



# Multiresolution algorithms for massively parallel molecular dynamics simulations of nanostructured materials <sup>☆</sup>

Rajiv K. Kalia <sup>a</sup>, Timothy J. Campbell <sup>a</sup>, Alok Chatterjee <sup>a</sup>, Aiichiro Nakano <sup>a,\*</sup>,  
Priya Vashishta <sup>a</sup>, Shuji Ogata <sup>b</sup>

<sup>a</sup> *Concurrent Computing Laboratory for Materials Simulations, Department of Physics & Astronomy and Department of Computer Science, Louisiana State University, Baton Rouge, LA 70803, USA*

<sup>b</sup> *Department of Applied Sciences, Yamaguchi University, 2557 Tokiwadai, Ube 755-8611, Japan*

## Abstract

Multimillion atom molecular-dynamics (MD) simulations are performed to investigate dynamics of oxidation of aluminum nanoclusters and properties and processes in nanostructured silicon carbide (n-SiC) and nanostructured amorphous silica (n-a-SiO<sub>2</sub>). The simulations are based on reliable interatomic interactions that include both ionic and covalent effects. The simulations are carried out on parallel architectures using highly efficient  $O(N)$  multiresolutions algorithms which include an adaptive load-balancing approach based on wavelets and a data-compression scheme based on fractals.

Results from the oxidation simulation reveal a passivating amorphous oxide layer of thickness  $\sim 40$  Å, which is in excellent agreement with experiments. The oxide layer is amorphous and has mixed tetrahedral,  $Al(O_{1/4})_4$ , and octahedral,  $Al(O_{1/6})_6$ , configurations. The average mass density in the oxide region is 75% of the bulk alumina density. Local stresses in the oxide scale are analyzed and their correlation with the dynamics of oxidation is determined.

Sintering, structural correlations, and mechanical behavior of n-SiC and n-a-SiO<sub>2</sub> are investigated. In the case of n-SiC, both experiment and simulation indicate the onset of sintering around 1500 K which is much lower than the sintering temperature for coarse-grained SiC. In both n-SiC and n-a-SiO<sub>2</sub>, pores are found to be self-similar. They have a fractal dimension close to 2 and their surface roughness exponents are  $\sim 0.5$ . Pair-distribution functions and bond-angle distributions reveal a crystalline core and an amorphous interface in the consolidated n-SiC. In the case of nanophase silica glasses, the short-range order (SRO) is similar to that in the bulk glass but not the intermediate-range order (IRO). In the nanophase system the first sharp diffraction peak (FSDP), the signature of IRO, has a much smaller height and is shifted toward smaller  $k$  relative to the FSDP in the bulk system. The elastic moduli of nanophase silica glasses scale with the density as  $\sim \rho^{3.5}$ ; the bulk, shear and Young's moduli of n-SiC scale as  $\sim \rho^\eta$ , where  $\eta$  is  $3.51 \pm 0.02$ ,  $3.29 \pm 0.06$ , and  $3.34 \pm 0.03$ , respectively. © 2000 Elsevier Science B.V. All rights reserved.

PACS: 81.05.YS

Keywords: Molecular dynamics simulations; Nanophase materials; Silica; Silicon carbide; Oxidation; Parallel computations

## 1. Introduction

Materials synthesized by consolidation of nanometer size atomic clusters are known as nanophase materials. In the mid 80's, a research group at Argonne

<sup>☆</sup> This paper is published as part of a thematic issue on Parallel Computing in Chemical Physics.

\* Corresponding author. E-mail: nakano@bit.csc.lsu.edu.

National Laboratory led by Dr. R. Siegel and a German group led by Dr. H. Gleiter began synthesizing them [1,2]. Since then, a wide variety of nanophase metals, alloys, and ceramics have been synthesized by various experimental methods. Prominent among them are the condensed-phase, gas-phase, and vacuum techniques [2].

Nanophase materials have unique physical properties that result from small grain sizes and the presence of a large number of atoms at interfaces. The structure of nanophase materials has been investigated by:

- (i) transmission electron microscopy, scanning electron microscopy, scanning tunneling microscopy, and field-ion microscopy;
- (ii) electron, X-ray, and neutron diffraction techniques; and
- (iii) Mössbauer, Raman, and positron annihilation spectroscopies.

The properties of nanophase materials are strongly influenced by their grain-boundary volume (which represents a significant fraction of their total volume).

Nanophase metals, intermetallic compounds, and ceramics are significantly better than their conventional coarse-grained counterparts [3–5]. For example, nanophase materials sinter at much lower temperatures than conventional coarse-grained solids. (In the case of nanophase  $\text{TiO}_2$  the sintering temperature is lowered by  $600^\circ\text{C}$  without any compacting or sintering aids.) This leads to considerable savings in the synthesis of such materials. Mechanical properties of consolidated nanophase materials show remarkable improvements in strength and hardness [6]. Cu and Pd assembled from 5 to 7 nm diameter clusters have hardness and yield strength values up to 500% greater than conventional metals [7,8]. This large increase in strength arises from the fact that the creation and motion of dislocations are hindered by small grains of nanophase metals.

