PERGAMON

International Journal of Solids and Structures 37 (2000) 379–389

INTERNATIONAL JOURNAL OF SOLIDS and STRUCTURES

www.elsevier.com/locate/ijsolstr

Hierarchical modeling in the mechanics of materials

E.B. Tadmor[a], R. Phillips[b,*], M. Ortiz[c]

[a]*Faculty of Mechanical Engineering, Technion — Israel Institute of Technology, Haifa 32000, Israel*
[b]*Division of Engineering, Brown University, Providence, RI 02912, USA*
[c]*Department of Aeronautics, California Institute of Technology, Pasadena, CA 91125, USA*

Abstract

Many problems in the mechanics of materials involve the operation of either multiple spatial or temporal scales simultaneously. As a result, an important thrust of recent work in this area has been the development of methods allowing for the consideration of multiple scales simultaneously. In this paper, we examine hierarchical approaches to modeling problems of this kind with special reference to the way in which information can be fed from one scale to the next in models of plasticity.

1. Introduction

In 1888, the British metallurgist W.C. Roberts-Austen reported on a series of experiments studying the effect of impurities on the tensile strength of gold. His results (plotted in Fig. 1) indicated a strong and regular dependence of the tensile strength on the size of the impurity. More than a century later, significant progress has been made in elucidating many aspects of plasticity and fracture, though we are still far from being able to predict data like that presented in Fig. 1 on the basis of a fundamental understanding of the interaction of the atoms making up a solid.

The traditional solid mechanics and engineering approach to problems of this variety has been to exploit empiricism and phenomenology. The material is treated as a black box and characterized by exhaustive testing. In this setting, curves such as those of Roberts-Austen are used as input into phenomenological constitutive models which are blind to the underlying defect structures responsible for the observed behavior. This approach has proven to be extremely efficient, leading from Hooke's elucidation of the law of elastic deformation to the modern theory of fracture mechanics.

In recent years, a dominant research theme has been to go beyond such phenomenology to a more

* Corresponding author Tel.: +1-401-863-2674; fax: +1-401-863-1501.
E-mail address: phillips@engin.brown.edu (R. Phillips)

