Reinhard Pippan
Peter Gumbsch

Editors

Multiscale Modelling of Plasticity and Fracture by Means of Dislocation Mechanics

CISM Courses and Lectures, vol. 522

SpringerWienNewYork

CISM COURSES AND LECTURES

The series presents lecture notes, monographs, edited works and proceedings in the field of Mechanics, Engineering, Computer Science and Applied Mathematics.
Purpose of the series is to make known in the international scientific and technical community results obtained in some of the activities organized by CISM, the International Centre for Mechanical Sciences.

INTERNATIONAL CENTRE FOR MECHANICAL SCIENCES

COURSES AND LECTURES - No. 522

CISM

MULTISCALE MODELLING OF PLASTICITY AND FRACTURE BY MEANS OF DISLOCATION MECHANICS

EDITED BY

REINHARD PIPPAN
AUSTRIAN ACADEMY OF SCIENCES, LEOBEN, AUSTRIA

PETER GUMBSCH
FRAUNHOFER-INSTITUT FÜR WERKSTOFFMECHANIK, IWM,
FREIBURG, GERMANY

SpringerWienNewYork

This volume contains 187 illustrations

All contributions have been typeset by the authors.

PREFACE

The main purpose of the

CISM-BIANCHI SESSION ON MULTISCALE MODELLING OF
PLASTICITY AND FRACTURE BY MEANS OF DISLOCATION
MECHANICS, Udine, July 4-8, 2005

was the discussion of the latest state of the modelling of plasticity and fracture of crystalline materials on the nano- and micro-scale. The present volume offers the revised lecture notes of all who contributed the course and one contribution of the CISM-course on Multiscale Modelling and Design of New Materials, held at the same time in Udine. Plasticity and fracture are multi-scale problems. The course was devoted to regimes from atomistic to meso-scale. The central part was the discrete dislocation mechanisms and neighbouring fields above and below molecular dynamics and crystal plasticity, respectively. The contributions give an introduction to the physical phenomena, the theoretical basis, the mathematical description and the different types of simulations. Different important practical problems such as plasticity of thin films, mechanics of polycrystalline materials, fatigue crack propagation, fracture of brittle and semi-brittle materials, formation of dislocation structures, etc. were considered as examples. All contributions start with text book like introductions, describing the different simulation techniques and discussing the latest state of possibilities and limitation of the different methods. A further important part in all contributions is the interconnection of the different hierarchical levels. The contributions in the present volume describe the different phenomena following the scale from the bottom up.

The first contribution is focused to molecular dynamics methods. It takes into account the theoretical background and gives practical examples to demonstrate its capabilities and limitations. The main topic is devoted to the fundamentals of the fracture processes. The second contribution is devoted to basic dislocation properties in an elastic continuum. The basics for accurate dislocation dynamic simulation of plasticity, the dislocation-dislocation and the dislocation self interaction is introduced in detail. The method is then applied to 2-dimensional and 3-dimensional problems. The third contribution gives an introduction to the linear elastic stress and a strain field of a

crack taking into account the presence of dislocations. Then the consequence of the discrete nature of plasticity is demonstrated with the example of a moderately cyclic loaded crack. The example is used to demonstrate the importance of the discrete dislocation consideration.

The fourth contribution deals also with discrete dislocation mechanics and its impact on fracture and fatigue crack propagation. Whereas the methods introduced in chapter three are limited to special geometries, in this part a method is introduced which permits a very general application of discrete dislocation mechanics. The fifth contribution deals with the formation of a dislocation pattern. It shows the newest results, how the collective behaviour of dislocations can be treated within a statistical physical framework

The sixth contribution gives an introduction to phenomenological crystal plasticity models, to their properties and their use either in component calculations or in microstructurally modelling. The last contribution, presented in the Multiscale Modelling and Design of New Materials-course is addressed to the key components of the computational homogenization scheme, i.e. the formulation of the microstructural boundary value problem and the coupling between the different levels based on the averaging theorems. With the detailed analyses of the different scales and bridging we do hope that the presented volume will contribute useful source of information on the fundamentals and the present-day techniques to model plastic deformation and fracture on the nano- and micro-scale.

It is the pleasure of the editors to express our gratitude to the Scientific Council of CISM for supporting the course as well as permitting these publications. We thank all the authors for their contributions and the fruitful discussions with the participants.

Peter Gumbsch and Reinhard Pippan

CONTENTS

Atomistic Simulation Methods and their Application on Fracture

Bernhard Eidel [*], Alexander Hartmaier [*] and Peter Gumbsch [†]

[*] Interdisciplinary Centre for Advanced Materials Simulation (ICAMS),
Ruhr-Universität Bochum, Stiepeler Strasse 129, 44801 Bochum, Germany
[†] Institut für Zuverlässigkeit von Bauteilen und Systemen, Universität Karlsruhe
and Fraunhofer Institut für Werkstoffmechanik, Wöhlerstr. 11, 79108 Freiburg,
Germany

Abstract *The present work on the molecular dynamics method covers the theoretical background of the method and gives practical examples to demonstrate its capabilities and limitations. The work focusses on topics which reveal fundamental mechanisms associated with fracture processes. Moreover, promising hybrid methods based on a concurrent atomistic/continuum coupling are reviewed since they combine accuracy and efficiency in a most favorable manner.*

1 Introduction

For many engineering questions connected to the mechanical properties of materials, one can of course profitably apply continuum mechanical descriptions of materials behaviour. However, when it comes to describing small specimens or when material specific questions need to be addressed, it is usually indispensable to investigate defect properties. For plastic deformation, discrete dislocation simulations (see (Kubin et al., 1992; Devincre and Roberts, 1996; Deshpande et al., 2003; Weygand et al., 2002; Weygand and Gumbsch, 2005) or the article by V. Mohles in this collection) can be applied. In these simulations the interaction of dislocation with interfaces determines the influence of the microstructures. These discrete defect based methods, however, require governing laws for the individual defect properties and more detailed descriptions for the short range interaction of defects. Such properties are difficult to obtain experimentally and are therefore usually investigated by atomistic methods.

Therefore, atomistic modelling is essential in advancing our understanding of the mechanical properties of materials. This is most obvious for the investigation of fracture processes. Materials behaviour with respect to fracture is of course ultimately determined by events on the atomic scale.

In the case of brittle fracture this connection is obvious, since the crack in a perfectly brittle material must be atomically sharp at its tip. The crack moves by breaking individual bonds between atoms and can therefore be regarded as a macroscopic probe for the atomic bonding. The transition from perfectly brittle to ductile behaviour similarly relies on atomistic processes since the multiplication or nucleation of dislocations at the crack tip is an indispensable ingredient of any modelling of the brittle to ductile transition. Atomistically the nucleation or multiplication of dislocations at the crack tip is identified with bond shearing events, which compete with the bond breaking events at crack extension.

The interaction of multiple defects, like cracks and dislocations, is still not routinely studied even with the simplest atomistic simulation methods since such studies necessarily require models that significantly extend in all three dimensions and consequently require the handling of millions of atoms. Since the application of atomistic techniques has great potential (Li et al., 2003; Kassner et al., 2005) but is still not widespread in the investigation of the mechanical properties of materials, it is probably advisable to give a short overview on the applicable methods and then to provide a few examples of successful application of atomistic modelling techniques to explain experimentally observed phenomena.

In the following section we give a general overview of the different descriptions for the atomic bonding, because they are the basis for all atomistic modelling. The next section describes the molecular dynamics method in some detail, also including the boundary conditions and visualisation of defects. This part is intended to give the reader, who is not familiar with atomistic simulation methods a feeling for the versatility of this approach. The capabilities and limitations of atomistic methods are not just related to the available description of the atomic bonding but also to the handling of boundary conditions and the analysis of the results. The subsequent section 4 deals with concurrent multiscale methods, which couple atomistic and continuum descriptions in a rather seamless manner, hence enabling the description of crystalline solids with atomistic accuracy but at smaller computational costs. The Finite Element Atomistic Method (FEAt) and variants of the Quasicontinuum (QC) method are reviewed and compared.

After this methods-oriented part, in Section 5 the phenomenology of fracture is discussed in the light of atomistic modelling. First, brittle fracture is put in the spotlight, which is the domain of the atomistic methods, as we have seen above. Second, it is shown that the analysis of the plastic zone around a crack tip is now within the scope of large scale atomistic simulations.

It is noted here that this article can neither provide a practical guide

for the use of atomistic methods nor can it claim to give an exhaustive overview of all the work where these methods provide insight into atomistic processes of fracture. It is rather meant to be a first introduction into the topic, highlighting a few illustrative examples and showing the possibilities and limitations of the most frequently used methods. Interested readers should also consult recent overviews on the subject (Li et al., 2003; Kassner et al., 2005) for more details and for additional references.

In order to limit the references to a useful number, we restrict ourselves to citations concerning models and numerical simulations to the literature dealing with fracture, where adequate. Furthermore, if the same authors published a number of articles on similar topics in fracture, the most recent work is usually given here since it will often provide guidance to previous work.

2   Description of Interatomic Bonds

In this section the different atomistic methods that have been used to model fracture processes are briefly introduced. The scope of this section is to give a general introduction into the basic idea behind the atomistic, *i.e.* non-continuum, method being used to model fracture of materials, to introduce the essential terms, and to provide the basic literature for further reading. No attempt is made to provide cooking recipes enabling the reader to implement such methods solely based on this text.

Before presenting the different methods to describe the atomic interaction, it is worth mentioning that atomistic modelling may be applied with very different intentions. In some cases, the atomistic simulations are used as testing grounds for ideas about the behaviour or energetics of defects. It may then be perfectly justified to use the simplest generic form of the interaction model. Alternatively, the goal may be to make quantitative or semiquantitaive predictions about the properties of specific materials. In this case one has to resort to materials-specific models or even quantum mechanical *ab initio* methods.