In ceramics, which are normally brittle, cluster assembly yields a different benefit. Ceramics can be rendered ductile when synthesized from clusters below 15 nm in diameter. The cause of enhanced ductility is believed to be grain-boundary sliding. Small grains also facilitate efficient impurity doping of these materials. Nanophase insulator and semiconductor materials can be doped at low temperatures allowing efficient introduction of impurity levels and therefore

their electrical, optical, and chemical properties can be tailored.

The chemical properties of nanophase materials are also quite unique. For example, the use of nanoparticles as catalysts to remove sulfur from car exhaust has been tested and found to be about five times more effective than conventional catalysts. This excellent performance is due to the high surface-to-volume ratio of small clusters in nanophase materials. Also the large porosity of these materials increases the adsorption area and makes them better catalysts [2,8–12].

This paper focuses on three kinds of systems:

- (i) nanophase silica consisting of amorphous  $\text{SiO}_2$  nanoclusters;
- (ii) nanophase SiC comprising crystalline nanoclusters; and
- (iii) passivation of aluminum nanoclusters with oxygen.

The last one is the basic entity for a unique kind of nanocomposite whose clusters have metallic interiors and ceramic surfaces. Such a nanocomposite has been synthesized by Sánchez-López et al. [13,14]. They have found that Al nanoclusters of various sizes (120–410 Å) form a 40 Å thick oxide scale regardless of the oxygen dosage. Recently other experimental studies of the passivation of ultrafine Al particles have also been reported. They reveal the formation of a 25 to 50 Å thick oxide scale in 500 to 700 Å size Al particles exposed to air. The oxide scale remains intact even after compaction. Thermogravimetric studies of small metallic particles show an oxide scale of thickness 30 to 40 Å in a 200 Å Al cluster [15]. In connection with ignition, Aumann et al. [16] have observed that Al particles with diameters between 240 and 650 Å oxidize with a square-root time dependence similar to that for flat Al samples. Additionally, they find that the increased surface area of Al particles lowers the oxidizing activation energy relative to that of flat Al samples.

Very little is known about the atomic-scale structure of oxidized nanoclusters. The structure of the oxide scale is believed to be amorphous, although there is no detailed structural analysis of the oxide scale. Recent studies indicate differences between the local structure of the oxide scales formed on Al nanoclusters and that on large Al surfaces. However, details of the differences are not known [14,16–18].

Starting with the pioneering works of Mott [19] and Cabrera and Mott [20], most of the theoretical studies on oxidation have focused on metal surfaces. In metals like magnesium, the oxide scale grows linearly with time. The resulting tensile stresses cause the oxide film to crack and become porous. In metals like Al, Cu, Fe, and Ti where the diffusion of ions or electrons through a non-porous oxide layer controls the oxide growth, a square root time dependence has been observed.

In this paper, we report the results of large-scale, parallel MD simulations of the oxidation of an aluminum nanocluster of radius 100 Å. The simulations take into account the effect of charge transfer between Al and O dynamically on the basis of the electronegativity equalization principle. The resulting Coulomb interaction is calculated with the  $O(N)$  fast multipole method [21,22]. We also use a multiple time-step approach to compute short-range and long-range forces efficiently [23]. The parallel MD implementation is based on highly efficient load-balancing and data-compression algorithms.

In the oxidation simulations, a passivating amorphous oxide layer of thickness  $\sim 40$  Å is formed during 466 ps of simulation time. The average mass density in the oxide region is 75% of the bulk alumina density. Local stresses in the oxide scale are analyzed and their correlation with the dynamics of oxidation is determined.

Large-scale, parallel MD simulations have also been performed to investigate sintering, structure, vibrational properties, and mechanical behavior of nanophase silica glasses consisting of glassy silica nanoclusters that are 70 Å in diameter. The system is sintered at a temperature of 1000 K and at pressures ranging from 0.6 to 2.4 GPa. During the sintering process, the structure of pores remains self-similar with a fractal dimension close to 2 and the surface roughness exponents of pores is  $\sim 0.5$ . The short-range order (SRO) in nanophase glasses is similar to that in the bulk glass. However, the intermediate-range order (IRO), as characterized by the first sharp diffraction peak (FSDP), is quite different in nanophase and bulk silica glasses. In the nanophase system, the FSDP has a much smaller height and is shifted toward smaller  $k$  relative to the FSDP in the bulk system. The low-energy modes in the vibrational density of states are significantly enhanced in the nanophase glass relative to those in the bulk system. The elastic moduli of

nanophase silica glasses scale with the density as  $M \sim \rho^{3.5}$ .

In this paper, we also present the results of the first joint MD and neutron scattering study of nanophase SiC (n-SiC). Both experiment and simulation indicate that sintering ensues above 1500 K which is much lower than the sintering temperature for coarse-grained SiC. In the MD simulations, the sintered n-SiC systems at different densities reveal self-similar pores whose fractal dimensions are  $\sim 2.4$  and surface roughness exponents  $\sim 0.45$ . The mean-squared displacements of both Si and C atoms in the nanoclusters indicate that the primary sintering mechanism is surface diffusion of atoms. Pair-distribution functions and bond-angle distributions in the consolidated n-SiC reveal that cores of nanoparticles are crystalline and the interfacial regions are amorphous. The elastic moduli of n-SiC increase with the density as  $\sim \rho^\eta$ , where  $\eta$  is  $3.51 \pm 0.02$ ,  $3.29 \pm 0.06$ , and  $3.34 \pm 0.03$  for the bulk, shear and Young's moduli, respectively.

