigma5(0\,\bar{2}\,\bar{1})$ boundary. The right pane shows the $\Sigma3(2\,2\,\bar{1})$ boundary.

polycrystalline systems at finite temperature. We focus on symmetric-tilt GBs, and constructed an algorithm to generate coarse-grained bicrystals using CSL based grain boundaries. Using FIRE to equilibrate the system, we predict and match the relaxed energies at $0K$ of fourteen bicrystals with MD data, providing with a simple and efficient way to calculate energies for other CSL based grain boundaries.

We further use the hotQC formulation to observe the evolution of relaxed energies of grain boundaries as a function of temperature, a novel study that has not been performed before. For the boundaries that are observed, the relaxed energy change shows a small dip as at room temperatures, providing additional validity to MD simulations. The relaxed GB shapes for these boundaries are consistent with the expected energy results and show little change from $0K$ to room temperature.

Chapter 6

CONCLUSIONS, DISCUSSION, AND OUTLOOK

Since its inception in 1996, the QC method has emerged as an effective scale-bridging technique and continues to improve with an active community of researchers adding new contributions every year [110]. Amelang [3] introduced the energy-based fully nonlocal QC method that had the advantage of seamlessly transitioning from the atomistic to the coarse-grained region (unlike most previous QC formulations) and vastly reduced spurious force artifacts by introducing the optimal summation rules [5]. This thesis has used the fully nonlocal QC formulation as a basis to advance the method in a number of directions, as outlined in the previous chapters. In this chapter, the important contributions provided by this thesis towards the advancement of the fully nonlocal QC method towards modeling of large-scale systems using automatic adaptivity at finite temperature are summarized. In addition, the current limitations of the method are also highlighted and future directions to be pursued are discussed.

6.1 Automatic Adaptivity

6.1.1 Highlights

Automatic model adaptivity has been implemented on two separate fronts, *(i)* neighborhood updates and *(ii)* mesh refinement. The novel algorithm for automatically updating the neighborhoods of sampling atoms enables simulations with large plastic deformation while guaranteeing accurate energy calculations. This algorithm has been successfully deployed in both 2D and 3D simulations and allows for coarse-graining of large-scale atomistic ensembles with low errors compared to MD, as shown from nanoindentation and nanovoid growth examples. Mesh refinement using an element-based constrained bisection algorithm has been implemented and fully captures the microstructure in simulations with severe deformation. Going beyond most prior QC techniques, this formulation is fully nonlocal and seamless. The reference mesh exists even in the fully refined domain and therefore enables seamless coarse-graining where needed.

6.1.2 Limitations

The limitations of the refinement algorithm due to the underlying Lagrangian formulation have been discussed, including the inability to undergo mesh coarsening once a defect has passed through the atomistic domain. Another limitation is that mesh refinement takes place on the *global* mesh consisting of all repatoms, and prevents massive parallelization of simulations as global communication has to be frequently conducted if repatoms are stored on distributed memory in a computer cluster. Currently, an urgent need for a 3D *local* mesh refinement algorithm (performing remeshing only where necessary without disturbing the remaining mesh) exists and we have attempted to solve it by implementing a *Contrained Advancing Front* (CAF) algorithm, described in Appendix C. However, the current implementation results in meshes of poor quality, and further work is needed before it can be suitably deployed in the QC formulation. A sample simulation is shown in Figure 6.1.

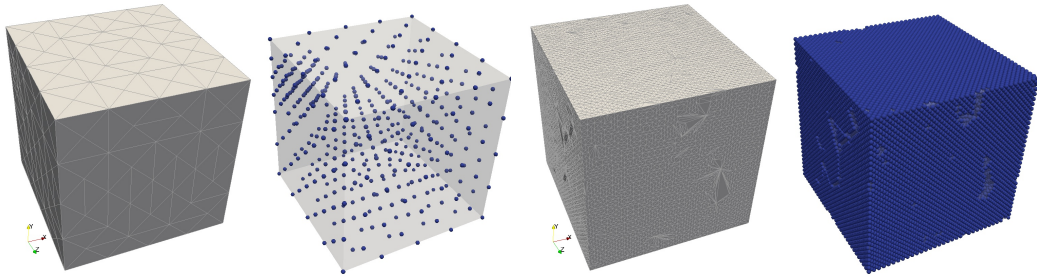

Figure 6.1: A simulation box containing $137,313$ underlying lattice points corresponding to the fcc lattice. At every refinement step, 0.1% of the total elements are chosen at random and refined to full atomistic resolution. The left pane shows the initial state of a mesh to be refined using the CAF algorithm. The initial mesh consists of 365 vertices and 1536 elements. The right pane shows the mesh after 95% of the total volume has been refined, which then contains $92,236$ vertices and $511,323$ elements.

6.2 Finite-temperature extension

6.2.1 Highlights

One of the highlights of this thesis is the extension of the fully nonlocal QC method to finite temperature. Temperature plays a key role in determining material behavior and a hotQC framework has been implemented by using the maximum entropy approach and averaging the thermal vibrations of atoms and evaluating mean positions and momenta using numerical quadrature. While previous max-ent hotQC formulations have restricted their analysis to numerical quadrature accurate up to

third-degree polynomials, we extend ours to fifth-degree polynomials. The thermal response of fcc metals to uniform thermal expansion has been demonstrated using a variety of EAM-based interatomic potentials that have not been previously reported, and our formulation is validated by comparisons to experimental behavior and LAMMPS based MD simulations. This data acts as a reference for a choice of interatomic potential for future finite-temperature applications and shows that the fifth-degree quadrature is essential for certain potentials to accurately simulate the system at higher temperatures.

The material response to temperature is also observed, viz. thermal softening. Using the optimal summation rules, it is shown that even with low repatom density (5%), accurate predictions of the bulk modulus, shear modulus, and uniaxial modulus of fcc copper as functions of temperature can be obtained, and validate the results with experimental data.

Finally, an important unknown in the QC method has been addressed, the temperature dependence of spurious force artifacts. Using a representative temperature of $100K$, the spurious artifacts have been evaluated along sharp and diffuse mesh interfaces, for both physical and thermal forces. The optimal summation rules show the lowest magnitudes of spurious forces (both physical and thermal) when compared to existing summation rules in other QC formulations. This implies that those summation rules are indeed optimal, even in a finite-temperature setting, which is of paramount importance for future simulations.

6.2.2 Limitations

The presented hotQC formulation has been deliberately limited to isothermal conditions in order to investigate and validate this important scenario. As has been done within the hotQC context, this model can be extended to account for mass and heat transport [120] as well as various other directions. However, this thesis seeks to prove that those extensions can rely on the fully nonlocal QC method. While the temperature is localized so that every repatom is assigned its own independent temperature, the temperature evolution is not solved for, and therefore the model is currently unable to analyze the influence of non-equilibrium heat flux occurring, e.g., during severe deformation or high-rate loading. This thesis lays the necessary groundwork to characterize the evolution of the discrete temperature field as has been done in [59], [87], and [120], and these avenues of research may be pursued in the future, leveraging its accomplishments.

6.3 Grain Boundaries

6.3.1 Highlights

The 3D fully nonlocal QC framework has been extended to general polycrystals. Most of the studies performed with the QC method since 1996 have been with single-crystal materials, with the polycrystalline studies limited to local QC (see e.g. [133, 134]) or quasi-2D grains (for e.g. [84, 102]). The 3D grains in this thesis have been constructed with seamless transition between the atomistic and coarse regions (as expected from a fully nonlocal method) and allow for coarse-graining within each grain.

Symmetry and coincidence sites have been used to solve meshing issues and construct stable coarse-grained CSL based bicrystals. While results of only fourteen such bicrystals have been shown, a foundation has been laid that allows construction of arbitrary boundaries based on appropriate GB features (required as inputs in Algorithm 3).

The optimal summation rules [5] have been implemented, allowing for the modeling of grain boundaries using free surfaces with significantly lower surface effects as compared to other QC formulations [4].

The hotQC formulation in this thesis allows for the evaluation of the grain boundary energy as well as the observation of the evolution of relaxed shapes of grain boundaries at finite temperature, which has not been reported by any other QC formulation in the literature.

6.3.2 Limitations

While the current hotQC formulation matches MD results, the computational overhead associated with the maintenance and updates of neighborhood verlet lists can act as a barrier in choosing QC over MD. QC excels when coarse-graining of the domain makes grain sizes go beyond the scope of MD, such as when calculating GB energies that require full resolution only near the boundary and do not require adaptive remeshing to provide atomistic resolution far away from the boundary. More is required (as discussed below) to extend this method to full adaptivity in 3D so that defect-GB interactions can be calculated effectively. We note that the finite-temperature extension presented here improves over MD simulations, as it is truly quasistatic without the need to resolve individual atomic thermal fluctuations.

6.4 Outlook

6.4.1 Dislocation-grain boundary interactions

Upon the successful implementation of the local 3D remeshing algorithm, one can put it to test by attempting to capture dislocation motion through grain boundaries. Our QC formulation is already equipped with distributed-memory parallel algorithms, allowing us to model nanoindentation in the presence of large-scale grain boundaries. Preliminary results are shown in Figure 6.2 and show dislocation loops interacting with the grain boundary. Local remeshing will significantly improve the

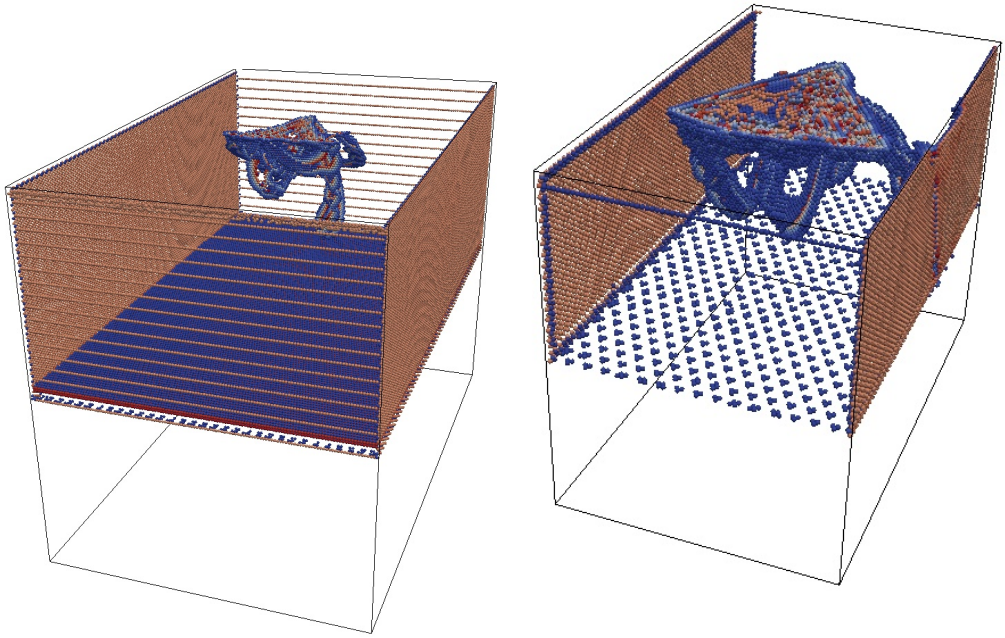

Figure 6.2: Nanoindentation using a pyramidal indenter of CSL based bicrystals of fcc Al modeled with the Mishin potential [137]. The left pane shows the formation of dislocation loops for a $\Sigma 11(\overline{1}\,\overline{1}3)$ boundary using a indenter force constant of $200 ev/\mathring{A}$. The simulation contains $1,700,175$ repatoms. The right pane shows the formation of dislocation loops for a $\Sigma 3(22\overline{1})$ boundary using a indenter force constant of $500 ev/\mathring{A}$. The simulation contains $493,633$ repatoms.

ability to increase grain sizes from the current limit of hundreds of nanometers by a few orders of magnitude (thereby approaching experimentally observed values).

6.4.2 Multi lattices using QC

So far we have limited our study to fcc metals, whereas the computational implementation is sufficiently general to model any crystalline system that can be described by a simple Bravais lattice. Morever, the nonlocal QC framework does not

come with a limitation to simple Bravais lattices and only requires that the initial set of atoms be defined on discrete lattice sites. Therefore, in principle, an extension to more complex Bravais lattices (including multi-lattices) such as the hcp lattice or even systems consisting of different materials is possible in our current framework.

