COMPUTERS IN PHYSICS

Parallel Computing

Massively parallel computer simulation of plane-strain elastic–plastic flow via nonequilibrium molecular dynamics and Lagrangian continuum mechanics

William G. Hoover
Department of Applied Science, University of California at Davis/Livermore and Precision Engineering Program, Lawrence Livermore National Laboratory, Livermore, California 94550

Anthony J. De Groot
Engineering Research Division, Lawrence Livermore National Laboratory, Livermore, California 94550

Carol G. Hoover
Methods Development Group, Mechanical Engineering Department, Lawrence Livermore National Laboratory, Livermore, California 94550

Application of massively parallel low-cost computers to the simulation of plane-strain elastic-plastic flow is discussed. Two different approaches, atomistic molecular dynamics and continuum mechanics, are applied to this problem. A hybrid scheme combining the two is also discussed.

INTRODUCTION

Low-cost massively parallel computers are ideally suited to the large-scale simulation of high-strain-rate elastic–plastic flows. The underlying numerical methods are straightforward and the results can be visualized relatively easily with color videotape. Low-cost parallel computers are revolutionizing computational physics and materials science. At Livermore we have applied two such machines, the SPRINT (systolic processor with reconfigurable interconnection network of transputers) and the Butterfly, to various simulation problems in irreversible solid-phase deformation using both atomistic molecular dynamics and Lagrangian continuum mechanics.[1-4]

Our long-standing interest in molecular simulation[5] has been broadened by the possibilities of much larger and more realistic simulations including hybrid numerical schemes which combine the atomistic and continuum viewpoints. Here, we describe both parent approaches, microscopic and macroscopic, and offer a preview of their combination.

Alder and Wainwright's pioneering molecular dynamics simulations were purely Newtonian, and used unrealistic hard-sphere and square-well force laws in order to save computer time. By 1959 machines were faster. Vineyard used nonequilibrium boundary conditions in simulating radiation damage in metals. Vineyard used purely repulsive force laws too. But these were more realistic, being continuous rather than impulsive. In 1964 Rahman[6] was able to simulate a *liquid* phase using continuous forces including both repulsions and attractions.

A systematic approach to *non*equilibrium molecular dynamics with continuous forces began about 1970 when Ashurst set out to simulate steady-state nonequilibrium flows of momentum and energy in nonequilibrium liquids. In these flows a few hundred bulk atoms were driven by smaller numbers of atoms confined to boundary heat and momentum reservoir regions.[5] By 1980 these nonequilibrium simulation methods were being used to generate and characterize strong dense-fluid shockwaves involving as many as 4800 atoms.[7]

In 1984 Abraham pioneered the simulation of *big* systems. This work, with up to 161 604 atoms,[8] was a conceptual breakthrough. Today "*massively parallel*" computation makes such a simulation a routine reality. Parallelism is achieved by spanning the problem, which can be either atomistic or continuum, with a computational grid, assigning the contents of the "zones" so defined to individual processors. *Million*-atom simulations[3,9] have already appeared. *Billion*-atom simulations should follow soon, with revolutionary impact on the ability of computing machines to simulate the processing of real materials.

Macroscopic continuum simulations are more familiar and widespread than microscopic atomistic ones. The finite-element and finite-difference approaches to engineering design were among the first applications of fast computers. These numerical approaches have been applied to challenging problems in the design of aircraft, buildings, ships, and weapons for half a century. Until recently, the relative complexity of atomistic systems coupled with persistent uncertainties in the underlying forces prevented a corresponding microscopic approach from contributing to engineering design. Through enhanced computational power and more imaginative models of atomic forces, the situation is changing rapidly. A combination of continuum

and atomistic simulations appears to be within our grasp. The trail to such combinations has already been blazed by various quasistatic studies of lattice defect structures.[10]

De Groot's SPRINT computer is shown in Fig. 1. SPRINT is a pioneering realization of the promise of low-cost parallel computation.[11] The 1988 version of this machine included 64 transputer processors and cost about $40 000 dollars to build. A million-atom time step requires about a minute, using the Lennard–Jones–spline potential discussed in Sec. II, and about 80 s for the embedded-atom modification appropriate to simple metals such as copper. The transputers in SPRINT execute instructions in FORTRAN, PASCAL, and C, and have a storage capacity of one-quarter million 32-bit words each. Current chips and the next generation of transputers are both an order of magnitude faster than the transputers in SPRINT. A straightforward scaling of the parallelism used in SPRINT should lead soon to teraflop computers.

These computational developments have been paralleled by revolutionary conceptual developments in understanding the solutions of the large systems of ordinary differential equations encountered in nonequilibrium statistical mechanics.[4] Here, we begin by reviewing the principles underlying both the atomistic and the continuum simulations. We then present a sampling of our deformation results and a preview description of our planned hybrid approach, which combines the microscopic and macroscopic techniques.
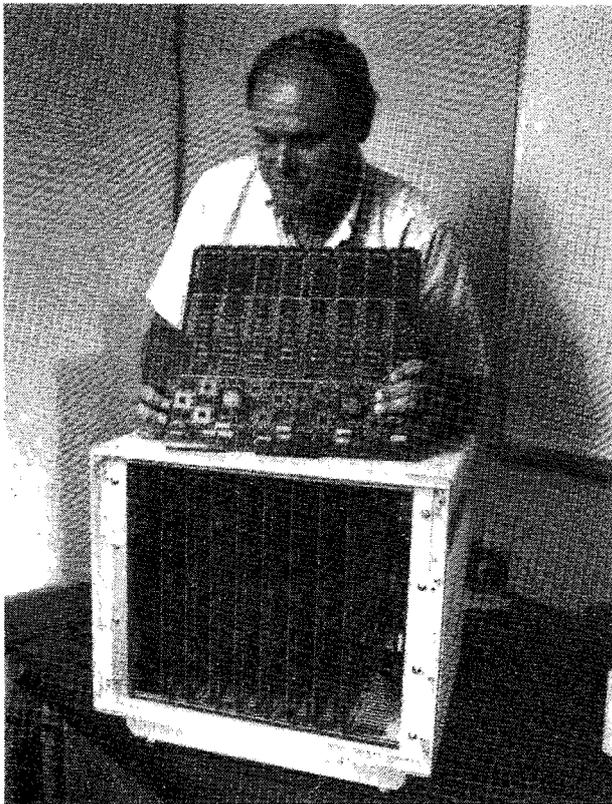

FIG. 1. The SPRINT computer at Livermore. This machine contains 64 transputers (small computers) and matches the performance of a modern CRAY-YMP supercomputer at greatly reduced cost.

I. FUNDAMENTAL PRINCIPLES OF COMPUTER SIMULATION[4]

Both the microscopic atomistic picture and the macroscopic continuum picture rely on *differential* equations of motion describing the time-and-space developments of the system coordinates. Of the two the atomistic picture *seems* simpler because it is based on *ordinary* differential equations of motion in which time is the only independent variable:

$$m\ddot{x} \equiv m\frac{d^2x}{dt^2} = F_{\text{atomic}} + F_{\text{boundary}}$$
$$+ F_{\text{constraint}} + F_{\text{driving}}. \tag{1}$$

The interparticle *atomic* forces F_{atomic} are the basic ingredient of the atomistic picture. Additional *boundary* and *driving* forces can be used to confine the system particles and to perform thermodynamic work on them. Likewise *constraint* forces can be used to force definite instantaneous or time-averaged values on time-and-space-dependent fluxes and state variables. Such constraints are necessary whenever a steady nonequilibrium process is considered. Extending and constraining Newtonian mechanics by adding thermostats is required *whenever* heat flow must be simulated. Here, we again[3] use Nosé–Hoover thermostat constraints[4] to impose a desired temperature on our system. As a bonus, this particular mechanical description of nonequilibrium heat flow also provides an exact microscopic demonstration of the macroscopic Second Law of Thermodynamics.[4,12] Only those flows of heat consistent with the Second Law can provide stable solutions of the Nosé–Hoover equations of motion (1); see Eq. (7) below.

The continuum picture is based on *partial* differential equations:

$$\frac{d\ln\rho}{dt} = -\nabla\cdot\mathbf{v}, \quad \rho\ddot{\mathbf{x}} \equiv \nabla\cdot\boldsymbol{\sigma} \equiv -\nabla\cdot\mathbf{P},$$
$$\rho\dot{e} = -\nabla\cdot\mathbf{Q} - \mathbf{P}{:}\nabla\mathbf{v}, \tag{2}$$

in which both time and space vary independently. Here, $\mathbf{x}$ is the inertial-frame coordinate, $\mathbf{v}$ is the velocity, $\mathbf{v}\equiv\dot{\mathbf{x}}$, $\mathbf{P}$ is the pressure tensor, and $\mathbf{Q}$ is the heat flux. In (2) ρ is the mass density and e is the energy per unit mass. In general, these hydrodynamic variables depend not only on time but also on space.

