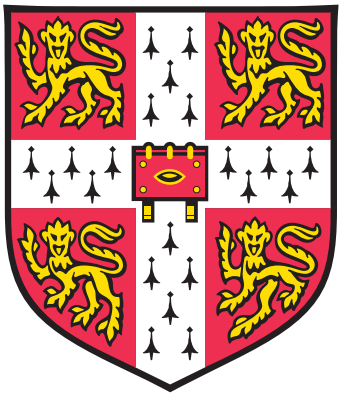

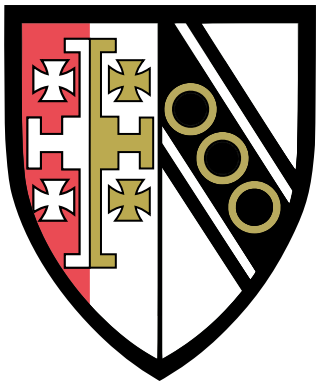

Numerical Constitutive Modelling for Continuum Mechanics Simulation

Oliver Strickson

Selwyn College

and

The Cavendish Laboratory,

University of Cambridge

September 2015

This dissertation is submitted for the degree of Doctor of Philosophy

Declaration

Cavendish Laboratory
September 2015

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University of similar institution except as declared in the Preface and specified in the text. It does not exceed the limit of sixty thousand words prescribed by the Degree Committee for Physics and Chemistry.

Oliver Strickson

Acknowledgements

This dissertation benefited greatly from the support of many individuals, and would not have been possible otherwise. Firstly, I wish to thank my academic supervisors, Prof. Emilio Artacho and Dr. Nikos Nikiforakis, who provided much valuable advice, mentorship and day-to-day encouragement. My industrial supervisor, Dr. Alan Minchinton of Orica Ltd. was always supportive when discussing new ideas for the project, and I also wish to thank him for many useful discussions during several research visits of Orica to the department.

I wish to thank my many colleagues and friends in the Laboratory for Scientific Computing at the Cavendish. Dr. Philip Blakely has been an invaluable source of C++ wisdom and programming advice.

Chapter 6 is the product of working closely with the group of Dr. Gábor Csányi in the Engineering Department, as well as Noam Bernstein of the United States Naval Research Laboratory. I wish to thank everyone in the group there for making me feel welcome, and particularly Dr. Albert Bartók-Pártay who was always happy to offer his advice on potential development and the workings of GAP and the QUIP code. This project also benefited from discussions with Prof. Keivan Esfarjani of MIT.

I wish to thank several members of the Shock Microstructure and Fracture group in the Cavendish Laboratory, for discussions from an experimentalist's perspective, and who also have variously engaged in the Orica collaboration. In particular, Dr. Andrew Jardine, Dr. David Williamson, Dr. Simon Kirk, and Dr. Stephen Walley for his vast knowledge of the shock-wave literature.

Parts of this work were carried out over several extended visits to CIC NanoGUNE, Donostia–San Sebastián, Spain, and I extend my thanks to everyone there who made my visits enjoyable and productive.

I am grateful to my parents, Susan and Anthony Strickson, who have always been supportive of my academic studies.

Finally, I gratefully acknowledge the financial support of my sponsors, Orica Ltd.

Contents

1. Introduction

Computer simulations of macroscopic physical systems inevitably depend on the materials involved, and good theoretical models of the materials themselves are essential in providing quantitative accuracy from these simulations.

Computational modelling techniques based on the numerical solution of approximations to the Schrödinger equation, in particular density functional theory (DFT), have been developed over the last few decades into mature, practical theories for computing ground-state properties of systems containing, at present, around a thousand atoms. This has been further aided by the parallel and continuing increases in computing power. Periodic boundary conditions along with Bloch's theorem allow the simulation of crystalline materials, and the determination of their bulk properties. For a variety of materials, density functional theory has proven capable of reproducing many ground state properties accurately, including lattice constants, elastic constants, phonon spectra and the location of phase transitions in pressure and temperature. See for example the reviews of Payne et al. (1992) and Hafner et al. (2006). It has become an important tool in physics, chemistry and materials science. The challenges and limitations of density functional theory are also well documented (Cohen et al., 2012).

Often effects across many disparate lengthscales or timescales, beyond the properties that can be determined from small-scale simulations of only thousands of atoms, are important in understanding a material's behaviour. A particular challenge on a computer, having finite resources of memory and time, is the presence in a problem of multiple length and time scales that differ by many orders of magnitude. Trying to capture several such effects at once is known as multiscale modelling. A recent review of techniques is given by Elliott (2011). Situations where a multiscale approach is productive include metal plasticity and hardening, surface diffusion, brittle crack propagation and many others. There is no catch-all solution for this class of problem, and it is an active area of research.

Macroscopic systems are typically modelled on the continuum scale, and for condensed-matter systems, this can be appropriate for systems between micrometres and kilometres

in size. Continuum mechanics treats a material as a continuous field of macroscopic, averaged quantities. Solutions (both analytical and numerical) to problems stated in these terms have been remarkably successful in describing and predicting many diverse physical phenomena, of the behaviour of fluids (Newtonian and otherwise, in various flow regimes (Hafez et al., 2010; Shang, 2004)), solids (linear and nonlinear elastic, as well as plasticity and fracture (Fung and Tong, 2001; Simo and Hughes, 2006)) and electrically-conducting fluids (magnetohydrodynamics, with applications to geophysics, astrophysics and plasma confinement (Ledvina et al., 2008; Tang and Chan, 2005)).

Constitutive relations are laws, either phenomenological or based on an underlying physical model, that represent the detailed processes that are not contained explicitly within a larger-scale model. They allow this model to be specialised to a particular system or material. The canonical example of a constitutive relation is the thermodynamic equation of state of a fluid.

Obtaining constitutive laws describing the bulk properties of crystalline materials from molecular systems is the common theme of the three investigations presented in this dissertation. We use 'first-principles' techniques where possible, an approach offering results which are predictive, with applicability to a wide class of materials, and with a systematic way to apply the techniques to any *particular* material of interest. Constitutive relations so-obtained can then be applied to problems in a continuum setting.

We use silicon throughout as a material to demonstrate these techniques. Silicon has long been a model material used for molecular dynamics simulations, and at the level of density functional theory, diamond-phase silicon is very efficient to compute, particularly for local atomic orbital based codes (such as SIESTA) due to being an insulator with a small number of valence electrons and a low coordination number. Despite this, it offers much interest in its own right, including a rich phase diagram with a structural phase-transition to a metallic phase, and also several challenges for its simulation, particularly relating to long phonon relaxation times. Elemental silicon has been extensively studied experimentally: a collection of properties is given by Hull (1999).

The first investigation (chapter 4) aims at developing an equation of state model for temperature dependent, anisotropic non-linear hyperelasticity. A method is presented for finding deformed states of a material on the same isentrope as a given starting configuration. The energies and stresses of a number of elastic deformations are sampled from DFT molecular dynamics using this method, over a given range of the seven-dimensional space of deformation and potential temperature. The complete energy surface within this range can then be reliably reconstructed using the technique of Gaussian process regression. This is a machine learning technique that has particular merit here due to its

ability to reconstruct a smooth surface without over-fitting. An equation of state model is then constructed for DFT silicon, and demonstrated within a finite-volume continuum elasticity simulation for several problems of interest involving shock waves, namely a shock interacting with a thermal gradient, the transmission and reflection of a shock due to a change in crystal orientation, and two impact test problems.

The second investigation (chapter 5) is concerned with the computation of properties of shock waves. Shock waves are used extensively to study matter at conditions of extreme pressure and temperature, and have been used to obtain some of the highest laboratory-attained pressures. They are useful for equation of state determination and are important dynamic phenomena in their own right, arising in aerodynamics, (Dolling, 2001) reactive flow (Dlott, 2011) and high-speed impact (Asay and Shahinpoor, 1993; Duvall and Graham, 1977).

We describe a simple annealing procedure to obtain the Hugoniot locus (states accessible by a shock wave) for a given material in a computationally efficient manner. We apply this method to determine the Hugoniot locus in bulk silicon from ab initio molecular dynamics with forces from density-functional theory, up to 70 GPa. In addition, we compute the Hugoniot locus for several empirical interatomic potentials modelling silicon. This lets us perform direct non-equilibrium molecular dynamics simulations of shock waves (which typically demand system sizes larger than we can afford with density functional theory), and compare with our indirect method in several circumstances arising in these cases. We also present a direct ab initio molecular dynamics simulation of an elastic shock wave in silicon. This gives a single, low-pressure point on the Hugoniot locus of DFT silicon, but also a profile of the shock wave, which is only available by a direct simulation, and the first obtained ab initio, to our knowledge.

The third and final investigation is into the computation of thermal conductivity from atomistic simulations. There are a number of established methods for thermal conductivity calculation, with varying computational requirements. For low temperatures, a solution of the phonon Boltzmann equation typically provides an accurate solution, with inputs being the second- and third-order force constants. These quantities can be computed from density functional theory (or density functional perturbation theory) affordably. In order to measure high-temperature conductivity, or the thermal resistance due to defects or nanostructures in a sample, a more direct method must be used, which precludes DFT due to the requirements of simulation time and system size for these other methods. Empirical interatomic potentials are much cheaper, and for these a thermal conductivity value can be obtained by a direct method. However, the thermal conductivity has proven challenging to reproduce, and a survey of potentials

for silicon reveals a spread of thermal conductivity values for several commonly-applied silicon potentials of at least a factor of four (at the same temperature).

Here, we produce a number of model interatomic potentials for silicon, using the Gaussian Approximation Potential (GAP) approach, as originally described in the work of Bartók et al. (2010, 2013b). This is a non-parametric technique for potential development, using Gaussian process regression to infer the energy of a given configuration of atoms, the means of representing and comparing atomic environments being a key component of the method. While their computational requirements are more than typical parameterized empirical potentials, they are several orders of magnitude less than DFT calculations.

The potential itself depends on a database of training configurations, and the potential can be improved systematically by increasing the size of the database, or including configurations close to those explored by a simulation of the system. We make use of this by producing several potentials from databases that are increasingly large and broad (in terms of the types of configuration represented), and computing the thermal conductivity from these. Each of these potentials reproduces the lattice parameter, energy minimum and elastic constants from DFT silicon, and is systematically improved such that the best of them reproduce the DFT value of phonon-Boltzmann conductivity to within a few percent.

2. Atomistic Modelling of Materials

The purpose of this chapter is to provide the theoretical background underpinning the methods used in the rest of the work for the computational modelling of condensed matter with atoms interacting through a Hamiltonian. The starting point is the quantum many-body problem, the solution of which gives the precise description of a system of atoms. Since this is intractable for any more than a few atoms, approximations must be used in practice. The first such described in section 2.1 is Density Functional Theory, as developed by Kohn and Sham (1965). We then discuss empirical potentials (section 2.2), before introducing Gaussian Approximation Potentials in section 2.3, a data-driven approach to potential development.

2.1. Density Functional Theory

This section is an overview of the method of Kohn–Sham Density Functional Theory (DFT), and some practicalities for using it for performing numerical simulation.

A good general reference is the textbook on electronic structure and numerical methods by Martin (2004). The review article of Jones and Gunnarsson (1989) gives a survey of early developments in density functional theory. The reader is referred to these sources for more detail on the methods described here.

2.1.1. The Quantum many-body problem

The evolution in time t of a physical system described by a wavefunction Ψ is governed by the Schrödinger equation,

$$i\frac{\partial \Psi}{\partial t} = \hat{H}\Psi, \tag{2.1}$$

where Ψ is the wavefunction of the system and $\hat{H}$ is the Hamiltonian operator. We work in atomic (or Hartree) units, where the reduced Planck constant is unity.

The stationary states have a definite energy E, which is an eigenvalue of the Hamiltonian:

$$\hat{H}\Psi = E\Psi. \tag{2.2}$$

The wave function represents our complete knowledge of a quantum system. Apart from a handful of special choices of Hamiltonian, it cannot be solved analytically.

For a system of N particles, the wave function can be expressed in the basis of eigenfunctions of the position operator of each particle (and perhaps other quantities, such as the particles' spins)

$$\Psi = \Psi(x_1, \ldots, x_N, s_1, \ldots, s_N). \tag{2.3}$$

In the general case, to represent this function approximately requires an amount of storage space that is exponential in the number of particles, before even considering how we might compute it.

More explicitly, suppose we have a single-particle quantum system and an M element basis $\{|1\rangle, \ldots, |M\rangle\}$ of state space, $\mathscr{H}$. The state of a system of N such particles exists in the tensor product space $\underbrace{\mathscr{H} \otimes \cdots \otimes \mathscr{H}}_{N \text{ terms}}$, and a general state may be expressed as the sum

$$\sum_{i_1,\ldots,i_N} c_{i_1 \ldots i_N} |i_1\rangle \otimes \cdots \otimes |i_N\rangle \quad 1 < i_1, \ldots, i_N \leq M. \tag{2.4}$$

To represent this state, we must keep track of the M^N quantities $c_{i_1,\ldots,i_N}$, an amount exponential in the number of particles.

One approach to reduce the complexity of this problem is to attempt to approximate the full expression as a low-rank tensor. Grasedyck et al. (2013) give some ideas on low-rank tensor approximation, and references to quantum mechanical applications.

Assuming that the particles are independent means that the Hamiltonian may be written as a sum $\hat{H} = \sum_{i=1}^{M} \hat{H}_i$, and so the eigenfunctions of $\hat{H}$ can always be separated into product states as

$$\bigotimes_{j=1}^{N} \left(\sum_{i=1}^{M} c_{ij} |i\rangle \right), \tag{2.5}$$

requiring merely $M \times N$ quantities, c_{ij}, linear in the number of particles. This would be suitable for Bosons.

Hartree–Fock approximates the general many-Fermion state by a single state of the form of eq. (2.4), choosing $c = \epsilon_{i_1 \ldots i_N} / \sqrt{N!}$ where ϵ is the N-dimensional Levi-Civita symbol. The expression is known as a *Slater determinant*. This is motivated by the fact that a wavefunction representing identical Fermions is antisymmetric under particle exchange,

that is, $\Psi(...,\boldsymbol{x}_i,...,\boldsymbol{x}_j...) = -\Psi(...,\boldsymbol{x}_j,...,\boldsymbol{x}_i...)$, which is also a defining property of the determinant. The Slater determinant is not the most general Fermionic wavefunction, however, and it is not usually possible to express the ground state as one either.

A system of N_e electrons and N_n atomic nuclei interacting with one-another in the non-relativistic limit has Hamiltonian (expressed in atomic units, where the electronic mass and charge are unity)

$$\hat{H} = -\frac{1}{2}\sum_i \nabla^2_{r_i} - \frac{1}{2}\sum_i \frac{1}{m_i}\nabla^2_{R_i} + \frac{1}{2}\sum_i^{N_e}\sum_{j\neq i}^{N_e}\frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} - \sum_i^{N_e}\sum_j^{N_n}\frac{Z_j}{|\boldsymbol{r}_i - \boldsymbol{R}_j|} + \frac{1}{2}\sum_i^{N_n}\sum_j^{N_n}\frac{Z_i Z_j}{|\boldsymbol{R}_i - \boldsymbol{R}_j|}. \quad (2.6)$$

An important simplifying assumption is provided by the approximation of Born and Oppenheimer (1927), which states that a molecular wavefunction can be written as a product of independent electronic and nuclear terms. The justification for this is that the electrons are much lighter than the nuclei (by three to five orders of magnitude), and so in the full coupled dynamics, the dynamics of the electrons are much more rapid. The assumption allows for the electrons to relax *instantaneously* to the positions of the nuclei. Under this assumption, the Hamiltonian for the N_e electrons moving in the potential of the nuclei is

$$\hat{H}_e = -\frac{1}{2}\sum_i \nabla^2_r + \frac{1}{2}\sum_i^{N_e}\sum_{j\neq i}^{N_e}\frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} - \sum_i^{N_e}\sum_j^{N_n}\frac{Z_j}{|\boldsymbol{r}_i - \boldsymbol{R}_j|}. \quad (2.7)$$

The nuclei enter only through their positions, $\boldsymbol{R}_i$, and so this is also known as the *frozen nuclei approximation*. The motions of the electrons can in this way be solved separately from those of the nuclei, which are often treated classically (for example, by molecular dynamics).It will prove useful to write the Hamiltonian as

$$\hat{H}_e = \hat{F} + \hat{V}, \quad (2.8)$$

where

$$\hat{F} = -\frac{1}{2}\sum_i \nabla^2_r + \frac{1}{2}\sum_i^{N_e}\sum_{j\neq i}^{N_e}\frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} \quad (2.9)$$

depends only on the positions of the electrons, and

$$\hat{V} = -\sum_{i}^{N_e}\sum_{j}^{N_n}\frac{Z_j}{|\boldsymbol{r}_i - \boldsymbol{R}_j|} \tag{2.10}$$

is a term representing the interaction of the electrons with the nuclei.

The **exchange energy** E_x, can be defined as the difference between the energy of the independent particle energy—the minimum energy obtainable from a state with the form of eq. (2.5)—and the Hartree–Fock energy. The **correlation energy** E_c, is defined as the difference between the exact energy and the Hartree–Fock energy. Sometimes it is useful to combine these into an **exchange–correlation energy**, $E_{xc} = E_x + E_c$. Kohn–Sham Density Functional Theory defines its own approximation to this combined term.

2.1.2. The variational approach

The energy eigenvalues of the stationary Schrödinger equation are stationary points of the energy, seen as a functional of the wavefunction, $E[\Psi] = \langle\Psi|\hat{H}|\Psi\rangle$.

Consider the stationary points of the energy functional, under the constraint of a normalized wavefunction $\langle\Psi|\Psi\rangle = 1$, imposed by a Lagrange multiplier E. The first variation operator δ vanishes:

$$\delta\left(\langle\Psi|\hat{H}|\Psi\rangle - E\langle\Psi|\Psi\rangle\right) = \langle\delta\Psi|\hat{H}|\Psi\rangle + \langle\Psi|\hat{H}|\delta\Psi\rangle - E\left(\langle\delta\Psi|\Psi\rangle + \langle\Psi|\delta\Psi\rangle\right) \tag{2.11}$$

$$= \langle\delta\Psi|\left[\hat{H}|\Psi\rangle - E|\Psi\rangle\right] + \left(\langle\delta\Psi|\left[\hat{H}|\Psi\rangle - E|\Psi\rangle\right]\right)^* = 0 \tag{2.12}$$

This is satisfied exactly when $|\Psi\rangle$ is an eigenstate of the Hamiltonian with eigenvalue E.

The Hellmann–Feynman theorem can be seen as a consequence of the variational principle.

Theorem 1 (Hellmann–Feynman)**.** *Consider a Hamiltonian $\hat{H}_\lambda$, that depends on a continuous parameter λ, and suppose that for a given λ it has a stationary wavefunction $|\Psi_\lambda\rangle$ with energy $E_\lambda := \langle\Psi_\lambda|\hat{H}_\lambda|\Psi_\lambda\rangle$. The derivative with respect to λ is given by*

$$\frac{\mathrm{d}E_\lambda}{\mathrm{d}\lambda} = \left\langle\frac{\mathrm{d}\hat{H}_\lambda}{\mathrm{d}\lambda}\right\rangle_{\Psi_\lambda}. \tag{2.13}$$

Proof. E is a function of λ (in some parameter space) and Ψ (a functional argument). The energy eigenvalue is $E_\lambda = E(\lambda, \Psi_\lambda)$.

We choose some particular $\lambda = \lambda_0$ at which to evaluate the derivative. By the chain rule,

$$\left.\frac{\mathrm{d}E_\lambda}{\mathrm{d}\lambda}\right|_{\lambda=\lambda_0} = \frac{\partial E}{\partial \lambda}(\lambda_0, \Psi_{\lambda_0}) + \int \mathrm{d}x \, \frac{\delta E}{\delta \Psi(x)}(\lambda_0, \Psi_{\lambda_0}) \left.\frac{\mathrm{d}\Psi_\lambda(x)}{\mathrm{d}\lambda}\right|_{\lambda=\lambda_0}. \tag{2.14}$$

The variational principle on the parametrized energy eigenvalues is

$$\frac{\delta E}{\delta \Psi(x)}(\lambda, \Psi_\lambda) = 0 \tag{2.15}$$

for any λ, and so the integral on the right hand side of eq. (2.14) vanishes.

$$\frac{\partial E}{\partial \lambda} = \frac{\partial}{\partial \lambda} \langle \Psi | \hat{H}_\lambda | \Psi \rangle = \left\langle \Psi \left| \frac{\mathrm{d}\hat{H}_\lambda}{\mathrm{d}\lambda} \right| \Psi \right\rangle \tag{2.16}$$

since here λ is a formal argument to E, and the other formal argument Ψ has no explicit λ dependence. ■

A similar argument can be used with any variational method (such as DFT) to show that it satisfies a similar relation to eq. (2.13), but with the energy eigenvalue (typically the ground state) arising from the Hamiltonian implied by the method, even if it is not a true solution of the Schrödinger equation.

Atomic forces, stresses and similar quantities that are derivatives of energy can be computed using this result.

2.1.3. Density Functional Theory

The defining notion of a density functional theory is that the total energy of the system is expressed as a functional of the electron density (which itself is a function of only a single position vector) rather than the wavefunction (a function of all N atomic position vectors).

Electron density is defined from the many-particle wavefunction, expressed in terms of the N electronic positions, as:

$$n(\boldsymbol{x}) = N \int \mathrm{d}\boldsymbol{x}_2 \ldots \int \mathrm{d}\boldsymbol{x}_N \, |\Psi(\boldsymbol{x}, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_N)|^2. \tag{2.17}$$

Suppose N electrons interact with each other and an external potential. We have the following two theorems due to Hohenberg and Kohn (1964):

Theorem 2 (Hohenberg and Kohn). *The electron density $n(\boldsymbol{x})$ determines this potential uniquely up to an additive constant.*

Proof. Consider two potentials that have some spatially-varying difference, that give rise to the same electron density, $n(\boldsymbol{x})$. Call these $v(\boldsymbol{x})$ and $v'(\boldsymbol{x})$, with ground-state energies E and E', Hamiltonians $\hat{H} = \hat{F} + \hat{V}$ and $\hat{H}' = \hat{F} + \hat{V}'$, and with $\hat{V}$ and $\hat{V}'$ the contributions to the Hamiltonian arising from the potentials, and $\hat{F}$ defined as eq. (2.9). The wavefunctions must be different, since they satisfy different Schrödinger equations. The fact that the ground state minimizes the energy gives:

$$E' = \langle \Psi' | \hat{H}' | \Psi' \rangle < \langle \Psi | \hat{H}' | \Psi \rangle = \langle \Psi | \hat{H} + \hat{V}' - \hat{V} | \Psi \rangle = E + \int \mathrm{d}\boldsymbol{x}\, n(\boldsymbol{x})[v(\boldsymbol{x}) - v'(\boldsymbol{x})]. \tag{2.18}$$

The second equality in the above equation arises because the Hamiltonians differ only by the contribution from the external potential. By the symmetry of the situation, we can exchange primed with unprimed quantities in this equation, but use the same (unprimed) density, which is assumed the same for both wavefunctions:

$$E < E' + \int \mathrm{d}\boldsymbol{x}\, n(\boldsymbol{x})[v'(\boldsymbol{x}) - v(\boldsymbol{x})]. \tag{2.19}$$

Adding the resulting inequalities eqs. (2.18) and (2.19) gives the contradiction $E + E' < E' + E$. We avoid this if there is a different density associated with v'.

The electron density therefore determines both the number of electrons and the external potential, which gives the full Hamiltonian. ■

Theorem 3 (Hohenberg and Kohn). *There is a functional of the density, $F[n]$, for a given external potential, giving the sum of the kinetic energy and interaction energy of the electrons. The density minimizing the functional*

$$E[n] = F[n] + \int \mathrm{d}\boldsymbol{x}\, n(\boldsymbol{x}) v(\boldsymbol{x}) \tag{2.20}$$

corresponds to the electronic ground-state electron density.

Proof. This is an alternative proof due to Levy (1979).[1] Define

$$F[n] = \min \langle \Psi | \hat{F} | \Psi \rangle = \langle \Psi_{\min} | \hat{F} | \Psi_{\min} \rangle, \tag{2.21}$$

[1] Hohenberg and Kohn originally proved that there exists an $F[n]$ that gives the kinetic energy and energy due to the electron-electron interactions, *but restricted to densities arising as the ground state of some external potential*. There are electron densities that cannot arise this way—see Levy (1979) for details. The argument of Levy is valid for densities arising from any antisymmetric wavefunction through eq. (2.17).

the last equality defining Ψ_{min}. The minimum is over all antisymmetric wavefunctions having a given density n. Then

$$E^{\Psi_{\text{min}}} = \langle \Psi_{\text{min}} | \hat{F} + \hat{V} | \Psi_{\text{min}} \rangle = F[n] + \int \mathrm{d}\boldsymbol{x}\, n(\boldsymbol{x}) v(\boldsymbol{x}). \tag{2.22}$$

For the ground state wavefunction Ψ_{GS} and corresponding density,

$$F[n_{GS}] = \langle \Psi_{\text{min}}^{n_{\text{GS}}} | \hat{F} | \Psi_{\text{min}}^{n_{\text{GS}}} \rangle \leq \langle \Psi_{GS} | \hat{F} | \Psi_{GS} \rangle, \tag{2.23}$$

by the definition of F. Adding the contribution from the external potential, and noting that this interacts with the system only through the density, we see that

$$\langle \Psi_{\text{min}}^{n_{\text{GS}}} | \hat{H} | \Psi_{\text{min}}^{n_{\text{GS}}} \rangle \leq E_{GS}. \tag{2.24}$$

But the ground state energy E_{GS} is the minimum obtainable, so the above inequality must in fact be an equality, and it follows that minimising E over n is equivalent to finding the ground state energy. ■

The second Hohenberg and Kohn theorem (thm. 3) allows the ground-state energy to be expressed as the minimum of some functional of the density. The first theorem means that any other ground state property is then available, in principle, since the density determines the Hamiltonian to within an arbitrary additive constant. The usefulness of DFT comes from the ability to approximate this functional well, and to find the minimizing density numerically.

2.1.4. Kohn–Sham Density Functional Theory

As before, the total energy functional is split into two terms, $E[n] = F[n] + V[n]$, where V is the external potential (depending on the positions of nuclei or any other external effects) and F is 'universal' in the sense that it is the same for any system of interacting electrons. This part of the functional can be written as the sum of several contributions, to be treated in turn:

$$F[n] = T[n] + E_H[n] + E_{xc}[n]. \tag{2.25}$$

The most straightforward term here is E_H, the **Hartree energy**, representing the Coulomb interaction of the electrons. Expressed in terms of the electron density this is

$$E_H[n] = \frac{1}{2} \iint \mathrm{d}\boldsymbol{x}\mathrm{d}\boldsymbol{x}' \, \frac{n(\boldsymbol{x}) n(\boldsymbol{x}')}{|\boldsymbol{x} - \boldsymbol{x}'|}. \tag{2.26}$$

T is the **kinetic energy**. The earlier Thomas–Fermi approach to DFT took T as an explicit, but approximate, functional of n. The development of Kohn and Sham was to retain the exact expression for the kinetic energy in terms of the wavefunction, which is some (albeit unknown) functional of the density. It is not possible to express the kinetic energy *directly* in terms of the density, but under the assumption of non-interacting single-electron states it is

$$T[n] = -\frac{1}{2}\langle\Psi[n]|\nabla^2|\Psi[n]\rangle = -\frac{1}{2}\sum_i^N \langle\psi_i|\nabla^2|\psi_i\rangle\,, \tag{2.27}$$

where the sum is over single electron wavefunctions. This is accomplished in Kohn–Sham DFT by diagonalizing the Hamiltonian. The meaning of the $|\psi_i\rangle$, and how they are calculated is discussed below.

The final term of eq. (2.25) is E_{xc}, the **exchange-correlation energy**. This term is unknown. It includes the difference between the Hartree energy and the true interaction energy of the electrons, and also the difference between the true kinetic energy and the kinetic energy in the independent electron approximation. If it were known exactly, the Kohn–Sham equations would be equivalent to finding the exact ground state of the many-electron, nonrelativistic Schrödinger equation, and therefore truly an *ab initio* theory.

As it stands, however, since the exchange-correlation functional is unknown, we must rely on approximations. Since these approximations are intended to work for a wide range of systems, DFT is sometimes called an *ab fere initio* ('almost from the beginning') method. It seems unlikely that it is possible to compute this term efficiently (any more so than solving the time-independent Schrödinger equation by some other 'hard' method, such as full configuration-interaction). In any case, this functional must exhibit some exotic features in order to describe the full range of phenomena arising in condensed matter physics. It must be non-local and discontinuous, for example. Schuch and Verstraete (2009) considered the problem of determining E_{xc} from a computational complexity aspect, and showed that if it were computable efficiently some problems considered hard for a quantum computer would have an efficiently-computable solution on a classical computer.

The variational principle applied to the energy (as a functional of the one-electron wavefunctions) states that

$$\frac{\delta\left(E - \sum_{jk}\epsilon_{jk}(\langle\psi_j|\psi_k\rangle - \delta_{jk})\right)}{\delta\psi_i} = 0, \tag{2.28}$$

where ϵ_{ij} are the Lagrange multiplier associated with the constraints on the wavefunctions

$$\langle\psi_i|\psi_j\rangle = \delta_{ij}. \tag{2.29}$$

We can expand this as follows:

$$\frac{\delta T}{\delta\psi_i} + \frac{\delta(E_H[n]+E_{xc}[n])}{\delta n}\frac{\delta n}{\delta\psi_i} - \frac{\delta}{\delta\psi_i}\sum_{jk}\epsilon_{jk}(\langle\psi_j|\psi_k\rangle - \delta_{jk}) = 0 \tag{2.30}$$

$$-\frac{1}{2}\nabla^2\psi_i + \hat{v}_{\text{KS}}\psi_i + \sum_j \epsilon_{ij}\psi_j = 0 \tag{2.31}$$

and through this equation we define the Kohn–Sham *effective potential*, $\hat{v}_{KS}$. Instead of the full (interacting) Hamiltonian, Kohn and Sham (1965) supposed that for any external potential v_{ext}, there is a system of non-interacting electrons in some other potential, that has the same electron density as the interacting system in v_{ext}. The non-interacting system has the single-electron states of eq. (2.27), $|\psi_i\rangle$, and allows us to write

$$\left(-\frac{1}{2}\nabla^2 + v_{KS}(\boldsymbol{x})\right)|\psi_i\rangle = \hat{H}_{KS}|\psi_i\rangle = \epsilon_i|\psi_i\rangle \tag{2.32}$$

where

$$v_{KS}(\boldsymbol{x}) = v_{\text{ext}}(\boldsymbol{x}) + \int \mathrm{d}\boldsymbol{x}'\,\frac{n(\boldsymbol{x}')}{|\boldsymbol{x}-\boldsymbol{x}'|} + v_{xc}[n](\boldsymbol{x}). \tag{2.33}$$

Here, v_{ext} is the potential due to the nuclei and other external effects, and

$$v_{xc}[n](\boldsymbol{x}) = \frac{\delta E_{xc}[n]}{\delta n(\boldsymbol{x})} \tag{2.34}$$

is the exchange-correlation potential. Kohn–Sham DFT is expressible as the eigenvalue problem eq. (2.32), of independent electrons interacting via the effective potential v_{KS}. Since this potential itself depends on the electron density, the whole system (potential and single-electron wavefunctions) must be solved together self-consistently, typically with an iterative method.

The Hamiltonian, $\hat{H}_{KS}$, will typically be represented in some basis (plane waves or localised atomic-orbital-like functions are two usual choices, see section 2.1.6). Finding the N lowest-lying eigenfunctions[2] $|\psi_i\rangle$ expressed in this basis involves diagonalizing the

[2] We have neglected spin-degeneracy—if it is included, we are interested in the $N/2$ lowest-energy states, with two electrons per state.

Hamiltonian matrix. Without exploiting any special structure, the number of operations required for the diagonalization is $O(M^3)$, where M is the number basis functions.

There are several approaches to approximating v_{xc}. Most of them rely on the observation that apart from the long-ranged Coulomb interaction of the electron density via the Hartree term, eq. (2.26), the interaction of the electrons is often quite localized (although there are many important exceptions).

The simplest approach, known as the **local density approximation (LDA)**, is to suppose that

$$v_{xc}[n](\boldsymbol{x}) = v_{xc}^{\mathrm{LDA}}(n(\boldsymbol{x})). \tag{2.35}$$

The most successful approaches rely on analytic expressions for the free electron gas, and define

$$E_{xc}^{\mathrm{LDA}}[n] = \int n(\boldsymbol{x})\epsilon_{xc}^{\mathrm{FEG}}(n(\boldsymbol{x}))\,\mathrm{d}\boldsymbol{x}, \tag{2.36}$$

with

$$E_{x}^{\mathrm{FEG}}[n] = -\frac{3}{4}\left(\frac{3}{\pi}\right)\int n(\boldsymbol{x})^{4/3}\,\mathrm{d}\boldsymbol{x}. \tag{2.37}$$

Even for this greatly simplified case, there is not an exact expression for the electron correlation energy. A number of approximate forms that reproduce certain known limits (e.g. low or high density) exist (Ceperley and Alder, 1980; Perdew and Wang, 1992; Perdew and Zunger, 1981; Vosko et al., 1980).

The **generalised gradient approximation (GGA)** allows for v_{xc} to be some function of both the electron density at a point and its gradient there:

$$v_{xc}[n](\boldsymbol{x}) = v_{xc}^{\mathrm{GGA}}(n(\boldsymbol{x}), \nabla n(\boldsymbol{x})). \tag{2.38}$$

A variety of commonly used ones can be found in: Becke (1988); Lee et al. (1988); Perdew et al. (1992, 1996); Zhang and Yang (1998).

Functionals of either variety are by construction incapable of reproducing long-ranged electron correlations important in some materials (for example, London dispersion or Mott-insulator transitions).

2.1.5. Pseudopotentials

The pseudopotential approach is important for the numerical solution of the Kohn-Sham equations. The key idea behind the pseudopotential approximation is that the core electrons of an atom are tightly bound to the nucleus and are not involved in chemical

bonding, so that the potential due to the nucleus can be replaced with one due to the combined effect of the nucleus and core electrons.

The advantage is twofold. First, it reduces the number of electrons in the system to be solved. Second, the orbitals of the valence electrons are no longer required to be orthogonal to the omitted core electron orbitals. This requirement is responsible for making the valence orbitals highly oscillatory near the nucleus, which is challenging to represent—see fig. 2.2.

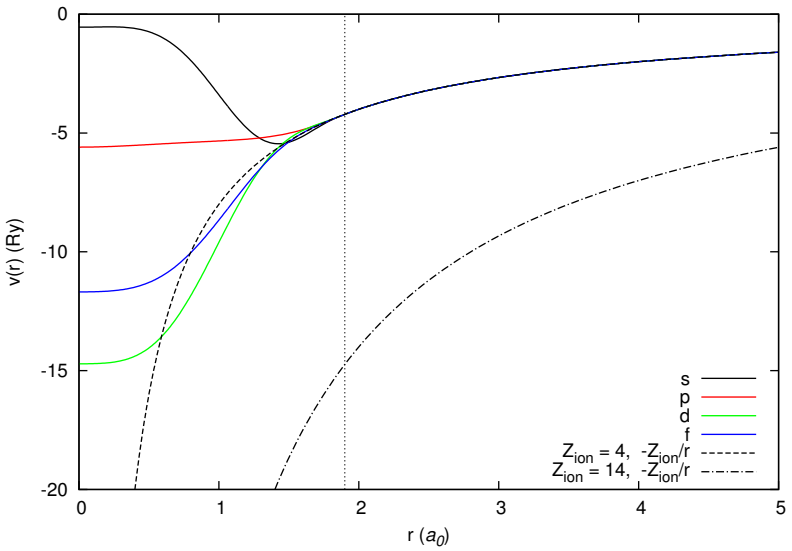

Figure 2.1.: Nuclear potential and pseudopotential for four angular momentum channels. The norm-conserving pseudopotential was generated for LDA silicon according to the Troullier–Martins method. The dashed line is the potential for a point nuclear charge of $+4e$, to which the pseudopotential is equal outside of the cutoff radius (vertical dotted line). The dash-dotted line is the all-electron potential for comparison.

A pseudopotential is generated by determining a cutoff radius, beyond which the wavefunctions of the valence electrons in some reference system agree, and within which the nuclear potential is modified to incorporate the effect of the core electrons. The potential arising this way is non-local and may depend on the angular momentum channel.

The pseudopotential eigenvalues should agree with the all-electron Hamiltonian, in the reference system and in other systems besides (to be transferable). Further requirements for good accuracy and transferability are are given by Hamann et al. (1979).

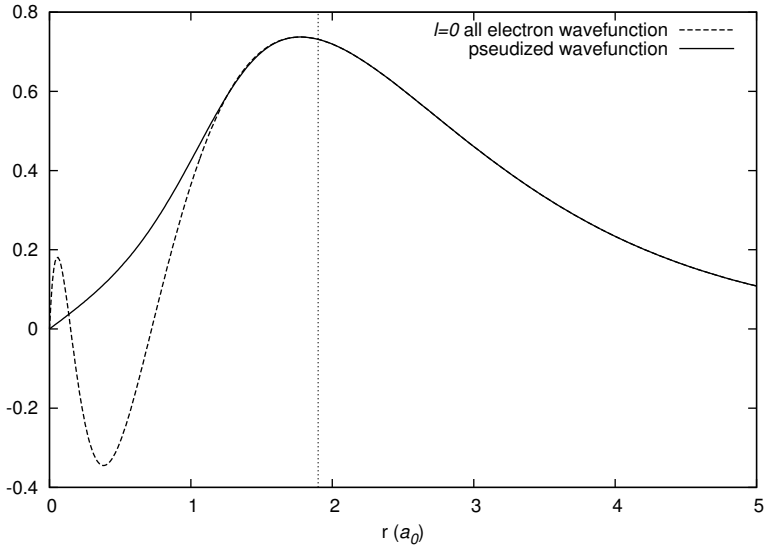

Figure 2.2.: The s channel ($n = 3$, $l = 0$) of the resulting wavefunction. The matching radius is indicated by the vertical dotted line.

One such idea due to Hamann et al. (1979) to aid transferability is to make the pseudopotential *norm-conserving*: the charge contained within the cutoff radius is the same as for the all-electron system.

The original formulation of Phillips and Kleinman (1959) involved computing two-point integrals, the cost of which scale like the square of the number of basis elements. An alternative formulation of Kleinman and Bylander (1982) allows linear scaling.

There are several practical schemes for generating pseudopotentials (e.g. Hamann et al., 1979; Kerker, 1980; Troullier and Martins, 1991; Vanderbilt, 1985).

2.1.6. Choice of basis

Numerically, the solution wavefunctions and density must be represented on a finite basis, rendering eq. (2.32) a matrix equation. The choice of basis should ideally result in a small error in ground state energy for a wide range of systems. This is always a compromise with computational cost. The two ways of reducing the error are to choose the shapes of the basis functions carefully, so that a good approximation can be obtained with relatively few basis functions, or to increase the size of the basis. Having a systematic way of improving the basis, so that it tends to a complete basis (every function is representable in the complete basis, and the error in representing it tends to zero with the size of the finite basis) is a valuable property.

A basis of plane waves is complete, so that a *finite* basis of plane waves can be improved systematically by adding higher frequencies to the basis. The pseudopotential approximation reduces the extent of the potential in Fourier space. This, along with fast numerical methods for plane wave calculations, have contributed to the success of plane-wave-based methods. Despite this, the dense density-matrix and Hamiltonian preclude linear scaling DFT.

An alternative approach is to use a basis of numerical atomic orbitals. Here, the shapes are chosen to represent the electronic configuration about an atom well. Because of this, good convergence can often be obtained with quite a small set of basis functions per atom, perhaps dozens. For example, a double-zeta polarized basis set (the nomenclature is described in the next section) for silicon, which is already well-converged for many purposes, has thirteen orbitals per atom. This keeps the dimensions of the Hamiltonian manageable, improving the speed of diagonalization. Numerical atomic orbitals with finite-support allow for linear-scaling methods to be used. A drawback in using this type of basis is the lack of obvious systematics for convergence.

The converged KS energy is variational in the basis set, which can aid choosing a basis.

Scheme of Junquera et al.

SIESTA can use local basis functions of arbitrary shape, and a basis set can be chosen to suit the requirements of the problem or user. We have used one such scheme of soft-confined spherical harmonics due to Junquera et al. (2001) and implemented in SIESTA as described by Soler et al. (2002). This scheme is summarized briefly below.

Each basis function (centred on atom I) is a product of a radial function and a spherical harmonic:

$$\phi_{Ilmn}(\boldsymbol{r}) = \phi_{Iln}(r_I) Y_{lm}(\hat{\boldsymbol{r}}_I), \tag{2.39}$$

where n labels the radial shape, l and m label the angular momentum, r_I is the position vector directed from atom I, $\boldsymbol{r}_I = \boldsymbol{r} - \boldsymbol{R}_I$, $r_I = |\boldsymbol{r}_I|$ and $\hat{\boldsymbol{r}}_I = \boldsymbol{r}_I / r_I$.

The radial shapes used are the eigenfunctions of a smooth confinement potential with $V \to \infty$ as $r \to r_c$, which ensures that the basis functions have finite support and continuous derivatives. The confinement potential used is:

$$V(r) = \begin{cases} 0 & r \leq r_i \\ V_0 \dfrac{\exp(-\frac{r_c - r_i}{r - r_i})}{r_c - r} & r_i < r < r_c \end{cases} \tag{2.40}$$

where r_i is known as the soft-confinement radius. Each angular momentum channel can use different parameters. Multiple radial shapes per channel can be used. One way of introducing an additional radial function based on another, is according to a split-valence scheme. This permits some additional freedom in radial dependence. Introduce a new radial shape by choosing a matching radius r_m, beyond which the new and original functions are identical, and defining $\phi_l(r) = r^l(a - b\,r^2)$ within the matching radius, choosing a and b by continuity of value and first derivative of the basis function at r_m. The procedure can be repeated as many times as desired. In the language of quantum chemistry, this gives 'multiple-ζ' bases. A single radial shape is single-ζ (SZ), an additional one double-ζ (DZ) and so on.

Additional angular freedom is afforded by *polarization orbitals*, which take an orbital of the highest angular momentum l in the basis, and generate a new one by applying a small electric field along the z direction. The polarized shape is determined from first-order perturbation theory. It can be shown that this contains an additional angular component of $l+1$ (and no higher), and has a radial shape obtained by solving an ordinary differential equation. Numerically, the radial basis shapes are evaluated as cubic splines interpolated from a fine radial grid. The full procedure is described by Soler et al. (2002).

2.1.7. Discretization

The total Hamiltonian (eq. (2.7)) can be rewritten (Soler et al., 2002) with the aim of eliminating long-range interactions, as

$$\hat{H} = \hat{T} + \sum_I \hat{V}^{KB} + \sum_I \hat{V}^{NA}(\boldsymbol{r}) + \delta V^H(\boldsymbol{r}) + V^{xc}(\boldsymbol{r}). \tag{2.41}$$

Here, $\hat{T}$ is the kinetic operator, $\hat{V}^{KB}$ is the nonlocal part of the pseudopotential, $\hat{V}_I^{NA}$ is the 'neutral atom' potential obtained by screening the local part of the pseudopotential by an arbitrary (local) valence charge distribution, which might be obtained by filling the first ζ atomic orbitals. The motivation for this is that this potential is exactly zero outside the range of the valence charge distribution. The term δV^H is the potential resulting from the assumed screening charge distribution and actual charge distribution. It integrates to zero and is typically small. The final term is the exchange-correlation. The matrix elements of the Hamiltonian, $\langle \phi_{Ilmn} | \hat{H} | \phi_{I'l'm'n'} \rangle$, are computed as follows. The kinetic and nonlocal pseudopotential terms depend only on the distance between atomic centres, and involve integrals over reciprocal space that can be precomputed on a fine grid. Once again, the details of the calculation are found in Soler et al. (2002). The fineness of

a regular k-grid is specified as the wave of highest wavenumber that can be represented on this grid: a wavenumber in reciprocal space is a real-space length, so the cutoff is given as a length, and this can be used to compare the quality of grids between different lattices.

The choice of grid for sampling over reciprocal space is addressed by Monkhorst and Pack (1976), who considered k-grids with their origin displaced with respect to the reciprocal-space lattice. Such a grid with n points along one dimension is able to compute the exact integral of functions which have up to the first n Fourier coefficients non-zero. In addition, for grids of even dimension, the number of points in the irreducible part of the Brillouin zone is smaller than a grid of the same dimension containing the origin. Calculation of the integrand need only be done at the points of the irreducible region, and the remaining points calculated by symmetry, reducing the computational cost of the calculation. Sometimes a good result can be obtained with a small number of k-points, or even a single k-point, if chosen carefully: the mean-value point of Baldereschi (1973) is such a choice. See fig. 2.3.

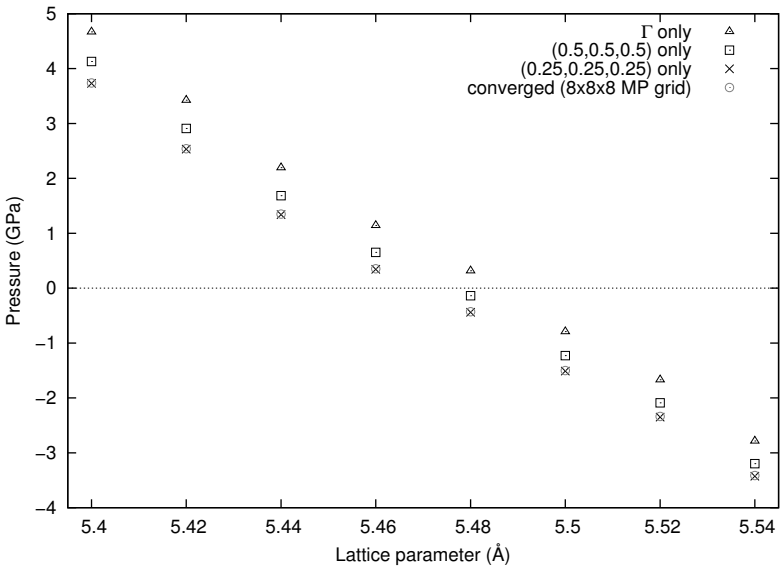

Figure 2.3.: The effect of the choice of k-grid on the lattice-parameter dependence of pressure in bulk silicon. Shown are several calculations of bulk silicon, using several special points around the Brillouin zone, and a highly converged grid. The point $(0.25, 0.25, 0.25)$ is the Baldereschi point.

The remaining matrix elements are evaluated on a three dimensional grid in real space, the fineness of which is measured by the **energy cutoff**: the highest energy plane-wave that can be resolved by the grid. The δV^H term is found by solving Poisson's equation on the grid, using either a fast Fourier transform or a multigrid solver. Once the potential

has been tabulated on the grid, the term $V(\boldsymbol{r})\phi^*_{Ilmn}(\boldsymbol{r})\phi_{I'l'm'n'}(\boldsymbol{r})\Delta r^3$ is accumulated into the matrix element for the Hamiltonian.

The issues of real- and reciprocal-space mesh generation is considered by Moreno and Soler (1992).

2.2. Empirical Interatomic Potentials

An alternative approach is to leave out the electrons from the calculation, and instead fit the total energy to an analytic expression, depending on pairs or otherwise small collections of neighbouring atoms. Once this is done, the resulting problem is greatly simplified, since the energy and forces are now evaluations of a simple analytic expression instead of the solution to a variational problem for the ground state of the electron density.

The clear advantage is that much larger systems can be considered, and when computing molecular dynamics trajectories, far longer times can be simulated. The drawback is that the models are less predictive and transferable, since a simple functional form forces strong assumptions to be made about the potential energy surface. Interatomic potentials are often fit to data for some specific systems of interest and break down for systems in other configurations. They can still be used to simulate many physical effects to qualitative accuracy.

There have been many empirical potentials developed for specific domains. The empirical potentials used in this work are described in detail in appendix C.

2.3. Gaussian Approximation Potentials

2.3.1. Introduction

Analytic interatomic potentials are computationally cheap to evaluate, while 'quantum-chemistry' methods, with their explicit treatment of electrons, are accurate, but prohibitively expensive for systems containing more than at most a few thousand atoms. A favourable feature of treating the electrons explicitly is that this tends to result in a more transferable method: it can be used to make predictions of systems different to the data (if there are any at all) used to parameterize it.

The development of Gaussian Approximation Potentials (GAP) was motivated by two considerations. First, that there exists a gap in accuracy between traditional interatomic potentials and models treating the electrons explicitly. Second, that it is possible to use an

accurate, quantum-mechanical method to obtain data from many small systems of atoms, that can then be fit by an analytic functional form to give an interatomic potential that is cheap to evaluate. The functional forms imposed by traditional interatomic potentials are often too restrictive to represent the energy landscape realistically.

Generating a useful potential from this process relies on being able to apply knowledge gained from simulating small systems over short timescales to the problem of describing much larger systems for longer times. This can be seen as a consequence of locality. In many cases, despite the 'global' nature of the Schrödinger equation, with each particle interacting with every other, long-ranged effects are unusual, and the behaviour of electrons (and therefore the atoms that interact with them) is determined by their immediate environment.

The GAP approach is driven by the wealth of data available from DFT, and other first principles simulation methods. By choosing a suitable way of representing local atomic environments, these can be used directly as training inputs for a potential based on Gaussian-process regression.

The potentials were first described in Bartók et al. (2010). A subsequent paper (Bartók et al., 2013b) reviewed ways of representing atomic environments, including bispectra and bond-order parameters, and introduced the Smooth Overlap of Atomic Positions (SOAP) method. Section 2.3.3 provides more detail.

The work cited so far is focused on developing improved interatomic potentials. The general approach of GAP can be applied to improve the accuracy of any one method based on a more accurate but more expensive method. Bartók et al. (2013a) applied a GAP method as a correction for DFT water, with the accurate calculations provided by coupled cluster calculations.

A related approach is that of Behler and Parrinello (2007), who used neural networks trained on DFT energies, with a local description of the atomic environment as input, based on radial Gaussians and cosines of bond angles.