An elementary implementation of a hcp-based QC formulation has been completed but needs significant testing and additions before it can be successfully deployed. A multi-lattice nonlocal QC formulation will be useful for modeling organic crystals, ferroelectrics, and phase transformations along with deformation mechanisms such as twinning, etc.

6.5 Final Comments

Overall, this thesis has contributed to the advancement of the 3D nonlocal QC method from a rather theoretical tool to a practically useful computational framework. This transformation has converted a method limited to single-crystals and zero temperature into a high-performance framework with the ability to model polycrystals at finite temperature. Multiple advances have been made towards automatic adaptivity, allowing QC to operate at length scales that were previously unattainable. There remain significant questions to be answered and improvements to be made, but the present work has improved the state-of-the-art in multiscale modeling and laid the foundation for various areas to be advanced by the scientific community in the future.

Appendix A

FINITE TEMPERATURE QUASICONTINUUM MISCELLANEOUS

A.1 Full Periodic Crystal Formulation

To validate the finite temperature QC model, we subject an infinite crystal to uniform thermal expansion, e.g. if we uniformly expand a fcc lattice, we expect the equilibrium configuration to depend on two scalar unknowns, the lattice parameter a and ω. This greatly simplifies the problem, and as every atom in the system is identical, we solve for the equilibrium configuration of exactly one atom. In the following, we formulate the simplified governing equations to be solved. The general minimization problem to be solved is (4.27), where we recall that the Helmholtz free energy is given as

$$F(\overline{\boldsymbol{q}},\omega,\theta)=\sum_{a=1}^{N}\left[3k_B\theta_a\log\frac{\hbar\omega_a}{k_b m_a\theta_a}+\langle V_a(\boldsymbol{q})\rangle\right].$$

For an infinite crystal at uniform temperature θ, each term in the summation is equal, so removing the index a gives

$$F(a,\omega)=3Nk_B\theta\log\frac{\hbar\omega}{k_b m\theta}+N\langle V(\boldsymbol{q})\rangle,$$

which motivates the term *energy per atom* given by

$$\tilde{F}(a,\omega)=\frac{F(a,\omega)}{N}=3k_B\theta\log\frac{\hbar\omega}{k_b m\theta}+\langle V(\boldsymbol{q}(a))\rangle.$$

$\langle V(\boldsymbol{q})\rangle$ is the phase average of the potential, and the dependence of the position $\boldsymbol{q}$ of the atom on the lattice parameter a comes from the fact that for a fcc lattice, the position of every atom can be obtained as the linear combination of the Bdravias lattice vectors which depend on the lattice parameter. From here, the equilibrium configuration can be obtained by solving the two scalar equations

$$\frac{\partial\tilde{F}}{\partial a}=0,\qquad\frac{\partial\tilde{F}}{\partial\omega}=0.$$

A.1.1 Minimizing with respect to lattice parameter

We use numerical Gaussian quadrature to calculate the phase average (Section 4.3). For the sake of clarity, we denote the quadrature sample p to be a dependent variable, as the atomic position $\boldsymbol{q}$ also depends on the quadrature sample (in addition to

a). This results in

$$\frac{\partial \tilde{F}}{\partial a} = \frac{\partial}{\partial a} \sum_{p=1}^{N_{QP}} W_p V(\boldsymbol{q}(a,p),\omega).$$

Using (2.11) for the potential and introducing the dependence of the lattice parameter and the quadrature sample, we write, for any site i,

$$\begin{aligned} V_i(\boldsymbol{q}(a,p)) = \sum_{k\in\mathscr{S}_i} \Phi(r_{ik}(a,p)) + \mathscr{F}(\rho_i) \qquad \text{with } \rho_i = \sum_{k\neq i} f(r_{ik}(a,p)), \\ \text{and } r_{ik} = |\boldsymbol{q}_i(a,p) - \boldsymbol{q}_k(a,p)|. \end{aligned} \tag{A.1}$$

Substituting this form of the potential in the equilibrium equation gives

$$\frac{\partial \tilde{F}}{\partial a} = \sum_{p=1}^{N_{QP}} W_p \left[\sum_{k\in\mathscr{S}_i} \left(\Phi^{'}(r_{ik}(a,p)) + \mathscr{F}^{'}(\rho_i) f^{'}(r_{jk}(a,p)) \right) \frac{\partial r_{ik}(a,p)}{\partial a} \right], \tag{A.2}$$

where

$$\frac{\partial r_{ik}(a,p,\omega)}{\partial a} = \frac{\partial}{\partial a} |\boldsymbol{q}_i(a,p) - \boldsymbol{q}_k(a,p)|.$$

From (4.29),

$$\boldsymbol{q}_i(a,p) = \overline{\boldsymbol{q}_i}(a) + \frac{\sqrt{2}\sigma}{\omega}\boldsymbol{x}_i(p), \qquad \boldsymbol{q}_k(a,p) = \overline{\boldsymbol{q}_k}(a) + \frac{\sqrt{2}\sigma}{\omega}\boldsymbol{x}_k(p),$$

which gives

$$\frac{\partial r_{ik}(a,p,\omega)}{\partial a} = \frac{1}{r_{ik}} \left((\boldsymbol{q}_i - \boldsymbol{q}_k) \cdot ((\overline{\boldsymbol{q}}_i^{'} - \overline{\boldsymbol{q}}_k^{'})) \right). \tag{A.3}$$

The $\overline{\boldsymbol{q}}_i^{'}$ stands for the partial derivative of the position with respect to the lattice parameter. Here, the assumption is that the position of any atom can be described as the linear combination of the Bdravais lattice vectors which depend on the lattice parameter. Substituting (A.3) back into (A.2), we obtain the final minimized expression with respect to a.

A.1.2 Minimization with respect to ω

Proceeding like the case of the lattice parameter a using (A.1) and noting that the dependence on ω in the potential term comes from the phase averaging, we write[1]

$$\frac{\partial \tilde{F}}{\partial \omega} = \frac{3k_B\theta}{\omega} + \sum_{p=1}^{N_{QP}} W_p \left[\sum_{k\in\mathscr{S}_i} \left(\Phi^{'}(r_{ik}(p,\omega)) + \mathscr{F}^{'}(\rho_i) f^{'}(r_{jk}(p,\omega)) \right) \frac{\partial r_{ik}(p,\omega)}{\partial \omega} \right], \tag{A.4}$$

[1] The positions still depend on the lattice parameter, but it is not explicitly shown because here we want to minimize with respect to ω.

where

$$\frac{\partial r_{ik}(p,\omega)}{\partial \omega} = \frac{1}{2r_{ik}}\frac{\partial}{\partial \omega}(\boldsymbol{q}_i(p,\omega)-\boldsymbol{q}_k(p,\omega))\cdot(\boldsymbol{q}_i(p,\omega)-\boldsymbol{q}_k(p,\omega)).$$

Once again from (4.29), we write,

$$\boldsymbol{q}_i(p,\omega) = \overline{\boldsymbol{q}_i} + \frac{\sqrt{2}\sigma}{\omega}\boldsymbol{x}_i(p), \qquad \boldsymbol{q}_k(p,\omega) = \overline{\boldsymbol{q}_k} + \frac{\sqrt{2}\sigma}{\omega}\boldsymbol{x}_k(p),$$

which gives

$$\frac{\partial \boldsymbol{q}_i(p,\omega)}{\partial \omega} = -\frac{\sqrt{2}\sigma}{\omega^2}\boldsymbol{x}_i(p), \qquad \frac{\partial \boldsymbol{q}_k(p,\omega)}{\partial \omega} = -\frac{\sqrt{2}\sigma}{\omega^2}\boldsymbol{x}_k(p).$$

This gives the required partial derivative as

$$\frac{\partial r_{ik}(a,p,\omega)}{\partial \omega} = \frac{-1}{r_{ik}}\left((\boldsymbol{q}_i-\boldsymbol{q}_k)\cdot\left(\frac{\sqrt{2}\sigma}{\omega^2}\boldsymbol{x}_i - \frac{\sqrt{2}\sigma}{\omega^2}\boldsymbol{x}_k\right)\right). \tag{A.5}$$

As in the case of the lattice parameter, substituting (A.5) into (A.4) yields the required minimized expression with respect to ω.

A.2 Units for parameters

Chapter 4 derives the governing equations to find equilibrium configurations in finite-temperature atomistics. In the numerical implementation of these equations, however, one must be careful to make sure that proper units of various quantities are used and we do not incur numerical errors. One example of such errors is when dividing two quantities which are orders of magnitude different. This section deals with the definition of various quantities used for the derivation of (4.33) and the different units that we use for our numerical implementation. Specifically, the goal of the section is to show that these units ensure that all quantities used in the hot-QC formulation make physical sense in a wide range of temperatures $(0-1000K)$. For the purpose of illustration, we use fcc copper as our material and the Extended Finnis Sinclair potential [29].

A.2.1 Basic units

- m_a - mass of an atom in atomic mass units. For copper, we have

$$\begin{aligned} m &= 63.546 a.m.u. \\ &= 63.546 \times 1.6726 \times 10^{-27} kg, \\ &= 63.546 \times 103.642641 = 6586.0572 \frac{eV fs^2}{\mathring{A}^2}. \end{aligned}$$

- θ - temperature of an atom in Kelvin.
- q - position of an atom is in $\mathring{A}$.
- k_B - Boltzmann constant It can be expressed in different units as follows:

$$\begin{aligned} k_B &= 1.380648 \times 10^{-23} \frac{J}{K}, \\ &= 8.6173303 \times 10^{-5} \frac{eV}{K}. \end{aligned}$$

For our applications, we will use the expression in electron volts.

A.2.2 Derived units

- **(Derived unit)** σ - The standard deviation of the momentum of an atom. From (4.24), we obtain the required units for σ. For a room temperature application, the value of σ is roughly:

$$\begin{aligned} \sigma &= \sqrt{k_B \theta m}, \\ &= \sqrt{8.6173303 \times 10^{-5} \frac{eV}{K} \times 300K \times 6586.0572 \frac{eV fs^2}{\mathring{A}^2}}, \\ &= 13.0485 \frac{eV fs}{\mathring{A}}. \end{aligned}$$

- ω - The ratio of the standard deviations of the momentum to the position. Using the obtained units for σ, we use (4.29) to obtain the units of ω_a:

$$\underbrace{\boldsymbol{q}(\boldsymbol{x}, \overline{\boldsymbol{q}})}_{\text{units of } \mathring{A}} = \frac{\sqrt{2}\sigma}{\omega} \underbrace{\boldsymbol{x}}_{\text{unitless}} + \underbrace{\overline{\boldsymbol{q}}}_{\text{units of } \mathring{A}.}$$

This means that the factor in front of the shift $\boldsymbol{x}$ has the unit of $\mathring{A}$. Further simplification gives

$$\begin{aligned} (\text{unit of})\omega &= (\text{units of }) \sigma \times \mathring{A}^{-1}, \\ &= \frac{eV fs}{\mathring{A}} \times \mathring{A}^{-1}, \\ &= \frac{eV fs}{\mathring{A}^2}. \end{aligned}$$

To obtain a more physical interpretation for this quantity, [60] and [119] relate it to the thermal vibrations of atoms. This motivates the need to express ω in

units of frequency. If we design a new quantity $\tilde{\omega}$ having units of frequency, we can relate it to ω as follows:

$$\begin{aligned}\tilde{\omega} &= \frac{\omega}{m},\\ (\text{units of})\tilde{\omega} &= \frac{\frac{eVfs}{\mathring{A}^2}}{\frac{eVfs^2}{\mathring{A}^2}}\\ &= \frac{1}{fs}.\end{aligned}$$

As we solve for ω in (4.33), it helps to have some intuition of what a typical value of ω is expected, specifically for solver convergence. In general, the frequencies of atomic vibrations are on the order of terahertz. Assuming a value of around $30tHz = 0.03PHz$ around room temperature, for copper,

$$\omega = 6586.0572 \times 0.03 = 197.58\frac{eVfs}{\mathring{A}^2}.$$

In our uniform thermal expansion studies from Section 4.7, we keep our initial guess for ω to be around this value for fcc copper and can similarly obtain values for aluminum to precondition the numerical solver.