If the spatial gradients are sufficiently small and fluctuations can be ignored, then the two approaches, microscopic and macroscopic, become equivalent. The atomistic description can be linked to the continuum one through the constitutive "equation-of-state" relations which distinguish one material from another. These relations describe the dependence of the comoving momentum and energy fluxes, $\mathbf{P}$ and $\mathbf{Q}$, on the strains, strain rates, and temperature gradients. To eliminate the effect of fluctuations, the microscopic behavior needs to be averaged, either over an "ensemble," equilibrium or steady-state, or, equivalently, over time. The simplest example of such an averaged macroscopic constitutive relation is Fourier's law relating the heat flux to the temperature gradient:

$$\mathbf{Q} = -\kappa\nabla T. \tag{3}$$

To begin, choose a form for the interatomic forces and consider the corresponding material in a state with fixed density and fixed average temperature. By carrying out atom-

istic simulations based on the equation of motion (1) and measuring the dependence of the heat flux vector $\mathbf{Q}$ on the temperature gradient ∇T, the macroscopic thermal conductivity κ defined by (3) becomes a known function of material and state. Likewise, by measuring stress, $\boldsymbol{\sigma} \equiv -\mathbf{P}$, as a function of density, energy, and shape, atomistic simulations can provide the constitutive information required to solve the continuum equation of motion. The simplest constitutive relation for stress describes an elastic solid:

$$\boldsymbol{\sigma} \equiv [-P_{eq}(\rho, e) + \lambda \nabla \cdot \mathbf{u}]\mathbf{I} + \eta[\nabla \mathbf{u} + \nabla \mathbf{u}^t], \qquad (4)$$

where the stress tensor $\boldsymbol{\sigma}$ is minus the pressure tensor $\mathbf{P}$, $\mathbf{u} \equiv \mathbf{x} - \mathbf{x}_{eq}$ is the *elastic displacement*, and λ and η are the elastic solid's *Lamé constants*. The equilibrium pressure P_{eq} as well as the Lamé constants, *all* depend on the *mass density ρ* and the *energy per unit mass e*.

In our own work in support of the Precision Engineering Program at the Lawrence Livermore Laboratory we have continued our study[3] of the prototypical *plastic*-flow experiment shown in Fig. 2, the two-dimensional analog of the standard three-dimensional hardness test. Plastic flow is an irreversible change of shape induced by stress anisotropy (shear stress). By including also the continuum approach, we explore both the correspondence between the microscopic and macroscopic approaches to simulation and the convergence of both methods to the large-system or fine-mesh limits.

The two-dimensional version of the problem was selected in order to simplify visualization. Three-dimensional laboratory hardness tests involve pressing a spherical or
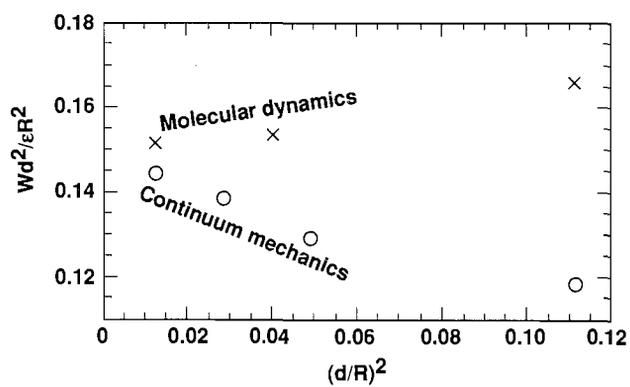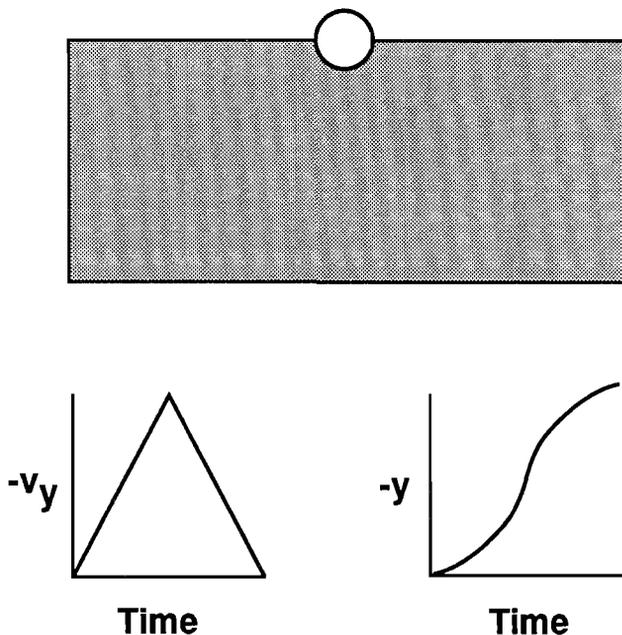

FIG. 2. Plane–strain indentation. The vertical indentor velocity v_y has a piecewise-linear time dependence. The maximum value of the indentor travel y is chosen equal to the indentor radius R. In the three-dimensional analog of this two-dimensional plane-strain problem, a half-space is indented by a cylindrical indentor with its axis parallel to the workpiece surface.

FIG. 3. Elastic work of indentation W for similar small-deformation simulations of an elastic workpiece using both molecular dynamics and continuum mechanics. The total indentor travel in each of these simulations is only $R/10$, where R is the indentor radius. The continuum simulations model *half* the workpiece with 30×26, 45×39, 60×52, and 90×78 zones. The molecular simulations model the *entire* workpiece with 60×30, 100×50, and 180×90 atoms.

prismatic indentation tool, the "indentor," into the workpiece being tested, with a known force. Typical tests are thermodynamically "quasistatic," so that the kinetic energy of the indentor can be ignored. The applied force per unit area of the resulting cavity is then a measure of the material's resistance to permanent plastic deformation. This phenomenological approach furnishes a useful reproducible characterization of solids' resistance to flow.

The correspondence of the two approaches, microscopic and macroscopic, is shown for the elastic case in Fig. 3. In that corresponding-states plot, d represents both the atomic diameter, at which the force vanishes, and the continuum zone width. The common value of the intercept, approximately 0.15, shows that the two *different* kinds of spatial discretizations of the continuum—interacting atoms or Lagrangian zones—lead to the *same* elastic modulus in the continuum limit. In this limit the atomic diameter and the Lagrangian mesh size are small relative to all macroscopic lengths. The results in the figure suggest that the deviation from the continuum result is inversely proportional to the number of degrees of freedom used to describe the problem.

II. ATOMISTIC SIMULATIONS

We have developed a computer model for the simulation of plane-strain indentation, with the basic geometry and boundary conditions indicated in Fig. 2, and began by studying a piecewise-smooth pair potential ϕ_{ljs} chosen to match the Lennard–Jones potential for separations less than the inflection point:

$$r < r_i \Rightarrow \phi_{ljs} = 4\epsilon\left(\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6\right)$$
$$\equiv \epsilon\left(\left(\frac{d}{r}\right)^{12} - 2\left(\frac{d}{r}\right)^6\right),$$
$$r_i < r < r_m \Rightarrow \phi_{ljs} = 3.292\,0028(\epsilon/\sigma^3)(r_m - r)^3$$
$$- 4.864\,890\,08(\epsilon/\sigma^2)(r_m - r)^2,$$

$$\frac{r_i}{\sigma} = \left(\frac{26}{7}\right)^{1/6} = 1.244\ 455\ 06,$$

$$\frac{r_m}{\sigma} = 1.737\ 051\ 787. \tag{5}$$

In our atomistic simulations the indentor is characterized as a smooth repulsive force center with an effective radius R. With an atomic "radius" equal to half the equilibrium spacing $d \equiv 2^{1/6}\sigma$, the force between an atom separated from the indentor by a distance r less than $R + [d/2]$ was calculated by computing an interaction potential $\phi_{\rm rep}(r + [d/2] - R)$ based on the purely repulsive part of the Lennard–Jones potential. In the present work we emphasize a particular weighted average of this "Lennard–Jones–spline" potential[13] $\phi_{\rm ljs}$ with an embedded-atom[14] interaction.[15] The relative weights we use to combine the pair and embedded-atom potentials furnish a useful description of the vacancy and defect energies exhibited by typical metals such as copper, gold, nickel, and silver:

$$\Phi \equiv (1/3)\Phi_{\rm ljs} + (2/3)\Phi_{ea}, \quad \Phi_{\rm ljs} = \Sigma\phi_{\rm ljs},$$

$$\Phi_{ea} = \Sigma\phi_{ea}(\rho) = \frac{ze\epsilon}{2}\Sigma(\rho_i \ln \rho_i), \quad \rho_i = \Sigma\rho_{ij},$$

$$\rho_{ij} = (1/ze)\left[(r_m^2 - r^2)/(r_m^2 - d^2)\right]^2. \tag{6}$$

The coordination number z is 6 in two dimensions and 12 in three. For either type of force law, Lennard–Jones–spline and $(1/3)$ Lennard–Jones–spline $+ (2/3)$ embedded-atom, we found[3] that the energy density W/R^2 required to indent a two-dimensional perfect crystal is about one-sixth the shear modulus for a plane-strain indentation with a circular indentor. Both the structure of the underlying equations of motion and the solution methods are no different in three dimensions where the present simulations correspond to periodic cylindrical indentations of a half-space.