2.1   Quantum Mechanics Based Descriptions of the Atomic Bonding

The most fundamental description of interatomic bonds that may have ionic, covalent, metallic or van-der-Waals character, or any mixture of these pure bond types, is given by explicitly dealing with all the electrons involved in the formation of the atomic bonds. The quantum mechanical description of a crystalline material is a classical problem of solid state physics and

appropriate tools have been developed there, which today are available for the use in many different application areas. In all these codes, the many-body problem of atomic nuclei and electrons must be solved for different geometries. Lattice periodicity is usually imposed. The solution is generally approached in three steps:

1. Within the Born-Oppenheimer approximation the motion of the atomic nuclei is (adiabatically) decoupled from the motion of the electrons. Since the mass and the inertia of the atomic nuclei is orders of magnitude larger than the mass of the electrons, the electronic quantum gas follows the motion of the nuclei almost instantaneously. Furthermore the electrons are always assumed to remain in their ground state with respect to the momentary position of the nuclei.

2. The quantum mechanical ground state of the inhomogeneous electron system in the system of the nuclei is determined then in an external electrostatic potential, given by the charge at the instantaneous position of the nuclei. In practical calculations, the many particle state of the interacting electrons is usually constructed from single particle states of non-interacting electrons in an effective (mean) field of all the other electrons. These single particle states are self consistently determined by iteratively solving the coupled set of single particle problems for all the electrons. The most successful approach to this problem has been *density functional theory (DFT)* (Hohenberg and Kohn, 1964; Kohn and Sham, 1965). Within the *local density approximation* (LDA) it is assumed that the exchange correlation of the electrons can be calculated based purely on the local electron density. This is sometimes not enough and generalised gradient approximations (GGA) have been introduced.

3. In the last step we now reintroduce the motion of the atomic nuclei. The ground state energy of the electron gas can be regarded as the potential energy for the nuclei in their configurational space. The atomic motion can then be simulated in this *adiabatic* potential. Local minima in this adiabatic potential reflect statically stable structures for the solid while saddle points are the static barriers for structural transformations. With the knowledge of the adiabatic potential, it is possible to determine kinetic, dynamic or statistical materials properties (e.g. phonon properties, transport properties, or phase equilibria) either by classically solving Newtons equations of motion, by minimising the total energy of a structure or by solving quantum mechanical problems for the nuclei (e.g. for zero point vibrations of light elements).

The main task of electron theory is to solve for the electronic ground state, the second step above, in a numerical way. There are several successful methods available to do so, based on plane waves, localised orbitals or a mixture of them. The Cambridge (Serial) Total Energy Package CASTEP, the Vienna Ab-initio Simulation Package VASP or the Mixed Basis Pseudo-Potential Code MBPP are standard packages for this task.

Since these quantum mechanical methods do not require any adjustable parameters and are therefore often termed *ab initio* or *first priciples* techniques. Their results provide highly accurate and predictive results on materials properties.

However, the methods are computationally very demanding and often require periodic structures, which are constructed by defining a unit cell comprising the atoms to be considered and repeating this unit cell periodically in all directions. This procedure is identical to the periodic boundary conditions that are used extensively in molecular dynamics simulations, as will be described below. However, the periodic cells in molecular dynamics simulations may contain several million atoms, while *ab initio* methods, due to their mathematical complexity are restricted to hundreds of atoms within the unit cells. This limitation is rather severe for applications in mechanical problems, because the generation of defects or the driving force on a defect, like the energy release rate during crack advance, all depend on the total elastic energy stored within the volume. Hence, if the volume under consideration is extremely small, the elastic strains in the volume have to be close to their theoretical limits in order to store sufficient elastic energy to drive the defects. Therefore, *ab initio* simulations of defect behaviour must always be carefully checked for size effects and artifacts caused by the restrictions on the effective volume. Notwithstanding these words of warning, it must be stated again that *ab initio* methods are the most fundamental methods available to describe behaviour of materials and thus possess the most predictive power. There are also some interesting and potentially fruitful ideas to combine *ab initio* methods with molecular dynamics simulations which will be described in section 4.6.

Before leaving this section it is worth mentioning that *ab initio* methods are not just applied to study some specific fracture problem, they also play an important role in the verification and the adjustment of simpler models of the atomic interaction. To develop material-specific (semi-) empirical interaction models, various adjustable parameters and sometimes even the functional form of the interaction model has to be chosen to reproduce available data. However, experiment will usually only provide data close to mechanical and thermodynamic equilibrium, whereas the interaction models will be used at grain boundaries or even in the highly strained region near

crack tips. Therefore it is advisable to adjust the parameters of the interaction models so that they perform well under these circumstances. Because of the lack of experimental data, *ab initio* calculations of large deformations or structural energy differences are often used instead. It therefore turned out to be very useful to build up extensive *ab initio* data bases for the development of simpler interaction models. While the first such data bases for aluminium (Ercolessi and Adams, 1994) or the intermetallic nickelaluminide B2-NiAl (Ludwig and Gumbsch, 1995) still constituted demanding calculations, such databases can today be generated routinely and very systematically.

2.2   Atomic Interaction Models, Potentials

While it is generally possible to use *ab initio* methods to study the dynamic evolution of an atomistic system, the computational burden is usually too high and (semi-)empirical potentials are therefore applied in molecular dynamics (MD) simulations. The atomic interaction as described by semi-empirical potentials always depends on the distance of interacting pairs of atoms and, for the more elaborate potentials, also on the bond angles and the local electron density. All these distance or angle dependent functions are represented either analytically or as tabulated functions of these parameters. The most simple pair potentials, like Lennard-Jones or the Morse potential, have just two or three free parameters that are used to change the characteristic properties of the potential like the lattice parameter, the bulk elastic modulus or the cohesive energy. More sophisticated potentials have a number of free parameters that yield much better results for material specific properties like lattice constant, sublimation energy, anisotropic elastic constants, vacancy formation energy, stacking fault energy, or whatever is felt necessary for a certain investigation. The free parameters of the potentials are usually fitted to material properties by calculating exactly these properties on rather small atomic ensembles and then varying the potential parameters until a reasonable match to experimental or *ab initio* data is reached. The interatomic potentials can roughly be categorized into three classes according to their level of approximation of the "real" quantum mechanical atomic interaction:

- angularly-dependent tight-binding or bond-order potentials
- multibody or embedded atom method (EAM) potentials
- pair potentials

The properties and characteristics of the different potentials are given in some detail below. For application purposes, the choice of a potential will usually involve a trade-off between the accuracy of the description of the atomic interaction from first principles calculations over semi-empirical po-

tentials to simple pair potentials and the size of the system that can be studied, *i.e.* the number of atoms that can be explicitly treated in the sample. The decision, whether higher accuracy in the description of the interatomic forces or larger system size, has to be made for each problem under consideration. For example, during fracture of covalently bonded systems like silicon, polymers or biological matter the behaviour of the entire system depends critically on the behaviour of the individual atomic bonds immediately in front of the crack tip (see Section 5.1). Here it becomes necessary to accurately evaluate the interatomic forces in order to obtain meaningful results. In contrast, for rather generic investigations of dynamic crack stability as a function of crack driving force, where the energy release rate and the stored elastic energy in the solid are decisive, it may be essential to have a reasonably large system (Gumbsch et al., 1997; Buehler et al., 2003). Even in these cases, however, materials-specific questions must again be treated with the more accurate DFT descriptions or coupled methods (Kermode et al., 2008).

Pair potentials. Pair potentials are the simplest form of interatomic potentials. This approach limits the interaction between two atoms to a dependence on their mutual distance, thus excluding completely any information about neighbouring atoms. Pair potentials in their most general form are written as

$$U_{\text{tot}} = \frac{1}{2} \sum_{i,j \neq i} V_{ij}(r_{ij}) + U(\Omega) \tag{1}$$

where V_{ij} is the pair potential, r_{ij} is the separation between atoms i and j and $U(\Omega)$ is a contribution to the cohesive energy which depends on the average volume per atom Ω. This latter term may mainly determine the cohesive energy, while the pair potential always determines the structural dependence.

Often, however, it is not attempted to mimic the properties of a particular material but more generic atomistic questions are of interest. Then, potentials like the Lennard-Jones potential or the Morse potential are applied. These have only few adjustable parameters, which are fitted to the nearest neighbour distance and the binding energy or the bulk modulus. Despite their simplicity and the lack of materials specificity, pair potentials can contribute significantly to our understanding of material behaviour and fracture in particular. Their simple structure permits investigation of generic effects, e.g. of the atomic size or of the role of elastic properties for a certain phenomenon. Other examples are when the consequences of the discrete crystal structure of solids shall be compared to predictions of

continuum models or when the influence of non-linearities in the elastic be-
haviour shall be mentioned that can ideally be compared with linear-elastic
solids described by a harmonic pair potential (Buehler et al., 2003). In these
examples use is made of the simplicity of pair potentials that allows clean
studies of single effects, while excluding all other influences that render the
behaviour of "real" materials so much more complex.

Of course, pair potentials cannot be used to describe situations where
directional bonds or bond angles play a role. They are also not well suited
to describe metals because the well established dependence of bond strength
on coordination cannot be represented. Furthermore, simple pair potentials
result in extremely low stacking fault energies, because they fail to distin-
guish the energy difference between face centred cubic (fcc) and hexagonal
close-packed (hcp) structures.

Embedded atom method (EAM) potentials. The most widely used
interaction models for metals are the EAM potentials. In this scheme the
volume-dependent term from Equation (1) is expressed as a local density-
dependent contribution to the total potential energy

$$U_{\text{tot}} = \frac{1}{2} \sum_{i, j \neq i} V_{ij}(r_{ij}) + \sum_i F(\rho_i) \tag{2}$$

with

$$\rho_i = \sum_{j \neq i} \rho_{ij}(r_{ij}) \tag{3}$$

where ρ_{ij} can be viewed as the contribution from atom j to the total elec-
tron density at atom i and F is the embedding energy associated with
placing atom i in this environment. Finally, V_{ij} is the pair potential contri-
bution to the potential energy of atom i. Different physical interpretations
of the terms are possible and consequently the functional forms and the
way in which the various parameters are determined may differ. Usually
the functions are adjusted to reproduce at least the lattice parameter, the
cohesive energy and the anisotropic elastic constants of the metal under
consideration. For alloys it is additionally necessary to adjust structural
energy differences of intermetallic phases and heats of solution (Ludwig and
Gumbsch, 1995). As mentioned above, it is also desirable to not only adjust
to such equilibrium crystal properties, but to also compare to some *ab initio*
data from atomic structures far from equilibrium or for large deformations.
Reliable and well tested potentials are available for the noble metals and
nickel as well as for some aluminium alloys (e.g. (Ercolessi and Adams,
1994; Ludwig and Gumbsch, 1995; Mishin et al., 2001)).