## 2. Molecular-dynamics simulations

Molecular-dynamics (MD) simulations for complex physical systems can be computationally very demanding: they may require large numbers of atoms, long simulation times, and computationally intensive interatomic interactions. For example, MD simulations of cluster-assembled nanophase materials require  $10^6$  to  $10^7$  atoms, since each cluster consists of  $10^3$  to  $10^5$  atoms. Preparation of well-thermalized nanophase systems and calculation of mechanical and thermal properties can require  $10^5$  to  $10^7$  time steps.

Molecular-dynamics simulations described in this paper are based on complex interatomic potentials. For nanophase silica and silicon carbide, these potentials combine two-body and three-body terms. The two-body terms include charge-transfer between atoms through screened Coulomb potentials, charge-dipole interactions, and steric repulsion between atoms. (In SiC, there is an additional dipole-dipole interaction term.) Covalent effects in silica and silicon carbide are taken into account through three-body bond-bending and bond-stretching terms.

To validate these potentials, MD simulations are performed to determine various structural and dynamical properties and the results are compared with exper-

imental measurements. In the case of silica, we have investigated the behavior of molten, crystalline and amorphous states and also of permanently densified amorphous SiO<sub>2</sub>. The MD results for pair-distribution functions, static structure factors, vibrational densities of states, bond-angle distributions, and elastic moduli in crystalline and amorphous SiO<sub>2</sub> are in good agreement with experimental results [24–26]. The interatomic potential for silicon carbide is validated by comparing the MD results with experimental measurements of lattice constant, cohesive energy, elastic moduli (C<sub>11</sub>, C<sub>12</sub>, C<sub>44</sub>), bulk modulus, and the phonon density-of-states of crystalline β-SiC.

Molecular-dynamics simulations of oxidation of aluminum nanoclusters are based on an interaction model [27] that can successfully describe a wide range of physical properties of Al and Al<sub>2</sub>O<sub>3</sub>. This so called electrostatic plus (ES+) model is capable of treating: (i) both metallic and ceramic systems; (ii) bond formation and bond breakage; and (iii) changes in charge transfer as the atoms move and their local environments are constantly altered. In the ES+ model, the potential energy of the system is expressed as the sum of an embedded-atom potential and an electrostatic potential (ES):  $\mathcal{V} = \mathcal{V}_{\text{EAM}} + \mathcal{V}_{\text{ES}}$ . The latter is expressed as,

$$\begin{aligned} \mathcal{V}_{\text{ES}} &= \sum_i v_i(q_i) + \frac{1}{2} \sum_{i \neq j} v_{ij}(\mathbf{r}_{ij}; q_i, q_j), \\ v_i(q_i) &= v_i(0) + \chi_i^0 q_i + \frac{1}{2} J_i^0 q_i^2, \\ v_{ij} &= (\mathbf{r}_{ij}; q_i, q_j) \\ &= \int d^3 r_1 \int d^3 r_2 \rho_i(\mathbf{r}_1; q_i) \rho_j(\mathbf{r}_2; q_j) / r_{12}, \end{aligned} \quad (1)$$

where the local atomic energy  $v_i(q_i)$  is a second-order Taylor expansion in the partial charge  $q_i$ ,  $\chi_i$  is the electronegativity [28], and  $J_i$  is associated with a self-Coulomb repulsion [29]. In Eq. (1),  $v_{ij}(\mathbf{r}_{ij}; q_i, q_j)$  represents the electrostatic interaction energy due to charge-density distributions  $\rho_i(\mathbf{r}; q_i)$  and  $\rho_j(\mathbf{r}; q_j)$  around atom  $i$  and  $j$  (including the nuclear point charge), respectively. It is possible to choose a realistic charge-density distribution function that would allow for directional dependence of bonds and local polarizability. For mathematical convenience, Streitz and Mintmire [27] chose  $\rho_i$  to be an extended form of a Slater-type 1s orbital:

$$\begin{aligned} \rho_i(\mathbf{r}; q_i) &= \mathcal{Z}_i \delta(\mathbf{r} - \mathbf{r}_i) + (q_i + \mathcal{Z}_i) \\ &\times \left( \frac{\zeta_i^3}{\pi} \right) \exp(-2\zeta_i |\mathbf{r} - \mathbf{r}_i|), \end{aligned} \quad (2)$$

where  $\zeta_i$  is the decay length for atomic orbitals and  $\mathcal{Z}_i$  is an effective core charge ( $0 < \mathcal{Z}_i < Z_i$ , with  $Z_i$  the total nuclear charge of an atom).

The parameters in the ES+ potential parameters are fitted to the bulk properties of both fcc aluminum and α-alumina crystal structures. The potential yields reasonable surface energies and relaxations for several low-index surfaces of α-alumina.

The unique feature of the ES+ model potential is that it explicitly includes dynamic charge transfer between anions and cations. Fig. 1 illustrates dynamic charge transfer in a MD simulation of oxygen on an aluminum surface. The method is based on a semi-empirical approach in which atomic charges are determined according to the electronegativity equalization condition [29,30]. In other words, as atomic positions change dynamically charge is redistributed in such a way as to cause the chemical potential or electronegativity to be equal at every atomic site.

