9 SCIENTIFIC HIGHLIGHT OF THE MONTH

Learn on the fly": a multiscale hybrid simulation method for material systems

Gábor Csányi[1], T. Albaret[2], M. C. Payne[1], A. De Vita[3,4]

[1] *Cavendish Laboratory, Madingley Road, Cambridge, CB3 0HE, United Kingdom*

[2] *Laboratoire de Physique de la Matiere Condensee et Nanostructures, Bat. Leon Brillouin, Campus de la Doua, Universite Claude Bernard Lyon 1, 43 Bld du 11 novembre 1918, 69622 Villeurbanne Cedex*

[3] *Physics Department, King's College London, Strand, London WC2R 2LS, United Kingdom*

[4] *INFM–DEMOCRITOS National Simulation Center and Center of Excellence for Nanostructured Materials (CENMAT), University of Trieste, Italy*

Abstract

We present an overview of a recently introduced hybrid molecular dynamics simulation scheme that links arbitrary atomistic techniques together in a truly seamless manner. Rather than constructing a new hybrid Hamiltonian that combines different models, we use a unique short–range classical potential and continuously tune its parameters to reproduce the atomic trajectories at the prescribed level of accuracy throughout the system.

This article is an extract from a Topical Review submission to Journal of Physics: Condensed Matter

1 Introduction

Natural phenomena occur on a variety of lengthscales, which not only roughly define scientific disciplines (physics and chemistry for the very small, biology for intermediate, geology for very long times and large lengthscales, and cosmology for the largest), but also delineate subfields within a given discipline, due to the very different experimental methods and theoretical models that are applicable at each scale. These form a multiscale hierarchy, in which parameters of larger scale models are measured or calculated on a smaller scale[1].

Yet in a large class of problems, the lengthscales cannot be separated in this way, the coupling between them is strong, and "bidirectional". This often happens when the microscopic phenomena are driven by some macroscopic force, so there is a complete feedback loop. Stress induced defect processes in solids are a good example, and brittle fracture is the prototypical problem.

In molecular biology, catalysis is often controlled by large–scale motion of macromolecules. If we wish to simulate these processes, we have to do that on more than one lengthscale *simultaneously*. Crucially, one cannot use the smallest–scale model to simulate the entire system, because it would be too expensive and hugely wasteful. A typical quantum mechanical molecular dynamics simulation can deal with hundreds of atoms, while the minimum system size that can capture the larger scale aspects can reach into hundreds of thousands of atoms. In the past decade there has been a growing effort to devise so–called *hybrid* simulation techniques, that seamlessly integrate a wide variety of different models, ranging from first principles methods to finite elements techniques, into a concurrent simulation[2, 3].
A pioneering work in this field was the *quasicontinuum method* [2], which successfully linked classical atomistic and continuum elasticity models. Here we concentrate on a lower level linkage: that of quantum mechanical to classical modelling. After mentioning general issues related to hybrid atomistic modelling, we discuss our scheme which was proposed in [4]. Finally, we present several validation tests and an example application to the brittle fracture in silicon.

2 Hybrid schemes

The objective is the following: given a large atomic system, some (perhaps small) regions of it need to be simulated with model A (which we take to be quantum mechanical), and the rest with model B (which we take to be classical atomistic). The major issues that have to be addressed are as follows.

- **Handshaking** The two models need to interact at the border that separates the two regions. If, as is natural, the system is partitioned by atoms into the regions A and B, how do each of the models react to having an artificial surface where we switch to using the other model?

- **Selection** How do we select which regions are to be treated by which model? In some cases, this might be straightforward, e.g. if we are investigating an impurity or a stationary defect. But in general, the region of interest where inherently quantum mechanical processes take place can appear, disappear and move, and we need to be able to track that.

- **Validation** Given any particular hybrid scheme, it has to be validated against fully accurate quantum mechanical calculations. A successful scheme will be insensitive to the precise location of the border between the two models, and all observables should be able to be converged by increasing the size of the quantum region.