The main advantage of EAM potentials is that they can be applied to inhomogeneous systems such as surfaces or cracks, since these potentials approximately incorporate the variation of bond strength with coordination. Decreasing the coordination (density) at an atom usually increases the strength of each of its bonds and decreases the bond length. In particular for fcc materials it has been easily possible to reproduce anisotropic elastic constants, sublimation energy, vacancy formation energy, surface and stacking fault energies with EAM potentials and EAM potentials can be applied to many different types of simulations. For body centred cubic (bcc) metals the method may still give a sufficiently precise description. For some of these metals, however, it has been shown that the EAM fails to give a valid description of the core structure of dislocations (Mrovec et al., 2004) and thus the Peierls stress to move the dislocation through the lattice. Therefore, these potentials cannot be used reliably to model plasticity in these materials. This may be seen as an indication of the importance of angular bonding characteristics and one may consequently resort to tight binding or bond order descriptions (Pettifor and Oleinik, 1999; Mrovec et al., 2004, 2007a).

Tight binding and bond order potentials.   The need for material specific models and for accuracy in atomistic simulations implies that a quantum mechanical description of the atomic interaction is required. The semi-empirical tight binding scheme is such an approximate quantum mechanical description of the energetics of systems of atoms (Harrison, 1980; Pettifor, 1995). The total energy of a system is given by:

$$U_{\text{tot}} = U_{\text{rep}} + U_{\text{bond}} \qquad (4)$$

where U_{rep} is a repulsive energy, generally given as a sum of pair potentials and U_{bond} is the bonding part of the energy. The latter is obtained by solving the eigenvalue problem for a given Hamiltonian, which is assumed to be fixed and not evaluated self-consistently as in the *ab initio* methods.

The Hamiltonian matrix elements are usually assumed to be rapidly decaying functions of the atomic separation, which have to be empirically adjusted to experimental data or results of *ab initio* calculations to give a material-specific model. The least clear but most important question in this context is the transferability of these matrix elements from ideal structures to highly distorted atomic environments.

The semi-empirical tight-binding scheme is in principle applicable to materials with various types of bonding but is most naturally suited to covalently bonded insulators and semiconductors as well as the transition metals and intermetallics .

The key computational aspect in solving a tight-binding model is the diagonalisation of the eigenvalue problem. This can of course be attempted in reciprocal space, where a whole set of very efficient methods is available, which however usually scale to the third power with the number of atoms in the system. Linear scaling can be achieved with real space methods of which the density matrix method for semiconductors and insulators and the bond order approach for metals and alloys are the most promising approaches. In the simplest approximations the latter can even be formulated analytically within the framework of bond order potentials (Pettifor and Oleinik, 1999) where the energy of the system is written as

$$U_{\text{tot}} = \frac{1}{2} \sum_{i,j \neq i} V_{ij}(r_{ij}) + \sum_{i\alpha,j\beta} H_{i\alpha,j\beta}\Theta_{i\alpha,j\beta} \tag{5}$$

where V_{ij} is the pair potential, depending on the type of atom and r_{ij} is the separation between atoms i and j. H is the Slater Koster hopping integral, which also depends on distance r_{ij} and where the Greek indices stand for the type of orbital. Θ is the bond-order matrix which gives the difference in the number of electrons in the bonding and antibonding states. Modern bond order potentials for carbon (Mrovec et al., 2007b) and some transition metals (Mrovec et al., 2004, 2007a) are just becoming available.

In the same spirit as the bond-order potentials, several simplified angularly-dependent potentials of Tersoff and Brenner-type have been developed for silicon and the hydrocarbons (see for example (Pastewka et al., 2008)). These potentials have seen continuous improvements but were notoriously difficult to adjust to bond breaking problems. Despite recent success with an explicit screening formulation for the hydrocarbons (Pastewka et al., 2008) giving up specific bonds in a way which does not require a self-consistent formulation of the electronic configuration remains an important and difficult problem.

3   The Molecular Dynamics Method

To continue we describe the pertinent methods used in atomistic simulations with some focus on the application in fracture processes and related problems. We start out in the present section with the Molecular Dynamics (MD) method. A short overview of the main ingredients like time integration and relaxation is given, furthermore the role of boundary and initial conditions is described and illustrated for mechanical problem sets in materials science. Section 3.6 gives a brief account of the visualisation and analysis of defects in atomistic simulations, which plays an ever more important role in computational materials science. Some popular and powerful

defect indicators are described. In Section 4 we analyse and compare two concurrent multiscale methods.

The main ingredients of an MD simulation are basically threefold:

(i) A model is needed for the interaction between the system constituents (e.g. atoms or molecules).

(iia) Time integration is required to advance the particle trajectories (positions and velocities) from time t to $t + \Delta t$.

(iib) Alternatively one may want to solve a stability problem which in an atomistic system requires an algorithm to relax the atomic coordinates to positions of vanishing forces.

(iii) An ensemble has to be chosen, for which boundary conditions and thermodynamic quantities like temperature, pressure or the number of particles are controlled.

3.1 Force Calculation

Forces are derived from the potential energy U that depends on the positions of all atoms. The description for the calculation of the energy can be based on different physical approximations as described above. The force acting on an atom i is given by taking the derivative of the potential energy with respect to the position vector $\boldsymbol{x}_i$ of atom i

$$\boldsymbol{f}_i = -\frac{dU(\boldsymbol{x})}{d\boldsymbol{x}_i}\,, \tag{6}$$

where $\boldsymbol{x}$ denotes the coordinates of all atoms. Once the force vector $\boldsymbol{f}_i$ acting on all atoms is known, the Newtonian equation of motion

$$\boldsymbol{f}_i = m_i \frac{d^2\boldsymbol{x}_i}{dt^2}\,, \tag{7}$$

can be integrated in time t to yield the motion of the atoms in space. The mass of the atom is given by m_i.

3.2 Integrating the Equations of Motion

Equation (7) constitutes a set of second-order ordinary differetial equations (ODEs), which can be strongly nonlinear. By converting them to first-order ODEs in the $6N$-dimensional space of $\{\boldsymbol{x}_N, \dot{\boldsymbol{x}}_N\}$, general numerical algorithms for solving ODEs such as the Runge-Kutta method could be applied. However, these general methods are rarely used in practice,

because the existence of a Hamiltonian allows for much simpler and even more accurate integration algorithms.

To represent other thermodynamic ensembles than the *micro-canonical* ensemble for which Equation (7) can be integrated directly, requires that Equation (7) is modified to create a dynamics in phase space that has the desired distribution density of e.g. a *canonical* or a *grand-canonical* ensemble (see e.g. (Frenkel and Smit, 2002)). The time-average of a single-point operator on such a trajectory then approaches the thermodynamic average.

An integrator serves the purpose of propagating particle positions and velocities over small time increments Δt.

$$\boldsymbol{x}^{3N}(t_0) \to \boldsymbol{x}^{3N}(t_0 + \Delta t) \to ... \to \boldsymbol{x}^{3N}(t_0 + L\,\Delta t) \tag{8}$$

The time step Δt has to be chosen such that the thermal oscillations of the atoms around their equilibrium positions are resolved in time. A typical frequency of this oscillation is the Debye frequency $\nu_D = c_t/a$, where c_t is the speed of transverse sound waves and a is the lattice parameter. A typical value for metals is $\nu_D \approx 10^{13}$ Hz. This implies that the typical time step for MD simulations has to be on the order of femtoseconds ($=10^{-15}$ s), which generally limits the method to simulations of fast processes such as brittle fracture or high-strain-rate plastic deformation.

Some popular time integration algorithms are the central difference algorithms: Verlet, velocity Verlet or leap-frog. They are shortly introduced below.

• **Verlet algorithm.** Assuming that the $\boldsymbol{x}^{3N}(t)$ trajectories are smooth, one may perform a third-order Taylor expansion of the positions $\boldsymbol{x}_i(t_0)$ forward ($\boldsymbol{x}_i(t_0 + \Delta t)$) and backward ($\boldsymbol{x}_i(t_0 - \Delta t)$) in time; their sum yields

$$\boldsymbol{x}_i(t_0 + \Delta t) + \boldsymbol{x}_i(t_0 - \Delta t) = 2\boldsymbol{x}_i(t_0) + \ddot{\boldsymbol{x}}_i(t_0)(\Delta t)^2 + O((\Delta t)^4). \tag{9}$$

Since $\ddot{\boldsymbol{x}}_i(t_0) = \boldsymbol{f}_i(t_0)/m_i$ can be evaluated given the atomic positions at $t = t_0$, $\boldsymbol{x}^{3N}(t + \Delta t)$ in turn may be approximated by,

$$\boldsymbol{x}_i(t_0 + \Delta t) = -\boldsymbol{x}_i(t_0 - \Delta t) + 2\boldsymbol{x}_i(t_0) + \frac{1}{m}\boldsymbol{f}_i(t_0)(\Delta t)^2 + O((\Delta t)^4). \tag{10}$$

Neglecting the $O((\Delta t)^4)$ term, we obtain a recursion formula to compute $\boldsymbol{x}^{3N}(t_0 + \Delta t)$. Although velocities are not needed in the recursion, they are often calculated since they are required for analysis of ensemble properties. They can be approximated by

$$\boldsymbol{v}_i(t_0) \equiv \dot{\boldsymbol{x}}_i(t_0) = \frac{1}{2\Delta t}[\boldsymbol{x}_i(t_0 + \Delta t) - \boldsymbol{x}_i(t_0 - \Delta t)] + O((\Delta t)^2). \tag{11}$$

This algorithm is not only one of the simplest, but also a good choice in general. It is fast, but not particularly accurate for long time steps, such that the forces on all particles must be computed rather frequently. It requires about as little memory as is at all possible. This is useful when very large systems are simulated. Verlet's short-term energy conservation is fair but, more important, it exhibits little long-term energy drift. This is related to the fact that the Verlet algorithm is time reversible and area preserving. In fact, although the Verlet algorithm does not conserve the total energy of this system exactly, strong evidence indicates that it does conserve a pseudo-Hamiltonian approaching the true Hamiltonian in the limit of infinitely short time steps.