0020-7683/00/$ - see front matter
PII: S0020-7683(99)00095-5

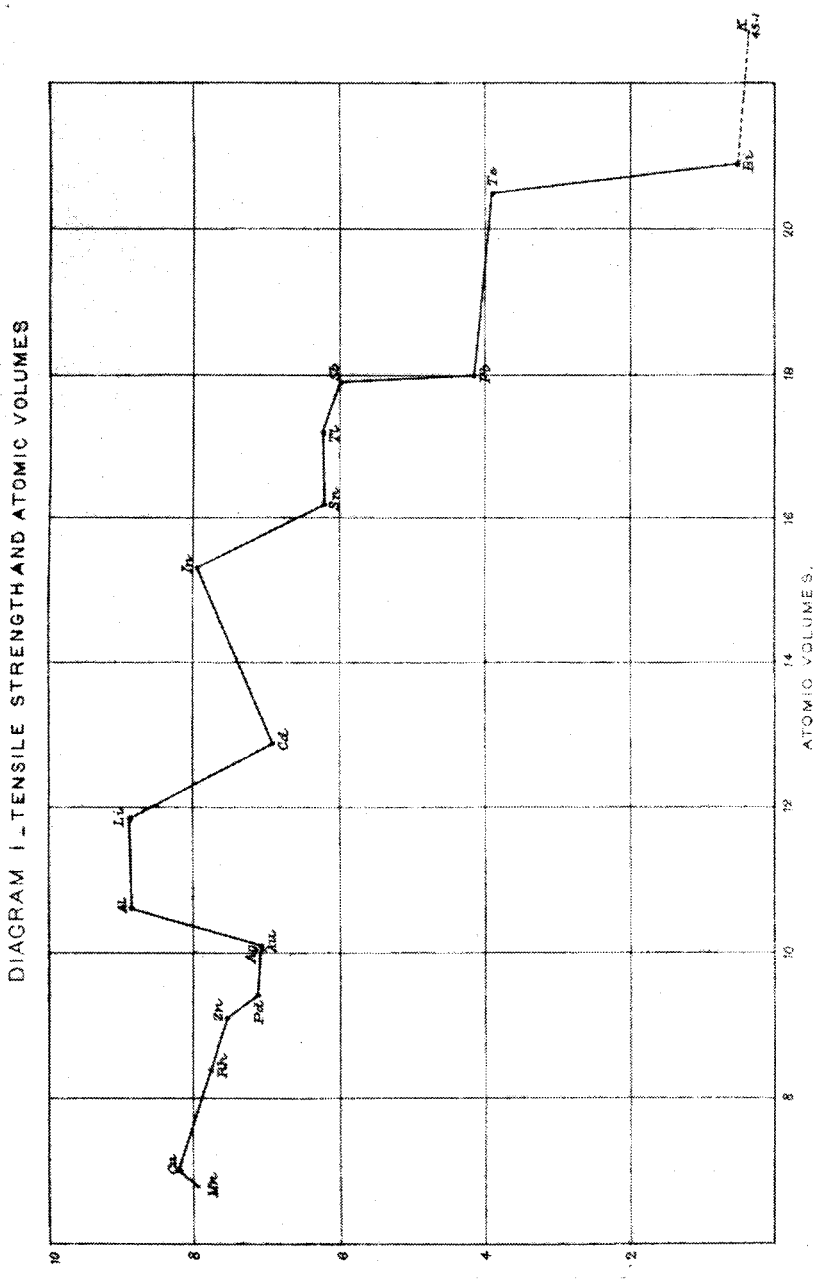

Fig. 1. Tensile strength of gold as a function of the atomic volume of the added impurity. Source: Roberts-Austen (1888).

fundamental understanding of material response. Inspiration for this approach has come on two main fronts. First, engineers continue to pursue the design of ever more complex devices and structures. This trend, coupled with financial and environmental constraints, has resulted in a strong drive towards optimal design. One of the key ways optimization may be achieved is by tailoring the materials making up a device to best serve the desired application. This trend began with the introduction of composite materials and may well continue with the creation of completely novel materials by direct engineering of their microstructure. To this end, phenomenology is of limited use since it has little predictive capability outside of its fitted range. Such efforts will, instead, require a comprehensive understanding of the effect of internal structure on macroscopic response. The second reason is tied to another important trend in engineering, that of miniaturization. This trend began with the miniaturization of electronic components and has been carried over into mechanics with the advent of microelectromechanical systems (MEMS). As the structural components in MEMS continue to shrink, with current device sizes of the order of microns, it is no longer possible to ignore a material's internal structure since it may be of the same length scale as the devices themselves. Similarly, the times over which processes of interest occur have been reduced with concomitant demands on the modeler.

Many disciplines stand to gain from an increased fundamental understanding of material behavior, from engineers engaged in the design of ceramic turbine blades for jet engines to materials scientists engaged in the construction of new alloys to nanoengineers crafting tomorrow's sub-micron motors. The task is formidable, however, as a result of the vast discrepancy between atomic time and length scales which control fundamental deformation mechanisms and the macroscopic observable response they engender. Consider, for example, the experiments of Roberts-Austen. What are the processes that take place inside a gold specimen containing a distribution of impurities as it is slowly pulled to failure?

In order to answer this question it is first necessary to consider the internal structure of such a specimen. While macroscopically it may appear homogeneous and isotropic, it actually possesses a complex internal structure. A typical tensile specimen is polycrystalline, i.e. composed of a very large number of grains separated by grain boundaries. The grains themselves, each of which is a single anisotropic crystal, contain an array of preexisting defects. At the nanoscale, we may expect vacancies, dislocations, stacking faults and twin boundaries. On a larger microstructural scale we may find microvoids, microcracks and larger cavities. The impurities may have segregated to grain boundaries or other defect sites, taken up interstitial sites in the lattice or coalesced to form small inclusions (second phase particles).

When the material is loaded, all of the above microstructural features react and interact in concert to yield those features observed macroscopically. Initially the response will be elastic and characterized primarily by lattice straining, which from an atomic perspective corresponds to bond stretching and bending. As the stress increases, some dislocations on preferred slip planes will begin to overcome the lattice resistance to their motion and move short distances before they are trapped again by entanglement with other dislocations, pile ups against interfaces and encounters with second phase particles. As the stress increases further, the material *yields* as the pinned dislocations overcome the obstacles, carrying with them increments of irreversible plastic deformation. Dislocation sources will be activated, increasing the dislocation density and making the motion of mobile dislocations increasingly more difficult — the material *hardens*. At the same time, voids may begin to nucleate at grain boundaries or at second phase particles. Vacancies diffuse and coalesce to these voids resulting in their growth. Grain boundary migration and sliding may contribute to the overall deformation as well. In addition, dislocation cell structures may form, affecting the hardening response and possibly acting as nucleation sites for microcracks at their walls. In the context of all of these processes, impurities will be present mediating many of them, such as the mobility of dislocations and grain boundaries. Temperature and strain rate will also play important roles in setting the response. These processes will continue until extensive damage leads to fracture and final failure.