The atomistic simulations depend upon *temperature* as well as force law and indentor velocity. In order to simulate *isothermal* deformation, we have incorporated Nosé–Hoover constraint forces[4] F_c in all out finite-temperature simulations. These constraint forces $F_c \equiv -\zeta p$ incorporate friction coefficients $\{\zeta\}$ which satisfy feedback relations forcing the time-averaged values of the corresponding kinetic energies $\{K(t)\}$ to agree with specified values $\{\langle K \rangle\}$:

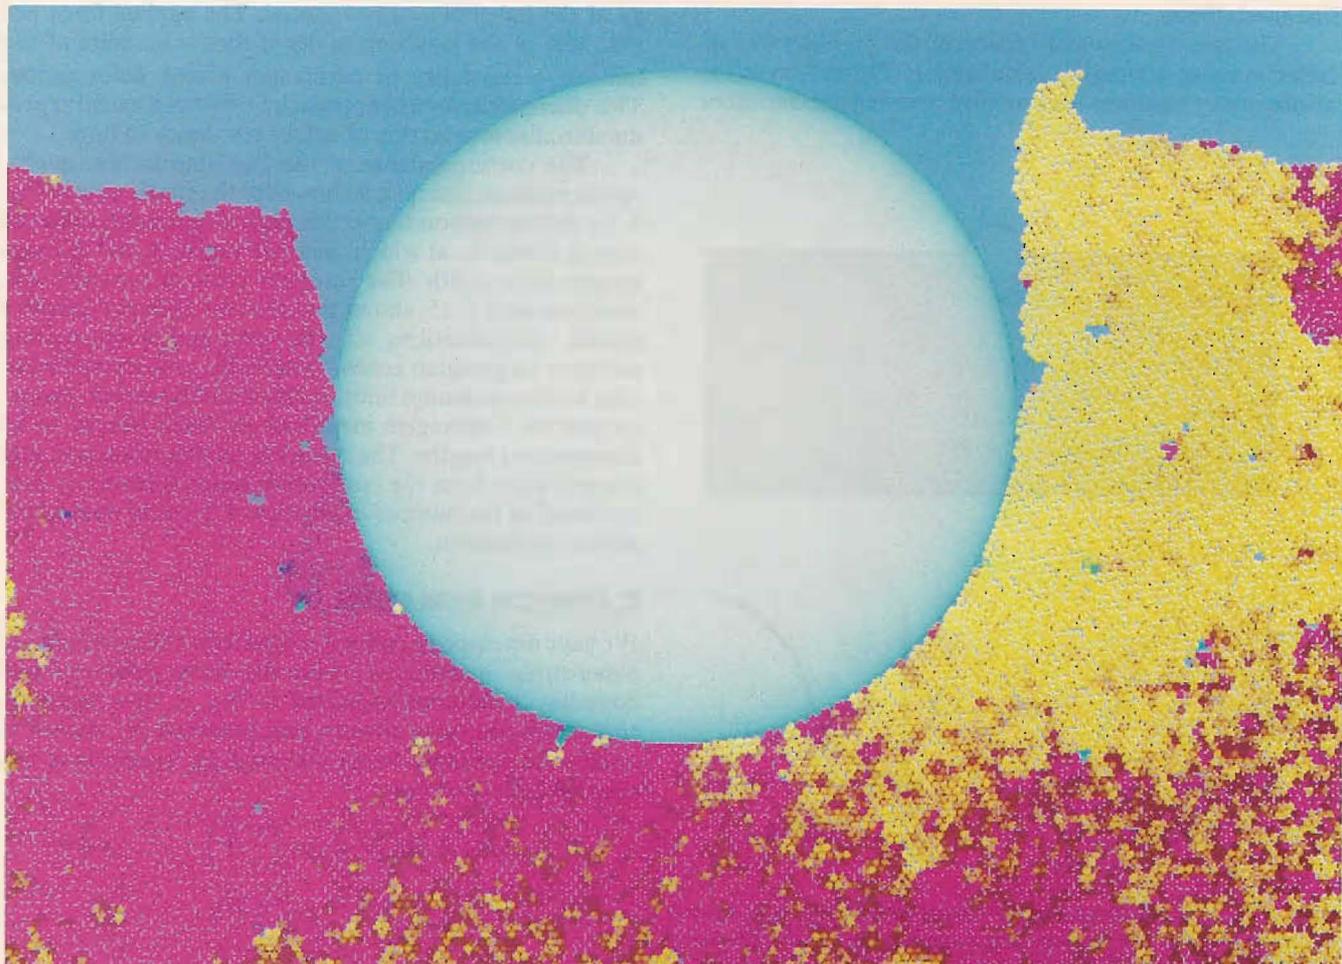$$\dot{p} = F - \zeta p, \quad \dot{\zeta} = [(K/\langle K \rangle) - 1]/\tau^2. \tag{7}$$

FIG. 4. Indentation of a 12 800-atom Lennard–Jones + embedded-atom workpiece at half the melting temperature. The indentor radius and depth of indentation are 40 d. The maximum indentation speed is $(\epsilon/m)^{1/2}$.

In finite-difference form the differential equations of motion become:[2,4]

$$m(x_+ - 2x_0 + x_-)/dt^2$$
$$= F_0 - m\zeta_0(x_+ - x_-)/(2dt),$$
$$\frac{(\zeta_+ - \zeta_-)}{2dt} = \frac{\Delta K_0/\langle K \rangle}{\tau^2},$$
$$K_0 = \frac{\Sigma m(x_+ - x_-)^2}{8dt^2}, \tag{8}$$

where the subscripts $\{-,0,+\}$ indicate the time sequence $\{t - dt, t, t + dt\}$. In two dimensions the N-body mean kinetic energy $\langle K \rangle$ is equal to NkT.

The atomistic simulations show an interesting dependence of cavity shape on force law, deformation rate, and crystal structure. We began our investigations with perfect crystals. We found, for the geometry shown in Fig. 1, that crystals containing 800 atoms or more were sufficiently large for the quasistatic work of indentation to be roughly described by the macroscopic yield strength. The *shape* of the indentation cavity is slower to converge, at least at the higher speeds, for which we simulated million-atom systems.

It is particularly interesting that the initial mirror symmetry of the atomistic indentation problem can be completely destroyed by thermal fluctuations. The underlying chaotic atomistic dynamics can have macroscopic consequences. This point is illustrated in Fig. 4, in which the shapes of the right and left sides of an initially symmetric indentation specimen are qualitatively different.

Closeups of the region near the indentor reveal that a disordered layer of atoms, many atomic diameters thick, forms along the margins of the indentor–workpiece interface.[3] It may be that this amorphous region is the two-dimensional analog of the Beilby layer[16] found at the surface of cold-worked three-dimensional metals. Despite the disordered appearance, thermodynamic melting seems to be ruled out by the relatively low temperatures involved, half the melting temperature in most of our work. Elastic-plastic *continuum* simulations, on the other hand, show a relatively featureless deformation.

The sixfold rotational symmetry of the triangular-lattice structure implies an isotropic *linear*-elastic response to *small* deformations, but the *large*-deformation *non*linear response of the same structure necessarily reflects the underlying orientation of the crystal planes. To evaluate the consequences of crystal structure, we have carried out a series of simulations on amorphous (polycrystalline) solids as well.

Amorphous atomistic solids are easy to make, following a suggestion furnished by Holian, who has synthesized and studied such solids at Los Alamos. He begins with a mechanically unstable lattice, such as the square lattice in two dimensions or a simple cubic lattice in three, and allows the unstable lattice to relax. Figures 5 and 6 show amorphous crystals grown in this way, both before and after plastic deformation. Following a suggestion from Voter, also at Los Alamos, and Holian, we have assigned a *color* to each atom in the amorphous crystal by working out the average inclination of its interacting neighbors with respect to the x axis. The averaged angle α, defined by the relation $\tan(\alpha) \equiv \Sigma \sin(6\theta)/\Sigma \cos(6\theta)$, is converted to a color. The factor 6 in the sums over interacting neighbors takes the sixfold rotational symmetry of the triangular-lattice structure into account and results in a "coloring" angle α, which ranges from $0-2\pi$.

An amorphous solid, composed of Lennard–Jones + embedded-atom material was relaxed for a few atomic vibrational periods prior to indentation, as shown in Fig. 5. Figure 6 is a closeup of a fully indented amorphous workpiece. Figure 7 shows the corresponding deformation of an initially perfectly crystalline triangular-lattice workpiece using exactly the same atomistic forces. Both cases have been colored to show the post-indentation grain structure.

Table I includes results of molecular dynamics simulations of both crystalline and amorphous solids, both using the same force laws and the same deformation schedule. We represent our results in terms of an effective microhardness or yield strength given by the specific work of deformation, W/R^2, where R is the indentor radius. The results marked n in the table correct earlier embedded-atom simulations[3] carried out at half the melting temperature. The present warm embedded-atom simulations use the correct 2.5% linear expansion from the cold-crystal lattice spacing. The previously published values for these entries mistakenly used a linear expansion only 0.75%, causing a significant error in the work of indentation.

In the *crystalline* case the data in Table I show that 3200 atoms provide a reasonable estimate of the specific work of indentation, in two dimensions. The number dependence is somewhat larger in the amorphous case, with the specific work of deformation for 12 800 atoms still exceeding the apparent large-system limit by nearly a factor of 2. The work of deforming the *amorphous* solid lies about 30% below the corresponding perfect-crystal deformation.

The main difference between the crystalline and amorphous simulations is that the appearance of the deformed solids is much less regular in the amorphous case. The data in Table I show that the work of deformation is insensitive to the details of the crystal structure. We believe that this condition reflects the highly disordered nature of the cold-worked solid inherent in large-scale plastic deformation. This fortunate circumstance should simplify the matching of the atomistic and continuum descriptions, our ultimate goal. We turn next to a continuum description of the same plane-strain deformation problem.

III. CONTINUUM SIMULATIONS

There are two common numerical approaches to continuum deformation problems. The "Eulerian" approach uses a fixed laboratory-frame coordinate system, through which the material flows. In this case it is natural to consider the *fixed-in-space* time variation of the field variables, density, stress, and energy: $\{\partial\rho/\partial t, \partial\sigma/\partial t, \partial e/\partial t\}$.

