

2.13

THE THEORY AND IMPLEMENTATION OF THE QUASICONTINUUM METHOD

E.B. Tadmor¹ and R.E. Miller²

¹*Technion–Israel Institute of Technology, Haifa, Israel*

²*Carleton University, Ottawa, ON, Canada*

While atomistic simulations have provided great insight into the basic mechanisms of processes like plasticity, diffusion and phase transformations in solids, there is an important limitation to these methods. Specifically, the large number of atoms in any realistic macroscopic structure is typically much too large for direct simulation. Consider that the current benchmark for large-scale fully atomistic simulations is on the order of 10^9 atoms, using massively paralleled computer facilities with hundreds or thousands of CPUs. This represents 1/10 000 of the number of atoms in a typical grain of aluminum, and 1/1 000 000 of the atoms in a typical micro-electro-mechanical systems (MEMS) device. Further, it is apparent that with such a large number of atoms, substantial regions of a problem of interest are essentially behaving *like a continuum*. Clearly, while fully atomistic calculations are essential to our understanding of the basic “unit” mechanisms of deformation, they will never replace continuum models altogether.

The goal for many researchers, then, has been to develop techniques that retain a largely continuum mechanics framework, but impart on that framework enough atomistic information to be relevant to modeling a problem of interest. In many examples, this means that a certain, relatively small, fraction of a problem require full atomistic detail while the rest can be modeled using the assumptions of continuum mechanics.

The quasicontinuum method (QC) has been developed as a framework for such mixed atomistic/continuum modeling. The QC philosophy is to consider the atomistic description as the “exact” model of material behaviour, but at the same time acknowledge that the sheer number of atoms make most problems intractable in a fully atomistic framework. Then, the QC uses continuum assumptions to reduce the degrees of freedom and computational demand without losing atomistic detail in regions where it is required.

The purpose of this article is to provide an overview of the theoretical underpinnings of the QC method, and to shed light on practical issues involved in its implementation. The focus of the article will be on the specific implementation of the QC method as put forward in Refs. [1–4]. Variations on this implementation, enhancements, and details of specific applications will not be presented. For the interested reader, these additional topics can be found in several QC review articles [5–8] and of course in the original references. The most recent of the QC reviews [5] provides an extensive literature survey, detailing many different implementations, extensions and applications of the QC. Also included in that review are several other coupled methods that are either direct descendants of the QC or are similar alternatives developed independently. For a detailed comparison between several coupled atomistic/continuum methods including the QC, the reader may find the review by Curtin and Miller [9] of interest.

A QC website designed to serve as a clearinghouse for information on the QC method has been established at www.qcmethod.com. The site includes information on QC research, links to researchers, downloadable QC code and documentation. The downloadable code is freely available and corresponds to the QC implementation discussed in this paper.

1. Atomistic Modeling of Crystalline Solids

In the QC, the point-of-view which is adopted is that there is an underlying atomistic model of the material which is the “correct” description of the material behaviour. This could, in principle, be a quantum-mechanically based description such as density functional theory (DFT), but in practice the focus has been primarily on atomistic models based on semi-empirical interatomic potentials. A review of such methods can be found, for example, in [10]. Here, we present only the features of such models which are essential for our discussion. We focus on *lattice statics* solutions, i.e., we are looking for equilibrium atomic configurations for a given model geometry and externally imposed forces or displacements, because most applications of the QC have used a static implementation. Recent work to extend QC to finite temperature and dynamic simulations shows promise, and can be found in Ref. [11].

We assume that there is some reference configuration of N atomic nuclei, confined to a lattice. Thus, the reference position of the i th atom in the model \mathbf{X}_i is found from an integer combination of lattice vectors and a reference (origin) atom position, \mathbf{X}_0

$$\mathbf{X}_i = \mathbf{X}_0 + l_i \mathbf{A}_1 + m_i \mathbf{A}_2 + n_i \mathbf{A}_3, \quad (1)$$

where (l_i, m_i, n_i) are integers, \mathbf{A}_j is the j th Bravais lattice vector.¹

The deformed position of the i th atom \mathbf{x}_i , is then found from a unique displacement vector \mathbf{u}_i for each atom.

$$\mathbf{x}_i = \mathbf{X}_i + \mathbf{u}_i. \quad (2)$$

The displacements \mathbf{u}_i , while only having physical meaning on the atomic sites, can be treated as a continuous field $\mathbf{u}(\mathbf{X})$ throughout the body with the property that $\mathbf{u}(\mathbf{X}_i) \equiv \mathbf{u}_i$. This approach, while not the conventional one in atomistic models, is useful in effecting the connection to continuum mechanics. Note that for brevity we will often refer to the field \mathbf{u} to represent the set of all atomic displacements $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N\}$ where N is the number of atoms in the body.

In standard lattice statics approaches using semi-empirical potentials, there is a well defined total energy function E^{tot} that is determined from the relative positions of all the atoms in the problem. In many semi-empirical models, this energy can be written as a sum over the energy of each individual atom. Specifically,

$$E^{\text{tot}} = \sum_{i=1}^N E_i(\mathbf{u}), \quad (3)$$

where E_i is the site energy of atom i , which depends on the displacements \mathbf{u} through the relative positions of all the atoms in the deformed configuration. For example, within the embedded atom method (EAM) [13, 14] atomistic model, this site energy is given by

$$E_i = U_i(\bar{\rho}_i) + \frac{1}{2} \sum_{j \neq i} V_{ij}(r_{ij}), \quad (4)$$

where U_i can be interpreted as an electron-density dependent embedding energy, V_{ij} is a pair potential between atom i and its neighbor j and $r_{ij} = \sqrt{(\mathbf{x}_i - \mathbf{x}_j) \cdot (\mathbf{x}_i - \mathbf{x}_j)}$ is the interatomic distance. The electron density at the position of atom i , $\bar{\rho}_i$, is the superposition of spherically averaged density contributions from each of the neighbors, ρ_j :

$$\bar{\rho}_i = \sum_{j \neq i} \rho_j(r_{ij}). \quad (5)$$

A similar site energy can be identified for other empirical atomistic models, such as those of the Stillinger–Weber type [15], for instance.

