Transactions of the Materials Research Society of Japan 29 [8] 3593-3597 (2004)

Multiscale Simulations of Fracture in Silicon

Noam Bernstein
Center for Computational Materials Science,
Naval Research Laboratory, Washington, DC, USA.

To help us understand the fracture properties of materials computer simulations must describe processes occurring on a wide range of length scales, with methods that are efficient enough and accurate enough to capture the important physics. I briefly present the original coupling of length scales (CLS) method, which combined tight-binding, empirical potential, and continuum elasticity descriptions of matter. I present a summary of the results, and the lessons we learned from the original implementation. I then present the dynamically coupled empirical-potential and tight-binding (DCET) method, which uses a more accurate approximation to compute the tight-binding forces. I show that the DCET approach gives brittle fracture in good agreement with experiment. To analyze the results I present a simple model for the energetics of the crack propagation process. The model indicates that the empirical potential simulations are ductile because the barrier suppresses brittle crack propagation until dislocations form and blunt the crack. Two length scales that control the fracture mode emerge from the model - one for bond breaking, and another for the shape of the crack tip. Differences in both length scales cause the qualitative differences in crack morphology between the DCET and the empirical potential results.

Keywords: Multiscale simulations, fracture, silicon, lattice trapping

1. INTRODUCTION

Structural materials are materials that are used to build structures, and must maintain their shape under mechanical load. If the load becomes too large, they fail in a variety of possible ways. For example the material can yield plastically, i.e. deform irreversibly, or fracture. The technological importance of structural materials is clear: We would like to be able to predict the fracture properties of materials, from first principles if possible, and control failure properties. However, there is also an interesting underlying scientific question: Why do different materials break in the particular ways that they do?

Two of the most common modes of failure are brittle and ductile fracture. In terms of the atomistic nature of materials these two modes of fracture are distinct, controlled by the details of the atomic motion at and near the crack tip. Brittle fracture is characterized by an atomically sharp crack tip. The sharp crack concentrates the applied stress, and the stress singularity remains as the crack propagates. In the simplest case there is no damage in the material except for the crack itself. This is why brittle fracture dissipates the least possible amount of energy during crack propagation - only the energy needed to expose the new crack surfaces.

In ductile fracture, the crack becomes blunt, and the material deforms non-reversibly. This occurrence of plasticity near the crack tip can take several forms, but the most common is the emission of dislocations from the crack tip. These dislocations cause the crack to blunt and therefore reduce the stress singularity at the crack tip. The motion of the dislocations also dissipates energy. This is why most structural materials are ductile - their failure dissipates a lot of energy, and preexisting cracks tend to become blunt and arrest.

In this work I simulate the fracture properties of silicon, which is an excellent model system for the temperature driven transition between brittle and ductile fracture [1]. Experiments on silicon take advantage of the availability of extremely pure and defect free samples, and simulations can use a range of techniques, including first-principles methods, semi-empirical quantum-mechanical methods, and empirical potentials. Our goal is a microscopic understanding of processes that occur at the crack tip. From experiment we can get mostly static or macroscopic information. We can measure the crack speed as a function of loading, and the critical loading at which crack propagation begin [2]. We can also examine the morphology of the newly exposed surface. However, atomic resolution information is only available for static or nearly static cracks, and only where they intersect the surface [3]. It is also extremely difficult to create a perfect system, entirely defect free and under perfectly controlled mechanical loading. Simulations can trivially provide perfect materials under perfectly controlled loading geometries, and atomistic resolution on femtosecond time scales. However, the system sizes and times that can be described are extremely limited. More importantly, the description of the interactions between the atoms is difficult, and the more accurate and reliable methods are very computationally intensive.

