

Concurrent coupling of length scales: Methodology and application

Jeremy Q. Broughton

Complex Systems Theory Branch, Naval Research Laboratory, Washington, D.C. 20375

Farid F. Abraham

IBM Research Division, Almaden Research Center, San Jose, California 95120

Noam Bernstein and Efthimios Kaxiras

Department of Physics, Division of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138

(Received 18 November 1998)

A strategic objective of computational materials physics is the *accurate* description of specific materials on length scales approaching the meso and macroscopic. We report on progress towards this goal by describing a seamless coupling of continuum to statistical to quantum mechanics, involving an algorithm, implemented on a parallel computer, for handshaking between finite elements, molecular dynamics, and semiempirical tight binding. We illustrate and validate the methodology using the example of crack propagation in silicon. [S0163-1829(99)01223-0]

I. INTRODUCTION

Historically, there has been a rich tradition of attempting to couple length and time scales in serial fashion. By this we mean that one set of calculations at a very fundamental level, and of high-computational complexity, is used to evaluate constants for use in a more approximate or phenomenological computational methodology at a longer length/time scale. A good example of this was the pioneering work of Clementi and coworkers¹ in the 1980's who used high-quality quantum-mechanical methods to evaluate the interaction of several water molecules. From this database, an empirical potential was parametrized for use in molecular-dynamics (MD) atomistic simulation. Such simulation was then used to evaluate the viscosity of water from the atomic autocorrelation functions. Finally, the computed viscosity was employed in a computational fluid dynamics calculation to predict tidal circulation in Buzzard's Bay, Massachusetts. This tour-de-force of computational physics was a powerful example of the sequential coupling of length and time scales: one series of calculations is used as input to the next up the length/time scale hierarchy. Another good example is in atmospheric and environmental science² where chemists use methods of high-computational complexity to evaluate reaction barriers of simple chemical reactions, which are then used in large rate equation coupled to spatial grid codes to determine and predict chemical meteorology. There are many other examples in the literature. But what unifies all these schemes is that an appropriate computational methodology is used for a given scale or task, whether it be the accuracy of quantum mechanics at the shortest scales, or rate theory and/or fluid dynamics at longest scales.

In contrast, there has been comparatively little effort devoted to the parallel coupling of different computational schemes for a simultaneous attack on a given problem; in our case, our interest dictates specific attention towards issues in materials or solid-state physics. We will focus specifically on the coupling of length scales (CLS). Perhaps such attention has been lacking because it necessitates significant computer

power that heretofore was lacking. The driver for attempting now to do so has been the advent of parallel computers. By dividing a problem into its natural components, each of which may be addressed by one or more processors, CLS becomes a more tractable proposition. The issue then becomes one of determining how to do the "handshaking" between the different regions. This is not just an algorithmic issue but also one that requires physical insight.

The reasons for wanting to couple length scales are not solely those predicted on computational power, however. The most important is that there is a large class of problems for which the physics is *inherently* multiscale; that is the different scales interact strongly to produce the observed behavior. Turbulence and crack propagation are good examples. It is necessary to know what is happening simultaneously in each region since one is strongly coupled to another. Another important realization is that once the issues of interfacing between the different regions have been solved, the overall algorithm becomes computationally very efficient. The reason is that one is using the right "tool" for the right part of the system, that tool having been optimized, historically, to solve a particular problem. In our case, the three regions of interest will be continuum mechanics, the implementation of which will be via finite elements (FE); atomistic statistical mechanics, implemented by molecular dynamics (MD); and mean-field quantum mechanics represented by semiempirical tight binding (TB). Since our interest pertains mainly to the solid state, we will confine our discussion to implementation of elasticity theory via finite elements. But perhaps the most important result from this general approach is that once the handshakes are made "seamless," the algorithm is not only efficient, it is also very *accurate*. "Accurate" implies that the dynamics of the phenomenon under study are indistinguishable whether they be determined from a length-scale-coupled system or from one of the same size comprising TB atoms *only*. The many-body behavior of specific materials really can be addressed on structurally important length scales.

Pioneering work on coupling atomistics to finite elements

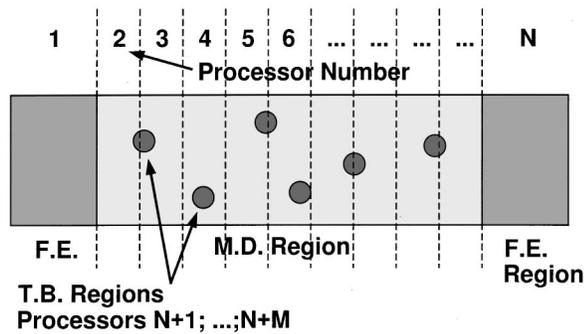


FIG. 1. Illustration of domain decomposition of pseudo-1D system showing coupling of length scales. There may be more than one processor per TB region.

has been performed by Tadmor, Ortiz, and Phillips.³ The embryonic ideas in their work were initiated before the advent of really accessible parallel computer power. Their interest focussed upon how to embed a defective system (such as a crack) within a continuum, but without the usual assumptions inherent in continuum codes in which a rather ad-hoc criterion is used for failure in a given region of space. In their formulation, the finite element mesh permeates the entire system, right down to atomic dimensions. An underlying atomistic Hamiltonian is used to determine the energy density of the system; a separate atomistic calculation is required for each cell in their finite element mesh. At present, their scheme is a zero temperature (T) relaxation technique. It is an adaptive methodology since mesh refinement occurs when defects propagate.

A complementary methodology has very recently been propounded by Rafii-Tabar and coworkers.⁴ It involves a stochastic coupling of a molecular-dynamics region to a finite element region. The system is propagated in time using a stochastic differential equation so as to produce something resembling Langevin dynamics. The molecular-dynamics region involves a few thousand atoms and represents a small part of the computational complexity of the simulation. However, initial results on crack propagation in a two-dimensional (2D) silver plate are encouraging.

In contrast, in this paper, we will advocate a rather different approach.^{5,6} We will be specifically interested in solid systems at finite temperature in which the trajectories are dictated by classical mechanics and for which the dynamical complexity of the system is of importance. The ensuing discussion centers on Fig. 1. We envisage a system in which the dynamically crucial part comprises a relatively small subset of all the atoms in the system. This part involves a disruption of the symmetry of the system. This is the region of broken bonds, or at least that region where bonds are breaking and which dictates the kinetics of the system. We imagine, for example, such a region being typified by a propagating crack tip. It is important to describe the energetics of this part of the system very accurately. Since it is a region where bonds are breaking it requires a quantum-mechanical description; empirical interatomic potentials are prone to be untrustworthy in such cases. We chose to describe this part of the system with a tight-binding Hamiltonian since this is one of the fastest quantum-mechanical algorithms that contains the basic physics of breaking bonds and from which it is simple to extract forces.

Since the computational complexity of tight binding is high (see below), it is not possible to describe the entire meso/macro system with such. It is important, however, that the region of dynamical breaking bonds be embedded in the correct atomistic statistical mechanical environment. Thus, we surround the TB region with an MD region in which the atoms interact with one another via a carefully parametrized empirical potential for the system of interest. The system, on this length scale, is less perturbed from equilibrium than the bond-ruptured region and we expect the empirical potential to work well. This MD region is required because it correctly captures the necessary thermal fluctuations and the pressure waves emitted by the bond rupturing and/or defective region. The MD region may also contain those areas of the system that are defective but for which the primary dynamics are no longer important. An example might be the surfaces of a crack entrained behind a propagating crack tip.

Lastly, although MD system sizes currently can run to hundreds of millions of atoms,^{7,8} when considering the meso and macro scales even these formidable calculations are unable to represent properly the rest of the environment of the dynamical system. Here, the precise statistical mechanics is less important than allowing the free passage of (usually long wavelength) energy into or out of the system. Such a coupling to the rest of the universe is important when studying problems where much of the physics is invested with long-range interactions, such as a long-range elastic field. Examples include delamination of ceramic/insulator interfaces, crack propagation, and dynamical micro electro mechanical systems. It is here that use of the basic tool of engineers is appropriate, namely finite elements, to solve for the displacement field of the system. In this far-field region of our envisaged system it is appropriate to embed the MD region in a continuum mechanical representation. Here, atoms are displaced only slightly from equilibrium and elasticity theory should work very well. Indeed, finite elements are the method of choice for this region. The FE algorithm is more computationally efficient than the MD since it deals only with the minimal degrees of freedom necessary to describe the correct physics. Consider, for a moment, large-scale simulations of crack propagation. Graphical representations of such systems focus mainly upon the high-energy parts of the system; namely the crack faces and emitted dislocations.^{7,9} In the far-field regions, little of interest is happening: many computer cycles are spent describing trajectories of atoms that do very little apart from vibrating around lattice sites. It is here that a mean-field approach is all that is required. Yet it is not appropriate to dispense with this region completely: again, atomistic simulations of crack propagation make the point. Large-scale MD simulation shows that long-wavelength pressure pulses radiate from the crack tip.¹⁰ These propagate to the edge of the simulation, bounce off the edges of the computational cell and reflect back towards the cell center thereby contaminating the very phenomenon under study. A FE region surrounding the MD region allows such pressure waves to propagate harmlessly into the continuum (the continuum region can be made large). Note that although Fig. 1 shows FE regions at the extreme left- and right-hand ends of our pseudo-1D topology (actually, a 3D simulation will be performed as demonstration) there is nothing to prevent decoration with FE regions

on all sides of the central MD region. Such decoration depends upon the physics of the system under study.