¹We omit a discussion of complex lattices with more than one atom at each Bravais lattice site. This topic is discussed in Refs. [5, 12].

In addition to the potential energy of the atoms, there may be energy due to external loads applied to atoms. Thus, the total potential energy of the system (atoms plus external loads) can be written as

$$\Phi(\mathbf{u}) = E^{\text{tot}}(\mathbf{u}) - \sum_{i=1}^N \mathbf{f}_i \mathbf{u}_i, \quad (6)$$

where $-\mathbf{f}_i \mathbf{u}_i$ is the potential energy of the applied load \mathbf{f}_i on atom i . In lattice statics, we seek the displacements \mathbf{u} such that this potential energy is minimized.

2. The QC Method

The goal of the static QC method is to find the atomic displacements that minimize Eq. (6) by approximating the total energy of Eq. (3) such that:

1. the number of degrees of freedom is substantially reduced from $3N$, but the full atomistic description is retained in certain “critical” regions,
2. the computation of the energy in Eq. (3) is accurately approximated without the need to explicitly compute the site energy of all the atoms,
3. the fully atomistic, critical regions can evolve with the deformation, during the simulation.

In this section, the details of how the QC achieves each of these goals are presented.

2.1. Removing Degrees of Freedom

A key measure of a displacement field is the deformation gradient \mathbf{F} . A body deforms from reference state \mathbf{X} to deformed state $\mathbf{x} = \mathbf{X} + \mathbf{u}(\mathbf{X})$, from which we define

$$\mathbf{F}(\mathbf{X}) \equiv \frac{\partial \mathbf{x}}{\partial \mathbf{X}} = \mathbf{I} + \frac{\partial \mathbf{u}}{\partial \mathbf{X}}, \quad (7)$$

where \mathbf{I} is the identity tensor. If the deformation gradient changes gradually on the atomic scale, then it is not necessary to explicitly track the displacement of every atom in the region. Instead, the displacements of a small fraction of the atoms (called representative atoms or “repatoms”) can be treated explicitly, with the displacements of the remaining atoms approximately found through interpolation. In this way, the degrees of freedom are reduced to only the coordinates of the repatoms.

The QC incorporates such a scheme by recourse to the interpolation functions of the finite element method (FEM) (see, for example, [16]). Figure 1 illustrates the approach in two-dimensions in the vicinity of a dislocation core. The filled atoms are the selected repatoms, which are meshed by a space-filling set of linear triangular finite elements. Any atom not chosen as a repatom, like the one labeled “A”, is subsequently constrained to move according to the interpolated displacements of the element in which it resides. The density of repatoms is chosen to vary in space according to the needs of the problem of interest. In regions where full atomistic detail is required, all atoms are chosen as repatoms, with correspondingly fewer in regions of more slowly varying deformation gradient. This is illustrated in Fig. 1, where all the atoms around the dislocation core are chosen as repatoms. Further away, where the crystal experiences only the linear elastic strains due to the dislocation, the density of repatoms is reduced.

This first approximation of the QC, then, is to replace the energy E^{tot} by $E^{\text{tot},h}$:

$$E^{\text{tot},h} = \sum_{i=1}^N E_i(\mathbf{u}^h). \quad (8)$$

In this equation the atomic displacements are now found through the interpolation functions and take the form


$$\mathbf{u}^h = \sum_{\alpha=1}^{N_{\text{rep}}} S_{\alpha} \mathbf{u}_{\alpha}, \quad (9)$$

where S_{α} is the interpolation (shape) function associated with repatom α , and N_{rep} is the number of repatoms, $N_{\text{rep}} \ll N$. Note that the formal summation over the shape functions in Eq. (9) is in practice much simpler due to the compact support of the finite element shape functions. Specifically, shape functions are identically zero in every element not immediately adjacent to a specific repatom. Referring back to Fig. 1, this means that the displacement of atom A is determined entirely from the sum over the three repatoms B, C and D defining the element containing A:

$$\mathbf{u}^h(\mathbf{X}_A) = S_B(\mathbf{X}_A) \mathbf{u}_B + S_C(\mathbf{X}_A) \mathbf{u}_C + S_D(\mathbf{X}_A) \mathbf{u}_D. \quad (10)$$

Introducing this kinematic constraint on most of the atoms in the body will achieve the goal of reducing the number of degrees of freedom in the problem, but notice that for the purpose of energy minimization we must still compute the energy and forces on the degrees of freedom by explicitly visiting every atom – not just the repatoms – and building its neighbor environment from the interpolated displacement fields. Next, we discuss how these calculations are approximated and made computationally tractable.

(a)



(b)




Figure 1. Selection of repatoms from all the atoms near a dislocation core are shown in (a), which are then meshed by linear triangular elements in (b). The density of the repatoms varies according to the severity of the variation in the deformation gradient. After Ref. [5]. Reproduced with permission.