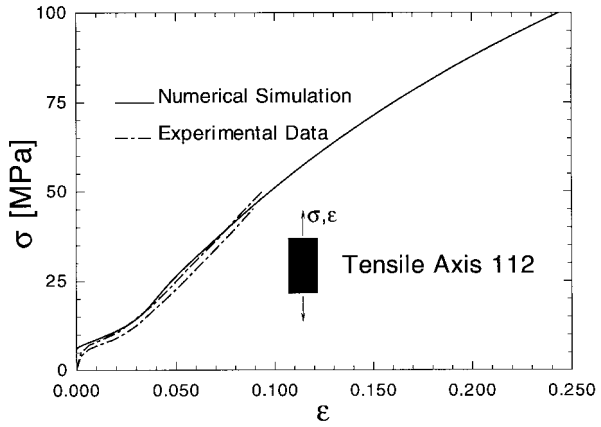

Fig. 2. Comparison of experimental measurement and numerical simulation for tension of a copper single crystal in a [112] direction. Source: Cuitiño and Ortiz (1992), reprinted with permission of IOP Publishing Limited.

The key observation from the above discussion is that the complex macroscopic response of a material results from the simultaneous operation of a host of different defect mechanisms occurring at widely different spatial and temporal scales. It will thus be very difficult, if not impossible, to formulate a single microscopically based theory for material response which includes all of these effects. A more natural alternative is to characterize the material through a hierarchy of models ranging from atomic-scale simulations to macroscopic continuum models. In this approach, each model will be active at a different scale, supplying parameters and motivating the phenomenology of the models above it in the hierarchy. To be successful, such an approach requires the close cooperation and collaboration of modelers and experimentalists working at the different scales. One thread in this endeavor is focused on developing a quantitative theory of single-crystal plasticity through a three tier hierarchy: atomic-scale simulation, dislocation dynamics and crystal plasticity theory. This hierarchy is discussed in detail in the next section. Section 3 concludes with a discussion of future research directions that are necessary to extend this hierarchy to larger scales as well as to deepen the theoretical understanding that serves as its foundation.

2. Hierarchical modeling — current research

Contemplation of a process such as ductile failure of a material is daunting in that it poses conceptual challenges over a vast range of length and time scales. One way of attempting to bridge these scales is hierarchically, where by starting with processes at the smallest length and time scales, an attempt is made to feed information across scales. The purpose of the present section is to examine a few steps in this hierarchy with special reference to one of the processes pertinent to ductile rupture, namely, the plastic deformation of a single-crystal grain. In particular, we will examine the connections between macroscopic crystal plasticity theory, mesoscopic dislocation dynamics and microscopic atomic-scale simulations, seeking channels for passing information between them.

2.1. Crystal plasticity

The end goal for the hierarchy of models we discuss here is the development of a robust continuum theory of crystal plasticity, placed on sound microscopic footing, free of phenomenology and capable of