This approach should not be viewed as universal panacea. It works well on a subset, albeit an important subset, of systems. These involve a central region that dictates the kinetic behavior of the system surrounded by a region only slightly perturbed from equilibrium. The MD region is chosen to be of sufficient size to allow all defects to form and/or propagate. The trajectory of the region of breaking bonds is tracked via a dynamical relocatable TB region. But once the defects reach the MD/FE interface, the algorithm, as propounded below, must terminate. A second restriction, upon which we have not yet touched, is that of timescale. As we illustrate below, the advent of parallel architecture machines allows access to meso and macroscale system sizes, but as yet the total elapsed time of the simulation is, and therefore those phenomena that may be accessed are, restricted to approximately one nanosecond. Time, unfortunately, is a sequential object and mere access to more processors does not solve, in any paradigmatic way, the issue of timescale in *dynamic* simulation. Progress in algorithms that address the timescale issue would revolutionize materials simulation.

Below, in Sec. II, we describe the principal components of our algorithm. Firstly, we outline the three independent methodologies of TB, MD, and FE. Secondly, we describe the handshaking between these regions, all of which must proceed in lock step. Section III illustrates and validates our methodology with the important example of crack propagation in silicon. Finally, Sec. IV gives our conclusions together with prognoses for implementations in other materials systems.

II. COMPUTATIONAL METHODOLOGIES

Before proceeding with specifics, some general points are worth making. In this discussion, the long axis of Fig. 1 will be designated y ; the other axis in the plane of the figure is x ; the axis into the paper will be z . The crux of this methodology is to allow the study of equilibrium and nonequilibrium dynamics of macroscale systems. In order to take meaningful ensemble averages, or in order to access timescales of use, it is necessary to propagate the system through times of order one nanosecond. A typical MD time step is, for chemically bonded systems, approximately 10^{-15} s: thus, on (say) 50 to 100 processors, one time step is about one second of wall-clock time. One million time steps may therefore be achieved in approximately ten days on a significant, but not unreasonable, fraction of a parallel machine.

The second point relates to the material with which we have chosen to illustrate our methodology: silicon. This is a material for which many very good empirical interatomic potentials¹¹ already exist: we chose the Stillinger-Weber¹² because of its computational simplicity. We chose silicon for other reasons too: its industrial importance for one, its locality of bonding for another. The principles outlined below will work for metals, and are the subject of on-going work.

Thirdly, our choice of tight binding¹³ to illustrate quantum-mechanical coupling is dictated mainly by the requirement of computational speed: the need to propagate a nontrivial number of atoms under the influence of this Hamiltonian within one second of wall-clock time. Other

Hamiltonians are possible in the quantum region, such as the moment methods (bond order) advocated by Pettifor and coworkers.^{14,15} We chose TB for its simplicity and intuitive appeal.

Lastly, our implementation of finite elements will be at the linear elastic level. Since FE is used only in the far-field region where atoms are perturbed only slightly from equilibrium, we view it as unnecessary to employ nonlinear elasticity theory. Further, since the systems of primary interest to us involve plane strain, our implementation will be for a 2D system. A 3D far-field region would be a straightforward extension.

As the early part of Sec. I implies, and as Fig. 1 illustrates, each of the three primary algorithms, and the regions of space that they describe, are distributed to different processors. The example of crack propagation in silicon was run on the IBM SP2 at the US Air Force High-Performance Computing Center at Maui. Each FE region is handled by a different processor; the MD region is domain decomposed across several computer nodes and the TB region is likewise spread over several processors. The code was written in FORTRAN with MPI for the message passing interface. Such code ports to most parallel architecture machines. The advantage of a pseudo-1D topology is that much message passing can be performed using the “shift” operator. Further, only data within interaction range (defined by the Hamiltonian) need be passed across boundaries (defined by the domain decomposition) between processors.

The overarching theme is that a Hamiltonian H_{Tot} will be defined for the entire system. Its degrees of freedom are atomic positions \mathbf{r} and their velocities $\dot{\mathbf{r}}$ for the TB and MD regions; and displacements \mathbf{u} and their time rates of change $\dot{\mathbf{u}}$ for the FE regions. (The velocities and conjugate momenta are simply related). Equations of motion for all the relevant variables in the system are obtained by taking the appropriate derivatives of this Hamiltonian. All variables can then be updated in lock step as a function of time using the same integrator. Thus, the entire time history of the system may be obtained numerically given an appropriate set of initial conditions. Conceptually, H_{Tot} may be written

$$\begin{aligned} H_{\text{Tot}} = & H_{\text{FE}}(\{\mathbf{u}, \dot{\mathbf{u}}\} \in \text{FE}) + H_{\text{FE/MD}}(\{\mathbf{u}, \dot{\mathbf{u}}, \mathbf{r}, \dot{\mathbf{r}}\} \in \text{FE/MD}) \\ & + H_{\text{MD}}(\{\mathbf{r}, \dot{\mathbf{r}}\} \in \text{MD}) + H_{\text{MD/TB}}(\{\mathbf{r}, \dot{\mathbf{r}}\} \in \text{MD/TB}) \\ & + H_{\text{TB}}(\{\mathbf{r}, \dot{\mathbf{r}}\} \in \text{TB}). \end{aligned} \quad (1)$$

This equation should be read as implying that there are three separate Hamiltonians for each subsystem as well as Hamiltonians that dictate the dynamics of variables in the handshake regions. “MD/TB” and “FE/MD” imply such handshake regions. Following a trajectory dictated by this Hamiltonian will result in a conserved total energy. This is an important feature of our computational approach since it ensured numerical stability.

A. Molecular dynamics

Molecular dynamics involves following the classical trajectories of atomic nuclei by integrating Newton’s laws of motion for the system. In the region labeled MD, the interatomic force law will be obtained from an empirical potential. In the region labeled TB, the force law will be extracted

from a quantum-mechanical Hamiltonian that solves the mean-field equations for the electrons in the system. In our example, we studied silicon for which we chose to use the Stillinger-Weber (SW) potential, V_{SW} , which involves both two-body and three-body interatomic terms

$$V_{\text{SW}} = \sum_{i < j} v^{(2)}(r_{ij}) + \sum_{i, (j < k)} v^{(3)}(\mathbf{r}_{ij}, \mathbf{r}_{ik}). \quad (2)$$

The sums are over atomic indices i, j, k . The exact form of these interactions v are given in Ref. 12. The sum in the three-body term is written such that i is the apex of the triplet.

Forces \mathbf{f} may be extracted from H_{SW} by taking derivatives with respect to the atomic coordinates. In order to obtain a trajectory through phase space for these atoms, we require an integrator for Newton's laws of motion. We chose the velocity-Verlet algorithm since it is easily augmented to handle multiple timescale MD.¹⁶ The following algorithm is iterated:

$$\begin{aligned} \dot{\mathbf{r}}_i \left(t + \frac{\Delta t}{2} \right) &= \dot{\mathbf{r}}_i(t) + \frac{\Delta t}{2m} \mathbf{f}_i(t), \\ \mathbf{r}_i(t + \Delta t) &= \mathbf{r}_i(t) + \Delta t \dot{\mathbf{r}}_i \left(t + \frac{\Delta t}{2} \right), \\ \mathbf{f}_i(t + \Delta t) &= \frac{\partial V_{\text{SW}}}{\partial \mathbf{r}_i(t + \Delta t)}, \\ \dot{\mathbf{r}}_i(t + \Delta t) &= \dot{\mathbf{r}}_i \left(t + \frac{\Delta t}{2} \right) + \frac{\Delta t}{2m} \mathbf{f}_i(t + \Delta t). \end{aligned} \quad (3)$$

At each iteration, each of the four steps is performed sequentially for every atom i in the system. After exiting the last step, the simulation time is incremented by Δt . For silicon, we used a time step of 5×10^{-16} s. The mass of the silicon atom m is 4.6639×10^{-26} kg. Evaluation of the SW energy and its attendant forces may be coded, by taking advantage of atomic neighbor tables, so that computer time scales as $O(N)$, where N is the number of atoms in the system. The two- and three-body terms truncate smoothly to zero just before the second-neighbor distance in zero-pressure diamond cubic structure silicon.