Appendix B

CONSTRUCTING CSL-BASED GRAIN BOUNDARIES

We present supplemental material that explains the design and implementation of grain boundaries used in Chapter 5. In particular, in Table B.1 we list the various grain boundaries used for the study along with the rotation matrices used for formulating the initial simulation box to perform calculations. In addition, Algorithm 3 outlines the steps used to generate a bicrystal simulation box for every grain boundary in Table B.1. The various notations used in the algorithm are explained below.

- $\boldsymbol{P}_{\text{fcc}}$ is the set of points constructed to lie on the sites for a fcc crystal lattice without any rotation. From a reference point (usually taken as the origin), points are constructed lying in integer multiples of vectors $a_0[0\,\frac{1}{2}\,\frac{1}{2}]$, $a_0[\frac{1}{2}\,0\,\frac{1}{2}]$ and $a_0[\frac{1}{2}\,\frac{1}{2}\,0]$, where a_0 is the lattice parameter of the fcc material being constructed at $0K$.

- $\boldsymbol{R}_{\text{misorient}}$ is the 3×3 rotation matrix that rotates one grain relative to the other and forms a CSL with the required Σ-value and the grain boundary normal. Table B.1 lists $\boldsymbol{R}_{\text{misorient}}$ for every boundary that we model. As we eventually choose the $y = 0$ plane as the grain boundary plane, we call the grain with a positive y-coordinate the "upper grain" and negative y-coordinate the "lower grain". This is the convention followed throughout Algorithm 3 and in Chapter 5. For a specific grain boundary, $\boldsymbol{R}_{\text{misorient}}$ is obtained from GBpy [82].

- $[\boldsymbol{P}_u, \boldsymbol{P}_l]$ are the sets of points obtained after the **misorient** function is applied to $\boldsymbol{P}_{\text{fcc}}$. As only the upper grain is misoriented, the points of the lower grain are the same as those already constructed in the global coordinate system, i.e. $\boldsymbol{P}_l = \boldsymbol{P}_{\text{fcc}}$. The points for the upper grain are obtained by rotating $\boldsymbol{P}_{\text{fcc}}$ with the misorientation rotation matrix, which is normalized column-wise first.

So, we first calculate

$$\begin{aligned}
\boldsymbol{R}_{\text{norm}}(:,1) &= \boldsymbol{R}_{\text{misorient}}(:,1) / \left\|\boldsymbol{R}_{\text{misorient}}(:,1)\right\|, \\
\boldsymbol{R}_{\text{norm}}(:,2) &= \boldsymbol{R}_{\text{misorient}}(:,2) / \left\|\boldsymbol{R}_{\text{misorient}}(:,2)\right\|, \\
\boldsymbol{R}_{\text{norm}}(:,3) &= \boldsymbol{R}_{\text{misorient}}(:,3) / \left\|\boldsymbol{R}_{\text{misorient}}(:,3)\right\|,
\end{aligned}$$

which then gives $\boldsymbol{P}_u = \boldsymbol{R}_{\text{norm}} \boldsymbol{P}_{\text{fcc}}$.

- $\boldsymbol{R}_{\text{basis change}}$ is the 3×3 rotation matrix that rotates points from both the grains which converts the grain boundary plane to $y = 0$, and the grain boundary normal to the y-axis. Table B.1 lists $\boldsymbol{R}_{\text{basis change}}$ for every boundary that is modeled, and is obtained from GBpy [82].

- $\left[\boldsymbol{P}_{u,\text{rot}}, \boldsymbol{P}_{l,\text{rot}}\right]$ are the sets of points obtained after the function **global_rotation** is applied to $[\boldsymbol{P}_u, \boldsymbol{P}_l]$. In the function, we first evaluate the actual transformation matrix $\boldsymbol{R}_{\text{transform}} = \boldsymbol{R}^{-1}_{\text{basis change}}$. Then the new set of points is obtained by

$$\begin{aligned}
\boldsymbol{P}_{u,\text{rot}} &= \boldsymbol{R}_{\text{transform}} \boldsymbol{P}_u, \\
\boldsymbol{P}_{l,\text{rot}} &= \boldsymbol{R}_{\text{transform}} \boldsymbol{P}_l.
\end{aligned}$$

  No further rotations are performed on the points.

- The final bicrystal is designed to be a convex box that contains two grains which are also convex boxes sharing a plan (the grain boundary plane). In order to allow a variety of sizes for every grain boundary, the global x, y, and z dimensions of this bicrystal-box are variables that need to be chosen carefully. In particular, these dimensions depend on the number of CSL cells in every dimension. n_1, n_2, and n_3 are the number of CSL cells in the x, y, and z dimension respectively.

- $\boldsymbol{R}_{\text{edge cut, l}}$ and $\boldsymbol{R}_{\text{edge cut, u}}$ are the 3×3 rotation matrices required for calculations of the final bicrystal-box dimensions. These are the effective rotations that points in the two grains have undergone from the initial state. They are obtained as

$$\begin{aligned}
\boldsymbol{R}_{\text{edge cut},l} &= \boldsymbol{R}_{\text{basis change}}, \\
\boldsymbol{R}_{\text{edge cut},u} &= \boldsymbol{R}_{\text{misorient}} \boldsymbol{R}_{\text{basis change}}.
\end{aligned}$$

- The final step in the box construction is to perform the **edge_cut** on points in both the grains to remove points that lie outside the intended final domain. We obtain x, y, and z limits for both grains, and only points belonging to the region between these limits are considered for the final bicrystal box construction. If a_0 is the lattice parameter of the material being modeled, the limits for both grains are as follows. For the lower grain:

$$
\begin{aligned}
x_{\text{lims},l} &= \left[-a_0 n_1 \left\|\boldsymbol{R}_{\text{edge cut},l}(:,1)\right\|, a_0 n_1 \left\|\boldsymbol{R}_{\text{edge cut},l}(:,1)\right\|\right], \\
y_{\text{lims},l} &= \left[-a_0 n_2 \left\|\boldsymbol{R}_{\text{edge cut},l}(:,2)\right\|, 0\right], \\
z_{\text{lims},l} &= \left[-a_0 n_3 \left\|\boldsymbol{R}_{\text{edge cut},l}(:,3)\right\|, a_0 n_3 \left\|\boldsymbol{R}_{\text{edge cut},l}(:,3)\right\|\right].
\end{aligned}
$$

  Once these limits are set, the final set of points chosen for the lower grain are obtained as the subset of $\boldsymbol{P}_{l,\text{rot}}$ that lie in between these limits. This is expressed as

$$
\boldsymbol{P}_{l,\text{final}} = \boldsymbol{P}_{l,\text{rot}} \in \left[x_{\text{lims},l}, y_{\text{lims},l}, z_{\text{lims},l}\right].
$$

  Similarly, the limits for the upper grain are

$$
\begin{aligned}
x_{\text{lims},u} &= \left[-a_0 n_1 \left\|\boldsymbol{R}_{\text{edge cut},u}(:,1)\right\|, a_0 n_1 \left\|\boldsymbol{R}_{\text{edge cut},u}(:,1)\right\|\right], \\
y_{\text{lims},u} &= \left[0, a_0 n_2 \left\|\boldsymbol{R}_{\text{edge cut},u}(:,2)\right\|\right], \\
z_{\text{lims},u} &= \left[-a_0 n_3 \left\|\boldsymbol{R}_{\text{edge cut},u}(:,3)\right\|, a_0 n_3 \left\|\boldsymbol{R}_{\text{edge cut},u}(:,3)\right\|\right],
\end{aligned}
$$

  and the final upper grain points are chosen as a subset of $\boldsymbol{P}_{u,\text{rot}}$ as

$$
\boldsymbol{P}_{u,\text{final}} = \boldsymbol{P}_{u,\text{rot}} \in \left[x_{\text{lims},u}, y_{\text{lims},u}, z_{\text{lims},u}\right].
$$

This gives the final set of points $\left[\boldsymbol{P}_{l,\text{final}}\boldsymbol{P}_{u,\text{final}}\right]$ that are used to construct a bicrystal box of $n_1 \times n_2 \times n_3$ CSL unit cells. To generate a coarse bicrystal, we start with coarse set of points $\boldsymbol{P}_{\text{fcc, coarse}}$ that are generated in multiples of the lattice parameter a_0, and the rest of the procedure follows in exactly the same way.

Case	Representation	$R_{\text{misorient}}$	$R_{\text{basis change}}$
1	$\Sigma5(0\bar{1}\bar{2})$	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & -2 \\ 0 & 3 & -1 \end{bmatrix}$	$\frac{1}{5}\begin{bmatrix} 5 & 0 & 0 \\ 0 & 4 & 3 \\ 0 & -3 & 4 \end{bmatrix}$
2	$\Sigma5(\bar{3}0\bar{1})$	$\begin{bmatrix} 0 & 1 & 0 \\ -3 & 0 & -1 \\ -1 & 0 & 3 \end{bmatrix}$	$\frac{1}{5}\begin{bmatrix} 5 & 0 & 0 \\ 0 & 4 & 3 \\ 0 & -3 & 4 \end{bmatrix}$
3	$\Sigma5(0\bar{2}\bar{1})$	$\begin{bmatrix} -1 & 0 & 0 \\ 0 & -2 & -1 \\ 0 & 1 & 2 \end{bmatrix}$	$\frac{1}{5}\begin{bmatrix} 5 & 0 & 0 \\ 0 & 4 & 3 \\ 0 & -3 & 4 \end{bmatrix}$
4	$\Sigma13(01\bar{5})$	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -5 \\ 0 & 5 & 1 \end{bmatrix}$	$\frac{1}{13}\begin{bmatrix} 13 & 0 & 0 \\ 0 & 12 & 5 \\ 0 & -5 & 12 \end{bmatrix}$
5	$\Sigma13(\bar{5}0\bar{1})$	$\begin{bmatrix} 0 & 1 & 0 \\ -5 & 0 & -1 \\ -1 & 0 & 5 \end{bmatrix}$	$\frac{1}{13}\begin{bmatrix} 13 & 0 & 0 \\ 0 & 12 & 5 \\ 0 & -5 & 12 \end{bmatrix}$
6	$\Sigma15(\bar{2}\bar{2}\bar{1})$	$\begin{bmatrix} 1 & 1 & -2 \\ -2 & -2 & -1 \\ -1 & 5 & 0 \end{bmatrix}$	$\frac{1}{15}\begin{bmatrix} 14 & 2 & -5 \\ 2 & 11 & 10 \\ 5 & -10 & 10 \end{bmatrix}$
7	$\Sigma15(\bar{2}\bar{1}2)$	$\begin{bmatrix} 1 & -2 & -1 \\ -2 & -1 & 2 \\ -1 & 0 & -5 \end{bmatrix}$	$\frac{1}{15}\begin{bmatrix} 14 & 2 & -5 \\ 2 & 11 & 10 \\ 5 & -10 & 10 \end{bmatrix}$
8	$\Sigma17(0\bar{5}\bar{3})$	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & -5 & -3 \\ 0 & 3 & -5 \end{bmatrix}$	$\frac{1}{17}\begin{bmatrix} 17 & 0 & 0 \\ 0 & 15 & 8 \\ 0 & -8 & 15 \end{bmatrix}$
9	$\Sigma17(305)$	$\begin{bmatrix} 0 & 1 & 0 \\ 3 & 0 & 5 \\ 5 & 0 & -3 \end{bmatrix}$	$\frac{1}{17}\begin{bmatrix} 17 & 0 & 0 \\ 0 & 15 & 8 \\ 0 & -8 & 15 \end{bmatrix}$
10	$\Sigma9(\bar{1}\bar{2}\bar{1})$	$\begin{bmatrix} -1 & 2 & 1 \\ -1 & -2 & -1 \\ 0 & -1 & 4 \end{bmatrix}$	$\frac{1}{9}\begin{bmatrix} 8 & 1 & -4 \\ 1 & 8 & 4 \\ 4 & -4 & 7 \end{bmatrix}$
11	$\Sigma9(\bar{1}\bar{1}2)$	$\begin{bmatrix} -1 & 1 & 2 \\ -1 & -1 & 2 \\ 0 & 4 & 1 \end{bmatrix}$	$\frac{1}{9}\begin{bmatrix} 8 & 1 & -4 \\ 1 & 8 & 4 \\ 4 & -4 & 7 \end{bmatrix}$
12	$\Sigma11(\bar{1}\bar{1}3)$	$\begin{bmatrix} -1 & 1 & -3 \\ -1 & -1 & 3 \\ 0 & 3 & 2 \end{bmatrix}$	$\frac{1}{11}\begin{bmatrix} 9 & 2 & -6 \\ 2 & 9 & 6 \\ 6 & -6 & 7 \end{bmatrix}$
13	$\Sigma25(0\bar{1}\bar{7})$	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & -7 \\ 0 & 7 & -1 \end{bmatrix}$	$\frac{1}{25}\begin{bmatrix} 25 & 0 & 0 \\ 0 & 24 & 7 \\ 0 & -7 & 25 \end{bmatrix}$
14	$\Sigma3(22\bar{1})$	$\begin{bmatrix} 2 & -1 & 2 \\ 2 & 2 & -1 \\ -1 & 2 & 2 \end{bmatrix}$	$\frac{1}{3}\begin{bmatrix} 0 & 1 & 2 \\ 1 & 1 & -1 \\ -1 & 1 & -1 \end{bmatrix}$

Table B.1: The different configurations constructed in this study. The second column shows the representation of the grain boundary in the bicrystal including the Σ-value and the grain boundary normal. $\boldsymbol{R}_{\text{misorient}}$ is the rotation needed to transform one of the grains with respect to other (in the global coordinate system) to form a CSL. $\boldsymbol{R}_{\text{basis change}}$ is the rotation that *both grains* go through so that the grain boundary plane becomes the plane $y = 0$ and the y axis becomes the boundary normal.