To correlate continuum results with atomistic simulations, which describe the flow of labeled particles in space, the comoving "Lagrangian" approach is more natural. In this case a *comoving* grid representing the material and subject to suitable external boundary conditions is followed in time and is described by the Lagrangian *comoving* (or "total" or "substantive") derivatives, $\{d\rho/dt, d\sigma/dt, de/dt\}$. In addition to the underlying conservation equations for mass, momentum, and energy,
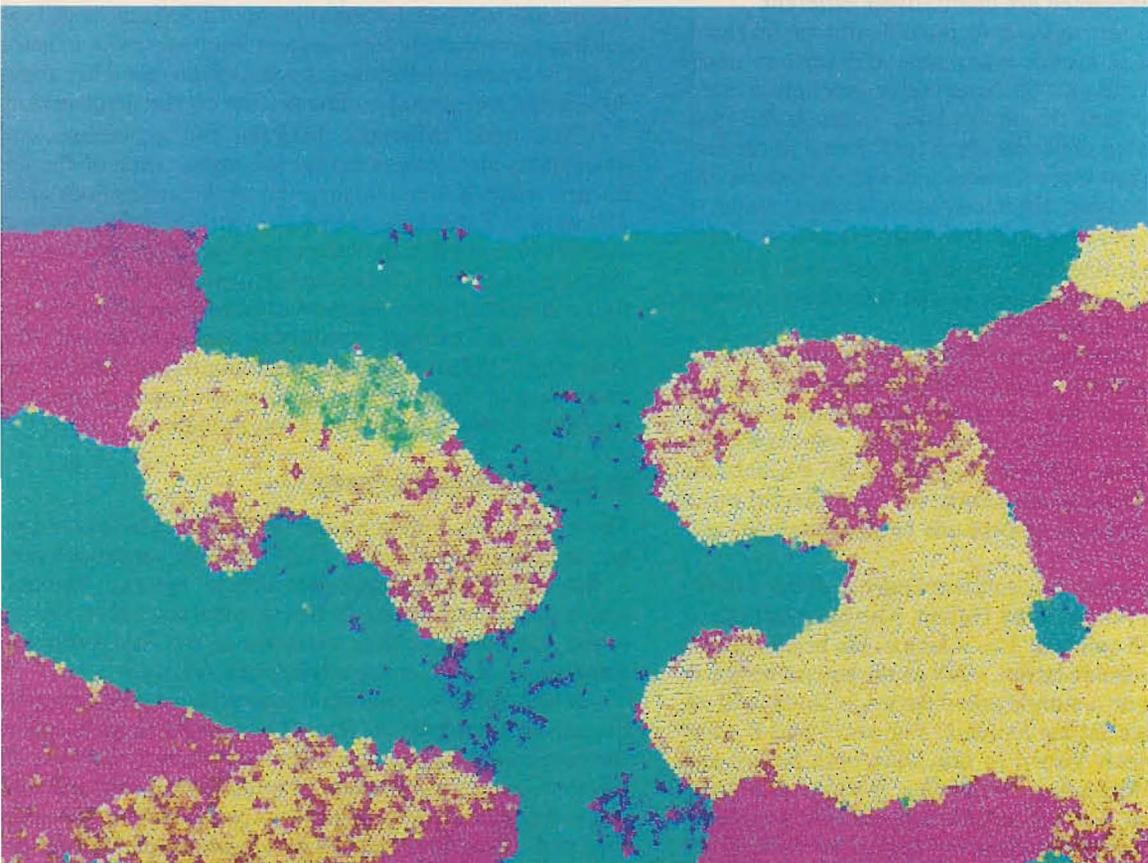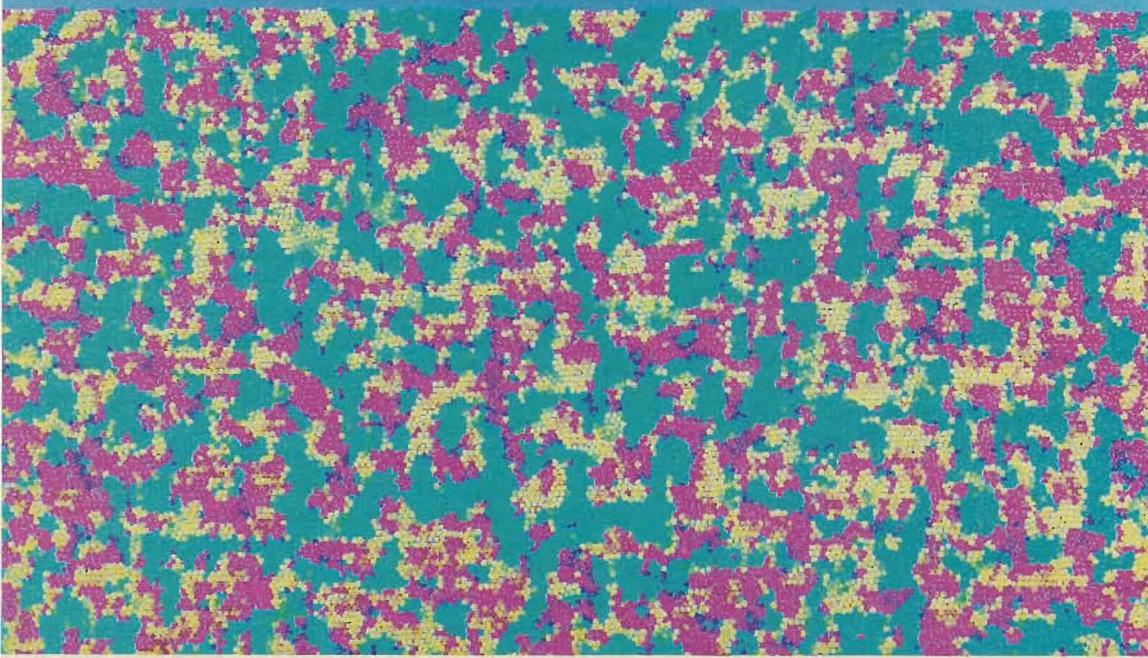
FIG. 5. Frames 1, 6, 25, and 100 from a color videotape showing the annealing of a mechanically unstable square lattice. Frame 100
amorphous configuration for an indentation simulation involving over a million atoms; see Fig. 6. The shading of each atom is determ
orientation of the atom's interacting neighbors. The interframe time interval is 72 time steps, 1.44 $d(m/\epsilon)^{1/2}$, so that the total
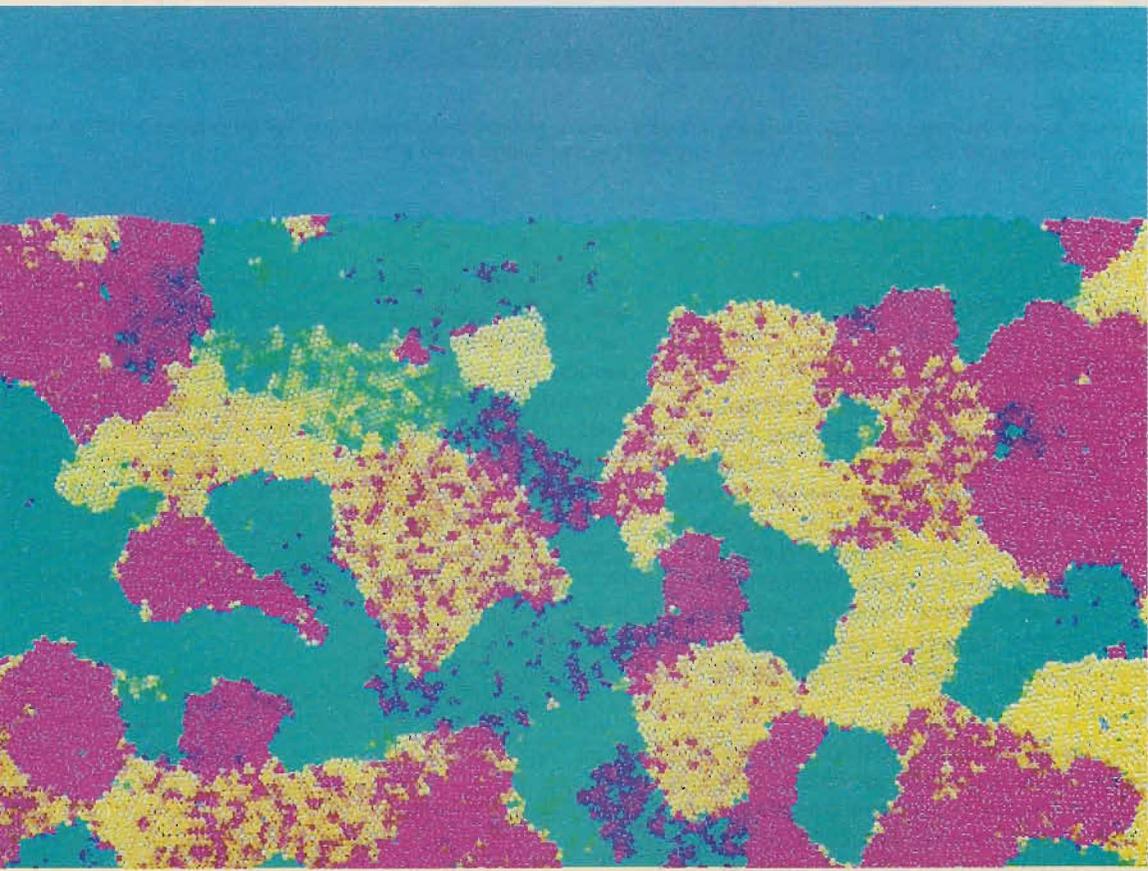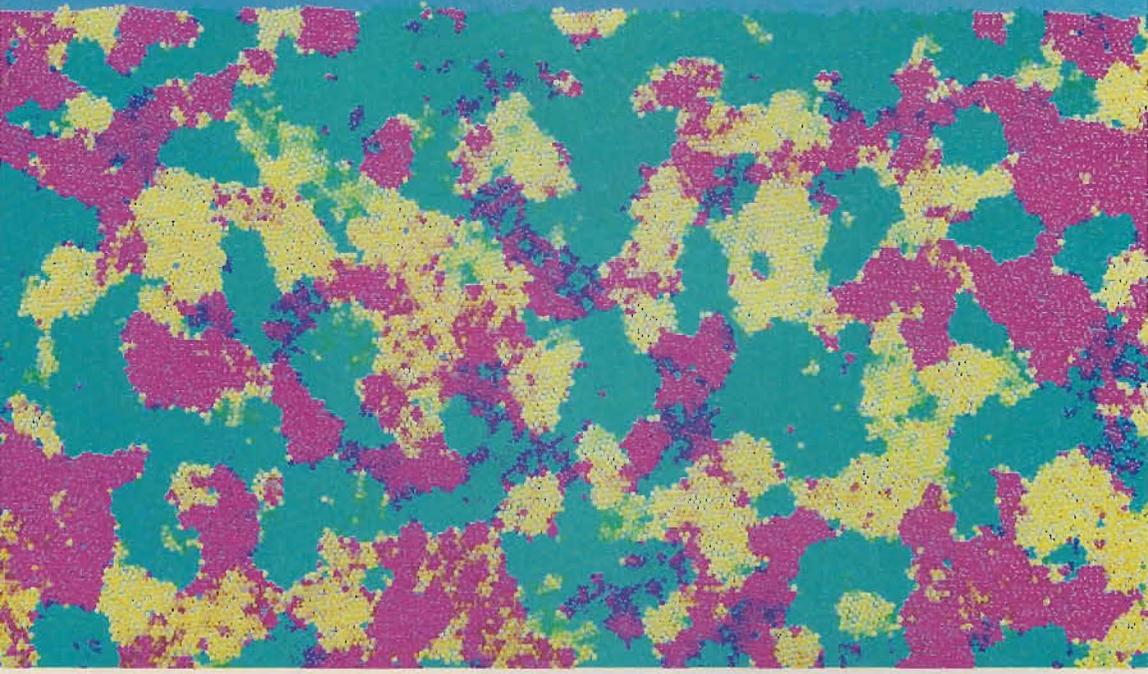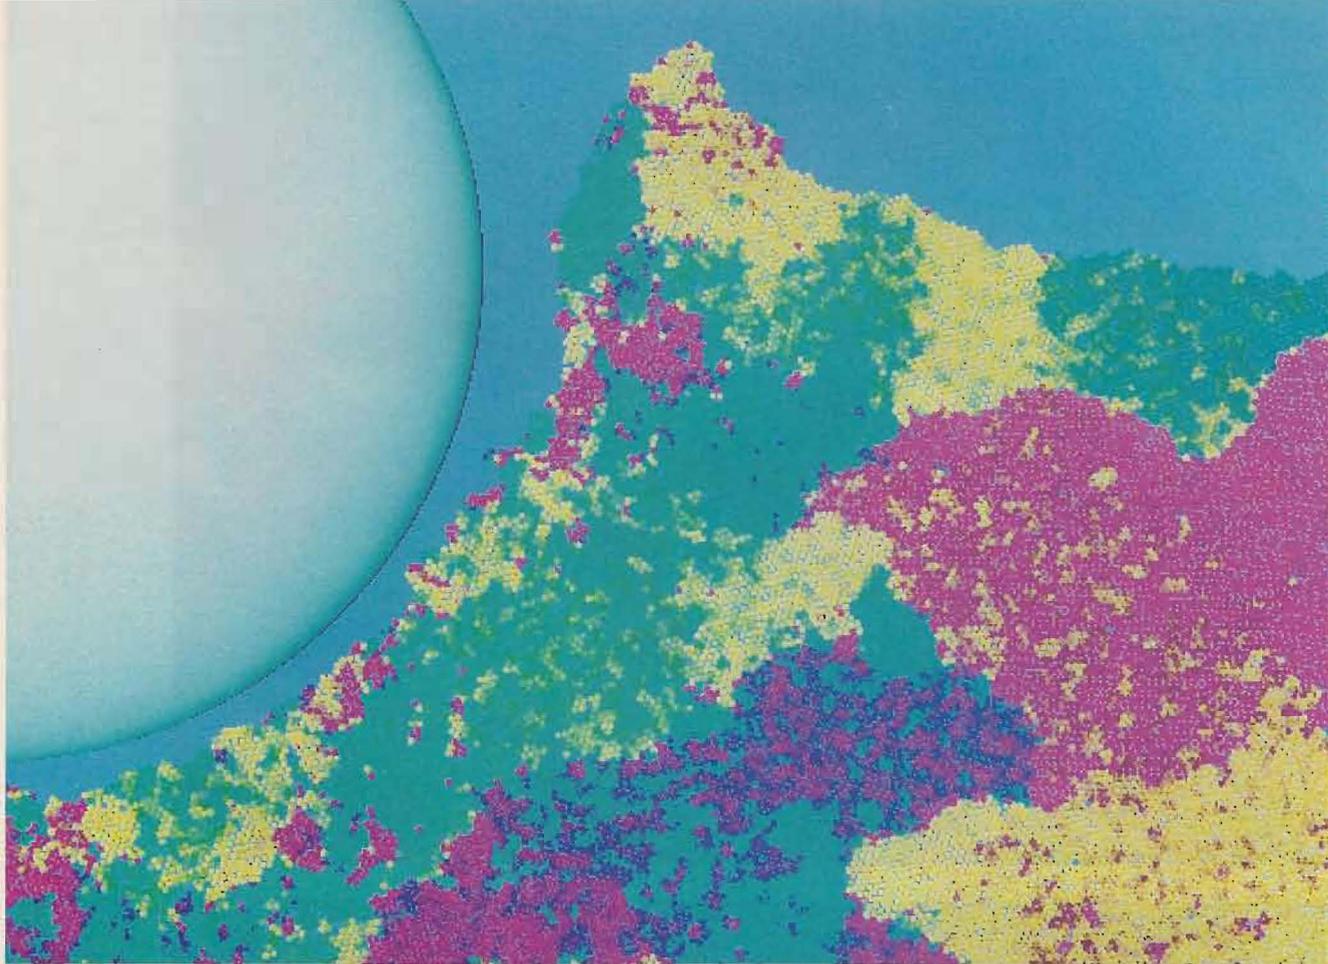
FIG. 5. (*Continued.*)