In Sec. 2 I present an overview of the concurrent coupling approach, and in Sec. 3 I present our original proof-of-principle method, and the lessons we

learned. In Sec. 4 I describe our current method and results of brittle fracture in silicon consistent with experiment. Finally, in Sec. 5 I present a model for lattice trapping that explains the differences between our ductile empirical potential simulations and the brittle fracture in the current concurrent coupling method, and summarize in Sec. 6

2. CONCURRENT COUPLING

The fracture process is controlled by processes that occur over a wide range of length scales. Each of these processes is best captured by different physical descriptions and different computational methods. At the shortest length scale, comprising a few hundred atoms near the crack tip, interatomic bonds are breaking and reforming. In this region we can regard the nuclei as classical point masses, but need a quantum-mechanical description of the electrons that mediate interatomic bonding. At somewhat longer length scales, bonds are highly strained but not broken. In this region we may still need an atomistic description, but the interactions between the atoms may be adequately given without an explicitly quantum-mechanical description. At longer length scales the atomistic nature of the material is irrelevant. The elastic deformation that drives the fracture process can be described by continuum elasticity. It is important to consider which of these aspects need to be included in a simulation, and how they can be combined.

In a concurrent coupling approach, a large system is described with a fast method that is somehow limited. Where this method fails a slower but more reliable method is used. For example in the case of fracture, we may need a quantum-mechanical description of bonding at the crack tip, but use an empirical potential to simulate the rest of the loaded system. However, this approach only works if there is a localized region where the slower method is needed. One obvious question about the concurrent coupling approach is when is it needed. Why can't the external loading conditions be represented by an appropriate choice of boundary conditions for the more reliable simulation method? The most important class of problems where this is the case is when there are dynamics, or changing boundary conditions. For example in the case of fracture, it is quite difficult to predict *a priori* the mechanical boundary conditions on a non-steadily moving crack in an arbitrary geometry.

3. THE ORIGINAL COUPLING OF LENGTH SCALES METHOD

3.1. Method

In the original Coupling of Length Scales (CLS) method [4, 5] the system was divided into three regions. Far from the crack we used a finite-element solution of linear continuum elasticity to describe the material. Closer to the crack we used molecular dynamics using the Stillinger-Weber empirical potential (EP). At the crack tip we used a simple quantum-mechanical model of bonding given by a tight-binding (TB) Hamiltonian for silicon. For computational efficiency we used some severe approximations in computing the forces on atoms in this region. To insure that simulations would be numerically stable for long times, the total energy of the system was well defined by construction and the dynamics for each region were updated in sync.

3.2. Results

Using this method we simulated crack propagation in systems that were about 4000 Å × 3600 Å × 11 Å. Our results showed elastic waves emanating from the crack tip, and considerable structure that developed in the fractured system. The crack was blunt, and voids and amorphous tendrils formed near the crack tip. The use of TB at the crack tip did not seem to influence the results as compared with simulations that just used finite-element and EP regions. Since this morphology is quite different from the conventional view of brittle fracture, it was not clear if our observations were consistent with the brittle nature of silicon at low temperatures. However, the disorder and plasticity near the crack tip never extended very far into the bulk, and it was not obvious that they could have been observed experimentally. Initially we believed that our observations lent support to Rice's speculation that silicon is intrinsically ductile, but low dislocation mobility keeps it brittle by preventing large-scale plasticity from occurring.

A more quantitative comparison of our simulations and experiments gave different conclusions. The critical loading for fracture to propagate, measured in terms of the energy release rate G, gives a quantitative measure of the energy dissipated. In our simulations we got values for the critical G, G_c, that ranged between 8 and 130 J/m^2, as did Marder *et al.* [6, 7]. The wide range depended on strain rate, orientation, and the precise definition of a "propagating" crack. The minimal energy to propagate a crack in silicon given by the Griffith criterion is two times the surface energy, or about 2 to 3 J/m^2. All of the extra energy is needed in the simulation to create the disorder around the crack tip. Careful experiments, on the other hand, showed that G_c is about 2.5 J/m^3 [2]. This value is very close to the Griffith criterion, and indicates that silicon is brittle in the conventional sense and that significant disorder, even if confined to the crack-tip region, is excluded.