2.2. Efficient Energy Calculations: The Local QC

In addition to the degree of freedom reduction described in Section 2.1, the QC requires an efficient means of computing the energy and forces without the need to visit every atom in the problem as implied by Eq. (8). The first way to accomplish this is by recourse to the so-called Cauchy–Born (CB) rule (see Ref. [17] and references therein), resulting in what is referred to as the *local* formulation of the QC.¹

The use of linear shape functions to interpolate the displacement field means that within each element, the deformation gradient will be uniform. The Cauchy–Born rule assumes that a uniform deformation gradient at the macro-scale can be mapped directly to the same uniform deformation on the micro-scale. For crystalline solids with a simple lattice structure,² this means that every atom in a region subject to a uniform deformation gradient will be energetically equivalent. Thus, the energy within an element can be estimated by computing the energy of *one* atom in the deformed state and multiplying by the number of atoms in the element. In practice, the calculation of the CB energy is done separately from the model in a “black box,” where for a given deformation gradient \mathbf{F} , a unit cell with periodic boundary conditions is deformed appropriately and its energy is computed. The strain energy density in the element is then given by

$$\mathcal{E}(\mathbf{F}) = \frac{E_0(\mathbf{F})}{\Omega_0}, \quad (11)$$

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 \mathbf{F} . Now the total energy of an element is simply this energy density times the element volume, and the total energy of the problem is simply the sum of element energies:

$$E^{\text{tot},h} \approx E^{\text{tot},h'} = \sum_{e=1}^{N_{\text{element}}} \Omega_e \mathcal{E}(\mathbf{F}_e), \quad (12)$$

where Ω_e is the volume of element e . The important computational saving made here is that a sum over all the *atoms* in the body has been replaced by a sum over all the *elements*, each one requiring an explicit energy calculation for only one atom. Since the number of elements is typically several orders of magnitude smaller than the total number of atoms, the computational

¹The term “local” refers to the fact that use of the CB rule implies that the energy at each point in the continuum will only be a function of the deformation at that point and not on its surroundings.

²A simple lattice structure is one for which there is only one atom at each Bravais lattice site. In a complex lattice with two or more atoms per site, the Cauchy–Born rule must be generalized to permit shuffling of the off-site atoms. See Ref. [12].

savings is substantial. The number of elements scales linearly with the number of repatoms, and so the local QC scales as $\mathcal{O}(N_{\text{rep}})$.

Note, however, that even in the case where the deformation is uniform within each element, the local prescription for the energy in the element is only approximate. This is because in the constrained displacement field \mathbf{u}^h , the deformation gradient varies from one element to the next. At element boundaries and free surfaces, atoms can have energies that differ significantly from that of an atom in a bulk, uniformly deformed lattice. Figure 2 illustrates this schematically for an initially square lattice deformed according to two different deformation gradients in two neighboring regions. The energy of the atom labeled as a “bulk atom” can be accurately computed from the CB rule; its neighbor environment is uniform even though some of its neighbors occupy other elements. However, the “interface atom” and “surface atom” are not accurately described by the CB rule, which assumes that these atoms see uniformly deformed bulk environments.

In situations where the deformation is varying slowly from one element to the next and where surface energetics are not important, the local approximation is a good one. Using the CB rule as in Eq. (11), the QC can be thought of as a purely continuum formulation, but with a constitutive law that is based on




Figure 2. On the left, the reference configuration of a square lattice meshed by triangular elements. On the right, the deformed mesh shows a bulk atom, for which the CB rule is exactly correct, and two other atoms for which the CB rule will give the wrong energy due to its inability to describe surfaces or changes in the deformation gradient. After Ref. [5]. Reproduced with permission.

atomistics rather than on an assumed phenomenological form. The CB constitutive law automatically ensures that the correct anisotropic crystal elasticity response will be recovered for small deformations. It is non-linear elastic (as dictated by the underlying atomistic potentials) for intermediate strains and includes lattice invariance for large deformations; for example, a shear deformation that corresponds to the twinning of the lattice will lead to a rotated crystal structure with zero strain energy density.

An advantage of the local QC formulation is that it allows the use of quantum-mechanical atomistic models that cannot be written as a sum over individual atom energies such as tight binding (TB) and DFT. In these models only the total energy of a collection of atoms can be obtained. However, for a lattice undergoing a uniform deformation it is possible to compute the energy density $\mathcal{E}(\mathbf{F})$ from a single unit cell with periodic boundary conditions. Incorporation of quantum-mechanical information into the atomic model generally ensures that the description is more transferable, i.e., it provides a better description of the energy of atomic configurations away from the reference structure to which empirical potentials are fitted. This allows truly first-principles simulations of some macroscopic processes such as homogeneous phase transformations.

2.3. More Accurate Calculations: Mixed Local/Non-Local QC

The local QC formulation successfully enhances the continuum FEM framework with atomistic properties such as nonlinearity, crystal symmetry and lattice invariance. The latter property means that dislocations may exist in the local QC. However, the core structure and energy of these dislocations will only be coarsely represented due to the CB approximation of the energy. The same is true for other defects such as surfaces and interfaces, where the deformation of the crystal is non-uniform over distances shorter than the cut-off radius of the interatomic potentials. For example, to correctly account for the energy of the interface shown in Fig. 2, the non-uniform environment of the atoms along the interface must be correctly accounted for. While the local QC can support deformations (such as twinning) which may lead to microstructures containing such interfaces, it will not account for the energy cost of the interface itself.