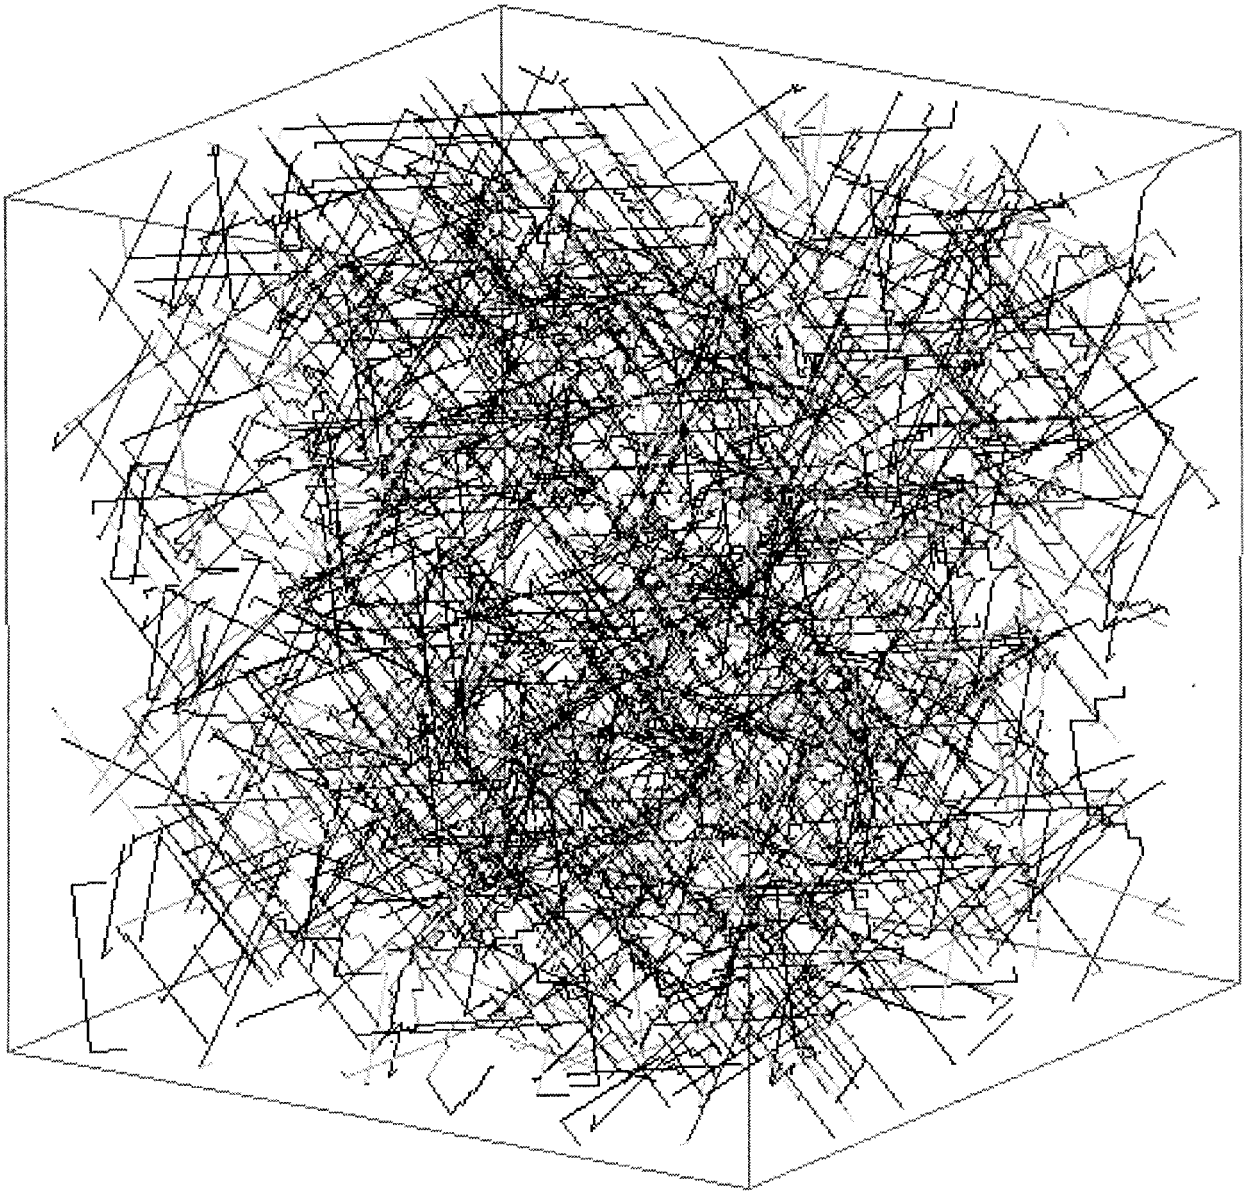

Fig. 3. Computational cell used within the context of dislocation dynamics models. The thin lines in the box are dislocation segments. Source: M. Fivel, private communication.