Despite this disagreement between our original simulations and experiment, we could at least draw some conclusions about the computational approach. By coupling continuum mechanics, EP, and TB we got stable, long-running simulations. However, neither the EPs nor the approximate solution of the TB forces were accurate enough to describe fracture in silicon. Clearly a more accurate solution of the forces in the TB region was needed. Finally, the time scales involved in the finite-element region and the TB regions were widely separated. In practice it is clear that the computational time for a simulation is dominated by the most accurate method. If an accurate solution of the forces in the TB region requires 100 s for a 1 fs time step, and we are willing to simulate for 1 week, that results in an overall simulated time $t = 6$ ps. There is a length scale l implied in this time

t, because information can only travel at the speed of sound in the material v

$$l = tv.$$

Given these order-of-magnitude times the length scale is 240 Å. Especially in the type of quasi-2D systems used here, it is easy to simulate a system of this size with EPs, without a need for a continuum elasticity region. We therefore developed a new version of the concurrent coupling method, focusing on a more accurate description of the TB region, and including only TB and EP regions.

4. DYNAMICALLY COUPLED EMPIRICAL POTENTIALS AND TIGHT-BINDING

4.1. Method

In the dynamically coupled EP and TB (DCET) approach we used a more accurate approximation to compute the forces in the TB region. The simulated system is divided into a large EP region, with a TB region surrounding the crack tip, and a boundary region where the two methods are coupled [8]. In the TB region we use a real-space Green's function method [9] with constraints at the boundary of the TB region. Mechanical coupling is provided by using the positions of atoms in the TB region to compute EP forces on atoms in the boundary and EP regions. The positions of atoms in the boundary region, which follow EP trajectories, act as a mechanical boundary condition for the atoms in the TB region. While the EP and TB regions acts as mechanical boundary conditions for each other, the forces are computed directly rather than as derivatives of a single total energy. This means that the total energy is not rigorously defined, and so numerical stability will have to be tested empirically.

4.2. Results

We used the DCET method with the environment dependent interatomic potential (EDIP) EP [10], and the nonorthogonal TB model for silicon by Bernstein and Kaxiras [11]. The system, which consisted of about 50000 EP atoms and 1000-2000 TB atoms, was initialized with the displacement field of a finite crack in an infinite plate under tensile loading. The boundaries along the loading axis were fixed, and the others periodic. The whole system was about 400 Å × 250 Å × 12 Å, with a [111] crack face and a $\langle 1\bar{1}0 \rangle$ crack front. The simulation shows that the crack propagates in a brittle manner, leaving behind an atomically smooth [111] surface (Fig. 1). The critical energy release rate for crack propagation, computed using the EP elastic constants, is almost exactly the Griffith criterion value of two times the surface energy, computed using TB. Clearly the DCET results are in much better agreement with experiment than our original CLS approach.

In is worth noting that the crack speed as a function of loading quickly rises to about 2.5 km/s, and then saturates. This behavior is quite similar to that seen in experiment. However, very similar speeds were seen for the original CLS simulations as well [5]. Clearly crack speed is not a sensitive test of the fracture mode.