● **Velocity-Verlet algorithm.** It starts with $\boldsymbol{v}^{3N}(t_0)$ and $\boldsymbol{x}^{3N}(t_0)$. One then evaluates

$$\boldsymbol{x}_i(t_0 + \Delta t) = \boldsymbol{x}_i(t_0) + \boldsymbol{v}_i(t_0)\Delta t + \frac{1}{2}\frac{\boldsymbol{f}_i(t_0)}{m_i}(\Delta t)^2 + O((\Delta t)^3), \qquad (12)$$

with $\boldsymbol{f}^{3N}(t_0 + \Delta t)$ evaluated from $\boldsymbol{x}_i(t_0 + \Delta t)$ one gets

$$\boldsymbol{v}_i(t_0 + \Delta t) = \boldsymbol{v}_i(t_0) + \frac{1}{2}\left[\frac{1}{m_i}\boldsymbol{f}_i(t_0) + \frac{1}{m_i}\boldsymbol{f}_i(t_0 + \Delta t)\right]\Delta t + O((\Delta t)^3), \ (13)$$

and has advanced by one step. This algorithm requires a little more computing but is very popular since it gives $\boldsymbol{x}^{3N}$ and $\boldsymbol{v}^{3N}$ simultaneously.

● **Leap-frog algorithm.** In the leap-frog algorithm, position and velocities are calculated with the same accuracies but are offset by $\Delta t/2$. It starts with $\boldsymbol{v}^{3N}(t_0 - \Delta t/2)$ and $\boldsymbol{x}^{3N}(t_0)$. Time integration is then first done on $\boldsymbol{v}$

$$\boldsymbol{v}_i(t_0 + \frac{1}{2}\Delta t) = \boldsymbol{v}_i(t_0 - \frac{1}{2}\Delta t) + \frac{1}{m_i}\boldsymbol{f}(t_0)\Delta t + O((\Delta t)^3), \qquad (14)$$

followed by integration of $\boldsymbol{x}$,

$$\boldsymbol{x}_i(t_0 + \Delta t) = \boldsymbol{x}_i(t_0) + \boldsymbol{v}_i(t_0 + \frac{1}{2}\Delta t) + O((\Delta t)^3). \qquad (15)$$

It can be shown that the leap-frog algorithm produces identical trajectories to the Verlet algorithm besides numerical rounding errors. It therefore has similar properties than the Verlet algorithm but of course provides coordinates and velocities at once.

For the description of *Predictor-Corrector Algorithms* and *Symplectic Integrators* we refer to standard text-books (Frenkel and Smit, 2002; Schlick, 2002; Rapaport, 2004).

3.3 Relaxation Algorithms

Mechanically stable configurations or thermally activated processes are extremely difficult if not impossible to assess with MD methods. Examples for such difficult-to-study processes in the context of fracture would be void formation and coalescence or motion of dislocations in lattices with high Peierls barriers (Marian et al., 2004). For these cases it is usually advisable to map out the energy landscape using relaxation algorithms and then to perform a metadynamics on the basis of such an energy landscape.

Mechanically stable equilibrium configurations correspond to minima of the total energy. Hence it is a most common task in computational materials science and solid state physics to find local or global minima of the potential energy, where a given initial configuration is the point of departure. To solve this task a variety of well-established optimisation methods are available (see e.g. in (Nocedal and Wright, 2006; Leach, 2001; Schlick, 2002)). The optimisation methods generally can be classified according to the highest order derivative used to minimise the (energy) functional. Hence, a non-derivative minimisation method like the simplex method can be considered as zeroth-order method. Zeroth-order methods are rarely used in molecular modelling since first derivatives of the energy (i.e. forces) are usually available and these methods then do not exploit all the available information. In first-order minimisation methods, the gradient of the energy indicates the direction to a minimum, its magnitude measures the steepness of the local slope. These methods are frequently used in molecular modelling. Prominent examples are the steepest descent method and variants of the conjugate gradient (CG) method. For a nice overview we refer to (Shewchuk, 1994). Most recent MD integration-based methods that also fall into this category are described below. Second order methods additionally use the second derivatives and thus the information of the local curvature to locate a minimum. Current state-of-the-art methods like the limited-memory version of the Broyden-Fletcher-Goldfarb-Shanno scheme (l-BFGS) (Nocedal and Wright, 2006) explicitly use only first order derivatives but accumulate information to obtain some approximate representation for the Hessian matrix to determine line search directions. Second order methods are by far not as robust as first order methods and can therefore only be used in molecular modelling if no structural changes occur.

MD-based methods, FIRE. Within the realm of molecular simulations, some methods have been proposed that serve the purpose of energy minimisation, that start out with MD and proceed in removing kinetic energy from the system, termed 'quenching'. This strategy has been successfully

applied in local minimisation as well as in global minimisation. In (Bitzek et al., 2006) a simple, yet powerful MD scheme for structural relaxation was proposed which belongs to this family of minimisers. Different from existing schemes this new algorithm crucially relies on inertia, it has therefore been named the Fast Intertial Relaxation Engine (FIRE). In the original paper the method's functional principle was explained by means of a blind skier searching for the fastest way to the bottom of a valley in an unknown mountain range described by the potential energy landscape $E(\boldsymbol{x})$ with $\boldsymbol{x} = (x_1, x_2)$. Assuming that the skier is able to retard and steer, the recommended strategy for the skier is to follow an equation of motion given by:

$$\dot{\boldsymbol{v}}(t) = 1/m\, \boldsymbol{F}(t) - \gamma(t)|\boldsymbol{v}(t)|[\hat{\boldsymbol{v}}(t) - \hat{\boldsymbol{F}}(t)], \qquad (16)$$

with the mass m, the velocity $\boldsymbol{v} = \dot{\boldsymbol{x}}$, the force $\boldsymbol{F} = -\nabla E(\boldsymbol{x})$, and hat denoting unit vector. The recommended strategy is that the skier introduces acceleration in a direction that is 'steeper' than the current direction of motion via the function $\gamma(t)$, if the power $P(t) = \boldsymbol{F}(t) \cdot \boldsymbol{v}(t)$ is positive, and in order to avoid uphill motion he simply stops as soon as the power becomes negative. $\gamma(t)$ must be chosen appropriately but should not be too large, because the current velocities carry information about the reasonable 'average' descent direction and energy scale. A discretised version of this equation in combination with an adaptive time step results in a minimisation scheme for multidimensional functions $E(x_1; \ldots x_M)$ which is competitive in speed with the fastest optimisers currently available (Bitzek et al., 2006), but has also other important features as we shall demonstrate.

The numerical treatment of the algorithm is simple. Any MD integrator can be used as the basis for propagation of the trajectories due to the conservative forces. The MD trajectories are continuously readjusted by two kinds of velocity modifications: (a) the above-mentioned immediate stop upon uphill motion and (b) a simple mixing of the global velocity and force vectors $\boldsymbol{v} \to (1-\alpha)\boldsymbol{v} + \alpha\hat{\boldsymbol{F}}|\boldsymbol{v}|$ resulting from an Euler-discretisation of the last term in Equation (16) with time step Δt and $\alpha = \alpha_{\text{start}}$. Both, Δt and α, are chosen adaptively on the fly.

The propagation rules for the FIRE algorithm can be summarised as follows (given: initial values for Δt, $\alpha = \alpha_{\text{start}}$ and the global vectors $\boldsymbol{x}$ and $\boldsymbol{v} = \boldsymbol{0}$):

1. MD integrator: calculate $\boldsymbol{x}$, $\boldsymbol{F} = -\nabla E(\boldsymbol{x})$ and $\boldsymbol{v}$ using any common MD integrator; check for convergence.

2. calculate $P = \boldsymbol{F} \cdot \boldsymbol{v}$.

3. set $\boldsymbol{v} \to (1-\alpha)\boldsymbol{v} + \alpha|\boldsymbol{v}|\hat{\boldsymbol{F}}$.

4. if $P > 0$ and the number of steps since P was negative is larger than $N_{\min}$, increase the time step $\Delta t \to \min(\Delta t f_{\mathrm{inc}}, \Delta t_{\max})$ and decrease $\alpha \to \alpha f_\alpha$.

5. if $P \leq 0$, decrease time step $\Delta t \to \Delta t f_{\mathrm{dec}}$, freeze the system $\boldsymbol{v} \to \boldsymbol{0}$, and set α back to α_{start}.

6. Return to MD.

In relaxation an accurate calculation of the atomic trajectories is not necessary, and the adaptive time step allows FIRE to increase Δt until either the largest stable time step $\Delta t_{\max}$ is reached, or an energy minimum along the current direction of motion ($P < 0$) is encountered. In the latter case the system is instantly frozen ($\boldsymbol{v} \to \boldsymbol{0}$) and the time step is substantially reduced in order to have a smooth restart. A short 'latency' time of $N_{\min}$ MD steps before accelerating the dynamics is important for the stability of the algorithm.

In (Bitzek et al., 2006) it is shown in several benchmark sets that FIRE is very competitive with sophisticated algorithms like the l-BFGS scheme and easily beats even advanced versions of the CG method (e.g. Polak-Ribière). The key advantage of the FIRE algorithm, however, is its extreme robustness. It finds the (local) minima even for extreme structural rearrangements like molecular folding and rotation or for atomic reconstructions in the core of crystalline defects and therefore lends itself ideally as a general purpose minimiser.