2.3.2. Representing atomic environments

When fitting an interatomic potential that depends only on nuclear positions, we must choose a set of descriptors for the atomic configuration. For example, in a central-force pair-potential, the descriptors are the pairwise distances between the atoms, $r_{ij} = |\boldsymbol{x}_i - \boldsymbol{x}_j|, 1 \leq i, j \leq N$. For this kind of potential we have the further restriction that the total energy is the sum of individual pairwise bonding energies (a term in the sum depending on both $r_{1,2}$ and $r_{3,4}$, for example, is not permitted, although it may be for a more general potential). A pairwise term involving a given atom m depends only on

r_{mj}, $1 \leq j \leq N$ (excluding r_{ij}, $i \neq m$). We divide each pairwise contribution to the energy between the two involved atoms somehow, perhaps equally, although the choice is arbitrary.

The descriptors r_{mj}, $1 \leq j \leq N$ are for the **local** environment, or atomic neighbourhood, of atom m, and after choosing a partitioning of atomic energies between these environments, we may write:

$$E = \sum_{m=1}^{N} E_m(r_{m1}, \ldots, r_{mN}). \tag{2.42}$$

We could impose a further restriction, that the local environment consists of only nearby atoms:

$$\frac{\partial E_m}{\partial r_{mi}} = 0, \; r_{mi} > r_{\text{cut}}. \tag{2.43}$$

We are interested in choosing descriptors for atomic neighbourhoods on which to base an interatomic potential. This is based on the observation that it is often possible to split the energy of a configuration up into local contributions. This is computationally useful because for a large system it leads to a tractable number of short-ranged descriptors, preferably linear in the system size.

A **complete** set of descriptors suffices to uniquely describe an atomic environment. That is, if the set of descriptors is not complete, there are some genuinely different configurations (perhaps with different energies) that have the same descriptors associated with them. For example, the pairwise distances between atoms do not comprise a complete set: Bartók et al. (2013b) give as an example two tetrahedra with the same edge lengths but different face angles.

It is clear that the atomic coordinates themselves provide a complete description of the configuration, but this description suffers from being **overcomplete**: there are sets of coordinates that describe equivalent atomic configurations. Equivalent configurations are invariants of the symmetries of the potential energy, since the potential energy surface completely determines the dynamics of the problem. That is, if for all collections of atomic positions $(\boldsymbol{x}_i)_{i=1}^N$ we have that

$$E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N) = E(S(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N)), \tag{2.44}$$

then $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N)$ and $S(\boldsymbol{x}, \ldots, \boldsymbol{x}_N)$ are equivalent configurations. The set of such transformations consists of uniform translations, (improper) rotations, spatial inversion and permutations of the coordinates. While this is true for the total energy as a function of

all the atomic positions, the invariants should also hold for every local environment, for any reasonable definition of local energy.

It is desirable to avoid overcomplete descriptors because the fitting process itself does not enforce these invariants, leading to spurious minima, as well as extra computational work due to the increase in the size of the training set required by the inclusion of many configurations that are otherwise known to be equal in energy. The dimensionality of the domain may also be larger.

2.3.3. The SOAP descriptor and kernel

This approach provides a measure of the similarity between two atomic environments given their positions, through a kernel. That is, instead of some descriptor values representing an atomic position, it allows us to determine whether two configurations are similar, and by how much. This is a different approach to first computing an intermediate descriptor to pass as an argument to a potential energy fitting function, and exploits the use of Gaussian process regression, where it can be used as the covariance kernel for potential energy.

First, define the atomic local neighbour density as

$$\rho(\boldsymbol{x}) = \sum_{i}^{N} \delta(\boldsymbol{x} - \boldsymbol{x}_i). \tag{2.45}$$

A measure of the similarity of two atomic environments ρ and ρ' is

$$S[\rho, \rho'] = \int \mathrm{d}\boldsymbol{x}\, \rho(\boldsymbol{x})\rho'(\boldsymbol{x}). \tag{2.46}$$

It is straightforward to see that this is invariant under permutations of atoms, but not rotations. Integrating over all possible rotations leads to

$$k[\rho, \rho'] = \int \mathrm{d}\hat{R} \left(\int \mathrm{d}\boldsymbol{x}\, \rho(\boldsymbol{x})\rho'(\hat{R}\boldsymbol{x}) \right)^{n}. \tag{2.47}$$

The power $n \geq 2$ is to preserve the angular information in the kernel. Taking $n = 1$ would allow the order of the integrals to be exchanged, giving

$$k[\rho, \rho'] = \int \mathrm{d}x\, \rho(\boldsymbol{x}) \int \mathrm{d}\hat{R}\, \rho'(\hat{R}\boldsymbol{x}). \tag{2.48}$$

The inner integral discards all of the angular information available from ρ'.

This kernel involves Dirac delta functions, and is therefore not suitable for numerical use. It is also not smooth as one atomic environment is varied. To correct these deficiencies, we can instead define the local atomic density in terms of Gaussians centred on each atom, as

$$\rho(\boldsymbol{x}) = \sum_{i}^{N} \exp\left(-\alpha|\boldsymbol{x} - \boldsymbol{x}_i|^2\right), \tag{2.49}$$

where α is chosen so that the Gaussians have a width of some fraction of the atomic spacing.

First, expand this in a basis of spherical harmonics and orthonormal radial functions, with coefficients c_{nlm},

$$\rho(\boldsymbol{x}) = \sum_{n,l,m} c_{nlm} g_n(x) Y_{lm}(\hat{\boldsymbol{x}}), \tag{2.50}$$

and define the power spectrum coefficients by

$$p_{nn'l} := \sum_{m} c_{nlm} c^*_{n'lm}. \tag{2.51}$$

The integral eq. (2.47) is evaluated by Bartók et al. (2013b) for $n = 2$ as

$$k[\rho, \rho'] = \sum_{n,n',l} p_{nn'l} p'_{nn'l}, \tag{2.52}$$

and converges to the correct value rapidly with the size of the basis.

The kernel is normalized and can be raised to a 'sensitivity' parameter ζ, that sharpens the overlap integral, making similar atomic environments relatively more significant to the prediction. The SOAP kernel is then

$$K^{\text{SOAP}}[\rho, \rho'] = \left(\frac{k[\rho, \rho']}{\sqrt{k[\rho, \rho] k[\rho', \rho']}}\right)^{\zeta}. \tag{2.53}$$

We choose a cutoff radius to make the environment short-ranged. The Gaussians can be weighted with a smooth transition function $f_{\text{cut}}(\boldsymbol{x}_i)$ so that atoms enter and leave the local environment without a discontinuity in the kernel, and so that nearby atoms are given more importance when computing the similarity between atomic environments about a point. The per-atom density is then computed as

$$\rho(\boldsymbol{x}) = \sum_{i}^{N} f_{\text{cut}}(\boldsymbol{x}_i) \exp\left(-\alpha|\boldsymbol{x} - \boldsymbol{x}_i|^2\right), \tag{2.54}$$

2.3.4. Multiple atomic species

We have so far assumed that all of the atoms are identical. The descriptor can also contain information about the species of the atom (atomic number), although we do not need this facility for the potentials considered later.

The idea is to modify the density, so it retains knowledge of the neighbouring species, represented as $|T_i\rangle$:

$$\rho(\boldsymbol{x}) = \sum_j^N f_{\text{cut}}(\boldsymbol{x}_j) \exp\left(-\alpha|\boldsymbol{x}-\boldsymbol{x}_j|^2\right)|T_j\rangle\,. \tag{2.55}$$

The similarity measure is computed as before, via eq. (2.46) and $\langle T_i|T_j\rangle = \delta_{T_iT_j}$. This reduces to the same value if there is a single atomic species. If there are several species present, the Dirac delta function means that environments are similar only when the same species are present at each location. Permitting cross-terms $\langle T_i|T_j\rangle \neq 0$, $i \neq j$, would allow for a measure of similarity between different atomic species.

2.3.5. Gaussian processes

The use of Gaussian processes in regression is discussed for a general situation in appendix B. The SOAP kernel described in the previous section is used as the covariance kernel on the training inputs.

Implementations of Gaussian process regression and GAP are included in the `libAtoms` and `QUIP` software packages (Bartók-Pártay et al., 2006–).

2.3.6. Local energy prediction

The training data consist of total energies, atomic forces and stresses of configuration, typically computed with DFT. Quantum mechanical methods do not provide a convenient local energy, so instead we must train on configurational total energies.

The total energy is the sum of atomic contributions

$$E^{(i)} = \sum_{\alpha=1}^{N_i} E_\alpha^{(i)} = \sum_{\alpha=1}^{N_i} \epsilon(q_\alpha^{(i)}), \tag{2.56}$$

where $\epsilon(q)$ is the energy of a particular atomic environment q, and α ranges over the atoms in a particular configuration labelled with i.

Covariance is bilinear, so the covariance of two total energy observations is

$$C_{ij} = \mathrm{Cov}(E^{(i)}, E^{(j)}) = \mathrm{Cov}\left(\sum_{\alpha=1}^{N_i} \epsilon(q_\alpha^{(i)}), \sum_{\beta=1}^{N_j} \epsilon(q_\beta^{(j)})\right) = \sum_{\beta=1}^{N_j}\sum_{\alpha=1}^{N_i} \mathrm{Cov}\left(\epsilon(q_\alpha^{(i)}), \epsilon(q_\beta^{(j)})\right). \tag{2.57}$$

That is, an element of the covariance matrix between two total energy observations can be computed as a sum over the covariance kernel (e.g., the SOAP kernel) computed for the environments of each pair of atoms present in the configurations.

Predictions are made using the expression for the mean of a Gaussian process, eq. (B.12) from appendix B:

$$\hat{\epsilon}(q') = \boldsymbol{k}^\mathsf{T} \boldsymbol{C}^{-1} \boldsymbol{E}, \tag{2.58}$$

where q' is the new environment for which a predicted energy is sought, $\boldsymbol{C}$ and $\boldsymbol{E}$ are the covariance matrix and total energy training observations as in eq. (2.57), and the ith component of $\boldsymbol{k}$ is $k_i = \mathrm{Cov}(E^{(i)}, \epsilon(q'))$. The predicted mean from a reduced-rank process, eq. (B.19), can be calculated in a similar way.

2.3.7. Atomic forces and the virial stress

The force on an atom is the derivative of total energy with respect to position. The force on an atom is by assumption only due to atoms in its local environment,

$$\boldsymbol{f}_\beta = -\sum_{\alpha=1}^{N} \frac{\partial E_\alpha}{\partial \boldsymbol{x}_\beta}. \tag{2.59}$$

The covariances between forces and local energies can be computed as derivatives of the energy covariance function, via eqs. (B.7) and (B.8).

An observed value of a component of the virial stress (times volume) is a linear combination of forces

$$V\sigma_{ij} = \sum_{\alpha=1}^{N} x_\alpha^i f_\alpha^j, \tag{2.60}$$

with spatial indices $1 \leq i, j \leq 3$, and α an index over the atoms in a configuration. The required covariance matrix elements once again follow from the bilinearity of covariance.

The derivative covariances for the SOAP kernel are quite complicated and are not given here. They can be found readily with computer algebra software.

2.4. Molecular Dynamics

The idea of a molecular dynamics (MD) method is to integrate Newton's equations of motion for the motion of atoms (or larger structural units of several atoms in the case of coarse grained methods—a recent review is given by Saunders and Voth (2013)), calculating the interatomic forces and thereby gaining knowledge about a system, either in a statistical sense (using an average over time or several replicas of the system to measure thermodynamic quantities) or in microscopic detail.

It is a good approximation to treat atomic nuclei as classical particles with the possible exceptions of hydrogen and deuterium (Bunker and Moss, 1977). The interatomic forces may be obtained from a number of techniques that differ greatly in cost and accuracy, including DFT and empirical potentials, discussed in the preceding sections.

Molecular dynamics is a mature field, and an overview of many techniques can be found in Allen and Tildesley (1987) and Frenkel and Smit (2002).

2.4.1. Velocity Verlet

For Hamiltonian systems, including systems of molecules, a numerical integrator is desired to conserve the Hamiltonian over times long compared to the integration timestep.

Trajectories in the interesting regions of phase space often have sensitive dependence on initial conditions. This ensures that any numerical errors due to the integrator will at first grow exponentially, and the accurate long-time prediction of atomic trajectories is impossible. This same feature is responsible for the rapid 'decoupling' of knowledge about the present and future configurations of a system, on which the results of statistical physics depend for justification. Rather than attempting the impossible task of finding accurate trajectories, it is more important that the integrator produces a trajectory that is a representative sample of phase space.

These properties (near-conservation of first integrals, realistic sampling) can be related to how well the numerical integration retains certain geometric features of the underlying differential equations, including time reversibility and symplecticity. See, for example, Hairer et al. (2003); Leimkuhler and Reich (2004).

Equation (2.61) is the velocity Verlet algorithm, which is a symplectic, time-reversible, second-order method that depends only on the atomic forces and no higher derivatives of

energy, which is used here and widely in MD simulation.

$$\begin{aligned} \boldsymbol{x}_{i+1} &= \boldsymbol{x}_i + \boldsymbol{v}_i\,\delta t + \frac{1}{2m}\boldsymbol{F}_i\,\delta t^2 \\ \boldsymbol{v}_{i+1} &= \boldsymbol{v}_i + \frac{1}{2m}(\boldsymbol{F}_i + \boldsymbol{F}_{i+1})\,\delta t \end{aligned} \tag{2.61}$$

$\boldsymbol{F}_i$ is the force on the particle at the ith timestep, which we suppose depends only on the positions of the other particles, and not their velocities. In practice, for each particle, the position and velocity are updated in place, the forces are updated between the position and velocity update step, and the value of the force for the previous timestep retained, to minimise storage requirements.

It may be shown (Hairer et al., 2003) that for velocity Verlet, the Hamiltonian at the nth timestep (H_n) is approximately conserved, and satisfies

$$|H_n - H_0| = C h^2 + C_N h^N t \quad \text{for} \quad 0 \leq t \leq h^{-N}. \tag{2.62}$$

3. Continuum Modelling of Solids

In the previous chapter, we summarized the atomistic approaches that we use, of microscopic particles interacting via a Hamiltonian. In this chapter, we turn to the macroscopic, finite-deformation, continuum-mechanical viewpoint, and give an overview of the background theory and relate quantities between the two viewpoints.

3.1. Basic Notions

In continuum mechanics, a body is made up of **material points**. Each is identified with the average position of a fixed collection of atoms, small enough that the change in shape of the region of space that they occupy is much smaller than the bulk motion of the material, but large enough that local equilibrium thermodynamic quantities may be defined there. That is, on the lengthscales typical of continuum simulation, a material point may be treated as a mathematical point, but on an atomic lengthscale, enough atoms are present that they may be treated according to statistical mechanics. This allows the definition of energy (and similar quantities, such as entropy and temperature) as a function of the material point, which are, strictly, global properties of the system as a whole.

Spatial or **Eulerian** quantities are those defined at fixed locations in space. **Material** or **Lagrangian** quantities are defined at a material point. In keeping with the literature, variables representing an Eulerian quantity are (with some exceptions) written in lower case, and those representing a Lagrangian quantity in upper case.

Let a material point in some reference configuration (typically the undeformed configuration of vanishing stress in solid mechanics) be denoted $\boldsymbol{X}$. The location of a material point depends on its initial location in the reference configuration according to the **motion**: $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{X}, t)$. The local deformation of the material is described by the **deformation gradient** tensor

$$F = \nabla_X \boldsymbol{x}. \tag{3.1}$$

The determinant expresses the volume change of the deformation, so that

$$\rho = \rho_0 / \det \boldsymbol{F}. \tag{3.2}$$

To ensure a bijection between Lagrangian and Eulerian coordinates, we require that F is invertible at every $\boldsymbol{X}$. The inverse deformation gradient is just

$$\boldsymbol{F}^{-1} = \nabla_x \boldsymbol{X}. \tag{3.3}$$

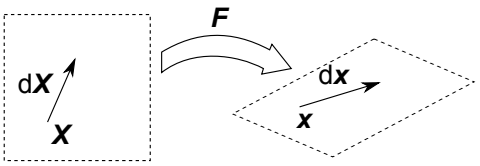

Figure 3.1.: Relationship between Lagrangian coordinate $\boldsymbol{X}$, Eulerian coordinate $\boldsymbol{x}$, and deformation gradient $\boldsymbol{F}$.

An elastic strain energy exists for each material point, and is given by

$$E^{\text{el}}(\boldsymbol{x}) := E^{\text{el}}(\boldsymbol{F}(\boldsymbol{x}), S(\boldsymbol{x})), \tag{3.4}$$

where $\boldsymbol{F}$ and S are the local deformation and entropy, and we will write $E(\boldsymbol{F}, S)$, with the spatial dependence implied, dropping the 'el' label where there are no other contributions to the strain. The important idea is that the local energy is a function of local deformation and entropy only.

We can relate the **stress** to the derivative of the strain energy with respect to deformation. There are several common measures of stress. We use the **Cauchy stress** throughout the thesis, but we first consider briefly the **Piola stress**. This has the simple definition of

$$P_{ij} = \frac{\partial E}{\partial F_{ji}}. \tag{3.5}$$

Note that the derivative is conventionally with respect to the transpose of $\boldsymbol{F}$.

The Piola stress relates forces on the deformed configuration to material area elements in the reference configuration. This is because the deformation gradient relates a displacement in the deformed configuration with the corresponding displacement in the the reference configuration.

The **Cauchy stress** appears naturally in many contexts because it relates quantities entirely in a material's current, deformed state, and does not rely on a reference configuration.

To compute it from the Piola stress, we must relate an area in the deformed configuration to the corresponding area in the reference configuration.

A small volume element $\mathrm{d}v$ in the deformed configuration is related to its corresponding volume element $\mathrm{d}V$ in the reference configuration as

$$\mathrm{d}v = (\det \boldsymbol{F})\,\mathrm{d}V. \tag{3.6}$$

A line element $\mathrm{d}\boldsymbol{l}$ is related to the reference configuration line element $\mathrm{d}\boldsymbol{L}$ as

$$\mathrm{d}\boldsymbol{l} = \boldsymbol{F}\mathrm{d}\boldsymbol{L}. \tag{3.7}$$

To see how an area element in the deformed configuration $\mathrm{d}\boldsymbol{a}$ is related to a reference area element $\mathrm{d}\boldsymbol{A}$, consider a volume element in the material configuration made from an area element $\mathrm{d}\boldsymbol{a}$ and a line element $\mathrm{d}\boldsymbol{l}$:

$$\mathrm{d}a_i \mathrm{d}l_i = \mathrm{d}v. \tag{3.8}$$

$$\mathrm{d}a_i F_{ij} \mathrm{d}L_j = (\det \boldsymbol{F})\,\mathrm{d}V \tag{3.9}$$

$$= (\det \boldsymbol{F})\,\mathrm{d}A_k \mathrm{d}L_k, \tag{3.10}$$

and

$$\mathrm{d}a_i F_{ij} = (\det \boldsymbol{F})\,\mathrm{d}A_k. \tag{3.11}$$

Finally, we obtain the expression relating area elements, as

$$\mathrm{d}\boldsymbol{a} = (\det \boldsymbol{F})\boldsymbol{F}^{-\mathsf{T}}\mathrm{d}\boldsymbol{A}. \tag{3.12}$$

The Cauchy stress relates the traction τ through a (real or imaginary) surface of a material body to the normal vector through the surface $\boldsymbol{n}$, as

$$\tau = \boldsymbol{n}^{\mathsf{T}}\boldsymbol{\sigma}. \tag{3.13}$$

The corresponding expression for the Piola stress is

$$\tau = \boldsymbol{N}^{\mathsf{T}} \boldsymbol{P}, \tag{3.14}$$

where $\boldsymbol{N}$ is the normal to the area element in the reference configuration. From eq. (3.12), we see that the stresses are related as

$$\boldsymbol{P} = (\det \boldsymbol{F})\, \boldsymbol{\sigma} \boldsymbol{F}^{-\mathsf{T}}. \tag{3.15}$$

Conservation of angular momentum implies that $\boldsymbol{\sigma}$ is symmetric. It can be shown (appendix D) that the Cauchy stress is the continuum limit of the atomistic virial stress.

The second law of thermodynamics for a material supporting general stress configurations, expressed in terms of a finite deformation, $\boldsymbol{F}$, is

$$\mathrm{d}E = V \sigma_{ik} F^{-1}_{jk}\, \mathrm{d}F_{ij} \; + T\,\mathrm{d}S. \tag{3.16}$$

The first term is analogous to '$-P\,\mathrm{d}V$' in a fluid, and is equal to it if the material can support only hydrostatic stress.

Since the strain energy of a material does not depend on body rotations performed after a deformation, specifying $E = E(\boldsymbol{F}, S)$ is redundant. A polar decomposition allows any invertible matrix to be written as the product $\boldsymbol{F} = \boldsymbol{R}\boldsymbol{U}$ of an orthogonal matrix $\boldsymbol{R}$ and a symmetric matrix $\boldsymbol{U}$, known as the **stretch tensor**. This can be seen as first applying uniform stretches along three orthogonal axes (represented by $\boldsymbol{U}$) before an arbitrary rotation of $\boldsymbol{R}$. Since any deformation can be described this way, without loss of generality E depends on $\boldsymbol{F}$ only through components of $\boldsymbol{U}$.

In practice, it is computationally cheaper to define the Cauchy–Green tensor, which is also symmetric,

$$\boldsymbol{C} = \boldsymbol{F}^{\mathsf{T}} \boldsymbol{F} = \boldsymbol{U}^{\mathsf{T}} \boldsymbol{R}^{\mathsf{T}} \boldsymbol{R} \boldsymbol{U} = \boldsymbol{U}^{\mathsf{T}} \boldsymbol{U} = \boldsymbol{U}^2. \tag{3.17}$$

Alternatively, and equivalently, we could use the **Finger tensor**

$$\boldsymbol{G} = \boldsymbol{C}^{-1} = \boldsymbol{F}^{-1} \boldsymbol{F}^{-\mathsf{T}}. \tag{3.18}$$

The Cauchy stress *cannot* be uniquely expressed in terms of either $\boldsymbol{U}$ or $\boldsymbol{C}$ however, since a body rotation of $\boldsymbol{R}$ after (an arbitrary) deformation has the effect of reorienting the stress, according to

$$\tilde{\boldsymbol{\sigma}} = \boldsymbol{R}^{\mathsf{T}} \boldsymbol{\sigma} \boldsymbol{R}, \tag{3.19}$$

even though both configurations are represented by the same right stretch tensor $\boldsymbol{U}$.

For an equation of state dependent on $\boldsymbol{G}$, as we will use later, we use eq. (3.16) and the chain rule to obtain

$$\sigma_{ij} = \frac{\rho_0}{\det \boldsymbol{F}} F_{ik} \frac{\partial E}{\partial G_{lm}} \frac{\partial G_{lm}}{\partial F_{jk}}, \tag{3.20}$$

and

$$\frac{\partial G_{ij}}{\partial F_{pq}} = -G_{iq} F^{-1}_{jp} - G_{jq} F^{-1}_{ip}. \tag{3.21}$$

The symmetry of $\boldsymbol{\sigma}$ is most easily shown in terms of $\boldsymbol{C}$,

$$\sigma_{ij} = 2\rho F_{il} \frac{\partial E}{\partial C_{lm}} F_{jm}, \tag{3.22}$$

which is symmetric provided $\frac{\partial E}{\partial C_{lm}}$ is also symmetric. This does not follow from the symmetry of $\boldsymbol{C}$, but since E depends on a symmetric argument, there is always a choice of E where this is true, $\frac{1}{2}(E(\boldsymbol{C}) + E(\boldsymbol{C}^\mathsf{T}))$.

For the special case of an isotropic material, both pre- and post-deformation rotations leave the strain energy invariant. Note that pre-deformation rotations also leave the stress invariant. In this case, we can express energy as a function of the *left* Cauchy–Green tensor

$$\boldsymbol{B} = \boldsymbol{F}\boldsymbol{F}^\mathsf{T} \tag{3.23}$$

or its inverse. A post-deformation rotation has the effect of a similarity transformation on $\boldsymbol{B}$. The invariants under this transformation are all functions of

$$\mathrm{I}_B = \operatorname{tr} \boldsymbol{B} \tag{3.24}$$

$$\mathrm{II}_B = \frac{1}{2}\left((\operatorname{tr}\boldsymbol{B})^2 - \operatorname{tr}(\boldsymbol{B}^2)\right) \tag{3.25}$$

$$\mathrm{III}_B = \det \boldsymbol{B}, \tag{3.26}$$

which are themselves known as the tensor invariants of a 3×3 symmetric matrix, and so E must take the form

$$E = E(\mathrm{I}_B, \mathrm{II}_B, \mathrm{III}_B, S). \tag{3.27}$$

In this case, we also have

$$\boldsymbol{\sigma} = 2\rho \boldsymbol{B} \left(\frac{\partial E}{\partial \boldsymbol{B}}\right)_S, \tag{3.28}$$

where in terms of $\boldsymbol{B}$,

$$\rho = \rho_0 / \sqrt{\det \boldsymbol{B}}. \tag{3.29}$$

3.2. Conservation Laws for Nonlinear Elasticity

Equations representing conservation of momentum and energy in an elastic material with no viscosity are

$$\partial_t(\rho v_i) + \partial_k(\rho v_i v_k - \sigma_{ik}) = 0 \tag{3.30}$$

$$\partial_t(\rho E) + \partial_k(\rho E v_k - v_i \sigma_{ik}) = 0. \tag{3.31}$$

Equations (3.30) and (3.31) are very similar to those appearing in the Euler equations of fluid dynamics, and can be derived in a similar manner; see Plohr and Sharp (1988). We have in addition the conservation of mass, giving an equation for the density,

$$\partial_t \rho + \partial_k(\rho v_k) = 0. \tag{3.32}$$

Together, eqs. (3.30) to (3.32) with an equation of state $E(p, \rho)$, and $\sigma_{ij} = -p\delta_{ij}$ would specify a fluid system.

In a solid system, the equation of state has one of the more general forms discussed in the previous section, and we must introduce an equation for the time dependence of the deformation gradient (or one of its analogues), which can be seen as a generalization of the density, found via eq. (3.2). We obtain such an equation from the symmetry of mixed partial derivatives:

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{F}(\boldsymbol{X}, t) = \frac{\partial}{\partial t}\boldsymbol{F}(\boldsymbol{X}, t) = \frac{\partial}{\partial t}\left(\frac{\partial \boldsymbol{x}}{\partial \boldsymbol{X}}\right) = \frac{\partial}{\partial \boldsymbol{X}}\left(\frac{\partial \boldsymbol{x}}{\partial t}\right) = \frac{\partial \boldsymbol{v}}{\partial \boldsymbol{X}}, \tag{3.33}$$

and we obtain a conservation law for the inverse deformation gradient:

$$\frac{\mathrm{d}\boldsymbol{F}}{\mathrm{d}t} = -\boldsymbol{F}\frac{\mathrm{d}\boldsymbol{F}^{-1}}{\mathrm{d}t}\boldsymbol{F} = \frac{\partial \boldsymbol{v}}{\partial \boldsymbol{X}} = \frac{\partial \boldsymbol{v}}{\partial \boldsymbol{x}}\boldsymbol{F} \tag{3.34}$$

$$\iff \quad \frac{\mathrm{d}\boldsymbol{F}^{-1}}{\mathrm{d}t} = -\boldsymbol{F}^{-1}\frac{\partial \boldsymbol{v}}{\partial \boldsymbol{x}} \tag{3.35}$$

$$\iff \quad \frac{\partial \boldsymbol{F}^{-1}}{\partial t} = -\frac{\partial \boldsymbol{F}^{-1}}{\partial \boldsymbol{x}}\boldsymbol{v} - \boldsymbol{F}^{-1}\frac{\partial \boldsymbol{v}}{\partial \boldsymbol{x}} \tag{3.36}$$

$$= -\frac{\partial}{\partial \boldsymbol{x}}(\boldsymbol{F}^{-1}\boldsymbol{v}), \tag{3.37}$$

or in index notation, with $\boldsymbol{g} = \boldsymbol{F}^{-1}$, as

$$\partial_t(g_{ij}) + \partial_j(g_{ik}v_k) = 0. \tag{3.38}$$

The conservation law may may alternatively be written in terms of $\boldsymbol{F}$ as

$$\partial_t(\rho F_{ij}) + \partial_k(\rho F_{ij}v_k - \rho F_{kj}v_i) = 0. \tag{3.39}$$

This can be proved using the conservation of mass, the Eulerian form of eq. (3.33) and Piola's identity (Plohr and Sharp, 1988),

$$\partial_i(A_{ij}/\det\boldsymbol{A}) = 0. \tag{3.40}$$

A compatibility condition is required so that for each i, F_{ij} is genuinely the gradient of (one component of) some deformation. A zero curl is sufficient to ensure this,

$$\nabla \times \boldsymbol{F}_i = 0 \tag{3.41}$$

where $\boldsymbol{F}_i$ is the ith row of $\boldsymbol{F}$. A non-zero value could correspond either to regions with voids, or material volumes that unphysically interpenetrate each other. If eq. (3.41) holds initially, it is true for all subsequent time. This is not necessarily the case for a numerical solution however, and the numerical method used should ensure that this quantity does not deviate much from zero. This can be achieved by using a source term in the evolution (as e.g. Miller and Colella, 2001).

An alternative statement, and appropriate for weak solutions in $\boldsymbol{F}$, is that following a closed loop in material coordinates should make a closed loop in spatial coordinates as well, so that the line integral of $\boldsymbol{F}$ should vanish around the curve. For differentiable $\boldsymbol{F}$, this is equivalent to eq. (3.41) by Stokes' theorem.

Equation (3.39) is in conservative form for each component ij, and strict conservation can be maintained with a suitable conservative numerical method (see below). There is no reason for the density itself to be conserved by a numerical method however, since errors in each component may not cancel. Miller and Colella (2001) solve this problem by explicitly including an equation for the density (eq. (3.32)) and enforcing consistency of the overdetermined system.

We now discuss the character of the solutions of this system of equations.

3.2.1. The Riemann problem

Similarity solutions (constant along paths of constant x/t) arise from **Riemann problem** initial conditions, consisting of two regions of piecewise-constant data separated by a discontinuity. The time evolution of this initial data comprises a number of separated, propagating waves, equal in number to the distinct characteristic speeds (fig. 3.2 shows an example). There are two possibilities for the shape of these waves.

A **fan** is an expanding similarity structure propagating at the local characteristic speed everywhere. The front propagates at the characteristic speed of the material ahead of it, and the characteristic speed decreases across it. These are normally associated with rarefaction waves since physically, decreasing the sound speed is achieved by lowering the density.

A **shock** is a steady similarity solution. With no higher order terms, these necessarily consist of discontinuities moving at a speed higher than the characteristic speed ahead of the shock and lower than behind it.

In addition, a **contact** is a wave with a jump in solution across the wave that does not change the characteristic speed, and they propagate at exactly the characteristic speed. Contact waves can be degenerate (multiple linearly independent wave-jumps with the same speed) without affecting the desirable strong hyperbolic property of the system (LeVeque, 2002, §16.2).

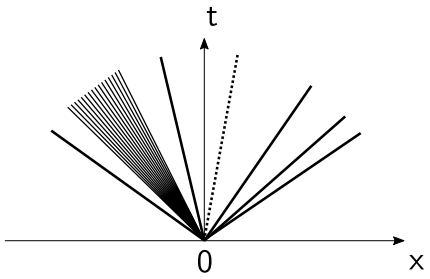

Figure 3.2.: Schematic of waves arising from one possible Riemann problem for the hyperelastic system. There is one rarefaction fan (second left-going wave, shown with many lines), five shock waves (solid lines) and a contact wave (the dotted line).

3.2.2. Characteristic structure of the solutions

A general hyperbolic system can be written in quasi-linear form as

$$w_t + \mathscr{A}(x)\, w_x = 0 \tag{3.42}$$

where $\boldsymbol{w}$ is the vector of primitive variables, and $\mathscr{A} = \partial \boldsymbol{f} / \partial \boldsymbol{w}$ is the Jacobian, with $\boldsymbol{f}$ the vector of conservative fluxes.

It can be demonstrated that the system of equations comprising eqs. (3.30), (3.31) and (3.39) is hyperbolic by showing that the eigenvalues of $\mathscr{A}$ are real. The vector of primitive variables in this case is $\boldsymbol{w} = (\boldsymbol{v}, \boldsymbol{F}, S)$. At this point, we switch to using full tensor notation, but only because it is more compact, and the discussion holds equally well for Cartesian tensors.

This system may be written compactly as follows

$$\begin{pmatrix} \dot{v}^l \\ \dot{F}^m_n \\ \dot{S} \end{pmatrix} = \begin{pmatrix} v^\alpha \delta^l_i & -A^{\alpha lk}_j & -B^{\alpha l} \\ -F^\alpha_n \delta^m_i & v^\alpha \delta^m_j \delta^k_n & 0 \\ 0 & 0 & v^\alpha \end{pmatrix} \begin{pmatrix} v^i{}_{;\alpha} \\ F^j_{k\,;\alpha} \\ S_{;\alpha} \end{pmatrix} \tag{3.43}$$

where $A^{ijl}_k = \frac{1}{\rho} \frac{\partial \sigma^{ij}}{\partial F^k_l}$, $B^{ij} = \frac{1}{\rho} \frac{\partial \sigma^{ij}}{\partial S}$, and the '0' elements of the matrix are zero tensors of the correct type (that is, with the correct number and placement of indices, and every component equal to zero).

This matrix equation should be interpreted as follows. Each tensor appearing as an element of the above matrix is a submatrix of an overall 13×13 matrix. Each index ranges over each spatial coordinate direction, $1 \leq i, j, k, l, m, n, \alpha \leq 3$. The overall matrix-vector product can be computed term by term, summing over repeated indices. The free indices then match those on the left-hand side of the equation. We should emphasize that this compact notation can always be written in the more conventional way of an explicit matrix by writing out each component, with the i, j, k appearing in the *matrix* indexing *rows* of a submatrix, and l, m, n indexing *columns*. The indices i, j, k appearing in the *vector* on the right-hand side, and l, m, n on the left-hand side would all index *rows* in the expanded form. When more than one index appears for a row index on a given term, they are 'flattened' in a consistent order, and treated as a row of length nine. The remaining index α is normally fixed along one axis and summed over explicitly when written in the expanded form.

The index α refers to a component in one direction of propagation. Henceforth, we consider a one-dimensional system along a direction $\boldsymbol{n}$, and we omit the index 'α' in the

above system for brevity, writing instead

$$v = n_\alpha v^\alpha$$
$$A^{ij}_k = n_\alpha A^{\alpha ij}_k$$
$$F_i = n_\alpha F^\alpha_i$$
$$\text{etc.}$$

The eigenvalues of the Jacobian are the linearized wave speeds. The characteristic equation is

$$\det \begin{pmatrix} (v-\lambda)\delta^l_i & -A^{lk}_j & -B^l \\ -F_n \delta^m_i & (v-\lambda)\delta^m_j \delta^k_n & 0 \\ 0 & 0 & v-\lambda \end{pmatrix} = 0. \tag{3.44}$$

Expanding the determinant in the last row, this becomes

$$(v-\lambda)\det \begin{pmatrix} (v-\lambda)\delta^l_i & -A^{lk}_j \\ -F_n \delta^m_i & (v-\lambda)\delta^m_j \delta^k_n \end{pmatrix} = 0. \tag{3.45}$$

It is useful to note the following identity for a block-partitioned $n \times n$ matrix:

$$\det \begin{pmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{pmatrix} = \det(\boldsymbol{A} - \boldsymbol{B}\boldsymbol{D}^{-1}\boldsymbol{C}) \det(\boldsymbol{D}), \tag{3.46}$$

where $\boldsymbol{A}$ is $k \times k$, $\boldsymbol{B}$ is $k \times (n-k)$, $\boldsymbol{C}$ is $(n-k) \times k$ and $\boldsymbol{D}$ is $(n-k) \times (n-k)$. From this identity, eq. (3.45) becomes

$$(v-\lambda)\det \left((v-\lambda)\delta^l_i - \frac{A^{lk}_i F_k}{(v-\lambda)} \right) \det \left((v-\lambda)\delta^m_j \delta^k_n \right) = 0, \tag{3.47}$$

and

$$(v-\lambda)^7 \det \left((v-\lambda)^2 \delta^i_j - \Omega^i_j \right) = 0 \tag{3.48}$$

where the *acoustic tensor* is defined as

$$\Omega^i_j = A^{ik}_j F_k. \tag{3.49}$$

We see that there is a sevenfold degenerate wave, with wave speed $\lambda = v$, and the remaining waves have speeds such that $(v-\lambda)^2$ is an eigenvalue of the acoustic tensor. This admits six

waves, with three either side of the wave with $\lambda = v$, provided that the three eigenvalues of the acoustic tensor are real and positive.

For the system to remain strongly hyperbolic, the waves with repeated eigenvalue should be linearly degenerate: a change in the solution parallel to an eigenvector $\boldsymbol{r}$ therefore should not affect the corresponding eigenvalue λ to first order:

$$(\boldsymbol{r}(\boldsymbol{q}) \cdot \nabla_{\boldsymbol{q}})\lambda(\boldsymbol{q}) = 0. \tag{3.50}$$

To show this, note that the repeated eigenvalue is just v, so look for eigenvectors of the form $(0, \boldsymbol{F}, S)$. Since a change in the solution by this amount does not affect the velocity, neither does it affect the wave speed. These must satisfy

$$A^{lk}_{j} F^{j}_{k} + B^{l} S = 0, \tag{3.51}$$

which is the kernel of a linear system of three equations in ten variables, which in the general (full rank) case, has seven dimensions, so seven eigenvectors of this form can indeed be found in general.

A sufficient condition for hyperbolicity is for

$$\frac{\partial \sigma^{\alpha i}}{\partial F^{j}_{\beta}} > 0 \tag{3.52}$$

for all i, j, α and β, which is just the condition of convexity on the equation of state.

The full set of eigenvectors are not given here, but are derived analytically in terms of the acoustic tensor for this system by Barton et al. (2009), and for a similar elastic-plastic system by Miller and Colella (2001).

3.3. The Equation of State

We have considered here an elastic formulation with an energy function, from which we can obtain stresses (and therefore the dynamical behaviour of the material) via eq. (3.20) or similar. The existence of such a strain energy function makes this a **hyperelastic** constitutive model. Several popular choices of constitutive model for nonlinear elasticity have used a **hypoelastic** formulation, where no such strain energy function exists. This is problematic because the energy of a deformed configuration depends on the precise path of deformation, which seems contrary to experience, even if Cauchy elasticity permits it.

Moreover, it is possible to find a closed cycle of deformation that results in a net increase in strain energy (Carroll, 2009, citing Rivlin).

For an anisotropic equation of state, the initial orientation of the material is important. This orientation is built into the equation of state.

$$\boldsymbol{F} = \boldsymbol{F}^{\text{el}}\boldsymbol{F}^{\text{orient}} \tag{3.53}$$

where $\boldsymbol{F}^{\text{orient}}$ is a rotation matrix. $\boldsymbol{F}^{\text{el}}$ is maintained by the system (that is, it is the '$\boldsymbol{F}$' appearing in eq. (3.39)), and satisfies the compatibility constraint, but the energy and stress depend on $\boldsymbol{F}$.

We do this by introducing the orientation as an advected quantity,

$$\partial_t(\rho F_{ij}^{\text{orient}}) + \partial_k(\rho v_k F_{ij}^{\text{orient}}) = 0. \tag{3.54}$$

The orthogonality of $\boldsymbol{F}^{\text{orient}}$ is preserved by the flow, since from the continuity equation, eq. (3.54) can be expressed, dropping the 'orient' label for brevity, as:

$$\partial_t F_{ij} + v_\alpha \partial_\alpha F_{ij} = 0, \tag{3.55}$$

so that

$$\begin{aligned}
\partial_t(F_{ki}F_{kj}) &= F_{ki}\partial_t F_{kj} + F_{ki}\partial_t F_{kj} && (3.56)\\
&= -v_\alpha(F_{ki}\partial_\alpha F_{kj} + F_{kj}\partial_\alpha F_{ki}) && (3.57)\\
&= -v_\alpha \partial_\alpha(F_{ki}F_{kj}). && (3.58)
\end{aligned}$$

If $F_{ki}F_{kj} = \delta_{ij}$ everywhere initially, then the final expression is zero, so it remains orthogonal at all later times. This constraint can be strongly enforced in a numerical scheme by mapping back to a nearby orthogonal matrix, to remove any drift from orthogonality (Higham, 2008, §2.6).

In a one-dimensional problem, with two crystal orientations present, we can keep track of a single scalar quantity representing the crystal orientation at a given location as a fractional rotation between the two orientations, in lieu of eq. (3.54).

This is a particularly simple special case of an interface problem in a multi-material system. The above description treats interfaces between orientations as diffuse, with a mixture model provided by the equation of state itself.

3.4. Eulerian finite-volume methods

In solid mechanics problems, the preferred way of representing the domain has traditionally been in the Lagrangian frame (computational cells deforming with the material) rather than the Eulerian frame (a grid of cells at fixed positions in space). There are a number of reasons to favour a finite-volume method in an Eulerian frame for certain solid problems, however.

First, the method does not break down under large deformations due to a highly distorted mesh, which can be a problem in a Lagrangian formulation. This can mean that a simulation domain contains cells with a small length along one or more dimensions, limiting the timestep and therefore reducing the accuracy and increasing the computational cost of the simulation.

Second, the field of shock-capturing methods in an Eulerian, finite-volume setting is mature and well-understood, and we are partly motivated by the solution to problems containing shock-waves. The principal advantage of a finite-volume scheme in this situation is that it is conservative and therefore the shock jump conditions are guaranteed to be satisfied. Non-conservative formulations suffer from incorrect wave speeds.

Third, simulations of fluids typically are solved in an Eulerian frame for the reasons given above. Using the same frame for the solid makes simulations of coupled fluid-solid problems convenient, although we do not consider this here.

The equation of state model developed in chapter 4 is independent of the method used to represent the domain or solve the hyperelastic system, and could also be used with a Lagrangian representation, but we assess the performance of the equation of state model using an Eulerian finite-volume code.

A general conservation law has form

$$\boldsymbol{q}_t + \boldsymbol{f}(\boldsymbol{q})_x = \mathbf{0} \tag{3.59}$$

where $\boldsymbol{q}$ is a vector of conserved variables and $\boldsymbol{f}$ is the vector of fluxes. To admit weak solutions (shock waves), where the derivatives are undefined, we instead write this as

$$\int_{x_1}^{x_2} [\boldsymbol{q}(x,t_2) - \boldsymbol{q}(x,t_1)]\,\mathrm{d}x + \int_{t_1}^{t_2} [\boldsymbol{f}(\boldsymbol{q}(x_2,t)) - \boldsymbol{f}(\boldsymbol{q}(x_1,t))]\,\mathrm{d}t = \mathbf{0} \tag{3.60}$$

in any region $[x_1, x_2]$ and $[t_1, t_2]$.

In a finite-volume representation (with a uniform grid Δx and timestep Δt) our data represent averages over the whole cell. At timestep n, and for cell i centred on x_i,

$$Q_i^n = \frac{1}{\Delta x}\int_{x_i-\Delta x/2}^{x_i+\Delta x/2} \mathbf{q}(x, t^n)\,\mathrm{d}x. \tag{3.61}$$

A forward-Euler update step for a finite-volume representation of eq. (3.59) is

$$Q_i^{n+1} = Q_i^n - \frac{\Delta t}{\Delta x}(f_{i+1/2}^n - f_{i-1/2}^n). \tag{3.62}$$

The numerical fluxes $f_{i\pm1/2}^n = f(q(x_i \pm \Delta x/2, t^n))$ are evaluated at the cell interfaces. This guarantees q is conserved over the whole domain.

Since the value of q is not known at the cell interfaces, only its average over the whole cell, the fluxes must be approximated there. The first-order centred (FORCE) numerical flux is defined as

$$f_{i+1/2}^{\mathrm{FORCE}}(Q_i, Q_{i+1}) = \frac{1}{2}(f_{i+1/2}^{\mathrm{RI}} + f_{i+1/2}^{\mathrm{LF}}) \tag{3.63}$$

where the Lax-Friedrichs flux is

$$f_{i+1/2}^{\mathrm{LF}} = \frac{1}{2}\left(f(Q_i) + f(Q_{i+1}) + \frac{\Delta x}{\Delta t}(Q_i - Q_{i+1})\right) \tag{3.64}$$

and the Richtmyer flux is

$$f_{i+1/2}^{\mathrm{RI}} = f(Q_{i+1/2}^{n+1/2}), \tag{3.65}$$

where,

$$Q_{i+1/2}^{n+1/2} = \frac{1}{2}\left(Q_i + Q_{i+1} + \frac{\Delta t}{\Delta x}\left(f(Q_i) - f(Q_{i+1})\right)\right). \tag{3.66}$$

The motivation for defining a flux in this way is that while still first order accurate in space, the error has been reduced by a constant factor in combining first and second order fluxes—equal weighting is the greatest reduction in error while still avoiding the undesirable feature of using a linear second-order method, of spurious oscillations introduced around large gradients in the solution. See Toro (2013, §§7.4.3, 18.2.1).

The flux depends on the state immediately to the left and right of the interface between cells i and $i+1$. In the expressions above, these are taken as the averages over the entirety of the cells, Q_i and Q_{i+1}. This amounts to assuming a piecewise constant reconstruction of the solution from the cell averages.

One approach to higher orders of spatial accuracy of the method is to perform a higher-order reconstruction of the solution, and evaluate the fluxes at the reconstructed

values at the cell interfaces. The numerical flux is then

$$f_{i+1/2} = f^{\text{FORCE}}_{i+1/2}(\tilde{Q}^L_{i+1/2}, \tilde{Q}^R_{i+1/2}), \tag{3.67}$$

where $\tilde{Q}^L_{i+1/2}$ is the reconstruction at the left-hand side of the interface.

An interface reconstruction affects only the *spatial* order of accuracy. Discretization in time should have the same order of accuracy for the method to have the same overall order.

3.4.1. WENO reconstruction

One such method of reconstruction is the weighted essentially non-oscillatory (WENO) approach described by Jiang and Shu (1996); Liu et al. (1994), a family of schemes of any desired order of accuracy for smooth solutions. This built on the earlier ENO scheme (Shu and Osher, 1988). The idea is to perform several piecewise polynomial interpolations of the cell-averaged solution (at different points of a stencil), before choosing the smoothest combination of these interpolants (in a certain sense, defined below). A review of high-order WENO schemes is given by Gerolymos et al. (2009).

The resulting schemes can be shown not to suffer from Gibbs' phenomenon at discontinuities. It may still produce some spurious oscillations, however, but ones that decay polynomially with resolution $O(\Delta x^l)$, instead of the $O(1)$ local error that characterizes Gibbs' ringing.

The WENO-5 reconstruction takes a five point stencil centred around cell i, and produces three reconstructions of the solution at $i + \frac{1}{2}$ using three subsets of the stencil points:

$$w^0 = \frac{1}{3}w_{i-2} - \frac{7}{6}w_{i-1} + \frac{11}{6}w_i \tag{3.68}$$

$$w^1 = -\frac{1}{6}w_{i-1} + \frac{5}{6}w_i + \frac{1}{3}w_{i+1} \tag{3.69}$$

$$w^2 = \frac{1}{3}w_i + \frac{5}{6}w_{i+1} - \frac{1}{6}w_{i+2}. \tag{3.70}$$

The final reconstruction is a linear combination of these intermediate reconstructions,

$$w_{i+1/2} = \omega_0 w^0 + \omega_1 w^1 + \omega_2 w^2. \tag{3.71}$$

The ideal weights are $\hat{\omega} = (0.1, 0.6, 0.3)$, which are used for a smooth solution. These are modified according to the smoothness indicators

$$\beta_0 = \frac{13}{12}(w_{i-2} - 2w_{i-1} + w_i)^2 \quad + \frac{1}{4}(w_{i-2} - 4w_{i-1} + 3w_i)^2 \tag{3.72}$$

$$\beta_1 = \frac{13}{12}(w_{i-1} - 2w_i \quad + w_{i+1})^2 + \frac{1}{4}(w_{i-1} - w_{i+1})^2 \tag{3.73}$$

$$\beta_2 = \frac{13}{12}(w_i \quad - 2w_{i+1} + w_{i+2})^2 + \frac{1}{4}(3w_i \quad - 4w_{i+1} + w_{i+2})^2 \tag{3.74}$$

as

$$\alpha_i = \frac{\hat{\omega}_i}{\epsilon + \beta^2} \tag{3.75}$$

and the weights used are found by normalizing

$$\omega_i = \frac{\alpha_i}{\sum \alpha_i}, \tag{3.76}$$

where the purpose of ϵ is to prevent division by zero, and can by small.

This reconstruction is nominally fifth-order accurate, but can degenerate into third-order accuracy for some (smooth) problems—see the example below. A fix to provide a fifth-order accurate reconstruction from the same stencil (and in general, nth order from WENO-n) while still retaining the favourable ENO property is provided by Henrick et al. (2005). This is known as mapped-WENO, or WENOM.

Even though oscillations near discontinuities are not as severe as a straightforward polynomial reconstruction, they are still present and can be intrusive. The MPWENO scheme of Balsara and Shu (2000) leaves the reconstruction on slowly varying solution regions untouched, but limits the value of the reconstruction near sharp changes in the solution, to obtain a monotonicity-preserving scheme.

Titarev and Toro (2004) note that for systems of equations, the reconstruction can perform quite badly unless it is done on the characteristic variables (eigenvectors of the Jacobian in the quasi-linear form), and not the conserved or primitive variables themselves. Obtaining the characteristic decomposition could be quite expensive (depending on the system) so may be not be worth the trade off. We note here that it is possible to use a reconstruction on non-characteristic variables in the hyperelastic system provided the limiting parameters are chosen quite conservatively. Similar conclusions were reached by Qiu and Shu (2002).

One advantage of using this approach is that reconstruction-based approaches to high order do not depend on details of the hyperbolic system, although monotonicity-preserving high-order methods are possible based on computing the fluxes.

WENO requires a strong-stability-preserving time integrator to be stable. See Gottlieb et al. (2001) for several such schemes (for linear and nonlinear problems).

3.4.2. Validation results

Several validation tests are given below for the hyperbolic solver. We use centred fluxes with WENO reconstruction.