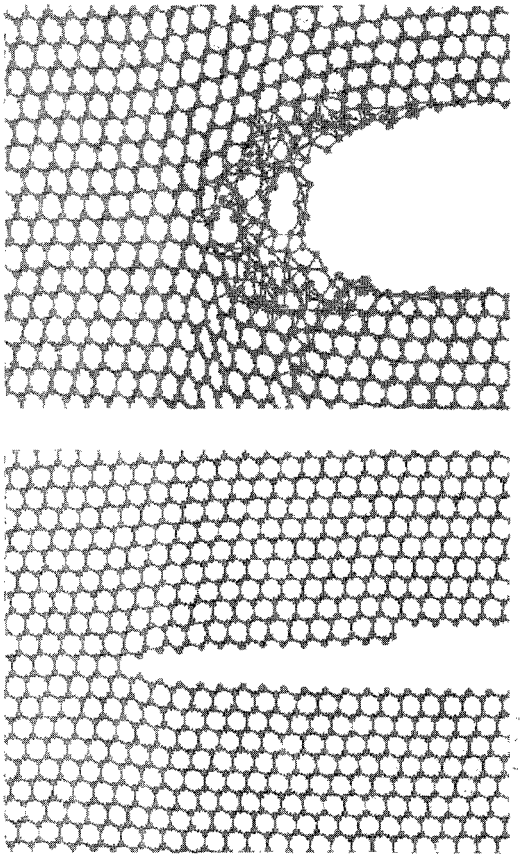

FIG. 1: Visualization of the simulated crack after propagation using the EDIP EP (top panel) and DCET (bottom panel) simulations.

	LDA	BK-TB[12]	EDIP[10]	SW
γ_s (111) ideal	1.7	1.0	1.1	1.4
γ_{us} glide relaxed	1.9	2.5	1.9	3.1
γ_{us} shuffle relaxed	1.7	1.1	1.3	0.8
γ_s/γ_{us} (glide)	0.90	**0.40**	0.59	**0.45**
γ_s/γ_{us} (shuffle)	1.02	**0.90**	**0.85**	1.71

TABLE I: Surface energy γ_s, unstable stacking energy γ_{us}, and Rice criterion ratios for first-principles DFT/LDA calculations, tight-binding, and the EDIP and SW empirical potentials.

4.3. Analysis

While the differences between the EP simulations (EDIP, SW, as well as the original CLS) and the DCET simulations are clear, the reasons for them are not. Presumably the accurate calculations of forces in the TB region is essential. What are the fundamental materials properties that are different in the EP and TB descriptions, and how do they control the fracture mode in the simulations?

One way to compare the tendency for brittle fracture vs. ductility is the Rice criterion [13, 14]. In this approach the Griffith criterion G is compared to the critical G for emitting dislocations. If the load for brittle fracture is lower than the load for emitting a dislocation, the material is brittle. If not, it is ductile. However, TB and the two EPs, SW and

EDIP, show very similar ratios for the surface energy (cleavage) and the unstable stacking energy (dislocation emission) (Table 4.3). In all cases the Rice criterion predicts that silicon should be brittle, although the EP simulations show ductile behavior.

The reason for the discrepancy between the EP simulations and the Rice criterion values for the EP is clear when one examines the critical stress for fracture. There is a large range of loadings where fracture is energetically favored (i.e. above the Griffith criterion G_c), but the crack does not propagate. Eventually the loading becomes so high that a dislocation is emitted. In the EP simulations the cracks are trapped by some energy barrier that prevents brittle fracture.

To understand this barrier we develop a simple model [8], related to the work of Curtin [15] and Perez [16]. The energy of the system is tracked as the crack propagates by one atomic spacing. The energy is separated into two components - a bond breaking energy, and an elastic relaxation energy. For the EP simulations energies can be computed directly. First a series of constrained relaxations are used to compute the total energy during the crack propagation. The bond-breaking energy as a function of distance is approximated by the energy per bond for separating two flat, periodic slabs to form two surfaces. The elastic energy is computed by subtracting the bond-breaking energy from the total energy during the crack propagation process. By repeating this process for different EPs at different loading a set of elastic energy curves are derived. We find that by appropriately scaling the energy and length scales, all of the elastic energy curves can be collapsed onto a single curve. The energy scale is the continuum mechanics energy release rate, and the length scale is the opening between two atoms immediately behind the crack tip. The existence of the universal elastic energy part indicates that, at least for the EP models, the crack propagation energy barrier can meaningfully be decomposed into these two parts.

For the EPs the total energy during the crack propagation process shows a clear and significant barrier (Fig. 2). At the Griffith criterion critical loading the final energy is nearly equal to the initial energy, consistent with the Griffith criterion definition. For the SW EP, at the loading where dislocation nucleation begins the barrier is still significant. For a modification of the SW potential that is brittle (but not a very good model of silicon) [17, 18], brittle fracture is initiated at the loading where the barrier goes to zero. There is one EP for silicon using the modified-embedded-atom method (MEAM), not used in the present work, that yields brittle fracture without modification [19]. The simulations by Swadener *et al.* did not focus on the value of G_c, but from their figures it appears that they need to apply loadings significantly above the Griffith criterion threshold to propagate the crack.