FIG. 6. Final configuration showing the indentation of an amorphous Lennard–Jones + embedded-atom workpiece. The temperature is half the melting temperature and the maximum indentor velocity, $(\epsilon/m)^{1/2}$, is about one-tenth the longitudinal sound velocity.

$$\frac{d \ln \rho}{dt} = - \nabla \cdot \mathbf{v}, \quad \rho \ddot{\mathbf{x}} \equiv \nabla \cdot \boldsymbol{\sigma} \equiv - \nabla \cdot \mathbf{P},$$

$$\rho \dot{e} = - \nabla \cdot \mathbf{Q} - \mathbf{P} : \nabla \mathbf{v}, \tag{2}$$

a well-posed continuum simulation requires constitutive relations expressing the pressure tensor and heat flux in terms of the field variables and their gradients: $\{\rho, \mathbf{v}, e, \nabla \rho, \nabla \cdot \mathbf{v}, \nabla e\}$. If the relationship between stress and strain were an elastic one, with $\boldsymbol{\sigma} = \boldsymbol{\kappa}(\boldsymbol{\epsilon}) : \boldsymbol{\epsilon}$, then a numerical simulation could be based on the simplest possible continuum representation, triangular zones. Within the ith triangular zone, the displacement vector $\mathbf{u}_i \equiv \mathbf{x} - \mathbf{x}_{eq}$ would then be taken to vary linearly in space, as given by the three vectors $\mathbf{A}_i$, $\mathbf{B}_i$, and $\mathbf{C}_i$ characterizing the zone

$$\mathbf{u} = \mathbf{A} + \mathbf{B}\mathbf{x} + \mathbf{C}\mathbf{y}. \tag{9}$$

In principle, in the absence of singularities, mesh refinement eventually reaches a microscale at which constant stress and strain are reasonable approximations within every zone. Thus this triangular-zone approach to elastic problems eventually converges to an accurate solution.

Plastic deformation is more complicated. The complexity reflects the atomic-scale discontinuities in flow due to the motion of dislocations.[17] When real materials are deformed beyond the "elastic limit" dislocations flow, and removing the load reveals that permanent "plastic" deformation has occurred, the original shape is not recovered. The simplest idealized continuum description of this deformational plastic flow ignores the dislocation flow and incorporates the concept of *plastic yielding*, with a plastic material undergoing shear-induced *inelastic* "plastic" deformation, on a characteristic timescale τ, whenever the shear stress reaches a yield stress $Y \equiv \sigma_{max}$, or, equivalently, whenever the elastic shear strain reaches a limiting value $\epsilon_{max} \equiv Y / \eta$:

$$\sigma_{shear} \equiv \left(\sigma_{xy}^2 + \frac{1}{4} (\sigma_{xx} - \sigma_{yy})^2 \right)^{1/2}$$