In order to correctly capture these details, the QC must be made non-local in certain regions. The energy of Eq. (8), which in the local QC was approximated by Eq. (12), must instead be approximated in a way that is sensitive to non-uniform deformation and free surfaces, especially in the limit where full atomistic detail is required.

We now make the ansatz that the energy of Eq. (8) can be approximated by computing only the energy of the repatoms, but we will identify each repatom as being either *local* or *non-local* depending on its deformation environment. Thus, the repatoms are divided into N_{loc} local repatoms and N_{nl} non-local repatoms ($N_{\text{loc}} + N_{\text{nl}} = N_{\text{rep}}$). The energy expression is then approximated as

$$E^{\text{tot},h} \approx \sum_{\alpha=1}^{N_{\text{nl}}} n_{\alpha} E_{\alpha}(\mathbf{u}_h) + \sum_{\alpha=1}^{N_{\text{loc}}} n_{\alpha} E_{\alpha}(\mathbf{u}_h). \quad (13)$$

The important difference between Eq. (8) and Eq. (13) is that the sum on *all* the atoms in the problem has been replaced with a sum on *only* the repatoms. The function n_{α} is a weight assigned to repatom α , which will be high for repatoms in regions of low repatom density and vice versa. For consistency, the weight functions must be chosen so that

$$\sum_{\alpha=1}^{N_{\text{rep}}} n_{\alpha} = N, \quad (14)$$

which further implies (through the consideration of a special case where every atom in a problem is made a repatom) that in atomically-refined regions, all $n_{\alpha} = 1$. From Eq. (14), the weight functions can be physically interpreted as the number of atoms represented by each repatom α .

The weight n_{α} for each repatom (local or non-local) is determined from a tessellation that divides the body into cells around each repatom. One physically sensible tessellation is Voronoi cells [18], but an approximate Voronoi diagram can be used instead due to the high computational overhead of the Voronoi construction. In practice, the coupled QC formulation makes use of a simple tessellation based on the existing finite element mesh, partitioning each element equally between each of its nodes. The volume of the tessellation cell for a given repatom, divided by the volume of a single atom (the Wigner–Seitz volume) provides n_{α} for the repatom. In typical QC simulations, non-local regions are fully refined down to the atomic scale, and so the weight of the non-local repatoms is one.

To compute the energy of a *local* repatom α , we recognize that of the n_{α} atoms it represents, n_{α}^e reside in each element e adjacent to the repatom. The weighted energy contribution of the repatom is then found by applying the CB rule within each element adjacent to α such that

$$E_{\alpha} = \sum_{e=1}^M \frac{n_{\alpha}^e}{n_{\alpha}} \Omega_0 \mathcal{E}(\mathbf{F}_e), \quad n_{\alpha} = \sum_{e=1}^M n_{\alpha}^e, \quad (15)$$

where $\mathcal{E}(\mathbf{F}_e)$ is the energy density in element e by the CB rule, Ω_0 is the Wigner–Seitz volume of a single atom and e runs over all elements adjacent to α .

Note that this description of the local repatoms is exactly equivalent to the element-by-element summation of the local QC in Eq. (12); it is only the way that the energy partitioning is written that is different. In a mesh containing only local repatoms, the two formulations are the same, but the summations have been rearranged from one over elements in Eq. (12) to one over the repatoms here.

The energy of each *non-local* repatom is computed from the deformed neighbor environment dictated from the current interpolated displacements in the elements. In essence, every atom in the vicinity of a non-local repatom is displaced to the deformed configuration, the energy of each non-local repatom in this configuration is computed from Eq. (4), and the total energy is the sum of these repatom energies weighted by n_α . For example, the energy of the repatom identified as an “interface atom” in Fig. 2 requires that the neighbor environment be generated by displacing each neighbor according to the element in which it resides. Thus, the energy of each non-local repatom is exactly as it should be under the displacement field \mathbf{u}^h , while the local approximation is used in regions where the deformation is uniform on the atomic scale. From this starting point, the forces on all the repatoms can be obtained as the appropriate derivatives of Eq. (13), and energy minimization can proceed.

When making use of the mixed formulation described in Eq. (13), it now becomes necessary to decide whether a given repatom should be local or non-local. This is achieved automatically in the QC using a *non-locality criterion*. Note that simply having a large deformation in a region does not in itself require a non-local repatom, as the CB rule of the local formulation will exactly describe the energy of any *uniform* deformation, regardless of the severity. The key feature that should trigger a non-local treatment of a repatom is a significant *variation* in the deformation gradient on the atomic scale in the repatom’s proximity. Thus, the non-locality criterion is implemented as follows. A cut-off, r_{nl} , is empirically chosen to be between two and three times the cut-off radius of the interatomic potentials. The deformation gradients in every element within this cut-off of a given representative atom are compared, by looking at the differences between their eigenvalues. The criterion is then:

$$\max_{a,b;k} |\lambda_k^a - \lambda_k^b| < \epsilon, \quad (16)$$

where λ_k^a is the k th eigenvalue of the right stretch tensor $\mathbf{U}_a = \sqrt{\mathbf{F}_a^T \mathbf{F}_a}$ in element a , $k = 1 \cdots 3$, and the indices a and b run over all elements within r_{nl} of a given repatom. The repatom will be made local if this inequality is satisfied, and non-local otherwise. In practice, the tolerance ϵ is determined empirically. A value of 0.1 has been used in a number of tests and found to give good results. The effect of this criterion is clusters of non-local atoms in regions of rapidly varying deformation.