For the DCET simulation the total energy is not defined, and therefore the energy during the crack propagation process can not be computed. However, the bond-breaking energy can be computed using TB, and the universal elastic energy can be rescaled. From these two pieces we can reconstruct the effective energy barrier to crack propagation (Fig. 2). We find that at the Griffith critical loading there is a very small barrier, and at the loading where fracture is initiated quasi-statically the barrier has gone to zero. This indicates that the model, and the separation of the energy barrier into bond-breaking and elastic-energy parts, applies to the DCET simulation as well.

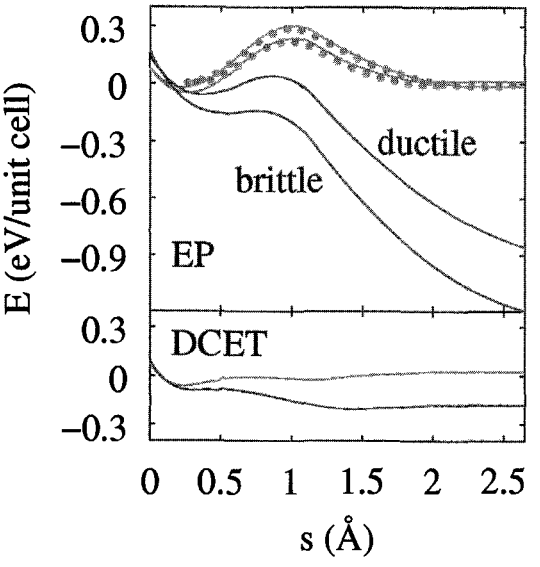

FIG. 2: Energy vs. crack opening distance for EPs (top panel) and DCET (bottom panel). Upper two curves in the bottom panel show energy barriers for SW and modified SW, with dots indicating direct calculation and lines indicating the prediction of the model. Lower two curves in the top panel indicate the model prediction for the barrier at the dislocation nucleation loading (SW, ductile) or crack propagation (modified SW, brittle). The upper curve in the bottom panel is at the Griffith critical loading, and the lower curve at the onset of crack propagation.

5. DISCUSSION

The EP simulations show large lattice trapping barriers to brittle fracture. For the ductile EPs this barrier is so large that crack propagation is completely suppressed, until at some large loading dislocation nucleation begins. It is clear from our barrier calculations that the modified SW EP is more brittle than SW, not because it is easier for brittle fracture to occur, but because it is harder for dislocations to nucleate. Experimental evidence concerning the barrier is indecisive. There is significant uncertainty is the experimental measurement of G_c [2], and no accurate way of measuring the surface energy. The best theoretical calculation for the surface energy, from density functional theory, indicate that the barrier, if any, is small. However, it is unclear if the theoretical value is accurate enough to exclude a barrier of around 5-10% of the surface energy, although the very large barriers predicted by SW are almost cer-

tainly incorrect.

The separation of the energy barrier to crack propagation into bond-breaking and elastic-energy contributions has several implications for our understanding of fracture. The energetics of breaking a bond at the tip of the crack are well described by the gradual separation of two flat, infinite slabs of material. This indicates that the bond-breaking process at the crack tip is not affected by the local strain gradients or the asymmetry created by the presence of the open crack. The elastic energy is universal up to the energy and length scale factors. The energy scale is accurately set by linear elasticity (it is the value of G), but the length scale, which comes from the opening of the crack immediately behind the tip, is not. The crack opening shape varies significantly between the DCET, EPs, and the continuum elasticity prediction. This indicates that linear elasticity holds except for the details of the shape of the crack tip.