Algorithm 3 Box Construction for a CSL Bicrystal

function box_construction
    Choose a grain boundary to construct
    Obtain $\boldsymbol{R}_{\text{misorient}}$ and $\boldsymbol{R}_{\text{basis change}}$ from Table B.1 for this boundary
    Generate $\boldsymbol{P}_{\text{fcc}}$: fcc lattice points in global coordinate system
    $[\boldsymbol{P}_u, \boldsymbol{P}_l]$ = **misorient**$(\boldsymbol{P}_{\text{fcc}}, \boldsymbol{R}_{\text{misorient}})$
    $\left[\boldsymbol{P}_{u,\text{rot}}, \boldsymbol{P}_{l,\text{rot}}\right]$ = **global_rotation** $\left(\boldsymbol{P}_u, \boldsymbol{P}_l, \boldsymbol{R}_{\text{basis change}}\right)$
    Choose integers n_1, n_2, n_3 for the number of cells of the CSL in the x, y and z directions respectively
    $\boldsymbol{R}_{\text{edge cut},l} = \boldsymbol{R}_{\text{basis change}}$,
    $\boldsymbol{R}_{\text{edge cut},u} = \boldsymbol{R}_{\text{misorient}} \times \boldsymbol{R}_{\text{basis change}}$,
    $\left[\boldsymbol{P}_{u,\text{final}}, \boldsymbol{P}_{l,\text{final}}\right]$ = **edge_cut** $\left(\boldsymbol{P}_{u,\text{rot}}, \boldsymbol{P}_{l,\text{rot}}, n_1, n_2, n_3, \boldsymbol{R}_{\text{edge cut},l}, \boldsymbol{R}_{\text{edge cut},u}\right)$
    return
end function

Applying misorientations to grains

function misorient($\boldsymbol{P}_{\text{fcc}}$, $\boldsymbol{R}_{\text{misorient}}$)
    $\boldsymbol{P}_l = \boldsymbol{P}_{\text{fcc}}$ (leave lower grain unchanged)
    $\boldsymbol{R}_{norm} = \text{normalize}\ (\boldsymbol{R}_{\text{misorient}})$
    $\boldsymbol{P}_u = \boldsymbol{R}_{norm} \times \boldsymbol{P}_{fcc}$ (rotate the upper grain points)
    return $[\boldsymbol{P}_u, \boldsymbol{P}_l]$
end function

Rotating both crystals to change GB normal

function global_rotation($\boldsymbol{P}_u, \boldsymbol{P}_l, \boldsymbol{R}_{\text{basis change}}$)
    $\boldsymbol{R}_{\text{transform}} = \text{inv}\ (\boldsymbol{R}_{\text{basis change}})$
    $\boldsymbol{P}_{u,\text{rot}} = \boldsymbol{R}_{\text{transform}} \times \boldsymbol{P}_u$
    $\boldsymbol{P}_{l,\text{rot}} = \boldsymbol{R}_{\text{transform}} \times \boldsymbol{P}_l$
    return $\left[\boldsymbol{P}_{u,\text{rot}}, \boldsymbol{P}_{l,\text{rot}}\right]$
end function

Applying edge cuts to complete the box

function edge_cut($\boldsymbol{P}_{u,\text{rot}}, \boldsymbol{P}_{l,\text{rot}}, n_1, n_2, n_3, \boldsymbol{R}_{\text{edge cut},l}, \boldsymbol{R}_{\text{edge cut},u}$)
    $a_0 = \text{Lattice Parameter of the fcc material}$
    $x_{\text{lims},u} = \pm a_0 n_1 \times \left\| \boldsymbol{R}_{\text{edge cut},u}(:,1) \right\|$
    $y_{\text{lims},u} = +a_0 n_2 \times \left\| \boldsymbol{R}_{\text{edge cut},u}(:,2) \right\|$
    $z_{\text{lims},u} = \pm a_0 n_3 \times \left\| \boldsymbol{R}_{\text{edge cut},u}(:,3) \right\|$
    $\boldsymbol{P}_{u,\text{final}} = \boldsymbol{P}_{u,\text{rot}} \in \left[x_{\text{lims},u}, y_{\text{lims},u}, z_{\text{lims},u}\right]$
    $x_{\text{lims},l} = \pm a_0 n_1 \times \left\| \boldsymbol{R}_{\text{edge cut},l}(:,1) \right\|$
    $y_{\text{lims},l} = -a_0 n_2 \times \left\| \boldsymbol{R}_{\text{edge cut},l}(:,2) \right\|$
    $z_{\text{lims},l} = \pm a_0 n_3 \times \left\| \boldsymbol{R}_{\text{edge cut},l}(:,3) \right\|$
    $\boldsymbol{P}_{l,\text{final}} = \boldsymbol{P}_{l,\text{rot}} \in \left[x_{\text{lims},l}, y_{\text{lims},l}, z_{\text{lims},l}\right]$
    return $\left[\boldsymbol{P}_{u,\text{final}}, \boldsymbol{P}_{l,\text{final}}\right]$
end function

Appendix C

AUTOMATIC ADAPTIVITY II. CONSTRAINED ADVANCING FRONT ALGORITHM

Chapter 2 discusses the formulation of our fully nonlocal QC method. As mentioned in earlier sections, when combined with automatic mesh adaptation techniques, this method does not require a-priori knowledge about where atomistic resolution will be needed during a simulation, since mesh refinement ensures that full atomistic resolution is deployed around evolving microstructural defects. However, going beyond all previous QC realizations, the fully-nonlocal scheme has introduced a number of significant challenges specifically associated with its adaptive remeshing capabilities: (*i*) in order to seamlessly refine down to the atomistic limit, every representative atom (i.e., every vertex in the mesh) is required to coincide with an atomic site of the underlying discrete crystal lattice; we will refer to this as the *discrete location constraint* (DLC). (*ii*) Existing vertices must be retained at their exact same location during remeshing (to guarantee the survival of lattice defects during remeshing). (*iii*) Massive parallelism of the implementation calls for remeshing schemes that can be performed locally, i.e., the geometric algorithm must offer control to remesh only small (well-defined) subdomains of the total simulation domain (thus avoiding refinement propagation). (*iv*) As in any remeshing technique, refinement criteria and geometric algorithms must be chosen to ensure sufficient mesh quality. (*v*) The special features of the nonlocal QC method require a careful updating/resampling procedure to define nodal and element data after remeshing. Chapter 3 of this thesis addresses points (*iv*) and (*v*) in a general three-dimensional (3D) setting [115]. However, all applications there are limited to 2D due to the lack of a suitable 3D remeshing algorithm that satisfies all of the above constraints. To this end, we have performed some elementary advances to close this gap by introducing the *Constrained Advancing Front* (CAF) algorithm for 3D mesh refinement in the nonlocal QC method. Various remeshing techniques and a myriad of implementations exist and find frequent use in computational solid mechanics, primarily in conjunction with the finite element (FE) method. Unfortunately, many of the above challenges do not exist in the FE framework: nodes do not have to lie on a discrete grid (the DLC is removed) and thus can be placed and moved to maximize mesh quality and to ensure local refinement. Also, the

principle of local action (unless in meshless or particle methods) makes parallelization more readily available than in the QC method, where long-range inter-element interactions cause complications (see [115]). For these reasons, traditional mesh refinement strategies are difficult to apply in our scenario.

For example, we aim to avoid algorithms that propagate remeshing into the surroundings in an uncontrollable manner, as seen, e.g., with the *Longest Edge Propagation Path* (LEPP) algorithm [94, 93, 95]. LEPP recursively searches for the longest edge within elements on the propagation path and can lead to a cascade of element refinements over large distances. An efficient, massively-parallel implementation, however, must avoid such uncontrolled cascading remeshing, since propagation paths will inevitably cross rank domain boundaries [48]. Introducing a simple cutoff in the propagation path has proven problematic, since the mesh quality suffers considerably. Furthermore, LEPP does not offer a choice in the number of new nodes being inserted, which may corrupt the one major advantage of the QC method, viz. the capability to locally confine costly atomistic resolution around defects benefits and to aggressively coarsen away from those defects. This, however, produces high gradients of element size away from the fully-resolved region, which is why LEPP is better suited in the FE context where a smooth element size distribution is preferred.

A local, *greedy, hill-climbing* algorithm was implemented previously as an alternative refinement scheme for the nonlocal QC method [3]. It localizes remeshing to only a few elements per refinement and aims to compensate for the imminent loss in mesh quality due to the extreme locality by considering all possible vertex insertions and bisections of which it chooses only the best option. While the algorithm is extremely fast, locality and the enforcement of the DLC decrease the mesh quality to an unacceptable level [3].

Delaunay Refinement [105] and *Constrained Delaunay Triangulation* (CDT) [23, 31] are further viable options that are both based on Delaunay triangulations. Here, for each element that must be refined, a small cavity is constructed, leaving the outer surface of the cavity as a set of faces to be preserved by the CDT. Unfortunately, common CDT algorithms presume the ability to insert so-called Steiner points in arbitrary locations on the constraining entities [27], which violates the DLC, affects locality, and requires special tools to re-establish the conformity of the mesh.

To overcome all of the above limitations, we introduce the CAF algorithm, which follows the original *Advancing Front* (AF) strategy of [39]. The remeshing scheme

is local (thus allowing for parallelization) and complies with all of the above geometric constraints; in particular, atoms must have conformal connectivity and fulfill the DLC. This new scheme deletes elements around the element to be refined and defines the thus-created inner surface of the surrounding mesh as the advancing front. After the insertion of new nodes, the key challenge is to rebuild the mesh by starting from the advancing front and completely closing the cavity that was created by deleting the previously existing elements. In essence, for a given set of nodes, a tetrahedral mesh is created based on the discrete underlying point set and the connectivity of the surrounding mesh.

The following work was done in collaboration with Lukas Munk as a part of his master's thesis. The description of the method and the algorithm are kept brief and serve as a summary of the implementation, and to complement Chapter 3 in the thesis. The complete details of the algorithm with applications to the quasicontinuum method will be published elsewhere.