The fact that the non-local repatoms tend to cluster into atomistically refined regions surrounded by local regions leads to non-local/local interfaces in the QC. As in all attempts to couple a non-local atomistic region to a local continuum region found in the literature, this will lead to spurious forces near the interface. These forces, dubbed “ghost-forces” in the QC literature, arise due to the fact that there is an inherent mismatch between the local (continuum) and non-local (atomistic) regions in the problem. In short, the finite range of interaction in the non-local region mean that the motion of repatoms in the local region will effect the energy of non-local repatoms, while the converse may not be true. Upon differentiating Eq. (13), forces on repatoms in the vicinity of the interface may include a non-physical contribution due to this asymmetry. Note that these ghost forces are a consequence of differentiating an approximate energy functional, and therefore they still are “real” forces in the sense that they come from a well-defined potential. The problem is that the mixed local/non-local energy functional of Eq. (13) is approximate, and the error in this approximation is most apparent at the interface. A consequence of this is that a perfect, undistorted crystal containing an artificial local/non-local interface will be able to lower its energy below the ground-state energy by rearranging the atoms in the vicinity of the interface. This is clearly a non-physical result.

In Ref. [3], a solution to the ghost forces was proposed whereby corrective forces were added as dead loads to the interface region. In this way, there is a well-defined contribution of the corrective forces to the total energy functional (since the dead loads are constant) and the minimization of the modified energy can proceed using standard conjugate gradient or Newton–Raphson techniques. The procedure can be iterated to self-consistency.

2.4. Evolving Microstructure: Automatic Mesh Adaption

The QC approach outlined in the previous sections can only be successfully applied to general problems in crystalline deformation if it is possible to ensure that the fine structure in the deformation field will be captured. Without *a priori* knowledge of where the deformation field will require fine-scale resolution, it is necessary that the method have an automatic way to adapt the finite element mesh through the addition or removal of repatoms.

To this end, the QC makes use of the finite element literature, where considerable attention has been given to adaptive meshing techniques for many years. Typically in finite element techniques, a scalar measure is defined to quantify the error introduced into the solution by the current density of nodes (or repatoms in the QC). Elements in which this error estimator is higher than some prescribed tolerance are targeted for adaption, while at the same time

the error estimator can be used to remove unnecessary nodes from the model. The error estimator of Zienkiewicz and Zhu [19], originally posed in terms of errors in the stresses, is re-cast for the QC in terms of the deformation gradient. Specifically, we define the error estimator to be

$$\varepsilon_e = \left[\frac{1}{\Omega_e} \int_{\Omega_e} (\bar{\mathbf{F}} - \mathbf{F}_e) : (\bar{\mathbf{F}} - \mathbf{F}_e) d\Omega \right]^{1/2}, \quad (17)$$

where Ω_e is the volume of element e , \mathbf{F}_e is the QC solution for the deformation gradient in element e , and $\bar{\mathbf{F}}$ is the L_2 -projection of the QC solution for \mathbf{F} , given by

$$\bar{\mathbf{F}} = \mathbf{S} \mathbf{F}_{\text{avg}}. \quad (18)$$

Here, \mathbf{S} is the shape function array, and \mathbf{F}_{avg} is the array of nodal values of the projected deformation gradient $\bar{\mathbf{F}}$. Because the deformation gradients are constant within the linear elements used in the QC, the nodal values \mathbf{F}_{avg} are simply computed by averaging the deformation gradients found in each element touching a given repatom. This is then interpolated throughout the elements using the shape functions, providing an estimate to the discretized field solution that would be obtained if higher order elements were used. The error, then, is defined as the difference between the actual solution and this estimate of the higher order solution. If this error is small, it implies that the higher order solution is well represented by the lower order elements in the region, and thus no refinement is required. The integral in Eq. (17) can be computed quickly and accurately using Gaussian quadrature. Elements for which the error ε_e is greater than some prescribed error tolerance are targeted for refinement. Refinement then proceeds by adding three new repatoms at the atomic sites closest to the mid-sides of the targeted elements. Notice that since repatoms must fall on actual atomic sites in the reference lattice, there is a natural lower limit to element size; if the nearest atomic sites to the mid-sides of the elements are the atoms at the element corners, the region is fully refined and no new repatoms can be added.

The same error estimator is used in the QC to remove unnecessary repatoms from the mesh. In this process, a repatom is temporarily removed from the mesh and the surrounding region is locally remeshed. If all of the elements produced by this remeshing process have a value of the error estimator below the threshold, the repatom can be eliminated.

3. Practical Issues in QC Simulations

In this section, we will use a specific, simple example to highlight the practical issues surrounding solutions using the QC method. The example to be

discussed is also provided with the QC download at qcmethod.com, and it is discussed in even greater detail in the documentation that accompanies that code.

3.1. Problem Definition

Consider the problem of a twin boundary in face-centered cubic (FCC) aluminum. The boundary is perfect but for a small step. A question of interest may be “how does this stepped boundary respond to mechanical load?” In this example, we probe this question by using the QC method to solve the problem shown in Fig. 3(a), where two crystals, joined by a stepped twin boundary, are sheared until the boundary begins to migrate due to the load. The result will elucidate the mechanism of this migration.