predicting the macroscopic response observed in tensile tests of single crystals. Much progress has already been made in this direction. Sophisticated theories of crystal plasticity exist where both experimental and theoretical considerations have led to microscopically motivated constitutive relations (Cuitiño and Ortiz, 1992; Bassani, 1993). See, for example, the comparison in Fig. 2 between an experimental tensile test and numerical predictions for a copper single crystal. In these theories, the kinematic description of the solid is characterized by the shear strains that are produced by the passage

of dislocations on particular slip planes. The shear strain rate on a system is posited to be a function of the resolved shear stress on that system as well as on the instantaneous flow stress on that system which must be overcome for dislocations to glide. Hardening of the crystal is incorporated through the assumption that the flow stress on a given system increases as a result of accumulated slip on that system as well as on all other systems. This is normally introduced through an instantaneous linear hardening relation which couples the rate of change of the flow stress to the strain rate on all systems through a proportionality factor referred to as the *hardening matrix*. The components of this matrix are not constant material properties, but rather they continuously change as the deformation evolves and depend on quantities such as the accumulated slip, flow stresses and dislocation densities on all systems. Predictive functional forms for these dependences have been developed based on arguments from statistical mechanics. One approach in the attempt to inform models of single crystal plasticity on the basis of mechanistic understanding is to seek an understanding of processes such as hardening directly on the basis of the dynamics of interacting dislocations.

2.2. Dislocation dynamics

The numerical simulation of the interaction of large numbers of dislocations and the patterns and structures they form is called *dislocation dynamics* (Devincre and Kubin, 1994; Bulatov et al., 1998). In these simulations, a computational cell is filled with dislocations on the different active slip systems which are then allowed to evolve according to a set of specified rules. Recently, these simulations have been extended to three dimensions offering a real possibility for obtaining quantitative information. At this point, such calculations are routinely made on computational cells with edge lengths of more than 10 microns. An example of such a computational cell and its complement of dislocations is shown in Fig. 3.

The key concept in dislocation dynamics simulations is the discretization of space and time. Discretization of space is achieved by dividing the dislocations into discrete segments which serve as the fundamental degrees of freedom in the simulation. Similarly, time is discretized into steps over which the dislocation patterns evolve according to well-defined rules. The interactions between dislocation segments are divided into two parts, long and short-range. Long-range interactions are treated elastically, based on classical dislocation theory. Short-range interactions such as dislocation annihilation and junction formation are introduced through a set of rules deduced from experimental observations. In addition, basic dislocation properties such as line tension, mobility and cross-slip probabilities must be accounted for.

The utility of such simulations lies in the fact that concepts treated phenomenologically on the crystal plasticity level may be studied explicitly within the confines of dislocation dynamics. For example, the crystal may be loaded in such a way that a single slip system becomes active, as a result of which the evolution of dislocation densities on all systems and the hardening of the active system can be 'measured' numerically. In addition, these simulations may shed light on the effect of dislocation patterning on hardening, an effect which is often neglected by crystal plasticity theories but is known to be significant. In the future, it may also be possible to incorporate other defects into the simulations such as vacancies, impurities and grain boundaries and to examine their influence on plastic response.

While these simulations are more fundamental in nature than crystal plasticity theory, they carry their own burden of phenomenology. For example, in current dislocation dynamics models, the laws governing the stress and temperature dependence of the dislocation velocity for a single dislocation are primarily phenomenological. Similarly, when two dislocations come into close proximity, they are able to form a junction which effectively binds them until a sufficient level of stress is reached to break the junction and send them on their respective ways. The rules for junction formation and strength are again phenomenological and constitute external input into models of dislocation dynamics. To gain

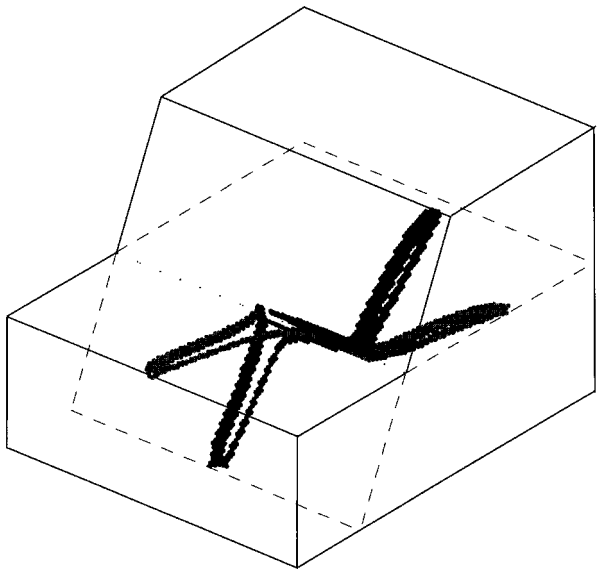

Fig. 4. Lomer–Cottrell junction computed using a quasicontinuum simulation. Source: D. Rodney, private communication.

insight into these quantities we must turn our eyes to a smaller scale and a lower level in the modeling hierarchy — that of atomic-scale simulation.