In MD simulations, the atomic charges,  $q_i$ , are determined at each time step by minimizing the electrostatic energy Eq. (1), subject to the charge-neutrality constraint,  $\sum_i q_i = 0$ . This constrained minimization is algebraically equivalent to the electronegativity equalization condition that the chemical potentials  $\partial \mathcal{V}_{\text{ES}} / \partial q_i$  be equal for all atoms. In practice the minimization involves concurrent solution of two sets of unconstrained linear equations using a conjugate gradient method [31]. A preconditioning scheme based on splitting the Coulomb-interaction matrix into short-range and long-range components has recently been proposed by Nakano [32]. This scheme is found to substantially improve the convergence and parallel efficiency by increasing data locality.

An alternative approach to the above variable charge method is to treat the atomic charges as dynamic variables in an extended Lagrangian [33]. Charges are given fictitious masses and velocities and then propagated with the atomic degrees of freedom in Newton's equations of motion. Typically, the timestep for the charge degree of freedom is much smaller than the time step for the atomic motion. Multiple time step integration schemes can be applied to speed up the computations [23].




Fig. 1. Snapshots of a variable-charge molecular-dynamics simulation of oxygen (large spheres) on an aluminum (small spheres) surface. The charge on an atom (represented by color) is allowed to vary with the environment.

The most compute-intensive part of the ES+ potential is the long-range Coulomb interaction. It is calculated with the Fast Multipole Method (FMM) developed by Greengard and Rokhlin [21,22]. In this hierarchical scheme, the MD cell is recursively divided in half along each of the Cartesian axes to obtain multiple sets of cells. In three-dimensions the entire structure can be viewed as an oct-tree structure with the MD cell as the root of the tree, corresponding to level  $l = 0$ . At each level in the tree there are  $8^l$  number of cells (each parent in level  $l - 1$  produces 8 children). The recursive decomposition continues to a level  $L$  at which point further refinements would make the cell size smaller than the cutoff distance for the short-range potential. In other words, at the highest level of refinement each cell corresponds to a link-cell in the link-cell list structure. Because the hierarchical decomposition in FMM is related to a tree structure, the cells at the highest level of refinement are often referred to as leaf-cells.

The FMM algorithm decomposes into an *upward pass* and a *downward pass*. The *upward pass* starts at the highest level of refinement  $L$  where the outer multipoles  $C_n^m$  for each leaf-cell (with respect to

the center of the leaf-cell) are computed from the equation:

$$C_n^m = \sum_{i=1}^N q_i r_i^n Y_n^{-m}(\theta_i, \phi_i). \quad (3)$$

Since the information about each atom is used only once, the computational cost is  $O(Np^2)$ , where  $p$  is the order of the multipole expansion. The outer multipoles  $D_l^j$  for each cell  $c$  in level  $L - 1$  are then computed from the outer multipoles  $C_n^m$  of  $c$ 's children in level  $L$  by using a transformation that translates each child's outer multipoles to the parent's center and then they are added together. This procedure is repeated for each of the successive levels in the tree until level 2 is reached. Since the maximum number of possible leaf-cells is  $N$  and each translation involves  $p^4$  operations the computational cost is  $O(Np^4)$ . At the end of the *upward pass* the outer multipoles for all cells at all levels are determined.

Before describing the *downward pass* it is necessary to define some of the terminology used in the FMM [21,22]. At a level  $l$  the *nearest neighbors* of a cell  $c$  are defined to be the set of 26 cells that share a boundary point with  $c$ . Two cells are said to *well-separated* if they are separated by at least one cell.

The *interaction set* for a cell  $c$  is defined to be the set of cells at the same level as  $c$  that are not *nearest neighbors* of  $c$  and whose parent cells are *nearest neighbors* of the parent of  $c$ . Note that at any level there are at most 189 entries in the *interaction list* of a cell.

The *downward pass* operates by consistently computing interactions between cells at the coarsest level possible. This is accomplished for a given cell by computing interactions with those cells which are *well-separated* and whose interactions have not been accounted for at the parent's level. The *downward pass* is initiated at the coarsest level by computing all the inner multipoles for each cell in the level. When free boundary conditions are used, the coarsest level that contains *well-separated* cells is  $l = 2$ . The inner multipoles for each cell  $c$  in the coarsest level are computed by converting the outer expansion of each cell that is *well-separated* from  $c$  to an inner expansion about the center of  $c$  using the outer-to-inner translation and then adding them together:

$$E_l^j = \sum_{n=0}^{\infty} \sum_{m=-n}^n O_{n+l}^{-m-j}(\mathbf{x}_2 - \mathbf{x}_1) D_n^m, \quad (4)$$

where

$$O_n^m(\mathbf{x}) = O_n^m(r, \theta, \phi) = \frac{(-1)^n i^{|m|} Y_n^M(\theta, \phi)}{A_n^m r^{n+1}}, \quad (5)$$

$$D_l^j = \sum_{n=0}^l \sum_{m=-n}^n I_{l-n}^{j-m}(\mathbf{x}_1 - \mathbf{x}_0) C_n^m, \quad (6)$$

$$I_n^m(\mathbf{x}) = I_n^m(r, \theta, \phi) = i^{-|m|} A_n^m r^n Y_n^m(\theta, \phi), \quad (7)$$

and

$$A_n^m = A_n^{-m} = \frac{(-1)^n}{\sqrt{(n-m)!(n+m)!}}, \quad (8)$$

In Eq. (5),  $Y_n^m$  are the spherical harmonics. The *downward pass* proceeds in a recursive manner, beginning at the next coarsest level, as follows:

- (i) for each cell  $c$  in level  $l$  the inner expansion belonging to  $c$ 's parent is converted to an inner expansion about  $c$ 's center using the inner-to-inner translation (4);
- (ii) for each cell  $c$  in level  $l$  the outer expansion of each cell in  $c$ 's interaction list is converted to an

inner expansion about  $c$ 's center using the outer-to-inner translation (4) and then they are added together;

- (iii) steps (i) and (ii) are repeated for level  $l + 1$  until they are completed for all levels including the leaf-level  $L$ .

Once steps (i) and (ii) have been completed at the leaf-level, the inner expansion for each leaf-cell will contain the interactions with all other *well-separated* leaf-cells. The number of operations required for steps (i) and (ii) is  $\leq 190Np^4$ . Finally, the far field contribution to the potential at each atomic position is computed from the inner expansion:

$$\mathcal{V}(\mathbf{x}) = \sum_{l=0}^{\infty} \sum_{j=-l}^l F_l^j I_l^j(\mathbf{x} - \mathbf{x}_3), \quad (9)$$

where

$$F_l^j = \sum_{n=l}^{\infty} \sum_{m=-n}^n I_{n-l}^{m-j}(\mathbf{x}_3 - \mathbf{x}_2) E_n^m. \quad (10)$$

This step requires  $\leq Np^2$  operations. Altogether, the computation for the *downward pass* scales as  $Np^4$ .

The *nearest neighbor* leaf-cell contributions are computed directly using the link-cell lists with the number of operations proportional to  $N/N_b$ , where  $N_b = N/8^L$  the average number of atoms per link-cell. From the discussion of the *upward pass*, *downward pass*, and direct calculation it can be seen that the computational complexity of the FMM algorithm is  $O(N)$ . For realistic system sizes serial implementation can still be impractical, thus motivating the need for parallel computation.

### 3. Parallel implementation

To implement the MD algorithm on parallel computers, we use a divide-and-conquer strategy based on domain decomposition [34]. We consider a system of  $N$  atoms contained in an MD cell. The partitioning of work among  $P$  processors (nodes) is accomplished as shown in Fig. 2. The MD cell is subdivided into  $P = P_x \times P_y \times P_z$  subsystems of equal volume. All the attributes (coordinates, velocities, accelerations, species, etc.) associated with atoms located within a particular subsystem are stored within the memory of the corresponding node. When atoms




Fig. 2. Domain-decomposition scheme in two dimensions for nine nodes. Arrows indicate the direction of message-passing. The dashed line indicates the copied boundary regions for node 5.

move out of a subsystem into a neighboring subsystem, the corresponding attributes (positions, velocities, etc.) are transferred using standard message-passing library routines [35].

The calculation of short-range interatomic forces in a subsystem is done using the link-cell list scheme. To calculate the forces on atoms within a subsystem, the coordinates of all the boundary atoms that reside in the 26 nearest-neighbor subsystems are required. The coordinates of boundary atoms are exchanged between nearest-neighbor nodes using *non-blocking* send and *blocking* receive operations. Newly received boundary atom coordinates are stored in augmented link-cell lists. Boundary-atom exchange between nearest-neighbor nodes is accomplished through six message-passing steps: north, south, east, west, upper, and lower. Boundary-atom coordinates that need to be exchanged with edge-sharing and corner-sharing neighbors are first passed to the appropriate face sharing neighbor and then forwarded to the appropriate destination node during the subsequent message-passing steps. Newton's third law can be used to reduce the number of message-passing steps by a factor of two. In this case, it is necessary to send the forces computed for boundary atoms back to the source node.

In the parallel FMM implementation, the unit MD cell is decomposed into  $P = P_x \times P_y \times P_z$  subsystems with the requirement that each  $P_\alpha$  be a power of 2.

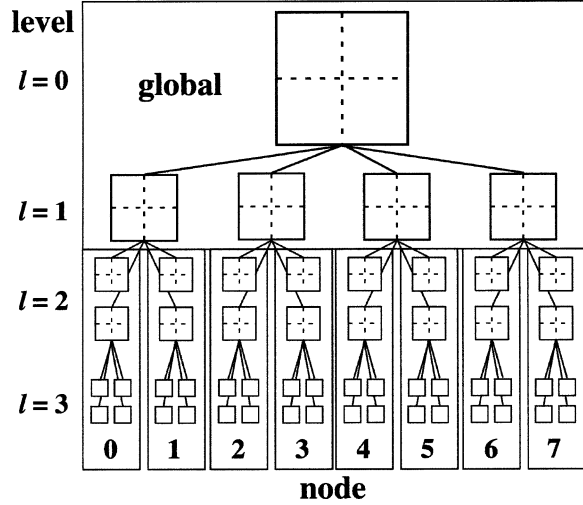


Fig. 3. Domain decomposition scheme for FMM in a two-dimensional system. In the lower levels, cells are local to a node. Cell information in the upper levels is made global to all nodes.