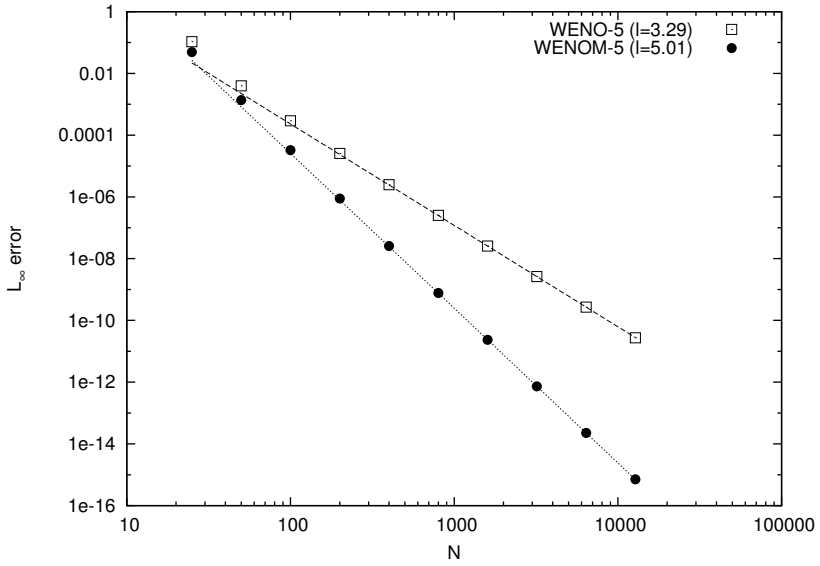

Figure 3.3.: Convergence of L_∞ error for WENO-5 and WENOM-5 for linear advection of the profile shown in the inset.

As a test of the WENO scheme, an initial profile given by

$$\sin(\pi x - \sin(\pi x)/\pi) \tag{3.77}$$

for $x \in [0, 1)$ and resolved with N cells is advected through a periodic domain. The fifth-order ℓ-SSPRK timestepping described in Gottlieb et al. (2001) was used. The error for the WENO-5 and WENOM-5 schemes is shown in fig. 3.3, achieving third and fifth order accuracy, as expected for this test case.

Two tests described in Barton et al. (2009) are performed for the hyperelastic system, using MPWENO-5 reconstruction of centred FORCE fluxes. The timestepping was third-order

Table 3.1.: Hyperelastic test problems from Barton et al. (2009).

	$\boldsymbol{u}_L$ (km/s)	$\boldsymbol{F}_L$	S_L (kJ/g/K)	$\boldsymbol{u}_R$ (km/s)	$\boldsymbol{F}_R$	S_R (kJ/g/K)
Test case 1	$\begin{pmatrix} 0 \\ 0.5 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 0.98 & 0 & 0 \\ 0.02 & 1 & 0.1 \\ 0 & 0 & 1 \end{pmatrix}$	10^{-3}	$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0.1 \\ 0 & 0 & 1 \end{pmatrix}$	0
Test case 2	$\begin{pmatrix} 2 \\ 0 \\ 0.1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ -0.01 & 0.95 & 0.02 \\ -0.015 & 0 & 0.9 \end{pmatrix}$	0	$\begin{pmatrix} 0 \\ -0.03 \\ -0.01 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0.015 & 0.95 & 0 \\ -0.01 & 0 & 0.9 \end{pmatrix}$	0

strong stability-preserving Runge-Kutta. Both test problems have Riemann-problem initial conditions of two uniform regions separated by a discontinuity at $x = 0.5$. They are given in table 3.1, assuming the Romenski equation of state parameterized for copper (eq. (A.5), table A.1). The results are shown in Figures 3.4 and 3.5 (pages 47 and 48).

Since the exact solutions involve discontinuities, it is possible to achieve at best globally first-order accurate numerical results. The solutions compare favourably with the results of Barton et al. (2009), despite using centred fluxes. For this reason, in the first test the jumps in transverse stress (but not in longitudinal stress) are captured over 7–8 cells instead of the six cells of Barton. The longitudinal shocks are captured equally well in both results, over three cells (varying with the strength of the wave). A number of artifacts due to the WENO reconstruction on the primitive variables are visible, particularly in the leftmost wave of the u_1 profile of fig. 3.4, and the overshoots on the two narrow regions between waves in the σ_{12} profile in the same figure.

The second test shows a strong heating effect at the central contact discontinuity, which reduces slowly with resolution. It must be run with a low CFL number (here 0.4) to avoid the onset of persistent oscillations in the entropy profile.

Figure 3.4.: Test case 1 from Barton et al. (2009, fig. 2), solved using MPWENO-5/FORCE, with a CFL number of 0.6. The black points are the solution on a 500 cell grid, and the red line a reference calculation on a highly converged grid, shown at time $t = 0.6\,\mu s$.

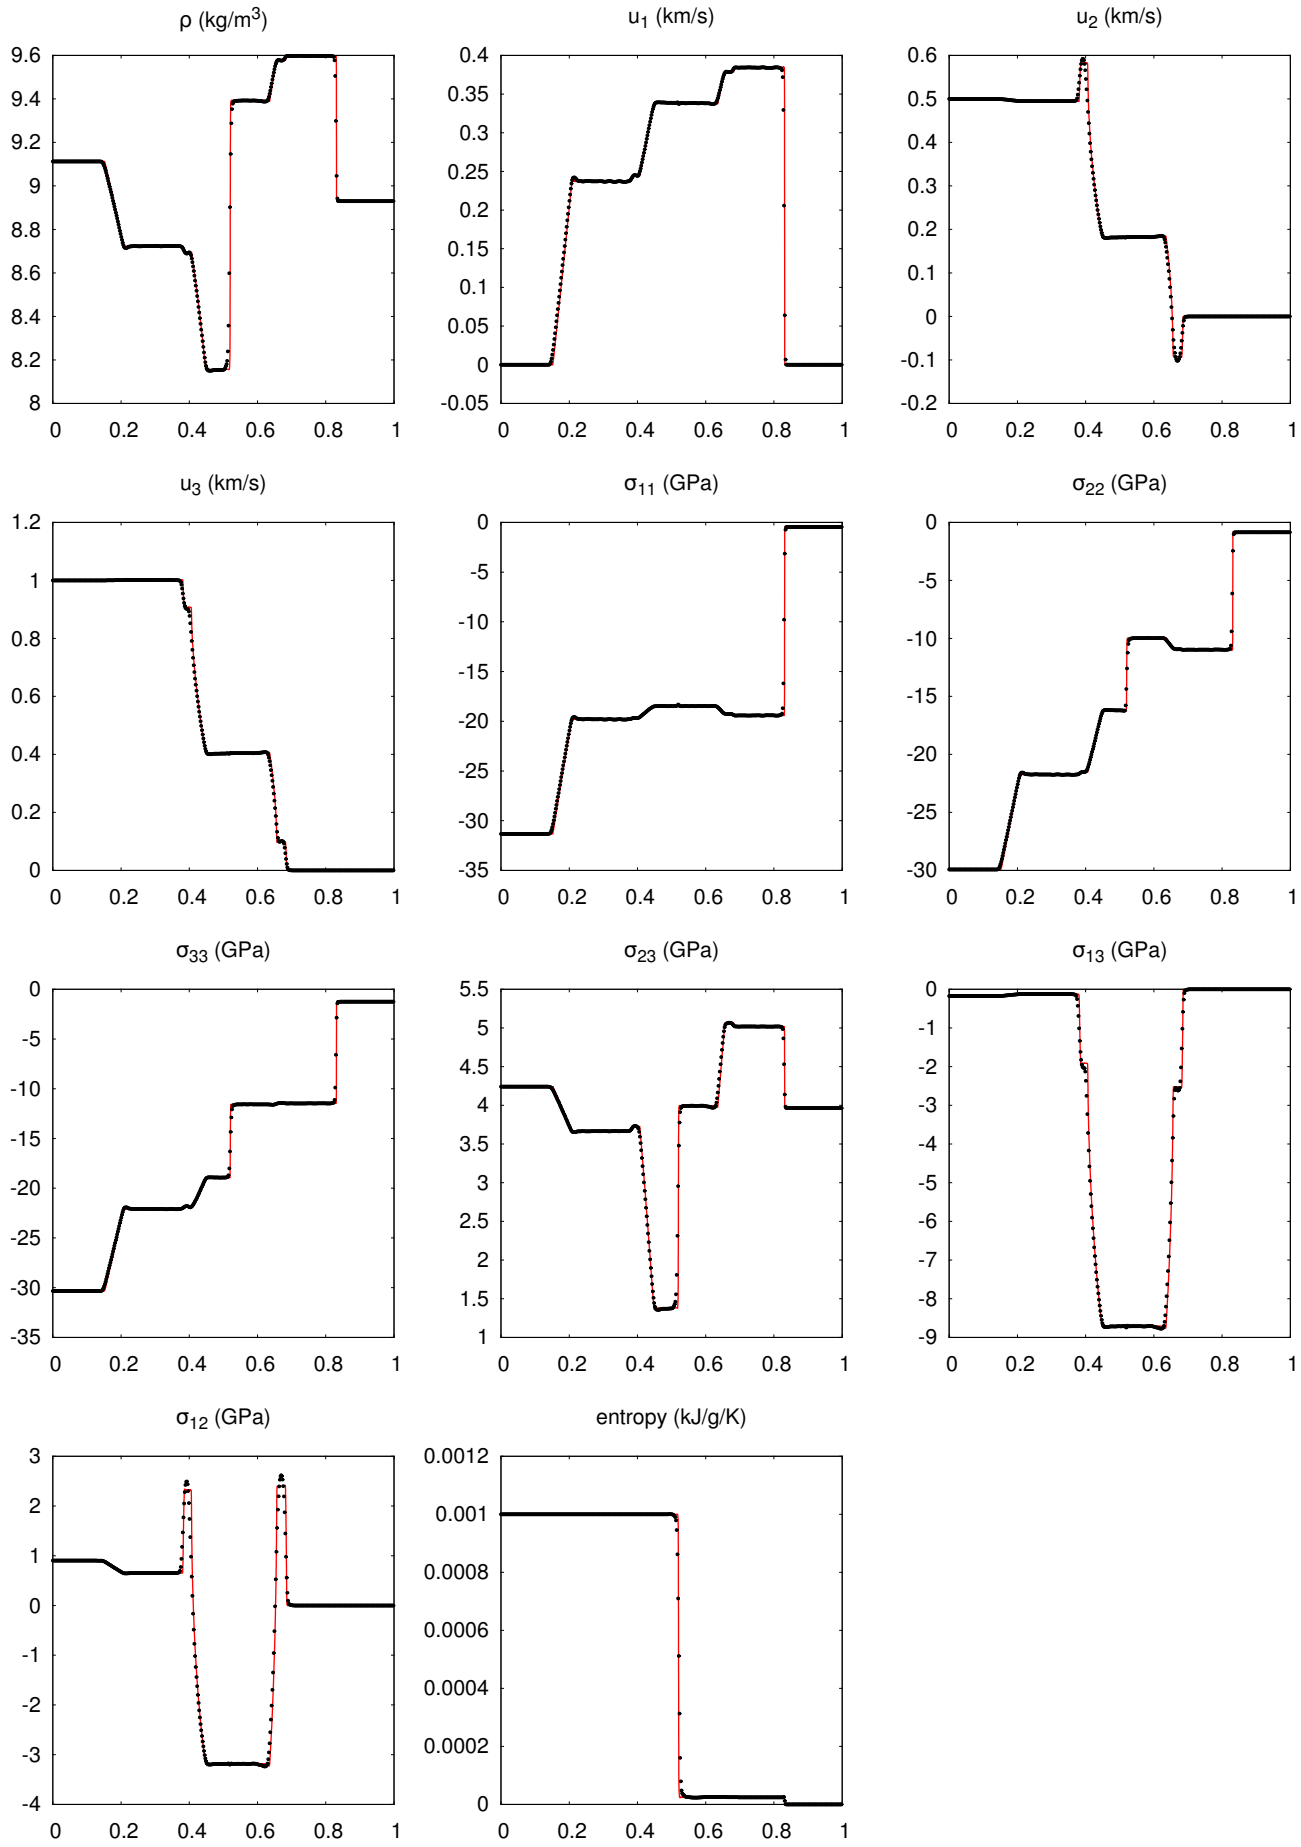

Figure 3.5.: Test case 2 from Barton et al. (2009, fig. 4), solved using MPWENO-5/FORCE, with a CFL number of 0.4. The black points are the solution on a 500 cell grid, and the red line a reference calculation on a highly converged grid, shown at time $t = 0.6\,\mu s$.

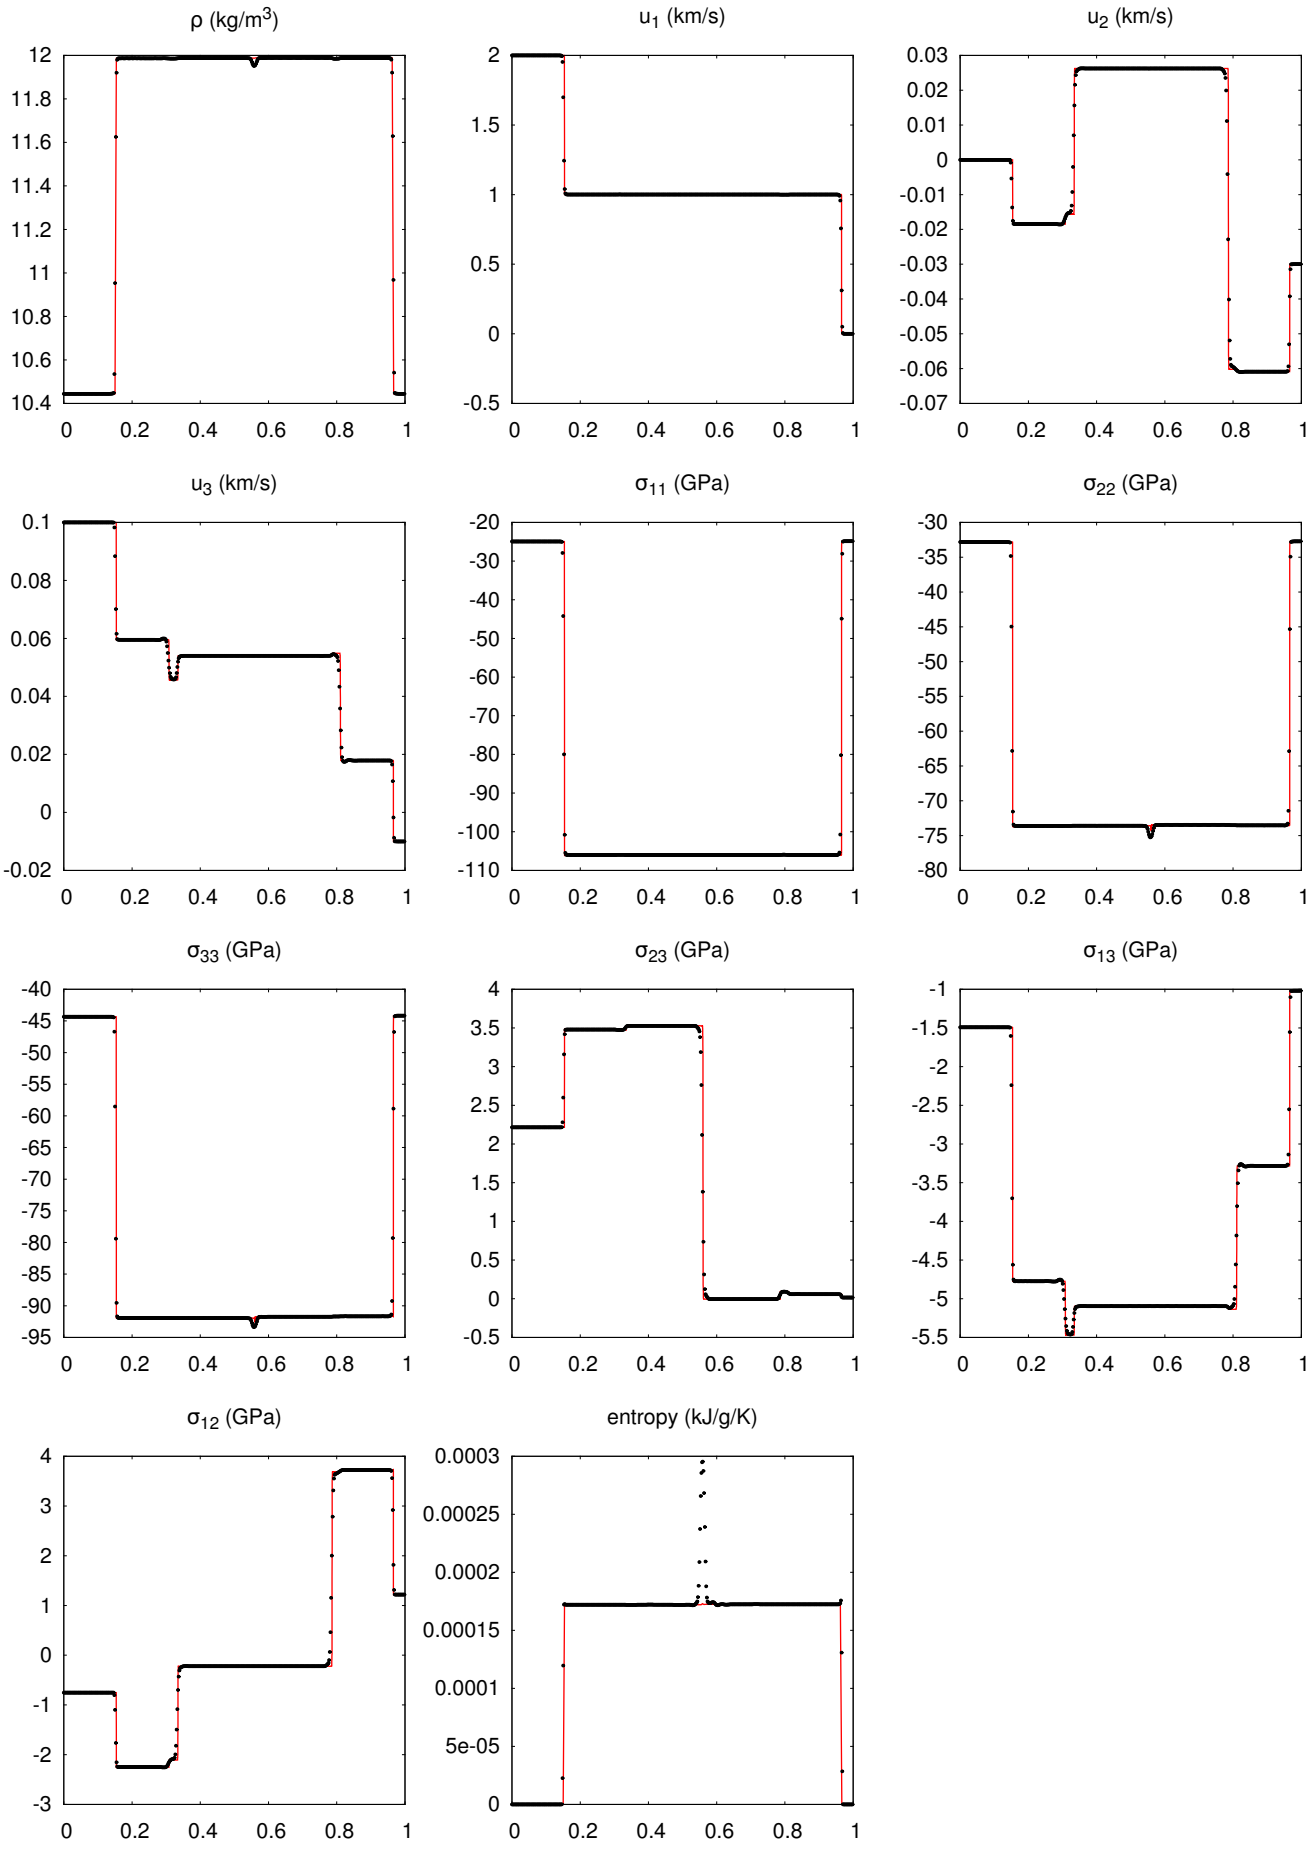

4. An Equation of State Model from First Principles Molecular Dynamics

4.1. Introduction and Motivation

An equation of state (EOS) is a central part of any continuum mechanics simulation, representing the properties of the specific system being simulated. For solid mechanics simulations, there are several regimes of interest, from low-strain elasticity at constant temperature, through to large strain and high strain-rate problems, where the material may have yielded or fractured. Often, both kinds of behaviour are present in the same problem. For example, a high speed impact may result in high stresses near the point of impact, but with elastic waves of low peak stress propagating far into the surrounding material. Beyond elasticity, effects such as plasticity and yield, work-hardening for ductile materials and fracture can be incorporated into a continuum solid-mechanics simulation, but a good elastic model is required upon which to build these other effects.

Coupled simulations of molecular dynamics and continuum mechanics provide another motivation for having a consistent material description for both frames, including the required anisotropy (Abraham et al., 1998; Curtin and Miller, 2003; Miller and Tadmor, 2009). Several of the methods described by the latter review articles can work with finite temperatures and nonlinear elasticity.

Equations of state as used in many branches of continuum mechanics (for fluids, elastic-solids and solids or fluids with more complex behaviour), are typically simple analytic forms or tables. The main (and considerable) advantage of this approach is that these forms may be evaluated quickly; even then, equation of state evaluations can still form the bottleneck of a computation. The disadvantage of this approach is that a given prescribed form may not be able to capture the shape of the deformation energy surface of a material well. Likewise, a tabular approach has difficulty being extended to the number of dimensions present for anisotropic elasticity.

Molecular dynamics provides a way of sampling the energy of a material as a function of, say, temperature and density, and therefore its equation of state. The material of interest is described by atoms interacting via a given potential. It is possible to use a first-principles technique such as density functional theory to compute the forces between the atoms and the equation of state is then truly a predictive, first-principles result.

The relation between continuum elasticity and a deforming atomistic lattice is made at zero temperature by the Cauchy–Born rule. This states that, for a lattice with a basis, the elastic strain energy of a deformation is the local minimum over the configuration of the basis atoms, for lattice vectors deformed from the original configuration by the deformation gradient (Weinan and Ming, 2007). At finite temperature, the new lattice vectors are obtained in the same way, but the basis atom positions are retained as degrees of freedom.

In this chapter, we demonstrate a method for determining the temperature-dependent hyperelastic equation of state for an anisotropic material from first-principles molecular dynamics. A method is considered for finding deformed states of a material on the same isentrope as a given starting configuration. The energies and stresses of a number of points are sampled from DFT molecular dynamics using this method, over a given range of the seven dimensional space of deformation and initial energy (of the undeformed configuration). The full energy surface within this range can then be reliably reconstructed using Gaussian process regression and stresses also found analytically. Validation of the reconstruction is provided by applying it to an analytic equation of state, and we perform convergence tests with empirical interatomic potentials. An equation of state model is constructed for DFT silicon, and demonstrated within a finite-volume continuum elasticity simulation.

We are particularly interested in applications involving shock waves, and this motivates many of the continuum initial-value problems considered in section 4.6. Henderson (2001) gives an overview of the effect of the equation of state on shock wave propagation, although this work is limited to the hyperbolic, convex case.

4.2. Related Work

Swift et al. (2001) consider a quasiharmonic phonon approach to producing an isotropic equation of state for silicon, relying on a thermodynamic integration of $E(V,T)$ to obtain the entropy for a complete EOS ($E(V,S)$). This approach works for large temperature and pressures (where the constant-entropy approximation of our approach would fail.) It requires simulating isochores down to absolute zero, where the thermodynamic inte-

grand has a (removable) singularity, perhaps creating some numerical difficulties. Their primary consideration in producing this EOS is, like us, continuum simulations involving shock waves, to high pressures, including phase changes. A similar approach is taken by Chentsov and Levashov (2012), who use (for a liquid) a sampling in density and temperature before solving an ODE to find the internal energy as a function of density and entropy. For a fluid phase, this amounts to solving

$$\left(\frac{\partial T}{\partial V}\right)_S = -\left(\frac{\partial S}{\partial V}\right)_T \Bigg/ \left(\frac{\partial S}{\partial T}\right)_V \tag{4.1}$$

$$= -\left(\frac{\partial P}{\partial T}\right)_V \Bigg/ \left(\frac{1}{T}\frac{\partial E}{\partial T}\right)_V \tag{4.2}$$

$$= \frac{-T}{\left(\frac{\partial E}{\partial T}\right)_V \Big/ \left(\frac{\partial P}{\partial T}\right)_V}. \tag{4.3}$$

but a similar procedure could also be used for a solid phase.

4.3. Isentropic Deformation

As discussed in chapter 3, we express deformations using the Finger tensor, $\boldsymbol{G}$. The Cauchy stress can be computed as a derivative of internal energy at constant entropy (eq. (3.20)),

$$\sigma_{ij} = \frac{\rho_0}{\det \boldsymbol{F}} F_{ik} \left(\frac{\partial E}{\partial G_{lm}}\right)_S \frac{\partial G_{lm}}{\partial F_{jk}}. \tag{4.4}$$

This suggests using $E(\boldsymbol{G}, S)$ as the most convenient form of EOS. Since we have no other need for the entropy, we do not calculate it, and instead define a reference energy E_0 as the total energy (kinetic and electronic) of the material when it is adiabatically brought to the reference configuration from the deformed configuration. This depends only on the entropy, so can be used to label the isentropes. The EOS can then be expressed as $E(\boldsymbol{G}, E_0)$, and

$$\left(\frac{\partial E}{\partial \boldsymbol{G}}\right)_{E_0} = \left(\frac{\partial E}{\partial \boldsymbol{G}}\right)_S. \tag{4.5}$$

We show that it is feasible to extract states along an isentrope directly with molecular dynamics in a slowly deforming box.

4.3.1. Slow deformation isentropes with molecular dynamics simulation

Molecular dynamics can be formulated as a Hamiltonian system, with generalized positions $\boldsymbol{q} = (\boldsymbol{q}^\alpha)_{\alpha=1}^N$ and momenta $\boldsymbol{p} = (\boldsymbol{q}^\alpha)_{\alpha=1}^N$, where α is an index over the atoms and $\boldsymbol{q}$ and $\boldsymbol{p}$ correspond to the absolute position and physical momenta of the particles. We work within periodic boundary conditions, specified by the matrix of box vectors $\boldsymbol{L} = (\boldsymbol{l}_1, \boldsymbol{l}_2, \boldsymbol{l}_3)$, and allow this to vary with time, $\boldsymbol{L} = \boldsymbol{L}(t)$. This means that the atoms interact via a periodic potential V, where for any i, and for any vector of integer coefficients, $\boldsymbol{a}$,

$$V(\boldsymbol{q}^1, \ldots, \boldsymbol{q}^{i-1}, \boldsymbol{q}^i + \boldsymbol{a}^\mathsf{T}\boldsymbol{L}(t), \boldsymbol{q}^{i+1}, \ldots, \boldsymbol{q}^N; \boldsymbol{L}(t)) = V(\boldsymbol{q}; \boldsymbol{L}(t)), \tag{4.6}$$

and a Hamiltonian involving these coordinates is

$$H = \sum_{\alpha=1}^{N} \frac{\boldsymbol{p}^\alpha \cdot \boldsymbol{p}^\alpha}{2m^\alpha} + V(\boldsymbol{q}; \boldsymbol{L}(t)). \tag{4.7}$$

The procedure depends on a slowly varied parameter in a Hamiltonian system preserving the entropy to first-order in the rate of variation of the parameter. To simplify the following argument, suppose that the Hamiltonian depends on only a single parameter L (instead of the matrix $\boldsymbol{L}$), which is varied slowly. The rate of increase of entropy of the system may be expressed as a series expansion in the rate of change of the parameter (see e.g. Landau and Lifshitz, 1980):

$$\frac{\mathrm{d}S}{\mathrm{d}t} = a_0 + a_1 \frac{\mathrm{d}L}{\mathrm{d}t} + a_2 \left(\frac{\mathrm{d}L}{\mathrm{d}t}\right)^2 + \ldots. \tag{4.8}$$

The coefficient a_0 is zero, since in thermodynamic equilibrium, the entropy will remain constant. The crucial observation is that a_1 is also zero, since if it were not, L could be varied so as to *decrease* the entropy.

The conclusion is that

$$\frac{\mathrm{d}S}{\mathrm{d}L} = \frac{\mathrm{d}S}{\mathrm{d}t} \Big/ \frac{\mathrm{d}L}{\mathrm{d}t} = a_2 \frac{\mathrm{d}L}{\mathrm{d}t}, \tag{4.9}$$

which may be made arbitrarily small by decreasing the rate of variation of the parameter L.

For the convenient evaluation of the potential, we wish to keep track of the location of the atoms within a single periodic cell with $0 < (\boldsymbol{L}^{-1}\boldsymbol{q}^\alpha)_i < 1$ for each atom α and each spatial coordinate i, and have these interact with the atoms in a number of periodic

image cells within a finite cutoff. Changing the box vectors may cause an atom to leave this central simulation cell. It could also leave this cell due to its motion taking it across the boundary of the cell. It is convenient, and common in molecular dynamics codes, to *remap* the position of the atom back into the central simulation cell. Note that the time evolution of the Hamiltonian system is expressed only in terms of the *unremapped* positions, so any remapping must be done only for the purpose of evaluating the potential.

Alternatively, we can use a canonical transformation to adjust the atom's velocity when mapping its position back into the central periodic cell, while preserving the Hamiltonian structure (and therefore eq. (4.9)). A canonical transformation is a change of coordinates from $(\boldsymbol{q}, \boldsymbol{p})$ to $(\boldsymbol{Q}, \boldsymbol{P})$ in the original Hamiltonian H that results in another Hamiltonian $\hat{H}$, and therefore the same dynamics. It can be shown that a generating function of the form $G(\boldsymbol{q}, \boldsymbol{P}, t)$ results in a canonical transform via the implicit relations

$$\boldsymbol{p} = \frac{\partial G}{\partial \boldsymbol{q}} \tag{4.10}$$

$$\boldsymbol{Q} = \frac{\partial G}{\partial \boldsymbol{P}} \tag{4.11}$$

$$\hat{H} = H + \frac{\partial G}{\partial t}. \tag{4.12}$$

Consider a generating function

$$G(\boldsymbol{q}, \boldsymbol{P}, t) = (\boldsymbol{q} - \boldsymbol{R}(t)) \cdot \boldsymbol{P}, \tag{4.13}$$

where $\boldsymbol{R}(t)$ is some time dependent translation vector of the atoms. For this choice,

$$\boldsymbol{p} = \boldsymbol{P} \tag{4.14}$$

$$\boldsymbol{Q} = \boldsymbol{q} - \boldsymbol{R}(t) \tag{4.15}$$

$$\hat{H}(\boldsymbol{Q}, \boldsymbol{P}) = H(\boldsymbol{q}(\boldsymbol{Q}, \boldsymbol{P}),\, \boldsymbol{p}(\boldsymbol{Q}, \boldsymbol{P})) - \dot{\boldsymbol{R}} \cdot \boldsymbol{P}. \tag{4.16}$$

The new momenta are the same as the original. Choosing

$$\boldsymbol{R}(t) = (\boldsymbol{a}_1^{\mathsf{T}} \boldsymbol{L}(t), \ldots, \boldsymbol{a}_N^{\mathsf{T}} \boldsymbol{L}(t)), \tag{4.17}$$

with $\boldsymbol{a}_i$ a vector of integers for every $1 < i < N$, the new positions are also the same as the original positions, but with each atom translated by a lattice vector. This can always be

done to result in $\boldsymbol{Q}^\alpha$ being inside the central periodic cell. Since the potential is periodic,

$$\hat{H}(\boldsymbol{Q},\boldsymbol{P}) = H(\boldsymbol{Q},\boldsymbol{P}) - \dot{\boldsymbol{R}}\cdot\boldsymbol{P}. \tag{4.18}$$

The system evolves in time via Hamilton's equations

$$\dot{\boldsymbol{Q}}^\alpha = \frac{\boldsymbol{P}^\alpha}{m^\alpha} - \dot{\boldsymbol{R}}(t) \tag{4.19}$$

$$\dot{\boldsymbol{P}}^\alpha = -\nabla_{Q^\alpha} V(\boldsymbol{Q};\boldsymbol{L}(t)). \tag{4.20}$$

The box deformation rate $\dot{\boldsymbol{L}}$ (and therefore $\dot{\boldsymbol{R}}$) is time dependent. The kinetic energy density and kinetic stress tensor are only meaningful when $\dot{\boldsymbol{R}} = 0$: these quantities should be computed in terms of the atomic velocities under zero deformation rate, $\boldsymbol{P}^\alpha/m^\alpha$.

For a constant deformation rate, adjusting the velocity by the deformation rate as described above, but not adjusting for this when computing the energy and stress results in an energy error proportional to the deformation rate, the same order as from eq. (4.9). For a solid undergoing elastic deformation, the number of components of atomic position outside of the central box is small, and so too is the deformation rate of the cell, so that the error in the kinetic energy is small.

4.3.2. Momentum scaling approach

Other approaches based on transforming a Hamiltonian are possible. In particular, the following alternative procedure. The box vectors $\boldsymbol{L}$ are varied slowly as before. Instead of leaving the absolute atomic positions unchanged (along with the remapping procedure if needed), we can alternatively *scale* the atomic positions along with the box. By Liouville's theorem, adjusting the position alone in general results in a change in phase space volume, which is proportional to the entropy. The correct way of achieving this is via another canonical transformation, with generating function

$$G(\boldsymbol{P},\boldsymbol{q}) = -\sum_{\alpha,i,j} F_{ij} q_j^\alpha P_i^\alpha, \tag{4.21}$$

leading to

$$Q^\alpha = -\frac{\partial G}{\partial \boldsymbol{P}^\alpha} = \boldsymbol{F}\boldsymbol{q}^\alpha \tag{4.22}$$

$$\boldsymbol{p}^\alpha = -\frac{\partial G}{\partial \boldsymbol{q}^\alpha} = \boldsymbol{F}^\mathsf{T}\boldsymbol{P}^\alpha \tag{4.23}$$

$$\boldsymbol{P}^\alpha = \boldsymbol{F}^{-\mathsf{T}}\boldsymbol{p}^\alpha. \tag{4.24}$$

4.3.3. Numerical test cases

As discussed above, the entropy change is asymptotically linear in the box deformation rate. In this section, we find the rate of change of entropy as a function of the deformation rate for several silicon systems, in order to verify that we obtain linear convergence in the entropy change when performing a numerical simulation of the deformation, and to demonstrate that for the system of interest the error is acceptable at practical deformation rates.

If the compression is indeed isentropic, $\mathrm{d}S = 0$ for the whole of any deformation path path, and so, by the second law of thermodynamics

$$\mathrm{d}E = \boldsymbol{\sigma}\boldsymbol{F}^{-\mathsf{T}} : \mathrm{d}\boldsymbol{F}. \tag{4.25}$$

The expression above holds in equilibrium, for ensemble averages of the quantities in it. To use it, after each small increment in box length, the stress should be allowed to relax to its equilibrium value by running the simulation unperturbed by box deformations.

Since we have microscopic control over the particles in our simulation, we may inadvertently cause the entropy to decrease (as a Maxwell's Demon), so we cannot necessarily rely on $\mathrm{d}S \geq 0$ for any conceivable process. For this reason, it is not enough to obtain convergence in just the total energy change (ΔE), since it might be the case that (for example) the entropy decreases under compression only to increase again on the expansion, for zero overall change.

Rearranging eq. (3.16), the entropy change is given by

$$\mathrm{d}S = \frac{\mathrm{d}E}{T} - \frac{1}{T}(\boldsymbol{\sigma}\boldsymbol{F}^{-\mathsf{T}}) : \mathrm{d}\boldsymbol{F}, \tag{4.26}$$

where we write ':' to mean double tensor contraction. This is true provided that these quantities are understood to mean ensemble averages.

Test case: uniaxial compression of an elastic solid

To demonstrate the consistency and numerical accuracy of the procedure, we perform the following test. We consider uniaxial compression of a 64 atom box of Tersoff silicon in cubic diamond structure initially at 300 K and zero stress. The final strain has 0.9 relative volume. This compression is then undone at the same rate until it returns to the starting volume. Numerically, this is done by changing the simulation box shape in small increments, with a relaxation time between each increment of 100 fs in this example, during which the system is evolved using Verlet integration, but is unperturbed by box changes.

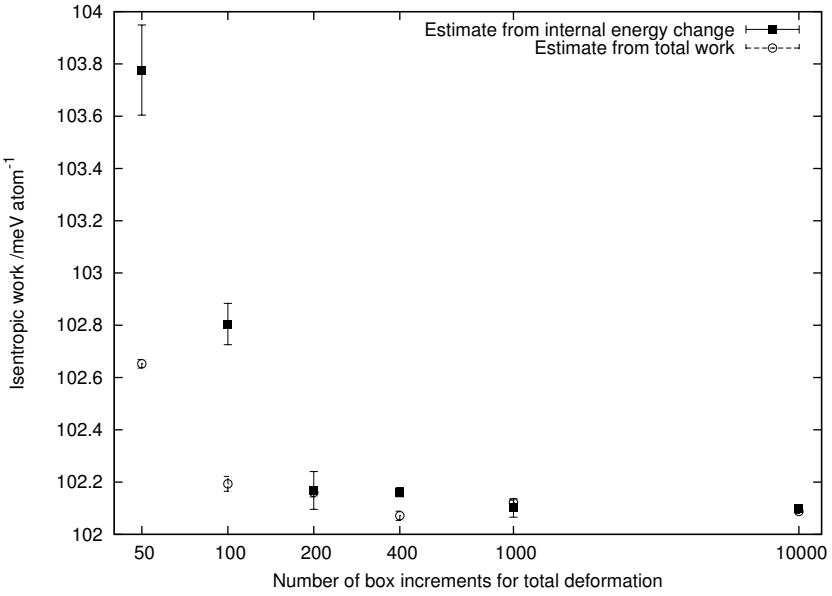

Figure 4.1.: Convergence of the estimated isentropic energy change of Tersoff silicon initially at 300 K and zero pressure, under uniaxial elastic compression to 0.9 relative volume, with the total time of the deformation. 100 fs was left between each box increment to allow the system to relax to equilibrium.

Figure 4.1 shows the convergence in the estimate of the energy of an isentropic uniaxial elastic compression. Correspondingly, fig. 4.2 shows the convergence of the internal energy change and total work performed of a complete compression/expansion cycle. It is possible to perform the deformation using 50 box increments (corresponding to a deformation time of 5 ps) to obtain an energy difference relative to the peak of 0.5%.

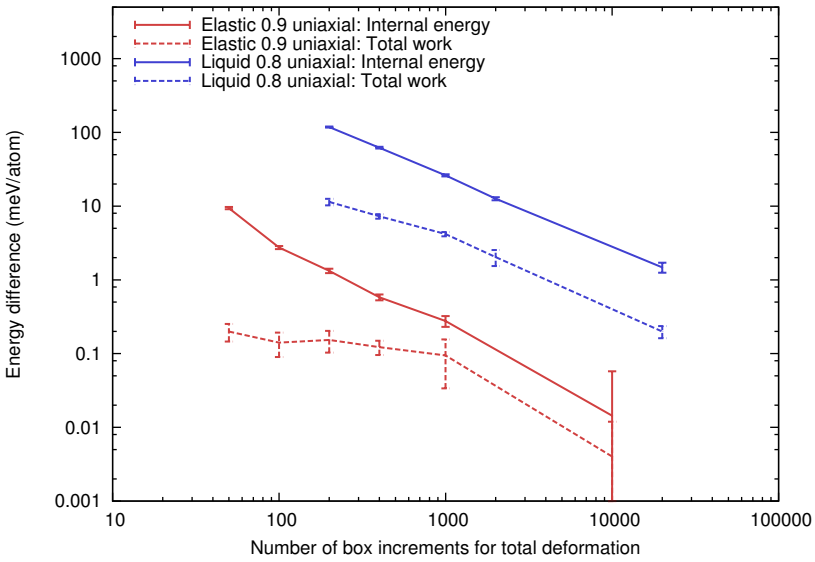

Figure 4.2.: The change in ensemble-averaged total energy and work done by the procedure described in the text, after a deformation to a strained state and a release to the starting state. For an isentropic process, this value should be zero. In the elastic case the converged value of the energy difference at maximum strain was 102 meV/atom, and in the liquid case it was 145 meV/atom.

Test case: Liquid Tersoff Silicon, 2500 K

As an additional test, starting from an undeformed configuration (which in this case is well-equilibrated liquid Tersoff silicon at 2500 K), the simulation box is uniaxially deformed to a volume compression of 0.8 at a constant rate. This compression is then undone at the same rate until it returns to the starting volume, as in the previous test. As before, this is done by changing the simulation box shape in small increments, with a relaxation time between increments of 100 fs. Note that if we are genuinely interested in isentropic compression of a fluid, it would be more efficient to use an isotropic compression, to avoid the artificial relaxation time to a hydrostatic stress configuration. However, we consider this test case precisely because the relaxation time effect provides a difficult situation for convergence. Note that long-time energy conservation from Verlet integration under these conditions demanded a time step of 0.1 fs.

In general, the observations from this simulation are very similar to those in the previous section, but a shorter relaxation time and overall simulation time is required for a comparable accuracy in the final estimated isentropic energy.

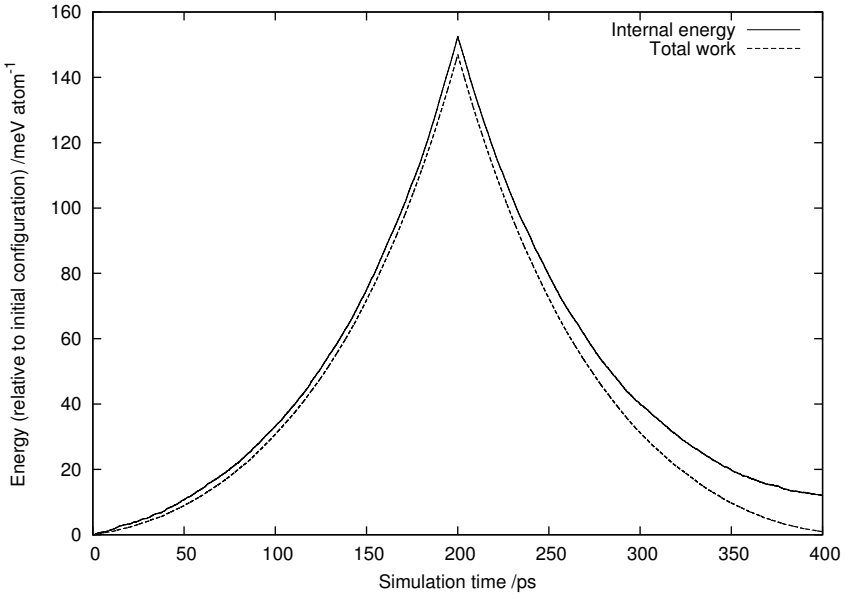

Figure 4.3.: Total energy and total work performed on the simulation box (computed as pressure integral—see text) for a simulation with deformation time 200 ps. The relaxation time between each volume change was 100 fs.

For this case, the relaxation of the virial stress from a series of successive small increments is shown in figs. 4.3 and 4.4. In this case, the overall rate of change is consistent with taking 2000 steps for the full deformation to 0.8 volume compression. It can be seen that equilibration after the step happens over a timescale $\sim$100 fs.

Using *shorter* relaxation times should also eventually converge, since each box increment causes a smaller initial perturbation in the virial stress, although the meaning of the virial stress when it has not yet converged is unclear.

The entropy rate is shown in fig. 4.5, which is roughly constant for this particular deformation process. The cause of the deviation from linearity for the slowest deformations is not clear.

Figure 4.2 shows the effect of the speed of the deformation on the total energy change and the corresponding integral of the virial stress, when 100 fs are left between successive steps. Both can be seen to converge to zero, although to obtain a difference of 1 meV/atom (0.5 % of the total), a very slow deformation is needed (2000 ps $\equiv 2 \times 10^7$ time steps).

To estimate the isentropic energy change from this figure, half the total energy change is subtracted from the internal energy at maximum deformation. Another estimate is obtained from using the total work instead of the internal energy. Both of these estimates

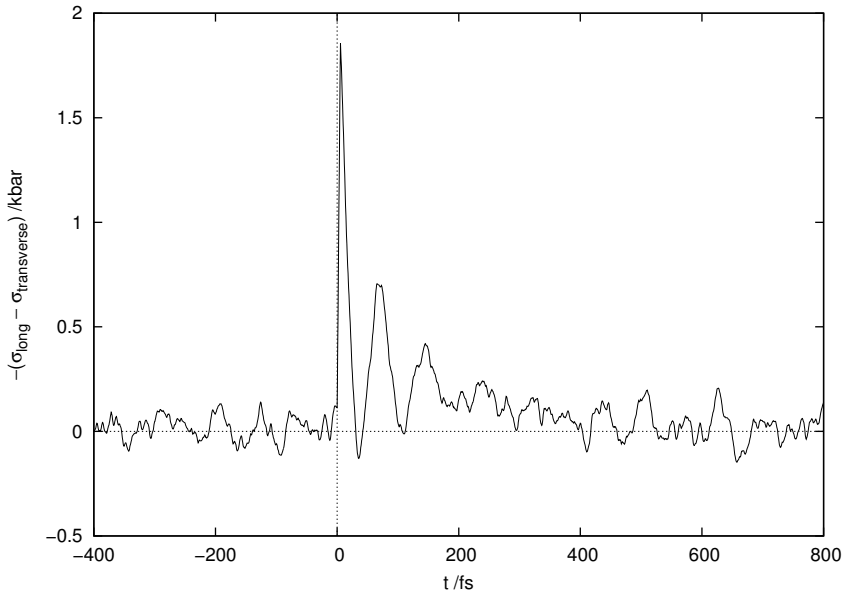

Figure 4.4.: Relaxation time of stress in a fluid subject to a uniaxial compression (to a relative volume of 0.999) over ten timesteps, starting at $t = 0$, according to the procedure described in the text. The figure shows the difference between the (negative) stress parallel to the deformation $\sigma_{\text{long}} = \sigma_{zz}$ and transverse to it $\sigma_{\text{transverse}} = \frac{1}{2}(\sigma_{xx} + \sigma_{yy})$. This is computed as an ensemble average over 5000 independent trajectories.

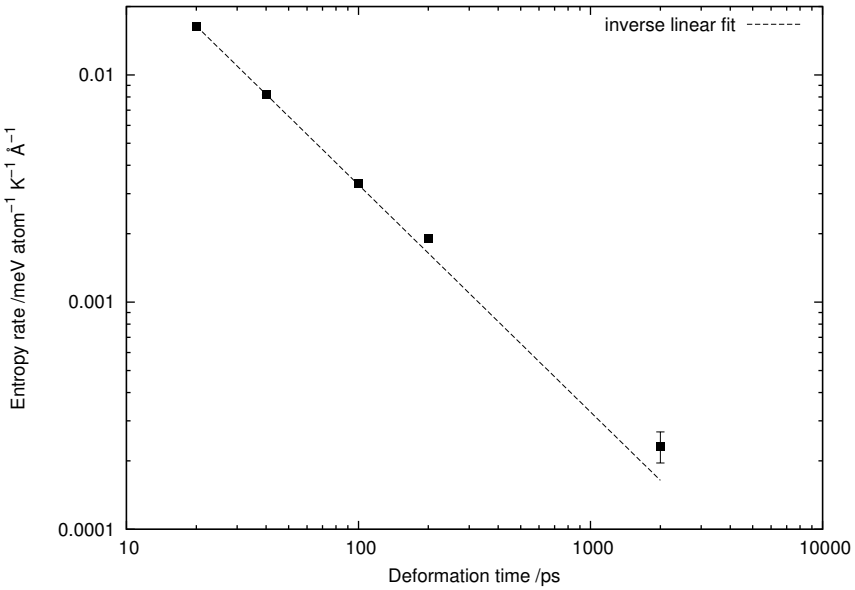

Figure 4.5.: The rate of increase of entropy, expressed per absolute length change in lattice parameter. Also shown is an inverse linear fit to the data. Error bars are only visible for the final point: the others are too small to be displayed on this plot. The relaxation time between each increment in volume was always 100 fs.

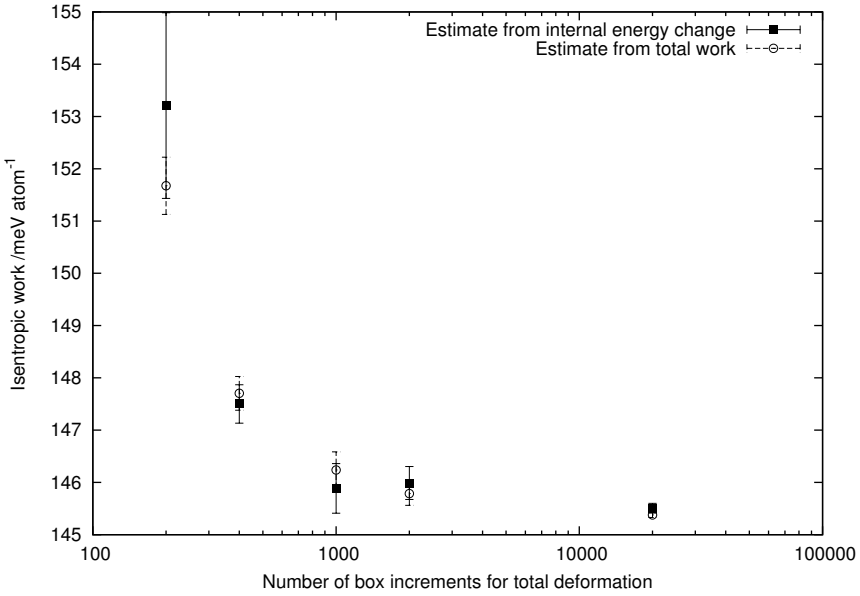

Figure 4.6.: Convergence of the estimated energy change under isentropic compression of the fluid to 0.8 relative volume, using the method above at various rates, always allowing 100 fs between each box increment. The box was deformed uniaxially, although the stress was isotropic after the relaxation time used.

are shown in fig. 4.6. Note the different scale to fig. 4.6, with generally a larger number of box increments taken for the deformation in the liquid case.

Shock loading to the same relative volume causes an internal energy increase of 211 meV/atom. This is actually slightly less than the fastest deformation used in the above procedure here (20 ps), which causes an internal energy increase of 214 meV/atom. (Note that the estimated isentropic work in this case is 153 meV/atom, which is still close to the converged value of 146 meV/atom).

4.3.4. Isentropic deformation summary—expected accuracy of the procedure

The relaxation time of 100 fs used in the previous examples results in long deformation times, suitable for empirical potentials, but too long to be practical for DFT simulation. However, we can rely on the observation that reducing the equilibration time between steps to 10 fs does not introduce much extra error for relatively rapid deformations. It does seem to make the method ultimately converge more slowly, for deformations over longer times, but in practice we can apply deformations to stresses of tens of GPa over ∼1 ps on 64 atom cells and achieve relative errors in the total energy of the strain of around 1%, and error in the strain energy computed by integrating the work of 0.2%. We observe that the behaviour is not strongly dependent on system size.