2.3. Atomic-scale simulation

The basic assertion of atomic level models is that ultimately the response of materials may be traced to the underlying motions and interactions of atoms (see Vitek, 1994 for an overview). In an atomistic simulation, the crystalline solid is represented explicitly as a collection of atoms occupying lattice sites in space. Based on the material in question, an appropriate model for atomic interactions is selected. Models ranging from fundamental and highly accurate quantum mechanical methods to empirical classical potentials are available. The choice depends on the complexity of the material in question and the size of the system that is being modeled. In general, the more accurate a method is over a broad range of configurations, the more computationally intensive it becomes. Currently, simulation cells used in both static energy minimization and finite temperature methods such as molecular dynamics range from those containing less than a hundred atoms for the most accurate descriptions to those with dimensions on the order of 0.1 micron for the simplest models of interatomic interaction.

Once the geometry is specified, initial conditions and boundary conditions are applied and the atoms are released and allowed to interact according to the specified model. Most simulations are built around two broad classes of models, statics and dynamics. The former is usually carried out at zero temperature and any loads are assumed to be quasistatic and the atomic level response is governed by the minimization of total potential energy — this is the method of *lattice statics* (LS). The other broad class of models consider dynamics at finite temperature where atomic motion is obtained by direct integration of Newton's second law — this is referred to as *molecular dynamics* (MD). Both methods have their advantages. LS, for example, may be used to obtain energy barriers for processes which, through the use of transition state theory, may lead to kinetic relations. MD may be used to study inherently dynamic processes such as crack propagation.

Just as one can construct a relation between dislocation dynamics and continuum crystal plasticity,

atomistic simulations can serve as a test bed for investigating and quantifying some of the *ad hoc* rules associated with dislocation dynamics simulations. The key virtue of atomistic approaches is that they explicitly account for the nonlinear effects found in the immediate vicinity of a given defect. As a result, the nonlinear interactions between defects, such as the formation of a junction by two dislocations, may be modeled directly and quantitative information, such as the junction's strength, may be extracted. The main limitation of this approach is that the number of atoms can become very large when treating realistic system sizes. A cubic micron of a generic metal contains on the order of 10^{11} atoms. This becomes especially problematic when the more accurate atomic interaction models requisite for quantitative detail are employed.

One possible remedy for this problem comes from the recognition that it is only in the immediate vicinity of the defect core that traditional continuum descriptions break down. A number of mixed atomistic/continuum methods have been set forth with the ambition of resolving the discrete lattice effects where needed, while allowing for some reduced description elsewhere. One particular example of this type of mixed atomistic/continuum modeling is the quasicontinuum method (Tadmor et al., 1996) within which the finite element method is used as a way of constraining the positions of vast numbers of atoms, effectively removing them from explicit consideration. This kinematic description is supplemented by the use of atomistic calculations as the basis of the material response, lending the method both the requisite spatial resolution and degree of freedom reduction to handle problems such as the formation of dislocation junctions. One example of this type of analysis is shown in Fig. 4. This figure shows the two different glide planes associated with the reacting dislocations as well as the highly energetic atoms present in the dislocation core regions of the reacting dislocations. The crucial observation is that along the line of intersection of the two glide planes, a junction segment is formed with a Burgers vector that is the vector sum of the Burgers vectors of the parent segments. In the discussion above, we have traced three steps in the chain of models connecting the fundamental response of a solid governed by the interacting atoms to the macroscopic ductile rupture of real materials. The hierarchy proceeds further in both directions. Starting with insights at the single crystal plasticity level, it is possible to invoke homogenization theories which attempt to intelligently average single crystal plasticity models to polycrystalline ensembles. Similarly, proceeding from the atomic scale elucidated above, one may proceed to the most fundamental quantum mechanical descriptions and use these to construct useful approximate descriptions of atomic interaction. Clearly, there is large scope for research at the different levels of this hierarchy, with the aim being to connect the different levels into a continuous viable model of material response.

3. Future research directions

The primary purpose of this article has been to recount one of the most exciting threads of research in the mechanics of materials, namely, the attempt to develop models of material response that are constructed on the basis of elementary mechanisms rather than curve fitting. The strategy proposed to face this challenge is to set up a hierarchy of modeling efforts in which models at one scale are used to inform those at another. In the previous section we gave a concrete example of this approach as it applies to the development of a quantitative theory of single crystal plasticity. In this section we go deeper, discussing directions for research that have received little attention so far, but which will have to be pursued in order to explain the behavior of real materials, such as the problem of ductile rupture with which we began this article.