3.4   Boundary and initial conditions

Initial conditions.   The integration of Newton's equations of motion requires an integrator and initial conditions (IC), namely $\boldsymbol{x}^{3N}(t = 0)$ and $\dot{\boldsymbol{x}}^{3N}(t = 0)$, the initial particle positions and velocities.

Generating $\boldsymbol{x}^{3N}(t = 0)$ for crystalline solids is easily done by a structure generator setting up a perfect crystal or an interface between two crystalline phases. Generating suitable velocity distributions or structures for a liquid or an amorphous solid is significantly more difficult. It can, however, often be circumvented by running the system from an artificial set-up for an extended equilibration time. One can for example melt a crystal and obtain the IC for an amorphous configuration then by quenching from the liquid.

Boundary conditions.   Boundary conditions can be classified into two major types: periodic boundary conditions (PBC) and isolated (or free) boundary conditions (IBC). IBC, in which surface atoms exhibit dangling bonds due to a lack of neighbours, are chosen for the analysis of surfaces, clusters and molecules. In addition, there can be extra forces or displace-

ments acting on the boundary atoms, resulting in so-called mechanical boundary conditions.

PBC are intended to mimic bulk solids or liquids, or the cores of very large systems, which are much larger than the simulated number of particles. The particles are contained within a primary simulation volume. This volume is sometimes referred to as the simulation box, unit cell or supercell.

When a particle leaves one side of this volume, it re-enters from the opposite side keeping the number of atoms in the central box constant. Atoms sitting in the vicinity of one side of the box, through periodic repetition of the box, are connected to the atoms on the other side of the box. The simulations therefore proceed as if the primary volume was surrounded in all directions by identical copies of itself to form a quasi-infinite volume. PBC for the two-dimensional case are illustrated in Figure 1. If PBC are used, the case that the cutoff radius is less than half the diameter of the periodic box is of special interest since in that case only the interaction of a given atom with the nearest periodic image of any other atom needs to be considered (*minimum image convention*). This case is displayed in Figure 1, where the dotted box comprises all nearest images of any other atom.
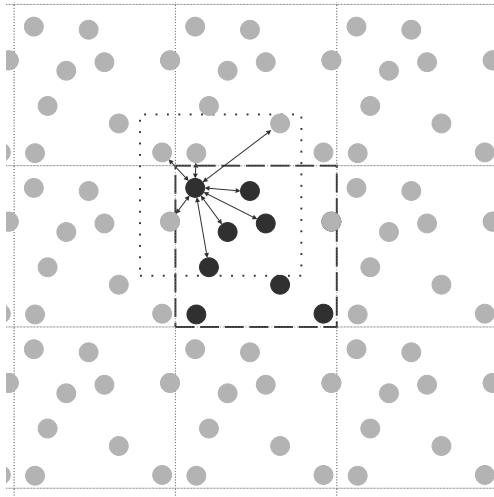

Figure 1. Periodic boundary conditions: the basic simulation volume (drawn in *black*) is repeated in all dimensions. Here, particles in neighbour volumes are drawn in *gray*. Interactions of one atom with its neighbours are indicated with arrows.

The cubic cell is the simplest geometry of an unit cell to visualise and

to program. Of course, any other shape can be chosen provided it fills all of space by translation operations. Five shapes satisfy this condition: the cube (and its close relation, the parallelepiped), the hexagonal prism, the truncated octahedron, the rhombic dodecahedron and the elongated dodecahedron, see e.g. (Leach, 2001).

Using PBC, surface effects can be eliminated from the simulation, but some inherent limitations dictate rules in applying them, see (Pöschel and Schwager, 2005).

(i) Since each particle or defect has an infinite number of images, there is always an infinite number of interactions with each other particle or defect.

(ii) If the chosen primary volume is too small, there appear correlations between opposite edges of the primary volume. If the spatial structures are of the same characteristic size as the system itself, a particle may interact directly or indirectly with itself across the primary volume. It is therefore not possible to achieve fluctuations that have a wavelength greater than the length of the cell.

The first item causes problems when long-range forces, such as electrostatic interactions or elastic defect-defect interactions are involved. Sophisticated techniques have been developed, such as Ewald-summation to simulate systems of charged or gravitating particles. For long-range electrostatic interactions it is frequently accepted that some long-range order will be imposed upon the system. Hence, the range of the interactions between particles in the system must generally be taken into account while planning a simulation. No problems will arise if the periodic box is large compared with the range over which the atomic interaction (e.g. of an EAM potential) acts. A box size greater than three times the cut-off radius of the potential is always sufficient. Specific measures have to be taken in the calculation of the forces if shorter box lengths shall be realized.

The second complication listed above is more substantial, since it is inherent to the system and cannot be solved by improved algorithms. Consequently, the basic volume size has to be chosen large enough to avoid undesired artificial correlations. Size scaling studies have to be performed to assess such effects. In favourable cases, the long range interactions of periodic arrays of defects are known and can be subtracted from the total energy of the unit cell.

Boundary conditions that mix IBC and PBC can be required in case that the simulated structure exhibits periodicity in some directions, but not in the others. Typical examples are one-dimensional structures like very long slabs or wires and systems where loads are applied in certain directions.

Mechanical loading through boundary conditions.   For the study of
the mechanical behaviour of solids like e.g. in fracture processes, a proper
definition of boundary conditions and their adequate application are nec-
essary. A simple approach to apply displacement boundary conditions is
realized via a domain decomposition, in that a domain is added at the
boundary with prescribed stress or displacements. In these domains, atoms
are not subject to the dynamics of Newton's equation of motion in an MD
simulation; instead, they follow prescribed displacements of the boundaries,
see e.g. (Buehler, 2008). Figure 2 schematically displays this approach
for the simulation of fracture. A crystalline slab is strained by means of
rigid, sufficiently thick boundary layers which are subject to prescribed dis-
placements, whereas the atoms in the interior follow Newton's equations of
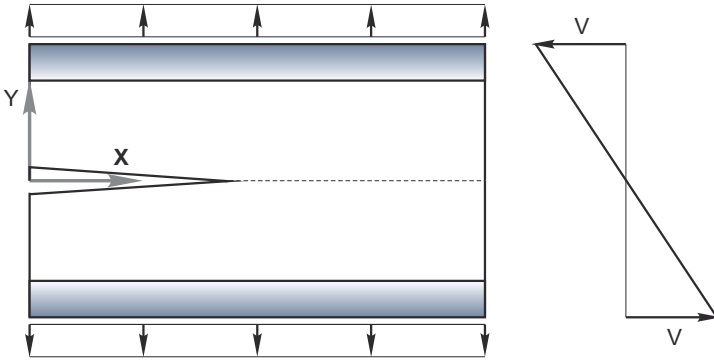motion.

Figure 2. Application of displacement boundary conditions: Atoms in
boundary domains (gray-shaded) follow prescribed boundary conditions and
are not subject to the equations of motion.

The application of pressure – or generally stress – instead of displace-
ment boundary conditions can be realized by utilizing appropriate ensem-
ble schemes such as the Parinello-Rahman scheme (Parinello and Rahman,
1980).

Since time steps in typical MD simulations are of the order of fem-
toseconds to keep track of thermal vibrations, the velocity of applied dis-
placement boundary conditions during dynamic straining is necessarily very
high. To apply increased mechanical loading, a linear velocity gradient can
be established prior to simulation to avoid shock wave generation from the
boundaries, see Figure 2, right. However, studying processes at low load-
ing rates is difficult. Sometimes, stress boundary conditions may allow to

simulate somewhat slower loading.

For well controlled loading situations, however, one has to resort to the application of constant displacement boundary conditions and study the system as it evolves in a given field. In the example of Figure 2, the application of a constant displacement u provides a constant (energetic) driving force for the propagation of the crack. With such boundary conditions one can reach the desired case of studying crack propagation at constant energy release rate G.

Since a propagating crack is expected to generate heat, a local temperature control (Finnis et al., 1991), resembling an electronic heat bath for the ions, is applied there. At the outer border of the model the coupling (damping) parameter is gradually increased to prevent reflections from the borders. The model is then first equilibrated at an applied strain corresponding to the Griffith load $G_0 = 2\gamma$. During the equilibration time the crack remains stationary. Thereafter the model is instantaneously strained further to a defined overload ΔG by scaling all displacements. For small overloads $\Delta G = 0.03 - 0.10\, G_0$ a short acceleration phase can be detected immediately after loading. This suggests that the crack has finite but very small inertia. After this acceleration phase the crack runs at constant velocity.

3.5   Stable Defects under Load

For the investigation of discrete defects like an individual crack or an individual dislocation one wants to supply the defect with the natural stress and strain fields it also would experience in a realistically large system. This is of course not possible with the rigid straight boundary conditions sketched in Figure 2. Instead a finite size model must be supplied in the border region with the forces or displacements determined from the continuum mechanical stress and strain fields of the defect. For the case of a sharp crack, the model is loaded by first applying the anisotropic linear elastic continuum solution (Sih and Liebowitz, 1968) for a fixed value of the stress intensity factor K to all atoms in the model. As a starting value for K the stress intensity factor K_G (Griffith load) is used. Then all atoms are relaxed to their equilibrium positions by a relaxation algorithm, except for the outermost atoms, which are held fixed at their initial positions. This configuration is then used for further incremental loading or unloading of the model, which is achieved by scaling all displacements and relaxing the whole model at each incremental step. Changes in the atomic configurations at the crack tip are determined by visual inspection of the relaxed configurations. The load at which one or more bonds are broken upon loading is taken as the upper critical stress intensity K_c^+. (See (Kohlhoff et al., 1991; Gumbsch,

1995) for details about the loading procedure, size scaling tests and the comparison to other methods.) Often periodic boundary conditions are applied along the crack front.

3.6   Visualisation and Analysis of Defects

A key issue in large-scale atomistic simulations is the automatic identification and visualisation of defects and microstructures.