4.4. Sampling and Reconstruction

Molecular dynamics simulations allow the thermoelastic energy surface to be sampled at a number of discrete points. These points must be interpolated to evaluate the energy of a particular arbitrary deformation at a given temperature. This requires the choice of a suitable interpolation method, and a procedure for choosing the sampling points.

We use Gaussian process regression for the interpolation (MacKay, 2003; Rasmussen and Williams, 2006), motivated by a number of considerations. First, its ability to handle multi-dimensional data. Second, the fact that (with a suitable covariance function) the interpolated function is smooth: we require the interpolant to have continuous second derivatives, since these appear in expressions for the wave speeds, and otherwise unphysical wave-splitting occurs. Third, it can incorporate derivative observations (obtained from the pressure) into the fit, and is able to conveniently predict derivatives of the interpolated function (and therefore pressures). At additional expense, we can obtain an estimate of the error in the prediction. We do not use this directly, but it is useful for validation purposes.

The sampling is performed by choosing the Finger tensor $\boldsymbol{G}$ (eq. (3.18)), and the uniformly at random over a problem-specific domain of interest, before converting this to a deformation gradient $\boldsymbol{F}$ (by a Cholesky decomposition), and thence to a target lattice ($\boldsymbol{FL}$). Random sampling in low dimensions suffers from the drawback of unevenly distributed points. There is a high probability of choosing two points that are too close to one another, compared with ideally-packed points. In seven dimensions, random sampling does not suffer as badly from this drawback, and the probability of two points being "too close" (that is, much closer than two points picked as the centres of neighbouring spheres from an optimal packing) is quite small (Conway and Sloane, 1995).

The sampling domain can be chosen to generously include the range over which the deviatoric part of the strain[1] is expected to be less than or equal to the yield criterion according to, for example, a continuum plasticity model, and with the isotropic part of the strain less than some bound. For the EOS given here, we chose simply to sample each component uniformly and independently over a range of $[0.9, 1.1]$ for the diagonal components and $[-0.3, 0.3]$ for the off-diagonal components. The initial energy E_0 was sampled by setting an initial temperature of between 0K and 900K.

The reference energy is the dominant contribution to the energy. To improve the fitting process, the energy is partitioned as

$$E(\boldsymbol{G}, E_0) = E_0 + E'(\boldsymbol{G}, E_0), \tag{4.27}$$

where E_0 is the energy of the undeformed configuration at the initial temperature, and E' is defined as the difference between the total energy and E_0 for a given $\boldsymbol{G}$. The interpolation is then applied to E'.

4.4.1. Details of the Gaussian Process

An overview of Gaussian process regression is given in appendix B.

We use a squared exponential covariance function eq. (B.3), between energy observations taken at inputs consisting of pairs of Finger tensor and initial energy, $\boldsymbol{x}^{(1)} = (\boldsymbol{G}^{(1)}, E_0^{(1)})$ and $\boldsymbol{x}^{(2)} = (\boldsymbol{G}^{(2)}, E_0^{(2)})$. This covariance function is:

$$C(\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}) = \sigma^2 \exp\left(-\frac{1}{2}\left[\sum_{i,j} \frac{(G_{ij}^{(1)} - G_{ij}^{(2)})^2}{r_{G\,ij}^2} + \frac{(E_0^{(1)} - E_0^{(2)})^2}{r_{E_0}^2}\right]\right) + \nu^2 \delta_{\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}}. \tag{4.28}$$

[1] The deviatoric part of a rank-two tensor ϵ is the tensor with its isotropic part subtracted: $\epsilon_{ij}^{\mathrm{Dev}} = \epsilon_{ij} - \frac{1}{3}\epsilon_{kk}\delta_{ij}$.

Covariances between value and derivative observations, and between two derivative observations, are the corresponding derivatives of this function.

The interpretation of the hyperparameters in the squared exponential covariance function is as follows: σ sets the overall scale of the inferred function, ν represents position-independent Gaussian noise in the outcomes that is independent of the inputs, and r_{ij} is the characteristic length over which the function values become decorrelated as the distance between the ijth component of the input varies, which may differ between components of the input. Separate noise hyperparameters are used for value and derivative observations.

4.4.2. Inverting the equation

When solving the hyperelastic equations (eqs. (3.30), (3.31) and (3.39)) numerically, E and $\boldsymbol{F}$ are known at a given timestep. In order to evaluate the flux, so as to determine these quantities at the next timestep, the stress is needed, and since this is a function of $\boldsymbol{F}$ and S (or equivalently, E_0), the equation of state must be inverted to find $E_0(E, \boldsymbol{F})$. We do this with a bounded Newton–Raphson method. The acceptable interval of E_0 is set as the bounds of the training data. If a Newton–Raphson step would take the iterated solution outside of this interval, a step of bisection search is used instead. This guarantees convergence to a solution (if it exists), and provides quadratic convergence close to it. If no solution exists an error condition can be triggered.

The error from the reconstruction is shown in fig. 4.7, from an EOS computed from molecular-dynamics trajectories from an empirical potential, allowing many more to be sampled. For the databases where gradient information is used, all six components of the gradient are included for one-sixth of the points in the database. It is observed in the figure that this is always beneficial, but much more so for small databases, where it can reduce the error by a factor of four.

The preceding discussion made no assumption about the crystal symmetry of the material under study. If the material has a symmetry beyond triclinic, we can make use of this knowledge by transforming a training or sampling point to the corresponding point in a single inequivalent region under the symmetry. This is primarily to remove from the fitting process the freedom to produce slightly different values for identical configurations, but also allows the domain to be sampled more densely for a given number of training points. It does not reduce the dimensionality of the space, however, but does save a constant factor (a factor of eight in the case of cubic symmetry). Miller (2004b) gives a minimal set of invariants for each crystal system.

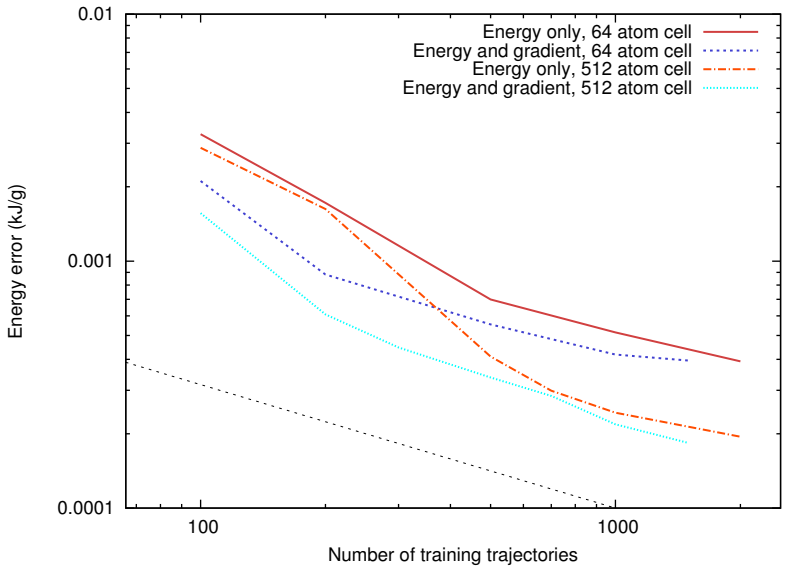

Figure 4.7.: Errors from the reconstruction. The energy scale for typical deformations considered here is 0.1 kJ/g. The errors are computed at independent test points not included in the training databases. These points (unlike those in the training sets) are themselves ensemble averages. For comparison, the faint dotted black line indicates the expected inverse square-root scaling.

4.5. Equation of state for DFT silicon

We produce an EOS for silicon as described by DFT-MD, and use it in a numerical simulation of a weak shock wave interacting with a change in density caused by a temperature gradient.

The EOS presented here was calculated with the SIESTA method and implementation of Density Functional Theory (Soler et al., 2002), using the GGA functional of Perdew et al. (1996).

The core electrons were described with a Troullier–Martins norm-conserving pseudopotential (Troullier and Martins, 1991) with a matching radius in each angular momentum channel of 1.89 a_0. The valence electrons were described with a basis of numerical atomic orbitals (Junquera et al., 2001) of double-ζ polarized type (representing 13 orbitals per atom). The basis was generated by fixing the longest orbital cutoffs at 7.0 a_0 and variationally optimising the other parameters in bulk diamond-phase silicon—the final basis parameters are given in table F.1.

The mesh used for integrals in real-space was well converged at a grid cutoff of 100 Ry. A 2^3 Monkhorst–Pack grid of points was used on the 64 atom simulations, to give an

effective cutoff length of 11 Å. We do not consider the thermal electronic contribution, which is expected to make a small contribution at the temperatures considered. An artificial electronic temperature of 300 K was used throughout.

Verlet integration modified as described above was used to follow an isentrope, with a timestep of 1 fs and forces from DFT. In all, 480 separate deformations of 64 atom silicon were performed, with each run for 2 ps. For each deformation, an initial configuration close to the target temperature was obtained by equilibration of the Tersoff potential on the intended undeformed state, before switching to DFT forces and continuing the integration. The DFT forces are integrated for 250 fs before starting the deformation, and for 250 fs afterwards, to obtain an averaged final temperature.

Table 4.1.: Maximum-likelihood hyperparameters for the Gaussian process equation of state for silicon. The noise hyperparameters for gradient observations are the same for each diagonal and off-diagonal component. There are no observations of the derivative with respect to E_0, so the corresponding noise term is omitted. Characteristic lengths $r_{G_{ij}}$ are likewise the same for each diagonal component and for each off-diagonal component.

ν_E (Ry)	$\nu_{dE/dG_{11}}$ (Ry)	$\nu_{dE/dG_{12}}$ (Ry)	σ (Ry)	$r_{G_{11}}$	$r_{G_{12}}$	r_{E_0} (Ry)
7.45×10^{-5}	0.749	0.330	38.2	0.49	0.36	142

After maximum-likelihood optimization of a Gaussian process containing 960 training points (all 480 energies and the full gradient for one-sixth of these points), the hyperparameters were as shown in table 4.1. Some curves for various paths through the EOS generated this way are shown in fig. 4.8.

Figure 4.9 shows the linear shear modulus at zero stress, as a function of temperature. It is systematically smaller than than experiment, given by Hull (1999) as 80.36 GPa at 0 K and 79.51 GPa at 300 K. Our value at zero temperature agrees with the GGA-DFT calculation of Lazar (2006). The temperature coefficient ($c_{44}^{-1}\mathrm{d}c_{44}/\mathrm{d}T$) of $-6.3 \times 10^{-5}\,\mathrm{K}^{-1}$ is roughly in line with the range of experimentally determined values quoted in the same reference, of between $-7.3 \times 10^{-5}\,\mathrm{K}^{-1}$ and $-1.0 \times 10^{-4}\,\mathrm{K}^{-1}$.

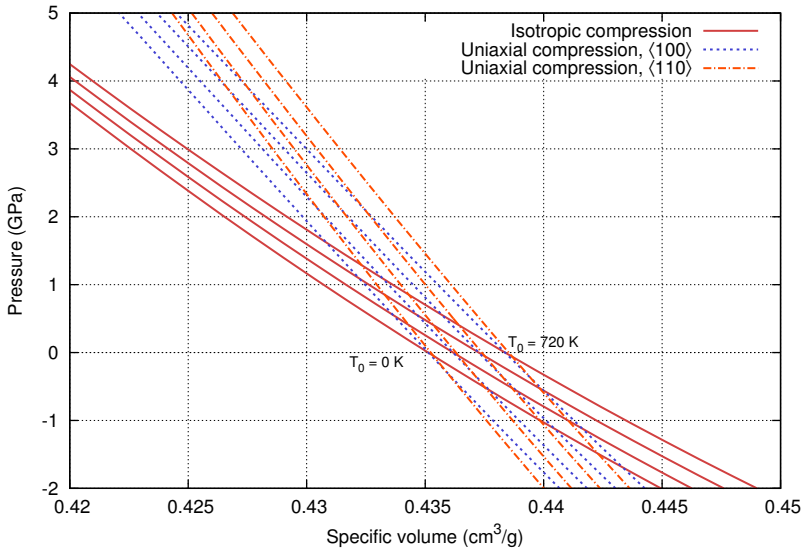

Figure 4.8.: Various compressive isentropes from the equation of state. Each isentrope has four curves, corresponding to initial temperatures of (left to right) 0, 240, 480 and 720 K.

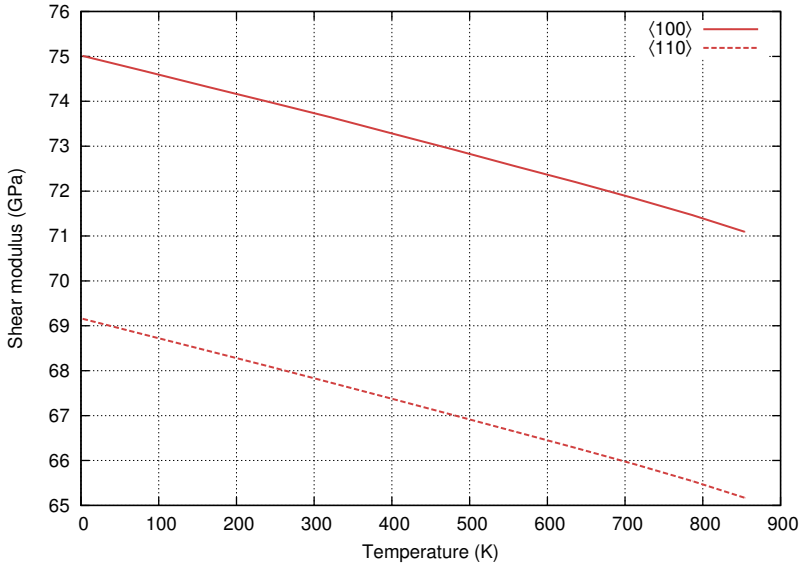

Figure 4.9.: The value labeled ⟨100⟩ gives the shear modulus ($G = c_{44}$) as a function of temperature. Also shown is the shear modulus for a rotated coordinate system, along ⟨110⟩.

4.6. Test problems

4.6.1. Seven-wave test case

This test case is a Riemann problem with appreciable jumps in all seven waves. It is a slightly modified version of the seven-wave test problem of Barton et al. (2009), for the silicon equation of state. The original parameters (applying to the Romenski equation of state for copper) are given as test case 1 in table 3.1: the only difference is that the entropy in this equation is mapped into an initial energy. The numerical solution is shown in fig. 4.10. The fastest rightmost wave is a shock, followed by two adjacent rarefaction waves (the split most clearly seen in u_2, and the shear components of the stress). Following this is the contact wave and two further adjacent rarefaction waves. The final wave is a very narrow rarefaction wave: the characteristic speed associated primarily with a longitudinal compression does not decrease much with a reduction in density.

4.6.2. Transverse impact test case

This test case is a modified version of those described by Miller (2004a) and Barton et al. (2009). It consists primarily of a longitudinal compressive wave, but with additional (relatively small) shear waves.

A spike in entropy and temperature is visible at the central contact, due to start-up error from a discontinuous initial condition. The overshoot most noticeable in the second wave of σ_{13} and u_3 is unphysical and due to the WENO reconstruction.

4.6.3. Thermal gradient

The initial density profile for the continuum elasticity simulation is generated from a temperature gradient from 710 K to 80 K over 1 cm by relaxing the stress to zero in each computational cell of the domain. A weak shock wave, with post-shock particle velocity of $0.12\,\mathrm{km\,s^{-1}}$ and longitudinal stress of $-2.25\,\mathrm{GPa}$, is maintained by the boundary condition on the left-hand side of the domain.

The density and stress profiles are shown in fig. 4.12 for times before, during and after the interaction with the density gradient. Most of the energy of the shock is transmitted, but a small reflected wave can be seen in this figure. The final stress state increases in magnitude slightly. The speed of the wave increases from $8.29\,\mathrm{km\,s^{-1}}$ to $8.46\,\mathrm{km\,s^{-1}}$, visible in fig. 4.13.

Figure 4.10.: Numerical solution to the seven-wave test case, at $t = 1\,\mu s$.

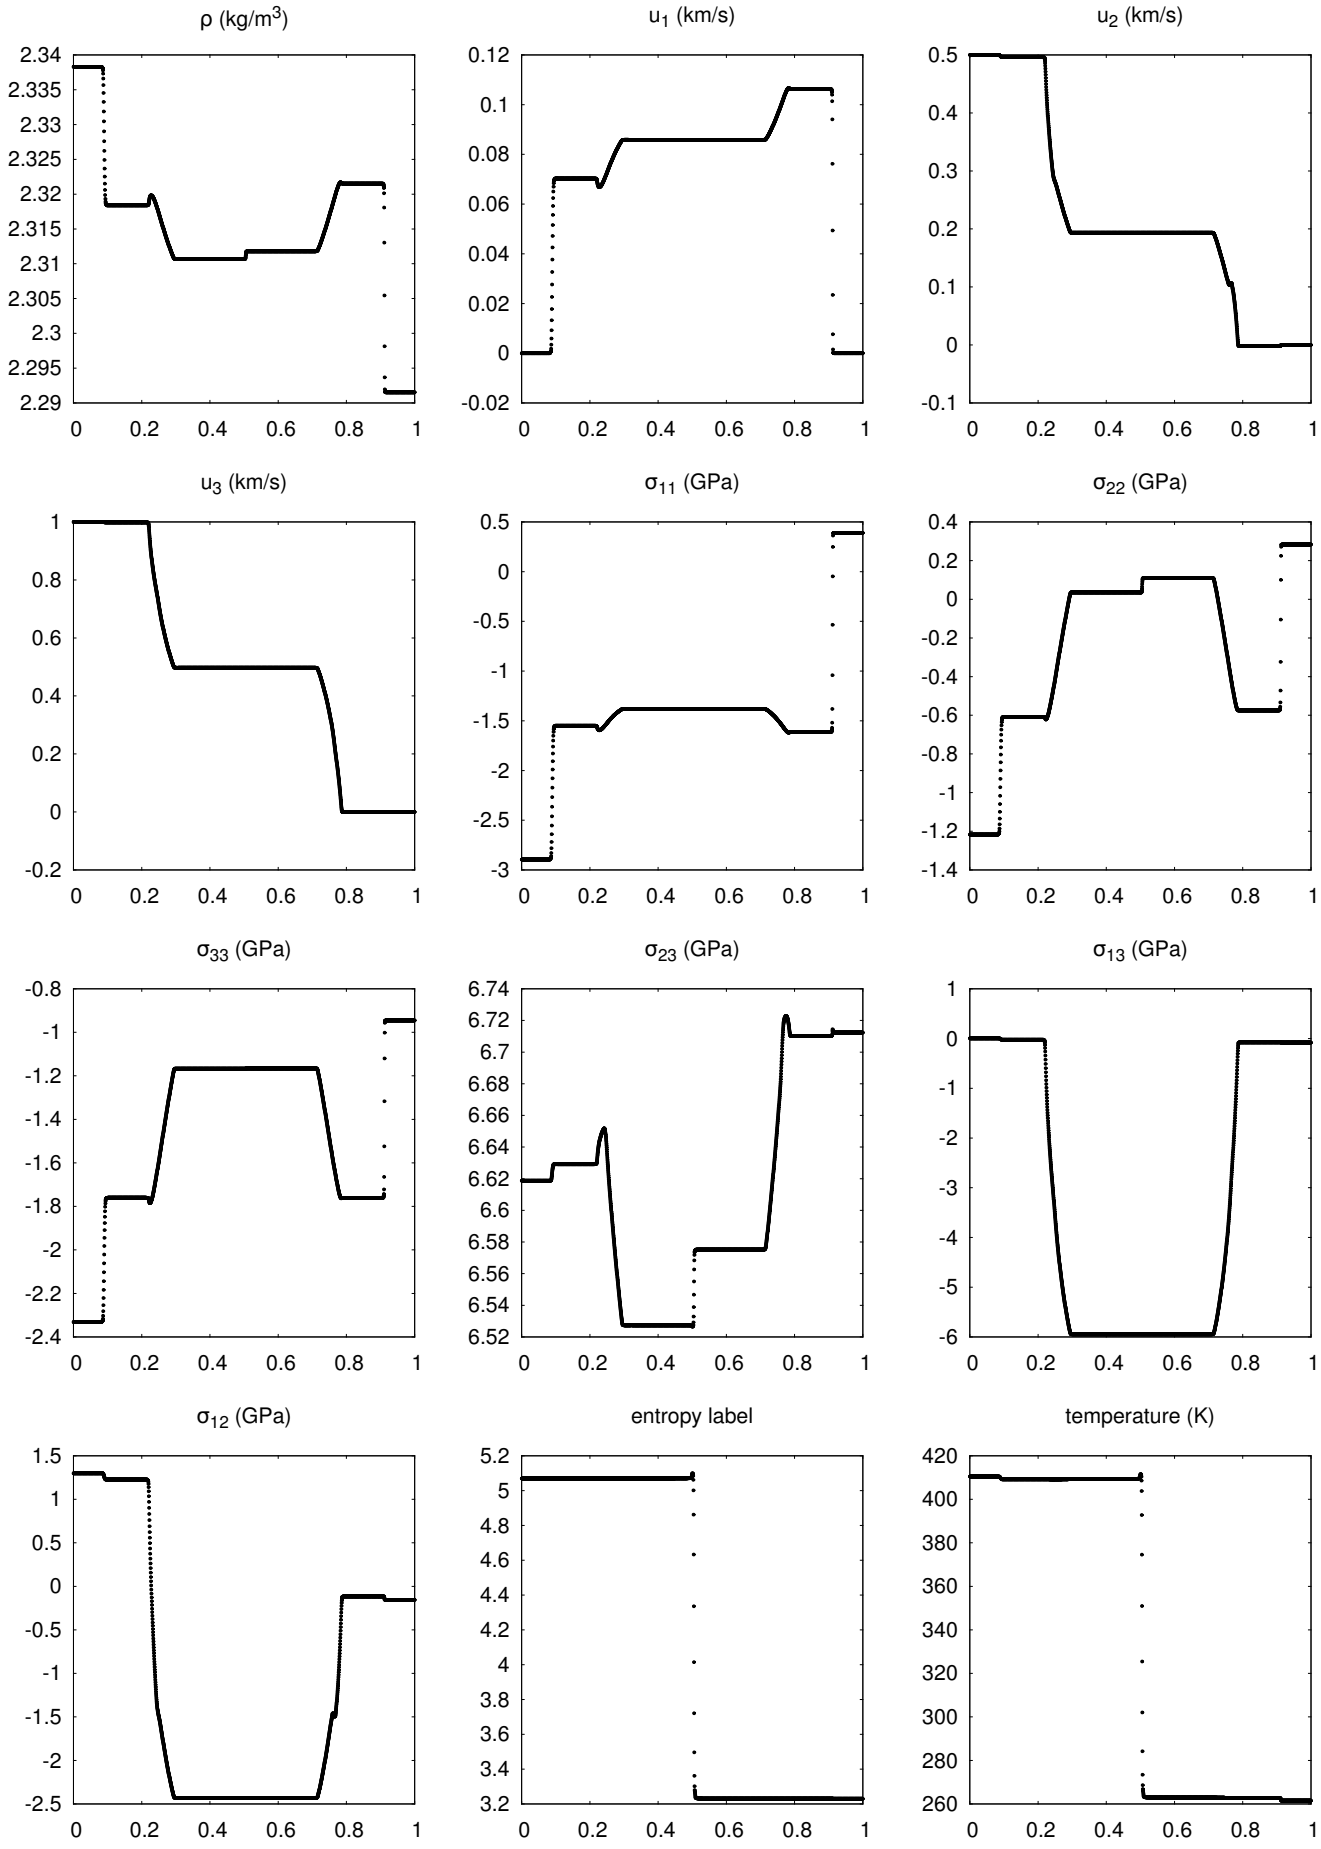

Figure 4.11.: Numerical solution to the transverse impact test case, at $t = 1\,\mu s$.

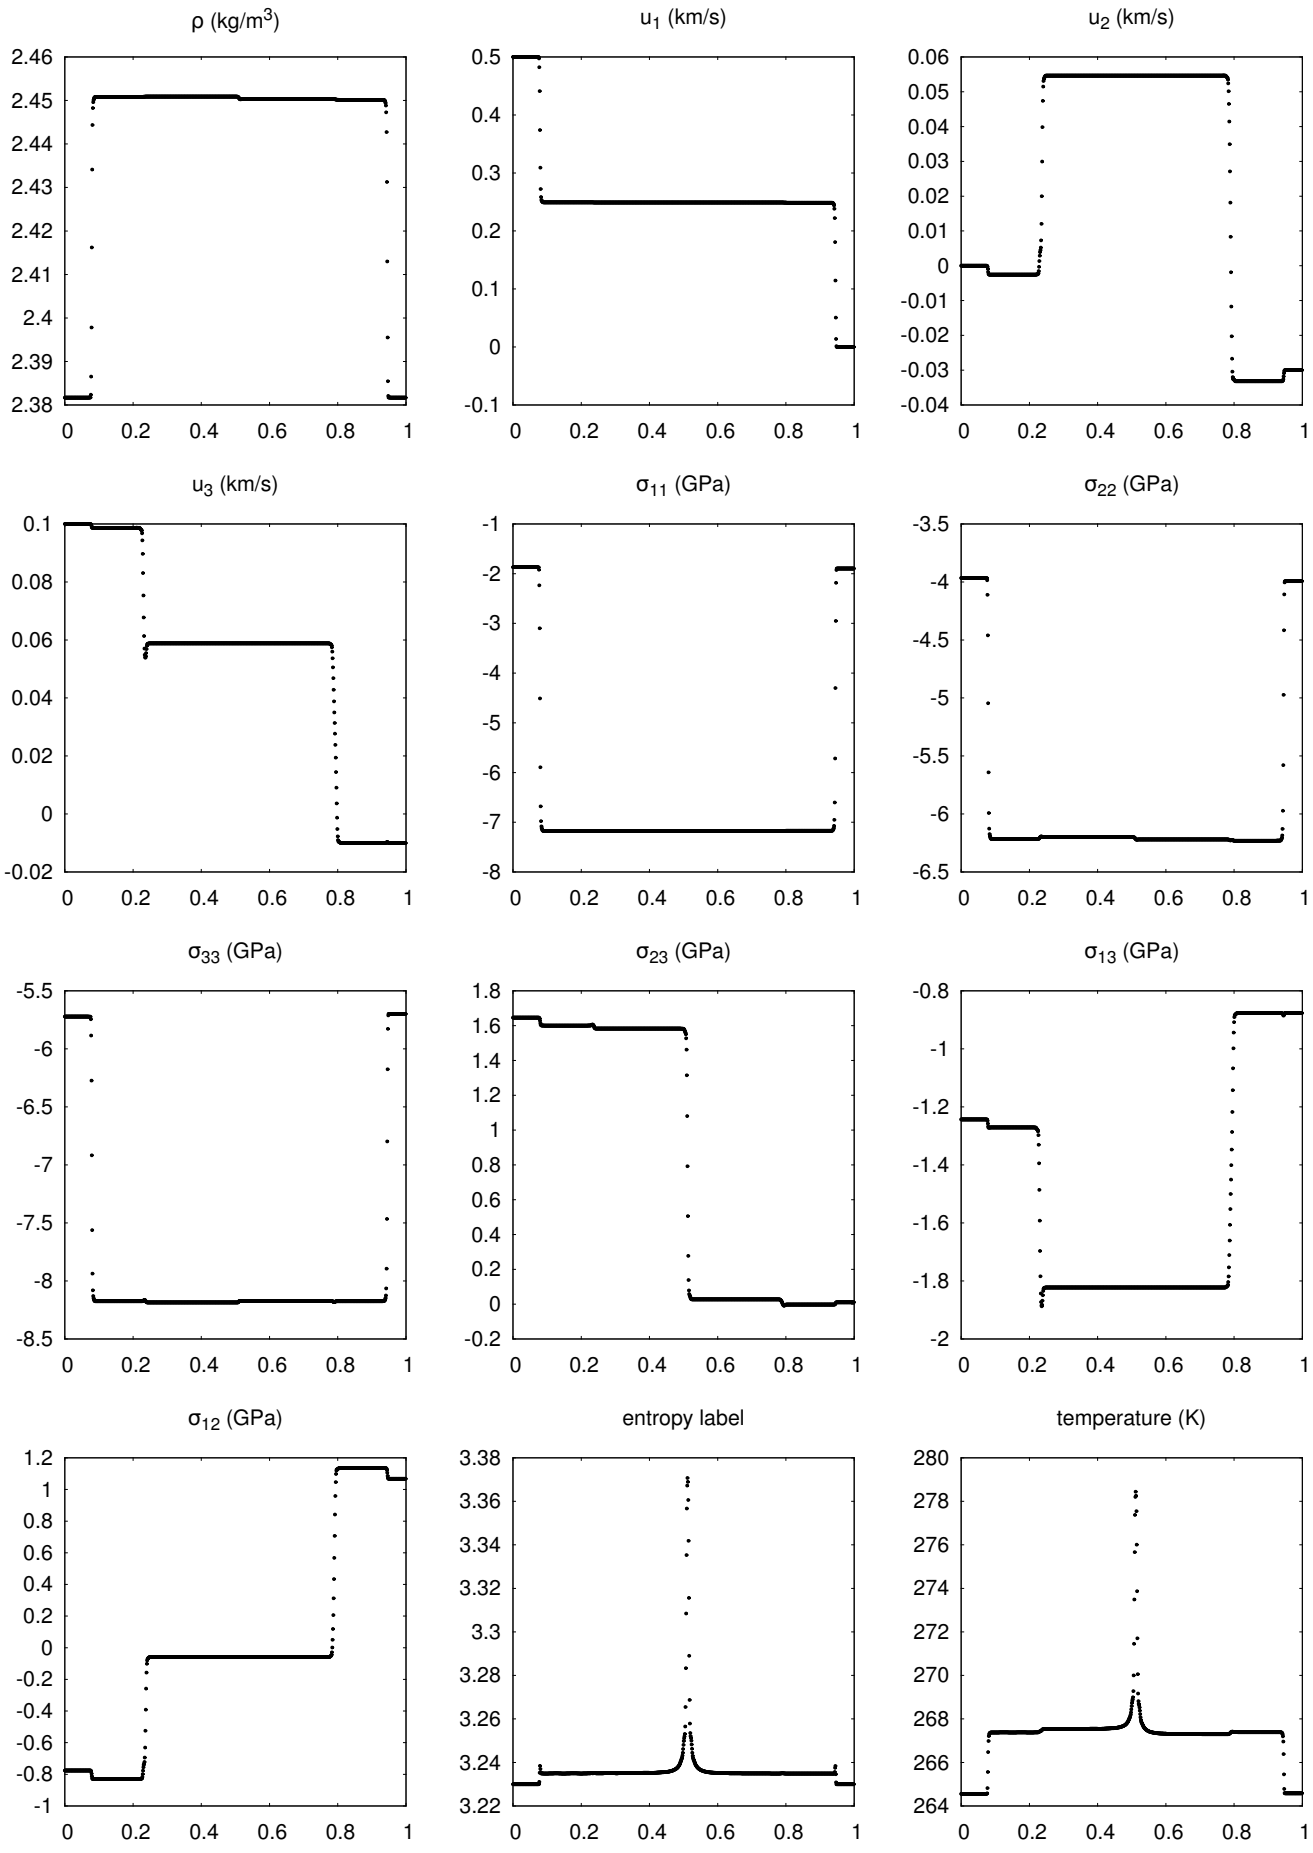

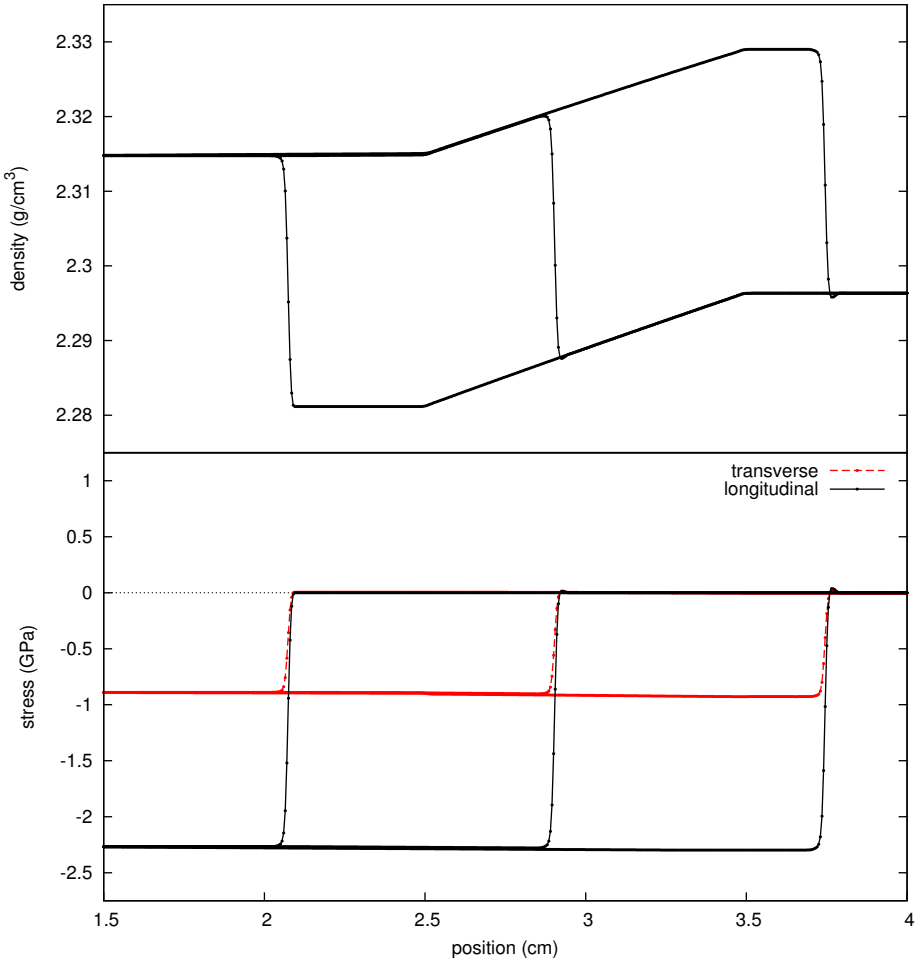

Figure 4.12.: Wave profiles in density and stress for the case of a shock impinging on a density gradient described in the text, at 1 μs intervals of simulation time.

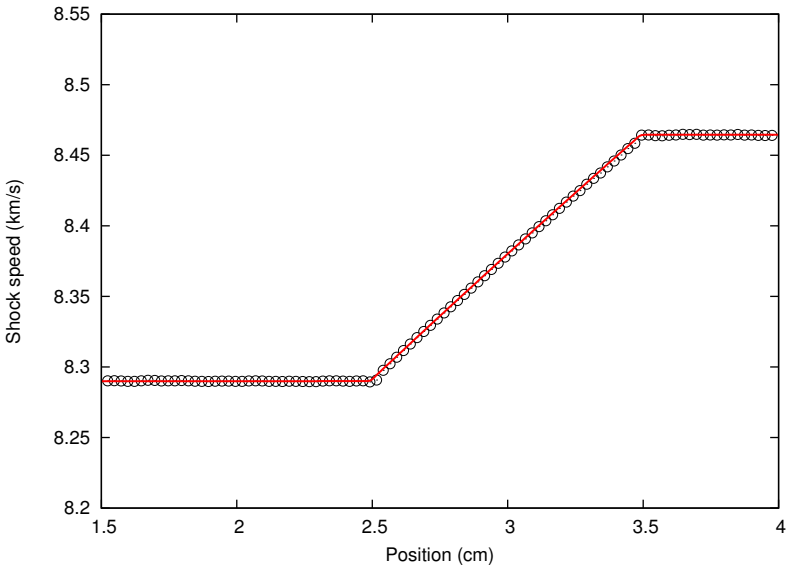

Figure 4.13.: The speed of the shock wave as it passes over the density gradient. Since in this case the shock is not steady, its speed is defined at the location where the particle velocity reaches half of its post-shock value. The circles are from direct simulation, and the line is the steady shock speed computed from the density at that shock location.

4.6.4. Grain boundary test case

The ability to simulate a polycrystalline structure is a possibility with an anisotropic equation of state. The grain structure of a material can be important in determining its mechanical properties (as well as other properties, including thermal conductivity where they contribute to thermal resistance). The grain-boundary orientation distribution for polycrystalline silicon is described by Ratanaphan et al. (2014). This motivates the following test problem.

The initial condition has the left half of the domain oriented with [100] in the x direction and the right half oriented as [110], achieved by choosing F_{orient} to be a rotation of 45° about the y-axis.

A shock of strength 3.86 GPa is maintained with a left-hand boundary condition of particle velocity 0.2 km/s and uniaxial deformation $F_{xx} = 0.976$ (with the identity at the zero stress configuration). The result of this is shown in fig. 4.14. The initial temperature was 290 K, and after the shock, this rises slightly (and nearly isentropically) to 293 K. Several artifacts are visible in the entropy and temperature panels of fig. 4.14: even though

they are small, the jump in solution is also small for this problem. The leftmost artifact is a start-up error due to the initial location of the shock.

As the shock interacts with the boundary, a small component is reflected, and the majority of the wave energy is transmitted. A uniaxial compression in the x direction in material oriented as [110] is not symmetric in the xy plane, and this is seen in the stress panel of fig. 4.14.

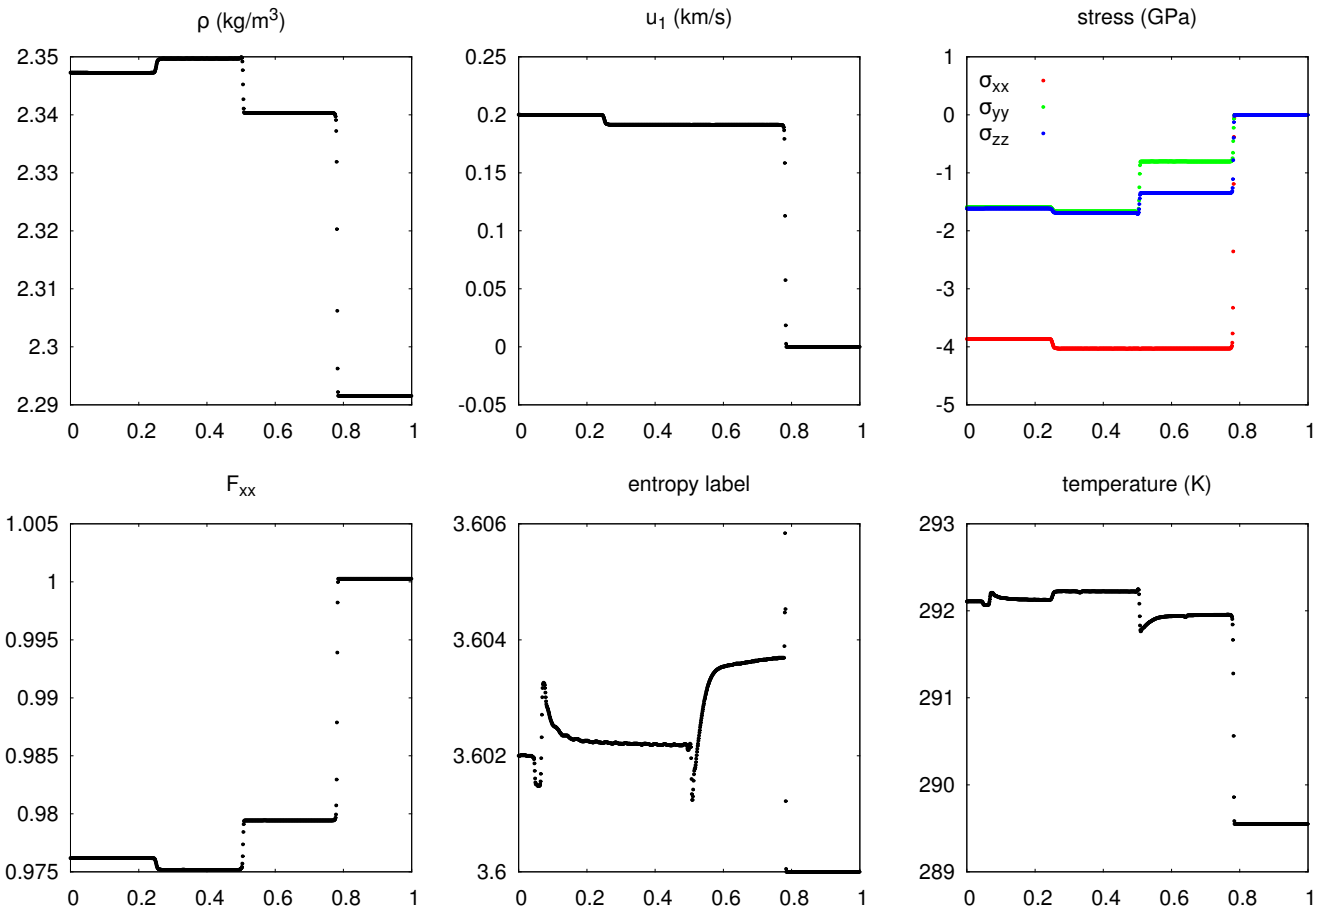

Figure 4.14.: Numerical results for test case GB1, solved using MPWENO-5/FORCE, with a CFL number of 0.6. The points are the solution on a 1000 cell grid at $t = 1.2\,\mu$s.

4.7. Conclusions

In summary,

- This procedure provides a method for reconstructing typical elastic strain-energy surfaces, using a few hundred training points as input, to an accuracy acceptable for continuum simulation, and introducing an error not greater than the underlying density functional theory calculations.
- DFT-MD simulations were performed for silicon to obtain strain energies and stresses, paramaterized by deformation gradient and another parameter dependent only on entropy for elastic deformations of a single crystalline phase.
- The reconstruction is done using Gaussian Process (GP) regression. While slower to evaluate than simple equation forms (see the note in appendix B), this is a trade-off against the systematic ability to generate equations of state for any crystalline material relatively simply, and with assurance that the surface from the regression is capable of accurately representing the sampled surface, given enough sampling points. The speed of evaluation may be improved by using a reduced-rank Gaussian process (appendix B.5) and, where available, using a general-purpose GPU as a co-processor to perform the equation of state evaluations (appendix B.6).
- The resulting equation of state has a complete form (energy as a function of deformation and entropy), and the stress from the fit is consistent with the virial stress from molecular dynamics.

4.7.1. Tabular equations of state

The equation of state we have presented here is, in a certain sense, tabular: the full data set is available to the Gaussian process regression used for interpolation. Several popular equations of state for fluids (for example, the extensive SESAME equation of state library (Lyon and Johnson, 1992)) are based on tabular data that are interpolated to calculate thermodynamic functions at a specific input.

For anisotropic elasticity however, a straightforward tabulation performs poorly because of the inherent dimensionality of the data. This also restricts any reasonably simple interpolation strategy to be multilinear, introducing discontinuities in stress, and in turn affecting the wave speed.

Even though more robust table-based interpolation approaches are possible (e.g. Barthelmann et al., 2000), we briefly demonstrate here why the straightforward ex-

tension of a tabular equation of state with multilinear interpolation, as is used successfully for fluid equations of state (or approximating a solid as such), is impractical for anisotropic elasticity.

The performance of this approach can be seen in figs. 4.15 and 4.16. These show the results of the initial value problem with data given as the second test case in table 3.1, using the Romenski equation of state for copper (appendix A and table A.1 therein). Compare with the high-resolution results of fig. 3.4. They show the numerical solution using the analytic form of equation of state with the same equation computed from two linearly interpolated tables: the first tabulating only the internal energy, but keeping the exact analytic stress (although not useful for a practical application), and the second adding a separate interpolated table for the stress. The equation of state is inverted numerically, according to the table. The numerical results are obtained with the FORCE flux. Even though the numerical method is quite diffusive, there are unwanted oscillations when the fully tabular EOS is used.

For this example, a uniform grid is used between 0.85 and 1.15 for the diagonal components of $\boldsymbol{C}$, -0.4 and 0.4 for the off diagonal components, and between 0.0 and 2×10^{-3} in entropy. The tables themselves are quite large: the 'coarse' tables store 15^7 double-precision floating point values (1GB), and the 'fine' tables, 20^7 values (10GB). Systematic improvement with the size of the tables is also slow.

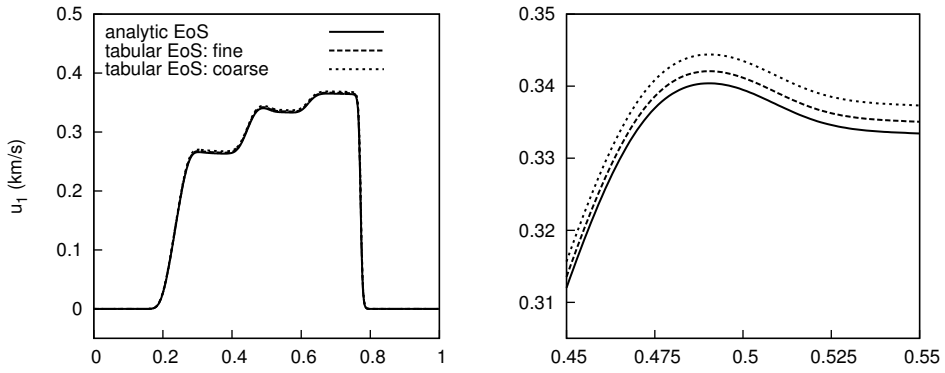

Figure 4.15.: One component of the solution to test case 1, using the Romenski equation of state for copper. Numerical solutions using the full analytic equation of state, and an equation of state with energy linearly interpolated from a table (but still using the analytic form of the stress) are shown, for two table resolutions. The coarse table with 15^7 points, and the fine table with 20^7 points.

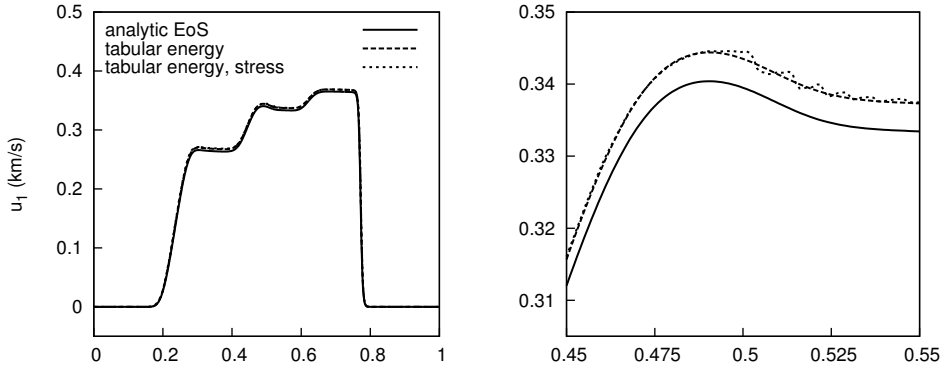

Figure 4.16.: One component of the solution to test case 1, using the Romenski equation of state for copper. The solution from the analytic equation of state is given, along with an equation of state with linearly interpolated energy (but analytic stress), and another with both tabular energy and stress.

5. Simulation of Shock Waves with Molecular Dynamics

5.1. Introduction

Shock waves in solids show a variety of behaviours, including yield and plasticity and structural phase transitions, which can lead to wave splitting and rarefaction shocks. Even the propagation of elastic waves can show strong nonlinearities and crystal-orientation dependence. Other factors, such as microstructure and porosity, complicate the behaviour further, although we do not consider these aspects in this work.

Simulations of shock waves have a long history (Holian, 2004), and in fact were some of the earliest to be performed on a computer. Direct simulations using empirical interatomic potentials are now feasible on a multi-billion atom scale on present hardware, which is large enough to observe detailed mechanisms of yield, plastic flow and shock interaction with nanostructures, directly (Kadau et al., 2006; Shekhar et al., 2013). Work with these potentials can give important insight and understanding, but a need for first-principles methods such as Density Functional Theory (DFT) exists in providing predictive power and accuracy. These methods must use more modest system sizes, of hundreds or thousands of atoms in the case of DFT.

Silicon has a rich phase diagram, with dense metallic phases rather different in character to the ambient diamond phase, making it an interesting and challenging object of simulation. In total, eleven stable or metastable phases of silicon are currently known (Mujica et al., 2003), and shock experiments have provided important data for constructing the phase diagram. See the discussion and reference in section 5.5.

We use direct, non-equilibrium simulations (section 5.3) as well as an indirect annealing method described in section 5.4 to give the Hugoniot loci of silicon as computed from a number of empirical potentials, and from Density Functional Theory. The Hugoniot locus is the set of thermodynamic states accessible from a given shock process. A single direct shock simulation is performed from first-principles. The fact that shock waves

can split into multiple waves due to phase transitions or yielding is taken into account here by specifying the strength of any preceding waves explicitly based on their yield strain. Points corresponding to uniaxial elastic compression along three crystal axes and a number of post-shock phases are given, including a plastically-yielded state, approximated by an isotropic stress configuration following an elastic wave of predetermined strength. The DFT results compare well to existing experimental data for shocked silicon.

5.2. Background

5.2.1. The Rankine–Hugoniot relations

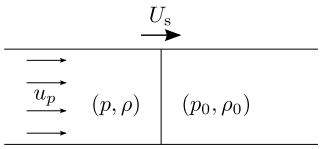

Figure 5.1.: A planar shock. U_s is the speed of the shock and u_p is the speed of the material behind it.

The Rankine–Hugoniot relations relate the thermodynamic state of a material on either side of a shock wave. A schematic of this situation is shown in fig. 5.1. The meanings of the symbols throughout this section are as follows: E is the specific internal energy, U is the velocity of the shock wave, u is the velocity of the material after the shock has passed (the 'particle velocity') and $v = 1/\rho$ is the specific volume (the reciprocal of the density). Subscripted variables refer to the pre-shocked state (by a zero for the initial state, a one before the second shock, etc.) and those without a subscript, to the post-shocked state. An equation below may be found to relate any three post-shock variables. It is assumed that $u_0 = 0$.

There are three independent jump conditions, corresponding to the conservation laws of energy, momentum and mass. There are five variables in a planar shock problem. Together, these lead to ten Rankine–Hugoniot relations, each relating three shocked variables. They are given in full below, as eqs. (5.1) to (5.10) in fig. 5.2. The equations apply to a steady-state flow in any material described with the above variables. The analogous equations for a solid have p replaced by the component of (the negative of) the stress normal to the shock front.

Material-specific properties are determined by an equation of state, reducing the number of independent quantities from two to one and giving a one-parameter family of states that may be reached by shocking from a given initial state. This is the *Hugoniot locus*

(sometimes just 'the Hugoniot'). The remaining parameter provides a measure of the strength of the shock.