C.1 Constraint Advancing Front Algorithm

Mesh refinement requires two main steps: *(i)* the insertion of new nodes, and *(ii)* their integration into the existing mesh. Both steps are severely constrained by the requirements of the nonlocal QC method. First, as explained in Section 3.3.1, every node must satisfy the **discrete lattice constraint** (DLC), i.e., for any nodal position $\boldsymbol{X}$ in the undeformed reference mesh we must ensure that

$$\exists\, i,j,k \in Z \quad \text{s.t.} \quad \boldsymbol{X} = i\boldsymbol{a}_1 + j\boldsymbol{a}_2 + k\boldsymbol{a}_3 \tag{C.1}$$

for given Bravais basis vectors $\{\boldsymbol{a}_1, \boldsymbol{a}_2, \boldsymbol{a}_3\}$. Consequently, as opposed to FEM codes, techniques like mesh smoothing in standard forms [18] are not applicable. This is particularly challenging when approaching full atomistic resolution. At this scale, discrete lattice sites can be relatively far apart from the ideal location that a traditional mesher would choose for a new node. In addition, the simple snapping of nodes to nearest lattice sites creates unacceptable inverted or zero-volume elements [3]. Second, for the sake of parallelization, mesh adaption in a body $\mathscr{B}$ must be confined to a small subdomain $\Omega \subset \mathscr{B}$, and existing nodes at $\boldsymbol{X} \in \Omega$ must be retained at their exact prior locations. Finally, the newly-created meshes must conform with the existing nodes and element connectivities on the boundary $\partial\Omega$.

Our new CAF algorithm, which is schematically shown in Fig. C.1 and whose overall structure is summarized in Algorithm 4, consists of three principal steps that

ensure compliance with the above requirements. For a given element to be refined, we

I. *define a cavity* Ω and *delete all elements* within Ω while retaining all existing nodes,

II. *insert admissible new nodes* in Ω for refinement,

III. *rebuild element connectivities* within Ω while recovering the connectivity of the enclosing mesh on $\partial\Omega$.

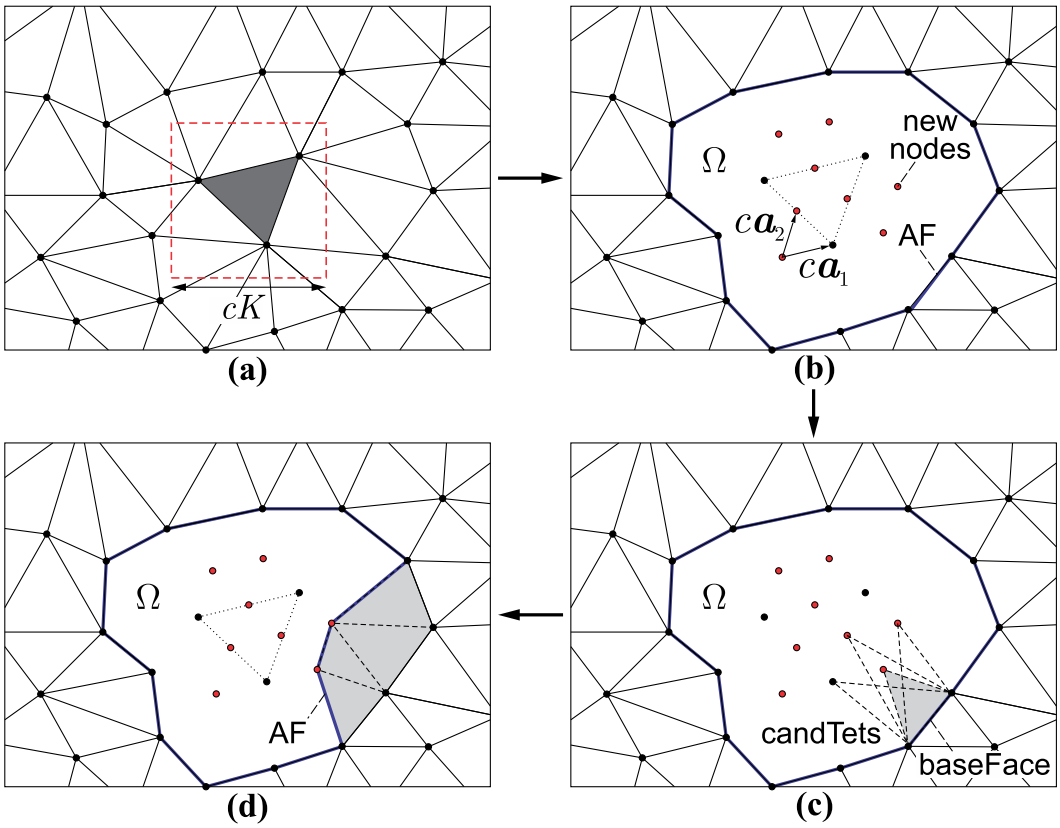

Figure C.1: Schematic illustrating the remeshing process in 2D: (a) original mesh with a highlighted element that requires refinement; surrounding elements are deleted. (b) The deleted elements form the cavity Ω and the advancing front (AF); all nodes are retained and new nodes are inserted. (c) For each face in the AF, candidate tetrahedra are identified and several checks are executed to identify the best one. (d) New elements are added and the AF is updated.

C.2 Last resort algorithm

If the above scheme does not succeed at rebuilding the mesh within the cavity in a conforming fashion, we apply a brute-force, *lastResort* algorithm to the remaining

Algorithm 4 Constrained Advancing Front

```
1:  for elementToRefine ∈ elementsToRefine do
2:      create the cavity Ω around elementToRefine
3:      insert new nodes on lattice sites
4:      initialize the advancing front: AF = ∂Ω \ (∂Ω ∩ ∂ℬ)
5:      while AF ≠ ∅ do
6:          take random baseFace ∈ AF
7:          find candidate vertices ∈ Ω near baseFace
8:          initialize set of candidateTets = ∅
9:          for each candidate vertex do
10:             if resulting candidateTet is Delaunay then
11:                 if resulting candidateTet is conformal then
12:                     add candidateTet to candidateTets
13:                 end if
14:             end if
15:         end for
16:         if candidateTets ≠ ∅ then
17:             from all candidateTets find the locally-optimal one;
18:             build optimal tet, update front, goto next baseFace
19:         else if |AF| > n_few faces then
20:             delete all created elements connected to baseFace
21:         else
22:             call lastResort algorithm
23:             if lastResort succeeds then
24:                 build tet, update front, goto next baseFace
25:             end if
26:         end if
27:         start over with this elementToRefine with a reduced cavity size
28:     end while
29: end for
```

front [45]. The recursive procedure, systematically tests a well-chosen subset of all possible meshing paths until it finds a valid tetrahedralization. This may involve deleting elements that have already been added within the cavity, leading to a less constrained cavity [91]. It is important to note that deletion will never target elements that did not belong to the initial cavity so the conformity with the surrounding mesh is preserved. As can be expected, this algorithm scales poorly with the number of faces in the AF and elements in the mesh. However, given a small remaining cavity, the algorithm will terminate whenever there exists a valid tetrahedralization in the subset. For further information the reader is referred to [45].

BIBLIOGRAPHY

[1] Marcel A. and Mitchell L. "Error Estimation and Atomistic-Continuum Adaptivity for the Quasicontinuum Approximation of a Frenkel-Kontorova Model". In: *Multiscale Modeling & Simulation* 7.1 (2008), pp. 147–170. DOI: 10.1137/070688559. eprint: http://dx.doi.org/10.1137/070688559. URL: http://dx.doi.org/10.1137/070688559.

[2] B. J. Alder and T. E. Wainwright. "Studies in Molecular Dynamics. I. General Method". In: *The Journal of Chemical Physics* 31.2 (1959), pp. 459–466. DOI: 10.1063/1.1730376. eprint: http://dx.doi.org/10.1063/1.1730376. URL: http://dx.doi.org/10.1063/1.1730376.

[3] J. S. Amelang. "A fully-nonlocal energy-based formulation and high-performance realization of the quasicontinuum method". PhD thesis. California Institute of Technology, 2015.

[4] J. S. Amelang and D. M. Kochmann. "Surface effects in nanoscale structures investigated by a fully-nonlocal energy-based quasicontinuum method". In: *Mechanics of Materials* in press (2015). ISSN: 0167-6636. DOI: http://dx.doi.org/10.1016/j.mechmat.2015.04.004. URL: http://www.sciencedirect.com/science/article/pii/S0167663615000861.

[5] J. S. Amelang, G. N. Venturini, and D. M. Kochmann. "Summation rules for the fully-nonlocal energy-based quasicontinuum method". In: *J. Mech. Phys. Solids* (2015). in press. ISSN: 0022-5096. DOI: http://dx.doi.org/10.1016/j.jmps.2015.03.007. URL: http://www.sciencedirect.com/science/article/pii/S0022509615000630.

[6] Behnam Amin-Ahmadi et al. "High resolution transmission electron microscopy characterization of fcc→9R transformation in nanocrystalline palladium films due to hydriding". In: *Applied Physics Letters* 102.7 (2013), p. 071911. DOI: 10.1063/1.4793512. eprint: https://doi.org/10.1063/1.4793512. URL: https://doi.org/10.1063/1.4793512.

[7] Hans C. Andersen. "Molecular dynamics simulations at constant pressure and/or temperature". In: *The Journal of Chemical Physics* 72.4 (1980), pp. 2384–2393. DOI: 10.1063/1.439486. eprint: http://dx.doi.org/10.1063/1.439486. URL: http://dx.doi.org/10.1063/1.439486.

[8] P.M. Anderson, J.P. Hirth, and J. Lothe. *Theory of Dislocations*. Cambridge University Press, 2016. ISBN: 9781316785102. URL: https://books.google.com/books?id=3IDBDQAAQBAJ.

[9] M. Ariza et al. "HotQC simulation of nanovoid growth under tension in copper". In: *International Journal of Fracture* 174 (1 2012), pp. 75–85. DOI: 10.1007/s10704-011-9660-4. URL: http://dx.doi.org/10.1007/s10704-011-9660-4.

[10] B. Eidel and A. Stukowski. “A variational formulation of the quasicontinuum method based on energy sampling in clusters”. In: *J. Mech. Phys. Solids* 57 (2009), pp. 87–108.

[11] Arash D. Banadaki and Srikanth Patala. “An efficient algorithm for computing the primitive bases of a general lattice plane”. In: *Journal of Applied Crystallography* 48.2 (Apr. 2015), pp. 585–588. DOI: 10.1107/S1600576715004446. URL: https://doi.org/10.1107/S1600576715004446.

[12] P.R.M. van Beers et al. “A multiscale model of grain boundary structure and energy: From atomistics to a continuum description”. In: *Acta Materialia* 82.Supplement C (2015), pp. 513–529. ISSN: 1359-6454. DOI: https://doi.org/10.1016/j.actamat.2014.08.045. URL: http://www.sciencedirect.com/science/article/pii/S1359645414006594.

[13] Marshall Bern and Paul Plassmann. “Mesh Generation”. In: *Handbook of Computational Geometry. Elsevier Science*. 2000, pp. 291–332.

[14] E. Bitzek et al. “Structural Relaxation Made Simple”. In: *Phys. Rev. Lett.* 97 (17 2006), p. 170201.

[15] E. Biyikli and A. C. To. “Multiresolution molecular mechanics: Adaptive analysis”. In: *Computer Methods in Applied Mechanics and Engineering* 305 (2016), pp. 682–702. ISSN: 0045-7825. DOI: http://dx.doi.org/10.1016/j.cma.2016.02.038. URL: http://www.sciencedirect.com/science/article/pii/S0045782516300688.

[16] C. L. Kelchner, S. J. Plimpton, and J. C. Hamilton. “Dislocation nucleation and defect structure during surface indentation”. In: *Phys. Rev. B* 58 (17 1998), pp. 11085–11088.

[17] Wei Cai et al. “Minimizing Boundary Reflections in Coupled-Domain Simulations”. In: *Phys. Rev. Lett.* 85 (15 Oct. 2000), pp. 3213–3216. DOI: 10.1103/PhysRevLett.85.3213. URL: https://link.aps.org/doi/10.1103/PhysRevLett.85.3213.

[18] Scott A. Canann, Michael B. Stephenson, and Ted Blacker. “Optismoothing: An optimization-driven approach to mesh smoothing”. In: *Finite Elements in Analysis and Design* 13.2 (1993), pp. 185–190. ISSN: 0168-874X. DOI: http://dx.doi.org/10.1016/0168-874X(93)90056-V. URL: http://www.sciencedirect.com/science/article/pii/0168874X9390056V.

[19] R. Car and M. Parrinello. “Unified Approach for Molecular Dynamics and Density-Functional Theory”. In: *Phys. Rev. Lett.* 55 (22 Nov. 1985), pp. 2471–2474. DOI: 10.1103/PhysRevLett.55.2471. URL: https://link.aps.org/doi/10.1103/PhysRevLett.55.2471.