We define a level  $L_g = \log_2[\max(P_x, P_y, P_z)]$  in the FMM tree. In the lower levels of the hierarchy,  $l \geq L_g$ , each cell is uniquely assigned to a node. For upper levels,  $l < L_g$ , the number of cells becomes smaller than the number of nodes. In this case, assignment of each cell to a node results in processors becoming idle. An alternative approach is to duplicate the multipole computations in the upper levels by setting the cells to be global. Fig. 3 illustrates this decomposition scheme for a two-dimensional system.

The *upward pass* and *downward pass* proceed as in the serial case. The calculation of multipole expansions for the leaf-cells is local to each node. Because a parent cell and its children reside on the same node, the outer-to-outer translations of the upward pass and the inner-to-inner translations of the downward pass do not require any communication between nodes. To compute the outer-to-inner translations from a cell's interaction set, the outer multipoles of two boundary layer cells must be copied from the nearest-neighbor nodes. This is accomplished through message-passing steps similar to those used in the short-range force calculations. The  $8^l/P$  cells on a node are augmented with outer multipoles copied from adjacent nodes to form an array consisting of the outer multipoles of  $(2^l/P_x + 2)(2^l/P_y + 2)(2^l/P_z + 2)$  cells at each layer.

The outer-to-inner translations at each of the lower levels ( $l \geq L_g$ ) are computed using the augmented set of outer multipoles. For upper levels ( $l < L_g$ ) the global set of outer multipoles is used.

Uniform spatial decomposition of a system containing an inhomogeneous distribution of atoms can result in unequal workloads among nodes. This degrades the efficiency of the parallel implementation based on conventional domain decomposition. Global structural rearrangement during a simulation can also lead to load-imbalance. Dynamic load-balancing schemes that adaptively repartition the workload are formulated to address these issues.

We have designed an adaptive, low-overhead load-balancing scheme that minimizes the communication and incorporates appropriate boundary conditions [36]. The main idea is to introduce an adaptive curvilinear-coordinate system,  $\xi$ , which is related to the dimensionless atomic coordinate,  $\mathbf{s}$ , by a mapping,  $\xi = \mathbf{s} + \mathbf{u}(\mathbf{s})$ . We have implemented both plane-wave and wavelet representations for the transformation:

$$\mathbf{u}(\mathbf{s}) = \sum_q \mathbf{a}_q \exp(i\mathbf{q} \cdot \mathbf{s}), \quad (11)$$

$$\mathbf{u}(\mathbf{s}) = \sum_{n,m} \mathbf{d}_{n,m} \Psi_{n,m}(\mathbf{s}). \quad (12)$$

We apply the uniform spatial decomposition in the curvilinear space. The coordinate transformation is then dynamically adapted to minimize the load-imbalance and communication costs using the simulated-annealing technique. Such a uniform decomposition in the curvilinear space generally results in curved partition boundaries in the Euclidean space. We find that a multiresolution analysis based on wavelets leads to a compact representation of  $\mathbf{u}(\mathbf{s})$ , and accordingly to fast convergence of the minimization procedure [37].

The load-balancer has been implemented on parallel computers and incorporated into the existing MD codes. In benchmark tests on 32 nodes, the load-balancer reduced the time to compute forces by a factor of 4.2, while the overhead due to the load-balancer was only 3.2% of the total execution time of the MD program (repartitioning with 5 simulated annealing iterations after 60 MD steps) [36]. We also

find that the load-balancer is highly scalable: for a  $1.04 \times 10^9$  atom MD simulation on 64 IBM SP2 nodes, the parallel efficiency is 92% [37].

#### 4. Data compression

A serious technological gap exists between the growth in processor power and that of input/output (I/O) speed. The I/O (including data transfer to remote archival storage devices) has thus become the bottleneck in large-scale MD simulations.

We have developed a scalable data-compression scheme to address the I/O problem. It uses octree indexing and sorts atoms accordingly on the resulting spacefilling curve [38] (see Fig. 4). Space-filling curve is a bijective mapping of a 1-dimensional array to 3-dimensional grid points, which preserves the spatial proximity of consecutive array elements in the 3-dimensional space. By storing differences between successive atomic coordinates, the I/O requirement with the same error tolerance level reduces from  $O(N \log N)$  to  $O(N)$ . This, together with




Fig. 4. A spacefilling curve based on octree indexing maps the 3-dimensional space into a sequential list, while preserving spatial proximity of consecutive list elements. (The panel shows a 2D example.)






Fig. 5. Snapshots showing the evolution of oxidation in a wedge cut from the Al nanocluster at 25 and 50 ps of simulation time. The larger spheres correspond to oxygen and smaller spheres to aluminum; color represents the sign and magnitude of the charge on an atom.

a variable-length encoding to handle exceptional values (outliers), reduces the I/O size by an order-of-magnitude (52 to 6 Bytes/atom) with a user-controlled error bound. The compression algorithm is also lightweight, and involves minimal computation. Since the source code is written in the standard C language and is portable to different architectures, it is readily integrated into an existing MD code as a plug-in.