Before we describe our scheme in detail, we have to address the question of which systems are suitable for hybrid simulation. Critically, the atomistic processes, as described by our quantum mechanical model, must be *local*. It is well known that in the vast majority of cases, quantum mechanics is local in the sense that the density matrix $\rho(\mathbf{r}, \mathbf{r}')$ is sparse, its elements decay fast as the separation between $\mathbf{r}$ and $\mathbf{r}'$ increases (this kind of locality we call *weak locality*). This is very useful for creating linear scaling electronic structure algorithms, but is not enough for a hybrid approach to be valid. What is needed is *strong locality*, which we define by

$$\frac{\partial^n}{\partial \mathbf{x}_j^n}\frac{\partial E_{\text{total}}}{\partial \mathbf{x}_i} \to 0 \quad \text{as} \quad |\mathbf{x}_i - \mathbf{x}_j| \to \infty \quad \forall\, n, i \neq j \tag{7}$$

where E_{total} is the total energy, and $\mathbf{x}_i$ and $\mathbf{x}_j$ are the atomic positions of atoms i and j. This guarantees that, if the quantum region is large enough, the trajectories that are important at the center of the quantum region are not affected by the fact that far away, the system is treated classically (and hence the microscopic trajectories will not be very accurate) and also that the trajectories in the quantum region can be computed with high accuracy by only considering a finite neighbourhood. The length scale of the decay in equation (7) thus gives a guide to how large the quantum region needs to be.

Most quantum mechanical systems either obey strong locality, or at least the parts of the Hamiltonian that do not, such as long range Coulomb forces and Van der Waals forces, can be very effectively treated purely classically, outside the hybrid framework.

3 A new scheme

We have recently proposed a novel approach to the handshaking problem[4]. While some of the analysis put forward will not be repeated here, we note that the main philosophical difference with the previous approaches is that we do not try to make a combined Hamiltonian from the separate model Hamiltonians $H_{\text{classical}}$ and H_{quantum}, rather, we focus on a purely local quantity, the force on each atom. Then, to make a seam between two models, one can compute the forces in the transition region using both models, and smoothly cross over from using one set of forces on one side of the seam to using the other set on the other side.

This can be done in a number of ways, for example, by averaging the forces, or by making the transition region large enough so that at the outer edge the two models yield practically identical forces. In practice however, we choose a third way. To see why, let us consider how the resulting forces are used in the simulation. Having given up the notion of a combined Hamiltonian, we do not have a strict conservation of the total energy anymore. What we do have, are forces on each atom at any given configuration, and thus we can perform molecular dynamics. However, doing MD directly using forces that come from a variety of different Hamiltonians is a recipe for disaster: for example, the bonds between atoms would not obey an "action–reaction" principle. So instead of varying the source Hamiltonian for the forces in space, we vary a uniformly valid Hamiltonian in time. We choose a suitably general classical Hamiltonian, and allow its parameters to be different for each atom, bond, angle etc. and also vary in time. So rather than being used directly, the above mentioned forces that are calculated using the different model Hamiltonians $H_{\text{classical}}$ and H_{quantum} are used to *tune* the parameters of our universal Hamiltonian at each time step: the parameters are varied to minimize the functional

$$\mathcal{F} = \sum \left| F_{\text{universal}} - F_{\{\text{classical, quantum}\}} \right|^2 . \tag{8}$$

The MD is then carried out using forces derived from the universal Hamiltonian. Thus the action–reaction principle is restored, and if the tuning is always successful, the resulting forces are extremely close to the separately calculated forces in each region. A very compact way to describe this scheme is that the various models that are assigned to the different parts of the system are being *instantaneously interpolated* by a universal potential. Of course, from a

practical point of view, a natural starting point for such a universal potential is an already existing well tested classical potential that gives an adequate description of near–equilibrium configurations for the system at hand.
A few comments are now in order.

- There is an alternative viewpoint from which the above scheme can be understood. Having assumed that we have a classical model that works well in most regions, let us choose this Hamiltonian as the universal one. We want to improve its accuracy in particular places in space and time and we do this by computing relevant new information in those places, using a quantum mechanical model. We then incorporate this new information into our classical trajectories by slightly adjusting the parameters that describe the classical model (only in the region where the relevant new information applies). Whichever viewpoint is more applicable might depend somewhat on one's personal preferences, but also on just how much the classical parameters had to be readjusted. For small changes, this second viewpoint is very appealing, but if the quantum region requires an alteration that is very drastic and different at every time step, it seems to us more intuitive to talk about interpolation.

- The scheme uses a time–dependent Hamiltonian to carry out the dynamics, and thus energy is not conserved. Although this is not a logical consequence of the above, nevertheless we have not yet found a formulation which obeys strict energy conservation when the quantum region changes during the simulation, i.e. when atoms regularly cross from the quantum to the classical region and vice versa. In practice a thermostat is used to absorb any small drift. This works well except for problems where heat exchange with a bath directly affects what we want to measure, e.g. relative enthalpies.