[20] Y. A. Chang and L. Himmel. "Temperature Dependence of the Elastic Constants of Cu, Ag, and Au above Room Temperature". In: *Journal of Applied Physics* 37.9 (1966), pp. 3567–3572. DOI: 10.1063/1.1708903. eprint: https://doi.org/10.1063/1.1708903. URL: https://doi.org/10.1063/1.1708903.

[21] Xiang Chen et al. "Recent Progress in the Concurrent Atomistic-Continuum (CAC) Method and its Application in Phonon Transport". In: (2017). DOI: 10.1557/mrc.2017.116. eprint: arXiv:1708.03210.

[22] Mohammed Cherkaoui and Laurent Capolungo. "Grain Boundary Modeling". In: *Atomistic and Continuum Modeling of Nanocrystalline Materials: Deformation Mechanisms and Scale Transition*. Boston, MA: Springer US, 2009, pp. 117–142. ISBN: 978-0-387-46771-9. DOI: 10.1007/978-0-387-46771-9_5. URL: https://doi.org/10.1007/978-0-387-46771-9_5.

[23] L. P. Chew. "Constrained Delaunay Triangulations". In: *Proceedings of the Third Annual Symposium on Computational Geometry*. SCG '87. Waterloo, Ontario, Canada: ACM, 1987, pp. 215–222. ISBN: 0-89791-231-4. DOI: 10.1145/41958.41981. URL: http://doi.acm.org/10.1145/41958.41981.

[24] Peter W. Chung. "Computational method for atomistic homogenization of nanopatterned point defect structures". In: *International Journal for Numerical Methods in Engineering* 60.4 (2004), pp. 833–859. DOI: 10.1002/nme.989. URL: http://dx.doi.org/10.1002/nme.989.

[25] John D. Clayton and Peter W. Chung. "An atomistic-to-continuum framework for nonlinear crystal mechanics based on asymptotic homogenization". In: *Journal of the Mechanics and Physics of Solids* 54.8 (2006), pp. 1604–1639. DOI: 10.1016/j.jmps.2006.02.004. URL: http://www.sciencedirect.com/science/article/pii/S0022509606000408.

[26] Valerie R. Coffman and James P. Sethna. "Grain boundary energies and cohesive strength as a function of geometry". In: *Phys. Rev. B* 77 (14 Apr. 2008), p. 144111. DOI: 10.1103/PhysRevB.77.144111. URL: https://link.aps.org/doi/10.1103/PhysRevB.77.144111.

[27] David Cohen-Steiner, Pierre Alliez, and Mathieu Desbrun. "Variational Shape Approximation". In: *ACM Trans. Graph.* 23.3 (Aug. 2004), pp. 905–914. ISSN: 0730-0301. DOI: 10.1145/1015706.1015817. URL: http://doi.acm.org/10.1145/1015706.1015817.

[28] Hans Conrad and Jagdish Narayan. "On the grain size softening in nanocrystalline materials". In: *Scripta Materialia* 42.11 (2000), pp. 1025–1030. ISSN: 1359-6462. DOI: https://doi.org/10.1016/S1359-6462(00)00320-1. URL: http://www.sciencedirect.com/science/article/pii/S1359646200003201.

[29] X. D. Dai et al. “Extended Finnin-Sinclair potential for bcc and fcc metals and alloys”. In: *Journal of Physics: Condensed Matter* 18 (2006), pp. 4527–4542.

[30] D. J. Diestler, Z. B. Wu, and X. C. Zeng. “An extension of the quasicontinuum treatment of multiscale solid systems to nonzero temperature”. In: *Journal of Chemical Physics* 121 (2004), pp. 9279–9282.

[31] Q. Du, M. D. Gunzburger, and L. Ju. “Constrained Centroidal Voronoi Tessellations for Surfaces”. In: *SIAM Journal on Scientific Computing* 24.5 (2003), pp. 1488–1506. DOI: 10.1137/S1064827501391576. eprint: http://dx.doi.org/10.1137/S1064827501391576. URL: http://dx.doi.org/10.1137/S1064827501391576.

[32] O. Hardouin Duparc et al. “High-resolution transmission electron microscopy observations and atomic simulations of the structures of exact and near Σ = 11, 332 tilt grain boundaries in nickel”. In: *Philosophical Magazine A* 80.4 (2000), pp. 853–870. DOI: 10.1080/01418610008212086. eprint: https://doi.org/10.1080/01418610008212086. URL: https://doi.org/10.1080/01418610008212086.

[33] L. M. Dupuy et al. “Finite-temperature quasicontinuum: Molecular dynamics without all the atoms”. In: *Physical Review Letters* 95 (2005), p. 060202.

[34] Ryan S. Elliott and Ellad B. Tadmor. *Knowledgebase of Interatomic Models Application Programming Interface*. https://openkim.org/kim-api. Online; accessed: 2017-11-20. 2011.

[35] F. Ercolessi and J. B. Adams. “Interatomic Potentials from First-Principles Calculations: The Force-Matching Method”. In: *EPL (Europhysics Letters)* 26.8 (1994), p. 583. URL: http://stacks.iop.org/0295-5075/26/i=8/a=005.

[36] Frédéric Feyel and Jean-Louis Chaboche. “FE2 multiscale approach for modelling the elastoviscoplastic behaviour of long fibre SiC/Ti composite materials”. In: *Computer Methods in Applied Mechanics and Engineering* 183.3 (2000), pp. 309–330. ISSN: 0045-7825. DOI: https://doi.org/10.1016/S0045-7825(99)00224-8. URL: http://www.sciencedirect.com/science/article/pii/S0045782599002248.

[37] R. P. Feynman. “Forces in Molecules”. In: *Phys. Rev.* 56 (4 Aug. 1939), pp. 340–343. DOI: 10.1103/PhysRev.56.340. URL: https://link.aps.org/doi/10.1103/PhysRev.56.340.

[38] Jacob Fish et al. “Concurrent AtC coupling based on a blend of the continuum stress and the atomistic force”. In: *Computer Methods in Applied Mechanics and Engineering* 196.45 (2007), pp. 4548–4560. ISSN: 0045-7825. DOI: https://doi.org/10.1016/j.cma.2007.05.020. URL: http://www.sciencedirect.com/science/article/pii/S0045782507002277.

[39] J. Frykestig. “Advancing Front Mesh Generation Techniques with Application to the Finite ElementMethod Pub. 94:10”. PhD thesis. Chalmers University of Technology, Göteborg, Sweden, 1994.

[40] Timothy C. Germann and Kai Kadau. “TRILLION-ATOM MOLECULAR DYNAMICS BECOMES A REALITY”. In: *International Journal of Modern Physics C* 19.09 (2008), pp. 1315–1319. DOI: 10.1142/S0129183108012911. eprint: http://www.worldscientific.com/doi/pdf/10.1142/S0129183108012911. URL: http://www.worldscientific.com/doi/abs/10.1142/S0129183108012911.

[41] Julia R. Greer and Jeff Th.M. De Hosson Jeff Th.M. “Plasticity in small-sized metallic systems: Intrinsic versus extrinsic size effect”. In: *Progress in Materials Science* 56.6 (Aug. 2011), pp. 654–724. ISSN: 00796425. DOI: 10.1016/j.pmatsci.2011.01.005.

[42] Julia R. Greer and Jeff Th.M. De Hosson. “Plasticity in small-sized metallic systems: Intrinsic versus extrinsic size effect”. In: *Progress in Materials Science* 56.6 (2011). Festschrift Vaclav Vitek, pp. 654–724. ISSN: 0079-6425. DOI: https://doi.org/10.1016/j.pmatsci.2011.01.005. URL: http://www.sciencedirect.com/science/article/pii/S0079642511000065.

[43] M. Gunzburger and Y. Zhang. “A quadrature-rule type approximation to the quasi-continuum method”. In: *Multiscale Model. Simul.* 8.2 (2010), pp. 571–590.

[44] E. O. Hall. “The deformation and ageing of mild steel”. In: *Proceedings of the Physical Society B* 64 (1951), pp. 742–753.

[45] Ziji Wu Hamid Ghadyani John Sullivan. “Boundary Recovery For Delaunay Tetrahedral Meshes Using Local Topological Transformations”. In: *Finite elements in analysis and design : the international journal of applied finite elements and computer aided engineering* 46 (2010), pp. 74–83. DOI: 10.1016/j.finel.2009.06.022.

[46] G. Hasson et al. “Theoretical and experimental determinations of grain boundary structures and energies: Correlation with various experimental results”. In: *Surface Science* 31.Supplement C (1972), pp. 115–137. ISSN: 0039-6028. DOI: https://doi.org/10.1016/0039-6028(72)90256-7. URL: http://www.sciencedirect.com/science/article/pii/0039602872902567.

[47] M. Iyer and V. Gavini. “A field theoretical approach to the quasi-continuum method”. In: *Journal of the Mechanics and Physics of Solids* 59.8 (2011), pp. 1506–1535. ISSN: 0022-5096. DOI: 10.1016/j.jmps.2010.12.002. URL: http://www.sciencedirect.com/science/article/pii/S0022509610002425.

[48] J. E. Savage J. G. Castanos. "PARED: a framework for the adaptive solution of PDEs". In: (1999), pp. 133–140. ISSN: 1082-8907. DOI: 10.1109/HPDC.1999.805291.

[49] E. T. Jaynes. "Information Theory and Statistical Mechanics". In: *Phys. Rev.* 106 (4 May 1957), pp. 620–630. DOI: 10.1103/PhysRev.106.620. URL: https://link.aps.org/doi/10.1103/PhysRev.106.620.

[50] R. A. Johnson. "Alloy models with the embedded-atom method". In: *Phys. Rev. B* 39 (17 June 1989), pp. 12554–12559. DOI: 10.1103/PhysRevB.39.12554. URL: http://link.aps.org/doi/10.1103/PhysRevB.39.12554.

[51] R. A. Johnson. "Analytic nearest-neighbor model for fcc metals". In: *Phys. Rev. B* 37 (8 Mar. 1988), pp. 3924–3931. DOI: 10.1103/PhysRevB.37.3924. URL: http://link.aps.org/doi/10.1103/PhysRevB.37.3924.

[52] Mark T. Jones and Paul E. Plassmann. "Adaptive refinement of unstructured finite-element meshes". In: *Finite Elements in Analysis and Design* 25.1–2 (1997). Adaptive Meshing, Part 1, pp. 41–60. ISSN: 0168-874X. DOI: http://dx.doi.org/10.1016/S0168-874X(96)00039-X. URL: http://www.sciencedirect.com/science/article/pii/S0168874X9600039X.

[53] W. K. Kim et al. "Hyper-QC: An accelerated finite-temperature quasicontinuum method using hyperdynamics". In: *Journal of the Mechanics and Physics of Solids* 63.0 (2014), pp. 94–112. ISSN: 0022-5096. DOI: http://dx.doi.org/10.1016/j.jmps.2013.10.001. URL: http://www.sciencedirect.com/science/article/pii/S0022509613002032.

[54] J. Knap and M. Ortiz. "An analysis of the quasicontinuum method". In: *J. Mech. Phys. Solids* 49.9 (Sept. 2001), pp. 1899–1923.

[55] D. M. Kochmann and G. N. Venturini. "A meshless quasicontinuum method based on local maximum-entropy interpolation". In: *Model. Simul. Mater. Sci. Eng.* 22 (2014), p. 034007.

[56] S. Kohlhoff, P. Gumbsch, and H. F. Fischmeister. "Crack propagation in b.c.c. crystals studied with a combined finite-element and atomistic model". In: *Philosophical Magazine A* 64.4 (1991), pp. 851–878. DOI: 10.1080/01418619108213953. eprint: http://dx.doi.org/10.1080/01418619108213953. URL: http://dx.doi.org/10.1080/01418619108213953.

[57] M L Kronberg and Wilson F H. "Secondary Recrystallization in Copper". In: *Transactions of the American Institute of Mining and Metallurgical Engineers* 185 (1949), pp. 501–514. ISSN: 1480064.