The compression algorithm for MD configurations (atomic positions, velocities, and other attributes) consists of the following major steps:

- (1) Convert all double-precision data to integers by dividing them by a user-specified error bound;
- (2) Compute the octree index,  $R_i$  of atomic positions for all the atoms,  $i = 1, \dots, N$  ( $R_i$  is obtained by interleaving the binary representation of  $x$ ,  $y$ , and  $z$  coordinates);
- (3) Sort the atoms in the increasing order of  $R_i$ ;
- (4) Differentiate the atomic-position data:  $\Delta R_i = R_i - R_{i-1}$ ;
- (5) Store differentiated  $\Delta R_i$ , velocities, and similar data using variable-length encoding.

## 5. Oxidation of Al nanocluster

The initial oxide growth is characterized by large stress gradients and rapid atomic diffusion. Fig. 5 shows the evolution of oxidation in a wedge cut from the system at 25 and 50 ps of simulation time. The thickness of the oxide increases linearly with time to 25 Å during the first 50 ps and subsequently it begins to saturate. The initial oxide growth is accompanied by a rapid increase in temperature ( $\sim 1500$  K) in the surface region (due to the energy released during Al–O bond formation). The growth of the oxide is due to the movement of oxygen towards the interior of the cluster and movement of aluminum towards the oxide surface. Movement of atoms through the oxide is consistent with the distribution of the local stresses. Fig. 6 shows the hydrostatic stress in a slice through the middle of the nanocluster at 50 ps. It can be seen that while large stress gradients exist throughout the nanocluster, the hydrostatic stresses in the oxide region are predominately negative causing movement of atoms into the oxide region. The hydrostatic stresses in the core region are predominately positive. The large stress gradients cause rapid movement of atoms throughout the

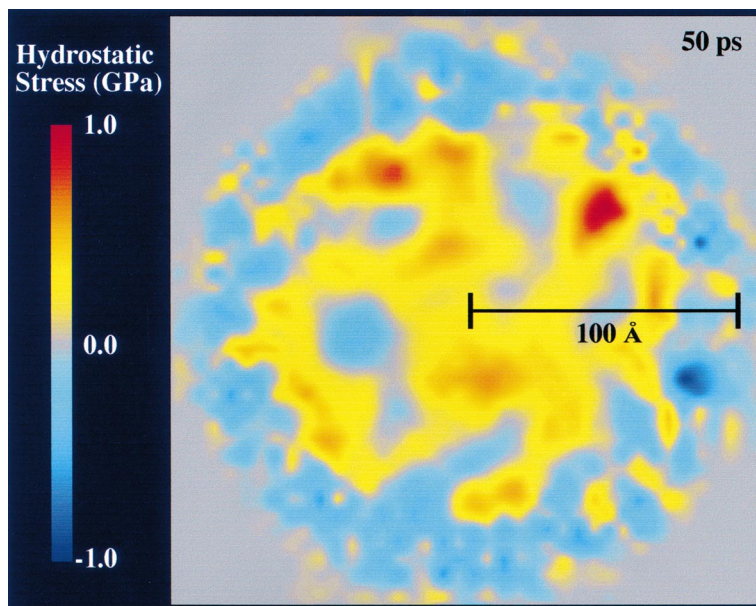


Fig. 6. Hydrostatic stress in a slice through the center of the oxidized Al nanocluster after 50 ps of simulation time. The stresses were calculated by averaging the atomic virial in  $6 \text{ \AA}$  voxels over a 1 ps interval.




Fig. 7. Snapshot of the oxidized Al nanocluster after 466 ps of simulation time. A  $90^\circ$  wedge is cut from the cluster to reveal the interior structure. The larger spheres correspond to oxygen and smaller spheres to aluminum; color represents the sign and magnitude of the charge on an atom.

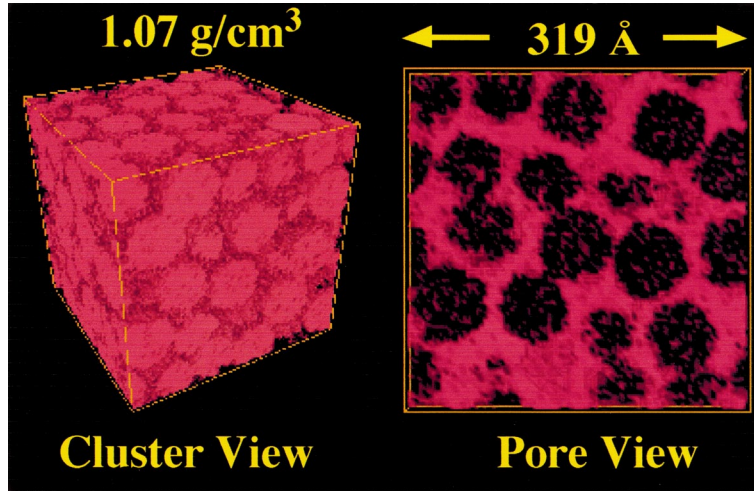


Fig. 8. Snapshot of the nanophase amorphous  $\text{SiO}_2$  system at a density of  $1.07 \text{ g/cm}^3$ . The figure on the left shows the nanoclusters in red. On the right is a  $20 \text{ \AA}$  slice showing the pore regions in red.

oxide. For example, around 50 ps diffusivities of Al and O are  $1.4 \times 10^{-4}$  and  $1.1 \times 10^{-4} \text{ cm}^2/\text{s}$ , respectively. Radial and tangential diffusivities remain about equal resulting in uniform growth of the oxide.