It is important to note that each point on the Hugoniot represents the result of the passage of a shock wave independent of the other points: it does not describe the thermodynamic path by which the material attains a particular final state, which will usually be different, where a process is slow enough that it may be defined at all—the strain rate due to the passage of a strong shock in a solid can be of the order of $10^8 \mathrm{s}^{-1}$ (e.g. Chhabildas and Asay, 1979).

The *Rayleigh line* is a straight line in the p–v plane, connecting the initial and final states. It is significant for some of the analysis below.

$$\frac{v}{v_0} + \frac{u}{U} = 1 \tag{5.1}$$

$$p - p_0 = \rho_0 U u \tag{5.2}$$

$$E - E_0 = \frac{1}{2}(p + p_0)(v_0 - v) \tag{5.3}$$

$$p - p_0 = \rho_0^2 U^2 (v_0 - v) \tag{5.4}$$

$$E - E_0 = \frac{1}{2} u^2 \left(\frac{p + p_0}{p - p_0} \right) \tag{5.5}$$

$$u^2 = (p - p_0)(v_0 - v) \tag{5.6}$$

$$E - E_0 = \frac{p^2 - p_0^2}{2\rho_0^2 U^2} \tag{5.7}$$

$$E - E_0 = \frac{1}{2} u^2 + \frac{p_0 u}{\rho_0 U} \tag{5.8}$$

$$E - E_0 = \frac{1}{2} u^2 + p_0 (v_0 - v) \tag{5.9}$$

$$E - E_0 = \frac{1}{2} U^2 \left(1 - \frac{v}{v_0} \right)^2 + p_0 (v_0 - v). \tag{5.10}$$

Figure 5.2.: The Rankine–Hugoniot relations, relating each set of three variables.

5.2.2. Elastic shock waves

For the lowest shock pressures, the mechanism of propagation is elastic strain. In the weak-shock limit of low particle velocity, the shock wave will be an isentropic,

pure-longitudinal elastic wave (the material is deformed uniaxially), with corresponding velocity

$$s_n = \sqrt{\frac{c_n}{\rho}} \tag{5.11}$$

where $c_n = \boldsymbol{c} \cdot \boldsymbol{n}$ is the fundamental elastic constant in the shock propagation direction, $\boldsymbol{n}$.

At a certain shock pressure, the material will yield dynamically. The maximum shock pressure a material may sustain without permanent deformation occurring is known as the *Hugoniot Elastic Limit* (HEL). The HEL varies greatly with material properties and also with microstructure. For example, in non-porous quartzite, 9GPa is typical, compared with 0.5GPa in porous rock of the same chemical composition (Ahrens et al., 1964). Sapphire has the largest observed HEL, of between 12 and 21GPa (Graham and Brooks, 1971). Many metals have an HEL of about 1GPa.

5.2.3. The Bethe-Weyl condition for the stability of a shock-wave

The argument here follows Gathers (1994, ch. 3).

We wish to determine a condition for the stability of a shock-wave, with respect to the wave splitting up into a series of shocks of intermediate strength that together take the material to the same final pressure. The final density and energy for a split shock and a single shock to the same pressure will in general be different. We will see that the shock is stable exactly when the Hugoniot locus in the p-v plane is concave upward.

Consider the shape of Hugoniot locus in the p–v plane. Its slope is $-v\left(\frac{\partial p}{\partial v}\right)_H$, which is related to the compressibility of the material on the Hugoniot, and where it is understood that p represents the component of stress normal to the shock front, and not necessarily the hydrostatic pressure. The compressibility tends to decrease with pressure, so the Hugoniot represented this way is generally concave upward. An exceptional situation may occur, in particular near phase transitions, where it is convex due to a cusp.

Suppose a shock from the initial state '0' to an intermediate state '1' is followed by a second shock to a state '2'. Variables will be labelled according to their respective state.

For the first shock, eq. (5.4) may be written:

$$U_1 = v_0\sqrt{\frac{p_1 - p_0}{v_0 - v_1}} \tag{5.12}$$

and for the second,

$$U_2 = u_1 + v_1\sqrt{\frac{p_2 - p_1}{v_1 - v_2}} \tag{5.13}$$

Note that $\frac{p_1-p_0}{v_0-v_1}$ is just the slope of the Rayleigh line connecting state 0 and state 1.

From equation (5.6),

$$u_1 = \sqrt{(p_1 - p_0)(v_0 - v_1)} \tag{5.14}$$

and substituting into the previous equation for U_2 gives

$$U_2 = \sqrt{(p_1 - p_0)(v_0 - v_1)} + v_1\sqrt{\frac{p_2 - p_1}{v_1 - v_2}} \tag{5.15}$$

If the first wave moves more slowly than the second, then the second will overtake it and a single shock structure will form.

$$\sqrt{(p_1 - p_0)(v_0 - v_1)} + v_1\sqrt{\frac{p_2 - p_1}{v_1 - v_2}} \geq v_0\sqrt{\frac{p_1 - p_0}{v_0 - v_1}}. \tag{5.16}$$

This equation may be simplified to give the condition for the stability of a single shock wave as

$$\frac{p_2 - p_1}{v_1 - v_2} \geq \frac{p_1 - p_0}{v_0 - v_1}, \tag{5.17}$$

for any intermediate state '1'.

This is equivalent to concavity,

$$\left(\frac{\partial^2 p}{\partial v^2}\right)_H \geq 0. \tag{5.18}$$

5.2.4. Strong shock waves: plasticity and phase changes

After yielding, a material has a stress distribution closer to isotropic, although there may be some remaining deviatoric stress. The Hugoniot locus in longitudinal stress therefore has a lesser slope immediately after the yield than before. The above argument leads us to conclude that a single shock wave would be unstable in this case, and it must split into an elastic precursor and a plastic shock for at least some wave strengths, as illustrated in figs. 5.3 and 5.4.

Similarly, for the case of a plastic shock, a pressure–volume Hugoniot is convex at the onset of a phase transition: if the change in slope is great enough, this causes the shock to split into a wave taking the material to the pressure at the onset of the phase transition, and a slower wave taking the material to its final state, which could be a coexistence of the two phases in the mixed-phase region, or existing fully in the new phase. Shock splitting

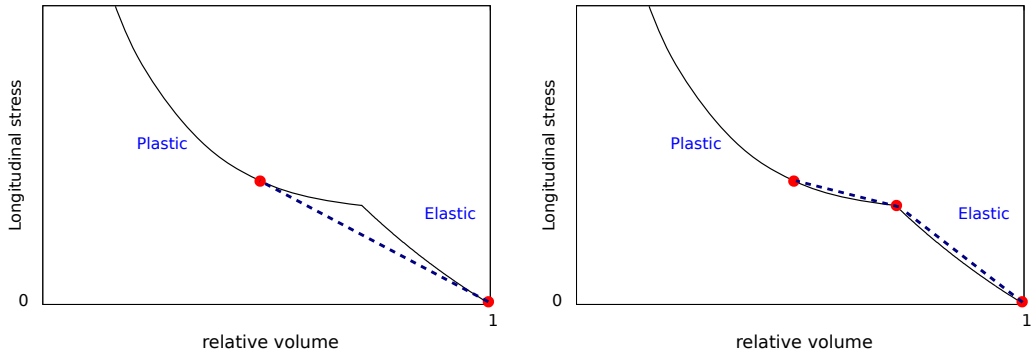

Figure 5.3.: A schematic of a convex pressure–volume Hugoniot leading to a split wave. The red points are initial and final shock states and the dark blue dashed lines are Rayleigh lines connecting them. The shock corresponding to the Rayleigh line in the left-hand figure is not observed: instead there are two shocks, with two separate Rayleigh lines indicated in the right-hand figure. The final state is typically at slightly lower pressure in the two-shock case, due to the reduced heating of two shocks to the same relative volume compared with a single shock.

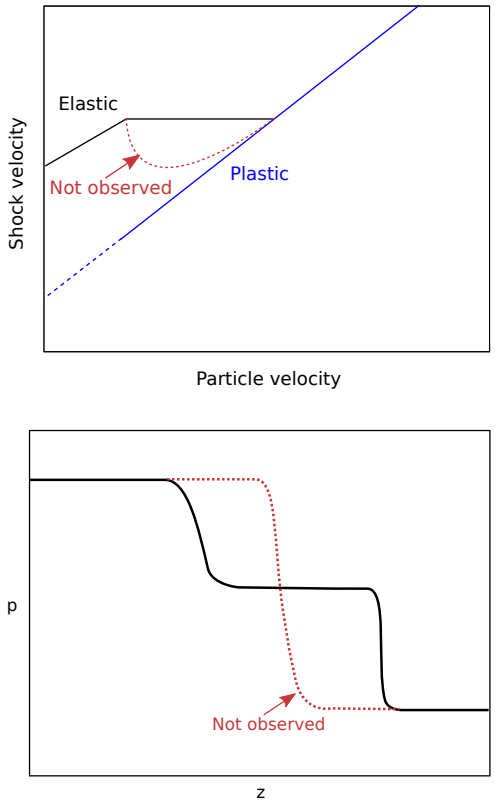

Figure 5.4.: A schematic of a split wave (shock-velocity–particle-velocity Hugoniot, top; wave profile, bottom) resulting from a convex Hugoniot.

on yielding and phase transitions have long been observed and understood (Duvall and Graham, 1977).

The plastic shock velocity–particle velocity Hugoniot locus should extrapolate to roughly the (calculated) isotropic sound speed limit at low strains, but due to residual stress, and the preceding elastic wave, the agreement might not hold well in all cases. The isotropic, or bulk sound speed is given by

$$c_{\text{bulk}} = \sqrt{\frac{K_0}{\rho}} \tag{5.19}$$

where K_0 is the bulk modulus in the reference configuration.

The Hugoniot locus through the mixed phase region can be constructed by considering the jump condition in enthalpy across the shock from the point ('1') at the onset of the transition to a point ('2') on the mixed Hugoniot

$$h_2 - h_1 = E_2 - E_1 + p_2 v_2 - p_1 v_1, \tag{5.20}$$

and on substituting eq. (5.3) for the jump in internal energy, this reduces to

$$h_2 - h_1 = \frac{1}{2}(p_2 - p_1)(V_2 + V_1). \tag{5.21}$$

The latent heat L of the phase transition results in a change in enthalpy, written according to the Clausius–Clapeyron equation as

$$\lambda L = -T \frac{\mathrm{d}p}{\mathrm{d}T}(V_1 - V_2), \tag{5.22}$$

where λ is the mass fraction of the second phase and the derivative is along the phase line.

Since the mixed region is not at constant pressure, there is an additional contribution to the enthalpy change from the difference in pressure and volume between the onset of the transition and the post-shock state. This leads to a linearized equation relating the pressure and volume changes on the phase-transition shock (Duff and Minshall, 1957),

$$p_2 - p_1 = (V_1 - V_2) \times \left[\beta v_1 + \left(\frac{1}{2T_1}(V_1 - V_2) - 2\alpha v_1 \right) \frac{\mathrm{d}T}{\mathrm{d}p} + \frac{c_p}{T_1} \left(\frac{\mathrm{d}T}{\mathrm{d}p} \right)^2 \right]^{-1}, \tag{5.23}$$

where β is the isothermal compressibility, α is the volumetric thermal expansion coefficient and c_p is the specific heat capacity at constant pressure. The derivative $\mathrm{d}T/\mathrm{d}p$ is once again along the phase boundary.

We require knowledge the onset of the transition in the p–v plane, which is not available from the single phase simulations alone (the simulated materials are capable of being superheated or supercooled substantially). This could be obtained from the point where the Hugoniot cuts the phase boundary obtained by some other method.

5.3. Direct Simulation of Shock Waves with Non-equilibrium Molecular Dynamics

The most direct approach to simulating a shock wave with molecular dynamics is to mimic experiment, and collide a 'flyer-plate' of one material into a 'sample' of another material. To simplify the situation slightly, simulating the 'flyer-plate' itself can be avoided by either representing it as a reflective boundary condition moving at the intended particle velocity relative to the sample (the *momentum mirror*), or by using atoms of some material constrained to move at a certain fixed velocity (the *constrained atom* approach). This latter method creates a soft potential reflecting the sample atoms, but does not allow energy to be lost to the flyer-plate, since their positions are not integrated dynamically. In both cases, the shock strength is maintained by moving the constraint at a constant velocity.

The typical shape of such a simulation is a large-aspect-ratio cuboidal box, with periodic boundary conditions in the transverse directions, and vacuum in the longitudinal direction.

Figure 5.5 shows a schematic of the two situations. There should be no effect on the results of a simulation due to using one method over the other for long enough simulations, since the steady shock state should be the same, and determined by the particle velocity. Differences at the start of the simulation could, for example, lead to one method triggering yielding sooner than the other. Since we are interested in small systems however, we are interested in the time it takes for the measured shock velocities to reach their converged, steady value in either case, which is found to occur rapidly.

The main disadvantage of the method is its computational cost, due to the large number of atoms involved. The scaling is also unfavourable, since simulating a shock for longer involves adding additional atoms in proportion to the simulation time, resulting in $O(N^2)$ time complexity at a minimum, and $O(N^4)$ time complexity in the case of cubic-scaling DFT. This method is the only way to observe time-dependent effects due to the rise of the

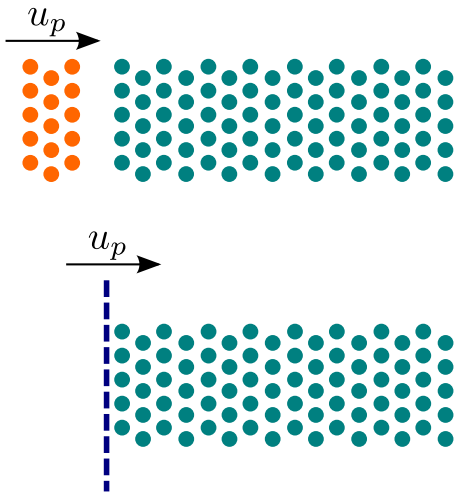

Figure 5.5.: Schematic diagram of two methods for generating a shock in an MD simulation. The teal circles represent atoms in the target (a two-dimensional cross section through a 3-d simulation). The top diagram shows 'flyer-plate' atoms in orange; the bottom diagram shows the momentum mirror. Periodic boundaries are used transverse to the shock direction.

shock, and its effect on, for example, nanostructures in the material, beyond the effect of the uniform post-shock stress. It is also the only means of observing the spatial profile of the shock wave.

A direct approach allows for spontaneous plasticity and phase transitions. As in other molecular dynamics simulations, there is likely to be a considerable delay in the onset of crystal plasticity or a phase transition, compared with any real sample of material, which will be much larger and inevitably contain defects to nucleate phase transitions. The periodicity of the domain can also restrict or inhibit plasticity or structural phase transitions, since these must then obey the same periodicity. Both of these problems can be resolved in part with a larger domain (since a larger region in phase space is being sampled by the simulation, and can reduce the energy cost of forming sub-domains of a new structure). Introducing defects by hand into the simulated domain is another strategy, but must clearly be done with care to be physical.

5.4. Annealing to the Hugoniot Locus

While the direct method allows the detailed structure and dynamics of a shock wave to be observed, including the detailed effects on any nanostructures present in the material, it is expensive, and only the simplest cases can be performed using DFT.

It is simple to check that a given equilibrium state lies on or close to the (single-shock) Hugoniot locus, which amounts to satisfying the Hugoniot relation for energy, eq. (5.3). This is a necessary condition for the state to be an allowed post-shock state, but the state so obtained might not be observed in a direct simulation in the case of plasticity or a phase change, where there would exist another state satisfying a Hugoniot relation (perhaps modified by the presence of multiple waves), but that is more favourable in free energy.

This suggests that a Hugoniot state can be determined dynamically from within a single molecular dynamics simulation by an annealing process, or some other modified dynamics to constrain the state to satisfy eq. (5.3). This is the approach taken by the Hugoniostat methods (Maillet et al., 2000; Ravelo et al., 2004) and the technique of Reed et al. (2003) The former simulations use modified Nosé–Hoover dynamics while the latter uses coupled dynamics of the atoms and simulation cell, whose Lagrangian involves the computed instantaneous shock speed, and varies the simulation cell uniaxially. One aim of these dynamics is to work on timescales comparable to shock-passage times, without the overhead of dealing with a direct non-equilibrium simulation.

If we are interested only in the final post-shock state, and are not interested in the (modified) dynamics while the constraint is being applied, we are free to use a method based on simple velocity rescaling, analogous to the procedure of Berendsen et al. (1984) which is what we propose here due to its increased efficiency in reaching the final state. A Berendsen thermostat is used with a variable target temperature computed from the instantaneous difference in energy between the total energy of the system, and the total energy that would be required to satisfy the energy Hugoniot relation, eq. (5.3), exactly, given the current instantaneous longitudinal stress.

The procedure is given explicitly below as algorithm 1. This may be combined with a further anneal to relax the pressure to a hydrostatic configuration if desired. Optionally, the box vectors may be gradually ramped between two states, which is most useful when the starting state of the simulation and the initial state of the Hugoniot locus are the same.

Even though Berendsen thermo- and barostats do not reproduce canonical statistics (Harvey et al., 1998), it is well known that they are much more efficient in annealing to a state at a desired temperature or pressure, compared with modified dynamics such as Nosé–Hoover. The same applies here, compared with the related Hugoniostat for shocks, and this justifies their use here, since we are interested only in the outcome of the anneal, not the intermediate dynamics. After the time-averaged state of the system closely

Algorithm 1 The Hugoniot-following procedure. The meaning of the variables used is as follows. E denotes an energy (refer to the sub and superscripts), V is the unit cell volume, σ the virial stress (includes a kinetic term), ς the instantaneous stress, $\boldsymbol{x}_n$, $\boldsymbol{v}_n$ and $\boldsymbol{F}_n$ the atomic positions, velocities and forces at the nth timestep ('unc' stands for 'uncorrected' and 'sca' for scaled), m is the mass of a given atom, τ_{relax} is the relaxation time described in the text, t and $\mathrm{d}t$ are the current time and timestep, and anything with a subscript '0' refers to its (time averaged) value in the unshocked state (which may be different from the starting state of the simulation).

procedure HUGONIOTANNEAL($E_0^{\text{tot}}, V_0, \sigma_0$)

compute $E^{\text{tot}}, \varsigma, \boldsymbol{F}_n$ from atomic positions $\boldsymbol{x}_n$

for all atoms **do** ▷ velocity Verlet

$$\boldsymbol{v}_n \leftarrow \boldsymbol{v}_{n-1} + \frac{\mathrm{d}t}{2m}(\boldsymbol{F}_{n-1} + \boldsymbol{F}_n)$$

$$\boldsymbol{x}_{n+1}^{(\text{unc})} \leftarrow \boldsymbol{x}_n + \mathrm{d}t\,\boldsymbol{v}_n + \frac{\mathrm{d}t^2}{2}\boldsymbol{F}_n/m$$

end for

$\sigma \leftarrow \frac{1}{V}\sum_{\text{atoms}} m\,\boldsymbol{v}_n \otimes \boldsymbol{v}_n + \varsigma$ ▷ compute the virial stress

$E^{\text{kin}} \leftarrow \sum_{\text{atoms}} \frac{1}{2} m\,\boldsymbol{v}_n \cdot \boldsymbol{v}_n$

▷ compute the target energy

$E^{\text{hug}} \leftarrow E_0^{\text{tot}} - \frac{1}{2}(\sigma^{33} + \sigma_0^{33})(V_0 - V)$

$E_{\text{target}}^{\text{kin}} \leftarrow E^{\text{kin}} + E^{\text{hug}} - E^{\text{tot}}$

$$r^2 \leftarrow \left(1 + \frac{\mathrm{d}t}{\tau_{\text{relax}}}\left(\frac{E_{\text{target}}^{\text{kin}}}{E^{\text{kin}}} - 1\right)\right)$$

for all atoms **do**

$\boldsymbol{v}_n^{(\text{sca})} \leftarrow r\,\boldsymbol{v}_n$ ▷ scale the velocities

▷ correct positions based on the scaled velocities

$\boldsymbol{x}_{n+1} \leftarrow \boldsymbol{x}_{n+1}^{(\text{unc})} + \mathrm{d}t\,(\boldsymbol{v}_n^{(\text{sca})} - \boldsymbol{v}_n)$

end for

$t \leftarrow t + \mathrm{d}t$, $n \leftarrow n + 1$

end procedure

satisfies the Hugoniot relation, the simulation can be restarted with Verlet dynamics to check if eq. (5.3) is indeed satisfied.

5.5. Silicon

The phase transition in silicon from the cubic diamond structure to the beta-tin structure, occurring at 12 GPa at room temperature, and undergoing a reduction in volume of 20 %, has been well established by static loading experiments, starting in the 1960s (Jamieson, 1963; Minomura and Drickamer, 1962). Evidence of at least one phase transition at similar pressures was then observed in shock-wave experiments, starting with Pavlovskii (1968). Gust and Royce (1970, 1971) found a three-wave structure for samples shocked in the ⟨100⟩ crystal direction and a four-wave structure when shocked in the ⟨110⟩ or ⟨111⟩ directions. In the latter cases, these waves were attributed to: an initial elastic precursor to the Hugoniot elastic limit of 5.5 GPa, followed by waves corresponding to a state of plastic yield and two successive phase transitions at 10 GPa and 13 GPa. Along ⟨100⟩, the higher elastic limit of 9 GPa obscures the first transition wave, and a single wave takes the material simultaneously to a new phase and to a state of hydrostatic stress.

The work of Goto et al. (1982) largely confirmed the findings of Gust and Royce (1971), although they observed a three-wave structure, regardless of crystal orientation, consistent with only a single phase transition at 13 GPa. Above the Hugoniot elastic limit, shock compression was found to result in a hydrostatic stress configuration, due to the complete loss of strength in the material.

More recently, and contrary to the earlier experimental work, Turneaure and Gupta (2007a,b) reported a single phase transition that is complete by 15.9 GPa. Shocks to these pressures show a much greater volume compression than the points attributed to belong to an extended mixed-phase region by both Gust and Royce (1971) and Goto et al. (1982) Here the phase transition is not complete until at least 30 GPa. This is attributed by Turneaure and Gupta (2007b) to release waves from the first two waves propagating back into the material before the arrival of the third wave, and altering the peak state. They avoid this eventuality by backing the silicon with a window made from lithium fluoride, a material with a good impedance match to silicon.

The *Imma* phase of silicon is found intermediate between the beta-tin and simple hexagonal phases, at low temperatures, and is stable between 13 GPa and 15 GPa at room temperature (McMahon et al., 1994). A recent simulation of directly shocked silicon using an empirical potential found a phase transition to a phase with *Imma* symmetry (Mogni et al., 2014), and proposed a mechanism for shear-stress relief. The Erhart and

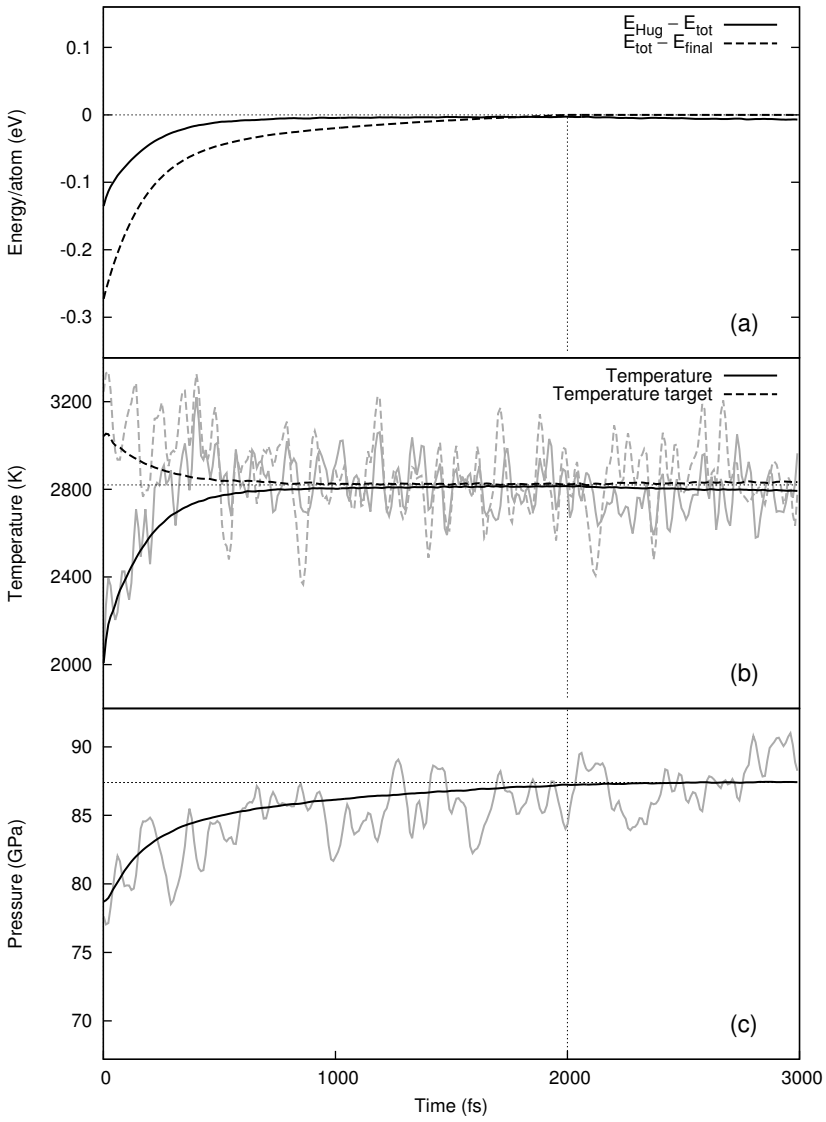

Figure 5.6.: Response of (a) total energy and the difference with the Hugoniot energy computed from eq. (5.3), (b) temperature and target temperature, and (c) pressure, to the Hugoniot anneal described here with a relaxation time of 100 fs. After 2000 fs, the anneal is switched off and the dynamics continued with Verlet integration. The response is averaged over 10 000 independent 216 atom Stillinger–Weber silicon systems, starting from a 2000 K liquid and annealed to the Hugoniot locus with an initial state of 300 K and zero stress. For comparison, the light grey lines are taken from a single trajectory—in the energy plot, this is indistinguishable from the mean.

Albe (2005) modification of the Tersoff (1986, 1988) potential used by these authors work tended to overestimate the stress in the dense phase. Theoretically, the energy and volume of these phases are close (Mujica et al., 2003). Swift et al. (2001) mention that a stable *Imma* phase is not found along the shock locus, so is not included in their equation of state.

5.5.1. Results from first-principles: annealing to the Hugoniot locus

The ab initio MD simulations described here were performed with the SIESTA method and implementation of Density Functional Theory (Soler et al., 2002), using the GGA functional of Perdew et al. (1996).

The core electrons were described with a Troullier–Martins norm-conserving pseudopotential (Troullier and Martins, 1991) with a matching radius in each angular momentum channel of 1.89 a_0. The valence electrons were described with a basis of numerical atomic orbitals (Junquera et al., 2001) of double-ζ polarized type (representing 13 orbitals per atom). The basis was generated by fixing the longest orbital cutoffs at 7.0 a_0 and variationally optimising the other parameters in bulk diamond-phase silicon—the final basis parameters are given in table F.1.

The mesh used for integrals in real-space was well converged at a grid cutoff of 100 Ry. The dense phases of silicon required several k-points to converge in energy, and in particular, for the cold compression curves of the various phases to converge in energy relative to one another. A 4^3 Monkhorst–Pack grid of points was used on the 64 atom simulations, to give an effective cutoff length of 21 Å. Note that this is a finer grid than used in the previous chapter.

The electronic temperature used in the DFT calculations should be consistent with the final temperature attained after the annealing process described below. The consistent forces for the ab initio molecular dynamics are the nuclear-position derivatives of the electronic free-energy as defined in Mermin's DFT (Mermin, 1965). All of the simulations reported below are for an electronic temperature of 300 K, except for the two highest-pressure liquid points, for which the electronic temperature was adjusted to coincide with the final (nuclear) temperature and re-run. The effect of the electronic temperature on the reported quantities was found to be quite small: the maximum difference in pressure for the hottest simulation between using a consistent electronic temperature and the initial 300 K is below 5 %.

The integration of the dynamics used the Born-Oppenheimer approximation with a timestep of 1 fs. These calculations assume instantaneous equilibrium of electrons and

ions behind the shock front. Note that this may fail behind a very strong shock wave, according to the experimental work of Celliers et al. (1992).

The calculated pressure–volume and shock-velocity–particle-velocity Hugoniot loci for the pure phases are compared to results from several experiments in figs. 5.9 to 5.11. The specific volume at zero pressure and 300 K for the PBE functional is $0.421\,\mathrm{cm}^3/\mathrm{g}$, which is smaller than the experimental value of $0.431\,\mathrm{cm}^3/\mathrm{g}$. Reduced volume is plotted in the figures: if specific volume were plotted, the DFT results would be offset by an amount corresponding to the difference. Both particle velocity and shock velocity scale with $\sqrt{1/\rho_0}$: the roughly 2% error in initial density results in a 1% error in both velocities. No adjustment is made to the particle velocity–shock velocity Hugoniot figure.

The curves for the elastic shocks are computed from a uniaxial box deformation along the indicated direction. The 'plastic' curve is for a split shock, with an elastic precursor to 6 GPa, taking the material to a hydrostatic stress configuration: this supposes that the material has no residual strength. The hydrostat in fig. 5.11 is for an unphysical shock process that relaxes the material to hydrostatic stress behind a single, unsplit shock wave. This permits comparison with the bulk speed of sound (the shock velocity for this wave should extrapolate to the bulk speed of sound at zero particle velocity.)

When comparing the hydrostat and the 'plastic' curve to the yielded phase, we assume that the yielding serves only to remove the deviatoric stress, and that the bulk response of the material is unaffected. We neglect the dissipative heating due to this effect.

The agreement with the experimental data for the elastic and plastic shocks is good, with the compressibility along ⟨100⟩, ⟨110⟩ matching well in value and ⟨111⟩ showing the correct trend (although underestimating the value). The close match between the experimental plastic shock pressures and the hydrostatic plastic shock calculated here supports the observation that the material loses all of its strength after yield.

The particle and shock velocities in fig. 5.11 are computed from the computed pressure and volume points using the Hugoniot relations

$$u_p^2 = (p - p_0)(V_0 - V) \tag{5.24}$$

$$U_s^2 = v_0^2 (p - p_0)/(V_0 - V), \tag{5.25}$$

where V is the absolute volume of the simulation and v is the specific volume. A linear fit to the elastic part of the shock-velocity–particle-velocity Hugoniot has coefficients given in table 5.1. The extrapolated value of the bulk sound speed of 6.51 km/s agrees very well with the value of 6.48 km/s calculated from the second order elastic constants (Gust and Royce, 1971; Hall, 1967).

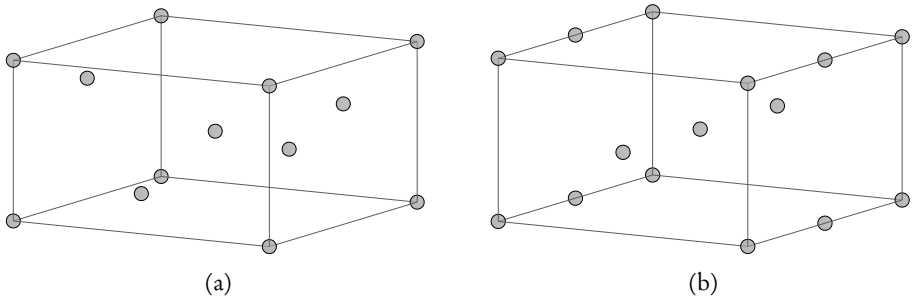

Figure 5.7.: Equivalent unit cells of the beta tin (a) and simple hexagonal (b) structures.

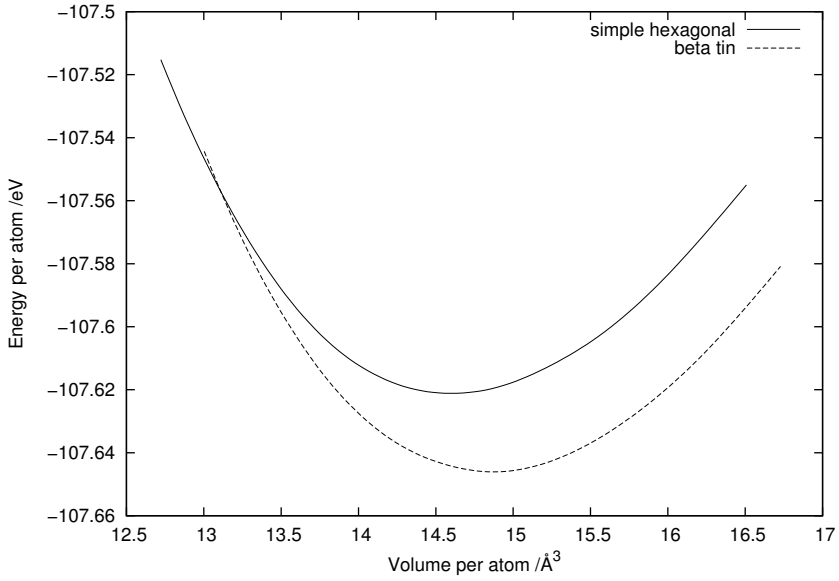

Figure 5.8.: Cold compression curves of beta tin and simple hexagonal phases of silicon from GGA/PBE. For convergence in transition pressure, a k-grid cutoff of at least 21 Å is needed. This figure used a cutoff of 42 Å.

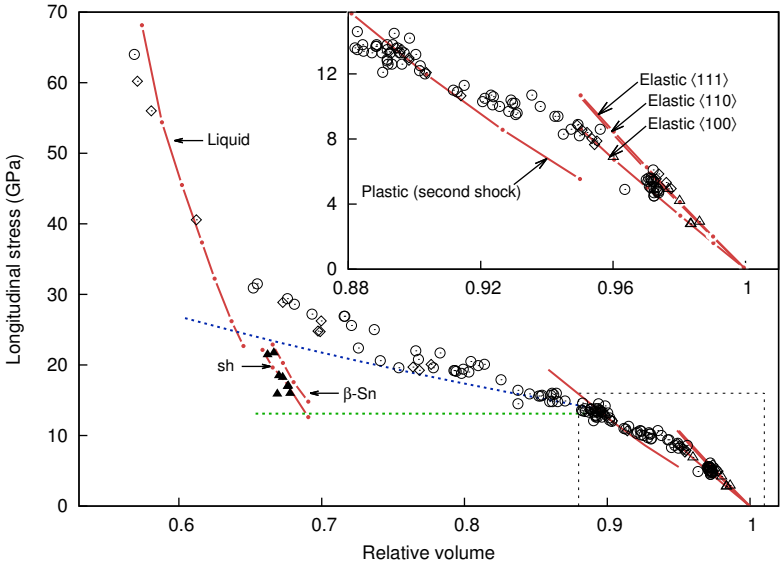

Figure 5.9.: Pressure–volume Hugoniot loci for silicon. The solid red lines in the figure are the DFT results from this work (with contained points indicating individual simulations), with an initial pre-shocked state of zero pressure and 300 K, with the final state in the indicated phase ('sh' for simple hexagonal). Estimated error is less than 5 % for the liquid and beta-tin phases, and is substantially smaller for the diamond phase. The symbols are experimental results from the literature: ○ Gust and Royce (1971), ◇ Goto et al. (1982), △ Turneaure and Gupta (2007a), ▲ Turneaure and Gupta (2007b). The dashed lines are approximations to the mixed-phase portion of the Hugoniot, for cubic diamond to: liquid (upper blue line) and beta tin (lower green line).

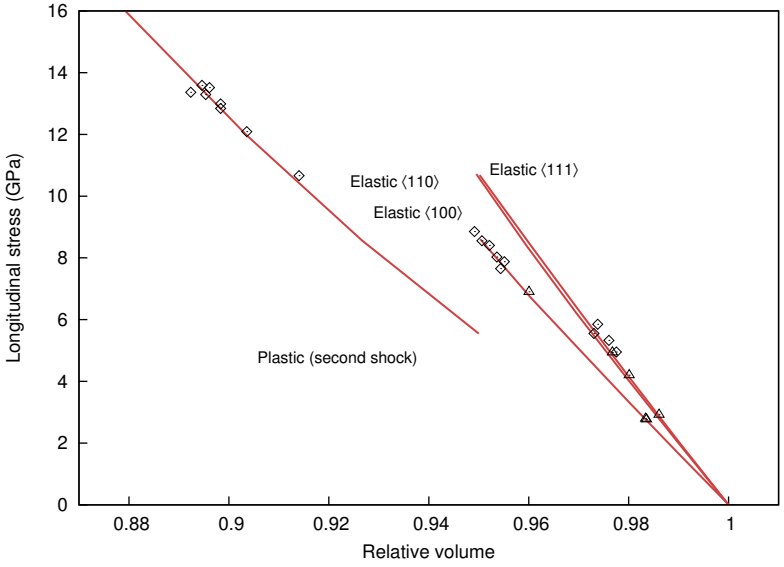

Figure 5.10.: Pressure–volume Hugoniot loci for silicon. This is a similar plot to fig. 5.9, with the meaning of the symbols and lines the same, emphasizing the small strain region of the Hugoniot locus and with the results of Gust and Royce (1971) omitted due to their larger variance.

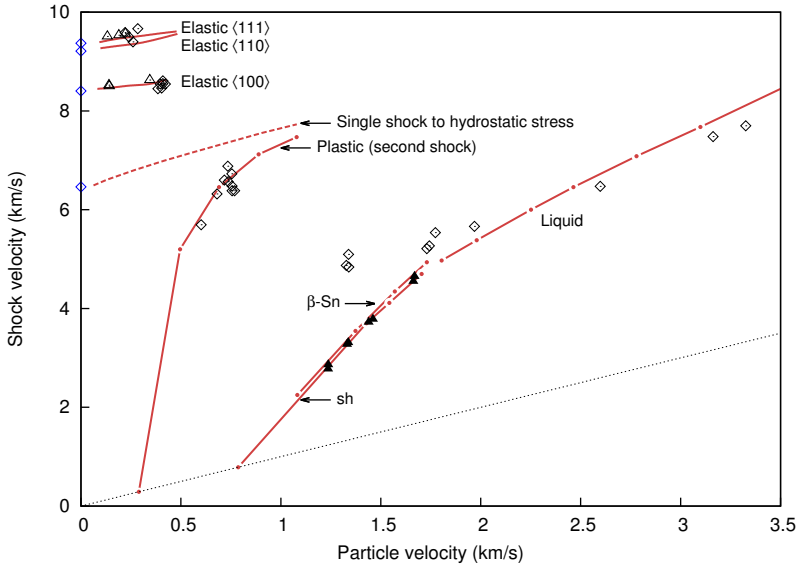

Figure 5.11.: Shock velocity–particle velocity Hugoniot loci for silicon. The DFT results (solid red lines and points) each correspond to an initial state of zero pressure and 300 K, with the final state in the indicated phase. The dashed line is for a single-shock process whose final state has a hydrostatic stress configuration. The meaning of the symbols is the same as in fig. 5.9, with the blue diamonds on the axis the elastic and bulk wave speeds from Goto et al. (1982) The dotted base line indicates equal shock and particle velocity, below which no viable shock should be recorded.

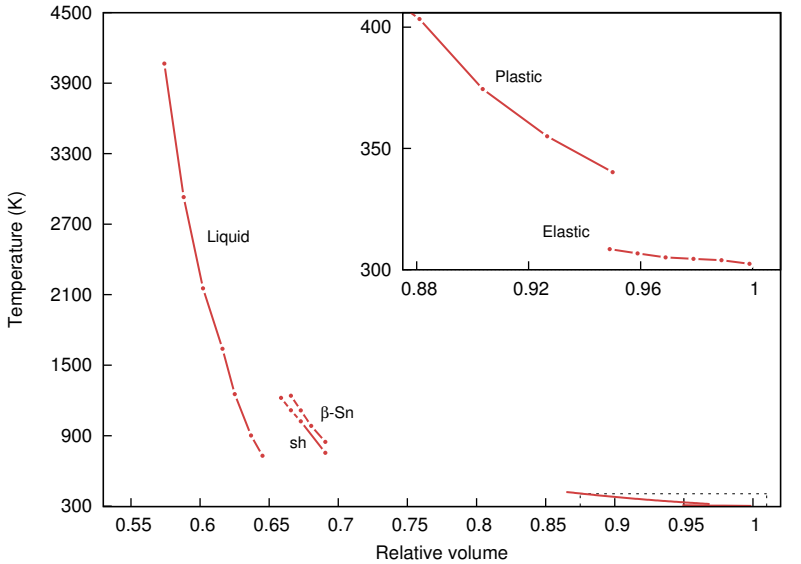

Figure 5.12.: Post-shock temperature as a function of volume for several final states. The DFT results (solid red lines) each correspond to an initial state of zero pressure and 300 K, with the final state in the indicated phase. The 'plastic' curve does not include the temperature rise due to dissipative heating. The meaning of the symbols is the same as in fig. 5.9

Table 5.1.: Coefficients of a linear fit the shock velocity for the elastic waves computed by the DFT anneal, $U_s = c_0 + s u_p$.

	⟨100⟩		⟨110⟩		⟨111⟩		Bulk	
	c_0 (km/s)	s (-)	c_0	s	c_0	s	c_0	s
This work	8.38	0.42	9.21	0.57	9.34	0.57	6.51	1.18
Goto et al. (1982)	8.42	0.32	9.24	1.01	9.39	0.98	-	-
Hall (1967)	8.43	-	9.13	-	9.34[a]	-	6.48[a]	-

[a] Calculated from the given elastic constants and density.

The β-Sn and simple hexagonal (sh) curves each correspond to a three-wave split shock structure, behind an elastic wave to the experimental elastic limit of 6 GPa,[1] and a second wave to a pressure close to the experimental location of the phase transition at 13.5 GPa.[2] For both of these waves, the computed volume for the ⟨100⟩ direction was used for the post-shock state. In general, it is quite insensitive to the precise location of the wave split, particularly for the elastic case, since the contribution to the energy change is much smaller than the 20 % volume reduction across the phase change. The final stress was hydrostatic. Since the c/a-ratio is free in the β-Sn and simple hexagonal structures, an additional relaxation step was used on the simulation box to impose a hydrostatic distribution of stress while simultaneously annealing to the Hugoniot. The β-Sn and simple hexagonal curves are close in pressure, temperature and shock velocity, with the experimental values closest to the simple hexagonal DFT Hugoniot. The computed pressures and temperatures of these points put them in stable region for the simple hexagonal structure on the silicon phase diagram (Kubo et al., 2008).

Part of the liquid Hugoniot corresponds to a three-wave shock structure, with the third wave reaching the final liquid state, behind a secondary wave to the onset of the melting transition and an elastic precursor wave. For the highest pressures, where the final wave has a velocity greater than that of the secondary wave of 6.83 km/s, it instead corresponds to a two-wave structure (behind only the elastic precursor). The largest shock pressures closely match the calculated liquid Hugoniot, with the simulated liquid being systematically slightly too stiff.

[1] This value is in fact rather small for the ⟨100⟩ shock, and would be more appropriate for the ⟨110⟩ or ⟨111⟩ shocks, with a measured HEL of 5.5 GPa.

[2] This value was used as it was the closest of the simulated points to the experimental value of 13.8 GPa.

Table 5.2.: Summary of values used at the onset of the cubic diamond to liquid phase transition. The phase line is as reported by Kubo et al. (2008). The other values are from Hull (1999), with α and c_p at 1600 K and ambient pressure, and β at 298 K and 13.8 GPa.

T (K)	dT/dP (KGPa^{-1})	α (K^{-1})	β (GPa^{-1})	c_p (Jg^{-1}K^{-1})
1683	62.4	4.5×10^{-6}	0.024	1.0

The predicted post-shock temperatures (given in fig. 5.12) indicate that these highest pressure points are likely to be liquid phase. The sixfold coordinated liquid lies close in p–v to the Hugoniot for the beta-tin phase, and so this phase transition does not exhibit the large mixed phase region as for the diamond to dense-phase silicon.

There is a considerable range of relative volume between the Hugoniot loci of the pure phases shown in fig. 5.9. The experimentally measured points in this region have a final state that is a mixture of two phases. Points on the mixed-phase region of the Hugoniot are on the intersection of the phase boundary for the two phases, as well as satisfying eq. (5.3).

We consider here two possible phase transitions starting from silicon in the cubic diamond structure: to a liquid, and to the beta-tin structure. In addition, we assume that the onset of either transition occurs at 13.5 GPa, close to the observed experimental value. The phase lines are experimental values, obtained by Kubo et al. (2008) This gives the two dashed lines appearing in fig. 5.9. The lower, green dashed line for diamond structure to beta-tin is nearly at constant pressure, since its slope is dominated by the steep phase-line of the transition $dT/dp = -1426$ K/GPa (Kubo et al., 2008). This is consistent with the experiment of Turneaure and Gupta (2007b) The upper, blue dashed line for melting the diamond structure is influenced most strongly by the compressibility β of the cubic diamond phase at the pressure and temperature of the onset. Representative literature values for the constants appearing in the above expression for the liquid are summarized in table 5.2. This line underestimates the experimentally observed slope seen by Gust and Royce (1971) and Goto et al. (1982) While the simulated temperature at this pressure is much too low for melting, the simulations of the 'plastically-yielded' state do not include dissipative heating and this could cause a considerable temperature rise above those reported in fig. 5.12.

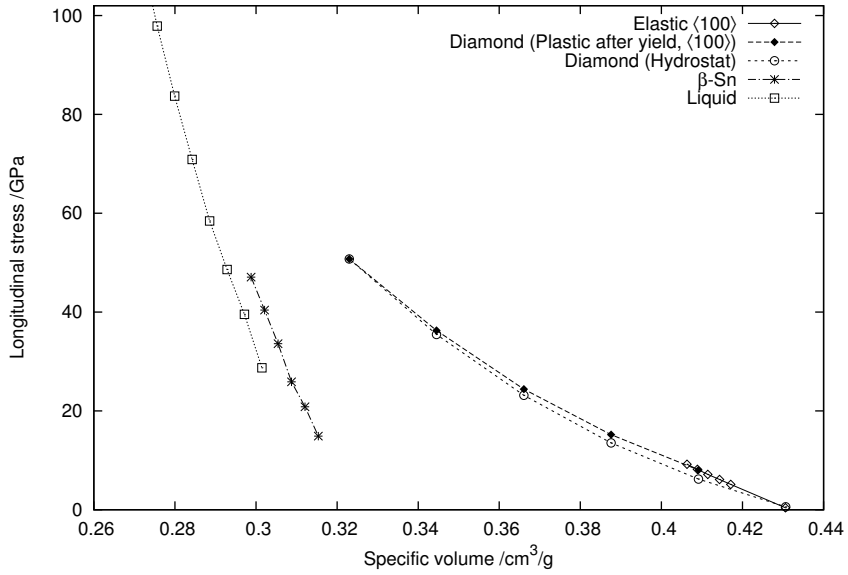

Figure 5.13.: p–v Hugoniot for Tersoff silicon (Tersoff, 1988, parameter set 'D') computed with the annealing method. Each set of points corresponds to an initial state of 0.6 GPa and 300 K, with the final state in the indicated phase.

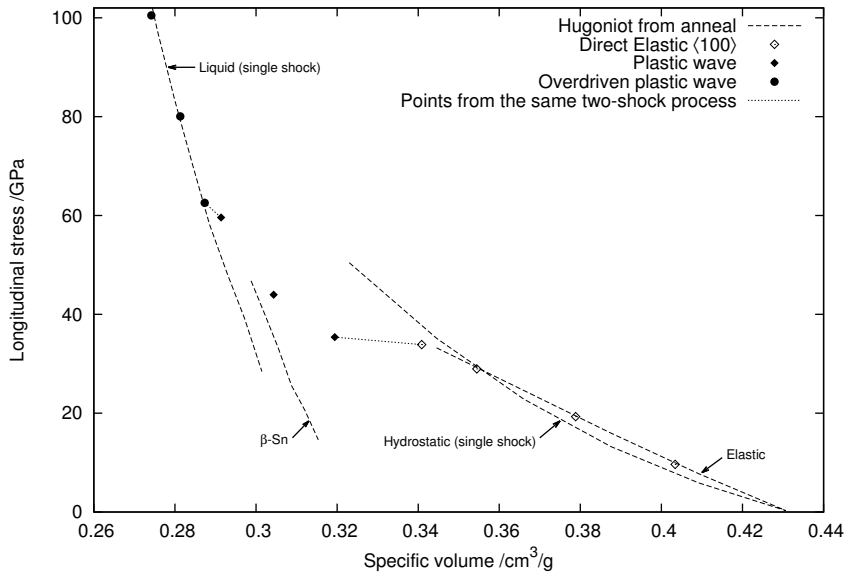

Figure 5.14.: Points from direct shock simulation of Tersoff silicon. The lines in the figure are taken from the anneal (fig. 5.13) for comparison. Elastic, two-wave elastic-plastic and overdriven single-wave results are shown.

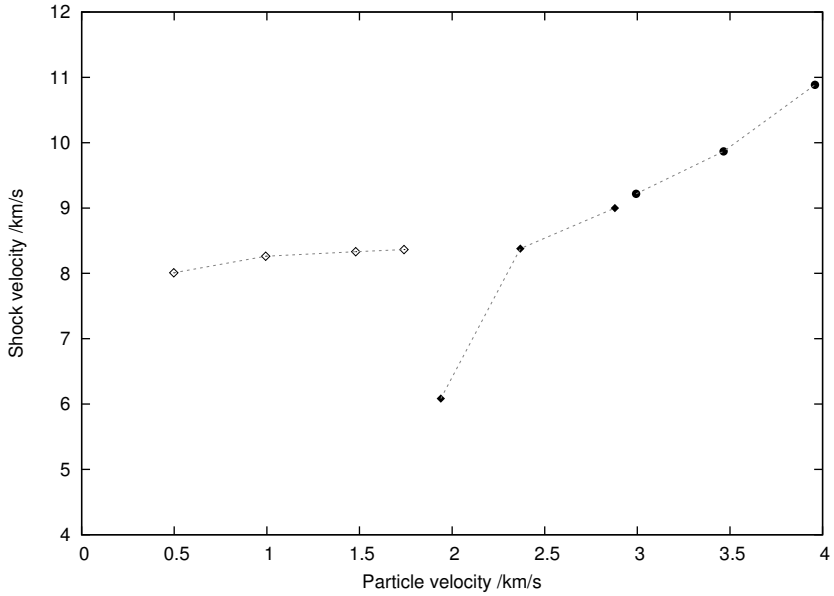

Figure 5.15.: Shock velocity–particle velocity Hugoniot locus for Tersoff silicon, showing an elastic wave and plastic behaviour with simultaneous phase transition to a sixfold coordinated structure. The symbols have the same meaning as in fig. 5.14.