$$> Y$$

$$\Rightarrow \frac{d\sigma_{shear}}{dt}$$

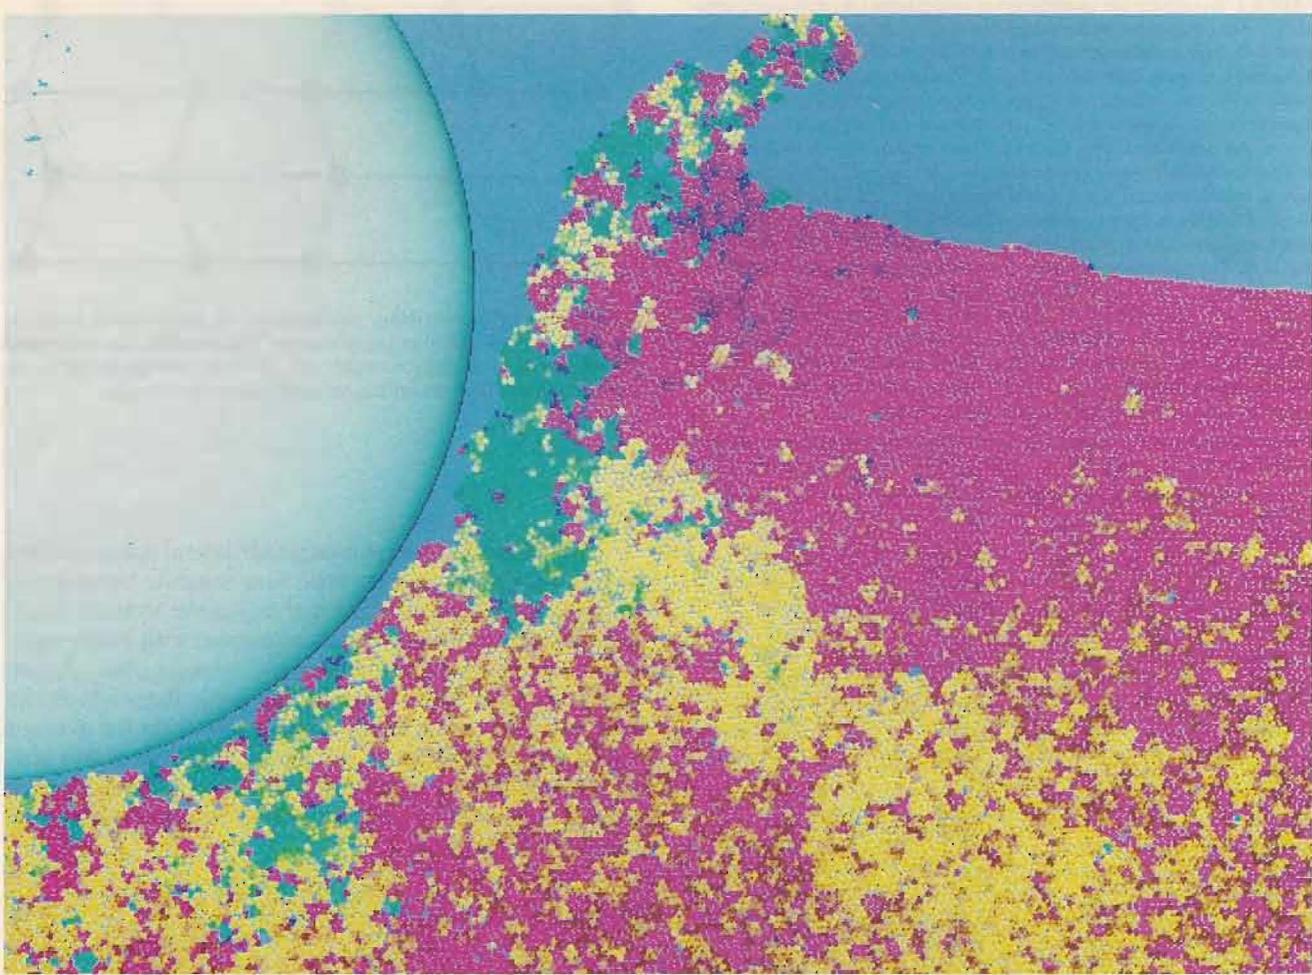$$= \frac{[Y - \sigma_{shear}]}{\tau},$$

FIG. 7. Indentation of a perfect Lennard–Jones + embedded-atom triangular-lattice crystal under boundary conditions identical to those of the amorphous-solid deformation shown in Fig. 6.

$$\sigma_{\text{shear}} \equiv \eta \epsilon_{\text{shear}}, \quad \epsilon_{\text{shear}} \equiv [\epsilon_{xy}^2 + (\epsilon_{xx} - \epsilon_{yy})^2]^{1/2},$$

$$\epsilon_{xx} \equiv \frac{\partial u_x}{\partial x}, \quad \epsilon_{xy} \equiv \frac{\partial u_x}{\partial y} + \frac{\partial u_y}{\partial x}, \quad \epsilon_{yy} \equiv \frac{\partial u_y}{\partial y}. \quad (10)$$

In traditional engineering texts, "yield strength" Y refers to the longitudinal stress/strain ratio at which plastic flow begins in extending a cylindrical bar with vanishing transverse stresses. For the case of equal Lamé constants $\eta = \lambda$, a brief (three-dimensional) calculation shows that the stress at which yield begins is $(5\eta/2)\epsilon_{xx}$, where $\epsilon_{xx} = -4\epsilon_{yy}$ is the corresponding maximum elastic longitudinal strain. Thus, for example, a yield value $\epsilon_{\text{shear}} = 0.02$ with $Y = 0.02\eta$ in our two-dimensional plane-strain calculations corresponds to a three-dimensional yield strain $\epsilon_{xx} = 0.008$ in three dimensions. In either case the "yield strength" would be $Y = 0.02\eta$.

The "invariant" combinations of stress and strain components just given are uniquely independent of the orientation of the x and y axes; that is, arbitrary rotation of the coordinate system leaves the numerical values of $[\sigma_{xy}^2 + (1/4)(\sigma_{xx} - \sigma_{yy})^2]^{1/2}$ and $[\epsilon_{xy}^2 + (\epsilon_{xx} - \epsilon_{yy})^2]^{1/2}$ unchanged. Thus the corresponding [von Mises] flow rule is said to be "isotropic." The physical relaxation time τ is of the order of nanoseconds to picoseconds and corresponds to the time period required for dislocations to move, or first to nucleate and then to move, so as to relieve shear stress. Because most flow processes are driven at subsonic rates well below the dislocation velocity it is reasonable to make the *simplest* choice, choosing the relaxation time τ equal to zero, so that the shear stress cannot exceed the shear strength Y. In order that the stress relaxation be independent of orientation, the plastic relaxation ratios are as follows:

$$[\Delta\sigma_{xx}, \Delta\sigma_{xy}, \Delta\sigma_{yy}]$$
$$\propto [(\sigma_{xx} - \sigma_{yy}), + 2\sigma_{xy}, (\sigma_{yy} - \sigma_{xx})].$$

Continuum simulations of plastic flow require a computational mesh capable of representing incompressible flows. The simple triangular zones that work so well for elastic problems are not sufficiently flexible to represent plastic flow.[18] To see this, imagine a region spanned by a highly refined grid of N points dividing the region into $2N$ contiguous triangles. Idealized constant-volume plastic flow would then require $2N$ constant-volume constraints, one for each triangle. Thus *triangular zones can be ruled out*

System size	R/d	$[\,mv_{max}^2/\epsilon\,]^{1/2}$	$(d^2/\epsilon)(work/R^2)$
3 200A	4	1.0	4.27
3 200C	4	1.0	6.13[n]
3 200C	4	0.1	2.96[n]
3 200C	4	0.01	2.80[n]
3 200C	4	0.001	2.36[n]
12 800A	8	1.0	4.42
12 800C	8	1.0	4.36[n]
12 800C	8	0.1	3.30[n]
12 800C	8	0.01	2.66[n]
12 800C	8	0.004	2.46[n]
12 800C	8	0.001	2.23
12 800C[a]	8	0.001	3.89
12 800C[b]	8	0.001	13.6
12 800CLJ	8	0.001	7.41
51 200A	16	1.0	3.31
51 200C	16	1.0	4.21
204 800A	32	1.0	2.75
204 800C	32	1.0	3.74
1 036 800A	72	1.0	2.49
1 036 800C	72	1.0	3.40[n]
1 036 800C[a]	72	1.0	5.24
1 036 800C[b]	72	1.0	13.3

[a] Calculation carried out at zero temperature.
[b] Calculation carried out at zero temperature using pure Lennard-Jones spline potential.
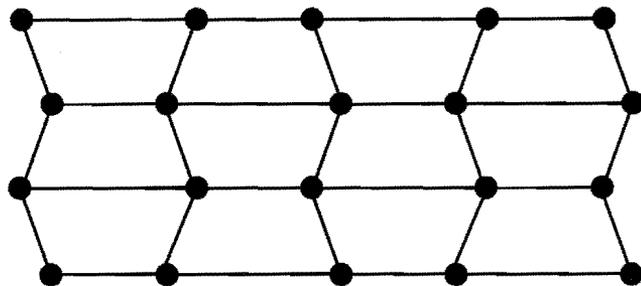

FIG. 8. Hourglass instability characteristic of quadrilateral zones in which stress is calculated as a zone average. This instability can be avoided by computing local stress averages, one for each zone quadrant, or by introducing artificial viscosities; we used local stress averages.

In our work, we have used quadrilateral zones in order to represent plasticity, but with four separate shear-stress tensors, one for each quadrant of each zone, to avoid hourglass instability. Each zone has associated with it four separate elastic shear-strain tensors $\{\epsilon_{xx} - \epsilon_{yy}, \epsilon_{xy} \equiv \epsilon_{yx}\}$, one for each of the four corners of the zone, from which the local shear deviation of the stress tensor from the average equilibrium stress can be calculated. The four strain tensors are calculated as integral averages of finite-difference representations of the following differential expressions for the shear strain rates:

$$\dot{\epsilon}_{12} \equiv \dot{\epsilon}_{xx} - \dot{\epsilon}_{yy} \equiv \frac{\partial v_x}{\partial x} - \frac{\partial v_y}{\partial y}, \quad \dot{\epsilon}_{44} \equiv \frac{\partial v_x}{\partial y} + \frac{\partial v_y}{\partial x}. \quad (11)$$

The integrals, averaged over quadrants of each zone, can be approximated by counterclockwise contour integrals:[19]

$$\langle \dot{\epsilon}_{12} \rangle = \left\langle \frac{\partial v_x}{\partial x} - \frac{\partial v_y}{\partial y} \right\rangle$$

$$\equiv \int \int \left(\frac{\partial v_x}{\partial x} - \frac{\partial v_y}{\partial y} \right) \frac{dx\,dy}{A}$$

$$= \frac{1}{A} \oint (v_x\,dy + v_y\,dx),$$

$$\langle \dot{\epsilon}_{44} \rangle = \left\langle \frac{\partial v_x}{\partial y} + \frac{\partial v_y}{\partial x} \right\rangle$$

$$\equiv \int \int \left(\frac{\partial v_x}{\partial y} + \frac{\partial v_y}{\partial x} \right) \frac{dx\,dy}{A}$$

$$= \frac{1}{A} \oint (v_y\,dy - v_x\,dx), \quad (12)$$

where the area of the quadrant is indicated by A. The velocity field in each zone is assumed to vary isoparametrically with x and y, so that the four strain-rate tensors can be expressed in terms of the velocities of each zone's nodal points. The bulk contribution to the stress tensor is computed using the *total* zone area and the *bulk* modulus.[3] For instance, the energy and pressure for a static lattice governed by a two-dimensional nearest-neighbor Lennard-Jones potential interaction are

because the number of constant-volume constraints and the number of "degrees of freedom" both become equal, to $2N$, as the mesh is refined by increasing N. The physical consequence of this equality is called "locking" and makes plastic flow impossible with simple triangular elements. The same stability of the triangular shape is responsible for the use of triangular strut arrays in truss bridges and towers. More complex elements are required to represent plastic flow. *Quadrilateral* zones [*one* for every mesh point as N is increased] furnish twice as many degrees of freedom as constraints and so provide the simplest model of incompressible flow.