For that purpose algorithms have been developed to identify, highlight and classify typical defect types, see e.g. (Li, 2007). Hence, pictures or movies based on these indicators realize the interface between experiments in the virtual laboratory and the scientist. They enable a transformation of vast data sets to structured information. Since visualisation criteria allow the judicious selection of regions of interest, they also serve a similar purpose as concurrent multiscale methods, namely to compress or reduce large data sets. In the following some of these visualisation criteria are put forward.

Energy method.   One way to extract crystal defects from their undisturbed neighbourhood is the energy method. In this method, atoms with an energy value larger than a defined threshold – or within a specified interval – are targeted for display, which exploits the fact that defects exhibit high-energy. The energy method has been successfully applied to visualise microcracks, dislocations, nanovoids and the like. For nanoindentation into (001) fcc aluminium the energy method selects dislocation loops slipping on (111) planes as visualised in Figure 3. Moreover it can be seen that the energy method accounts for the energy difference between free surfaces and bulk material. In order to avoid free surfaces preventing analysis, a subregion of interest is chosen for visualisation. A shortcoming of the energy-method is that it cannot be applied at elevated temperatures. Furthermore, crystalline defects like stacking faults are hard to see using the energy method. In these cases it is favorable to use defect-indicators which take into account the structure of the crystal lattice and its symmetries.

Centrosymmetry parameter.   The centrosymmetry parameter (CP) as introduced by (Kelchner et al., 1998) is defined for an fcc atom according to

$$P = \sum_{i=1}^{6} |\boldsymbol{r}_i + \boldsymbol{r}_{-i}|^2 \; , \qquad (17)$$

where vectors $\boldsymbol{r}_i$ and $\boldsymbol{r}_{-i}$ correspond to the six pairs of next neighbours lying at opposite sites with respect to the considered atom in the lattice. By definition, the CP is zero for an atom in the bulk of a perfect material subject
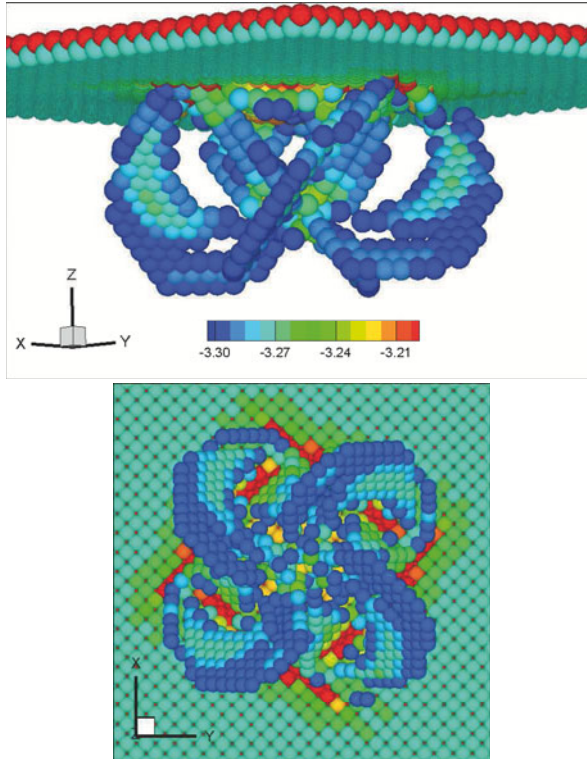
Figure 3. Simulation of nanoindentation into (001) fcc Al at zero temperature: the internal energy criterion extracts the free surface and the dislocation loops.

to purely homogeneous elastic deformations. The deviation of P from zero therefore measures the strength of disturbed centrosymmetry at a lattice site. Opposed to the Energy Method, the CP enables identification and classification of defects like free surfaces, partial dislocations and stacking faults by a certain number.

Since the CP – opposed to the Energy-Method – is related to the crystal's structure, it is invariant to thermal fluctuations and hence also applicable at finite temperatures.

Slip vector analysis. The slip vector analysis (SVA) was proposed by (Zimmerman et al., 2001) in the context of MD simulation of nanoindenta-
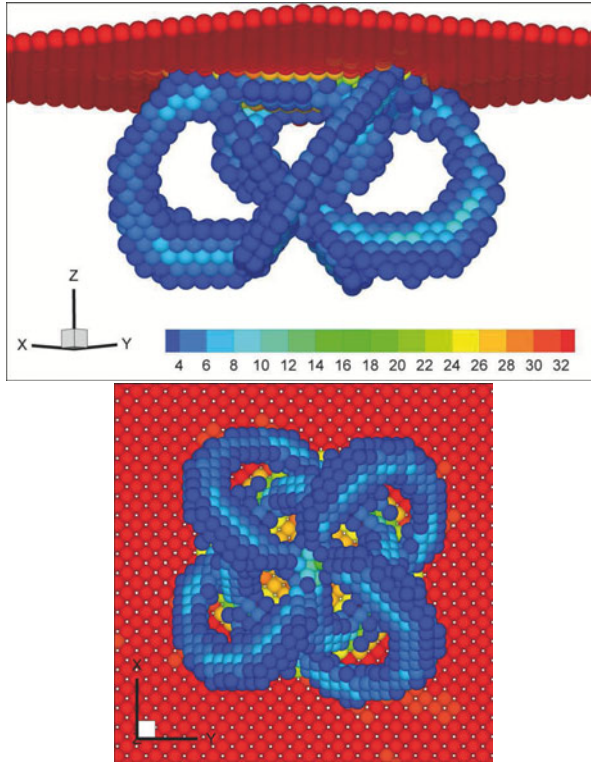
Figure 4. Simulation of nanoindentation into (001) fcc Al: the centrosymmetry parameter extracts the free surface and the dislocation loops.

tion; the slip vector of atom α is defined as

$$s^{\alpha} = -\frac{1}{n_s} \sum_{\beta \neq \alpha}^{n_{\alpha}} \left\{ \boldsymbol{r}^{\alpha\beta} - \boldsymbol{R}^{\alpha\beta} \right\} \; , \tag{18}$$

where n_s is the number of slipped neighbours, n_{α} is the number of nearest neighbours, $\boldsymbol{r}^{\alpha\beta}$ and $\boldsymbol{R}^{\alpha\beta}$ are the vector differences between atom α and atom β in the current and the reference configuration, respectively.

The particular advantage of the SVA over CP is, that Burgers vectors are directly accessible and thus, SVA realizes the bridge to crystallography. Furthermore, the slip vector approach can be applied to very different types of defects in crystalline solids, whereas CP measures only disturbed centrosymmetry.
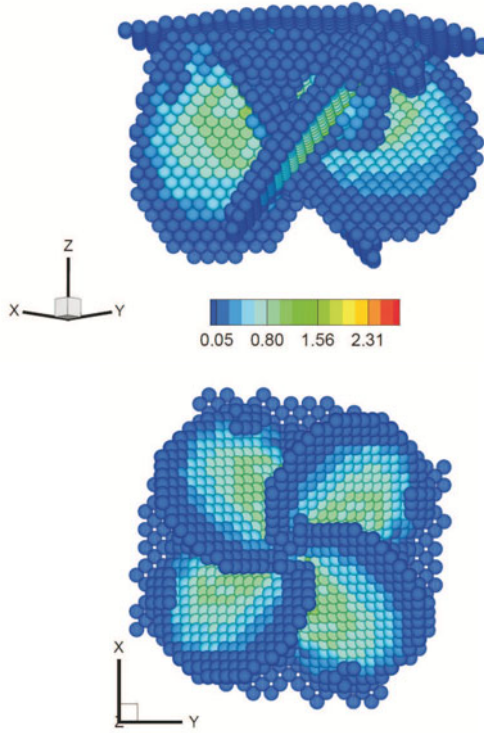
Figure 5. Simulation of nanoindentation into (001) fcc Al: The slip vector analysis, L_2 norm of $\boldsymbol{s}^\alpha$, extracts the dislocation loops.

Other criteria to identify and visualise defects in crystalline solids are the common neighbour analysis (Honeycutt and Andersen, 1987) and the Bond Angle Analysis (Ackland and Jones, 2006). The bond angle analysis is particularly useful to determine the local coordination, distinguishing fcc, hcp, bcc, and other relatively close-packed structures.

4   Concurrent Multiscale Methods

4.1   Introduction and Classification of Multiscale Methods

Despite the considerable achievements by means of molecular simulations, which have been sped up by the ever more increasing computer power, the range of applicability of atomistic models and methods is still rather lim-

ited in that they have not yet reached the typical time and length scales of engineering applications. The reason is that quite disparate time and length scales have to be considered; for MD the maximum time step is dictated by the frequency of thermal vibrations, hence in the order of femtoseconds, whereas a process like e.g. crack propagation may occur in the order of seconds. The spatial problem is not less demanding, since the length scale at the bottom is in the range of atomic spacings, hence of nanometres, whereas the world of engineering problems lives in the range of some centimetres - and beyond. For that reason many efforts have been undertaken to overcome the time scale and length scale dilemma by *coarse-graining* approaches. The accurate, at best *seamless* information passing from a bottom scale to a corresponding coarse-grained scale –and eventually backwards– is one of the key challenges in computational materials science.

In the following we will consider only the coupling of length scales and not time scales, and doing this, we will moreover restrict to the atomistic-to-continuum scale-coupling at zero temperature. Multiscale methods can be generally separated into two main categories, *hierarchical/sequential* and *concurrent*.

- *Hierarchical or sequential multiscale methods.* In this concept, material information on the atomic scale is generated and passed to a larger length scale. A simple example of this concept is the atomistic calculation of material parameters like elastic constants, thermal expansion coefficients, hardening moduli which are then used as input in continuum constitutive equations. The embedded atom method (EAM) itself can be seen as an example of a hierarchical multiscale method, since the parameters of EAM are determined in ab-initio calculations. Another instance is the identification of parameters in traction-separation laws of atomic debonding, which are employed in cohesive zone finite elements.

  This concept of information-passing from a small-scale model to a larger scale model is relatively simple and cheap. Since the material's small-scale response is parametrised for a fixed set of tests to feed the larger-scale model, the larger-scale model cannot account for the full complexity of the material's behaviour in situations far from the test bed, where the model parameters were calibrated. In order to overcome these limitations a second branch of multiscale methods has emerged.