5.5.2. Results for the Tersoff and Stillinger–Weber potentials

Figure 5.13 shows a pressure–volume Hugoniot for silicon as described by the Tersoff potential (Tersoff, 1986, 1988), also given in appendix C.2. Note immediately that the dense β-Sn and liquid phases overestimate the pressure by a great deal compared with experiment, and both are much stiffer. This is due to the dominant term in the expression for the Tersoff potential at close range (eq. (C.7)), and with the potential far outside of the density range for which it was parameterized. The liquid curve is a second shock, with an initial elastic precursor. The β-Sn phase curves shown are computed for tertiary shocks, following the elastic precursor and plastic wave.

The annealed Hugoniot for the Tersoff potential can be compared with points computed directly using the same potential, in fig. 5.14. The direct points were obtained by shocking a sample of $10 \times 10 \times 100$ unit cells in the [001] direction using the flyer-plate approach of section 5.3.

Tersoff is much softer than Stillinger–Weber silicon (Stillinger and Weber, 1985), described below. This is especially true in the 30 GPa regime, due to high pressure liquid being denser than the diamond phase, although it is still somewhat stiffer than experiment. The elastic points computed directly, and the points from the shocks resulting in melting,

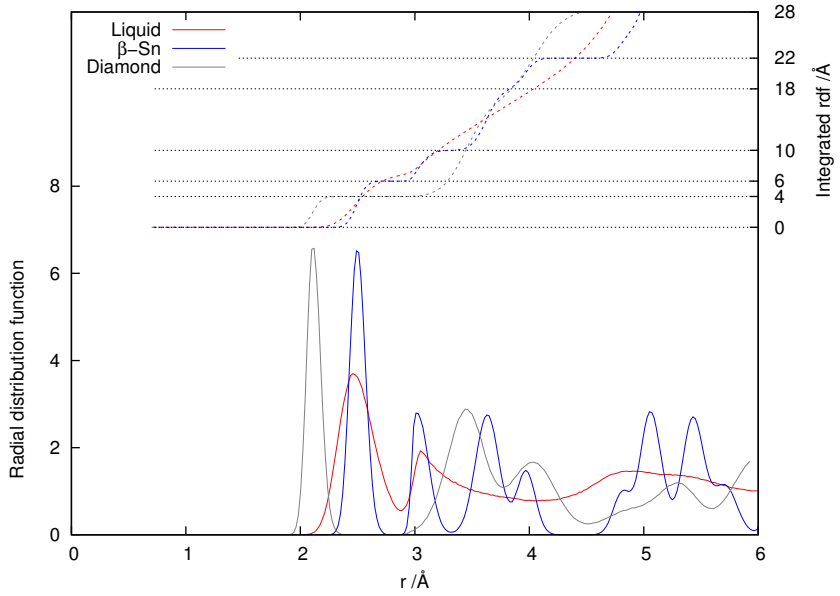

Figure 5.16.: Radial distribution of liquid tersoff silicon

match the annealed result. Typically for this type of simulation, the elastic precursor drives the material to a much higher stress than its experimental HEL, and this also increases with the shock strength.

An unusual feature of Tersoff silicon apparent from these simulations is that for a large enough shock pressure, the uniaxial elastic Hugoniot lies beneath the hydrostat. That is, the material is softer under a uniaxial compression than an isotropic one. It is only significant in these direct simulations because of the very large elastic limit. The plastic split-wave profiles show a sharp elastic precursor followed by a broader yield. In these cases, yielding resulted in a denser substance than the hydrostatically compressed diamond phase.

Figure 5.17 is the p–v Hugoniot using the potential of Stillinger and Weber (1985) (SW), given in appendix C.1. The potential itself is designed to favour a tetrahedral coordination, and there are known problems with describing the liquid and amorphous phases (Vink et al., 2001). It is nevertheless one of the most used and heavily studied potentials, including its phase behaviour (e.g. Broughton and Li, 1987).

The annealed Hugoniot for the SW potential can be compared with points computed directly, in fig. 5.18. The direct points were obtained by shocking a sample of $10 \times 10 \times 100$ unit cells in the [001] direction using the flyer-plate approach, as with Tersoff described above. The lines in this plot are the Hugoniot anneal data from fig. 5.17. The lowest extent of the liquid line is the lowest pressure of a single-shock process where the resulting

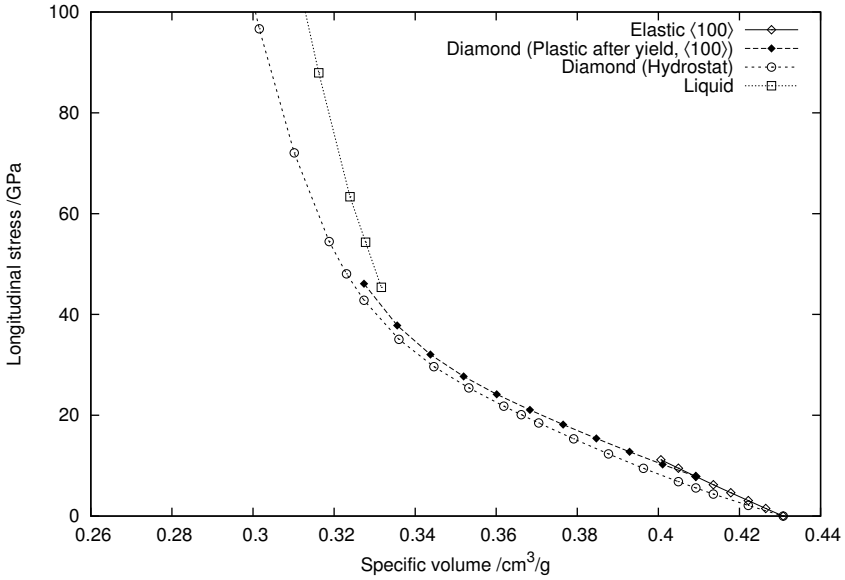

Figure 5.17.: p–v Hugoniot for Stillinger–Weber silicon (Stillinger and Weber, 1985). Each set of points corresponds to an initial state of zero pressure and 300 K, with the final state in the indicated phase.

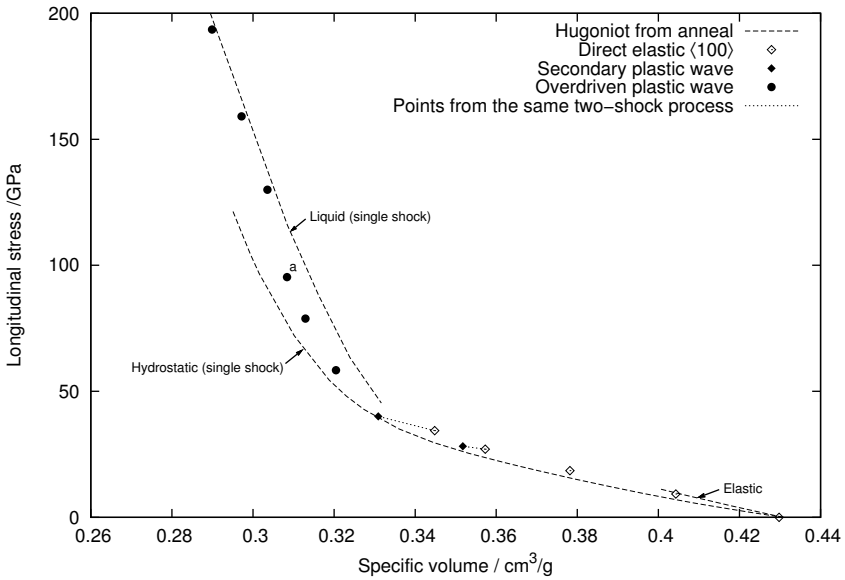

Figure 5.18.: The p–v Hugoniot of SW silicon, with the points from a direct shock simulation. The lines in the figure are taken from the anneal (fig. 5.17) for comparison. Elastic, two-wave elastic-plastic and overdriven single-wave results are shown. The point labelled 'a' is in a mixed-phase region, and is discussed in the text. See also fig. 5.20.

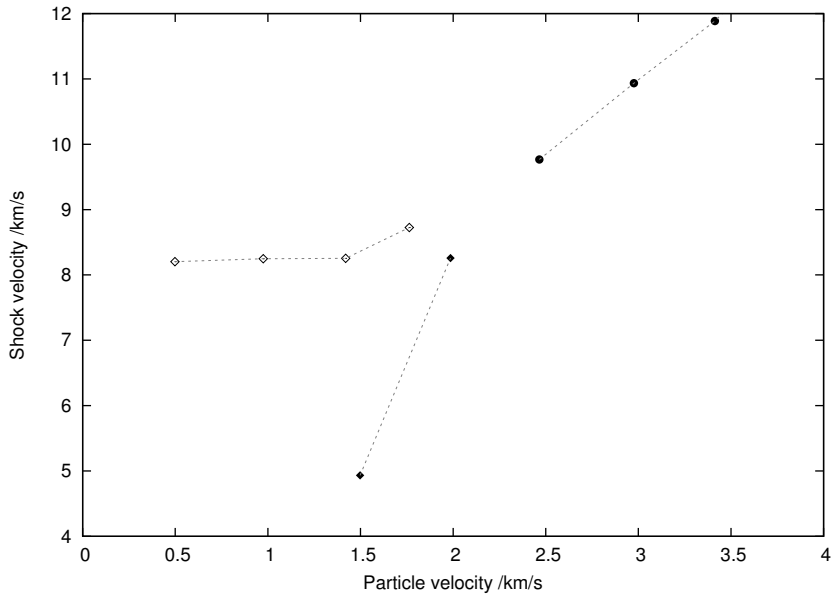

Figure 5.19.: Shock velocity–particle velocity Hugoniot locus for SW silicon, showing an elastic wave and plastic behaviour with simultaneous phase transition to a sixfold coordinated structure. The symbols have the same meaning as in fig. 5.18.

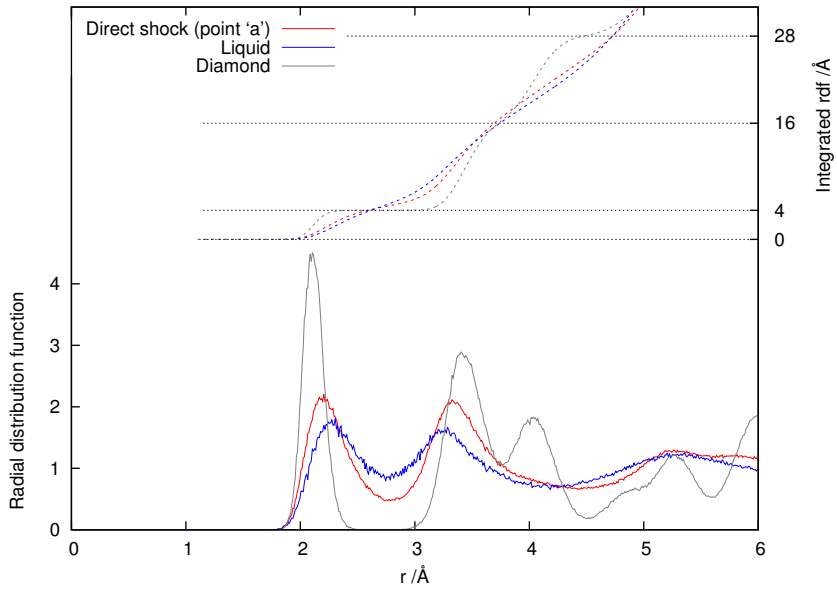

Figure 5.20.: Radial distribution functions for the point labelled 'a' in fig. 5.18, along with the RDF for diamond structure silicon and liquid silicon at the same specific volume on the Hugoniot locus.

structure does not spontaneously solidify. Likewise, above the given extent of the hydrostatic line the diamond structure spontaneously melted.

The data extracted from the shock MD trajectory were the shock speed (as the mean velocity of the point on the front where the mean particle velocity of a slice normal to the shock direction was 20% of its post-shock value), the particle velocity (as its mean value well behind the shock, and before any secondary shock), the shock stress (in a similar way, using the virial stress, and correcting for the bulk motion behind the shock).

We can check that the shock and particle velocities computed from the stress–volume relation and the Hugoniot relations eqs. (5.4) and (5.6) are consistent with those measured directly in the simulation. In general, for the size of simulation used here, the elastic shocks are consistent to within 0.5%, the strongest single shocks, involving a phase transformation, to 1%. The worst case for the direct simulation has split waves of very different velocities. The worst disparity was 12%, for the slowest secondary plastic wave in Tersoff silicon.

Note that due to a slightly different specific volume used for the initial structure, the elastic points lie just below those from the anneal (by about 0.4GPa, corresponding to a slightly lower wave speed.

The final split wave in fig. 5.18 is the limiting case of the plastic wave travelling just slower than the elastic one, as can be seen considering their slope, or from fig. 5.18. When repeating this for a larger simulation ($40 \times 40 \times 200$ unit cells), the yielding occurred after an initial delay, but eventually caught up with the elastic front, forming a single extended shock front, shown in fig. 5.21. This steady shock consists of an initial sharp front (with considerable dispersion) accounting for most of the change in particle velocity, and leaving the material in an initial uniaxially compressed state. This is followed by a gradual rise in particle velocity accompanied by permanent plastic deformation of the material.

In this figure, to obtain a more resolved profile, we suppose that the average profile is slowly evolving compared with the speed of the shock, and average the velocity across many frames, shifted according to the average shock speed. The result is for a system of $100 \times 100 \times 200$ unit cells (sixteen million atoms). This is an average over 40 frames with 250 timesteps between each frame, after the profile has been allowed to evolve for 10ps.

Some features to note:

- The HEL at about 30GPa is much larger than the experimental value.

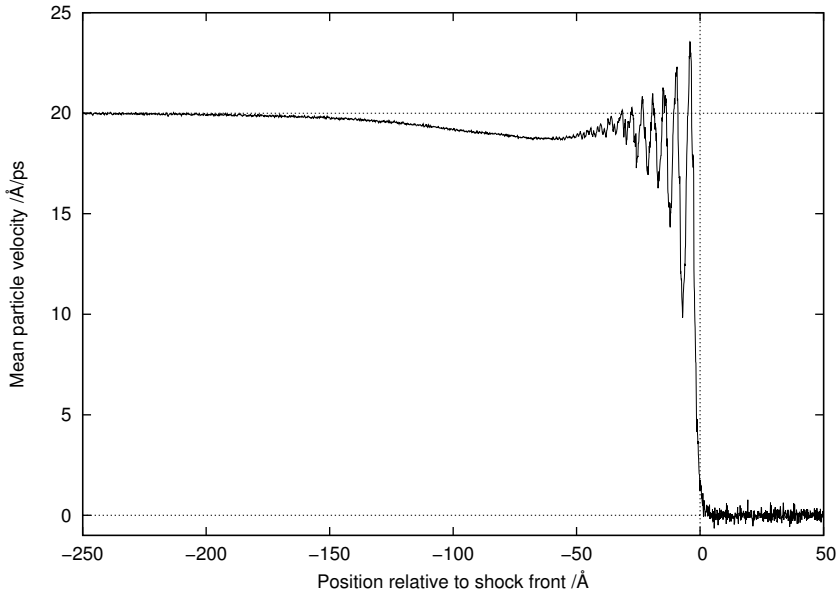

Figure 5.21.: Mean 'steady' velocity profile for the limiting single wave plastic shock, relative to the shock front. The figure was produced by identifying the mean shock speed and averaging over a number of frames of the simulation translated at this velocity to keep the shock position stationary at zero. The velocity on the vertical axis is relative to the pre-shocked material.

- The shock pressure of the precursor shock varies strongly with the final shock pressure. The range of volume compression over which a two-wave profile results is relatively small.
- The points resulting from a yielded diamond structure lie just above the calculated hydrostatic curve, as would be expected from both a residual shear strength and heating from dissipative yield. The stress profile behind the stronger shocks is completely hydrostatic.
- As the shock strength is increased, the direct locus deviates from the hydrostat, and for the strongest shocks lies on the liquid curve.
- The point labelled 'a' in fig. 5.18 has a structure intermediate between a diamond and liquid. The final structure has separated regions of yielded diamond structure and amorphous structure very close to the liquid structure whose RDF is given in fig. 5.20.
- The shock profile shows considerable dispersion at the elastic wave-front.
- The high pressure liquid phase of silicon shows relatively little deviation from a tetrahedral structure, as evidenced by the RDF in fig. 5.20, although there is some fivefold coordination.

5.5.3. Direct NEMD shock simulation from first-principles

A silicon sample, consisting of $2 \times 3 \times 20$ conventional unit cells of silicon (960 atoms), was obtained from a equilibrated Tersoff simulation, with the lattice parameter scaled from the 5.43 Å of Tersoff to 5.49 Å, which was the 300 K lattice parameter found from SIESTA. This was done leaving the fractional coordinates of the atoms, and their absolute velocities, unchanged. The resulting system was found to be close to zero pressure and an equilibrium velocity distribution in SIESTA, but was further annealed using a Berendsen thermostat for just under 500 fs, and then integrated for a further 500 fs using Verlet.

The functional, pseudopotential and basis set were as described in section 5.5.1. The real space grid also used a cutoff of 100 Ry. However, only a Γ-point calculation was possible (even a single non-gamma-point calculation, such as at the Baldereschi point, would have entailed considerable extra computation).

The shock wave was then initiated in the well-equilibrated structure. To generate a shock in the material, the first surface layers of atoms (one conventional unit cell deep, 48 atoms in total) were fixed to move at a constant velocity in the z direction (and zero

velocity in the transverse directions), and independently of the dynamics. The particle velocity was chosen to be 0.36 km/s, which resulted in an elastic shock of just over 8 GPa when using the Tersoff potential. The experimental Hugoniot elastic limit was found to be 9±1 GPa in the ⟨100⟩ direction (Gust and Royce, 1970), occurring at a volume compression of 0.946, so this choice of shock strength for the simulation is just before the experimental yield.

A snapshot of the simulation when the shock is about halfway through the domain is shown in fig. 5.22. Also shown is the extracted location of the shock, taken to be the location where the mean z-velocity in a moving window first exceeds 2 nm/ps. The shock location extracted in this way, against time, is shown in fig. 5.23, from which the shock speed can be found to be $8.39 \pm 0.013\,\mathrm{km\,s^{-1}}$. The shock reached the end of the domain after a simulation time of 1.2 ps. The total CPU time for the simulation was 32 000 core-hours.

As before, to obtain a more resolved profile than the instantaneous snapshot of fig. 5.22, we suppose that the average profile is slowly evolving compared with the speed of the shock, and average the velocity across many frames, each shifted according to the average shock speed, resulting in fig. 5.24. This means that the part of the profile closest to the shock has smallest error, since it is present in every frame, while parts further away consist of averages over fewer and fewer frames. Very finely resolved features, such as those appearing in fig. 5.21, are not visible unfortunately, due to the much smaller simulation size, although we can conclude that the shock front is less than 2 nm thick, resulting in a rise time of 0.24 ps, or a strain rate of $2 \times 10^{11}\,\mathrm{s^{-1}}$.

Finally, we provide fig. 5.25, combining the shock speed measurements from the direct point with the first-principles annealing method, and the results from the empirical potentials. Note that the value for the direct point is lower than the annealed points due to the lack of a fully converged k-grid, although this is seen to agree with the anneal when performed using an equivalent k-grid. The Γ point offers enough resolution in the longitudinal direction of the simulation cell, of 20 unit cells, but not in the much narrower transverse dimensions, of 2 and 3 unit cells.

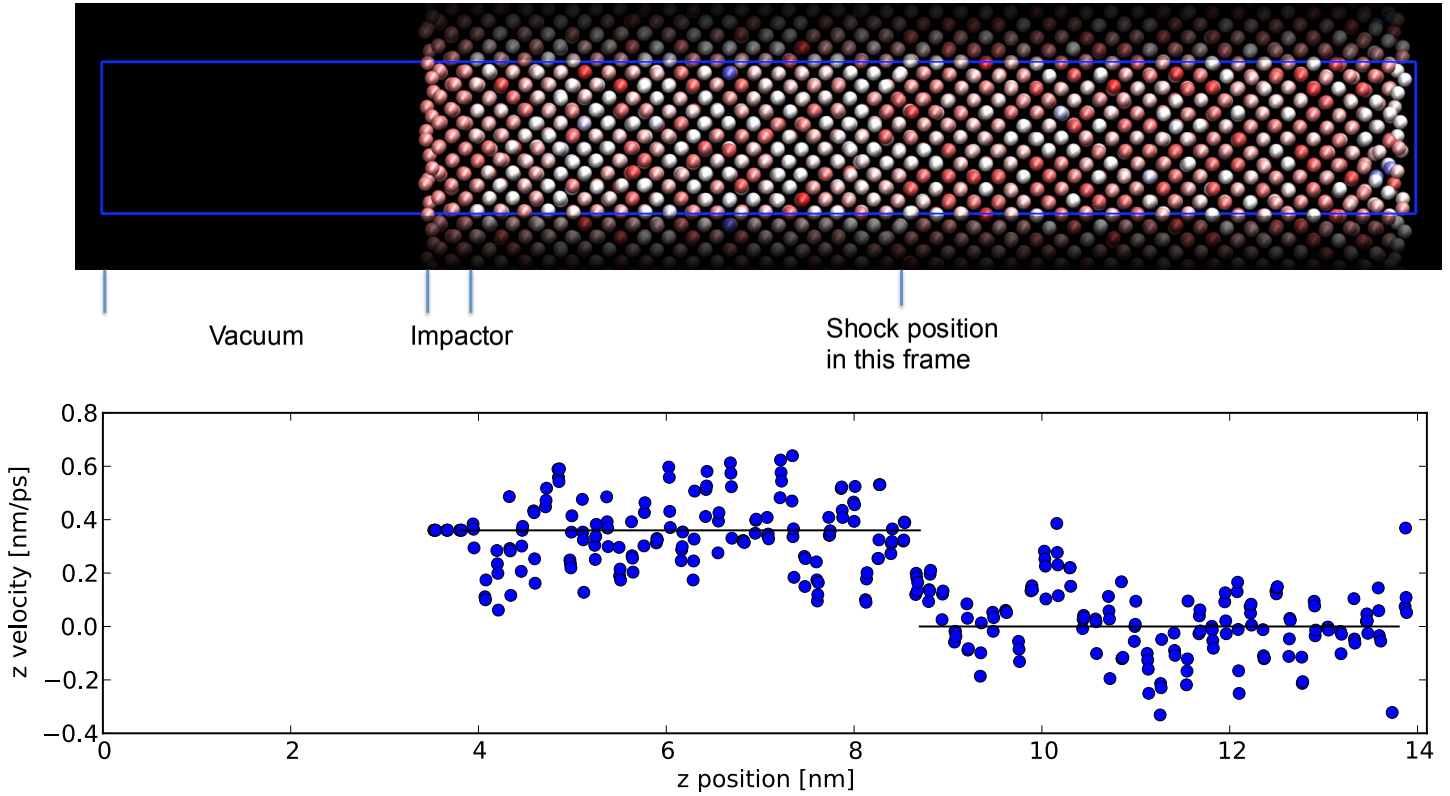

Figure 5.22.: A cross-section in yz of one frame of the directly simulated DFT elastic shock, which is travelling parallel to the z direction. The top figure shows the positions of the atoms within the domain, coloured by their instantaneous z-velocity (red to blue from slow to fast). The blue outline in this figure represents the boundary of the domain. Periodic boundary conditions were used and the periodic image atoms are shown faintly either side of the box. The bottom figure is the mean velocity of the atoms projected onto the z axis. The identified shock location and the mean velocity either side are indicated, along with the impactor atoms.

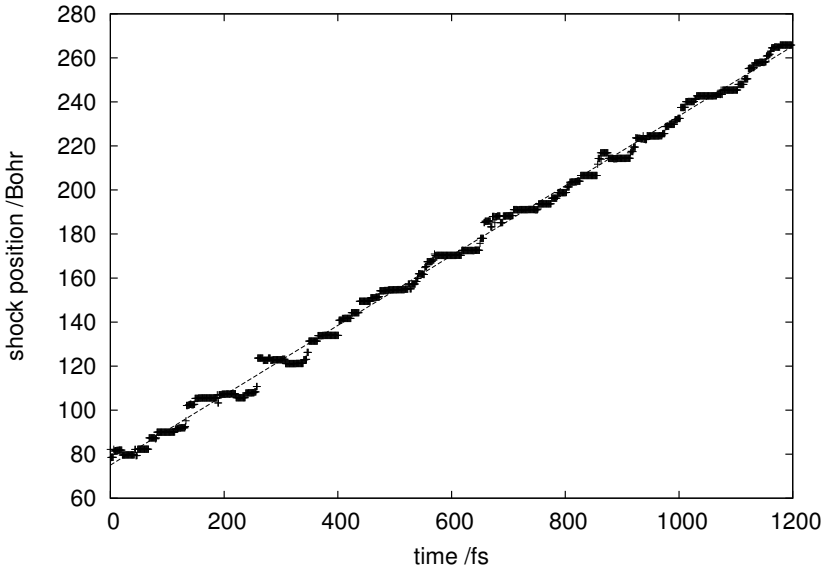

Figure 5.23.: The position of the shock wave with simulation time, for the direct DFT elastic shock simulation. The steady shock speed is extracted as $8.39 \pm 0.013\,\mathrm{km\,s^{-1}}$.

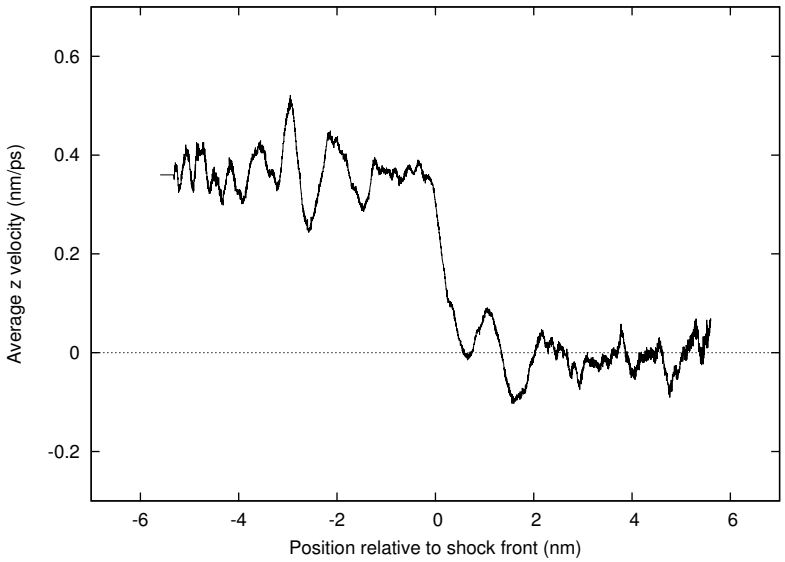

Figure 5.24.: Mean 'steady' velocity profile for the direct DFT elastic shock simulation, relative to the shock front. The figure was produced by identifying the mean shock speed and averaging over a number of frames of the simulation translated at this velocity to keep the shock position stationary at zero. The velocity on the vertical axis is relative to the pre-shocked material.

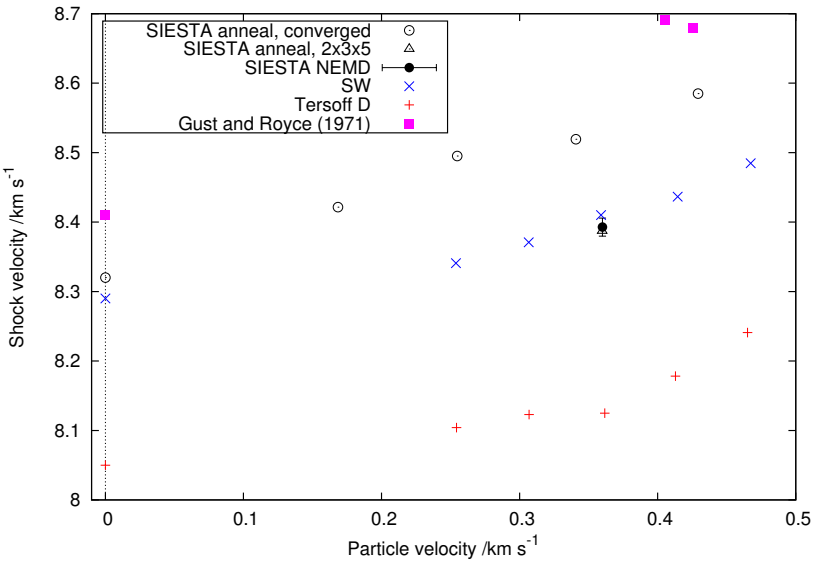

Figure 5.25.: Summary of elastic wave speeds for the empirical interatomic potentials considered here, and the first-principles results (direct and annealed). The single point representing the direct shock simulation in SIESTA (solid circle) partially obscures the annealed point (triangle) which used an equivalent sampling of reciprocal space.

5.6. Conclusions

In conclusion, we have described a simple annealing method and shown that it may be used to obtain a state on the Hugoniot locus of a pure phase of a material with several condensed phases efficiently, from first-principles. An approximation relying on the slope of the phase boundary can be used to obtain the part of the Hugoniot corresponding to coexistence between two phases.

In the case of silicon, the results computed using this procedure with the forces described using density functional theory match existing experimental data very well for pressures up to 60GPa, the limit of available experimental data. We have provided a prediction of the shock temperatures of silicon over this pressure range. This study supports the conclusions of the experimental work in general, that silicon after yield supports no deviatoric stress, and of Turneaure and Gupta (2007b) that the first observed phase transition along the shock locus is likely to be to simple hexagonal.

The method was also applied to two empirical interatomic potentials commonly used to simulate silicon, which can then be shocked directly. The direct method and the anneal produce equivalent points on the Hugoniot for the pure phases, and the direct simulations also give points in mixed-phase regions. In addition, we provided a direct-shock calculation for DFT silicon, giving a single point on the elastic part of the Hugoniot locus, which matches an equivalent point computed by the annealing method, and also the wave profile, which for the elastic shock is very sharp.

It was observed that the wave profile for elastic shock waves (single waves or precursor shocks) exhibit some degree of dispersion. This can be noticed in a single frame of the longitudinal velocity profile of the shock wave as an initial overshoot above the averaged post-shock wave speed, but is made most visible by taking a steady time-average over many frames of the simulation. The theory of dispersive waves in lattices is covered by Ercolani (1994).

6. Computation of Lattice Thermal Conductivity from Molecular Dynamics Simulation

6.1. Introduction and Motivation

In a solid, heat is conducted by interacting lattice vibrations (phonons) and electrons. For insulators, the electronic contribution to the thermal conductivity is negligible, while in metals it is dominant (Ho et al., 1978; Srivastava et al., 1970). In a semiconductor, the electronic contribution becomes important at high temperature. In silicon it contributes around 10 % of the total value at 1000 K (Shanks et al., 1963), and < 0.1 % under 700 K.

Engineering uses of materials for their thermal conductivity are widespread and well-known, with materials with particularly high conductivity found where required to transmit large quantities of heat (high-power electrical components and heat sinks, e.g. Ghosh et al. (2008)), and low values used for thermal insulation (for example, thermal barrier coatings on gas-turbine parts, (Padture et al., 2002)), and for efficient thermoelectrics (Tritt and Subramanian, 2006).

The aim of this chapter is to compute the temperature-dependent lattice thermal conductivity from molecular dynamics simulation. It is found to be a demanding test of an empirical potential, with a large variation in lattice thermal conductivity from many existing empirical interatomic potentials that represent the same material. We investigate the ability of a Gaussian Approximation Potential (GAP) to reproduce the thermal conductivity.

We choose silicon as a test material, due to the extensive body of work on potential development and electronic structure calculation for silicon in general, as well as being well studied for thermal conductivity. The high thermal conductivity of silicon makes obtaining an accurate value difficult, since this implies long phonon mean free paths, which must be captured in a simulation.

An overview of two methods commonly used to compute thermal transport properties of materials is given below. Section 6.4 describes a method based on solving the Phonon

Boltzmann Transport Equation (PBTE). The Green–Kubo method of linear response is considered in section 6.5. A method not considered here is a direct non-equilibrium simulation. These are described in detail by Broido et al. (2005) and He et al. (2012).

The computational requirements for the direct non-equilibrium and linear-response methods can be demanding for materials with high conductivity. Direct non-equilibrium methods typically require systems with (at least) thousands of atoms, simulated for hundreds of picoseconds. Methods based on linear response can use smaller systems, but often need several nanoseconds of data. The requirements depend on the thermal conductivity itself, with more conductive materials requiring larger and longer simulations to take into account the longer relaxation times and lengths present. Because of this, both techniques are currently beyond the reach of DFT for silicon. First-principles forces (such as DFT) can be used to compute force constants for PBTE, although this is limited to relatively low temperatures.

Table 6.1 summarizes a selection of experimental measurements of thermal conductivity for high temperature (>300 K) silicon reported in the literature. A non-uniform isotopic composition provides an additional mechanism for scattering phonons. Experiments for both the natural isotopic distribution (given in table 6.2) and highly pure silicon-28 are described. Those for naturally occurring silicon agree within around 5 %, which is the error given by Glassbrenner and Slack (1964). The results of this study were systematically slightly larger than the others, by 5 to 10 % across the whole temperature range. Ruf et al. (2000) reported a thermal conductivity for pure silicon-28 that is anomalously large, 60 % larger than the natural distribution. This was later noted by the original authors as being in error (Ruf et al., 2003) and confirmed by Inyushkin et al. (2004), who obtained a value only 10 ± 2 % larger than naturally occurring silicon. Including an extra scattering term in Boltzmann transport equation calculations (with force constants from DFT) reproduces this 10 % difference (Ward et al., 2009).

Many thermal conductivity calculations have been performed using empirical interatomic potentials for silicon. He et al. (2012) offer a review of published experimental and simulated thermal conductivities for silicon, germanium and silicon germanide, and present simulations for the Tersoff potential from PBTE, Green–Kubo MD and direct non-equilibrium MD. At the level of Boltzmann transport, Broido et al. (2005) considered three commonly used potentials for silicon (Tersoff (1986, 1988), Stillinger and Weber (1985) and the environment dependent potential of Bazant et al. (1997) and Justo et al. (1998)). They found that none of the three potentials reproduced the experimental conductivity between 0 and 300 K, and that in addition, agreement between Grüneisen parameters and thermal expansion was not a good indicator of agreement between thermal conductivities.

Table 6.1.: Summary of high-temperature experimental and first-principles thermal conductivities of silicon from the literature.

Reference	Method	System	Temp. (K)	Therm. cond. ($Wm^{-1}K^{-1}$)
Shanks et al. (1963)	experiment	Si nat.	300	142
			400	97
			600	58
			800	40
			1000	30
Glassbrenner and Slack (1964)	experiment	Si nat.	300	156±8
			400	105±5
			600	64±3
			800	43±2
			1000	31±1.5
Ho et al. (1972)	experiment	Si nat.	300	148
			400	99
			600	62
			800	42
			1000	31
Ruf et al. (2000)	experiment	Si-28 (99.86 %)	300	237±8 [a]
Kremer et al. (2004)	experiment	Si nat.	300	144±3 [b]
		Si-28 (99.96 %)	300	157±3 [b]
Inyushkin et al. (2004)	experiment	Si nat.	300	143
		Si-28 (99.98 %)	300	156
Ward et al. (2009)	DFT LDA, PBTE-3rd	Si nat.	296	135
			350	112
		Si-28	296	146
			350	119
Garg (2011)	DFT LDA[c], PBTE-3rd	Si-28	300	136

[a] Inyushkin et al. (2004) suggest that this result is an overestimate.

[b] An average across three samples at three separate laboratories

[c] Garg, private communication

Table 6.2.: Naturally occurring silicon nuclides

nuclide	mass (u)	natural abundance (%)
^{28}Si	27.977	92.2
^{29}Si	28.976	4.7
^{30}Si	29.974	3.1

Schelling et al. (2002) compared direct (non-equilibrium), and Green–Kubo (equilibrium) methods for computing the lattice thermal conductivity of Stillinger–Weber silicon at 1000 K, including a consideration of finite-size scaling effects for both methods. Abs da Cruz et al. (2011) computed the high temperature thermal conductivity (500 to 950 K) from several empirical potentials using the direct non-equilibrium method, the results ranging over more than a factor of two between potentials at the same temperature.

Broido et al. (2007) computed an iterative solution to the exact PBTE for bulk silicon and germanium, with force constants derived from density functional perturbation theory, and report excellent agreement with experiment for both materials, over a temperature range of 0 to 300 K. Several solutions to various levels of approximation of the PBTE have been determined in the literature. A common approach is to assume that each phonon mode decays towards the equilibrium distribution with a characteristic relaxation time. A number of authors have taken this approach, including Callaway (1959), and more recently by Kazan et al. (2010).

6.2. Continuum Picture

Denote the **heat flux** at a point in a continuum body by $\boldsymbol{q}$ (units $\mathrm{W\,m^{-2}}$). Fourier's law of heat conduction states that the heat flux depends linearly on the gradient of the temperature, T, and that heat flows down temperature gradients:

$$q_i = -\kappa_{ij}(\nabla T)_j. \tag{6.1}$$

The positive quantity κ_{ij} is independent of the temperature gradient, but in general will depend on the local temperature. For a material in which the heat conductivity is not direction dependent, the tensor κ_{ij} may be replaced by a scalar conductivity, κ. The remainder of this chapter assumes that this is the case.

Fourier's law leads directly to the parabolic heat equation. From continuity of energy at constant pressure, with h the specific enthalpy,

$$\frac{\partial h}{\partial t} + \frac{1}{\rho}\nabla \cdot \boldsymbol{q} = 0, \tag{6.2}$$

and the definition of heat capacity

$$\left(\frac{\partial h}{\partial T}\right)_p = c_p \tag{6.3}$$

we obtain

$$\frac{\partial T}{\partial t} - \frac{1}{c_p \rho}\nabla \cdot (\kappa \nabla T) = 0. \tag{6.4}$$

If κ does not vary in space (other than through its dependence on temperature), then after linearizing we can write

$$\frac{\partial T}{\partial t} - \alpha \nabla^2 T = 0 \tag{6.5}$$

with **heat diffusivity** $\alpha = \kappa/(\rho c_p)$.

6.3. Heat Conduction and Phonons

A model of interacting phonons explains many features of the temperature dependence of thermal conductivity, especially at low temperature (Ashcroft and Mermin, 1976; Srivastava, 1990).

Phonons are the normal modes of an infinite crystal whose internal energy is harmonic in deviations from the equilibrium atomic positions:

$$U = U_{\text{eq}} + \frac{1}{2}\sum_{\substack{\boldsymbol{R},\boldsymbol{R}' \\ \mu,\nu \in \{x,y,z\}}} u_\mu(\boldsymbol{R}) \mathsf{D}_{\mu\nu}^{\boldsymbol{R}\boldsymbol{R}'} u_\nu(\boldsymbol{R}'), \tag{6.6}$$

for some constants $\mathsf{D}_{\mu\nu}^{\boldsymbol{R}\boldsymbol{R}'}$, where $\boldsymbol{u}(\boldsymbol{R})$ is the deviation of the atom whose equilibrium position is $\boldsymbol{R}$ and U_{eq} is the internal energy at equilibrium.

This is often a good approximation to a real crystal, provided the deviations $\boldsymbol{u}$ are small. The thermal conductivity in such a perfectly harmonic crystal would be infinite,[1] however, and so we must also consider anharmonic terms.

[1] Since thermal transport would then obey a wave equation.

Perturbing the harmonic potential (6.6) with third and fourth[2] order anharmonic terms of the form (third order)

$$\frac{1}{6}\sum_{\substack{R,R',R''\\ \mu,\nu,\lambda}} \mathsf{E}^{R,R',R''}_{\mu,\nu,\lambda}\, u_\mu(\boldsymbol{R})u_\nu(\boldsymbol{R}')u_\lambda(\boldsymbol{R}'') \tag{6.7}$$

does not greatly affect normal modes of finite displacement, provided these terms are small, but can be seen as allowing transitions to occur between different normal modes. This is *phonon scattering*. Conservation of energy and crystal momentum govern the scattering processes that can occur. In particular, conservation of crystal momentum states that the crystal momentum before and after a scattering event must be related by

$$\sum n_{ks}\boldsymbol{k} = \sum n'_{ks}\boldsymbol{k} + \boldsymbol{K} \tag{6.8}$$

where $\boldsymbol{K}$ is a reciprocal lattice vector. If $\boldsymbol{K}$ is zero, the analogy with real momentum is direct and the process is said to be *normal*. Otherwise, it is said to be *umklapp*.[3] It can be shown (Srivastava, 1990) that it is only umklapp processes that disrupt any heat current present, so without them a crystal would have infinite thermal conductivity.

In this way, phonons can be treated in the same way as gas molecules in kinetic theory and a suitable Boltzmann equation derived whose scattering terms may be calculated from the third-order force constants E appearing in equation (6.7), and perhaps also the fourth-order constants. Once these are known, the linearized Boltzmann equation can be solved iteratively to obtain the thermal conductivity. This approach is suitable for low to moderate temperatures, and is used by Broido et al. (2005), but becomes less effective at higher temperatures when the deviations of the atoms from their crystal lattice sites are large and even-higher-order anharmonic terms become important.

Figure 6.1 illustrates the temperature dependence of the thermal conductivity of a typical crystal. The peak usually appears at a temperature of several tens of Kelvin.

Phonons obey Bose–Einstein statistics:

$$n_{ks} = \frac{1}{\exp(\hbar\omega_s(\boldsymbol{k})/k_B T) - 1}. \tag{6.9}$$

At high temperatures, this is asymptotically proportional to T for any given $\boldsymbol{k}$. The high temperature heat capacity asymptotes to a constant value, and the mean free path

[2] The number of allowed scattering interactions involving fourth order phonons makes up for their smaller magnitude, so they are often included. Also, a cubic potential results in an unstable Hamiltonian.

[3] Ger. *umklappen*, to fold down.

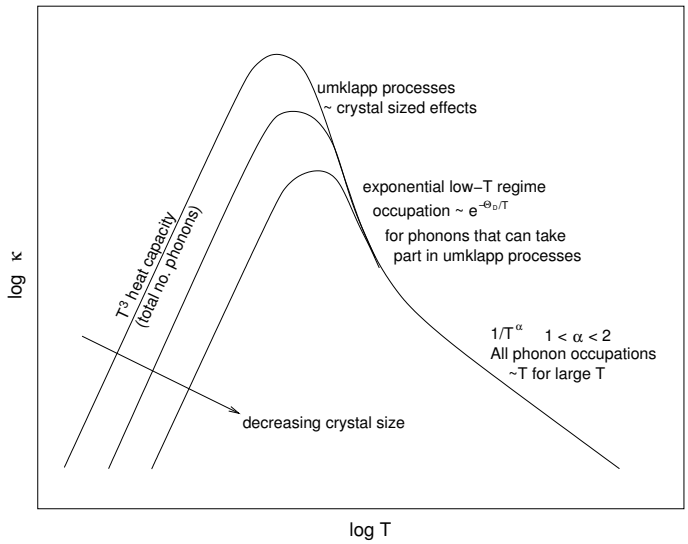

Figure 6.1.: Typical temperature dependence of the lattice thermal conductivity of a typical crystalline material.

of a phonon will vary inversely with the number of phonons present with which it may scatter. It is found that this results in an approximate power law for the thermal conductivity of $\kappa \propto 1/T^{\alpha}$ with $1 < \alpha < 2$.

For low temperatures, atomic displacements are small and harmonic and not much phonon scattering occurs. The conductivity is ultimately limited by the (constant) length of the crystal, or by other constant-length effects, such as the concentration of impurities. The low temperature heat capacity varies as T^3, and this is reflected in the thermal conductivity.

At these very low temperatures, only low-frequency, small wave-number phonons are present in considerable number, both before and after any collision, so umklapp processes occur only with very small probability. As the temperature increases, the occupation numbers of phonons able to participate in umklapp processes grows approximately as

$$\frac{1}{\exp(\Theta_D/T)-1} \approx \exp(-\Theta_D/T), \tag{6.10}$$

where Θ_D is the Debye temperature, and this is reflected in figure 6.1 as the sharp drop-off in κ after the T^3 behaviour. The peak occurs where the mean free path of phonons due to umklapp processes is of the same magnitude as the crystal dimensions.

6.4. The Boltzmann Transport Equation

This section follows Ward et al. (2009) and Li et al. (2014) in expressing the linearized Boltzmann equation iteratively. It is implemented in the latter authors' `ShengBTE` code.

The Boltzmann transport equation describes how the phase-space density of a thermodynamic system changes in time out of equilibrium. In general, this can be stated

$$\frac{\mathrm{d}f_\lambda}{\mathrm{d}t} = \left(\frac{\partial f_\lambda}{\partial t}\right)_{\text{scatter}} + \left(\frac{\partial f_\lambda}{\partial t}\right)_{\text{diffusion}} \tag{6.11}$$

where $\lambda = (j, \boldsymbol{q})$ with j the branch index and $\boldsymbol{q}$ the wavenumber.

For phonons, $f(\lambda, \boldsymbol{x}, t)$ is the distribution function of phonons, and for steady state, the transport and scattering terms cancel,

$$\left(\frac{\partial f_\lambda}{\partial t}\right)_{\text{scatter}} + \nabla T \cdot \boldsymbol{v}_\lambda \frac{\partial f_\lambda}{\partial T} = 0, \tag{6.12}$$

where $\boldsymbol{v}_\lambda$ is the group velocity of the phonon mode labelled by λ. For small temperature gradients, we can express the difference from the equilibrium distribution linearly,

$$f_\lambda = f_\lambda^0 + g_\lambda \tag{6.13}$$

and

$$g_\lambda = -\boldsymbol{F}_\lambda \cdot \nabla T \frac{\partial f_\lambda^0}{\partial T}, \tag{6.14}$$

where f^0 is the equilibrium phonon distribution, and $\boldsymbol{F}_\lambda$ is unknown. It is useful to express it as

$$\boldsymbol{F}_\lambda = \tau_\lambda^0 (\boldsymbol{v}_\lambda + \boldsymbol{\Delta}), \tag{6.15}$$

where we define $\boldsymbol{\Delta}$ later on. The *relaxation time approximation* (RTA) is equivalent to setting $\boldsymbol{\Delta} = 0$ in the above equation. The relaxation times τ_λ^0 are calculated in terms of the scattering rates (calculated in the equilibrium distribution) as

$$\frac{1}{\tau_\lambda^0} = \frac{1}{N}\left(\sum_{\lambda',\lambda''} \Gamma^{0,+}_{\lambda,\lambda',\lambda''} + \frac{1}{2}\sum_{\lambda',\lambda''} \Gamma^{0,-}_{\lambda,\lambda',\lambda''}\right). \tag{6.16}$$

$\Gamma^{\pm}$ are terms representing three-phonon scattering processes, with a single phonon λ scattering into two others (+) and two phonons λ' and λ'' colliding to form a single phonon (−), hence the factor of half multiplying this last term to avoid double counting.

These are given in terms of f as

$$\Gamma^{+}_{\lambda,\lambda',\lambda''} = \frac{\hbar\pi}{4}\frac{f_{\lambda'} - f_{\lambda''}}{\omega_\lambda \omega_{\lambda'} \omega_{\lambda''}} |V^{+}_{\lambda,\lambda',\lambda''}|^2 \, \delta(\omega_\lambda + \omega_{\lambda'} - \omega_{\lambda''}) \tag{6.17}$$

$$\Gamma^{-}_{\lambda,\lambda',\lambda''} = \frac{\hbar\pi}{4}\frac{f_{\lambda'} + f_{\lambda''} + 1}{\omega_\lambda \omega_{\lambda'} \omega_{\lambda''}} |V^{-}_{\lambda,\lambda',\lambda''}|^2 \, \delta(\omega_\lambda - \omega_{\lambda'} - \omega_{\lambda''}), \tag{6.18}$$

where the Dirac deltas serve to enforce conservation of energy.

The terms $V^{\pm}_{\lambda,\lambda',\lambda''}$ are three-phonon matrix elements. We use the notation introduced earlier: $\boldsymbol{R}$ is the position of an atom in the unit cell, $\boldsymbol{R}'$ and $\boldsymbol{R}''$ are atoms positions in supercell. Here, $\boldsymbol{u}$ is an eigenfunction, with branch index and wavenumber λ, and α, β, γ spatial indices. The matrix elements are computed as follows from the third order elastic constants:

$$V^{\pm}_{\lambda,\lambda',\lambda''} = \sum_{\substack{\boldsymbol{R},\boldsymbol{R}',\boldsymbol{R}'' \\ \alpha,\beta,\gamma}} \mathsf{E}^{\boldsymbol{R},\boldsymbol{R}_1,\boldsymbol{R}_2}_{\alpha,\beta,\gamma} \frac{u^{\lambda}_{\alpha}(\boldsymbol{R}) u^{\lambda'}_{\beta}(\boldsymbol{R}') u^{\lambda''}_{\gamma}(\boldsymbol{R}'')}{\sqrt{M_{\boldsymbol{R}} M_{\boldsymbol{R}'} M_{\boldsymbol{R}''}}}. \tag{6.19}$$

The second term of eq. (6.15), $\boldsymbol{\Delta}$, is defined as

$$\boldsymbol{\Delta}_\lambda = \frac{1}{N}\sum_{\lambda',\lambda''} \Gamma^{+}_{\lambda,\lambda',\lambda''} \left(\frac{\omega_{\lambda''}}{\omega_\lambda} \boldsymbol{F}_{\lambda''} - \frac{\omega_{\lambda'}}{\omega_\lambda} \boldsymbol{F}_{\lambda'} \right) + \frac{1}{2N}\sum_{\lambda',\lambda''} \Gamma^{-}_{\lambda,\lambda',\lambda''} \left(\frac{\omega_{\lambda''}}{\omega_\lambda} \boldsymbol{F}_{\lambda''} + \frac{\omega_{\lambda'}}{\omega_\lambda} \boldsymbol{F}_{\lambda'} \right). \tag{6.20}$$

Additional terms can be included to represent other sources of scattering, such as impurities, or from an isotope distribution.

This equation can be solved iteratively for $\boldsymbol{F}$ until convergence to a prescribed tolerance. It can be shown that the lattice thermal conductivity is then given by

$$\boldsymbol{\kappa} = \frac{1}{k_B T^2 V N} \sum_{\lambda} f^0_\lambda (f^0_\lambda + 1)(\hbar\omega_\lambda)^2 \boldsymbol{v}_\lambda \otimes \boldsymbol{F}_\lambda. \tag{6.21}$$

6.5. Linear Response and the Green–Kubo Relation

A linear transport equation is of the form

$$J = LF \tag{6.22}$$

where J is the current density of some quantity of interest and F is the thermodynamical force conjugate to it. L is the linear transport coefficient.