- The parameter tuning of the universal Hamiltonian need not necessarily be carried out at every time step. One can use the same parameters for a small number of steps, and only then compute the forces with the expensive quantum mechanical model, which are then used to retune the parameters. This procedure can be formalized in a predictor–corrector scheme as follows. From a point $\mathbf{R}^0$ in phase space, having parameters α^0 (let us denote this by $(\mathbf{R}^0, \alpha^0)$, make a number of MD steps to arrive at point $(\mathbf{R}^1, \alpha^0)$. Here, quantum calculations and optimization takes place, so having got new parameters we are at $(\mathbf{R}^1, \alpha^1)$, this is the end of the "predictor" part. The MD could be continued from this point, but then the trajectory would suffer an abrupt change in its second derivative, because of the discontinuity in the Hamiltonian parameters, α. So instead, we go back to the previous point $(\mathbf{R}^0, \alpha^0)$ and redo the dynamics, interpolating the parameters between α^0 and α^1. At the end of the interpolation we arrive at a slightly different point in phase space, $(\mathbf{R}^2, \alpha^1)$, completing the "corrector" part. The cycle could be repeated by recomputing the quantum mechanical forces, getting new parameters α^2 and going back to $(\mathbf{R}^0, \alpha^0)$ to reinterpolate. In practice, we limit ourselves to one cycle, thereby ensuring a continuous variation of the parameters along the trajectory.

- The alternative viewpoint above highlights a possible computational simplification. Far away from the quantum region, the potential parameters would not be changing very much

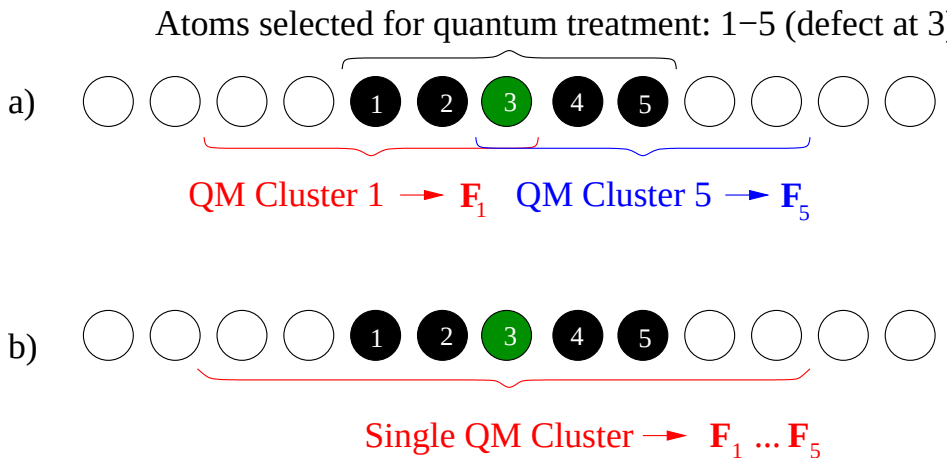

Figure 1: Given a region of interest, e.g. a point defect (green), we declare that this and other atoms within a certain distance should be treated with quantum mechanical accuracy (black). There are two ways of computing the accurate forces that we use in the fitting process, a) by carving a temporary cluster separately for each QM atom, and keeping only the force on the central atom from each QM calculation, or b) by computing a *larger* cluster quantum mechanically, and only keeping the forces on the originally selected QM atoms, discarding the forces near the (artificial) surface.

(or at least the precise details of the trajectories of far–away atoms do not affect processes in the quantum region due to strong locality). So they might as well be left as they are, and not included in the parameter optimization. This speeds up the optimization by a significant factor. If we have several quantum regions which are well separated, the optimization of the parameters near each region could be carried out in parallel.

- Because of strong locality, the forces on the atoms even in the quantum region can be computed independently. To do this, for each required atom, one carves out a spherical cluster centered on the given atom, and computes the forces, retaining only the force on the central atom (see Figure 1a). This temporary cluster can be chemically terminated (e.g. by hydrogen atoms) on the surface for improved accuracy (or smaller cluster radius for any given accuracy). This termination is not strictly necessary, but makes for a faster algorithm. We emphasize that the boundary of these clusters plays a very different role from the boundary of the quantum region which causes a lot of the difficulties in the traditional embedding schemes. Because each cluster calculation is independent, they can be carried out in parallel requiring minimal bandwidth for inter–processor communication. This possibility makes the scheme particularly suitable for running on massively parallel computer–farms, which are typically connected by a comparatively slow communication network.