- *Concurrent multiscale methods.* Opposed to the aforementioned class of methods, concurrent multiscale methods explicitly couple models on different length scales. Corresponding computational frameworks allow the running of simulations on disparate length scales in parallel,

i.e. *concurrently*. For that aim the simulation domain is decomposed into different regions, where fully atomistic resolution is retained in critical regions, where deformation strongly varies and where inelastic deformations occur. In regions with weakly varying, purely elastic deformations, continuum constitutive laws are typically employed along with a *coarse-graining* as realized by the finite element method. The overall goal is to achieve a result as accurate as necessary and as effective as possible.

Concurrent multiscale methods are typically used for problem sets where inelastic deformations localise in regions of confined size. These regions typically are embedded into other regions, which deform elastically and which form the largest portion of the entire simulation domain. Since crack propagation is a prominent example in this class, research in the mechanics of fracture has driven the development of concurrent multiscale models and methods, and vice versa.

A number of review articles for concurrent multiscale methods coupling different length scales is available, see e.g. (Ortiz and Phillips, 1999; Miller and Tadmor, 2002; Curtin and Miller, 2003; Miller and Tadmor, 2009), and with a special focus on fracture processes, (Abraham et al., 2000).

In the following we describe two prominent concurrent multiscale methods, the Finite Element Atomistic Method (FEAt) and the Quasicontinuum (QC) Method.

4.2   The Finite Element Atomistic (FEAt) Method

The Finite Element Atomistic (FEAt) Method as introduced in (Kohlhoff, 1990; Kohlhoff et al., 1991) is based on the decomposition of a crystal into generally three different domains as visualised in Figure 6:

A : a lattice region with fully atomistic resolution, where interatomic potentials are employed.

C : a continuum region, discretized by finite elements, where a local, but nonlinear continuum constitutive law of anisotropic elasticity is employed.

T : a transition region in between A and C, where the coupling between the local continuum and the nonlocal lattice is mediated. In this region the framework of Kröner's nonlocal elasticity theory is used. It is a continuum theory, which takes the finite range of internal forces into account and therefore can be seen as a continuation of the lattice, (Kröner, 1963).
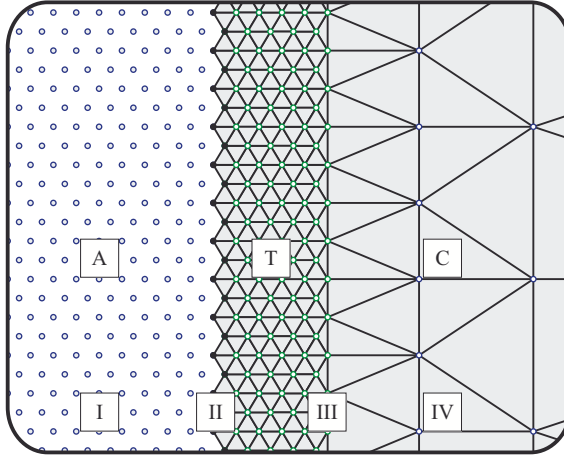
Figure 6. Scale transition scheme in FEAt.

Coupling conditions. The coupling between the atomistic and continuum regions in FEAt is in the first place based on the (i) strong kinematical compatibility and on (ii) stress/force compatibility.

Kinematical coupling. The kinematical compatibility in the transition T region is realized as follows; the transition region where the lattice and the continuum overlap is divided into two zones, II and III, Figure 6. Each zone provides the displacement boundary conditions for the other zone. Zone II, which reduces to a surface in a three-dimensional model, supplies the boundary conditions for the continuum, region IV. The FE nodes on this surface coincide with the atoms of the lattice and move with them. Conversely, in zone III the atoms, which constitute the outer shell of the lattice, are made to move in accord with the FE nodes with which they coincide. A one-to-one correspondence of nodes and atoms throughout the transition region is obtained in this way. This is why this coupling is generally termed *strong compatibility* in the classification of domain decompositions in multiscale models.

The width of zone III must be at least equal to the cut-off length of the potentials used to describe the atomistic core region. Note that the transition zone III must be $2R_{cut}$ thick when the underlying atomistic model is of the EAM type, and thus atomic forces (derivatives of energy) depend on the electron-density at an atom and at an atom's neighbour R_{cut} away, the latter of which depends on a neighbour's neighbour, up to another R_{cut} away.

Force coupling.   The design of the model ensures equality of the displacement fields in the lattice and the continuum throughout the transition region, and thus equality of strains at the interface. However, since the use of forces has been explicitly avoided in this coupling scheme, any direct interaction between the stress fields of the two media is prevented. This means that force equilibrium or equality of stress between the lattice and the continuum is not established a priori. Additional conditions are necessary. In FEAt, the coupling condition is on the force level, whereas other concurrent multiscale methods perform the coupling on the energy level. To define these coupling conditions in a consistent fashion, the elastic energy $E = \hat{E}(\varepsilon)$ is expanded into a Taylor series about the state of zero strain under the assumption of zero stress,

$$E(\varepsilon) = E(0) + \left.\frac{\partial E}{\partial \varepsilon_{ij}}\right|_0 + \frac{1}{2}\left.\frac{\partial^2 E}{\partial \varepsilon_{ij}\partial \varepsilon_{kl}}\right|_0 \varepsilon_{ij}\varepsilon_{kl} + \frac{1}{6}\left.\frac{\partial^3 E}{\partial \varepsilon_{ij}\partial \varepsilon_{kl}\partial \varepsilon_{mn}}\right|_0 \varepsilon_{ij}\varepsilon_{kl}\varepsilon_{mn} + \ldots .$$
(19)

For the equality of stresses, the strains and all coefficients in this series must be equal in the atomistic region and in the continuum. Since the strains have been made equal by means of the strong compatibility of displacements, equality of the stresses amounts to the requirement, that the elastic constants of the continuum equal those defined by the interatomic potential in the atomistic region:

$$\mathbb{C}_{ij} = \left.\frac{\partial E}{\partial \varepsilon_{ij}}\right|_0 , \quad \mathbb{C}_{ijkl} = \left.\frac{\partial^2 E}{\partial \varepsilon_{ij}\partial \varepsilon_{kl}}\right|_0 , \quad \mathbb{C}_{ijklmn} = \left.\frac{\partial^3 E}{\partial \varepsilon_{ij}\partial \varepsilon_{kl}\partial \varepsilon_{mn}}\right|_0 . \quad (20)$$

Since the reference state of the series in Equation (19) is a homogeneous deformation, $\varepsilon = 0$, it is indeed permissible to assign to the continuum the second- and higher order elastic constants as defined by Equation (20) and which are derived from the interatomic potential. The first-order elastic constants in the continuum are zero by definition, which imposes the restriction on the potential that it must provide zero stress in a perfect lattice. Within the framework of *local* and *linear* elasticity theory, equilibrium between the lattice and the FE continuum is fulfilled, if terms up to the second order are matched. FEAt accounts for elastic nonlinearity in that, additionally, elastic constants of third order, are adapted as well.

Two approximations in this coupling scheme are made. First, that the series in Equation (19) has to be cut off at some stage. Second, that there is a transition at the discrete interface from interatomic and hence finite-range forces to continuum Cauchy-type stresses. The latter approximation introduces a discontinuity in the nonlocal part of the stress tensor whenever

the strain gradient at the interface does not vanish. As a consequence, the magnitudes of the strain and of the strain gradient at the atomic-continuum influences the quality of the approximation in the model and the applicability to various situations.

Since in the transition region, II and III, finite element mesh nodes and lattice sites as well as their degrees of freedom coincide by virtue of strong compatibility, dispensable degrees of freedom are elimininated by condensation before the solution process. In (Kohlhoff, 1990) it was observed, that the global finite element stiffness matrix is not symmetric, which indicates that the forces are not conservative and hence that the governing equations are not derived from a variational principle. This is due to the fact, that the FEAt model does not start out from a well-defined total energy for the entire coupled problem but rather effects the coupling between the atomistic and continuum domains on the force/stress level as described above.

The fundamental difference between force coupling and energy coupling, its implications and consequences, are analyzed in Section 4.4 where we compare two variants of the QC method.

Application to fracture simulation. The method was used in (Kohlhoff et al., 1991) to analyse crack propagation on cleavage and non-cleavage planes in bcc crystals, using potentials for iron and tungsten as examples. The results explain why both, the {100} and {110} planes, are cleavage planes in bcc metals and why cleavage on {100} is easier than on the close-packed {110} planes.

4.3   The Quasicontinuum-Method Based on the Cauchy-Born Rule

The quasicontinuum method is a prominent example of a bottom-up, concurrent multiscale method aiming at a seamless link of atomistic with continuum length scales. This aim is achieved by three main building blocks which are common to each of the existing QC-versions, the QC based on Cauchy-Born elasticity, (Tadmor et al., 1996), (Shenoy et al., 1999), and two variants of a fully nonlocal QC method, (Knap and Ortiz, 2001), (Eidel and Stukowski, 2009):

  (i) a coarse-graining of fully atomistic resolution via kinematic constraints in order to reduce the number of degrees of freedom. Full atomistic resolution is retained where necessary.
 (ii) an approximation of the energy/forces in coarse-grained regions via numerical quadrature which avoids the explicit computation of the site energy of all the atoms.
(iii) adaptivity, i.e. spatially adaptive resolution, is necessary to automatically balance accuracy and efficiency. It must be directed by a suitable
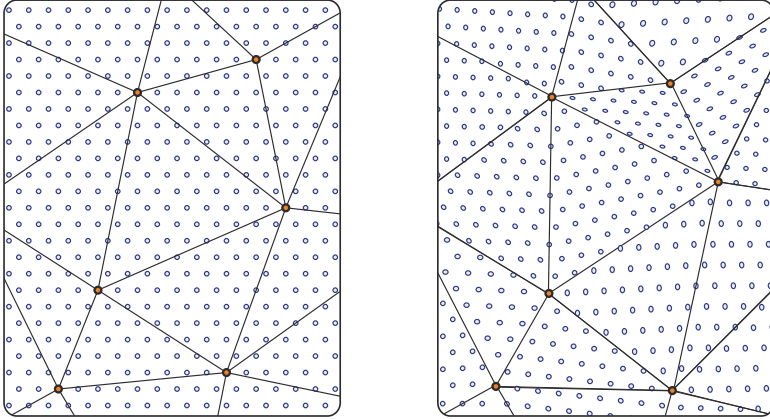
refinement indicator.