B. Tight binding

Semiempirical tight binding involves (a) an *ansatz* for the total energy of the system and (b) a parametrization of the integrals that occur in mean-field treatments of the electronic structure of a system. The total energy of the system is written as

$$V_{\text{TB}} = \sum_{n=1}^{N_{\text{occ}}} \varepsilon_n + \sum_{i < j} v^{\text{rep}}(r_{ij}). \quad (4)$$

The sum is over all occupied states N_{occ} up to the Fermi level. Foulkes and Haydock¹³ have provided sound justifications for the form of this equation. It is similar in form to the full density-functional expression for the total energy, but the double counting of the Coulomb and exchange-correlation

terms inherent in the eigenvalue sum are approximated by the sum over all pairs of atoms of the repulsive interatomic potential v^{rep} . The eigenvalues $\{\varepsilon\}$ corresponding to the one-electron states of a first principles Hartree-Fock or density-functional calculation are obtained from a nonorthogonal one-electron Hamiltonian

$$[\mathbf{H}] \Psi_n = \varepsilon_n [\mathbf{S}] \Psi_n$$

$$\Psi_n = \sum_{i\alpha} c_{i\alpha}^n \phi_{i\alpha}, \quad (5)$$

where the matrix elements within $[\mathbf{H}]$ and $[\mathbf{S}]$ are obtained by reducing the equivalent integrals within an extensive database of first principles calculations to parametric form

$$H_{i\alpha j\beta} \equiv \langle \phi_{i\alpha} | \hat{H} | \phi_{j\beta} \rangle = h_{\alpha\beta}(\mathbf{r}_{ij}) \quad (6)$$

$$S_{i\alpha j\beta} \equiv \langle \phi_{i\alpha} | \phi_{j\beta} \rangle = s_{\alpha\beta}(\mathbf{r}_{ij}).$$

The one-electron wave functions $\{\Psi\}$ are expanded as a linear combination of atomic basis functions $\{\phi\}$. n labels the orbital number while α and β label the basis functions (in the minimal basis of silicon, these represent s , p_x , p_y , and p_z atomic orbitals). The size of the $[\mathbf{H}]$ and $[\mathbf{S}]$ matrices are therefore $(4N \times 4N)$. Their matrix elements are parametrized in two center approximation by pairwise functions h and s . These functions smoothly truncate to zero near 5 \AA , which is between the third- and fourth-neighbor distances in silicon. Note that since all integrals are represented by pair functions (v^{rep} , h , and s) the exact form of the basis functions are not required: the one-electron states $\{\Psi\}$ are represented within this formalism only by the sets of coefficients $\{c\}$ in their expansion. The functions v^{rep} , h , and s are obtained by fitting to a database involving the experimental indirect band gap of the diamond cubic structure and the total energies of crystalline and defective diamond cubic and β -tin silicon at different densities. The parameters for this fit are given by Bernstein and Kaxiras.¹⁷ For most purposes, this nonself-consistent TB Hamiltonian describes bulk, amorphous, and surfaces properties of silicon very well.

Solving for the coefficients $\{c\}$ is a generalized eigenvalue problem. For given set of atomic coordinates, the coefficients are found by diagonalization. One-electron states are occupied up to the Fermi level. A small amount of Fermi-level broadening is used. Forces are then found from the derivative of the TB energy with respect to displacement of the nuclei

$$\begin{aligned} \mathbf{f}_i^{\text{TB}} &= - \left[\sum_{n=1}^{N_{\text{occ}}} \sum_{\alpha} c_{i\alpha}^n \sum_{j\beta} c_{j\beta}^n \left(\frac{\partial H_{i\alpha j\beta}}{\partial \mathbf{r}_i} - \varepsilon_n \frac{\partial S_{i\alpha j\beta}}{\partial \mathbf{r}_i} \right) \right] \\ &\quad - \sum_{j \neq i} \frac{\partial v^{\text{rep}}(r_{ij})}{\partial \mathbf{r}_i}. \end{aligned} \quad (7)$$

Derivatives of the coefficients with respect to atomic positions do not occur. Such terms vanish identically by orthogonality. Knowing the forces, atomic coordinates may be advanced through time using exactly the same algorithm as that used for the SW system; namely that given by Eq. (3). The same time step may be used since Einstein frequencies

in both cases are very similar; indeed SW parameters could be adjusted to ensure equality.

The above describes an $O(N^3)$ algorithm: brute force diagonalization is $O(N^3)$, which parallelizes poorly. This scaling may be improved somewhat by implementing a fictitious Lagrangian¹⁸ but this at best is $O(N^2)$. There is much discussion in the literature about $O(N)$ schemes for electronic structure.^{19–22} We chose not to implement such yet in our CLS algorithm since (a) such methods often have problems with situations in which states wander across the Fermi level and (b) the crossover of improved efficiency from the $O(N^3)$ to the $O(N)$ schemes occurs at system sizes above several hundred atoms.²³ Diagonalization wins at the lower atom numbers. Our stipulation of one second of wall-clock time per time step places our choice firmly in the $O(N^3)$ system-size regime. As $O(N)$ schemes improve, this choice is likely to change. For the present, then, we use the diagonalization route since it is ‘‘black box’’ and robust.

C. Finite elements

Since finite elements will be used to describe the far-field region of the simulation, linear elasticity theory will be used to develop the FE equations of motion. The total elastic energy of a solid, in the absence of tractions and body forces, is then given by

$$H_{\text{FE}} = V_{\text{FE}} + K_{\text{FE}},$$

$$V_{\text{FE}} = \frac{1}{2} \int d\Omega \left[\sum_{\mu, \nu, \lambda, \sigma=1}^3 \epsilon_{\mu\nu}(\mathbf{r}) C_{\mu\nu\lambda\sigma} \epsilon_{\lambda\sigma}(\mathbf{r}) \right], \quad (8)$$

$$K_{\text{FE}} = \frac{1}{2} \int d\Omega \rho(\mathbf{r}) \dot{\mathbf{u}}^2(\mathbf{r}).$$

Here, Ω is the volume of the system and V_{FE} is the Hookian potential energy term. It involves the symmetric strain tensor ϵ quadratically multiplying the elastic constant tensor C . The subscripts, $\mu, \nu, \lambda, \sigma$ denote Cartesian directions. The second term, the kinetic energy, K_{FE} , involves the time rate of change of the displacement field $\dot{\mathbf{u}}$ and the mass density ρ . The strains and displacements are simply related

$$\epsilon_{\mu\nu} = \left[\frac{\partial u_\mu}{\partial r_\nu} + \frac{\partial u_\nu}{\partial r_\mu} \right]. \quad (9)$$

They are viewed as being continuous variables and so the total energy of the system is an integral of these quantities over the volume of the sample. The FE algorithm involves dividing the volume into cells: often tetrahedra in 3D and triangles in 2D.²⁴ An exhaustive literature exists on how to choose these cells including automatic mesh generation and adaptive meshing.²⁵ We will describe our mesh in the next subsection. If the values of the displacements and their time derivatives are known at the vertices of these cells, then interpolation functions may be used to determine the values of these variables everywhere within each cell. A common choice, made because of computational simplicity, is linear interpolation inside each cell. The result is that the displacement fields are represented in piecewise smooth fashion. In our example of crack propagation, we will be interested in plane strain situations for which a 2D far-field region suf-

fices. We will confine the rest of the description therefore to triangular cells. Equation (8) can now be approximated by

$$H_{\text{FE}} = \frac{1}{2} \sum_m^{N_{\text{cell}}} \sum_{p,q=1}^6 [u_p^m K_{pq}^m u_q^m + \dot{u}_p^m M_{pq}^m \dot{u}_q^m], \quad (10)$$

where $[\mathbf{K}]$ and $[\mathbf{M}]$ are local stiffness and mass matrices, respectively. The displacements and their time rates of change are now only defined at the apices of each triangle. The cell index is denoted by m . The total number of FE cells is N_{cell} . The (p, q) sum runs over the (3×2) Cartesian directions associated with the apices of a 2D triangular cell. Thus, the total FE energy is given as a sum of products of local matrices. The stiffness matrix is given by

$$[\mathbf{K}]^m = \frac{L}{4A_m} [\mathbf{B}^T][\mathbf{C}][\mathbf{B}]^m, \quad (11)$$

where A^m is the area of the m th triangle and $[\mathbf{C}]$ is the reduced (3×3) elastic constant matrix and $[\mathbf{B}]$ is the matrix of coordinate differences of the apices of the FE mesh. L is the thickness of the material in the third dimension. The values of the elements of $[\mathbf{C}]$ depend upon the orientation of the system; they are functions of the three fundamental elastic constants of silicon, namely C_{11} , C_{12} , and C_{44} , which are given for zero-temperature SW silicon by Balamane *et al.*²⁶ and by Ray.²⁷ We employed the mean of their results. A , L , and $[\mathbf{B}]$ are time independent

$$[\mathbf{B}^T]^m = \begin{pmatrix} b_1^m & 0 & a_1^m \\ 0 & a_1^m & b_1^m \\ b_2^m & 0 & a_2^m \\ 0 & a_2^m & b_2^m \\ b_3^m & 0 & a_3^m \\ 0 & a_3^m & b_3^m \end{pmatrix}, \quad (12)$$

where

$$b_l^m = y_{l+1}^m - y_{l+2}^m$$

$$a_l^m = x_{l+2}^m - x_{l+1}^m. \quad (13)$$

Here, 1 denotes the cyclic apex index, running from 1 to 3, and x and y denote the 2D FE mesh coordinates. Displacements $\{\mathbf{u}\}$ are defined with respect to these coordinates.