The transport coefficient can be calculated from the autocovariance of the corresponding current (Kubo, 1957):

$$L = \beta V \int_0^\infty \mathrm{d}s \, \langle J(0) J(t) \rangle_{F=0}. \tag{6.23}$$

For the particular case of thermal conductivity, this becomes

$$\kappa = \frac{V}{3k_B T^2} \int_0^\infty \mathrm{d}t \, \langle \boldsymbol{q}(0) \cdot \boldsymbol{q}(t) \rangle. \tag{6.24}$$

By ergodicity of the system, we may replace the ensemble average appearing in equation (6.24) with a time average of a single realisation of the time evolution, provided it is long enough:

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T \mathrm{d}\tau \, \boldsymbol{q}(\tau) \cdot \boldsymbol{q}(\tau + t) = \langle \boldsymbol{q}(0) \cdot \boldsymbol{q}(t) \rangle. \tag{6.25}$$

We estimate these quantities by using a discrete sampling from an MD simulation of the instantaneous heat flux at equilibrium $\{\boldsymbol{q}_i\}$, where i labels the time step, and replacing the integrals with finite sums

$$R_n = \frac{1}{3M} \sum_{i=0}^{M} \boldsymbol{q}_i \cdot \boldsymbol{q}_{i+n}, \tag{6.26}$$

for large M.

This requires a definition of the heat flux $\boldsymbol{q}$ in terms of the microscopic degrees of freedom of the system, which is missing from the discussion in a continuum setting above, where temperature is a distinct degree of freedom from local velocity (which we took to be zero in any case).

Define

$$\boldsymbol{q} = \frac{1}{V} \frac{\mathrm{d}}{\mathrm{d}t} \left(\sum_i^N \boldsymbol{r}_i E_i \right) \tag{6.27}$$

where E_i is the energy of the i^{th} particle, and

$$E_i = e_i + \frac{1}{2} m_i \boldsymbol{v}_i \cdot \boldsymbol{v}_i. \tag{6.28}$$

There is no physically distinguished definition of e_i: the only energy rigorously defined is that for the whole system. For a pair potential, we might assign half of the interaction energy of a pair of particles to each of them. For many body potentials there are several 'reasonable' definitions.

We now show that we can often express eq. (6.27) in terms of the local energy and a local stress, provided that we choose a definition of local stress to be consistent with the local energy.

To make use of eq. (6.27), apply the time derivative to obtain

$$\boldsymbol{q} = \frac{1}{V}\sum_i \left(\boldsymbol{v}_i E_i + m_i \boldsymbol{r}_i(\boldsymbol{v}_i \cdot \dot{\boldsymbol{v}}_i) + \boldsymbol{r}_i \dot{e}_i\right) \tag{6.29}$$

$$= \frac{1}{V}\sum_i \left(\boldsymbol{v}_i E_i + \boldsymbol{r}_i(\boldsymbol{v}_i \cdot \boldsymbol{f}_i) + \boldsymbol{r}_i \dot{e}_i\right) \tag{6.30}$$

$$= \frac{1}{V}\sum_i \left(\boldsymbol{v}_i E_i + \sum_j \boldsymbol{r}_i \left(-\boldsymbol{v}_i \cdot \frac{\partial e_j}{\partial \boldsymbol{r}_i} + \boldsymbol{v}_j \cdot \frac{\partial e_i}{\partial \boldsymbol{r}_j}\right)\right) \tag{6.31}$$

where we used

$$\dot{e}_i = \sum_j \boldsymbol{v}_j \cdot \frac{\partial e_i}{\partial \boldsymbol{r}_j} \tag{6.32}$$

and

$$\boldsymbol{f}_i = -\sum_j \frac{\partial e_j}{\partial \boldsymbol{r}_i}. \tag{6.33}$$

Exchanging indices:

$$\boldsymbol{q} = \frac{1}{V}\sum_i \boldsymbol{v}_i E_i - \frac{1}{V}\sum_{i,j} (\boldsymbol{r}_j - \boldsymbol{r}_i)\left(\boldsymbol{v}_i \cdot \frac{\partial e_j}{\partial \boldsymbol{r}_i}\right) \tag{6.34}$$

$$= \frac{1}{V}\sum_i \boldsymbol{v}_i E_i - \frac{1}{V}\sum_i \boldsymbol{v}_i \sum_j (\boldsymbol{r}_j - \boldsymbol{r}_i) \otimes \frac{\partial e_j}{\partial \boldsymbol{r}_i} \tag{6.35}$$

The last term is a sum of a tensor product of two vectors over each atom. We now compare this quantity to the virial stress.

Define the **virial stress** as

$$\boldsymbol{S} = -\frac{\partial e}{\partial \boldsymbol{F}}, \tag{6.36}$$

where $\boldsymbol{F}$ is the deformation gradient defined in chapter 3.

$$\boldsymbol{L}' = \boldsymbol{F}\boldsymbol{L} \tag{6.37}$$

In the following, the Latin indices refer to atoms and Greek to spatial coordinates.

$$S_{\alpha\beta} = -\frac{\partial e}{\partial F_{\alpha\beta}} \tag{6.38}$$

$$= -\sum_{i}^{N}\sum_{\gamma}^{3}\frac{\partial r_{i\gamma}}{\partial F_{\alpha\beta}}\frac{\partial e}{\partial r_{i\gamma}} \tag{6.39}$$

$$= -\sum_{i,j}^{N}\sum_{\gamma}^{3}\frac{\partial r_{i\gamma}}{\partial F_{\alpha\beta}}\frac{\partial e_j}{\partial r_{i\gamma}} \tag{6.40}$$

$$= -\sum_{i,j}^{N} r_{i\beta}\frac{\partial e_j}{\partial r_{i\alpha}} \tag{6.41}$$

$$\boldsymbol{S} = -\sum_{i,j}^{N} \boldsymbol{r}_i \otimes \frac{\partial e_j}{\partial \boldsymbol{r}_i}. \tag{6.42}$$

We might also define a per-atom virial:

$$\boldsymbol{S}_i = -\sum_{j}^{N} \boldsymbol{r}_i \otimes \frac{\partial e_j}{\partial \boldsymbol{r}_i}. \tag{6.43}$$

As for the local energy, this partitioning is quite arbitrary, the only requirement is that its sum over all of the atoms in the domain is $\boldsymbol{S}$. The way we have defined it here however, is not consistent with the form appearing in the expression for the heat flux below.

Note that the expression for virials appearing for pair potentials is sometimes written

$$\boldsymbol{S}_i = -\frac{1}{2}\sum_{j}^{N}(\boldsymbol{r}_i - \boldsymbol{r}_j) \otimes \boldsymbol{f}_{ij} \tag{6.44}$$

where $\boldsymbol{f}_{ij}$ is the force on atom i due to atom j, and the expression relies on antisymmetric forces:

$$\boldsymbol{f}_{ij} = -\boldsymbol{f}_{ji}. \tag{6.45}$$

In the general situation, we can use

$$\boldsymbol{f}_{ij} = \frac{\partial e_j}{\partial \boldsymbol{r}_i}, \tag{6.46}$$

but this does not necessarily satisfy eq. (6.45).

Returning to the fluxes, eq. (6.35) becomes

$$\boldsymbol{q} = \frac{1}{V}\left(\sum_i \boldsymbol{v}_i E_i + \sum_i \boldsymbol{v}_i \cdot \boldsymbol{S}_i\right), \tag{6.47}$$

if we define

$$\boldsymbol{S}_i = \sum_j^N (\boldsymbol{r}_j - \boldsymbol{r}_i) \otimes \frac{\partial e_j}{\partial \boldsymbol{r}_i}. \tag{6.48}$$

This quantity summed over all of the atoms in the domain gives the total virial, as is required of a good local virial, provided that the local energy is a function only of the vector displacements to the other atoms, $\boldsymbol{r}_{ij} = \boldsymbol{r}_j - \boldsymbol{r}_i$.

$$e_i = e_i(\{\boldsymbol{r}_{ij}\}_{j=1, j\neq i}^N), \tag{6.49}$$

which has one fewer degree of freedom than the general case.

To see this, consider the term of the sum,

$$\sum_{i,j}^N \boldsymbol{r}_j \otimes \frac{\partial e_j(\{\boldsymbol{r}_{jk}\})}{\partial \boldsymbol{r}_i} = \sum_j^N \boldsymbol{r}_j \otimes \sum_i^N \frac{\partial e_j(\{\boldsymbol{r}_{jk}\})}{\partial \boldsymbol{r}_i} \tag{6.50}$$

And (no summation unless explicit):

$$\sum_i \frac{\partial e_j(\{\boldsymbol{r}_{jk}\})}{\partial \boldsymbol{r}_i} = \sum_{\substack{ijk \\ k\neq j}} \frac{\partial e_j}{\partial \boldsymbol{r}_{jk}} \frac{\partial \boldsymbol{r}_{jk}}{\partial \boldsymbol{r}_i} \tag{6.51}$$

$$= \sum_{\substack{ijk \\ k\neq j}} \frac{\partial e_j}{\partial \boldsymbol{r}_{jk}} \frac{\partial}{\partial \boldsymbol{r}_i}(\boldsymbol{r}_j - \boldsymbol{r}_k) \tag{6.52}$$

$$= \sum_{\substack{ijk \\ k\neq j}} \frac{\partial e_j}{\partial \boldsymbol{r}_{jk}} (\delta_{ij} - \delta_{ki}) \tag{6.53}$$

$$= \sum_{\substack{ik \\ k\neq j}} \frac{\partial e_i}{\partial \boldsymbol{r}_{ik}} - \sum_{\substack{ij \\ i\neq j}} \frac{\partial e_j}{\partial \boldsymbol{r}_{ji}} = 0 \tag{6.54}$$

This is true of many potentials, but in particular it is not true for force-constant potentials (imagine all of the atoms being shifted by the same constant amount away

from their equilibrium positions: there would be no force on any atom, but the local energy of each atom would increase by the same amount.

Equation (6.47) still holds however, provided we interpret the term $\boldsymbol{S}_i$ correctly, no longer as local virials.

6.6. The Potentials

A number of Gaussian approximation potentials were trained on databases of varying composition, using the SOAP descriptor; see section 2.3. Table 6.3 gives a breakdown of the structures contained in each database. A training configuration is a snapshot of the system of atoms (with a pre-determined structure, or as a frame of a molecular dynamics simulation) along with a total energy and forces. Each configuration represents a combination of many atomic environments (for example, the two-atom unit diamond structure cells give two environments per configuration). The k-means clustering scheme is used to choose certain 'distinguished points' to reduce the total of these to a managable number, but constrained to have a minimum number of environments represented for each configuration type.

The forces and energies of the configurations were obtained ab initio with the CASTEP DFT code (Clark et al., 2005), using the PW91 functional (Perdew and Wang, 1992) and a highly converged plane-wave basis (with an energy cutoff of 250 eV) and k-grid (with a cutoff of at least 20 Å). All but the single unit cell structures were sampled using molecular dynamics, by subsampling a trajectory computed using coarser DFT parameters. The SOAP parameters and hyperparameters are given in table 6.4.

The 'target' DFT value of the thermal conductivity was known when training the potentials. To avoid cherry-picking training databases with particularly favourable values, we kept the potentials searched systematically improving via their size and the use of additional structures.

Table 6.5 shows the performance of the potentials compared with DFT, including the thermal conductivity from the Boltzmann transport equation at 300 K. This was calculated using the `ShengBTE` code (Li et al., 2014), on a converged grid of 30^3 k-points. Figure 6.2 shows the temperature dependent thermal conductivity for the two most converged potentials.

Clearly the phonon spectrum, Grüneisen parameters and thermal conductivity are related to the accuracy to which the potential can reproduce the second- and third-order force constants, on which they all depend. There is improvement in most of these quantities as the potential is systematically enlarged. The first potential has a very small training

Table 6.3.: The training databases and validation results. Shown are the number of training configurations used, broken down by type and size (upper part of the table). Each of these represents multiple atomic environments: for example, the two-atom unit diamond structure cells give two environments per configuration. Each of the potentials apart from Pot. 1 used the *k*-means clustering scheme for reducing the rank of the covariance matrix, with distinguished points chosen from the training data, using a minimum of 1000 from each configuration type (counting all of the cubic diamond entries together). The total number of environments used by this scheme are given in the last row of the table.

Training configuration	atoms	Pot. 1	Pot. 2	Pot. 3	Pot. 4	Pot. 5
cubic diamond	2	104	104	104	104	104
	16	-	-	220	220	220
	54	-	110	110	110	110
	128	-	-	55	55	55
liquid	64	-	-	-	63	63
	128	-	-	-	7	7
amorphous	64	-	-	-	31	31
	216	-	-	-	16	16
beta-tin	2	-	-	-	-	60
	16	-	-	-	-	220
	54	-	-	-	-	110
	128	-	-	-	-	55
simple-hexagonal	1	-	-	-	-	110
	8	-	-	-	-	30
	27	-	-	-	-	30
	64	-	-	-	-	53
number of represented environments		208	2000	2000	5000	5000

ν_{energy}	(eV/atom)	0.003
ν_{force}	(eV/Å)	0.2
ν_{virial}	(eV/atom)	0.2
$1/\alpha$	(Å)	0.5
r_{cut}	(Å)	5.0
l_{max}		12
n_{max}		10
ζ		4

Table 6.4.: Hyperparameters used for the SOAP kernel. The ν are the noise hyperparameters, $1/\alpha$ is the width of the Gaussian used to represent the atomic density in eq. (2.54), and r_{cut} the cutoff radius. l_{max} and n_{max} are the maximum values to use for the sum in eq. (2.52); ζ appears in eq. (2.53).

database, and although it produces reasonable lattice parameter and elastic constants, it cannot reproduce the other quantities accurately. The second potential has better second- and third-order force constants overall, resulting in smaller errors in phonon spectrum and Grüneisen coefficient, although this does not translate into a better thermal conductivity, which is still dramatically underestimated. Introducing liquid structures into the training database more than doubles the computed thermal conductivity, which is improved further still by the addition of other structures (amorphous, beta-tin and simple-hexagonal).

This addresses an important aim of this part of the project: to construct a potential model capable of reproducing the thermal conductivity to first-principles accuracy. Even though the conductivity obtained as a solution to the Boltzmann transport equation is approximate, and is only applicable up to some fraction of the Debye temperature, there is good agreement between DFT and experiment over a considerable temperature range, deviating in the expected way from experiment for temperatures of 600 K and higher, and it is possible to obtain convergence in PBTE conductivity, along with other harmonic and anharmonic properties, using this method of constructing a potential.

The restriction to low temperatures does not apply to the Green–Kubo approach, which we now consider. At low temperatures these methods should agree when fully converged in system size and time. In practice, it was found that it is often difficult to observe an agreement between them, and they often differ markedly. Size scaling effects were considered by Sellan et al. (2010), Esfarjani et al. (2011) and He et al. (2012), who note that by using small systems, as is necessary with the Green–Kubo method, this

Table 6.5.: Validation results for the potentials

Property	(unit)	Pot. 1	Pot. 2	Pot. 3	Pot. 4	Pot. 5	DFT
Relaxed lattice parameter	(Å)	5.459	5.461	5.461	5.461	5.461	5.461
Energy/atom at lattice minimum	(eV)	−163.178	−163.178	−163.178	−163.178	−163.178	−163.178
Elastic constants: c_{11}	(GPa)	158	157	152	152	150	154
c_{12}	(GPa)	61	59	58	56	56	56
c_{44}	(GPa)	78	77	75	75	73	75
Phonon spectrum L_2 error	(THz)	9.854	2.217	2.207	2.062	1.346	-
Phonon spectrum L_∞ error	(THz)	27.553	4.672	4.341	4.606	2.887	-
Grüneisen gamma L_2 error		1.199	0.561	0.385	0.314	0.318	-
Grüneisen gamma L_∞ error		3.421	2.335	1.607	1.424	1.281	-
3rd-order PBTE therm. cond. @300 K	($\mathrm{W\,m^{-1}\,K^{-1}}$)	47	42	112	123	141	139

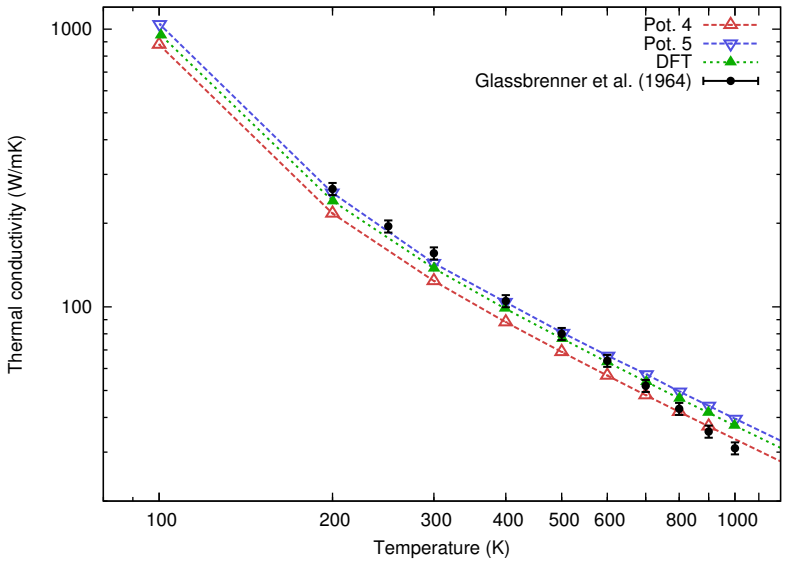

Figure 6.2.: Thermal conductivity from the iterative solution to the phonon-Boltzmann transport equation, for the two best GAP potentials, the result directly from the DFT on which they were trained, and the experiment of Glassbrenner and Slack (1964).

removes long-wavelength phonons with long mean free paths, and this could significantly reduce the conductivity.

We now give a particularly well-converged result in time for Stillinger–Weber silicon, with isotopically pure ^{28}Si. A temperature of 600 K was chosen as a compromise between limiting the thermal conductivity for relatively rapid convergence of the autocorrelation, and allowing comparison with the PBTE. It is found that despite obtaining a larger value with a smaller error bar than other reported results, it does not agree with the third-order phonon-transport result.

We obtain 25.6 μs of total simulation data, comprising 512 independent 50 ns trajectories, for both system sizes. This is considerably more than other simulations described in the literature, and allows a converged value to be obtained for the integrated heat-flux autocorrelation up to several nanoseconds, and allows for a more reliable fit to the autocorrelation.

Figure 6.3 shows the heat-flux autocorrelation for the 10^3 unit cell system. Even though it rapidly becomes quite small, the long-time tail is important for convergence of the integral, shown in fig. 6.4.

It is not possible to use systems much larger than 10^3 for this length of simulation. This small system discretizes the phonons that can be represented in the simulation domain,

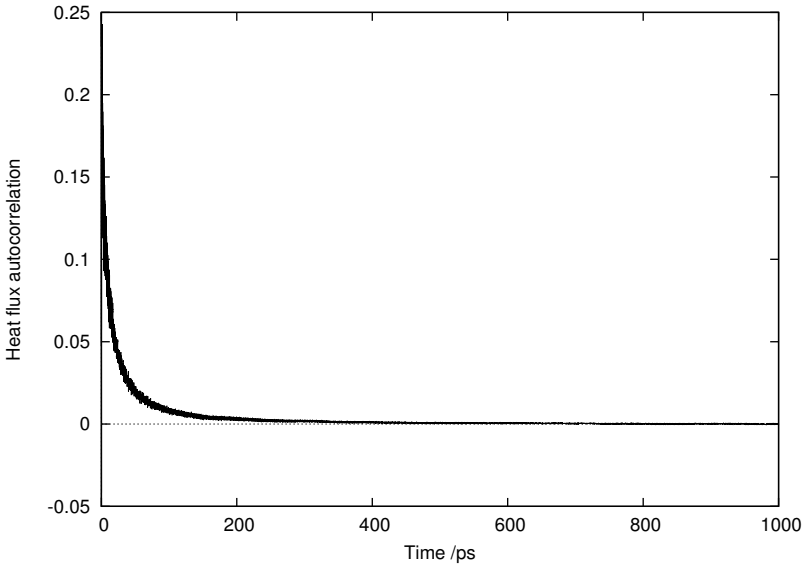

Figure 6.3.: The heat-flux autocorrelation for the Stillinger–Weber potential, for the 10^3 unit cell system at 600 K.

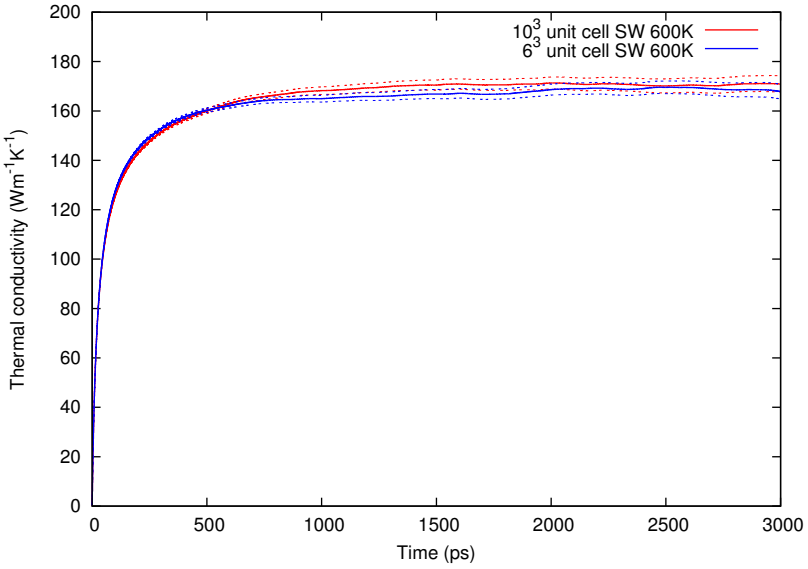

Figure 6.4.: Integrated heat-flux autocorrelation for the Stillinger–Weber potential, up to 3 ns, for two system sizes at 600 K. The dashed lines are standard error computed from the independent trajectories.

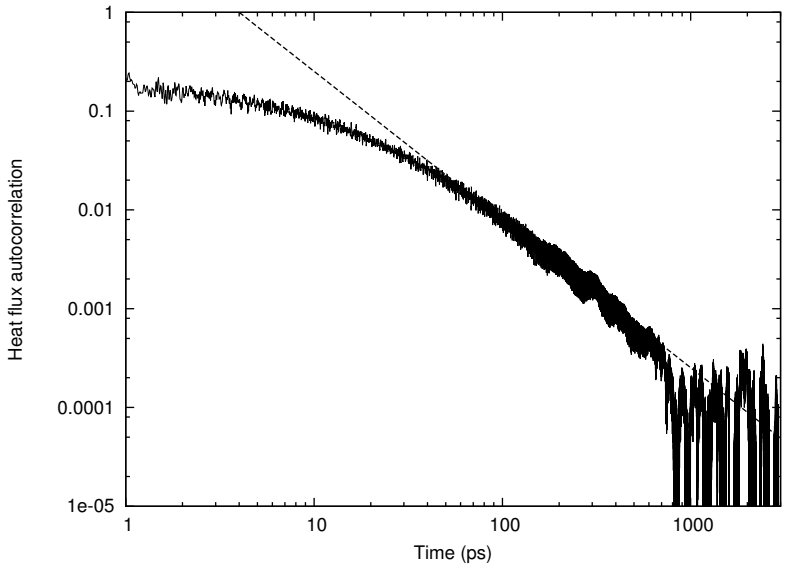

Figure 6.5.: The heat-flux autocorrelation for the Stillinger–Weber potential, for the 10^3 unit cell system at 600 K. The dashed line is a power law fit to the tail, with exponent -1.5.

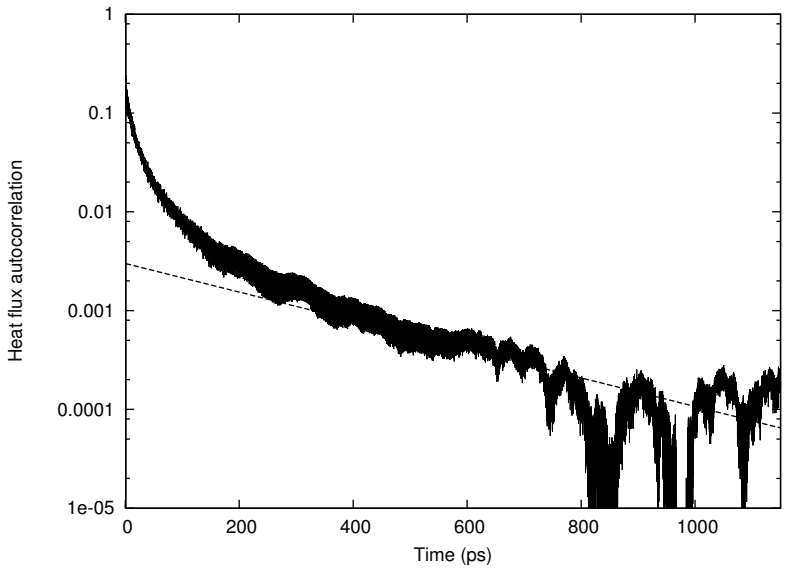

Figure 6.6.: The heat-flux autocorrelation for the Stillinger–Weber potential, for the 10^3 unit cell system at 600 K. The dashed line is an exponential fit to the tail, with time constant 300 ps.

and importantly, eliminates long wavelength phonons entirely. There is some reason to believe that these contribute a lot to the heat transport, as they typically have low scattering rates (within the relaxation time approximation), so to converge with system size could require very large simulations indeed. We do not find a significant difference in the final value between the 6^3 unit cell and 10^3 unit cell system.

Integrating the heat-flux autocorrelation between 0 and 1 ns (and discarding the integral of the tail after 1 ns) gives an estimate for the thermal conductivity (via eq. (6.24)) as $166 \pm 4\,\mathrm{W\,m^{-1}\,K^{-1}}$. After this point, the error begins to increase more rapidly than the value, although this does not imply convergence to within the error bar.

Figures 6.5 and 6.6 show the autocorrelation fit to a power law and to an exponential. The power law has exponent -1.5 and the exponential has a time constant of 300 ps. Adding the above estimate of the integral to 1 ns to the integral (to ∞) of these fitting functions gives a thermal conductivity of $185\,\mathrm{W\,m^{-1}\,K^{-1}}$ for the power law of, and $175\,\mathrm{W\,m^{-1}\,K^{-1}}$ for the exponential. The error in the estimate in this case is determined by the quality of this fit at long times, and is very sensitive for the power law (since the exponent is close to -1). In either case, the additional contribution to the conductivity from integrating the tail is significant.

We note that this value is somewhat larger than other literature values for Stillinger–Weber silicon computed with the same method ($90 \pm 20\,\mathrm{W\,m^{-1}\,K^{-1}}$ Volz and Chen (2000); $120\,\mathrm{W\,m^{-1}\,K^{-1}}$ Schelling et al. (2002, extrapolated)). This is likely to be in part because of the longer integration time possible with our data. For comparison, the experimental value at this temperature is $64\,\mathrm{W\,m^{-1}\,K^{-1}}$, although we do not expect Stillinger–Weber to reproduce this value particularly well.

Using third-order phonon-Boltzmann transport from force constants generated with the same system gives $300\,\mathrm{W\,m^{-1}\,K^{-1}}$, which is in agreement with the value of Broido et al. (2005). The PBTE solution can include contributions from very long wavelength phonons (dependent on the $\boldsymbol{k}$-grid), but neglects higher-order phonon scattering processes. Both of these effects could inflate the conductivity figure over the Green–Kubo result, only the first desirably. A similar effect was demonstrated for the Tersoff potential by He et al. (2012), who found a conductivity from a relaxation time approximation to the PBTE to be 26 % higher than the Green–Kubo method.

6.7. Summary

We have presented a Gaussian approximation potential model for silicon capable of reproducing the low and intermediate temperature thermal conductivity of the underlying

density functional theory on which it was fit, as determined by the full iterative solution to the phonon-Boltzmann transport equation. This in turn agrees well with experiment. It is found that including a wide enough range of structures in the training database is essential in producing an accurate value, even if these are apparently quite different from the structure under consideration.

The issue of using molecular dynamics simulations in equilibrium to obtain a value of thermal conductivity was considered, although it is found to differ from the PBTE solution. A value particularly well-converged in time was given for the Stillinger-Weber potential.

7. Conclusions and Further Work

Below, we summarize the main conclusions of the thesis, and describe several directions for further work.

Hyperelastic equation of state modelling from first principles

We have developed a framework for constructing a temperature-dependent, anisotropic equation of state for hyperelasticity from first-principles calculations, accurately and robustly. Although the computational requirements are somewhat greater than commonly used equations of state for nonlinear elasticity, the approach is general and could be applied to many materials straightforwardly. The method was showcased using some one-dimensional problems (in the full three-dimensional elastic system), including a shock interacting with a thermal gradient, and a change in crystal orientation in the material. A further development of the application of this method would be the extension to two or three dimensions. A particular problem of interest where this could be applied is in the simulation of a polycrystal under deformation.

Shock waves from first principles

We have described an annealing method well-suited to first-principles simulations to obtain the Hugoniot locus of a material with molecular dynamics, and demonstrated this for DFT silicon. We were able to construct the Hugoniot of a number of post-shock phases, and these compare well to experiment up to the available range of experimental data of 60 GPa. States are given for elastic shocks, an approximation to plastically-yielded states, the simple-hexagonal phase, the beta-tin phase and a liquid. These calculations support the experimental observations that silicon loses all strength after yielding, and that the first phase transition along the Hugoniot locus is likely to be a simple hexagonal phase.

In addition, we perform several direct shock simulations using empirical interatomic potentials, obtaining a number of points in regions of the Hugoniot not accessible to the

anneal. We have performed a direct simulation of an elastic shock wave, with forces from DFT, to just below the Hugoniot elastic limit, and observe a very short rise time for this shock.

There is considerable interest in observing a much stronger shock wave in DFT silicon directly. The high pressure beta-tin and liquid phases of silicon are metallic, with a locally six-fold coordination. In addition to the structural phase transformation, an electronic structure method would allow the insulator to metal transition to be observed directly, via the local band-structure. This system requires a k-point sampling beyond just the Γ-point used for the elastic shock, which is crucial to simulate a metal well. This requirement already increases the computational requirements, further complicated by the resulting dense structure after the passage of the shock, which is more challenging to converge in SCF than the initial diamond structure. Since the simulation domain must remain small, a shock-wave above the split-wave limit is required. Although some effort was made in obtaining convergence (at all, and then quickly) and running this simulation in a reasonable time, it remains ongoing work to complete it. Along with the elastic shock, these would be the first ab initio shock simulations we are aware of.

Thermal conductivity

We have developed a systematically-improved empirical interatomic potential model for silicon, using the approach of Gaussian approximation potentials. This applies a Bayesian inference technique commonly used in the machine learning community to the development of accurate and transferable potentials from an underlying structure database of energies and forces, which can be computed from a first-principles method.

We showed that by using this approach it is possible to obtain a potential capable of reproducing the temperature-dependent thermal conductivity, as computed from the phonon Boltzmann transport equation, to first-principles accuracy, but with a potential having much lower computational cost. The advantage of this approach is that while the thermal conductivity is important in many contexts, methods for computing the *high temperature* conductivity, or the thermal resistance of atomic-scale structures, are too demanding for first-principles methods such as density functional theory.

Even though only diamond structure silicon was considered, it was found that including other structures in the training database, seemingly quite different from those encountered when computing the force constants, was necessary to obtain a converged result. This does not sacrifice other material properties computed from the potential, which can also be seen to converge. This is already an improvement on potentials based on simple functional forms, where once certain properties are fixed, there remains limited scope for

improving the others. This is a benefit arising the underlying Gaussian process, whose shape is quite arbitrary compared with parameterized forms.

Reconciling the low-temperature phonon Boltzmann approach with the Green–Kubo method, which can be applied at any temperature, seems to be a challenging problem, and exposes many subtleties. These are discussed in a number of places in the literature (e.g. Esfarjani et al., 2011; He et al., 2012). An attempt at comparing the Green–Kubo conductivity as obtained by the GAP potentials and DFT, by fitting to a force-constant potential up to fourth order, based on the work of Esfarjani et al. (2011), is currently underway.

A. The Romenski Equation of State

The specific internal energy is written in terms of the specific entropy and the invariants of the inverse right Cauchy–Green tensor, as equations (A.1–3.26), but in terms of $\boldsymbol{C} = \boldsymbol{B}^{-1}$:

$$\mathrm{I}_C = \operatorname{tr} \boldsymbol{C} \tag{A.1}$$

$$\mathrm{II}_C = \frac{1}{2}\left((\operatorname{tr} \boldsymbol{C})^2 - \operatorname{tr}(\boldsymbol{C}^2)\right) \tag{A.2}$$

$$\mathrm{III}_C = \det \boldsymbol{C}. \tag{A.3}$$

It is decomposed into two terms,

$$E(\mathrm{I}_C, \mathrm{II}_C, \mathrm{III}_C, S) = U(\mathrm{III}_C, S) + W(\mathrm{I}_C, \mathrm{II}_C, \mathrm{III}_C), \tag{A.4}$$

where U is the thermal and hydrostatic energy density and W is the elastic energy due to shear deformation. These are given as

$$U(\mathrm{III}_C, S) = \frac{K_0}{2\alpha^2}\left(\mathrm{III}_C^{\alpha/2} - 1\right)^2 + c_v T\, \mathrm{III}_C^{\gamma/2} \tag{A.5}$$

$$W(\mathrm{I}_C, \mathrm{II}_C, \mathrm{III}_C) = \frac{B_0}{2}\, \mathrm{III}_C^{\beta/2}(\mathrm{I}_C^2/3 - \mathrm{II}_C), \tag{A.6}$$

with T defined in terms of S assuming a constant heat capacity,

$$T = T_0(\exp(S/c_v) - 1). \tag{A.7}$$

Here, c_v is the specific heat capacity at constant volume, K_0 is the bulk modulus and B_0 is the elastic P-wave modulus.

An alternative parameterization in terms of b_0 and c_0, the bulk and longitudinal speeds of sound, has

$$B_0 = b_0^2 \tag{A.8}$$

$$K_0 = c_0^2 - \frac{4}{3} b_0^2 \tag{A.9}$$

Table A.1.: A parameterization of the Romenski equation of state for copper.

ρ_0 (gcm^{-3})	b_0 (kms^{-1})	c_0 (kms^{-1})	c_v ($\mathrm{kJg}^{-1}\mathrm{K}^{-1}$)	T_0 (K)	α	β	γ
8.93	2.1	4.6	3.9×10^{-4}	300	1.0	2.0	3.0

B. Gaussian Process Regression

B.1. Introduction

Here, we give an overview of Gaussian process regression, providing the details of the methods used in section 2.3 and chapter 4, and serving to define some notation. For a broader overview of Gaussian process regression, the reader is referred to MacKay (2003, ch. 45). Approaches to reduced-rank regression (appendix B.5) are covered by Quiñonero Candela and Rasmussen (2005) and Snelson and Ghahramani (2007).

A **Gaussian process** is a stochastic process $\{t_n | n \in I\}$, n ranging over an index set I, where each finite collection of N variables $\boldsymbol{t} = (t^{(n_1)}, t^{(n_2)}, \ldots, t^{(n_N)})$ has a joint Gaussian distribution

$$P(\boldsymbol{t}|\mu, C) = \frac{1}{Z} \exp\left(-\frac{1}{2}(\boldsymbol{t}-\mu)^\mathsf{T} C^{-1}(\boldsymbol{t}-\mu)\right) \tag{B.1}$$

The parameters μ and C are the (vector) mean and $N \times N$ covariance matrix of the distribution, and Z is a normalising factor[1], included so that the probability density P integrates to one over the domain of $\boldsymbol{t}$ (Hoel et al., 1987, sec. 4.2).

Gaussian process regression is the use of a Gaussian process prior in a Bayesian function-inference problem. Consider that when performing an interpolation of some known data points, $\left\{(\boldsymbol{x}^{(i)}, t^{(i)})\right\}$, to a new pair of values, $(\boldsymbol{x}', t')$, where $\boldsymbol{x}'$ may be different from any of the other $\boldsymbol{x}^{(i)}$, and t' is unknown, we are interested in the conditional probability distribution

$$P(t'|x', \{(\boldsymbol{x}^{(i)}, t^{(i)})\}). \tag{B.2}$$

We assume that the *joint* probability distribution of any collection of outputs is given by eq. (B.1), and therefore that $\boldsymbol{t}$ is a Gaussian process. We must then find the above *conditional* probability distribution.

Assuming that an observation t is generated from some underlying function of the input variables $\boldsymbol{x}$, for another input location $\boldsymbol{x}'$, an estimate of the value of the function

[1] Given by $Z = 1/\sqrt{(2\pi)^N \det C}$.

there, t', is given by the mean of the distribution. The variance gives a measure of the uncertainty in this estimate. The use of Gaussians makes finding explicit expressions for the mean and variance of the above distribution straightforward, and these values also determine the (Gaussian) distribution completely.

Whether we consider this reconstruction an interpolation or a regression depends on whether we consider the observations $t^{(i)}$ exact values of an underlying function (and therefore points that the interpolant are constrained to pass through), or whether this function is subject to statistical noise in obtaining the output values (where there is no such constraint).

We will sometimes write $\boldsymbol{X}$ for the collection of all of the inputs $(\boldsymbol{x}^{(i)})_{i=1}^{N}$, and $\boldsymbol{t}$ for the corresponding outputs $(t^{(i)})_{i=1}^{N}$. An individual input $\boldsymbol{x}^{(i)}$ may be a vector. It is also possible for the outputs t_i to be vector-valued, but here we restrict them to scalars.

B.2. The covariance function

The covariance matrix determines how the outputs t depend on one another through eq. (B.1), and is required to be positive definite. Since each output depends on a vector of inputs, it is reasonable to assume that the elements depend on the input vectors through a **covariance kernel** as $C_{ij} = C(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(j)})$. A stronger assumption is that the covariance kernel is **stationary**: the correlation between two observations depends only on the separation of their inputs, $\boldsymbol{x}^{(ij)} = \boldsymbol{x}^{(i)} - \boldsymbol{x}^{(j)}$.

A choice appropriate for many circumstances, also used for aspects of this work is the square-exponential covariance function

$$C(\boldsymbol{x}, \boldsymbol{y}) = \sigma^2 \exp\left(-\frac{1}{2}\sum_i \frac{(x_i - y_i)^2}{r_i^2}\right) + \nu^2 \delta_{\boldsymbol{x},\boldsymbol{y}}. \tag{B.3}$$

Some others are described by Gibbs and Mackay (1997). The SOAP kernel, eq. (2.52), has been discussed already. This form of kernel has a zero-mean prior on the output: far from any input data points, the kernel will revert to predicting a value of zero for the output[2].

The interpretation of the hyperparameters in the squared exponential covariance function is as follows: σ sets the overall scale of the inferred function (σ^2 is the prior variance of the inferred function value at a point, before any noise is considered), ν represents position-independent Gaussian noise in the outcomes that is independent of the inputs, and r_i is the characteristic length over which the function values become

[2] Consider two uncorrelated observations, with a large separation. Their covariance is just the square of the mean, which for this function is zero

decorrelated as the distance between two inputs along dimension i varies, which may differ between input dimensions. A large value of r_i (compared with the scale of the inputs) indicates that the corresponding input dimension has little influence over the function value, while a small value of r_i indicates that the function is highly sensitive to changes in the input along that dimension. The noise parameter ν has the added role numerically of improving the condition number of the covariance matrix when all of the inputs are highly correlated. For this reason, a small, positive value can be used for ν when computing the covariance matrix, even when it might otherwise be zero.

In both of the applications we consider, observations of function derivatives are available. Making use of derivative information is a matter of specifying the covariance between pairs of derivatives, and between derivative and value measurements. For the squared-exponential kernel, we have the following covariances: between a derivative (with respect to the ith component) and value,

$$-\frac{\partial C}{\partial x_i}(\boldsymbol{x},\boldsymbol{y}) = \frac{\partial C}{\partial y_i}(\boldsymbol{x},\boldsymbol{y}) = \frac{x_i - y_i}{r_i^2}\sigma^2 \exp\left(-\frac{1}{2}\sum_k \frac{(x_k - y_k)^2}{r_k^2}\right) \tag{B.4}$$

and between two derivative outcomes (with respect to x_i and x_j),

$$\frac{\partial^2 C}{\partial x_i \partial y_j}(\boldsymbol{x},\boldsymbol{y}) = \left(\frac{1}{r_i r_j} - \frac{1}{2}\frac{x_i - y_i}{r_i^2}\frac{x_j - y_j}{r_j^2}\right)\sigma^2 \exp\left(-\frac{1}{2}\sum_k \frac{(x_k - y_k)^2}{r_k^2}\right). \tag{B.5}$$

As the notation indicates, these are just the corresponding partial derivatives of the covariance function. Ignoring the noise term for the time being, this follows by the chain rule: an output $t^{(i)}$ depends only on $\boldsymbol{x}^{(i)}$, so

$$C_{ij} = \mathrm{Cov}(t(\boldsymbol{x}^{(i)}), t(\boldsymbol{x}^{(j)})), \tag{B.6}$$

where $t(\boldsymbol{x})$ is the random variable representing the value of the process at $\boldsymbol{x}$. Differentiating gives

$$\frac{\partial C_{ij}}{\partial \boldsymbol{x}^{(i)}} = \mathrm{Cov}\left(\frac{\partial t}{\partial \boldsymbol{x}}(\boldsymbol{x}^{(i)}), t(\boldsymbol{x}^{(j)})\right) \tag{B.7}$$

$$\frac{\partial^2 C_{ij}}{\partial \boldsymbol{x}^{(i)} \partial \boldsymbol{x}^{(j)}} = \mathrm{Cov}\left(\frac{\partial t}{\partial \boldsymbol{x}}(\boldsymbol{x}^{(i)}), \frac{\partial t}{\partial \boldsymbol{x}}(\boldsymbol{x}^{(j)})\right), \tag{B.8}$$

which are the required derivative covariances.

The noise has been excluded from the above expressions. This is appropriate since the derivative of t refers to the function underlying the value observations, and this is assumed to be differentiable, which it is no longer if it contains inherent Gaussian noise. The derivatives observations may be available through other means to the values, with a different noise level.

The covariance matrix is a prior on the space of functions considered by the regression model. It typically depends on a number of parameters. For a square-exponential kernel, these are r_i, σ, and the noise ν. These are known as **hyperparameters** (as distinct from plain parameters), and affect only the prior distribution. The hyperparameters determine the prior distribution on the space of fitting functions. As the size of the training set increases, the prior has less influence and the posterior distribution converges to a distribution independent of the prior, although perhaps slowly. This contrasts with parameters in the usual sense, which are varied to best represent the data[3]. Gaussian process regression is sometimes referred to as a nonparametric method, since it does not have any variable parameters of this kind.

B.3. Inference

Given an $N \times N$ covariance matrix $\boldsymbol{C}$, computed from the N ordinates $\boldsymbol{x}$, we now consider the task of inferring the output value t' for an input $\boldsymbol{x}'$. Starting from the known joint Gaussian distribution, and using the property of the Gaussian process, we show that the desired conditional distribution of the output value is Gaussian, with a mean and variance that can be readily computed.

Bayes theorem applied to eq. (B.2) is

$$P(t'|\boldsymbol{t}) = \frac{P(\boldsymbol{t}, t')}{P(\boldsymbol{t})}. \tag{B.9}$$

The numerator of this expression is another joint Gaussian with a $(N+1) \times (N+1)$ covariance matrix, which may be written

$$\begin{pmatrix} \boldsymbol{C} & \boldsymbol{k} \\ \boldsymbol{k}^\mathsf{T} & \kappa \end{pmatrix}, \tag{B.10}$$

where $\boldsymbol{k}$ has ith component $k_i = C(\boldsymbol{x}^{(i)}, \boldsymbol{x}')$, and $\kappa = C(\boldsymbol{x}', \boldsymbol{x}')$.

[3] Hyperparameters may enter into the prior as parameters on the model itself, even when not explicitly working in a Bayesian setting as we are here. For example, in a parametric linear least-squares polynomial fit, the degree of the polynomial could be considered a hyperparameter.

The inverse of this block matrix (see e.g. Barnett, 1979) is

$$\begin{pmatrix} C^{-1} + mC^{-1}\boldsymbol{k}\boldsymbol{k}^{\mathsf{T}}C^{-\mathsf{T}} & -mC^{-1}\boldsymbol{k} \\ -m\boldsymbol{k}^{\mathsf{T}}C^{-\mathsf{T}} & m \end{pmatrix}, \tag{B.11}$$

where $m = (\kappa - \boldsymbol{k}^{\mathsf{T}}C^{-1}\boldsymbol{k})^{-1}$. This is the covariance matrix for the joint Gaussian distribution of all of the observations. All of the observation values apart from t' are known, and the probability distribution of t' given the remaining observations $\boldsymbol{t}$ is Gaussian with mean

$$\hat{t}' = \boldsymbol{k}^{\mathsf{T}}C^{-1}\boldsymbol{t} \tag{B.12}$$

and variance

$$\mathrm{Var}(t') = \kappa - \boldsymbol{k}^{\mathsf{T}}C^{-1}\boldsymbol{k}. \tag{B.13}$$

B.4. Choosing the hyperparameters

Denote the hyperparameters in a particular covariance model by θ. Rather than taking them as given, fixed values, treat them as additional quantities to be inferred from the data. We are now interested in the distribution $P(\theta|\boldsymbol{x}, \boldsymbol{t})$, and the calculation of eq. (B.2) should be replaced by the integral

$$\int_\theta P(t'|\boldsymbol{x}', X, \boldsymbol{t}, \theta)\, P(\theta|X, \boldsymbol{t})\, \mathrm{d}\theta. \tag{B.14}$$

In practice, and especially for large systems, the cost of performing such an integral is prohibitive. It is also in many cases unnecessary, since its value is dominated by the most likely values of θ.

By Bayes' theorem, express

$$P(\theta|\boldsymbol{t}) = P(\boldsymbol{t}|\theta)P(\theta) \tag{B.15}$$

and find the θ that maximises the likelihood $P(\boldsymbol{t}|\theta)$.

The logarithm of the likelihood is

$$\log P(\boldsymbol{t}|\theta) = \frac{1}{2}\boldsymbol{t}^{\mathsf{T}}C^{-1}\boldsymbol{t} - \frac{1}{2}\log\det C - \frac{N}{2}\log 2\pi \tag{B.16}$$

and its derivative is

$$\frac{\partial \log P}{\partial \theta_i} = \frac{1}{2} \boldsymbol{t}^\mathsf{T} \boldsymbol{C}^{-1} \frac{\partial \boldsymbol{C}}{\partial \theta_i} \boldsymbol{C}^{-1} \boldsymbol{t} - \frac{1}{2} \operatorname{tr} \boldsymbol{C}^{-1} \frac{\partial \boldsymbol{C}}{\partial \theta_i}. \tag{B.17}$$

The logarithm of the likelihood is used numerically due to the rapid decay of the likelihood away from its maximum. Numerical computation of log det can be achieved directly and stably.

Knowledge of the derivative allows local maxima to be found using gradient-based optimisation algorithms such as the method of conjugate gradients. Alternatively, to avoid implementing an analytic gradient of each covariance function, a derivative-free optimisation method[4] can be used, at the cost of slower convergence (but only affecting the training time).

For part of this work, we used the `NLopt` optimisation library (Johnson, 2014), which implements several gradient-based and derivative-free methods.

B.5. Reduced-rank approximation

A Gaussian process with n training points has a prediction cost that is $O(n)$, the cost of a dot product of two vectors of length n. Variance predictions are $O(n^2)$, the cost of a matrix-vector product, and training is $\Omega(n^2)$ and in practice not much better than $O(n^3)$, the cost of inverting the covariance matrix.

A reduced-rank Gaussian process approximates the full rank-n covariance matrix with a matrix of lower rank, m. The motivation for this is that it reduces the value-prediction complexity to $O(m)$ and the variance-prediction complexity to $O(m^2)$. The training complexity involves $O(nm^2)$ operations for the inversion, but must also take into account the procedure used to approximate the full matrix with one of lower rank.

One particular reduced-rank approximation involves choosing m distinguished input locations, and approximating the full matrix as

$$\boldsymbol{C} \approx \boldsymbol{C}_* = \boldsymbol{K}_{nm} \boldsymbol{K}_{mm}^{-1} \boldsymbol{K}_{mn}. \tag{B.18}$$

Here, $\boldsymbol{K}_{mn}$ is the covariance matrix involving the m distinguished inputs and the n training points, $\boldsymbol{K}_{nm}$ is its transpose, and $\boldsymbol{K}_{mm}$ is the covariance matrix involving just the m distinguished inputs. We avoid computing all of the elements of $\boldsymbol{C}_*$, and do not need to so in order to make predictions.

[4] This is to be preferred to a finite-difference gradient used in a gradient-based algorithm.

The noise term of eq. (B.3) is excluded when calculating the $\boldsymbol{K}$ matrices of various shapes. The noise on the inputs introduces an additional diagonal term onto $\boldsymbol{C}$, but a similar term would also be desirable for stability of expressions involving $\boldsymbol{K}_{mm}^{-1}$. A small diagonal term can be added to $\boldsymbol{K}_{mm}$ as well, but this can no longer be identified straightforwardly with input-independent Gaussian noise, and serves only to aid numerical stability.

Under this approximation, the predicted mean of an observation at $\boldsymbol{x}'$ is

$$\hat{t}' = \boldsymbol{k}_m \boldsymbol{Q}^{-1} \boldsymbol{K}_{mn} \boldsymbol{t} \tag{B.19}$$

and the variance is

$$\mathrm{Var}(t') = \kappa - (\boldsymbol{k}_m^\mathsf{T} \boldsymbol{K}_{mm}^{-1} \boldsymbol{k}_m - \nu^2 \boldsymbol{k}_m^\mathsf{T} \boldsymbol{Q}^{-1} \boldsymbol{k}_m), \tag{B.20}$$

where $\boldsymbol{Q} = \nu^2 \boldsymbol{K}_{mm} + \boldsymbol{K}_{mn} \boldsymbol{K}_{nm}$ and $\boldsymbol{k}_m$ has components $C(\boldsymbol{x}^{(i)}, \boldsymbol{x}')$ with i ranging over the indices of the m distinguished inputs. The expected errors in these predictions against those of the full-rank process are discussed by Quiñonero Candela and Rasmussen (2005).