The implementation of the QC method used to solve this problem has been described as “two and a half” dimensional to emphasize that, while it is not a fully 3D model it is also not simply 2D. Specifically, the reference crystal structure is 3D, and all the underlying atomistic calculations (both local and non-local) consider the full, 3D environment of each atom. However, the deformation of the crystal is constrained such that the three components of displacement, u_x , u_y and u_z are functions only of two coordinates x and y . This allows, for example, both edge and screw dislocations, but forces the line direction of the dislocations to be along z . For the reader who is familiar with purely atomistic simulations, this is equivalent to imposing periodic boundary conditions along the z direction, and then using a periodic cell with the

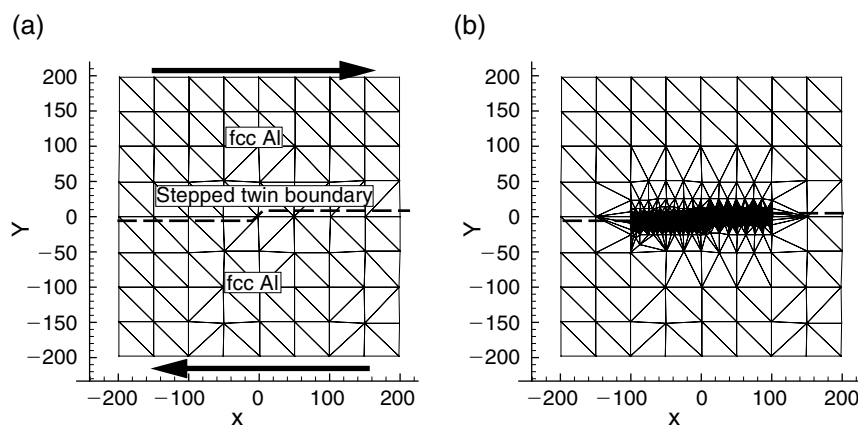


Figure 3. (a) Initial coarse mesh used to define the simulation volume and (b) the final mesh after the automatic adaption.

minimum possible thickness along z to produce the correct crystal structure. We sometimes refer to this as a “2D” implementation for brevity, but ask that the reader bears in mind the true nature of the model. The use of a 2D implementation of the QC to study this problem is appropriate given its geometry. However, fully 3D implementations of the QC exist and these must be used for many problems of interest (see examples in Ref. [5]).

The starting point for a QC simulation is a crystal lattice, defined by an origin atom and a set of Bravais vectors as in Eq. (1). To allow the QC method to model polycrystals, it is necessary to define a unique crystal structure within each grain. The shape of each grain is defined by a simple polygon in 2D. Physically, it makes sense that the polygons defining each grain do not overlap, although it may be possible to have holes between the grains. In our example, it is easy to see how the shape of the two grains could be defined to include the grain boundary step. Mathematically, the line defining the boundary should be shared identically by the two grains, but this can lead to numerical complications; for example in checking whether two grains overlap. Fortunately, realistic atomistic models are unlikely to encounter atoms that are less than an Angström or so apart, and so there exists a natural “tolerance” in the definition of these polygons. For example, a gap between grains of 0.1 \AA will usually provide sufficient numerical resolution between the grains without any atoms falling “in the gap” and therefore being omitted from the model.

In the QC implementation, the definition of the grains is separate from the definition of the actual volume of material to be simulated. This simulation volume is defined by a finite element mesh between an initial set of repatoms. Each element in this mesh must lie within one or more of the grain polygons described above, but the finite element mesh need not fill the entire volume of the defined grains. It is useful to think of the actual model (the mesh) being “cut-out” from the previously defined grain structure. For our problem, a sensible choice for the initial mesh is shown in Fig. 3(a), where the grain boundary lies approximately (to within the height of the step) along the line $y = 0$. Elements whose centroid lie above or below the grain boundary are assumed to contain material oriented according to the lattice of the upper or lower grain, respectively.

Since our interest here is atomic scale processes along the grain boundary, it is clear that the model shown in Fig. 3(a), with elements approximately 50 \AA in width, will not provide the necessary accuracy. Thus, we can make use of the QC’s automatic adaption to increase the resolution near the grain boundary. The main adaption criterion, as outlined earlier, is based on error in the finite element interpolation of the deformation gradient. However, there will initially be no deformation near the grain boundary and thus no reason for automatic adaption to be triggered. It is therefore necessary to force the model to adapt in regions that are inhomogeneous at the atomic scale for reasons other than deformation. To this end, we can identify certain segments of the

grain boundary as “active” segments. Any repatom within a prescribed distance of an active segment will be made non-local. This further implies that the elements touching this repatom will be targeted for refinement, since we require that $n_\alpha = 1$ for all non-local repatoms. The effect of such a technique is shown in Fig. 3(b), where the segment of the boundary between $x = -100$ and 100 \AA was defined to be active. The result is that the grain boundary structure is correctly captured in the vicinity of the step, as well as for some distance on either side of the step.

3.2. Solution Procedure

In the static QC implementation, the solution procedure amounts to minimization of the total energy (elastic energy plus the potential energy of the applied loads, see Eq. (6)) for a given set of boundary conditions (applied displacements or forces on certain repatoms). However, problems solved using the QC method are typically highly nonlinear, and as such their energy functional typically includes many local minima. In order to find a physically realistic solution, it is necessary to use a quasi-static loading approach, whereby boundary conditions are gradually incremented, the energy is minimized, and the minimum energy configuration is used in generating an initial guess to the solution after the subsequent load increment. Again, we can refer to the specific example of the stepped twin boundary to make this more clear.