[58] L.P Kubin and G Canova. "The modelling of dislocation patterns". In: *Scripta Metallurgica et Materialia* 27.8 (1992), pp. 957–962. ISSN: 0956-716X. DOI: https://doi.org/10.1016/0956-716X(92)90456-O. URL: http://www.sciencedirect.com/science/article/pii/0956716X92904560.

[59] Y. Kulkarni. "Coarse-graining of atomistic description at finite temperature". PhD thesis. California Institute of Technology, 2007. URL: http://thesis.library.caltech.edu/4498/.

[60] Y. Kulkarni, J. Knap, and M. Ortiz. "A variational approach to coarse graining of equilibrium and non-equilibrium atomistic description at finite temperature". In: *Journal of the Mechanics and Physics of Solids* 56 (2008), pp. 1417–1449.

[61] R. LeSar, R. Najafabadi, and D. J. Srolovitz. "Finite-temperature defect properties from free-energy minimization". In: *Phys. Rev. Lett.* 63 (6 Aug. 1989), pp. 624–627. DOI: 10.1103/PhysRevLett.63.624. URL: https://link.aps.org/doi/10.1103/PhysRevLett.63.624.

[62] B. Li, F. Habbal, and M. Ortiz. "Optimal transportation meshfree approximation schemes for fluid and plastic flows". In: *International Journal for Numerical Methods in Engineering* 83.12 (2010), pp. 1541–1579. ISSN: 1097-0207. DOI: 10.1002/nme.2869. URL: http://dx.doi.org/10.1002/nme.2869.

[63] Xingjie Helen Li et al. "Theory-based benchmarking of the blended force-based quasicontinuum method". In: *Computer Methods in Applied Mechanics and Engineering* 268.0 (2014), pp. 763–781. ISSN: 0045-7825. DOI: http://dx.doi.org/10.1016/j.cma.2013.10.007. URL: http://www.sciencedirect.com/science/article/pii/S0045782513002594.

[64] Computational Solid Mechanics Lagoratory LSMS. https://lsms.epfl.ch/page-52155-en.html. 2017.

[65] K. Lu and M.L. Sui. "An explanation to the abnormal Hall-Petch relation in nanocrystalline materials". In: *Scripta Metallurgica et Materialia* 28.12 (1993), pp. 1465–1470. ISSN: 0956-716X. DOI: https://doi.org/10.1016/0956-716X(93)90576-E. URL: http://www.sciencedirect.com/science/article/pii/0956716X9390576E.

[66] M. S. Daw and M. I. Baskes. "Embedded-atom method: Derivation and application to impurities, surfaces, and other defects in metals". In: *Phys. Rev. B* 29 (12 1984), pp. 6443–6453.

[67] J. Marian et al. "Finite-temperature extension of the quasicontinuum method using Langevin dynamics: entropy losses and analysis of errors". In: *Modelling and Simulation in Materials Science and Engineering* 18.1 (2010), p. 015003.

[68] M.I. Mendelev and A.H. King. "The interactions of self-interstitials with twin boundaries". In: *Philosophical Magazine* 93.10-12 (2013), pp. 1268–1278. DOI: 10.1080/14786435.2012.747012. eprint: http://dx.doi.org/10.1080/14786435.2012.747012. URL: http://dx.doi.org/10.1080/14786435.2012.747012.

[69] M.I. Mendelev et al. “Analysis of semi-empirical interatomic potentials appropriate for simulation of crystalline and liquid Al and Cu”. In: *Philosophical Magazine* 88.12 (2008), pp. 1723–1750. DOI: 10.1080/14786430802206482. eprint: http://www.tandfonline.com/doi/pdf/10.1080/14786430802206482. URL: http://www.tandfonline.com/doi/abs/10.1080/14786430802206482.

[70] Nicholas Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092. DOI: 10.1063/1.1699114. eprint: http://dx.doi.org/10.1063/1.1699114. URL: http://dx.doi.org/10.1063/1.1699114.

[71] Christian Miehe, Jörg Schröder, and Jan Schotte. “Computational homogenization analysis in finite plasticity Simulation of texture development in polycrystalline materials”. In: *Computer Methods in Applied Mechanics and Engineering* 171.3 (1999), pp. 387–418. ISSN: 0045-7825. DOI: https://doi.org/10.1016/S0045-7825(98)00218-7. URL: http://www.sciencedirect.com/science/article/pii/S0045782598002187.

[72] M.J. Mills et al. “High-resolution transmission electron microscopy of grain boundaries in aluminum and correlation with atomistic calculations”. In: *Ultramicroscopy* 40.3 (1992), pp. 247–257. ISSN: 0304-3991. DOI: https://doi.org/10.1016/0304-3991(92)90121-Y. URL: http://www.sciencedirect.com/science/article/pii/030439919290121Y.

[73] Y. Mishin et al. “Structural stability and lattice defects in copper: ¡i¿Ab initio¡/i¿, tight-binding, and embedded-atom calculations”. In: *Phys. Rev. B* 63 (22 May 2001), p. 224106. DOI: 10.1103/PhysRevB.63.224106. URL: http://link.aps.org/doi/10.1103/PhysRevB.63.224106.

[74] T.G. Nieh and J. Wadsworth. “Hall-Petch relation in nanocrystalline solids”. In: *Scripta Metallurgica et Materialia* 25 (1991), p. 955.

[75] NIST. http://www.ctcms.nist.gov/potentials. 2016.

[76] F. C. Nix and D. MacNair. “The Thermal Expansion of Pure Metals: Copper, Gold, Aluminum, Nickel, and Iron”. In: *Phys. Rev.* 60 (8 Oct. 1941), pp. 597–605. DOI: 10.1103/PhysRev.60.597. URL: https://link.aps.org/doi/10.1103/PhysRev.60.597.

[77] Hiroshi Ogawa, Fumihiro Wakai, and Yoshio Waseda. “Molecular Dynamics Simulation of the Model Grain Boundary Structure of Polycrystalline Materials”. In: *Molecular Simulation* 18.3 (1996), pp. 179–192. DOI: 10.1080/08927029608024122. eprint: http://dx.doi.org/10.1080/08927029608024122. URL: http://dx.doi.org/10.1080/08927029608024122.

[78] David L. Olmsted, Stephen M. Foiles, and Elizabeth A. Holm. “Survey of computed grain boundary properties in face-centered cubic metals: I. Grain boundary energy”. In: *Acta Materialia* 57.13 (2009), pp. 3694–3703. ISSN: 1359-6454. DOI: https://doi.org/10.1016/j.actamat.2009.04.

007. URL: http://www.sciencedirect.com/science/article/pii/S1359645409002274.

[79] "On the determination of molecular fields. —II. From the equation of state of a gas". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 106.738 (1924), pp. 463–477. ISSN: 0950-1207. DOI: 10.1098/rspa.1924.0082. eprint: http://rspa.royalsocietypublishing.org/content/106/738/463.full.pdf. URL: http://rspa.royalsocietypublishing.org/content/106/738/463.

[80] C. Ortner and H. Wang. "A posteriori error control for a quasi-continuum approximation of a periodic chain". In: *IMA Journal of Numerical Analysis* (2013). DOI: 10.1093/imanum/drt011. eprint: http://imajna.oxfordjournals.org/content/early/2013/10/05/imanum.drt011.full.pdf+html. URL: http://imajna.oxfordjournals.org/content/early/2013/10/05/imanum.drt011.abstract.

[81] W. C. Overton and J. Gaffney. "Temperature Variation of the Elastic Constants of Cubic Elements. I. Copper". In: *Physical Review* 98 (May 1955), pp. 969–977. DOI: 10.1103/PhysRev.98.969.

[82] Srikanth Patala and Banadaki Arash D. https://pypi.python.org/pypi/GBpy. 2015.

[83] Lejcek Pavel. *Grain Boundary Segregation in Metals*. Springer, Berlin, Heidelberg, 2010. ISBN: 978-3-642-12505-8. DOI: https://doi.org/10.1007/978-3-642-12505-8.

[84] V. Péron-Lührs, F. Sansoz, and L. Noels. "Quasicontinuum study of the shear behavior of defective tilt grain boundaries in Cu". In: *Acta Materialia* 64.Supplement C (2014), pp. 419–428. ISSN: 1359-6454. DOI: https://doi.org/10.1016/j.actamat.2013.10.056. URL: http://www.sciencedirect.com/science/article/pii/S1359645413008197.

[85] N. J. Petch. "The cleavage strength of polycrystalsl". In: *Journal of the Iron and Steel Institute* 174 (1953), pp. 25–28.

[86] Rob Phillips. *Crystals, Defects and Microstructures: Modeling Across Scales*. Cambridge University Press, 2001. DOI: 10.1017/CBO9780511606236.

[87] M. Ponga, M. Ortiz, and M. P. Ariza. "Finite-temperature non-equilibrium quasi-continuum analysis of nanovoid growth in copper at low and high strain rates". In: *Mechanics of Materials* in press (2015). ISSN: 0167-6636. DOI: http://dx.doi.org/10.1016/j.mechmat.2015.02.007. URL: http://www.sciencedirect.com/science/article/pii/S016766361500054X.

[88] S Qu et al. “A finite-temperature dynamic coupled atomistic/discrete dislocation method”. In: *Modelling and Simulation in Materials Science and Engineering* 13.7 (2005), p. 1101. URL: http://stacks.iop.org/0965-0393/13/i=7/a=007.

[89] R. E. Miller and E. B. Tadmor. “A unified framework and performance benchmark of fourteen multiscale atomistic/continuum coupling methods”. In: *Modelling Simul. Mater. Sci. Eng.* 17 (2009), p. 053001.

[90] A. Rahman. “Correlations in the Motion of Atoms in Liquid Argon”. In: *Physical Review* 136 (Oct. 1964), pp. 405–411. DOI: 10.1103/PhysRev.136.A405.

[91] A. Rassineux. “Generation and optimization of tetrahedral meshes by advancing front technique”. In: *International Journal for Numerical Methods in Engineering* 41.4 (1998), pp. 651–674. ISSN: 1097-0207. DOI: 10.1002/(SICI)1097-0207(19980228)41:4<651::AID-NME304>3.0.CO;2-P. URL: http://dx.doi.org/10.1002/(SICI)1097-0207(19980228)41:4%3C651::AID-NME304%3E3.0.CO;2-P.

[92] W. T. Read and W. Shockley. “Dislocation Models of Crystal Grain Boundaries”. In: *Phys. Rev.* 78 (3 May 1950), pp. 275–289. DOI: 10.1103/PhysRev.78.275. URL: https://link.aps.org/doi/10.1103/PhysRev.78.275.

[93] Maria-Cecilia Rivara. “Lepp-bisection algorithms, applications and mathematical properties”. In: *Applied Numerical Mathematics* 59.9 (2009). Second Chilean Workshop on Numerical Analysis of Partial Differential Equations (WONAPDE 2007), pp. 2218–2235. ISSN: 0168-9274. DOI: http://dx.doi.org/10.1016/j.apnum.2008.12.011. URL: http://www.sciencedirect.com/science/article/pii/S0168927408002146.

[94] M.-C. Rivara. “Algorithms for refining triangular grids suitable for adaptive and multigrid techniques”. In: *International Journal for Numerical Methods in Engineering* 20.4 (1984), pp. 745–756. ISSN: 1097-0207. DOI: 10.1002/nme.1620200412. URL: http://dx.doi.org/10.1002/nme.1620200412.

[95] P. A. Rodriguez and M.-C. Rivara. “Multithread Lepp-Bisection Algorithm for Tetrahedral Meshes”. English. In: *Proceedings of the 22nd International Meshing Roundtable*. Ed. by Josep Sarrate and Matthew Staten. Springer International Publishing, 2014, pp. 525–540. ISBN: 978-3-319-02334-2. DOI: 10.1007/978-3-319-02335-9_29. URL: http://dx.doi.org/10.1007/978-3-319-02335-9_29.

[96] R. E. Rudd and J. Q. Broughton. “Coarse-grained molecular dynamics: Nonlinear finite elements and finite temperature”. In: *Physical Review B* 72 (2005), p. 144104.