The mass matrix $[\mathbf{M}]$ is handled rather differently. In principle the kinetic energy density varies across any given cell. However, to anticipate the next subsection, it will be necessary to reduce the FE mesh in the FE/MD handshake region to coincide with the perfect atomic lattice. At the atomic level each atom has kinetic energy and is apportioned accordingly. Thus for our FE mesh, we chose to use the ‘‘lumped-mass’’ approximation, which reduces for the smallest mesh size to the atomic limit. One third of the mass in each cell is apportioned to each apex. The kinetic energy is thus given by

$$K_{\text{FE}} = \sum_{t=1}^{N_{\text{mesh}}} M^t (\dot{\mathbf{u}}^t)^2,$$

$$M^t = \rho L \sum_{m=1}^{N_{\text{cell}}} \sum_{l=1}^3 \delta_{lm_l} \frac{A_m}{3}, \quad (14)$$

where t labels the FE mesh points, of which there are N_{mesh} total, and m_l labels the mesh point index at each of the three apices of cell m . The $\{\dot{\mathbf{u}}\}$ are vectors of length two since they relate to a mesh point.

Forces for the displacement degrees of freedom in Eq. (10) are obtained by taking positional derivatives. Displacements and their time rate of change may be obtained as a function of time, for given boundary conditions, using the same update algorithm, Eq. (3), and time step as that used for the MD and TB. The force due to the m th cell is

$$\mathbf{f}_{\text{FE}}^m = [\mathbf{K}]^m \mathbf{u}^m. \quad (15)$$

In keeping with Eq. (10), the vectors in Eq. (15) are of length six. The total force associated with a mesh point is then the sum of the contributions from each of the cells with apices in common with that point. Finally

$$\mathbf{f}_{\text{FE}}^t = M^t \ddot{\mathbf{u}}^t. \quad (16)$$

D. Handshaking FE/MD

The two principal issues for handshaking between FE and MD regions are (a) the overlap of the mesh with the atoms and (b) the form of the Hamiltonian.

In the first case, we generalize an idea due to Kohlhoff and coworkers.²⁸ An imaginary surface is drawn between the FE and MD regions. Within the range of the MD interatomic potential from this surface, FE mesh points are located at ideal lattice sites. As long as there is no diffusion, atoms and mesh points remain on either side of this interface. However, the distinction between the two becomes academic: atomic motion may be viewed as displacement around a lattice (mesh) site, and the displacement field may be viewed as motion of an atom away from its perfect site. This methodology would work also for amorphous systems; all that is required is a one-to-one mapping of a mesh point to an atom site. Moving away from the handshake region into the FE region, the mesh spacing may be made larger. This is the principal reason that the FE algorithm is computationally efficient. The largest spacing depends upon the physics that we wish to capture; for example the largest spacing determines the shortest wavelength phonons that we wish to propagate unimpeded through the FE region.

In our illustration of the CLS method, we chose to examine brittle fracture in silicon; specifically we oriented our rectilinear system such that it has (100) faces on all sides. The FE region is represented as a 2D system, which nevertheless handles the third dimension in plane strain; Eq. (11) has the parameter L to represent the thickness of the sample. Thus, the projection of a diamond cubic lattice onto a (100) plane is required to generate the 2D mesh. Figure 2 shows this, as well as the triangulation we used for the FE cells. At the edges of the computational cell, the triangulation is able to wrap around; i.e., the FE region can be made periodic in

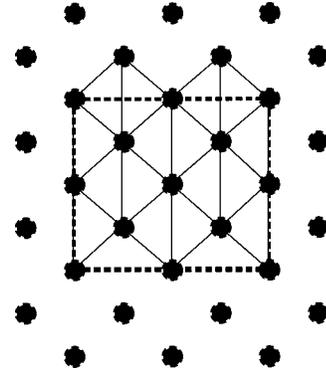


FIG. 2. Triangulation of 2D unit cell. Unit cell denoted by dashed lines. Triangles marked with continuous lines. Note continuous lines on left and right but not top and bottom of unit cell boundary. Eight mesh points and 16 triangles per unit cell.

exactly the same way as the MD. For a periodic system, there are twice as many cells as there are mesh points. Away from the handshake region and into the FE region, we expanded the mesh along one dimension (the long axis in Fig. 1) while keeping the mesh spacing constant in the second. The function chosen for this expansion was of hyperbolic tangent form. Thus near the handshake region there is no expansion of the atomic mesh and far away from the handshake region the spacing asymptotes to a constant multiple of the atomic lattice parameter. In the crack propagation example below, this multiple was ten. The transition region spanned a couple of hundred Angstroms.

Turning now to the form of the handshake Hamiltonian, where we differ from the prior work of Kohlhoff and co-workers²⁸ is in the dynamics of the handshake region. We found it very important to define a conservative Hamiltonian so as to ensure symplectic time evolution of the atomic and displacement trajectories within the handshake region. This is key to the successful implementation of the *dynamic* coupling of length scales. We note that it is academic whether to say either that atom sites propagate into the FE region or that FE mesh points propagate into the MD region. The discussion is best understood with reference to Fig. 3. In conceptualizing this Hamiltonian, we imagine that two different materials sit on either side of an interface; in one case it is FE silicon and in the other it is SW silicon. The cross terms (i.e., the handshake Hamiltonian) to first order can be approximated by a mean of the two descriptions. All FE triangles that cross the interface contribute half their weight to the Hamiltonian. Any triangle that is fully in the MD region contributes zero weight. Similarly, any SW interaction (two body or three body) which crosses the interface contributes half its usual weight. Any SW interaction between mesh points, all of which are fully on the FE side of the interface, contributes zero weight. The SW energy formulation that concentrates upon atomic coordinates $\{\mathbf{r}\}$ and the FE energy formulation that concentrates upon displacements $\{\mathbf{u}\}$ can be used throughout the interface because of the indistinguishability of what are atoms or mesh points. The one-to-one mapping of atoms to nodes is not required at distances greater than twice the SW pair cutoff away from the interface in the FE region. This is the distance of greatest three-body range. Figure 3 indicates diagrammatically those interactions that contribute to the handshake Hamiltonian. Thus

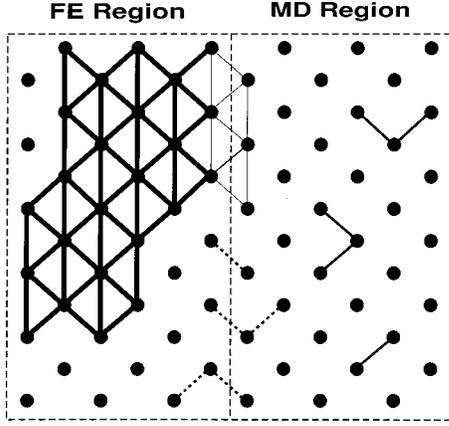


FIG. 3. Illustration of FE/MD handshake Hamiltonian. FE cells contributing fully to overall Hamiltonian (unit weight) marked with heavy lines. FE cells contributing to handshake Hamiltonian (half weight) represented by light lines. Two and three-body terms (dotted lines) of SW interaction that cross boundary also carry half weight. Continuous lines represent unit weight SW interactions. Representative examples in each case.

$$V_{\text{FE/MD}} = \frac{1}{4} \sum_{m^I=1}^{\#\text{cross}} \sum_{p,q=1}^6 u_p^{m^I} K_{pq}^{m^I} u_q^{m^I} + \frac{1}{2} \left[\sum_{(i<j)^I} v^{(2)}(\mathbf{r}_{(ij)^I}) + \sum_{[i,(j<k)^I]} v^{(3)}(\mathbf{r}_{(ij)^I}, \mathbf{r}_{(ik)^I}) \right]. \quad (17)$$

Here, the superscript I implies those interactions that cross the FE/MD boundary. Indeed, $V_{\text{FE/MD}}$ is only defined for interactions that cross the boundary. Otherwise, the other terms in Eq. (1) define the forces at mesh points and atoms. The formalism in Eq. (17) is meant to imply that *any* one atom of the triplet in the three-body terms can be on an opposite side of the interface to the other two.