Some treatments of quadrilateral zones can exhibit stability problems. An elementary treatment using quadrilateral zones leads to "hourglass" instability, a form of shear deformation shown in Fig. 8. Because the hourglass deformation, with $\epsilon_{xy} \propto xy$, has a zone-averaged shear strain zero, there are no restoring forces, and a simple zone-averaged quadrilateral simulation is unstable.
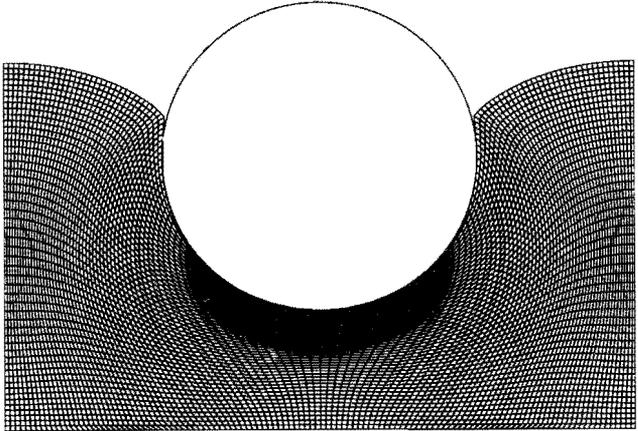
FIG. 9. Finite-difference indentation of an elastic–plastic workpiece using quadrilateral zones and the nearest-neighbor Lennard–Jones equation of state. This 3120-zone simulation has a maximum indentor speed $0.25(\epsilon/\text{m})^{1/2}$, and indentor radius $30d$, where d is the zone width, and $\epsilon_{\text{max}} = 0.04$.

$$\frac{\Phi}{N} \equiv 3\epsilon\left[\left(\frac{\rho}{\rho_0}\right)^6 - 2\left(\frac{\rho}{\rho_0}\right)^3\right], \qquad (13)$$

$$P_{\text{eq}} V \equiv 18\epsilon\left[\left(\frac{\rho}{\rho_0}\right)^6 - \left(\frac{\rho}{\rho_0}\right)^3\right].$$

In this pair-potential case the two Lamé constants are equal and are related to the two-dimensional bulk and shear moduli as follows:

$$\lambda \equiv \eta \equiv G \equiv \frac{B_{2D}}{2} = \frac{\epsilon}{V}\left[63\left(\frac{\rho}{\rho_0}\right)^6 - 36\left(\frac{\rho}{\rho_0}\right)^3\right]. \quad (14)$$

In each zone these contributions to the stress are used to accelerate the nodes. The full calculation proceeds by computing $\{\dot{r}, \ddot{r}, \dot{\epsilon}, \ddot{\epsilon}\}$ from $\{r, \dot{r}, \epsilon, \dot{\epsilon}\}$ by following these steps.

(1) Fix initial conditions, specifying r, $\dot{r}$, ϵ, and $\dot{\epsilon}$, and the boundary conditions. Typically the initial velocities, strains, and strain rates all vanish, except at the indentor-workpiece boundary.

(2) Compute new coordinates r from nodal and boundary velocities $\dot{r}$.

(3) Compute new elastic zone strain tensors ϵ, and stress tensors σ, from the strain-rate tensors $\dot{\epsilon}$.

(4) Calculate the nodal accelerations $\ddot{r}$ from neighboring zones' dilational stresses and each zone quadrant's corner-shear-stress tensor.

(5) Compute zone strain-rate tensors taking plastic strain relaxation into account so that no shear stress exceeds the elastic limit.

(6) Return to step (2) until the simulation is complete.

The differential equations for r, $\dot{r}$, ϵ, and $\dot{\epsilon}$ can be solved by using the classic fourth-order Runge–Kutta integration or by the simpler centered-difference approach. We have used both methods but finally adopted the centered-difference approach because it minimizes storage. For the appearance of a typical simulation, see Fig. 9.

The boundary conditions for the continuum simula-

tion include a "free" boundary, at the workpiece top, vertical "roller" boundaries along the central symmetry line and the vertical boundary, and a "fixed" boundary at the workpiece base. The interaction between the workpiece and the indentor was frictionless, described by an elastic purely repulsive Hooke's-Law pair interaction between the indentor mass and the horizontal surface nodes. The resulting vertical component of the summed indentor forces were then integrated with respect to displacement, so as to measure the work of indentation, given in Table II.

The numerical convergence of this continuum approach seems slower than the atomistic one. Without considerable testing, we cannot be sure that our numerical technique converges to the solution of the continuum Eqs. (2) and (10). In Fig. 10 (top), we show the dependence of the calculated effective yield strength W/R^2 on the mesh size d. Just as in the elastic case, the coarse-mesh data indicate a quadratic mesh dependence. We subsequently discovered that much *finer* meshes, shown in Fig. 10 (bottom) and with as many as 600 000 Lagrangian zones, indicate an asymptotic *logarithmic* dependence on zone size, probably a consequence of a stress singularity at the indentor. Typical data points are included in Table II and

TABLE II. Results for the indentation of specimens using Lagrangian finite-difference approximations to the equation of state. The maximum shear strain is ϵ_{max}. In these continuum simulations, symmetry was used. The number of $d{\times}d$ Lagrangian zones spanning one-half the workpiece area is specified. A typical time step is 0.005 or $0.01\,(\text{m}d^2/\epsilon)^{1/2}$. Calculations use an indentor force constant 100. Because square zones of sidelength d and mass m were used, with unit mass density $\rho(d^2/\text{m})$, the work values in this table have been multiplied by $1.1547 = (4/3)^{1/2}$ and are directly comparable to the atomistic results given in Table I. The distance traveled by the indentor is equal to the indentor radius R. The computer time required for the last entry in the table was approximately 400 h.

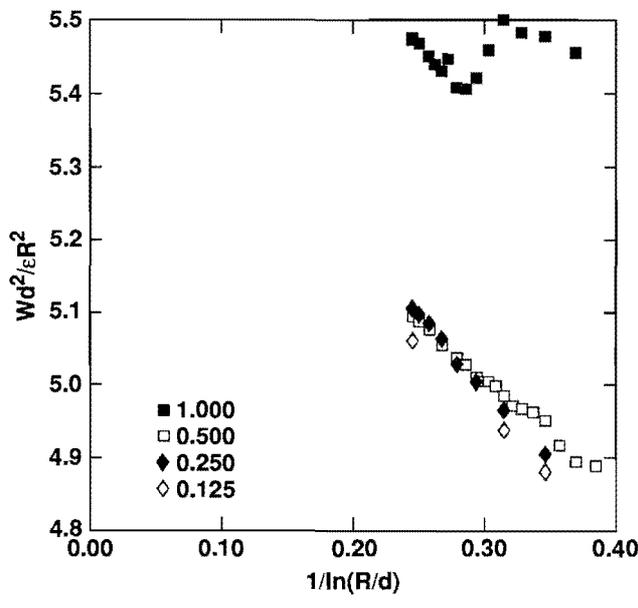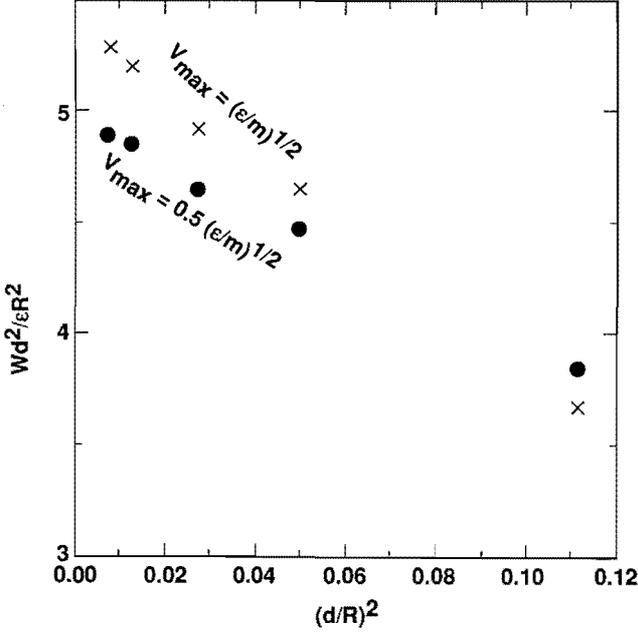| Zones	R/d	ϵ_{max}	$[m v_{\text{max}}^2/\epsilon]^{1/2}$	$(\text{Work}/\epsilon)(d^2/R^2)$
780	3	0.02	1.00	3.67
780	3	0.02	0.50	3.84
780	3	0.04	0.50	8.30
780	3	0.08	1.00	18.40
780	3	0.08	0.50	17.96
1 755	4.5	0.02	1.00	4.63
1 755	4.5	0.02	0.50	4.46
3 120	6	0.02	1.00	4.92
3 120	6	0.02	0.50	4.65
3 120	6	0.04	0.50	9.48
7 020	9	0.02	1.00	5.20
7 020	9	0.02	0.50	4.85
12 480	12	0.02	1.00	5.27
12 480	12	0.02	0.50	4.88
12 480	12	0.04	0.50	9.91
49 920	24	0.02	1.00	5.50
49 920	24	0.02	0.50	4.99
49 920	24	0.02	0.25	4.97
199 680	48	0.02	1.00	5.45
199 680	48	0.02	0.50	5.08
199 680	48	0.02	0.25	5.08
296 595	58.5	0.02	0.125	5.10
312 000	60	0.02	1.00	5.48
312 000	60	0.02	0.50	5.10
449 280	72	0.02	0.50	5.11
611 520	84	0.02	0.50	5.13