3.1. Bridging temporal scales

Until now, our discussion has focused on the difficulty of linking mechanisms that are operating at different length scales. However, as we have alluded to repeatedly, there are also a host of outstanding questions concerning the appropriate treatment of problems in the temporal domain. Many processes that occur in materials, most notably those involving mass transport, take place over time scales that are many orders of magnitude larger than the vibrational periods of atoms. At the same time, gross macroscopic diffusion, which plays an important role in phenomena such as creep, owes its existence to these atomic level processes. A research theme of profound current importance concerns the construction of fundamentally based models of microstructural evolution in a way that respects the disparate temporal scales presented by these problems. A less well-developed research area, but clearly an exciting one, concerns the construction of hierarchies like those sketched above, but now in the temporal domain, beginning with the atomic scale jump processes that characterize diffusion and passing to the long time creep processes that plague materials in a variety of engineering applications.

3.2. Chemistry effects

Real materials are, for the most part, not the mono-atomic solids discussed in Section 2, but rather complex systems containing many different atomic species. The addition of a small amount of some impurity to a material can dramatically affect properties such as its yield strength, fracture toughness and tensile strength (as demonstrated in Fig. 1). The choice of different material combinations to yield desirable properties is the fundamental concern of alloy design where complex mixtures of impurities are concocted to obtain desired properties through experience, trial-and-error-based empirical approaches and intuition. In addition, real materials operate in environments which, over time, can lead to changes in material properties, e.g. through hydrogen embrittlement. In order to develop predictive fundamental models which account for chemistry effects it will be necessary to investigate the role played by impurities in processes such as dislocation motion, grain boundary mobility, void nucleation and the strength and structure of grain boundaries.

3.3. Microstructural evolution

In addition to the complexity introduced into material behavior as a result of the presence of more than a single chemical species, material response is often dictated by structure at the microstructural scale. In practice, what often separates the response of one alloy from another is the difference in their microstructures. Thus the design of a new alloy (material) sometimes corresponds to the design of a novel microstructure. In addition to the dream of using the computer to suggest new material combinations, it is also clear that the processing that gives rise to particular microstructures requires modeling as well (see Rappaz and Gandin, 1993 for an example). Such modeling requires a rigorous understanding of processes such as recovery, recrystallization, grain growth, and coarsening. These are long time processes which are also strongly influenced by chemistry effects and will thus require advances in both of the previous categories.

3.4. Interactions of multiple defects

An allied problem to that of microstructural evolution is that of determining the effects of a given microstructure on the physical properties of the material. Recent years have seen significant progress in the treatment of isolated defects from a fundamental perspective. It is now possible to obtain the structure and energetics of most of the key defects in the mechanics of materials. In addition, in

conjunction with transition state theory, it is also possible to construct plausible insights into the kinetics of defect motion. However, a wide variety of open problems remain concerning how various defects conspire in their mutual interactions and their interactions with the underlying microstructure to produce the observed macroscopic response. From the standpoint of the hierarchical perspective adopted here, many of the key challenges to be faced in closing the hierarchy come in two principal areas: 1) the treatment of multiple defects simultaneously in a statistically consistent way, and 2) the analysis of lattice defects from the perspective of their ability to form and interact with features of the material at the level of its microstructure.

3.5. Nanotechnology

Most of the focus of this article has been on the development of predictive models of material response. However, it should be noted that much can be learned even before a full understanding is reached. One example where this is true is seen in the field of nanotechnology where the components are small enough that individual mechanisms are important and hence it is not necessary to rely on the type of statistical averaging alluded to above. For example, it is rapidly becoming possible to use atomic-scale methods to model entire MEMS devices. Such approaches can be used to study the effect of free surfaces and interfaces on such devices as well as to identify failure mechanisms. It will also be possible to use atomistics to validate higher level continuum theories currently used to design MEMS, identify the limits of their applicability and possibly suggest appropriate phenomenology at higher scales.