In the present implementation of the CLS algorithm, the FE is 2D and the MD is 3D. As we have said, the third dimension of the FE region is treated in mean field. Thus, in $V_{\text{FE/MD}}$, x and y displacements of atoms on the MD side of the FE/MD boundary that contribute to the elastic energy are obtained by averaging over all equivalent atoms in the depth z . In similar fashion, in determining the SW energy contribution to the handshake Hamiltonian, we replicate all x and y displacements in the third dimension on the FE side of the boundary by assuming that atoms sit at ideal lattice sites in that dimension. The overall Hamiltonian remains conservative.

In making the FE/MD interface seamless, two other issues confront the definition of energy. They both involve reference state; one is potential energy, the other is thermal energy. The SW potential is referenced to infinitely separated atoms. The FE potential is referenced to a $T=0$ unstrained lattice. For the purposes of graphical analysis, therefore, a constant offset energy that does not affect the dynamics is added to each FE mesh point. The $T=0$ energy density for SW silicon at zero pressure is -4.33444 eV/atom. The offset energy is computed for every FE point using an equation entirely analogous to that used to compute mass in the ‘‘lumped-mass’’ approximation except that instead of a mass

density, we now use the SW energy density [see Eq. (14)]. As before, this scheme ensures the correct limiting behavior as the mesh spacing is reduced to atomic dimensions. For atoms in the handshake region, for systems with unusual orientation where the offset is nontrivial to estimate atom by atom, a $T=0$ calculation with zero strain for the coupled FE/MD system can be performed. The offset may thereby be calculated to maintain the energy/atom constant through the interface. This is easily achieved by virtue of periodicity and symmetry.

Turning now to the thermal energy, the work of Rudd and Broughton²⁹ indicates that the square of the time rate of change of the displacements in the FE region are related to the temperature. In coarsening the FE grid upon moving away from the handshake region, atomic degrees of freedom are lost: the FE algorithm involves an average over these. Thus, to bring the atomic and continuum thermal energies onto an equivalent footing, the total FE thermal energy may be written again by means of an offset. These corrected energies are denoted by a prime

$$K'_{\text{FE}} = \frac{3}{2} (N_{\text{atom}} - N_{\text{mesh}}) k_B T + K_{\text{FE}} + \frac{1}{2} N_{\text{mesh}} k_B T, \quad (18)$$

$$V'_{\text{FE}} = \frac{3}{2} (N_{\text{atom}} - N_{\text{mesh}}) k_B T + V_{\text{FE}} + \frac{1}{2} N_{\text{mesh}} k_B T.$$

N_{atom} is the number of atoms contained within an equivalent 3D volume. k_B is Boltzmann’s constant. Equipartition has been invoked. We further assume that the background temperature does not vary during the simulation. The first term therefore accounts for the missing degrees of atomic freedom while the last term augments the 2D FE plane-strain simulation for the missing third dimension in its degrees of freedom. As before, these offsets do not affect the dynamics of the system and the thermal corrections can be apportioned to each mesh point in like manner to that described above for the zero-temperature FE potential energy. For finite temperature simulations, the $\{\mathbf{u}\}$ degrees of freedom are thermalized to a Maxwellian distribution. Also, the appropriate elastic constants for that temperature should be used in the FE equations of motion so as to make the MD and FE regions seamless and compatible. Further, since this methodology requires a continuation of ideal lattice sites into the FE/MD handshake region so as to determine mesh coordinates, the appropriate lattice parameter for given temperature should be used.

Lastly, in this subsection, we discuss the issue of dissipation in the FE region. The continuum representation of matter used here involves linear elasticity theory. This is a harmonic theory. Thus, vibrational modes of given $\{k, \omega\}$ relationship, which depend upon the long-wavelength elastic constants of the medium, propagate undamped. In order to thermalize short-wavelength phonons propagating through regions where the mesh spacing changes and also to allow energy to be dissipated in the FE region, the FE degrees of freedom are weakly coupled to a Brownian heat bath whose dynamics are set to the temperature at which the simulation is being performed. In effect, this couples the phonon modes of the FE region. The force used in the third step of the velocity Verlet algorithm [see Eq. (3)] now includes random and dissipative terms

$$\mathbf{f}_{\text{FE}}^t = -\frac{\partial V_{\text{FE}}}{\partial \mathbf{u}^t} + \rho^G(T, y^t) - \xi(y^t) M^t \dot{\mathbf{u}}^t, \quad (19)$$

where ρ^G is a Gaussian random variable and ξ is a friction coefficient. The two are related by the fluctuation dissipation theorem. Specifically the variance σ^G of the Gaussian is given by:

$$\sigma^G = \sqrt{\frac{2\xi M^t k_B T}{\Delta t}}. \quad (20)$$

In order to perturb the dynamics of the active zone (i.e., MD and TB) minimally, ξ was made a function of the (time invariant) FE mesh y coordinate. ξ was linearly ramped from zero in the handshake region to finite value (we chose 0.1 for the crack propagation study) at the extremal outer edge of the FE regions.

E. Handshaking MD/TB

In contradistinction to the FE/MD handshake algorithm, where a plane *between* rows of atoms was defined, the MD/TB handshaking takes place conceptually across a plane *consisting* of atoms. This different approach is necessitated because it is difficult to apportion (localize) energy in a computationally efficient way to specific bonds in an electronic structure calculation. The total energy is a property of the entire system. Attempts to define a 50/50 Hamiltonian, such as was used for the FE/MD interface, run into issues of orthonormality such that Eq. (7) is no longer sufficient: derivatives of the electronic coefficients with respect to the atomic coordinates are required and these must be computed numerically.

The ensuing description of the handshake Hamiltonian is best understood with reference to Fig. 4. A principal reason for choosing a semiconductor like silicon to demonstrate the CLS algorithm is that covalent bonds are local objects. ‘‘Dangling’’ bonds may be ‘‘tied off’’ with univalent atoms. Simply put, the region chosen for TB description is terminated with univalent atoms. We call these ‘‘silogens’’ to represent the fact that they behave like monovalent (i.e., hydrogenic) silicon atoms. The $[\mathbf{H}]$ and $[\mathbf{S}]$ matrix elements and the repulsive pair potential v^{rep} that couple these atoms to the silicon atoms within the interior of the TB region were chosen to (a) maintain electroneutrality (as measured by Mulliken charges) on both the silogens and silicons, (b) locate the silogen potential energy minimum at a Si-Si distance, *not* a Si-H distance (c) provide a bond energy equal to a single Si-Si bond and (d) provide a longitudinal force constant equal to that of silicon. At the perimeter of the TB region, the silogens are constrained to sit at the silicon sites of the MD region. In many cases, this means that more than one silogen sits at a given site. [At a Si(100) surface, terminating dangling bonds with silogens necessitates placing two at each empty silicon site.] Thus, there are no matrix elements nor v^{rep} terms that couple *any* of the silogens to one another. Operationally, a circle is drawn around an inner set of atoms; these are designated TB silicons. Then any atom outside this circle, but within range of an inner atom, is designated as a silogen. The range criterion used was the mean of the first (r_0) and second neighbor $[\sqrt{(8/3)}r_0]$ distances of the equilibrium silicon lattice.

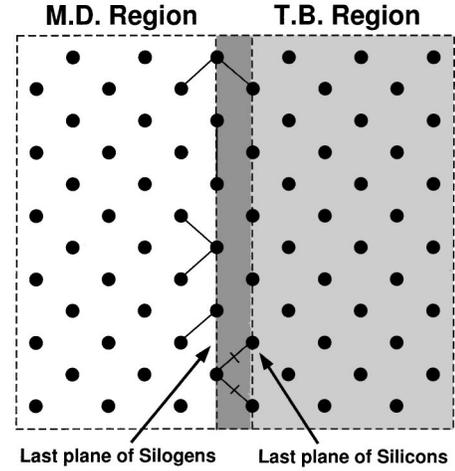


FIG. 4. Illustration of MD/TB handshake Hamiltonian. Outer perimeter of TB region terminated with monovalent silogens (see text) constrained to sit at silicon sites. TB Hamiltonian diagonalized for sum of light gray plus dark gray regions: Si-Si matrix elements employed in light gray area, Si-Silogen matrix elements used in dark gray region. Two- and three-body SW interactions contributing to handshake Hamiltonian designated by full lines. Broken lines represent noncontributing SW three-body term. Only representative SW examples are shown.