  One can always fall back on a single large quantum mechanical calculation, as depicted in Figure 1b. The cluster should then be large enough that the forces on atoms near its boundary can be discarded (ie not used in the parameter optimization), thus again avoiding the problems that the introduction of an artificial surface might lead to.

Let us review the approximations that the scheme involves. Firstly, somewhat trivially, parts of

the system are treated with a less accurate classical model. This is a controllable approximation, since the size of the quantum mechanically treated region can be varied at will; this can be done in all embedding schemes. We expect a smooth variation of physical observables as the size of the quantum region is varied because the changing boundary of the quantum region does not bias the calculation: information coming from near the surface of the temporarily carved clusters is thrown away (see Figure 1). Secondly, if we compute quantum mechanical forces using separate clusters, the force on the central atom should be converged with the cluster radius, or at least the error should be within acceptable limits. For example, even in a first principles *ab initio* quantum mechanical calculation, if we do not expect to get forces more accurately than 0.1 eV/Å, it would be a waste of effort to converge to better than this with the cluster radius. The same consideration sets the target accuracy of the parameter optimization of our universal Hamiltonian. Below we give typical errors for all these approximations.

Although we lost the ability to do strictly microcanonical simulations, this is offset by other gains. The generality of the force–fitting means that any quantum engine can be plugged into the algorithm as a black box, without a need to alter the software. All physical observables can be computed from the MD trajectories, and we expect successes in problems where finite temperature is important, and the controlling quantity is the free energy. Moreover, since averages from different thermodynamical ensembles approach each other in the limit of large samples, long time averages from non–microcanonical molecular dynamics trajectories are just as good as those from microcanonical runs.

It turns out that the present scheme allows us to make the QM region very small and mobile, i.e. we can continuously select which atoms to treat more accurately during the simulation. This brings the obvious benefit of making the calculations run faster, but highlights the need for robust algorithms for determining where the QM region should be. It is clear that this will depend on the nature of the scientific question one is asking. It is not in general possible to correct the classical trajectories everywhere they go wrong, since tests show that even at modest displacements from equilibrium, classical models may predict forces that are wrong by as much as 1 eV/Å as compared with quantum mechanical forces.

In practice, this means that, for solid phases, the selection of the quantum region has to be implemented using geometric and topological criteria. For this, we found it advantageous to keep a running average of atomic positions to filter out the fast optical phonons, and assess the local geometry based on the these averaged positions,

$$\tilde{\mathbf{x}}(t) \equiv \frac{1}{1/2 + \tau/\Delta t} \sum_{n=0}^{\infty} e^{-n\Delta t/\tau} \mathbf{x}(t - n\Delta t), \tag{9}$$

where Δt is the time step of the molecular dynamics, and τ is the time–constant of the averaging. We found the averaged coordinates are very well suited to finding the location of point defects and locating bonds that are being broken.

We now present each step of the scheme in order, indicating numerical values for parameters that we have been using. For the sake of concreteness, we take as an example the diffusion of a vacancy in crystalline silicon. This is not a problem which a priori stands out as needing an embedded simulation, but precisely because of this, it can provide validation.

1. **Initialization** Start with the a 3x3x3 cubic unit cell (215 atoms and a vacancy), with

atoms near their bulk equilibrium positions, and with a reasonable choice of initial parameters for the classical potential, in our case we choose the Stillinger–Weber potential[5]. The velocities are randomized and rescaled to the simulation temperature, after which a thermostat is used to maintain an NVT ensemble.

2. **Extrapolation** As in standard molecular dynamics (MD), use the potential with fixed parameters to generate five 1 fs time steps of the system trajectory using the velocity Verlet algorithm.

3. **Testing** In the latest configuration, the local validity of the classical potential needs to be assessed on a site by site basis, and a selected subset of atoms is flagged for quantum treatment. As discussed above, this selection cannot be made fully automatically, because it depends on the scientific question being asked. Presently we are interested in the diffusion of a point defect, so it is natural to describe the region surrounding the defect quantum mechanically. To identify where the vacancy is, we use equation (9) with a time–constant of 100 fs, and match the averaged atomic coordinates with their nearest lattice points. Having located the site of the missing atom, we flag all atoms that are within 7.5 Å of its lattice point (about 100 atoms).