The common view of the fracture process, where the only length scale is given by the bond breaking distance, is replaced by a more complex picture. The length scale governing the crack-tip shape is also important. In the DCET results, the bond breaking distance is 2-3 times larger than for the EPs, and the crack-opening distance 25-40% smaller. Both length scales combine to reduce the lattice trapping barrier in the DCET simulations as compared with the EP results.

6. SUMMARY

The process of fracture depends on an interplay between the bond breaking process, deformation near the crack tip, and the long range elastic load. Concurrent multiscale simulations provide an efficient way of combining methods that describe these different physical regimes. While our original CLS approach showed that the three regions can be simulated by coupling tight binding, empirical potentials, and finite elements, it proved to fall short of the required level of accuracy. The DCET method, which grew out of the original proof-of-principle implementation but focused on only on a more accurate evaluation of the tight-binding and empirical-potential regions, has shown that it is possible to simulate dynamic brittle fracture in silicon. Analysis of the results and comparison to empirical-potential simulations shows that lattice trapping during brittle fracture can be considerable, and that two mechanisms are involved. One, the bond-breaking process, has a length scale that has long been known to be important. The other, elastic relaxation controlled by the shape of the crack tip, brings in a second length scale that has previously been discounted. It is the interplay between these two mechanisms that controls the lattice trapping energy barrier to brittle fracture.

Acknowledgments

I'd like to thank my collaborators on the original CLS approach, E. Kaxiras, F. F. Abraham, and J. Q. Broughton, and on the DCET work, D. W. Hess. This work was supported by the Office of Naval Research, the Naval Research Laboratory, and the DOD HPCMPO CHSSI program. Computer time was provided by the DOD HPCMPO Challenge program and the high-performance computing program.

[1] J. Samuels and S. G. Roberts, Proc. Roy. Soc. London A **421**, 1 (1989).
[2] J. A. Hauch, D. Holland, M. P. Marder, and H. L. Swinney, Phys. Rev. Lett. **82**, 3823 (1999).
[3] F. Celarie, S. Prades, D. Bonamy, L. Ferrero, E. Bouchaud, C. Guillot, and C. Marliere, Phys. Rev. Lett. **90**, 075504 (2003).
[4] F. F. Abraham, J. Q. Broughton, N. Bernstein, and E. Kaxiras, Europhys. Lett. **44**, 783 (1998).
[5] J. Q. Broughton, F. F. Abraham, N. Bernstein, and E. Kaxiras, Phys. Rev. B **60**, 2391 (1999).
[6] F. F. Abraham, N. Bernstein, J. Q. Broughton, and D. Hess, Mat. Res. Soc. Bull. **25**, 27 (2000).
[7] D. Holland and M. Marder, Adv. Mat. **11**, 793 (1999).
[8] N. Bernstein and D. W. Hess, Phys. Rev. Lett. (2003), in press.
[9] N. Bernstein, Europhys. Lett. **55**, 52 (2001).
[10] J. F. Justo, M. Z. Bazant, E. Kaxiras, V. V. Bulatov, and S. Yip, Phys. Rev. B **58**, 2539 (1998).
[11] N. Bernstein and E. Kaxiras, Phys. Rev. B **56**, 10488 (1997).
[12] E. Kaxiras and M. S. Duesbery, Phys. Rev. Lett. **70**, 3752 (1993).
[13] J. R. Rice and R. Thomson, Phil. Mag. **29**, 73 (1974).
[14] J. R. Rice, J. Mech. Phys. Solids **40**, 239 (1992).
[15] W. A. Curtin, J. Mater. Res. **5**, 1549 (1990).
[16] R. Perez and P. Gumbsch, Phys. Rev. Lett. **84**, 5347 (2000).
[17] D. Holland and M. Marder, Phys. Rev. Lett. **80**, 746 (1998).
[18] D. Holland and M. Marder, Phys. Rev. Lett. **81**, 4029 (1998).
[19] J. G. Swadener, M. I. Baskes, and M. Nastasi, Phys. Rev. Lett. **89**, 085503 (2002).

(Received October 31, 2003; Accepted April 23, 2004)