In Fig. 4, matrix elements that couple atoms across the light gray region are of Si-Si form. Atoms coupled across the dark gray region use Si-Silogen matrix elements. Parameters of the former are given by Bernstein and Kaxiras.¹⁷ The latter, using the same formalism, are

$$\epsilon_s = -7.661518eV,$$

$$V_{ss\sigma} = -1.6967418eV, \quad (21)$$

$$V_{sp\sigma} = 3.8704886eV,$$

$$S_2(r_{ij}) = \frac{S_{ss\sigma}(r_{ij}) - \sqrt{3}S_{sp\sigma}(r_{ij})}{2}.$$

Bernstein and Kaxiras¹⁷ provide a prescription for relating the S overlap and the V terms above. The TB $[\mathbf{H}]$ and $[\mathbf{S}]$ generalized eigenvalue problem is solved for the entire silicon plus silogen system. Forces are extracted as defined by Eq. (7).

All that remains is to determine which SW two- and three-body terms are required to couple the silogen atoms to the MD region. Examples are shown in Fig. 4. Since silogens are not coupled to one another in the TB region, SW terms that account for such are required. All SW pair terms between a silogen-silicon and either a silicon atom in the MD region or another silogen-silicon are included. All SW triplets that include at least one MD silicon to one silogen-silicon pair are also included. The forces that arise on the silogen-silicons from these terms are added to the forces arising from the TB Hamiltonian on these atoms. Thus, Eq. (1) should more correctly be written

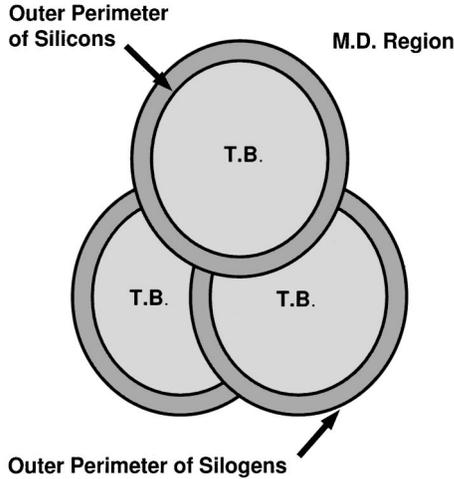


FIG. 5. Illustration of “clover-leaf” overlapping TB regions embedded in MD region. Atomic force is function of overlap (see text).

$$\begin{aligned}
 H_{\text{Tot}} = & H_{\text{FE}}(\{\mathbf{u}, \dot{\mathbf{u}} \in \text{FE}\}) + H_{\text{FE/MD}}(\{\mathbf{u}, \dot{\mathbf{u}}, \mathbf{r}, \dot{\mathbf{r}} \in \text{FE/MD}\}) \\
 & + H_{\text{MD}}(\{\mathbf{r}, \dot{\mathbf{r}} \in \text{MD}\}) + H_{\text{MD/TB}}(\{\mathbf{r}, \dot{\mathbf{r}} \in \text{MD/TB}\}) \\
 & + H_{\text{TB}}(\{\mathbf{r}, \dot{\mathbf{r}} \in \text{MD/TB}\}), \quad (22)
 \end{aligned}$$

where the penultimate term involves only SW interactions crossing the boundary and the last involves a TB calculation for the combined silicon plus silogen system.

In the absence of a dynamic allocation of the TB region to those parts of the system where bonds are breaking during a simulation, the above prescription produces a conservative Hamiltonian. Unfortunately, for many systems (such as the crack propagation example below), the hundred or so atoms, whose forces may currently be updated using a non-orthogonal TB Hamiltonian in one second of wall-clock time, do not comprise as large a region as we might wish. Part of the problem may be solved using periodic boundary conditions. Our example of crack propagation used a slab two unit cells deep; thus, the TB region is actually a cylinder. There may be other systems, such as a void within the bulk of a material, where periodicity is not appropriate and a spherical region must be used. Another part of the problem can be ameliorated by using more than one processor per TB region to perform the diagonalization, but unfortunately such algorithms presently are not efficient on coarse-grained scalable architecture computers.

Instead, we chose to represent the region of breaking bonds by a “clover leaf” of TB regions. Figure 5 gives an illustration of three overlapping TB regions. In the crack propagation simulation, described later, eight overlapping regions were used. Each of these regions is diagonalized separately. Each is handled by a separate processor. As in Fig. 4, the inner TB couplings are denoted by light gray while the matrix elements coupling TB silicons to TB silogens are represented by dark gray. After forces on each atom are obtained for each TB region separately, the force to be used in the velocity Verlet [see Eq. (3)] update, is obtained via an average over the different regions: where there is no overlap of TB regions, use the same prescription as for a single TB region; where light gray silicons overlap, use the mean

value; where a silogen of one TB region overlaps a silicon of another, use the silicon value. The number of atoms that are propagated using TB forces is therefore less than the total number within all the “clover-leaf” regions. These rules are intuitive and although it is not now possible to write an overall conservative Hamiltonian for the entire system, we have found from experience that the atomic trajectories (as measured by the lack of anomalies in local kinetic energies) are well behaved. For a system without periodicity, such as the decoration of TB regions around a void, the “clover-leaf” metaphor can easily be generalized to something akin to a “raspberry.”

Finally in this subsection, the allocation algorithm is described: that is the algorithm whereby the TB “clover leaf” is made to track a region of breaking bonds. The energy and force algorithm *as implemented* in the MD and TB regions proceeds by calculating the SW energy for *all* atoms in the MD processors. TB processors calculate not only TB energies and forces but also those SW forces that must be subtracted from those double counted in the MD processors. The result is that SW energies, by suitable apportioning of two- and three-body terms, are available for all atoms. These are then used to discriminate different regions. The apex of a crack is found, for example, by locating the atom, with a potential energy greater (more positive) than 60% of the bulk cohesive potential energy, furthest to the left or right of the center of the system (see Fig. 6). The central TB region of the “clover leaf” is then placed at that atom. Such placement does not have to be performed at every time step—in our implementation, it is done every ten steps.

F. Seamless FE/MD/TB

The foregoing discussion indicates how the simulation can be made seamless. The TB region, since it is the region described at the most microscopic level, should determine the elastic constants and the atomic force fields used elsewhere in the system. Thus, firstly, a pure TB simulation is performed for a small number of atoms representing the bulk system (at given temperature and pressure). By appropriate deformation of the computational cell, the elastic constants are extracted. By movement of one atom within the cell, a local “Einstein oscillator” force constant can be found. The SW parameters for silicon may then be adjusted to reproduce the same quantities. The elastic constants from the TB region are also used for the $[\mathbf{K}]$ stiffness matrix of the FE region. Lastly, the parameters used for the Si-Silogen matrix elements are adjusted so that displacement of a silogen-silicon in the coupled system gives rise to the same “Einstein oscillator” force constant of the pure bulk system. Work is presently proceeding to fulfill this prescription for full seamless integration.

In the crack propagation example given below the elastic constants of the SW and FE region are made identical. However, the SW and TB elastic constants are close but slightly different. Nevertheless, the results indicate that even here, the seamless objective is close to reality.

III. RESULTS

Our application is the rapid brittle fracture of a silicon slab flawed by a microcrack at its center and under uniaxial