4. **Quantum mechanical calculations** We compute accurate forces on all flagged atoms using an empirical tight binding method[6] and the carved–cluster approach described above. The cluster radius is set to 7 Å, this yields about 150 atoms in each cluster, including terminating hydrogens. The termination atoms are put in place of the Si atoms that fell just outside of the given cluster and their positions are adjusted to the equilibrium Si–H distance along the direction of each Si–Si bond that was cut. The Hellman–Feynman force on the central atom is then computed after a direct diagonalization of the Hamiltonian. This force is shown to converge quickly with cluster radius and to be relatively insensitive to the precise termination strategy. We determined in a separate test (0.5 ps LOTF simulation of a vacancy diffusion event at 1400K) that a 7 Å radius is enough so that the RMS deviation from the the exact TB forces (computed with periodic boundary conditions) is < 0.1 eV/Å. Constructing clusters of similar size including all Si atoms up to a chosen neighbour shell rather than up to a fixed radius yields analogous results.

5. **Force fitting** We tune the parameters of the classical potential until it reproduces the accurate forces, by minimizing (8) using the conjugate gradients technique. The starting values for the parameters are always reset to be the same as those of original unmodified potential. For atoms that do not have a quantum mechanical force calculated on them, the current classical force is used as the target of the optimization. Note that we require the evaluation of the classical potential for a large number of different parameters sets, all at the same atomic configuration. We have devised special routines that reformulate the potential to be of the form

$$E = M_0(\mathbf{R}) + M_1(\mathbf{R})\alpha + M_2(\mathbf{R})\alpha^2,$$

where only the matrices M_i depend on the atomic coordinates and are computed once before the optimization starts; α and α^2 indicate the vector of the parameters and their

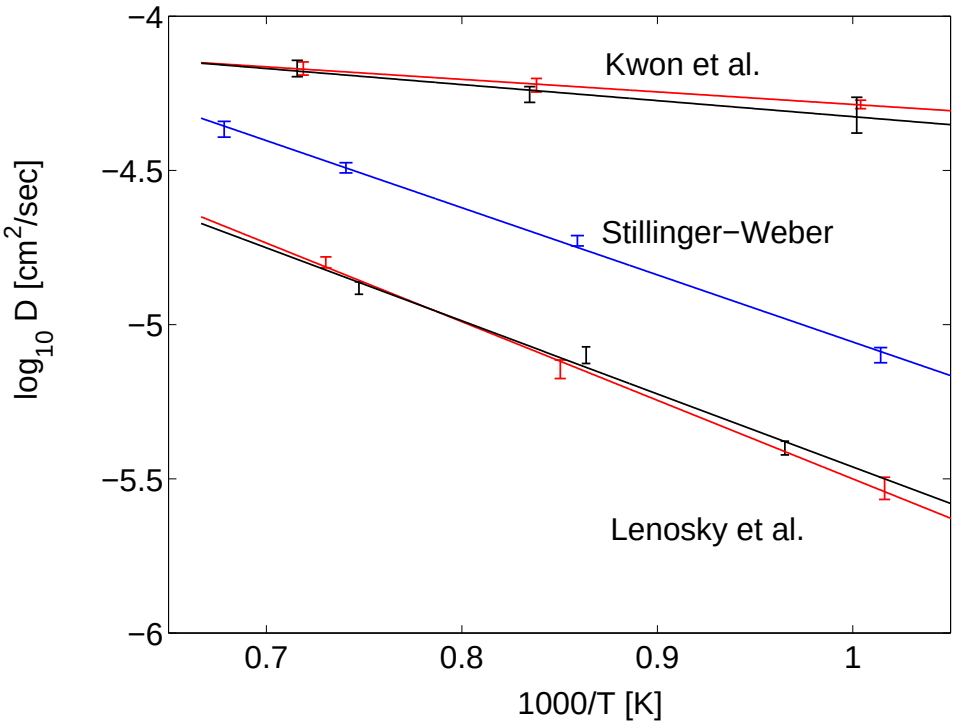

Figure 2: Arrhenius plot of the silicon vacancy diffusivity obtained by the present scheme fitted "on the fly" on different TB models (red), compared with the results of fully TB simulations (black) and fixed-parameter SW potential (blue).

squares, respectively. The forces can be written in a similar way and thus during the optimization, most of the CPU time is spent in BLAS libraries. Tests show that parameters related to atoms which are further than a few neighbour shells outside the QM region do not change appreciably with the fit. The optimization is therefore limited to the two–body and three–body terms of the potential involving atoms in a spherical region of 12 Å radius centered on the diffusing defect. About 30 conjugate gradient steps are sufficient to perform a converged fit, the maximum residual deviation from the target forces being typically less than 0.01 eV/Å.