Figure 7. Finite-element discretisation of a crystal in the QC-method in the (left) undeformed configuration and (right) in the deformed configuration. Atoms within elements smoothly follow the mesh nodes by linear interpolation.

To set the stage, we consider a crystal in d-dimensional space consisting of a set $\mathscr{L} \subset \mathbb{Z}^d$ of atoms, that are initially located on a Bravais lattice spanned by lattice vectors $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_d$. Their coordinates in the initial configuration read $\boldsymbol{X}_{\boldsymbol{l}} = \sum_{i=1}^{d} l^{(i)} \boldsymbol{A}_i$, $\boldsymbol{l} \in \mathscr{L} \subset \mathbb{Z}^d$. The corresponding atomic coordinates in the current configuration are denoted by vector $\boldsymbol{x}_{\boldsymbol{l}}$.

Upscaling via coarse-graining. In regions of weakly varying elastic deformation it is sufficient to consider the movement of some judiciously selected, *representative* atoms (rep-atoms), $\mathscr{L}_h \subset \mathscr{L}$. Only these atoms keep their independent degrees of freedom, whereas all other atoms, $\overline{\mathscr{L}_h} = \mathscr{L} \backslash \mathscr{L}_h$, are forced to follow via kinematic constraints borrowed from the finite element method: $\boldsymbol{x}_{\boldsymbol{l}} = \sum_{\boldsymbol{j} \in \mathscr{L}_h} \boldsymbol{x}_{\boldsymbol{j}} \varphi_{\boldsymbol{j}}(\boldsymbol{X}_{\boldsymbol{l}})$, $\boldsymbol{l} \in \overline{\mathscr{L}_h}$. FE shape functions, $\varphi_{\boldsymbol{j} \in \mathscr{L}_h}$, exhibit the properties, $\sum_{\boldsymbol{j} \in \mathscr{L}_h} \varphi_{\boldsymbol{j}}(\boldsymbol{X}_{\boldsymbol{i}}) = 1 \; \forall \; \boldsymbol{i} \in \mathscr{L}$ (partition of unity), and, $\varphi_{\boldsymbol{j}}(\boldsymbol{X}_{\boldsymbol{j}}') = \delta_{\boldsymbol{j}\boldsymbol{j}'} \; \forall \; \boldsymbol{j}, \boldsymbol{j}' \in \mathscr{L}_h$ (compact support). The use of (here: linear) shape functions for interpolation requires the generation of a triangulation with representative atoms as mesh nodes. Figure 7 schematically displays the discretisation of the crystal into finite elements.

The interpolation of nodal displacements implicitly introduces a continuum assumption into the QC method. Notwithstanding, this first approx-

imation is purely kinematical in nature, since no constitutive assumptions are made.

The approximation step of discretisation or coarse-graining reduces the number of arguments in the exact total potential, $E^{\mathrm{tot}}(\{\boldsymbol{x_i}|\ \boldsymbol{i} \in \mathscr{L}\}) \to E^{\mathrm{tot}}(\{\boldsymbol{x_i}\mid \boldsymbol{i} \in \mathscr{L}_h\}) =: E^{\mathrm{tot},h}$, and thus reduces the number of unknowns in the computation. Both existing QC methods have this approximation step in common, but differ in the way further approximations are made.

Next, we focus on the QC version based on Cauchy-Born elasticity, in Section 4.4 the fully nonlocal QC versions are described and compared.

Efficient energy/force calculation: the local QC. After thinning-out dispensable degrees of freedom via the kinematic constraints in terms of linear finite-element shape functions, the first QC-version as proposed by (Tadmor et al., 1996) accomplishes an efficient energy/force calculation in the continuum region by recourse to the so-called *Cauchy-Born (CB) Rule* -hence QC-CBR- resulting in what is referred to as the *local* formulation of the QC.

The CBR postulates that when a monatomic crystal is subjected to a small linear displacement of its boundary, then all atoms will follow this displacement, see (Born and Huang, 1998), (Ericksen, 1983), (Zanzotto, 1996). The CBR is schematically illustrated in Figure 8 for a crystalline cantilever undergoing elastic bending deformation.
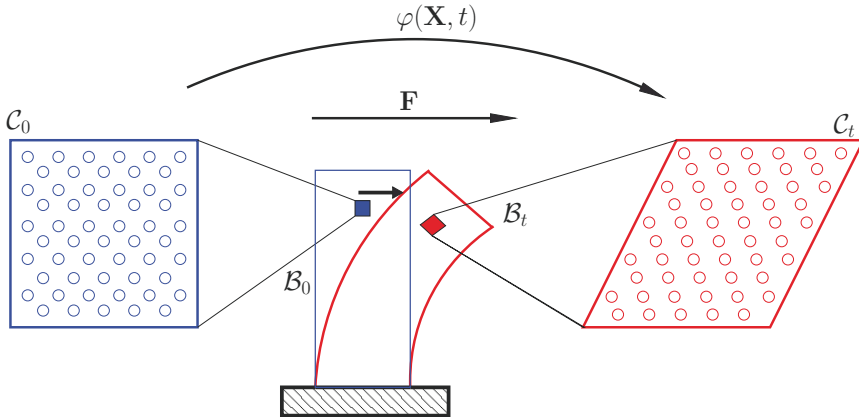

Figure 8. The Cauchy-Born-rule assuming a homogeneous deformation state in small representative volumes.

The CBR is applied in the QC method in that the continuum defor-

mation gradient as a macro-scale quantity is directly mapped to a *uniform* deformation of a small volume on the atomistic or nano-scale. For crystalline solids with a simple lattice structure the assumption of locally homogeneous deformation state implies that every atom in a region subject to a uniform deformation gradient will be energetically equivalent. As a consequence, calculating the energy within a specific finite element can be approximated by computing the energy of only one single atom in the deformed state and multiplying this figure by the number of atoms in the specific finite element. Within the QC-computational framework, the calculation of the CB energy is done separately in a subroutine; for a given deformation gradient $\boldsymbol{F}$ the lattice vectors in a unit cell with PBC is deformed according to $\boldsymbol{F}$

$$\boldsymbol{a}_i = \boldsymbol{F}\boldsymbol{A}_i \,, \tag{21}$$

where $\boldsymbol{A}_i$ and $\boldsymbol{a}_i$ are the lattice vectors in the undeformed configuration and in the deformed configuration, respectively.

The deformed lattice vectors enter the employed potential for energy calculation, such that the CBR enables the free energy of a deformed crystalline body (as a function of lattice vectors) to be expressed alternatively as a function of the deformation gradient $\boldsymbol{F}$. The corresponding strain energy density in the element is then given by

$$\mathcal{E} = \frac{E_0(\boldsymbol{F})}{\Omega_0} \,, \tag{22}$$

where Ω_0 is the unit cell volume (in the reference configuration) and E_0 is the energy of the unit cell when its lattice vectors are distorted according to $\boldsymbol{F}$. Now the total energy of a finite element is this energy density times the element volume, the total energy of the problem is simply the sum of all element energies:

$$E^{tot,h} = \sum_{i=1}^{N_{element}} \Omega_i \mathcal{E}(\boldsymbol{F}_i) \,, \tag{23}$$

where Ω_i is the volume of element i.

Linear interpolation functions in tetrahedral finite elements require only one single Gauss-point for numerical quadrature and therefore imply a constant deformation gradient per element as visualised in the right of Figure 7. Note, that the locally constant deformation gradient in the finite element matches the assumption of a locally constant deformation gradient in the CBR. As a consequence, the application of the CBR implies that in the energy calculation the summation over the number of lattice sites boils down

to the number of finite elements $N_{element}$, see Equation (23). Since the crystal is in general subject to inhomogeneous deformations, the element-wise constant deformation gradient is an approximation and so is the calculated energy via the CBR. In settings where the deformation is varying slowly and the element size is adequate with respect to the variations of the deformation, this type of energy calculation is sufficiently accurate and very effective.

For a mathematical analysis on the range of validity of the CBR we refer to (Friesecke and Theil, 2002), where it is found that the CBR fails for relatively small elastic deformations. An extension of the classical linear CBR to high order is proposed in (Sunyk and Steinmann, 2003).

Nonlocal QC. In nonlocal regions, which can be eventually refined to fully atomistic resolution, the energy E_α of an atom residing on a mesh node α is calculated by numerical quadrature. Specifically the new approximate energy takes the form

$$E^{tot,h} = \sum_{\alpha=1}^{N_{rep}} n_\alpha E_\alpha(\boldsymbol{u}_h) \,, \tag{24}$$

where $\boldsymbol{u}_h$ represents the finite element nodal displacements. The computational saving is that the summation of all the atoms is replaced by a sum over all representative atoms N_{rep}. In the line of numerical quadrature, n_α is the weight function for repatom α which requires for consistency

$$\sum_{\alpha=1}^{N_{rep}} n_\alpha = N \,. \tag{25}$$

Hence, n_α is the number of atoms represented by atom α, which implies in the limiting case of fully atomistic resolution $n_\alpha = 1$.

Mixed local-nonlocal QC. In order to combine the high accuracy of the nonlocal formulation with the efficiency of the local formulation, the former is employed in critical regions, where atomic scale accuracy is required, where the latter formulation is employed in regions where the deformation is changing relatively slowly on the atomic scale. As a result, the QC-CBR runs both formulations concurrently in a single simulation.

As in the energy-based nonlocal QC, the coupled approach is based on the ansatz that the energy can be approximated by computing only the energy of the repatoms. In the coupled approach however, each repatom is judiciously selected as being either local or nonlocal depending on its