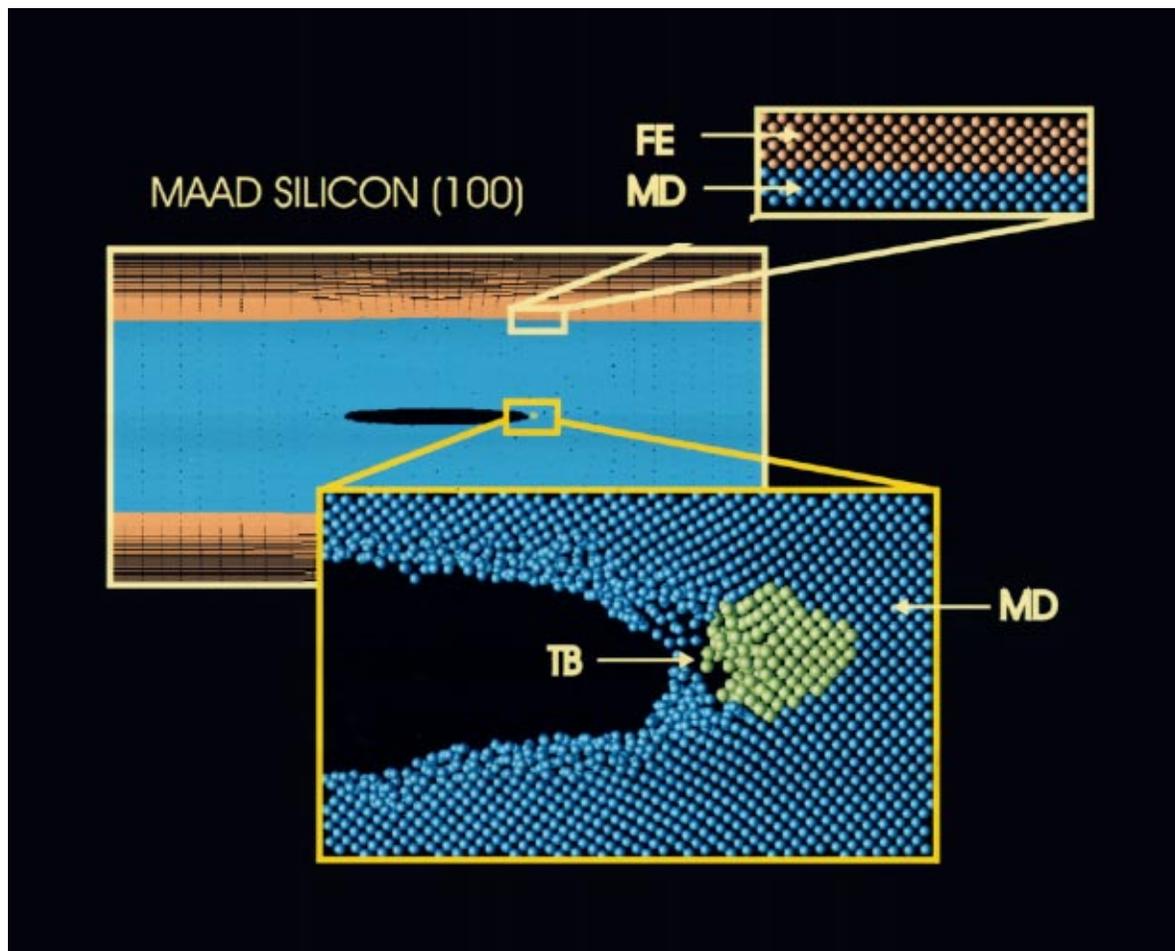


FIG. 6. (Color) The geometrical decomposition of the silicon slab into the five different dynamic regions of the simulation: the continuum finite-element region (FE); the atomistic molecular-dynamics region (MD); the quantum tight-binding region (TB); the FE-MD ‘‘handshaking’’ region; and the MD-TB ‘‘handshaking’’ region. The image is the simulated silicon slab, with expanded views of the FE-MD (orange nodes-blue atoms) interface and the TB (yellow atoms) region surrounded by MD (blue) atoms. Note that the TB region surrounds the crack tip with broken-bond MD atoms trailing behind this region. The acronym, MAAD, implies ‘‘macroatomistic *ab initio* dynamics.’’

tension. This example was chosen to both illustrate and validate our CLS scheme. Figure 6 shows the geometrical decomposition of the silicon slab into the five different dynamic regions of the simulation. The MD region was spatially domain decomposed onto 24 processors. Each FE region was handled by its own processor. We tracked the path of the crack and placed the center of the TB region at the apex of the crack. This is where bond breaking occurs; it is the region that is crucial to determining the kinetics of the crack propagation process. For the extended regions of bond rupture (see Fig. 6), we used a ‘‘clover leaf’’ of eight overlapping TB regions, each being cylindrical and distributed to a different processor. The exposed notch faces were x - z planes with (100) faces, with the notch pointed in the (010) direction. There were 258 048 mesh points in each FE region, 1 032 192 atoms in the MD region, and around 280 unique atoms in the TB region. Each of the eight TB regions is a cylinder with radius of 5.43 Å in the yz plane. The lengths of the MD region were 10.9 Å (the slab thickness and periodic), 521 Å (before the pull, in the direction of pull), and 3,649 Å (the primary direction of propagation and periodic). The full pull length of the FE+MD system was

5,602 Å. The entire system including the FE represented 11 093 376 atoms. The time for a TB force update was 1.5 s, that for the MD update was 1.8 s and that of the FE was 0.7 s. We could, thus, afford to double the size of the FE region in order to accomplish complete computational load balancing but without any sacrifice of wall-clock time. The TB

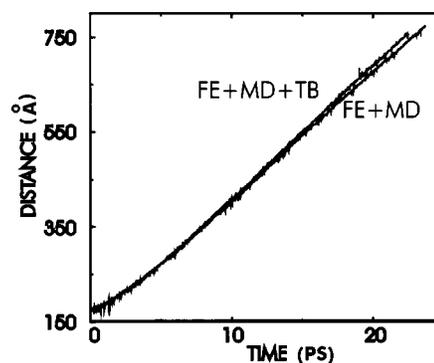


FIG. 7. The distance versus time history of the two crack tips, one having TB atoms always centered at the immediate failure region. The distance is in Angstroms, and the time is in picoseconds.

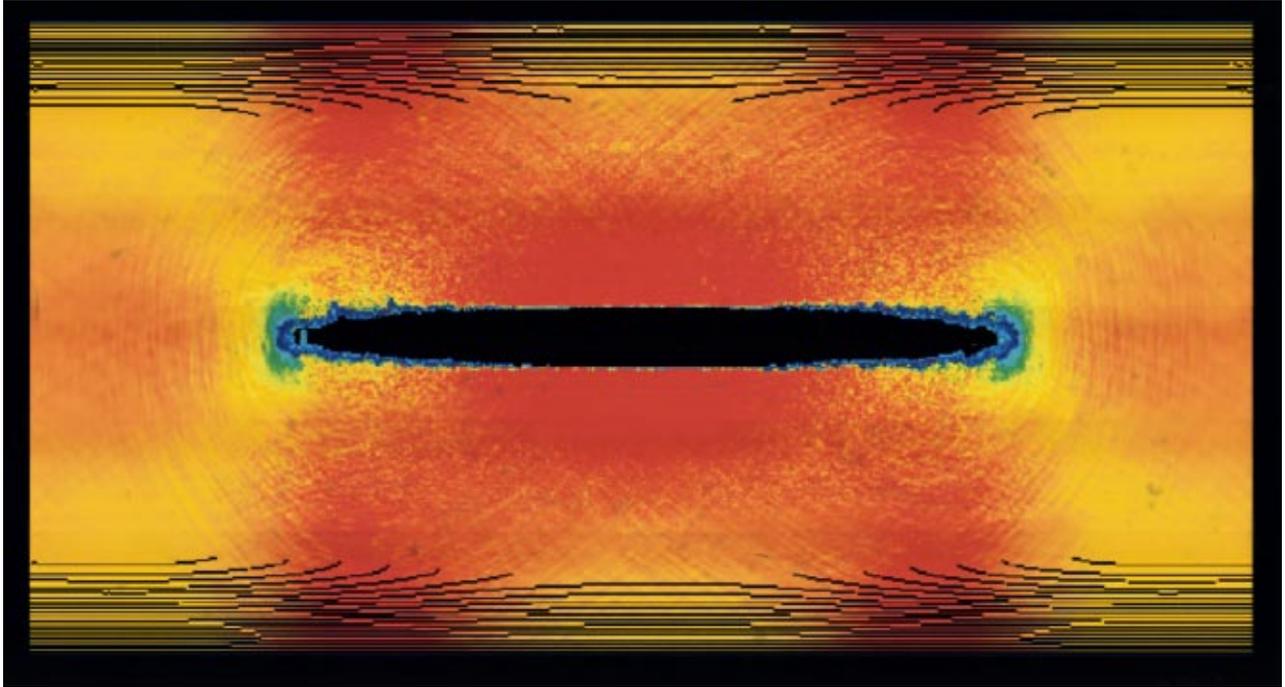


FIG. 8. (Color) Stress waves propagating through the slab using a finely tuned potential energy color scale at a point in time after the asymptotic crack speed has been achieved. Blue represents high stress; red represents low stress. Yellow is intermediate.

region was relocated every 10 time steps.

The rectilinear computational cell comprised (100) faces on all sides. The reduced elastic constant matrix [see Eq. (11)] for this geometry was obtained by averaging the results of Balamane, Halieioğlu, and Tiker²⁶ and Ray²⁷ for the zero-temperature C_{11} , C_{12} , and C_{44} elastic constants of SW silicon

$$[C] = \begin{pmatrix} 1.578 \times 10^6 & 0.7930 \times 10^6 & 0.0 \\ 0.7930 \times 10^6 & 1.578 \times 10^6 & 0.0 \\ 0.0 & 0.0 & 0.6365 \times 10^6 \end{pmatrix}, \quad (23)$$

where the units are megabar.

The slab was initialized at zero temperature, and a constant strain rate was imposed on the outermost FE boundaries defining the opposing horizontal faces of the slab. Further, a linear velocity gradient was applied within the slab, which resulted in an increasing internal strain with time. The solid failed at the notch tip when the solid had been stretched by $\sim 1.5\%$. The imposed strain rate was set to zero at the onset of crack motion. Figure 7 presents the distance versus time history of the two crack tips, one having the TB atoms dynamically centered at the immediate failure region. The propagating cracks rapidly achieved a limiting speed (2770 m/s) equal to 85% of the Rayleigh speed, the sound speed of the solid silicon surface. The two distance-time histories are very similar. In hindsight, this might have been expected since the elastic modulus of silicon calculated from the empirical SW potential and from TB are very similar up to the mechanical stability limit of the bulk solid. More importantly, this indicates that the handshaking between the MD region and the TB region was reliable.