6. **Interpolation** Now return the state of the system to that before the extrapolation and rerun the dynamics, interpolating the potential parameters between the old and the new values. We use a cubic spline interpolation, so the time–derivatives of the parameters are zero at the fit–points.

7. **Return to 2.**

4 Validation

Using this procedure, we computed the diffusivity for the Si vacancy[7] by running a 1 ns constant temperature (CT) simulation of a 215 atom periodic system at $T = 1400$K and obtained a value of $(2.1 \pm 0.4) \times 10^{-6}$ cm^2/sec. This compares well with the value of $(2.3 \pm 0.2) \times 10^{-6}$ cm^2/sec, calculated using the (same) tight binding scheme[6] for the full system. For comparison, the Stillinger–Weber potential gives $(3.4 \pm 1.0) \times 10^{-5}$ cm^2/sec), a much larger value. It must be

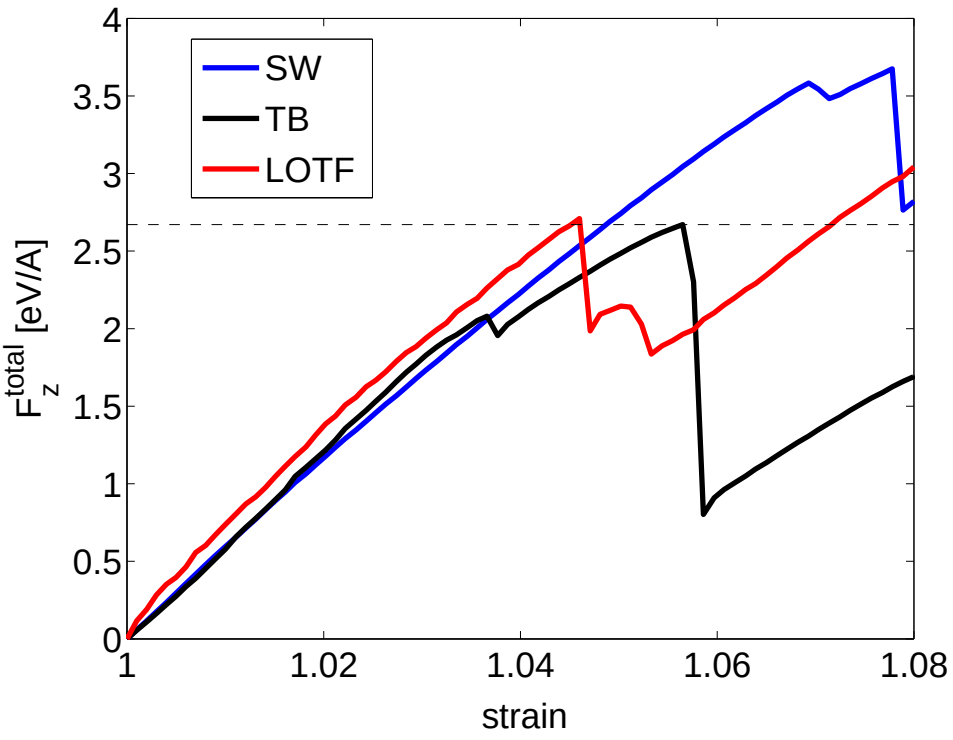

Figure 3: Stress–Strain curves of a silicon nanobar with a vacancy, using three different methods: Stillinger–Weber (blue), tight–binding (black), and LOTF(red).

stressed that the diffusivity is a quantitative test of the embedding method, and should not be taken to be an application where embedding is really necessary. A 215 atom silicon system is most efficiently simulated by standard techniques.

This is even more true of our next test, where we look at the temperature dependence of the vacancy diffusivity. To enable the large number of runs that a good comparison requires, we have chosen to use a 2x2x2 cubic unit cell with 63 atoms. Here, the quantum forces used for the parameter fitting in the LOTF run were calculated using the full periodic periodic system, rather than with clusters.