FIG. 10. Variation of the work of indentation W as a function of size. The data shown correspond to indentation of a Lennard–Jones–spline material with a yield stress 0.02η and maximum indentor velocities $1.0(\epsilon/m)^{1/2}$, $0.5(\epsilon/m)^{1/2}$, $0.25(\epsilon/m)^{1/2}$, and $0.125(\epsilon/m)^{1/2}$, where ϵ is the Lennard–Jones well depth. The data shown at the top correspond to coarser meshes with half the workpiece described by 26×30, 39×45, 52×60, 78×90, and 104×120 Lagrangian zones of sidelength d. The data shown at the bottom include up to 296 595 mesh points; see table II for selected results.
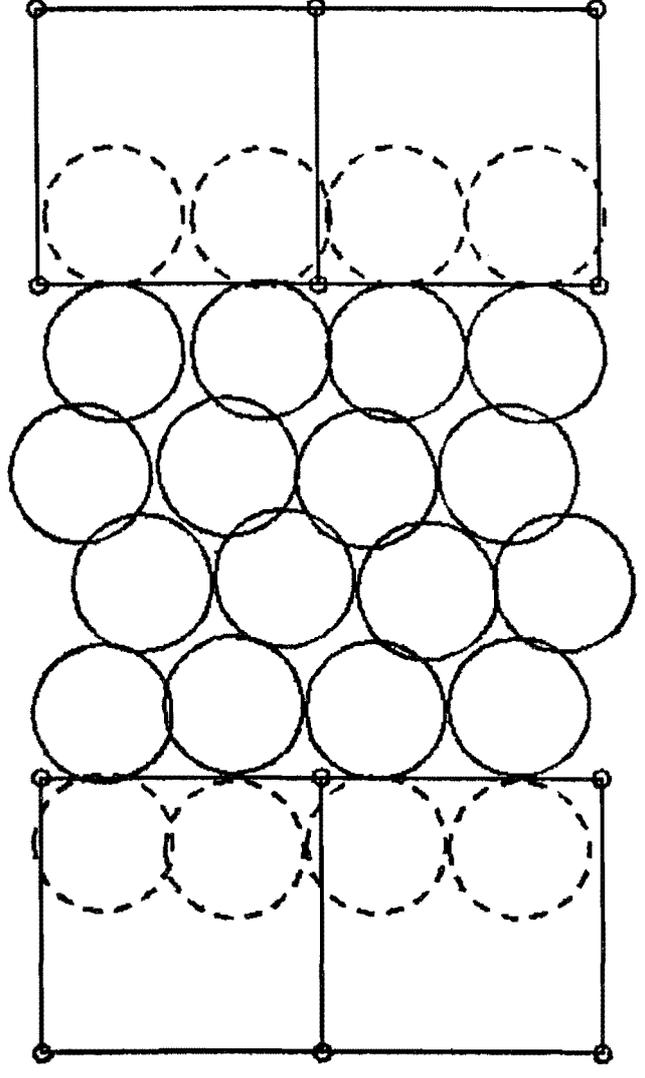

FIG. 11. Snapshot from a hybrid simulation with 16 atoms and four Lagrangian zones. The impulses due to the atoms are applied to the boundaries of the continuum zones. The perpendicular interaction of the atoms is governed by "image atoms," shown dashed.

suggest an overall uncertainty in the extrapolated continuum limit of order 5%. Our results are consistent with a correspondence between the continuum and atomistic simulations with a two-dimensional effective yield strain of order 0.03 to 0.04. The results for the amorphous and crystalline workpieces are not significantly different in this regard.

IV. COMBINED APPROACHES

As an extension of this work, we are exploring hybrid numerical methods combining the atomistic and continuum approaches to large-scale massively parallel simulation. Just as in Lagrangian continuum mechanics, we consider describing a deforming material with an initial checkerboard arrangement of zones. However, some zones are atomistic, with internal motions characterized by an atomistic time scale, and some are continuous, characterized by a sound traversal time. The boundaries between such contiguous regions are characterized as *reflecting* on the atomistic side, with a mass chosen to match the acoustic impedance of the neighboring continuum material. The atomistic impulse transmitted to each boundary between continuum time steps is reflected in the displacements governing the development of the continuum stress tensor; see Fig. 11. Within the atomistic zones, and near the continuum boundary, the mean atomistic velocities *parallel* to that

boundary are controlled with a Nosé–Hoover thermostat, so as to impose a no-slip condition. In the direction *perpendicular* to the continuum boundary, the nearby atoms interact with image atoms in order to provide a realistic surface energy.

V. SUMMARY

Massively parallel computers, typified by SPRINT, make possible routine simulations of deformation problems characterized by *millions* of degrees of freedom. In the case of plastic flow, such simulations indicate relatively rapid convergence of the microscopic systems' flow properties to the macroscopic limit, at least in the absence of stress singularities. The results for crystalline and amorphous solids are roughly similar. Comparison with the continuum simulations implies a rough two-dimensional plastic yield strain about 0.04. Continuum simulations require relatively sophisticated flow models to match the details of the atomistic results, such as the shape of the indentation cavity.

VI. ACKNOWLEDGMENTS

VII. ABOUT THE COVER

We have recently carried out parallel *three*-dimensional simulations of indentation. We modeled amorphous silicon with Stillinger–Weber forces using a 4096-atom amorphous periodic module provided by Frederick O. Wooten. The cover illustrates two of the amorphous 32 768-atom indentations, one with a smooth-faced tetrahedral indentor, the other with a rough face-centered-cubic atomistic indentor. The cover illustrations were created with ray-tracing and image composition software written by Michael J. Allison and Charles W. Grant of the Electrical Engineering Department at the Lawrence Livermore National Laboratory.

REFERENCES

1. The SPRINT computer is described by A. J. De Groot, S. R. Parker, and E. M. Johansson, in *SVD and Signal Processing; Algorithms, Applications and Architectures*, edited by E. F. Deprettere (North-Holland, Amsterdam, 1988).
2. Finite-difference algorithms are discussed by B. L. Holian, A. J. De Groot, W. G. Hoover, and C. G. Hoover, Phys. Rev. A **41**, 4552 (1990).
3. W. G. Hoover, A. J. De Groot, C. G. Hoover, I. Stowers, T. Kawai, B. L. Holian, T. Boku, S. Ihara, and J. Belak, Phys. Rev. A **42**, 5844 (1990).
4. W. G. Hoover, *Computational Statistical Mechanics* (Elsevier, Amsterdam, 1991).
5. W. G. Hoover and W. T. Ashurst, Adv. Theoret. Chem. **1**, 1 (1975); see also W. T. Ashurst, Ph.D. dissertation, University of California at Davis/Livermore (1974).
6. A. Rahman, Phys. Rev. A **136**, 405 (1964).
7. B. L. Holian, W. G. Hoover, B. Moran, and G. K. Straub, Phys. Rev. A **22**, 2798 (1980).
8. F. F. Abraham, W. E. Rudge, D. J. Auerbach, and S. W. Koch, Phys. Rev. Lett. **52**, 445 (1984).
9. W. C. Swope and H. C. Andersen, Phys. Rev. B **41**, 7042 (1990).
10. J. E. Sinclair, in *Computer Simulation in Physical Metallurgy*, edited by G. Jaeucci (Reidel, Boston, 1985).
11. M. M. Waldrop, Science **240**, 286 (1988).
12. H. A. Posch, W. G. Hoover, and B. L. Holian, Ber. Bunsen-Ges. Phys. Chem. **94**, 250 (1990); W. G. Hoover, H. A. Posch, B. L. Holian, M. J. Gillan, M. Mareschal, and C. Massobrio, Mol. Simul. **1**, 79 (1987).
13. B. L. Holian and D. J. Evans, J. Chem. Phys. **78**, 5147 (1983).
14. The embedded-atom concept was originated by three Sandia scientists, Murray Daw, Mike Baskes, and Steve Foiles; see S. M. Foiles, M. I. Baskes, and M. S. Daw, Phys. Rev. B **33**, 7983 (1986).
15. B. L. Holian, A. F. Voter, N. J. Wagner, R. J. Ravelo, S.-P. Chen, W. G. Hoover, C. G. Hoover, J. E. Hammerberg, and T. D. Dontje, Phys. Rev. A **43**, 2655 (1991).
16. D. Scott, W. W. Seifert, and V. C. Westcott, Sci. Am. **230**(5), 88 (1974).
17. For references and related calculations, see W. G. Hoover, A. J. C. Ladd, and B. Moran, Phys. Rev. Lett. **48**, 1818 (1982).
18. J. G. Nagtegaal, D. M. Parks, and J. R. Rice, Comp. Methods Appl. Mech. Eng. **4**, 153 (1974).
19. M. Wilkins, J. Comp. Phys. **5**, 406 (1970).