3.6. Coupled problems

An entire range of materials applications rely on the coupling between mechanical properties and other features such as electrical, magnetic, thermal or optical properties. An important trend reflecting this coupling is the development of *smart* materials. These are materials that can function as both sensors and actuators depending on the situation (Newnham, 1997). The physics underlying such materials usually involves couplings between different physical phenomena. For example, there is *electromechanical* coupling in piezoelectric and electrostrictive materials where the mechanical response of the material is closely tied to its electrical properties. Magnetostrictive materials exhibit a coupling between mechanical and magnetic fields. Shape memory alloys capitalize on a *thermomechanical* effect, whereby the material regains its undeformed shape after being heated, and the list goes on (see Newnham, 1997 for details). There can also be a coupling between mechanical loading and the internal structure of the material, as in stress-induced phase transformations, stress-induced metal–insulator transformations and phonon and quantum effects which can influence the elastic constants of the material. This class of problems demonstrates unequivocally the growing role of interdisciplinary insights in the modeling of materials.

3.7. Nanoscale experimentation

We have identified a number of key areas in which modeling in the mechanics of materials still presents challenges. Since much of the work considered is primarily computational, another key thrust in this area has to be careful experimental validation of these models. Though the use of computers in mechanics has provided many insights, it also makes the task of determining the correctness of a given model more difficult. Consequently, the final and perhaps the most important step in the hierarchy outlined above is to bring the predictions of such models into direct contact with the experiments themselves (see Yu and Pennycook, 1997 for a discussion of experimental techniques at the nanoscale).

The fundamental perspective of this article is that we have reached an exciting crossroads in the

mechanics of materials which reflects the confluence of threads coming from materials science, condensed matter physics and the more traditional discipline of solid mechanics. One of the key issues to be faced at this time concerns the construction of models of material response based on fundamental mechanistic understanding rather than enlightened empiricism. It appears that real progress in pushing materials to new extremes in size, speed and versatility will require hierarchical modeling of a wide variety of processes involving not only the mechanical properties of materials, but the coupling of mechanical response to other properties.

Acknowledgements

The authors are grateful to J. Mullin for her help and valuable comments and to M. Fivel and D. Rodney for the use of their figures.

References

Bassani, J.L., 1993. Plastic flow of crystals. Adv. Appl. Mech 30, 191–258 (This article provides a comprehensive, in depth, discussion of crystal plasticity).

Bulatov, V., Abraham, F.F., Kubin, L., Devincre, B., Yip, S., 1998. Connecting atomistic and mesoscale simulations of crystal plasticity. Nature 391, 669–671 (This article concerns that part of the modeling hierarchy connecting dislocation dynamics to atomistic simulation).

Cuitiño, A.M., Ortiz, M., 1992. Computational modelling of single crystals. Modelling. Simul. Mater. Sci. Eng. 1, 225–263 (A review article discussing computationally tractable models of crystalline plasticity).

Devincre, L., Kubin, L.P., 1994. Simulations of forest interactions and strain hardening in fcc crystals. Modelling Simul. Mater. Sci. Eng. 2, 559–570 (This paper spells out the ways in which dislocation dynamics codes may be used to examine hardening in crystals).

Newnham, R.E., 1997. Molecular mechanisms in smart materials. MRS Bulletin 22 (5), 20–33 (This article gives an introduction to the field of smart materials, explaining what they are, how they work and what applications they are good for).

Rappaz, M., Gandin, Ch.-A., 1993. Probabilistic modelling of microstructure formation in solidification processes. Acta Met. 41, 345–360 (Rappaz and Gandin address a hierarchical approach to microstructural evolution with the coupling of a finite element approach and a cellular automaton).

Roberts-Austen, W.C., 1888. On certain mechanical properties of metals considered in relation to the periodic law. Phil. Trans. R. Soc. Lond. A 179, 339–349 (This paper describes Roberts-Austen experiments on the effect of added impurities on the tensile strength of gold).

Tadmor, E.B., Ortiz, M., Phillips, R., 1996. Quasicontinuum analysis of defects in solids. Phil. Mag. A73, 1529–1563 (This paper demonstrates how atomistic simulation may be coupled with finite element interpolation to construct a mixed atomistic and continuum description of deformation in solids).

Vitek, V., 1994. Atomistic modeling: an overview. In: The Encyclopedia of Advanced Materials, 1st ed. Pergamon, Oxford, pp. 153–161 (An overview of atomistic modeling, describing methods for characterizing atomic response and giving examples of typical simulations).

Yu, E.T., Pennycook, S.J., 1997. Nanoscale characterization of materials. MRS Bulletin 22 (8), 17–62 (The entire issue is dedicated to experimental techniques at the atomic and nano-scales).