Figure 2 shows our results for the vacancy diffusivity in Si as a function of temperature, obtained by CT simulations using the fixed–parameter SW potential, two different TB schemes[8, 9] and our hybrid LOTF scheme. We note that the large difference between the results obtained by fitting the scheme on the two different quantum models is due to the different predictions of these models in accurate quantitative computations, and faithfully reproduced by the hybrid method. This emphasizes the fact that the present scheme can at best be expected to reproduce the results of the QM model that it is given, but can in no way improve its accuracy.

A different quantitative test is provided by the example of stressing a silicon nanobar. Starting with a 2x2x10 supercell of the cubic unit cell, we introduce a vacancy, and relax the structure to equilibrium. Then the atomic coordinates are rescaled parallel to the bar axis, and again relaxed to equilibrium while holding the ends fixed. In the hybrid simulation, the QM region included atoms within 4 Å of the vacancy. In this case, rather than using clusters, the tight–binding forces were calculated by considering a cluster that was a bit larger than the QM region (by 5 Å on each side), and again terminated with hydrogen atoms where bonds were broken in the

construction of the cluster.
Figure 3 shows the resulting stress–strain curves (the sum of the z components of the forces on the fixed atoms at one end of the bar is proportional to the stress). The hybrid simulation and the fully quantum simulation give the same critical stress level (dashed line) at which the bar fails mechanically near the vacancy, while the classically computed critical stress is significantly higher. This example also illustrates that the microscopic trajectories are, of course, not necessarily reproduced: the curves have minor jogs and kinks due to small changes in the surface reconstruction of the nanobar, which do not affect the critical stress.

5 Brittle failure

Having validated our scheme in small systems where it was possible to compute observables using a fully quantum mechanical method, we now turn to a system which has been epitomized as the benchmark for hybrid schemes, brittle failure of covalent crystals. When single crystals of silicon or diamond are strained beyond breaking point, they relieve the stress by opening cracks. A universal observation is that these cracks are atomically sharp and the resulting newly opened surfaces are almost atomically flat[10]. The system is evidently multiscale, and the behaviour on the different lengthscales are highly correlated: the large stress concentration near the crack tip induces the breakage of atomic bonds, and it is the successively broken bonds that let the crack tip (and the stress field) advance forward as the crack propagates.
To date, all purely classical simulations have failed to reproduce a basic feature of the system, namely that the cracks tips remain atomically sharp as the crack opens. This is principally due to the overestimation of lattice trapping effects[11], which results in an excess of applied stress with respect to the critical value to get the crack moving in the first place. This excess of elastic energy is dissipated by tip blunting and local melting.
Other efforts at developing hybrid schemes have also concentrated on cracks[3]. The principal advantages of our scheme is that the quantum region can be made very small (300-600 atoms for a silicon crack tip) and is allowed to follow the tip as the crack propagates, resulting in a computational speedup without compromising the accuracy of the calculation.
Figure 4 shows three snapshots of our simulation of the $(111)[1\bar{1}0]$ crack system, as it propagates in the $[11\bar{2}]$ direction at 0 K. We used the *ab initio* SIESTA[12] package to treat atoms in the moving quantum region. Apart from showing smooth atomically flat surfaces opening, using such a sophisticated QM engine revealed that the surface reconstruction is different on the upper and lower surfaces, with the upper surface showing the 2×1 Pandey reconstruction, as indicated by the alternating pentagons and heptagons. The reason why only one surface shows this reconstruction can be understood by considering the dynamics of atoms near the crack tip. The atoms, shown in black, that were part of a hexagon are forced towards each other as the crack tip passes, thus driving the formation of a pentagon. The corresponding atoms on the lower surface (shown in gray) are forced away from each other, preventing pentagon formation.