A more powerful signature of seamless coupling, one which represents a validation of the method, is depicted in Fig. 8. We note that stress waves pass from the MD region to the FE regions with no visible reflection at the FE/MD interface; i.e., the coupling of the MD region to the FE region appears seamless. Further, there are no obvious discontinuities at the MD/TB interface; this observation remains true even at higher spatial magnifications. A discussion of the fracture physics, outside the purpose of the present paper, has been presented in an earlier publication.^{5,6}

IV. CONCLUSIONS

We have described an algorithm that successfully couples length scales. It is a finite temperature, dynamic, parallel algorithm. The general approach is applicable to many material types. The applications of such a methodology will be many and varied. The work described above represents an illustration of what is possible, but developmental work still remains which we now briefly discuss:

More thought will be required for treating metals since (a) bonds are less localized than in silicon and (b) termination of the surface ‘‘dangling bonds’’ of the active TB region with monovalent species is no longer a good approximation. As we have said, this methodology is not wedded to TB: other fast quantum mechanical formulations, as long as they pass the one second of wall-clock time per time step test, may be more appropriate. Indeed, as computational hardware performance improves, the choice of which quantum-mechanical scheme to place at the heart of this CLS scheme will almost certainly change. An example might be the kinetic energy functional³¹ local-density approximation methods advocated by Carter, Madden, and co-workers,³⁰ which require no diagonalization and which parallelize efficiently. Such

schemes would require no “clover-leaf” overlapping of quantum mechanical regions as we have, out of necessity, implemented here. A further difficulty for metals involves thermal conductivity and dissipation through mixing of electronic states at the Fermi level; something that an empirical potential cannot capture. Perhaps the empirical potential region will have to be augmented with auxiliary degrees of freedom to enable correct handshaking conditions for this added level of complexity.

The handshaking between each neighboring region works well under the conditions of our example. We expect long-wavelength phonons to propagate with minimal back scattering through the MD/FE interface. However, in systems where significant short-wavelength energy is emitted from the central region, we expect that the present handshaking methodology will have to be augmented. The reason is that the shorter wavelength, high-frequency vibrations cannot be supported by the larger mesh spacings inherent in the outer part of the FE region and will be scattered back into the MD region. Such may be mitigated by adding a random and dissipative heat bath to the FE degrees of freedom. Other schemes^{29,32} to mitigate these effects including adding Nose-Hoover chain thermostats³³ to the FE mesh and developing a coarse grained MD formulation for the handshake region, are being pursued. Indeed, such coarse graining³² can be made to go beyond linear elasticity so that (a) dissipation occurs naturally and (b) constant pressure algorithms in the far-field regions are viable.

The present algorithm dynamically tracks the crack tip with a TB region. Our example of brittle fracture in silicon is appropriate for demonstration purposes since it is clear where the TB region should be placed. In the case of dislocation generation at the tip, indicative of ductility and typical

in metals, a dynamic allocation of processors as and when required to focus on these areas will be required. This is the subject of on-going work.

Turning now to long term goals, in our present formulation we require an MD/FE handshake region in which there is no diffusion and in which no defects propagate. It may well be that, as research on CLS gains momentum, a hybrid of the present formulation and the adaptive grid work of Tadmor, Ortiz, and Phillips³ (discussed in the introduction) will evolve. Such would be a powerful computational tool.

Finally, as a philosophical observation, we note that the algorithm we have espoused links not only length scales but also disciplines. The TB region was employed to study breaking bonds; that is the realm of chemistry. The MD region was implemented to describe the statistical mechanics of the system; this is the forte of physics. And the FE region that was used to couple the atomistics to the rest of the universe is the traditional methodology of engineers. This marriage of disciplines and its concomitant negation of traditional barriers may well prove to be the true power of this approach. This formalism represents the beginnings of computational atomistic engineering.

ACKNOWLEDGMENTS

We acknowledge support from the USAF Maui High-Performance Computing Center, which is administered by the University of New Mexico. F.F.A. thanks H. Gao and J.Q.B. thanks R. E. Rudd for discussions. J.Q.B. wishes to acknowledge support of the ONR and DARPA and to thank DOD’s HPCMO for a support. Lastly, we all wish to thank NSP and the ITP at Santa Barbara for support and hospitality during the spring of 1997.

-
- ¹E. Clementi, Philos. Trans. R. Soc. London, Ser. A **326**, 445 (1988).
- ²H. Elbern, Atmos. Environ. **31**, 3561 (1997).
- ³E. B. Tadmor, M. Ortiz, and R. Phillips, Philos. Mag. A **73**, 1529 (1996).
- ⁴H. Raffi-Tabar, L. Hua, and M. Cross, J. Phys.: Condens. Matter **10**, 2375 (1998).
- ⁵F. F. Abraham, J. Q. Broughton, N. Bernstein, and E. Kaxiras, Europhys. Lett. **44**, 783 (1998).
- ⁶F. F. Abraham, J. Q. Broughton, N. Bernstein, and E. Kaxiras, Comput. Phys. **12**, 538 (1998).
- ⁷F. F. Abraham, D. Schneider, B. Land, D. Lifka, J. Skovira, J. Gerner, and M. Rosendrantz, J. Mech. Phys. Solids **45**, 1461 (1997).
- ⁸M. E. Bachlechner, A. Omeltchenko, A. Nakano, R. K. Kalia, P. Vashishta, I. Ebbsjo, A. Madhukar, and P. Messina, Appl. Phys. Lett. **72**, 1969 (1998).
- ⁹F. F. Abraham and J. Q. Broughton, Comput. Mater. Sci. **10**, 1 (1998).
- ¹⁰F. F. Abraham, D. Brodbeck, R. Rafey, and W. E. Rudge, Phys. Rev. Lett. **73**, 272 (1994).
- ¹¹M. Z. Bazant, E. Kaxiras, and J. F. Justo, Phys. Rev. B **56**, 8542 (1997).
- ¹²F. H. Stillinger and T. A. Weber, Phys. Rev. B **31**, 5262 (1985).
- ¹³W. M. C. Foulkes and R. Haydock, Phys. Rev. B **39**, 12 520 (1989).
- ¹⁴P. Alinaghian, P. Gumbsch, A. J. Skinner, and D. G. Pettifor, J. Phys.: Condens. Matter **5**, 5795 (1993).
- ¹⁵M. Aoki, Phys. Rev. Lett. **71**, 3842 (1993).
- ¹⁶M. Tuckerman, B. J. Berne, and G. J. Martyna, J. Chem. Phys. **97**, 1990 (1992).
- ¹⁷N. Bernstein and E. Kaxiras, Phys. Rev. B **56**, 10 488 (1997).
- ¹⁸F. S. Khan and J. Q. Broughton, Phys. Rev. B **39**, 3688 (1989).
- ¹⁹X.-P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10 891 (1993).
- ²⁰M. S. Daw, Phys. Rev. B **47**, 10 895 (1993).
- ²¹P. Ordejon, D. A. Drabold, R. M. Martin, and M. P. Grumbach, Phys. Rev. B **51**, 1456 (1995).
- ²²F. Mauri, G. Galli, and R. Car, Phys. Rev. B **47**, 9973 (1993).
- ²³C. S. Jayanthi, S. Y. Wu, J. Cocks, Z. L. Xie, M. Menon, and G. Yang, Phys. Rev. B **57**, 3799 (1998).
- ²⁴T. J. R. Hughes, *The Finite Element Method* (Prentice-Hall, Englewood Cliffs, NJ, 1987).
- ²⁵See *Modeling, Mesh Generation, and Adaptive Numerical Methods for Partial Differential Equations*, edited by I. Babuska, J. E. Flaherty, W. D. Henshaw, J. E. Hopcroft, J. E. Oliger, and T. Tezduyar, The IMA Volumes in Mathematics and its Applications Vol. 75 (Springer-Verlag, Berlin, 1995).

- ²⁶H. Balamane, T. Halicioglu, and W. A. Tiller, *Phys. Rev. B* **46**, 2250 (1992).
- ²⁷J. R. Ray, *Comput. Phys. Rep.* **8**, 109 (1988).
- ²⁸S. Kohlhoff, P. Gumbsch, and H. F. Fischmeister, *Philos. Mag. A* **64**, 851 (1991).
- ²⁹R. E. Rudd and J. Q. Broughton, *Phys. Rev. B* **58**, 5893 (1998).
- ³⁰S. Watson, B. J. Jesson, E. A. Carter, and P. A. Madden, *Europhys. Lett.* **41**, 37 (1998).
- ³¹L.-W. Wang and M. P. Teter, *Phys. Rev. B* **45**, 13 196 (1992).
- ³²R. E. Rudd and J. Q. Broughton (unpublished).
- ³³G. J. Martyna, M. L. Klein, and M. Tuckerman, *J. Chem. Phys.* **97**, 2635 (1992).