[97] R.E. Rudd and J.Q. Broughton. "Concurrent Coupling of Length Scales in Solid State Systems". In: *physica status solidi (b)* 217.1 (2000), pp. 251–291. ISSN: 1521-3951. DOI: 10.1002/(SICI)1521-3951(200001)217:1<251::AID-PSSB251>3.0.CO;2-A. URL: http://dx.doi.org/10.1002/(SICI)1521-3951(200001)217:1%3C251::AID-PSSB251%3E3.0.CO;2-A.

[98] Robert E. Rudd and Jeremy Q. Broughton. "Coarse-grained molecular dynamics and the atomic limit of finite elements". In: *Phys. Rev. B* 58 (10 Sept. 1998), R5893–R5896. DOI: 10.1103/PhysRevB.58.R5893. URL: http://link.aps.org/doi/10.1103/PhysRevB.58.R5893.

[99] Jean-Paul Ryckaert, Giovanni Ciccotti, and Herman J.C Berendsen. "Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes". In: *Journal of Computational Physics* 23.3 (1977), pp. 327–341. ISSN: 0021-9991. DOI: https://doi.org/10.1016/0021-9991(77)90098-5. URL: http://www.sciencedirect.com/science/article/pii/0021999177900985.

[100] Plimpton S. http://lammps.sandia.gov/papers.html. 1997.

[101] Plimpton S. http://lammps.sandia.gov/. 1997.

[102] Debin Shan, Lumeng Wang, and Lin Yuan. "Effect of the Σ5(310)/[001]ϑ = 53.1° grain boundary on the incipient yield of bicrystal copper: A quasicontinuum simulation and nanoindentation experiment". In: *Journal of Materials Research* 28.5 (2013), pp. 766–773. DOI: 10.1557/jmr.2012.424.

[103] V. Shenoy, V. Shenoy, and R. Phillips. "Finite temperature quasicontinuum methods". In: *Materials Research Society Symposium Proceedings* 538 (1999), pp. 465–471.

[104] V.B. Shenoy et al. "An adaptive finite element approach to atomic-scale mechanics - the quasicontinuum method". In: *Journal of the Mechanics and Physics of Solids* 47.3 (1999), pp. 611–642. ISSN: 0022-5096. DOI: http://dx.doi.org/10.1016/S0022-5096(98)00051-9. URL: http://www.sciencedirect.com/science/article/pii/S0022509698000519.

[105] J. R. Shewchuk. "Delaunay Refinement Algorithms for Triangular Mesh Generation". In: *Comput. Geom. Theory Appl.* 22.1-3 (May 2002), pp. 21–74. ISSN: 0925-7721. DOI: 10.1016/S0925-7721(01)00047-5. URL: http://dx.doi.org/10.1016/S0925-7721(01)00047-5.

[106] L.E. Shilkrot, R. E. Miller, and W. A. Curtin. "Multiscale plasticity modeling: coupled atomistics and discrete dislocation mechanics". In: *Journal of the Mechanics and Physics of Solids* 52 (2004), pp. 755–787.

[107] Frank H. Stillinger and Aneesur Rahman. "Improved simulation of liquid water by molecular dynamics". In: *The Journal of Chemical Physics* 60.4 (1974), pp. 1545–1557. DOI: 10.1063/1.1681229. eprint: http://dx.

doi.org/10.1063/1.1681229. URL: http://dx.doi.org/10.1063/1.1681229.

[108] A.H Stroud. *Approximate calculation of multiple integrals by A.H. Stroud*. Vol. 1. J. Johnson, 1971. ISBN: 0130438936.

[109] A.P. Sutton and J. Chen. "Long-range Finnis-Sinclair potentials". In: *Philosophical Magazine Letters* 61 (3 1990), pp. 139–146.

[110] E. B. Tadmor and R. E. Miller. http://qcmethod.org/. 2015.

[111] E. B. Tadmor, M. Ortiz, and R. Phillips. "Quasicontinuum analysis of defects in solids". In: *Philos. Mag. A* 73.6 (June 1996), pp. 1529–1563.

[112] E. B. Tadmor et al. "Nanoindentation and incipient plasticity". In: *Journal of Materials Research* 14 (1999), pp. 2233–2250. DOI: 10.1557/JMR.1999.0300.

[113] E. Tadmor et al. "Finite-Temperature Quasi-Continuum". In: *Applied Mechanics Reviews* 65.1 (2013), p. 010803. DOI: 10.1115/1.4023013. URL: http://hal.archives-ouvertes.fr/hal-00849048.

[114] Z. Tang et al. "Finite-temperature quasicontinuum method for multiscale analysis of silicon nanostructures". In: *Physical Review B* 74 (2006), p. 064110.

[115] I. Tembhekar et al. "Automatic adaptivity in the fully nonlocal quasicontinuum method for coarse-grained atomistic simulations". In: *International Journal for Numerical Methods in Engineering* 110.9 (2017).
Contributions: Ishan Tembhekar performed all the numerical experiments and simulations and participated in the writing of the manuscript., pp. 878–900. ISSN: 1097-0207. DOI: 10.1002/nme.5438. URL: http://dx.doi.org/10.1002/nme.5438.

[116] The CGAL Project. *CGAL User and Reference Manual*. 4.7. CGAL Editorial Board, 2015. URL: http://doc.cgal.org/4.7/Manual/packages.html.

[117] M. A. Tschopp and D. L. Mcdowell. "Asymmetric tilt grain boundary structure and energy in copper and aluminium". In: *Philosophical Magazine* 87.25 (2007), pp. 3871–3892. DOI: 10.1080/14786430701455321. eprint: https://doi.org/10.1080/14786430701455321. URL: https://doi.org/10.1080/14786430701455321.

[118] V. B. Shenoy et al. "Quasicontinuum models of interfacial structure and deformation". In: *Phys. Rev. Lett.* 80 (1998), pp. 742–745.

[119] G. Venturini. "Topics in Multiscale Modeling of Metals and Metallic Alloys". PhD thesis. California Institute of Technology, 2010.

[120] G. Venturini et al. “Atomistic long-term simulation of heat and mass transport”. In: *Journal of the Mechanics and Physics of Solids* 73.0 (2014), pp. 242–268. ISSN: 0022-5096. DOI: http://dx.doi.org/10.1016/j.jmps.2014.09.008. URL: http://www.sciencedirect.com/science/article/pii/S002250961400194X.

[121] G. Venturini et al. “Thermal Expansion Behavior of Al and Ta using a finite-temperature extension of the quasicontinuum method”. In: *International Journal for Multiscale Computational Engineering* 10.1 (2012), pp. 1–11.

[122] Loup Verlet. “Computer ”Experiments” on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules”. In: *Phys. Rev.* 159 (1 July 1967), pp. 98–103. DOI: 10.1103/PhysRev.159.98. URL: http://link.aps.org/doi/10.1103/PhysRev.159.98.

[123] Gregory J. Wagner and Wing Kam Liu. “Coupling of Atomistic and Continuum Simulations Using a Bridging Scale Decomposition”. In: *J. Comput. Phys.* 190.1 (Sept. 2003), pp. 249–274. ISSN: 0021-9991. DOI: 10.1016/S0021-9991(03)00273-0. URL: http://dx.doi.org/10.1016/S0021-9991(03)00273-0.

[124] Gui Jin Wang, A.P. Sutton, and V. Vitek. “A computer simulation study of ¡001¿ and ¡111¿ tilt boundaries: the multiplicity of structures”. In: *Acta Metallurgica* 32.7 (1984), pp. 1093–1104. ISSN: 0001-6160. DOI: https://doi.org/10.1016/0001-6160(84)90013-0. URL: http://www.sciencedirect.com/science/article/pii/0001616084900130.

[125] X. Wang and X. Guo. “Quasi-Continuum Model for the Finite Deformation of Single-Layer Graphene Sheets Based on the Temperature-Related Higher Order Cauchy-Born Rule”. In: *Journal of Computational and Theoretical Nanoscience* 10.1 (2013), pp. 154–164. DOI: doi:10.1166/jctn.2013.2672. URL: http://www.ingentaconnect.com/content/asp/jctn/2013/00000010/00000001/art00024.

[126] Ulrike G. K. Wegst et al. “Bioinspired structural materials”. In: *Nat Mater* 14.1 (Jan. 2015), pp. 23–36. URL: http://dx.doi.org/10.1038/nmat4089.

[127] Mingjian Wen. *Linear thermal expansion coefficient of fcc Al at room temperature under zero pressure*. https://openkim.org/cite/TE_957040092249_000. Online; accessed: 2017-11-20. 2015.

[128] J. M. Wernik and S. A. Meguid. “Coupling atomistics and continuum in solids: status, prospects, and challenges”. In: *International Journal of Mechanics and Materials in Design* 5.1 (Sept. 2008), p. 79. ISSN: 1573-8841. DOI: 10.1007/s10999-008-9087-x. URL: https://doi.org/10.1007/s10999-008-9087-x.

[129] J M Winey, Alison Kubota, and Y M Gupta. “A thermodynamic approach to determine accurate potentials for molecular dynamics simulations: thermoelastic response of aluminum”. In: *Modelling and Simulation in Materials Science and Engineering* 17.5 (2009), p. 055004. URL: http://stacks.iop.org/0965-0393/17/i=5/a=055004.

[130] S.P. Xiao and T. Belytschko. “A bridging domain method for coupling continua with molecular dynamics”. In: *Computer Methods in Applied Mechanics and Engineering* 193.17-20 (2004). ¡ce:title¿Multiple Scale Methods for Nanoscale Mechanics and Materials¡/ce:title¿, pp. 1645–1669. ISSN: 0045-7825. DOI: 10.1016/j.cma.2003.12.053. URL: http://www.sciencedirect.com/science/article/pii/S004578250400026X.

[131] Shuozhi Xu et al. “Sequential slip transfer of mixed-character dislocations across Σ3 coherent twin boundary in FCC metals: a concurrent atomistic-continuum study”. In: *Npj Computational Materials* 2 (Jan. 2016), 15016 EP -. URL: http://dx.doi.org/10.1038/npjcompumats.2015.16.

[132] Q. Yang, Biyikli. E., and A.C. To. “Multiresolution molecular mechanics: Statics”. In: *Computer Methods in Applied Mechanics and Engineering* 258.0 (2013), pp. 26–38. ISSN: 0045-7825. DOI: http://dx.doi.org/10.1016/j.cma.2013.01.014. URL: http://www.sciencedirect.com/science/article/pii/S004578251300025X.

[133] W. Yu and Z. Wang. “Interactions between edge lattice dislocations and Σ11 symmetrical tilt grain boundary: comparisons among several FCC metals and interatomic potentials”. In: *Philosophical Magazine* 94.20 (2014), pp. 2224–2246. DOI: 10.1080/14786435.2014.910318. eprint: http://dx.doi.org/10.1080/14786435.2014.910318. URL: http://dx.doi.org/10.1080/14786435.2014.910318.

[134] Wenshan Yu, Zhiqiang Wang, and Shengping Shen. “Edge dislocations interacting with a Σ11 symmetrical grain boundary in copper upon mixed loading: A quasicontinuum method study”. In: *Computational Materials Science* 137.Supplement C (2017), pp. 162–170. ISSN: 0927-0256. DOI: https://doi.org/10.1016/j.commatsci.2017.05.033. URL: http://www.sciencedirect.com/science/article/pii/S0927025617302732.

[135] Min Zhou. “A new look at the atomic level virial stress: on continuum-molecular system equivalence”. In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 459.2037 (2003), pp. 2347–2392. ISSN: 1364-5021. DOI: 10.1098/rspa.2003.1127. eprint: http://rspa.royalsocietypublishing.org/content/459/2037/2347.full.pdf. URL: http://rspa.royalsocietypublishing.org/content/459/2037/2347.

[136] X. W. Zhou, R. A. Johnson, and H. N. G. Wadley. “Misfit-energy-increasing dislocations in vapor-deposited CoFe/NiFe multilayers”. In: *Phys. Rev. B* 69

(14 Apr. 2004), p. 144113. DOI: 10.1103/PhysRevB.69.144113. URL: https://link.aps.org/doi/10.1103/PhysRevB.69.144113.

[137] Rajendra R. Zope and Y. Mishin. "Interatomic potentials for atomistic simulations of the Ti-Al system". In: *Phys. Rev. B* 68 (2 July 2003), p. 024102. DOI: 10.1103/PhysRevB.68.024102. URL: https://link.aps.org/doi/10.1103/PhysRevB.68.024102.