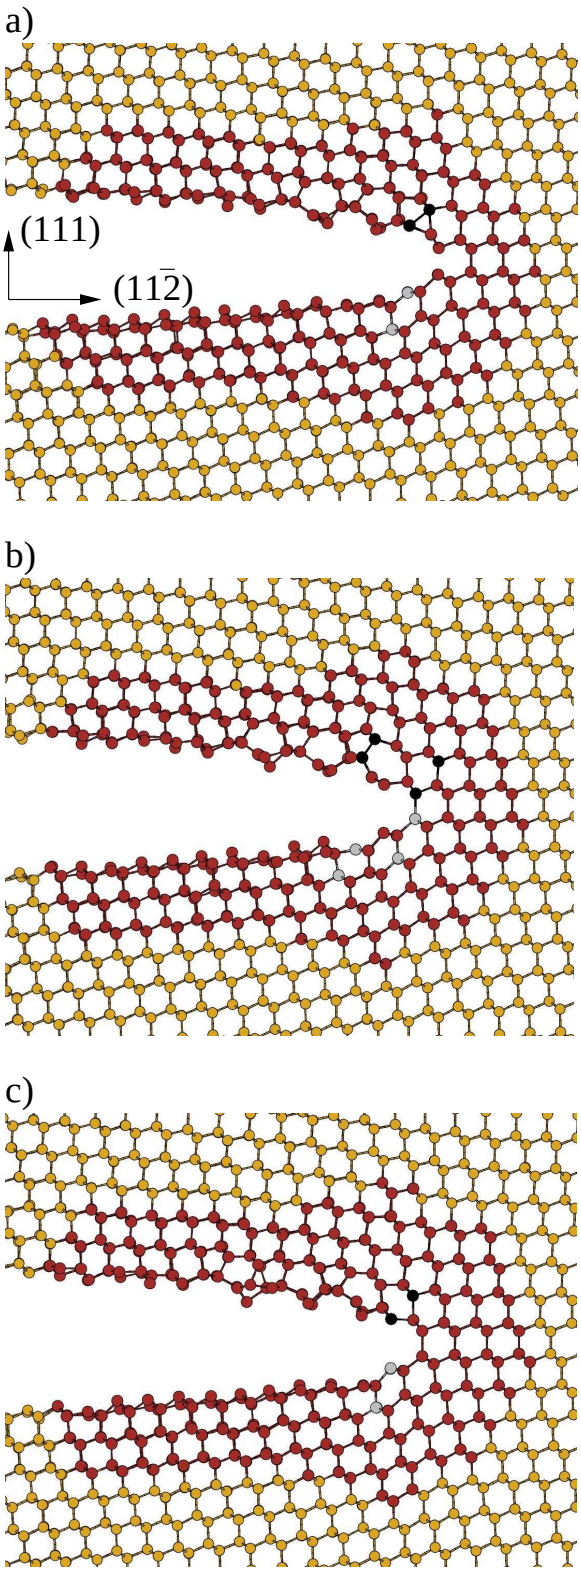

Figure 4: Three snapshots of the opening (111)[1$\bar{1}$0] crack system simulated using LOTF. Atoms in red were treated quantum mechanically using the *ab initio* SIESTA package. The quantum region follows the crack tip as it moves from left to right. Note the pentagons and heptagons of the 2×1 Pandey reconstruction on the top surface. See text for an an explanation of the dynamical aspects.

References

[1] R. M. Nieminen, J. Phys. Cond. Mat. **14** 2859 (2002).

[2] E. B. Tadmor *et al.* Philos.Mag. **73** 1529 (1996) ; E. B. Tadmor *et al.*, Langmuir **12** 4529 (1996).

[3] F.F.Abraham *et al.*, Comp.Phys. **12**,538 (1998); F.F.Abraham *et al.*, Mat. Res. Soc. Bulletin **25**, 27 (2000); C. L. Rountree *et al.* Annu. Rev. Mater. Res. **32** 377 (2002); S. Ogata *et al.* J. Appl. Phys. **95** 5316 (2004); S. Ogata and R. Belkada, Comp. Mat. Sci. **30** 189 (2004).

[4] Gábor Csányi *et al.*, Phys. Rev. Letts **93** 175503 (2004); A.DeVita and R.Car,MRS Symp.Proc. **491**,473 (1998).

[5] F.Stillinger and T.Weber,Phys.Rev.B **31**,5262 (1985).

[6] D.R.Bowler *et al.*, J.Phys.Cond.Matt. **10**, 3719 (1998).

[7] M.Tang *et al.*, Phys.Rev.B **55**, 14279 (1997).

[8] I.Kwon *et al.*, Phys. Rev. B **49**, 7242 (1994).

[9] T.J.Lenosky *et al.*, Phys. Rev B **55**, 1528 (1997).

[10] S. Kohlhoff *et al.*, Phil.Mag.A**64**,851 (1991).

[11] N. Bernstein and D. W. Hess, Phys. Rev. Lett. **91**, 025501 (2003).

[12] D.Sanchez-Portal *et al.*, Int.J.Q.Chem. **65**,453 (1997).