B.5.1. Choosing the distinguished inputs

The reduced-rank approximation with lowest elementwise error is the truncated singular-value decomposition, obtained by leaving out the smallest elements of the diagonal matrix in the full decomposition (the Eckart–Young theorem; see e.g. Van Huffel and Vandewalle, 1991, p.31). The advantage of the approach described above (sometimes known as a skeleton decomposition) over one based on singular values, is that the rows of $\boldsymbol{K}_{mn}$ in the approximation are the same as rows of the full-rank matrix, and only some of these need to be computed—a truncated singular value decomposition would require $O(n^2)$ space.

The best approximation with the form of eq. (B.18) maximizes the determinant of the intersection matrix $\boldsymbol{K}_{mm}$ (Chiu and Demanet, 2013). For a covariance matrix this is achieved by choosing uncorrelated training points, which are usually far apart in input space, and this admits an algorithm such as k-means clustering, described below.

B.5.2. k-means clustering

The method of k-means clustering (MacQueen, 1967) can be used with the reduced-rank Gaussian process to find a good set of distinguished input vectors heuristically. The algorithm partitions the n input points into m clusters whose variance is minimized. The

means of the clusters can be used as the input points.

$$\min_{S_i} \sum_{i=1}^{m} \sum_{x \in S_i} |x - \mu_i|^2, \tag{B.21}$$

where

$$\mu_i = \frac{1}{|S_i|} \sum_{x \in S_i} x \tag{B.22}$$

are the means of the clusters.

The algorithm proceeds iteratively, by assigning each point to the set with the closest mean to it at the current iteration, and then recomputing the mean for the next iteration. The initial partitioning can be chosen randomly and the minimizer obtained from several initial partitions used to avoid a poor local minimum.

B.6. A GPU implementation

In a continuum simulation, many independent equation of state evaluations are needed across the domain, and these can be made in parallel for a given timestep. A Gaussian process prediction is made by performing a predetermined number of covariance function evaluations followed by a dot product with a fixed vector. In the dense case, this involves computing the vector $\boldsymbol{k}$ in eq. (B.12), followed by a dot product with the precomputed vector $C^{-1}\boldsymbol{t}$ (or for the reduced-rank case, with the vector $Q^{-1}K_{mn}\boldsymbol{t}$). Since this latter vector is the same for every prediction, several predictions can be made at once as a matrix-vector product. These features make it a high-throughput task amenable to a GPU implementation. This is done because a single dot product is unlikely to be able to reach the full performance of the GPU, due to the relative cost of the transfer to and from the device.

The idea of our implementation is to package up a number of equation of state inputs, send them to a kernel function implementing Gaussian process prediction, which performs the covariance function evaluations (one per thread) followed by the matrix-vector product on the GPU. This latter step is performed by CUBLAS, a GPU library for performing linear algebra operations (Nvidia Corporation, 2014).

Performance data is reported in table B.1. Double precision is used throughout—single precision gives wildly inaccurate results (>100% relative error) in realistic cases due to the typically poor condition number of the covariance matrix C. There is a constant overhead associated with OpenMP in the multithreaded CPU implementation, and a

similar overhead in the GPU implementation for the launch of the kernel function which includes transfer of data to and from the device. These overheads can be seen to become less significant for the larger runs. For comparison, the Romenski equation of state eq. (A.5) took 0.4 μs per evaluation on the same CPU.

Table B.1.: Comparison of timing data for our CPU and GPU implementations of (full-rank) Gaussian process regression with a Gaussian covariance function [eq. (B.3)]. Times are the best of ten runs.

Database size	Number of evaluations	Time per evaluation (μs)		
		CPU[a], single thread	CPU[a], 16 threads	GPU[b,c]
100	1	95.1	5570	1000
	10	21.9	456	100
	100	14.5	50.5	10.2
	1000	13.7	8.01	1.26
	10 000	13.6	2.32	0.360
	20 000	13.6	1.53	0.320
2000	1	371	4690	1080
	10	284	626	108
	100	274	93.2	14.3
	1000	273	23.1	4.92
	10 000	274	19.1	4.04
	20 000	274	18.4	3.99

[a] Intel Xeon E5-2650 v2 with 16 cores at 2.60 GHz, ideal double precision performance of 333 GFLOPS

[b] NVIDIA K20, ideal double precision performance of 1170 GFLOPS

[c] including transfer time to and from the device for the input and result but not for the (constant) training data

C. Empirical Interatomic Potentials

This chapter gives a detailed description of the empirical interatomic potentials used for this thesis (GAP and DFT are described separately in their own chapters).

C.1. The Stillinger–Weber potential

This potential was described by Stillinger and Weber (1985), and is a three body potential, with a term favouring tetrahedral coordination. The form is quite simple, making it cheap to evaluate.

$$E/\epsilon = \sum_{i \leq j} V_2(r_{ij}/\sigma) + \sum_{i \leq j \leq k} V_3(\boldsymbol{r}_i/\sigma, \boldsymbol{r}_j/\sigma, \boldsymbol{r}_k/\sigma) \tag{C.1}$$

where

$$V_2(r) = \begin{cases} A(Br^{-p} - r^{-q})\exp[(r-a)^{-1}], & r < a \\ 0, & r \geq a, \end{cases} \tag{C.2}$$

and

$$V_3(\boldsymbol{r}_i, \boldsymbol{r}_j, \boldsymbol{r}_k) = \sum_{\mathrm{cyc}(i,j,k)} h(r_{ij}, r_{ik}, \theta_{jik}) \tag{C.3}$$

with

$$h(r_{ij}, r_{ik}, \theta_{jik}) = \begin{cases} \lambda \exp[\gamma(r_{ij}-a)^{-1} + \gamma(r_{ik}-a)^{-1}] & \\ \qquad \times \left(\cos\theta_{jik} + \frac{1}{3}\right)^2 & r_{ij} < a \text{ and } r_{ik} < a \\ 0 & \text{otherwise.} \end{cases} \tag{C.4}$$

In the above expressions, a is a cutoff radius, and the potential tends smoothly to zero at this distance. The form of h is chosen to favour tetrahedral angles between atoms (with $\cos\theta_{jik} = -\frac{1}{3}$). There are nine adjustable parameters including the cutoff and choice of

energy and distance units. For silicon, we use the original parameterization of Stillinger and Weber (1985), given in table C.1.

Table C.1.: Parameters for silicon, for the potential of Stillinger and Weber (1985).

ϵ (eV)	σ (Å)	a	λ	γ	A	B	p	q
2.1683	2.0951	1.80	21.0	1.20	7.049556277	0.6022245584	4	0

C.2. The Potential of Tersoff

This potential was first descibed by Tersoff (1986), and improved parameters were given subsequently by (Tersoff, 1988, 1989), with better elastic properties and for multicomponent systems. It is designed to describe tetrahedral semiconductors well, including carbon, silicon, and germanium.

The form of the potential is

$$E = \frac{1}{2}\sum_{i\neq j} V_{ij} \tag{C.5}$$

$$V_{ij} = f_C(r_{ij})\left[a_{ij} f_R(r_{ij}) + b_{ij} f_A(r_{ij})\right] \tag{C.6}$$

$$f_R(r) = A\exp(-\lambda_1 r) \tag{C.7}$$

$$f_A(r) = -B\exp(-\lambda_2 r) \tag{C.8}$$

$$f_C(r) = \begin{cases} 1, & r \leq R-D \\ \frac{1}{2} - \frac{1}{2}\sin\left[\frac{\pi}{2}(r-R)/D\right], & R-D < r \leq R+D \\ 0, & r > R+D \end{cases} \tag{C.9}$$

$$b_{ij} = (1+\beta^n \zeta_{ij}^n)^{-1/2n} \tag{C.10}$$

$$\zeta_{ij} = \sum_{k\neq\{i,j\}} f_C(r_{ik}) g(\theta_{ijk}) \exp[\lambda_3^3 (r_{ij} - r_{ik})^3] \tag{C.11}$$

$$g(\theta) = 1 + c^2/d^2 - c^2/[d^2 + (h - \cos\theta)^2] \tag{C.12}$$

$$a_{ij} = (1+\alpha^n \eta_{ij}^n)^{-1/2n} \tag{C.13}$$

$$\eta_{ij} = \sum_{k\neq\{i,j\}} f_C(r_{ij}) \exp[\lambda_3^3 (r_{ij} - r_{ik})^3], \tag{C.14}$$

where r_{ij} is the distance between atoms i and j and θ_{ijk} is the angle between $\boldsymbol{r}_{ij}$ and $\boldsymbol{r}_{jk}$. The thirteen adjustable parameters (with units where applicable) are A (eV), B (eV), $\lambda_{1,2,3}$ (Å^{-1}), α, β, n, c, d, h, R (Å) and D (Å).

We use the paramaterization for silicon of Tersoff (1989), given in table C.2.

Table C.2.: Parameters for silicon, for the potential of Tersoff (1989).

A	(eV)	1.8308×10^3
B	(eV)	4.7118×10^2
λ_1	(Å^{-1})	2.4799
λ_2	(Å^{-1})	1.7322
λ_3	(Å^{-1})	0.0
α		0.0
β		1.1×10^{-6}
n		7.8734×10^{-1}
c		1.0039×10^5
d		1.6217×10^1
h		-5.9825×10^{-1}
R	(Å)	2.7
S	(Å)	3.0

D. The Virial Stress

We show here that the virial stress is the (equilibrium) microscopic expression for the Cauchy stress in the thermodynamic limit. For further discussion of stress in microscopic systems and reconciling this with the continuum expression, see Admal and Tadmor (2010).

Consider the following time derivative

$$\frac{\mathrm{d}}{\mathrm{d}t}\sum^{N} \boldsymbol{r}\boldsymbol{p} = \sum^{N} \boldsymbol{v}\boldsymbol{p} + \sum^{N} \boldsymbol{r}\dot{\boldsymbol{p}} \tag{D.1}$$

$$= \sum^{N} \boldsymbol{v}\boldsymbol{p} + \sum^{N} \boldsymbol{r}\boldsymbol{f} \tag{D.2}$$

$$= 2\mathscr{T} + \mathscr{W} \tag{D.3}$$

where $\mathscr{T}$ is the kinetic tensor and $\mathscr{W}$ is known as the virial tensor.

Splitting the virial term into an 'internal' part and and 'external' part, we identify the external part of the per-atom forces with a continuum traction, which it becomes in the limit of large system size:

$$\mathscr{W}^{\mathrm{ext}} = \int_{\partial\mathscr{V}} \boldsymbol{r}\boldsymbol{\sigma}\cdot\mathrm{d}\boldsymbol{A}. \tag{D.4}$$

This expression is written in terms of the continuum Cauchy stress, and relies on the result that the traction through a surface is the sum of forces on the atoms that we cannot account for within the system.

By the divergence theorem (under the summation convention),

$$\mathscr{W}^{\mathrm{ext}}_{ij} = \int_{\mathscr{V}} \partial_k(r_i\sigma_{kj})\,\mathrm{d}V = \int_{\mathscr{V}} (r_i\partial_k\sigma_{kj} + \sigma_{ij})\,\mathrm{d}V = V\sigma_{ij}, \tag{D.5}$$

where the last result holds in equilibrium, with the stress constant throughout the system.

In the continuum limit, for a system in equilibrium, we therefore have

$$\frac{1}{V}\left(\frac{\mathrm{d}}{\mathrm{d}t}\sum^{N} \boldsymbol{r}\boldsymbol{p} - 2\mathscr{T} - \mathscr{W}^{\mathrm{int}}\right) \overset{\text{continuum}}{\longrightarrow} \boldsymbol{\sigma}, \tag{D.6}$$

where the limit takes $V \to \infty$ holding N/V constant.

Finally, we define the *instantaneous virial stress* as

$$\boldsymbol{\sigma}^{\mathrm{virial}} = -\frac{1}{V}\left(2\mathscr{T} + \mathscr{W}^{\mathrm{int}}\right) \tag{D.7}$$

Taking a (finite) time average, we obtain

$$\overline{\boldsymbol{\sigma}^{\mathrm{virial}}} = \boldsymbol{\sigma} - \frac{1}{T}\left(\frac{1}{V}\sum^{N} \boldsymbol{r}\boldsymbol{p}\right)\Bigg|_{0}^{T}. \tag{D.8}$$

In many circumstances, the averaged quantity on the right hand side of this equation tends to zero as T grows large.

Equation (D.7) may alternately be written

$$\boldsymbol{\sigma}^{\mathrm{virial}} = -\frac{1}{V}\sum_{\alpha}^{N}(\boldsymbol{f}^{\alpha}\boldsymbol{r}^{\alpha} + m^{\alpha}\boldsymbol{v}^{\alpha}\boldsymbol{v}^{\alpha}) \tag{D.9}$$

where the sum is over the atoms, indexed by α.

E. The Pair Correlation Function

In a system of particles, the *pair correlation function*, or *radial distribution function*, $g(r)$, is a measure of the average density variation with distance from a given reference particle or species.

It is defined through

$$\rho^{(1)}(r) = \rho g(r) \tag{E.1}$$

where ρ is the bulk density of the system, and

$$\rho^{(1)}(r) = \int_{|r|=r} \rho^{(1)}(\boldsymbol{r})\,\mathrm{d}\boldsymbol{r}. \tag{E.2}$$

$\rho^{(1)}(\boldsymbol{r})$ is the distribution of density with spatial position.

In terms of the probability density of a configuration,

$$\rho^{(1)}(\boldsymbol{r}) = N \underbrace{\int \cdots \int}_{N-1} \mathbb{P}(\boldsymbol{r}, \boldsymbol{r}_2, \ldots, \boldsymbol{r}_N)\,\mathrm{d}\boldsymbol{r}_2 \cdots \mathrm{d}\boldsymbol{r}_N, \tag{E.3}$$

and for a system of particles whose potential energy depends on their configuration as $U(\{\boldsymbol{r}_i\})$,

$$\mathbb{P}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_N) = \frac{1}{Z} \exp\left(-\frac{U(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_N)}{k_B T}\right), \tag{E.4}$$

where Z is the usual partition function, and so, with system volume V,

$$g(r) = \frac{V}{Z} \int_{|\boldsymbol{r}|=r} \underbrace{\int \cdots \int}_{N-1} \exp\left(-\frac{U(\boldsymbol{r}, \boldsymbol{r}_2, \ldots, \boldsymbol{r}_N)}{k_B T}\right) \mathrm{d}\boldsymbol{r}_2 \cdots \mathrm{d}\boldsymbol{r}_N\,\mathrm{d}\boldsymbol{r}. \tag{E.5}$$

Calculation of $g(r)$ numerically can be done by evaluating the integral eq. (E.2) discretely, averaging over particles and time.

F. LCAO Basis for Silicon

Table F.1.: Basis parameters for a double-ζ polarized basis set for silicon, according to the soft-confinement scheme of Junquera et al. (2001). The first ζ for $l = 0$ and $l = 1$ was set to $7.0\,a_0$, and the other parameters were variationally optimized. There are thirteen basis functions per atom. For the purposes of basis generation, an effective ionic charge of -0.464 was used, which was also variationally optimized. The cutoff radii of the first and second zeta functions are $r(\zeta_1)$ and $r(\zeta_2)$, and r_i is the confinement potential's internal radius. V_0 is the soft-confinement prefactor.

n	l	r_i (a_0)	$r(\zeta_1)$ (a_0)	$r(\zeta_2)$ (a_0)	V_0 (Ry)
3	0	4.970	7.000	4.377	15.426
3	1	3.831	7.000	4.091	4.696
3	2	0.031	4.554	-	11.969

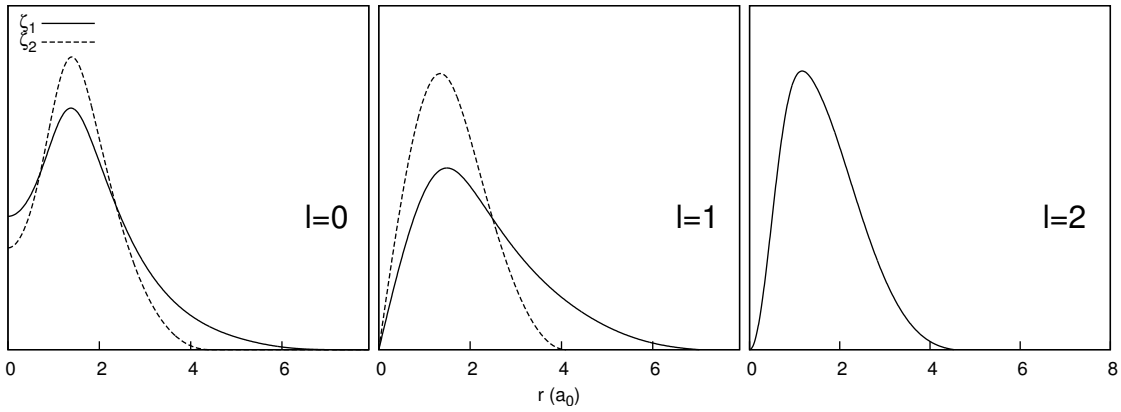

Figure F.1.: The radial basis shapes resulting from the basis parameters in table F.1.

Bibliography

F. F. Abraham, J. Q. Broughton, N. Bernstein, and E. Kaxiras. Spanning the continuum to quantum length scales in a dynamic simulation of brittle fracture. *EPL (Europhysics Letters)*, 44(6):783, 1998. doi: 10.1209/epl/i1998-00536-9.

C. Abs da Cruz, K. Termentzidis, P. Chantrenne, and X. Kleber. Molecular dynamics simulations for the prediction of thermal conductivity of bulk silicon and silicon nanowires: Influence of interatomic potentials and boundary conditions. *J. Appl. Phys.*, 110(3):034309–034309, 2011.

N. C. Admal and E. Tadmor. A Unified Interpretation of Stress in Molecular Systems. *J. Elast.*, 100(1-2):63–143, 2010. ISSN 0374-3535.

T. Ahrens, V. Gregson, and S. R. Institute. *Shock compression of crustal rocks: Data for quartz, calcite, and plagioclase rocks*. Stanford Research Institute, 1964.

P. Allen and D. Tildesley. *Computer Simulation of Liquids*. Oxford science publications. Clarendon Press, 1987. ISBN 9780198553755.

J. Asay and M. Shahinpoor, editors. *High-Pressure Shock Compression of Solids*. Shock Wave and High Pressure Phenomena. Springer-Verlag New York, 1993. ISBN 978-0-387-97964-9. doi: 10.1007/978-1-4612-0911-9.

N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Holt-Saunders, 1976.

A. Baldereschi. Mean-value point in the brillouin zone. *Phys. Rev. B*, 7:5212–5215, Jun 1973. doi: 10.1103/PhysRevB.7.5212.

D. S. Balsara and C.-W. Shu. Monotonicity preserving weighted essentially non-oscillatory schemes with increasingly high order of accuracy. *J. Comput. Phys.*, 160 (2):405–452, 2000. ISSN 0021-9991. doi: 10.1006/jcph.2000.6443.

S. Barnett. *Matrix Methods for Engineers and Scientists*. McGraw-Hill Book Company, 1979. ISBN 9780070840843.

V. Barthelmann, E. Novak, and K. Ritter. High dimensional polynomial interpolation on sparse grids. *Advance. Comput. Math.*, 12(4):273–288, 2000. ISSN 1019-7168. doi: 10.1023/A:1018977404843.

A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi. Gaussian Approximation Potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, 104 (13):136403, 2010.

A. P. Bartók, M. J. Gillan, F. R. Manby, and G. Csányi. Machine-learning approach for one- and two-body corrections to density functional theory: Applications to molecular and condensed water. *Phys. Rev. B*, 88:054104, Aug 2013a. doi: 10.1103/PhysRevB.88.054104.

A. P. Bartók, R. Kondor, and G. Csányi. On representing chemical environments. *Phys. Rev. B*, 87:184115, May 2013b. doi: 10.1103/PhysRevB.87.184115.

A. P. Bartók-Pártay, S. Cereda, G. Csányi, J. Kermode, I. Solt, W. Szlachta, C. Varnai, and S. Winfield. `libAtoms+QUIP`: software library for carrying out molecular dynamics simulations., 2006–. URL `http://www.libatoms.org`.

P. Barton, D. Drikakis, E. Romenski, and V. Titarev. Exact and approximate solutions of riemann problems in non-linear elasticity. *J. Comput. Phys.*, 228(18):7046–7068, 2009. ISSN 0021-9991. doi: 10.1016/j.jcp.2009.06.014.

M. Z. Bazant, E. Kaxiras, and J. F. Justo. Environment-dependent interatomic potential for bulk silicon. *Phys. Rev. B*, 56:8542–8552, Oct 1997. doi: 10.1103/PhysRevB.56.8542.

A. D. Becke. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A*, 38:3098–3100, Sep 1988. doi: 10.1103/PhysRevA.38.3098.

J. Behler and M. Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, Apr 2007. doi: 10.1103/PhysRevLett.98.146401.

H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak. Molecular dynamics with coupling to an external bath. *J. Chem. Phys.*, 81(8):3684–3690, 1984. doi: 10.1063/1.448118.

M. Born and R. Oppenheimer. Zur Quantentheorie der Molekeln [On the quantum theory of molecules]. *Ann. Phys. (Leipzig)*, 389(20):457–484, 1927.

D. A. Broido, A. Ward, and N. Mingo. Lattice thermal conductivity of silicon from empirical interatomic potentials. *Phys. Rev. B*, 72(1):014308, 2005.

D. A. Broido, M. Malorny, G. Birner, N. Mingo, and D. A. Stewart. Intrinsic lattice thermal conductivity of semiconductors from first principles. *Appl. Phys. Lett.*, 91 (23):231922, 2007. doi: 10.1063/1.2822891.

J. Q. Broughton and X. P. Li. Phase diagram of silicon by molecular dynamics. *Phys. Rev. B*, 35:9120–9127, Jun 1987.

P. Bunker and R. Moss. The breakdown of the born-oppenheimer approximation: the effective vibration-rotation hamiltonian for a diatomic molecule. *Mol. Phys.*, 33(2): 417–424, 1977. doi: 10.1080/00268977700100351.

J. Callaway. Model for lattice thermal conductivity at low temperatures. *Phys. Rev.*, 113:1046–1051, Feb 1959. doi: 10.1103/PhysRev.113.1046.

M. M. Carroll. Must Elastic Materials be Hyperelastic? *Math. Mech. Solids*, 2009.

P. Celliers, A. Ng, G. Xu, and A. Forsman. Thermal equilibration in a shock wave. *Phys. Rev. Lett.*, 68:2305–2308, Apr 1992. doi: 10.1103/PhysRevLett.68.2305.

D. M. Ceperley and B. J. Alder. Ground state of the electron gas by a stochastic method. *Phys. Rev. Lett.*, 45:566–569, Aug 1980. doi: 10.1103/PhysRevLett.45.566.

A. Chentsov and P. Levashov. Isentropic compression of deuterium by quantum molecular dynamics. *Contrib. Plasma Phys.*, 52(1):33–36, 2012.

L. C. Chhabildas and J. R. Asay. RiseâĂŘtime measurements of shock transitions in aluminum, copper, and steel. *Journal of Applied Physics*, 50(4):2749–2756, 1979. doi: 10.1063/1.326236.

J. Chiu and L. Demanet. Sublinear randomized algorithms for skeleton decompositions. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1361–1383, 2013.

S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. Probert, K. Refson, and M. C. Payne. First principles methods using castep. *Zeitschrift für Kristallographie*, 220(5-6):567–570, 2005.

A. J. Cohen, P. Mori-Sánchez, and W. Yang. Challenges for density functional theory. *Chem. Rev.*, 112(1):289–320, 2012. doi: 10.1021/cr200107z.

J. Conway and N. Sloane. What are all the best sphere packings in low dimensions? *Discrete Comput. Geom.*, 13(1):383–403, 1995. ISSN 0179-5376. doi: 10.1007/BF02574051.

W. A. Curtin and R. E. Miller. Atomistic/continuum coupling in computational materials science. *Modelling Simul. Mater. Sci. Eng.*, 11(3):R33, 2003. doi: 10.1088/0965-0393/11/3/201.

D. D. Dlott. New developments in the physical chemistry of shock compression. *Ann. Rev. Phys. Chem.*, 62(1):575–597, 2011. doi: 10.1146/annurev.physchem.012809.103514.

D. S. Dolling. Fifty years of shock-wave/boundary-layer interaction research: what next? *AIAA journal*, 39(8):1517–1531, 2001. doi: 10.2514/2.1476.

R. E. Duff and F. S. Minshall. Investigation of a shock-induced transition in bismuth. *Phys. Rev.*, 108:1207–1212, Dec 1957. doi: 10.1103/PhysRev.108.1207.

G. E. Duvall and R. A. Graham. Phase transitions under shock-wave loading. *Rev. Mod. Phys.*, 49:523–579, Jul 1977. doi: 10.1103/RevModPhys.49.523.

J. A. Elliott. Novel approaches to multiscale modelling in materials science. *Int. Mater. Rev.*, 56(4):207–225, 2011. doi: 10.1179/1743280410Y.0000000002.

N. Ercolani, editor. *Singular Limits of Dispersive Waves*. Series B: Physics Vol. 320, NATO ASI Series. Springer, 1994. ISBN 9780306446283.

P. Erhart and K. Albe. Analytical potential for atomistic simulations of silicon, carbon, and silicon carbide. *Phys. Rev. B*, 71:035211, Jan 2005. doi: 10.1103/PhysRevB.71.035211.

K. Esfarjani, G. Chen, and H. T. Stokes. Heat transport in silicon from first-principles calculations. *Phys. Rev. B*, 84(8):085204, 2011.

D. Frenkel and B. Smit. *Understanding molecular simulation: from algorithms to applications*, volume 1. Academic Press, 2002.

Y. Fung and P. Tong. *Classical and Computational Solid Mechanics*. Advanced series in engineering science. World Scientific, 2001. ISBN 9789810241247.

J. Garg. Thermal conductivity from first-principles in bulk, disordered, and nanostructured materials, 2011.

G. Gathers. *Selected topics in shock wave physics and equation of state modeling*. World Scientific, 1994.

G. Gerolymos, D. Sénéchal, and I. Vallet. Very-high-order WENO schemes. *J. Comput. Phys.*, 228(23):8481–8524, 2009. ISSN 0021-9991. doi: 10.1016/j.jcp.2009.07.039.

S. Ghosh, I. Calizo, D. Teweldebrhan, E. P. Pokatilov, D. L. Nika, A. A. Balandin, W. Bao, F. Miao, and C. N. Lau. Extremely high thermal conductivity of graphene: Prospects for thermal management applications in nanoelectronic circuits. *Appl. Phys. Lett.*, 92(15):151911, 2008. doi: 10.1063/1.2907977.

M. Gibbs and D. J. C. Mackay. Efficient implementation of gaussian processes. Technical report, Cavendish Laboratory, University of Cambridge, UK, 1997.

C. Glassbrenner and G. A. Slack. Thermal conductivity of silicon and germanium from 3 K to the melting point. *Phys. Rev*, 134(4A):A1058–A1069, 1964.

T. Goto, T. Sato, and Y. Syono. Reduction of shear strength and phase-transition in shock-loaded silicon. *Jpn. J. Appl. Phys.*, 21(6A):L369–L371, 1982. doi: 10.1143/JJAP.21.L369.

S. Gottlieb, C.-W. Shu, and E. Tadmor. Strong stability-preserving high-order time discretization methods. *SIAM Review*, 43(1):89–112, 2001. doi: 10.1137/S003614450036757X.

R. Graham and W. Brooks. Shock-wave compression of sapphire from 15 to 420 kbar. The effects of large anisotropic compressions. *J. Phys. Chem. Solids*, 32(10):2311–2330, 1971. ISSN 00223697. doi: 10.1016/S0022-3697(71)80226-3.

L. Grasedyck, D. Kressner, and C. Tobler. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen*, 36(1):53–78, 2013.

W. Gust and E. Royce. Dynamic yield strengths of light armor materials. Technical report, California Univ., Livermore. Lawrence Radiation Lab., 1970.

W. Gust and E. Royce. Axial yield strengths and two successive phase transition stresses for crystalline silicon. *J. Appl. Phys.*, 42(5):1897–1905, 1971. doi: 10.1063/1.1660465.

M. Hafez, K. Ōshima, and D. Kwak. *Computational Fluid Dynamics Review 2010*. World Scientific, 2010. ISBN 9789814313360.

J. Hafner, C. Wolverton, and G. Ceder. Toward computational materials design: The impact of density functional theory on materials research. *MRS Bulletin*, 31:659–668, 9 2006. ISSN 1938-1425. doi: 10.1557/mrs2006.174.

E. Hairer, C. Lubich, G. Wanner, et al. Geometric numerical integration illustrated by the Stormer-Verlet method. *Acta Numerica*, 12:399–450, 2003.

J. J. Hall. Electronic effects in the elastic constants of *n*-type silicon. *Phys. Rev.*, 161: 756–761, Sep 1967. doi: 10.1103/PhysRev.161.756.

D. R. Hamann, M. Schlüter, and C. Chiang. Norm-conserving pseudopotentials. *Phys. Rev. Lett.*, 43:1494–1497, Nov 1979. doi: 10.1103/PhysRevLett.43.1494.

S. C. Harvey, R. K.-Z. Tan, and T. E. Cheatham. The flying ice cube: Velocity rescaling in molecular dynamics leads to violation of energy equipartition. *J. Comput. Chem.*, 19(7):726–740, 1998. ISSN 1096-987X. doi: 10.1002/(SICI)1096-987X(199805)19:7<726::AID-JCC4>3.0.CO;2-S.

Y. He, I. Savić, D. Donadio, and G. Galli. Lattice thermal conductivity of semiconducting bulk materials: atomistic simulations. *Phys. Chem. Chem. Phys.*, 14(47):16209–16222, 2012.

L. E. Henderson. General laws for propagation of shcok waves through matter. In O. I. Gabi Ben-Dor and T. Elperin, editors, *Handbook of Shock Waves, Volume 1*. Academic Press, 2001.

A. K. Henrick, T. D. Aslam, and J. M. Powers. Mapped weighted essentially non-oscillatory schemes: Achieving optimal order near critical points. *J. Comput. Phys.*, 207(2):542–567, 2005. ISSN 0021-9991. doi: 10.1016/j.jcp.2005.01.023.

N. Higham. *Functions of Matrices: Theory and Computation*. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics, 2008. ISBN 9780898716467.

C. Y. Ho, R. W. Powell, and P. E. Liley. Thermal conductivity of the elements. *Journal of Physical and Chemical Reference Data*, 1(2):279–421, 1972. doi: http://dx.doi.org/10.1063/1.3253100.

C. Y. Ho, M. W. Ackerman, K. Y. Wu, S. G. Oh, and T. N. Havill. Thermal conductivity of ten selected binary alloy systems. *J. Phys. Chem. Ref. Data*, 7(3):959–1178, 1978. doi: 10.1063/1.555583.

P. Hoel, S. Port, and C. Stone. *Introduction to Stochastic Processes*. Waveland Press, 1987. ISBN 9780881332674.

P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136:B864–B871, Nov 1964. doi: 10.1103/PhysRev.136.B864.

B. Holian. Molecular dynamics comes of age for shockwave research. *Shock Waves*, 13 (6):489–495, June 2004. ISSN 0938-1287. doi: 10.1007/s00193-004-0226-5.

R. Hull. *Properties of Crystalline Silicon*. EMIS datareviews series. INSPEC, the Institution of Electrical Engineers, 1999. ISBN 9780852969335.

A. V. Inyushkin, A. N. Taldenkov, A. M. Gibin, A. V. Gusev, and H.-J. Pohl. On the isotope effect in thermal conductivity of silicon. *physica status solidi (c)*, 1(11):2995–2998, 2004. ISSN 1610-1642. doi: 10.1002/pssc.200405341.

J. C. Jamieson. Crystal structures at high pressures of metallic modifications of silicon and germanium. *Science*, 139(3556):762–764, 1963. doi: {10.1126/science.139.3556.762}.

G.-S. Jiang and C.-W. Shu. Efficient implementation of weighted ENO schemes. *J. Comput. Phys.*, 126(1):202–228, 1996. ISSN 0021-9991. doi: 10.1006/jcph.1996.0130.

S. G. Johnson. The `NLopt` nonlinear optimization package (version 2.4.2), 2014. URL `http://ab-initio.mit.edu/nlopt`.

R. O. Jones and O. Gunnarsson. The density functional formalism, its applications and prospects. *Rev. Mod. Phys.*, 61(3):689–746, July 1989. ISSN 0034-6861. doi: 10.1103/RevModPhys.61.689.

J. Junquera, O. Paz, D. Sánchez-Portal, and E. Artacho. Numerical atomic orbitals for linear-scaling calculations. *Phys. Rev. B*, 64:235111, Nov 2001. doi: 10.1103/PhysRevB.64.235111.

J. a. F. Justo, M. Z. Bazant, E. Kaxiras, V. V. Bulatov, and S. Yip. Interatomic potential for silicon defects and disordered phases. *Phys. Rev. B*, 58:2539–2550, Aug 1998. doi: 10.1103/PhysRevB.58.2539.

K. Kadau, T. C. Germann, and P. S. Lomdahl. Molecular dynamics comes of age: 320 billion atom simulation on BlueGene/L. *Int J Mod Phys C*, 17(12):1755–1761, 2006. doi: 10.1142/S0129183106010182.

M. Kazan, G. Guisbiers, S. Pereira, M. Correia, P. Masri, A. Bruyant, S. Volz, and P. Royer. Thermal conductivity of silicon bulk and nanowires: Effects of isotopic composition, phonon confinement, and surface roughness. *Journal of Applied Physics*, 107(8):083503–083503, 2010.

G. P. Kerker. Non-singular atomic pseudopotentials for solid state applications. *J. Phys. C: Solid State Phys.*, 13(9):L189, 1980. doi: 10.1088/0022-3719/13/9/004.

L. Kleinman and D. M. Bylander. Efficacious form for model pseudopotentials. *Phys. Rev. Lett.*, 48:1425–1428, May 1982. doi: 10.1103/PhysRevLett.48.1425.

W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133–A1138, Nov 1965. doi: 10.1103/PhysRev.140.A1133.

R. Kremer, K. Graf, M. Cardona, G. Devyatykh, A. Gusev, A. Gibin, A. Inyushkin, A. Taldenkov, and H.-J. Pohl. Thermal conductivity of isotopically enriched ^{28}Si: revisited. *Solid State Commun.*, 131(8):499–503, 2004. ISSN 0038-1098. doi: 10.1016/j.ssc.2004.06.022.

A. Kubo, Y. Wang, C. E. Runge, T. Uchida, B. Kiefer, N. Nishiyama, and T. S. Duffy. Melting curve of silicon to 15 GPa determined by two-dimensional angle-dispersive diffraction using a Kawai-type apparatus with X-ray transparent sintered diamond anvils. *Journal of Physics and Chemistry of Solids*, 69(9):2255–2260, 2008. ISSN 0022-3697. doi: http://dx.doi.org/10.1016/j.jpcs.2008.04.025.

R. Kubo. Statistical-mechanical theory of irreversible processes. I. General theory and simple applications to magnetic and conduction problems. *J. Phys. Soc. Jpn.*, 12(6): 570–586, 1957.

L. Landau and E. Lifshitz. *Statistical Physics*, volume 5 of *Course of Theoretical Physics*. Elsevier Science, third edition, 1980.

P. Lazar. *Ab initio modelling of mechanical and elastic properties of solids*. PhD thesis, Faculty of Physics, University of Vienna, 2006.

S. Ledvina, Y.-J. Ma, and E. Kallio. Modeling and simulating flowing plasmas and related phenomena. *Space Sci. Rev.*, 139(1-4):143–189, 2008. ISSN 0038-6308. doi: 10.1007/s11214-008-9384-6.

C. Lee, W. Yang, and R. G. Parr. Development of the colle-salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B*, 37:785–789, Jan 1988. doi: 10.1103/PhysRevB.37.785.

B. Leimkuhler and S. Reich. *Simulating Hamiltonian Dynamics*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2004. ISBN 9780521772907.

R. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002. ISBN 9781139434188.

M. Levy. Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v-representability problem. *PNAS*, 76(12):6062–6065, 1979.

W. Li, J. Carrete, N. A. Katcho, and N. Mingo. ShengBTE: A solver of the Boltzmann transport equation for phonons. *Comput. Phys. Commun.*, 185(6):1747–1758, 2014. ISSN 0010-4655. doi: 10.1016/j.cpc.2014.02.015.

X.-D. Liu, S. Osher, and T. Chan. Weighted essentially non-oscillatory schemes. *J. Comput. Phys.*, 115(1):200–212, 1994. ISSN 0021-9991. doi: 10.1006/jcph.1994.1187.

S. Lyon and J. Johnson. SESAME: The Los Alamos National Laboratory Equation of State Database. Technical Report LA-UR-92-3407, Los Alamos National Laboratory, 1992.

D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003. ISBN 9780521642989.

J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, Calif., 1967. University of California Press.

J.-B. Maillet, M. Mareschal, L. Soulard, R. Ravelo, P. S. Lomdahl, T. C. Germann, and B. L. Holian. Uniaxial Hugoniostat: A method for atomistic simulations of shocked materials. *Phys. Rev. E*, 63:016121, Dec 2000. doi: 10.1103/PhysRevE.63.016121.

R. Martin. *Electronic Structure: Basic Theory and Practical Methods*. Cambridge University Press, 2004. ISBN 9780521782852.

M. I. McMahon, R. J. Nelmes, N. G. Wright, and D. R. Allan. Pressure dependence of the Imma phase of silicon. *Phys. Rev. B*, 50(2):739, 1994. doi: 10.1103/PhysRevB.50.739.

N. D. Mermin. Thermal properties of the inhomogeneous electron gas. *Phys. Rev.*, 137: A1441–A1443, Mar 1965. doi: 10.1103/PhysRev.137.A1441.

G. Miller. An iterative riemann solver for systems of hyperbolic conservation laws, with application to hyperelastic solid mechanics. *J. Comput. Phys.*, 193(1):198–225, 2004a. ISSN 0021-9991. doi: 10.1016/j.jcp.2003.08.005.

G. Miller and P. Colella. A high-order eulerian godunov method for elastic–plastic flow in solids. *J. Comput. Phys.*, 167(1):131–176, 2001. ISSN 0021-9991. doi: 10.1006/jcph.2000.6665.

G. H. Miller. Minimal rotationally invariant bases for hyperelasticity. *J. Appl. Math.*, 64 (6):2050–2075, 2004b. doi: 10.1137/S0036139903438776.

R. E. Miller and E. B. Tadmor. A unified framework and performance benchmark of fourteen multiscale atomistic/continuum coupling methods. *Modelling Simul. Mater. Sci. Eng.*, 17(5):053001, 2009. doi: 10.1088/0965-0393/17/5/053001.

S. Minomura and H. Drickamer. Pressure induced phase transitions in silicon, germanium and some III–V compounds. *J. Phys. Chem. Solids*, 23(5):451–456, 1962. doi: 10.1016/0022-3093(92)90093-Y.

G. Mogni, A. Higginbotham, K. Gaál-Nagy, N. Park, and J. S. Wark. Molecular dynamics simulations of shock-compressed single-crystal silicon. *Phys. Rev. B*, 89(6): 064104, 2014. doi: 10.1103/PhysRevB.89.064104.

H. J. Monkhorst and J. D. Pack. Special points for brillouin-zone integrations. *Phys. Rev. B*, 13:5188–5192, Jun 1976. doi: 10.1103/PhysRevB.13.5188.

J. Moreno and J. M. Soler. Optimal meshes for integrals in real- and reciprocal-space unit cells. *Phys. Rev. B*, 45:13891–13898, Jun 1992. doi: 10.1103/PhysRevB.45.13891.

A. Mujica, A. Rubio, A. Muñoz, and R. J. Needs. High-pressure phases of group-IV, III–V, and II–VI compounds. *Rev. Mod. Phys.*, 75:863–912, Jul 2003. doi: 10.1103/RevModPhys.75.863.

Nvidia Corporation. cuBLAS User Guide. *NVIDIA Corporation, Santa Clara, California*, 2014. URL `http://docs.nvidia.com/cuda/cublas`.

N. P. Padture, M. Gell, and E. H. Jordan. Thermal barrier coatings for gas-turbine engine applications. *Science*, 296(5566):280–284, 2002. doi: 10.1126/science.1068609. URL `http://www.sciencemag.org/content/296/5566/280.abstract`.

M. Pavlovskii. Formation of metallic modifications of germanium and silicon under shock loading. *Sov. Phys. Solid State*, 9:2514–2518, Feb 1968.

M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos. Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients. *Rev. Mod. Phys.*, 64:1045–1097, Oct 1992. doi: 10.1103/RevModPhys.64.1045.

J. P. Perdew and Y. Wang. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B*, 45:13244–13249, Jun 1992. doi: 10.1103/PhysRevB.45.13244.

J. P. Perdew and A. Zunger. Self-interaction correction to density-functional approximations for many-electron systems. *Phys. Rev. B*, 23:5048–5079, May 1981. doi: 10.1103/PhysRevB.23.5048.

J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh, and C. Fiolhais. Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation. *Phys. Rev. B*, 46:6671–6687, Sep 1992. doi: 10.1103/PhysRevB.46.6671.

J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.*, 77:3865–3868, Oct 1996. doi: 10.1103/PhysRevLett.77.3865.

J. C. Phillips and L. Kleinman. New method for calculating wave functions in crystals and molecules. *Phys. Rev.*, 116:287–294, Oct 1959. doi: 10.1103/PhysRev.116.287.

B. J. Plohr and D. H. Sharp. A conservative eulerian formulation of the equations for elastic flow. *Advances in Applied Mathematics*, 9(4):481–499, 1988. ISSN 0196-8858. doi: 10.1016/0196-8858(88)90025-5.

J. Qiu and C.-W. Shu. On the construction, comparison, and local characteristic decomposition for high-order central WENO schemes. *J. Comput. Phys.*, 183(1):187–209, 2002. ISSN 0021-9991. doi: 10.1006/jcph.2002.7191.

J. Quiñonero Candela and C. E. Rasmussen. A unifying view of sparse approximate gaussian process regression. *J. Mach. Learn. Res.*, 6:1939–1959, Dec. 2005. ISSN 1532-4435.

C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. Adaptative Computation and Machine Learning Series. MIT Press, 2006. ISBN 9780262182539.

S. Ratanaphan, Y. Yoon, and G. S. Rohrer. The five parameter grain boundary character distribution of polycrystalline silicon. *J. Mater. Sci.*, 49(14):4938–4945, 2014. ISSN 0022-2461. doi: 10.1007/s10853-014-8195-2.

R. Ravelo, B. L. Holian, T. C. Germann, and P. S. Lomdahl. Constant-stress Hugoniostat method for following the dynamical evolution of shocked matter. *Phys. Rev. B*, 70:014103, Jul 2004. doi: 10.1103/PhysRevB.70.014103.

E. J. Reed, L. E. Fried, and J. D. Joannopoulos. A Method for Tractable Dynamical Studies of Single and Double Shock Compression. *Phys. Rev. Lett.*, 90:235503, Jun 2003. doi: 10.1103/PhysRevLett.90.235503.

T. Ruf, R. Henn, M. Asen-Palmer, E. Gmelin, M. Cardona, H.-J. Pohl, G. Devyatych, and P. Sennikov. Thermal conductivity of isotopically enriched silicon. *Solid State Commun.*, 115(5):243–247, 2000.

T. Ruf, R. Henn, M. Asen-Palmer, E. Gmelin, M. Cardona, H.-J. Pohl, G. Devyatych, and P. Sennikov. Erratum to "thermal conductivity of isotopically enriched silicon" [solid state commun., 115 (2000) 243–247]. *Solid State Commun.*, 127(3):257, 2003. ISSN 0038-1098. doi: 10.1016/S0038-1098(03)00367-3.

M. G. Saunders and G. A. Voth. Coarse-graining methods for computational biology. *Ann. Rev. Biophys.*, 42(1):73–93, 2013. doi: 10.1146/annurev-biophys-083012-130348.

P. Schelling, S. Phillpot, and P. Keblinski. Comparison of atomic-level simulation methods for computing thermal conductivity. *Phys. Rev. B*, 65(14):144306, 2002.

N. Schuch and F. Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Physics*, 5(10):732–735, 2009.

D. P. Sellan, E. S. Landry, J. E. Turney, A. J. H. McGaughey, and C. H. Amon. Size effects in molecular dynamics thermal conductivity predictions. *Phys. Rev. B*, 81: 214305, Jun 2010. doi: 10.1103/PhysRevB.81.214305.

J. Shang. Three decades of accomplishments in computational fluid dynamics. *Progress in Aerospace Sciences*, 40(3):173–197, 2004. ISSN 0376-0421. doi: 10.1016/j.paerosci.2004.04.001.

H. R. Shanks, P. D. Maycock, P. H. Sidles, and G. C. Danielson. Thermal conductivity of silicon from 300 to 1400°k. *Phys. Rev.*, 130:1743–1748, Jun 1963. doi: 10.1103/PhysRev.130.1743.

A. Shekhar, K.-i. Nomura, R. K. Kalia, A. Nakano, and P. Vashishta. Nanobubble collapse on a silica surface in water: Billion-atom reactive molecular dynamics simulations. *Phys. Rev. Lett.*, 111:184503, Oct 2013. doi: 10.1103/PhysRevLett.111.184503.

C.-W. Shu and S. Osher. Efficient implementation of essentially non-oscillatory shock-capturing schemes. *J. Comput. Phys.*, 77(2):439–471, 1988. ISSN 0021-9991. doi: 10.1016/0021-9991(88)90177-5.

J. Simo and T. Hughes. *Computational Inelasticity*. Interdisciplinary Applied Mathematics. Springer New York, 2006. ISBN 9780387227634.

E. Snelson and Z. Ghahramani. Local and global sparse gaussian process approximations. In *International Conference on Artificial Intelligence and Statistics 11*, pages 524–531, 2007.

J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal. The SIESTA method for ab initio order-N materials simulation. *J. Phys.: Condens. Matter*, 14(11):2745, 2002. doi: 10.1088/0953-8984/14/11/302.

B. N. Srivastava, S. Chatterjee, S. K. Sen, and D. K. Chakraborty. Lattice thermal conductivity of some copper and silver alloys. *J. Phys. C: Solid State Phys.*, 3(2S):S169, 1970. doi: 10.1088/0022-3719/3/2S/307.

G. Srivastava. *The Physics of Phonons*. Adam Hilger, 1990. ISBN 9780852741535.

F. H. Stillinger and T. A. Weber. Computer simulation of local order in condensed phases of silicon. *Phys. Rev. B*, 31(8):5262, 1985.

D. C. Swift, G. J. Ackland, A. Hauer, and G. A. Kyrala. First-principles equations of state for simulations of shock waves in silicon. *Phys. Rev. B*, 64:214107, Nov 2001. doi: 10.1103/PhysRevB.64.214107.

W. M. Tang and V. S. Chan. Advances and challenges in computational plasma science. *Plasma Phys. Contr. F.*, 47(2):R1, 2005. doi: 10.1088/0741-3335/47/2/R01.

J. Tersoff. New empirical model for the structural properties of silicon. *Phys. Rev. Lett.*, 56:632–635, Feb 1986. doi: 10.1103/PhysRevLett.56.632.

J. Tersoff. Empirical interatomic potential for silicon with improved elastic properties. *Phys. Rev. B*, 38:9902–9905, 1988. doi: 10.1103/PhysRevB.38.9902.

J. Tersoff. Modeling solid-state chemistry: Interatomic potentials for multicomponent systems. *Phys. Rev. B*, 39:5566–5568, Mar 1989. doi: 10.1103/PhysRevB.39.5566.

V. Titarev and E. Toro. Finite-volume WENO schemes for three-dimensional conservation laws. *J. Comput. Phys.*, 201(1):238–260, 2004. ISSN 0021-9991. doi: 10.1016/j.jcp.2004.05.015.

E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer Science & Business Media, 2013. ISBN 9783662034903.

T. M. Tritt and M. A. Subramanian. Thermoelectric materials, phenomena, and applications: A bird's eye view. *MRS Bulletin*, 31:188–198, 3 2006. ISSN 1938-1425. doi: 10.1557/mrs2006.44. URL `http://journals.cambridge.org/article_S088376940000991X`.

N. Troullier and J. L. Martins. Efficient pseudopotentials for plane-wave calculations. *Phys. Rev. B*, 43:1993–2006, Jan 1991. doi: 10.1103/PhysRevB.43.1993.

S. J. Turneaure and Y. Gupta. Inelastic deformation and phase transformation of shock compressed silicon single crystals. *Appl. Phys. Lett.*, 91(20):201913, 2007a. doi: 10.1063/1.2814067.

S. J. Turneaure and Y. Gupta. X-ray diffraction and continuum measurements in silicon crystals shocked below the elastic limit. *Appl. Phys. Lett.*, 90(5):051905, 2007b. doi: 10.1063/1.2436638.

S. Van Huffel and J. Vandewalle. *The Total Least Squares Problem: Computational Aspects and Analysis*. Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics, 1991. ISBN 9780898712759.

D. Vanderbilt. Optimally smooth norm-conserving pseudopotentials. *Phys. Rev. B*, 32:8412–8415, Dec 1985. doi: 10.1103/PhysRevB.32.8412.

R. Vink, G. Barkema, W. van der Weg, and N. Mousseau. Fitting the Stillinger–Weber potential to amorphous silicon. *J. Non-Cryst. Solids*, 282(2-3):248–255, 2001. ISSN 0022-3093.

S. G. Volz and G. Chen. Molecular-dynamics simulation of thermal conductivity of silicon crystals. *Phys. Rev. B*, 61:2651–2656, Jan 2000. doi: 10.1103/PhysRevB.61.2651.

S. H. Vosko, L. Wilk, and M. Nusair. Accurate spin-dependent electron liquid correlation energies for local spin density calculations: a critical analysis. *Can. J. Phys.*, 58 (8):1200–1211, 1980. doi: 10.1139/p80-159.

A. Ward, D. A. Broido, D. A. Stewart, and G. Deinzer. Ab initio theory of the lattice thermal conductivity in diamond. *Phys. Rev. B*, 80:125203, Sep 2009. doi: 10.1103/PhysRevB.80.125203.

E. Weinan and P. Ming. Cauchy–born rule and the stability of crystalline solids: static problems. *Archive for Rational Mechanics and Analysis*, 183(2):241–297, 2007. doi: 10.1007/s00205-006-0031-7.

Y. Zhang and W. Yang. Comment on "Generalized Gradient Approximation Made Simple". *Phys. Rev. Lett.*, 80:890–890, Jan 1998. doi: 10.1103/PhysRevLett.80.890.