Our desire, in this example, is to study the effect of applying a shear strain to the stepped twin boundary. Specifically, we may be interested in knowing the critical shear strain at which the boundary begins to migrate and to understand the mechanism of this migration. We begin by choosing a sensible strain increment to apply, such that the incremental deformation will not be too severe between minimization steps. For this example, the initial guess, \mathbf{u}_0^{n+1} , used to solve for the relaxed displacement, \mathbf{u}^{n+1} , of load step $n + 1$ is given by

$$\mathbf{u}_0^{n+1} = \mathbf{u}^n + \Delta \mathbf{F} \mathbf{X}, \quad (19)$$

where \mathbf{u}^n is the relaxed, minimum energy displacement field from load step n , $\mathbf{u}^0 = 0$, and the matrix $\Delta \mathbf{F}$ corresponding to pure shear along the y direction is

$$\Delta \mathbf{F} = \begin{bmatrix} 1 & \Delta\gamma & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (20)$$

Thus, a shear strain increment of $\Delta\gamma$ is applied, the outer repatoms are held fixed to the resulting displacements, and all inner repatoms are relaxed until the

energy reaches a minimum. Then, another strain increment is superimposed on these relaxed displacements and the process repeated. After n load steps, a total macroscopic shear strain of $\gamma = n\Delta\gamma$ has been applied to the outer boundary of the bi-crystal.

The energy minimization can be performed using several standard approaches, such as the conjugate gradient (CG) or the Newton–Raphson (NR) methods (both of which are described, for example, in Ref. [20]). The CG method has the advantage over the NR technique in that it requires only the energy functional and its first derivatives with respect to the repatom positions (i.e., the forces). The NR method requires a second derivative, or “stiffness matrix” that is not straightforward to derive or to code in an efficient manner. Once correctly implemented, however, the NR method has the advantage of quadratic convergence (compared to linear convergence for the CG method) once the system is close to the energy minimizing configuration.

By monitoring the applied force (measured as the sum of forces in the y -direction applied to the top surface of the bi-crystal) versus the accumulated shear strain, γ , it can be observed that there is an essentially linear response for the first six load steps, and then a sudden load drop from step six to seven. This jump corresponds to the first inelastic behaviour of the boundary, the mechanism of which is shown in Fig. 4. In Fig. 4(a), a close-up of the relaxed step at an applied strain of $\gamma = 0.03$ is shown, while Fig. 4(b) shows the relaxed configuration after the next strain increment at $\gamma = 0.035$. The mechanism of this boundary motion is the motion of two Shockley partial dislocations from the corners of the step along the boundary. This can be seen clearly by observing the finite element mesh between the repatoms in Fig. 4(c). Because the mesh is triangulated in the reference configuration, the effect of plastic slip is the shearing of a row of elements in the wake of the moving dislocations.

One challenge in modeling dislocation motion in crystals at the atomic scale is evident in this simulation. In crystals with a low Peierls resistance like the FCC crystal modelled here, dislocations will move long distances under small applied stresses. In this simulation, the Shockley partials which nucleated at the step move to the ends of the region of atomic-scale refinement. In order to rigorously compute the equilibrium position of the dislocations, it would be necessary to further adapt the model. The presence of the dislocation in close proximity to the larger elements to the left of the fully refined region will trigger the adaption criterion, as well as increase the number of repatoms that are non-local according to the non-locality criterion defined earlier. This will allow the dislocations to move somewhat further upon subsequent relaxation. In principle, this process of iteratively adapting and relaxing can be repeated until the dislocations come to its true equilibrium, which in this example would be at the left and right free surfaces of the bi-crystal.

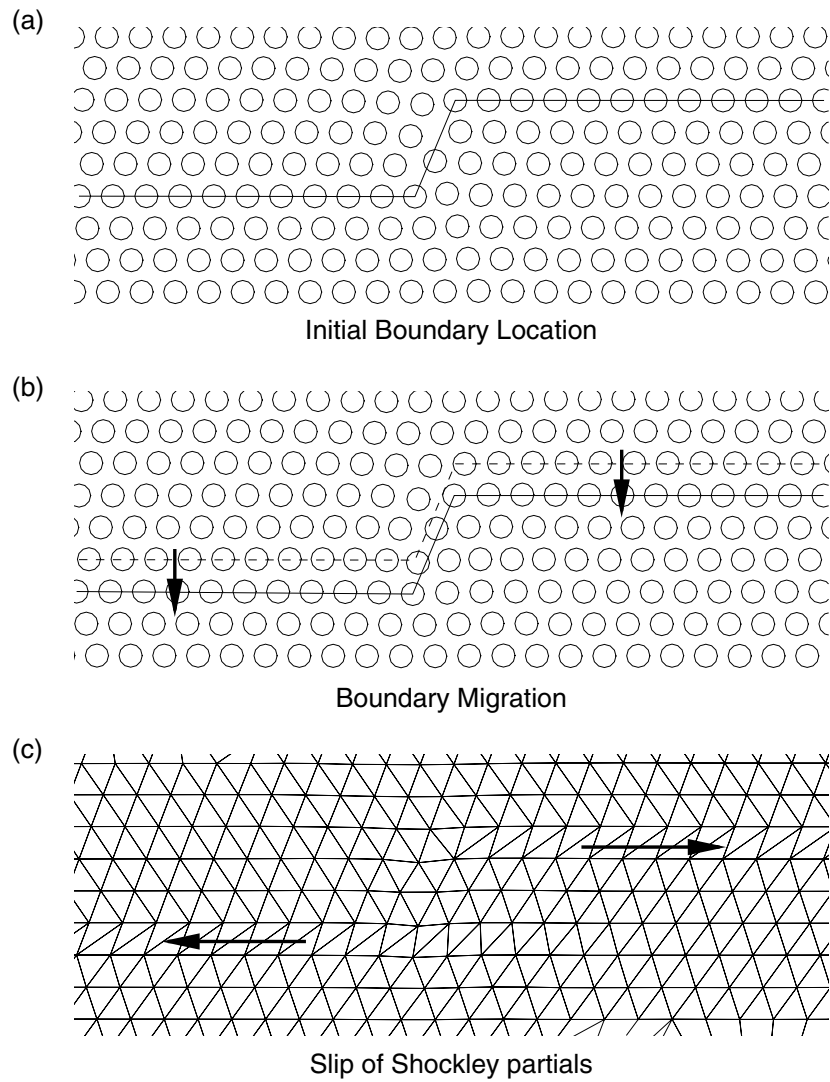


Figure 4. Mechanism of migration of the twin boundary under shear. (a) Before migration. (b) After migration (c) Deformed mesh showing the motion of Shockley partial dislocations.

In practice, however, we may not be interested in the full details of where this dislocation comes to rest, if we are willing to accept some degree of error in the simulation. Specifically, the fact that the dislocation is held artificially close to the step may effect the critical load level at which subsequent migration events occur. The compromise is made for the sake of computational speed, which will be significantly compromised if we were to iteratively adapt and relax many times for each load step.

4. Summary

This review has summarized the theory and practical implementation of the QC method. Rather than provide an exhaustive review of the QC literature (which can already be found, for example, in Ref. [5]), the intent has been to provide a simple overview for someone interested in understanding one implementation of the QC method. More specific details, including free, open-source code and documentation, can be found at www.qcmethod.com.

References

- [1] E.B. Tadmor, M. Ortiz, and R. Phillips, "Quasicontinuum analysis of defects in solids," *Phil. Mag. A*, 73, 1529–1563, 1996a.
- [2] E.B. Tadmor, R. Phillips, and M. Ortiz, "Mixed atomistic and continuum models of deformation in solids," *Langmuir*, 12, 4529–4534, 1996b.
- [3] V.B. Shenoy, R. Miller, E. Tadmor, D. Rodney, R. Phillips, and M. Ortiz, "An adaptive methodology for atomic scale mechanics: the quasicontinuum method," *J. Mech. Phys. Sol.*, 47, 611–642, 1998a.
- [4] V.B. Shenoy, R. Miller, E.B. Tadmor, R. Phillips, and M. Ortiz, "Quasicontinuum models of interfacial structure and deformation," *Phys. Rev. Lett.*, 80, 742–745, 1998b.
- [5] R.E. Miller and E.B. Tadmor, "The quasicontinuum method: overview, applications and current directions," *J. of Computer-Aided Mater. Design*, 9(3), 203–231, 2002.
- [6] M. Ortiz, A.M. Cuitino, J. Knap, and M. Koslowski, "Mixed atomistic continuum models of material behavior: the art of transcending atomistics and informing continua," *MRS Bull.*, 26, 216–221, 2001.
- [7] D. Rodney, "Mixed atomistic/continuum methods: static and dynamic quasicontinuum methods," In: A. Finel, D. Maziere, and M. Veron (eds.), *NATO Science Series II*, Vol. 108, "Thermodynamics, Microstructures and Plasticity," Kluwer Academic Publishers, Dordrecht, 265–274, 2003.
- [8] M. Ortiz and R. Phillips, "Nanomechanics of defects in solids," *Adv. Appl. Mech.*, 36, 1–79, 1999.
- [9] W.A. Curtin and R.E. Miller, "Atomistic/continuum coupling methods in multi-scale materials modeling," *Model. Simul. Mater. Sci. Eng.*, Vol. 11(3), R33–R68, 2003.
- [10] A. Carlsson, "Beyond pair potentials in elemental transition metals and semiconductors," *Sol. Stat. Phys.*, 43, 1–91, 1990.
- [11] V. Shenoy, V. Shenoy, and R. Phillips, "Finite temperature quasicontinuum methods," *Mater. Res. Soc. Symp. Proc.*, 538, 465–471, 1999.
- [12] E. Tadmor, G. Smith, N. Bernstein, and E. Kaxiras, "Mixed finite element and atomistic formulation for complex crystals," *Phys. Rev. B*, 59, 235–245, 1999.
- [13] M. Daw and M. Baskes, "Embedded-atom method: derivation and application to impurities, surfaces, and other defects in metals," *Phys. Rev. B*, 29, 6443–6453, 1984.
- [14] J. Norskov and N. Lang, "Effective-medium theory of chemical binding: application to chemisorption," *Phys. Rev. B*, 21, 2131–2136, 1980.
- [15] F. Stillinger and T. Weber, "Computer-simulation of local order in condensed phases of silicon," *Phys. Rev. B*, 31, 5262–5271, 1985.

- [16] O.C. Zienkiewicz, *The Finite Element Method*, vols. 1–2, 4th edn. McGraw-Hill, London, 1991.
- [17] J. Ericksen, In: M. Gurtin (ed.), *Phase Transformations and Material Instabilities in Solids*, Academic Press, New York.
- [18] A. Okabe, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Wiley, Chichester, England, 1992.
- [19] O.C. Zienkiewicz and J. Z. Zhu, “A simple error estimator and adaptive procedure for practical engineering analysis,” *Int. J. Numer. Meth. Eng.*, 24, 337–357, 1987.
- [20] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in FORTRAN: The Art of Scientific Computing*, 2nd edn. Cambridge University Press, Cambridge, 1992.