The oxidation simulation was continued to 466 ps. During this time a  $40 \text{ \AA}$  thick oxide scale with an average mass density of  $2.9 \text{ g/cm}^3$  (about 75% of the crystalline  $\alpha\text{-Al}_2\text{O}_3$  density) is formed, as shown in Fig. 7. This is in agreement with oxidation experiments in which a 20 to  $50 \text{ \AA}$  thick oxide scale forms on aluminum nanoclusters 100 to  $700 \text{ \AA}$  in size. Analysis of the oxide region shows how structural correlations vary as we pass through the oxide from the metal-oxide interface to the oxide-environment interface. In the metal-oxide region the Al–O bond length is  $1.81 \text{ \AA}$ . It increases slightly in the interior and oxide-environment regions where the oxygen density is higher. The coordination numbers for Al are 3.1 in the metal-oxide interface region, 3.9 in the interior of the oxide scale, and 4.3 in the oxide-environment interface. Analysis of the O–Al–O bond-angle distribution for each region in the oxide shows that there are two distinct peaks at  $90^\circ$  and  $109^\circ$ , indicating mixed octahedral,  $\text{Al}(\text{O}_{1/6})_6$ , and tetrahedral,  $\text{Al}(\text{O}_{1/4})_4$ , configurations. These structural results are consistent with experimental results for amorphous

alumina which are known to consist of tetrahedrally and octahedrally coordinated aluminum [39,40].

## 6. Nanophase silica and silicon carbide

We have performed large-scale parallel MD simulations to investigate sintering, structure, and mechanical properties of nanophase silicon carbide (3 nm clusters) and nanophase silica (7 nm clusters), see Fig. 8. In addition, a neutron scattering experiment has been performed at the Intense Pulsed Neutron Source Division (IPNS), Argonne National Laboratory to investigate the sintering behavior of nanocrystalline silicon carbide.

Experimental and MD results for the sintering of nanophase silicon carbide (n-SiC) are compared in Fig. 9 which shows the experimentally observed particle size (solid circles) and the molecular-dynamics results for the rate of bond formation between the nanoparticles (solid squares with lines) as a function of sintering temperature [41]. The MD simulation is able to provide detailed picture of the onset of sintering through a dramatic change in the rate of bond formation around the experimental sintering temperature of 1500 K. Both MD simulation and neutron scattering data indicate that the sintering temperature in n-




Fig. 9. Neutron-scattering results (solid circles, left axis) indicating the onset of sintering in nanophase SiC by a rapid increase in the average particle size above 1500 K. Comparison is made with molecular-dynamics results (solid squares, right axis) for the rate of bond formation between nanoclusters as a function of temperature.

SiC is considerably lower than that (2100–2400 K) for coarse-grained SiC.

MD simulations provide further details regarding the structural changes in micropores during sintering and consolidation. Pore analysis is performed by dividing the MD box into small cubic cells (length  $\sim 0.4$  nm) and identifying contiguous unoccupied cells using the breadth-first search approach. The morphology of pores is determined from the variations of the average pore radius  $R$  and the interface width  $W$  (a measure of the pore-surface roughness) with the pore volume  $V$ . Fig. 10 shows a log–log plot of  $R$  and  $W$  as a function of the pore volume  $V$  for the n-SiC at  $2.83$  g/cm<sup>3</sup>. The mean pore radius  $R$  scales as  $V^\eta$  with  $\eta = 0.42$  giving a fractal dimension of  $2.4$  ( $d = 1/\eta$ ). The interface width  $W$  scales as  $V^\mu$  with  $\mu = 0.47$ . Similar results are found for nanophase silica, where  $\eta = 0.47$  ( $d = 1/\eta = 2.1$ ) and  $\mu = 0.51$ . Within the statistical uncertainty, we find no difference in the fractal dimension or the roughness exponent of pores in n-SiC or nanophase silica at different densities. A similar fractal dimension ( $d = 2.0$ ) has been observed in densified silica aerogels using small-angle neutron scattering [42].

There has been considerable debate over the atomic configuration in the vicinity of grain boundaries of consolidated nanophase materials. While some experiments concluded that atoms in the grain boundary region were gas-like showing no organized structure




Fig. 10. Log–log plot of the average pore radius ( $R$ ) and interface width ( $W$ ) as a function of the pore volume ( $V$ ) for nanophase SiC. Solid lines are the best least-squares fits.

[43,44], other experiments claimed the presence of short-range order similar to that in liquids or amorphous solids [45,46]. To address these issues we have performed a complete structural characterization of nanophase silica glasses at various densities and of the fully consolidated n-SiC using large scale MD simulations.

We find that in nanophase silica glasses the short-range order (SRO) is similar to that of bulk silica glass. Both the nanophase and bulk silica glass contain corner-sharing  $\text{Si}(\text{O}_{1/4})_4$  tetrahedral units. Dramatic changes in the first sharp diffraction peak (FSDP) in the neutron-scattering static structure factor,  $S_n(k)$ , are observed for nanophase silica glasses. The FSDP yields information about the intermediate-range order (IRO) in glasses, and in the past decade it has been the focus of many experimental and computer simulation studies [47,48]. Fig. 11 shows the MD results for  $S_n(k)$  in the region of the FSDP for bulk silica glass and nanophase silica glasses at various densities. As seen in Fig. 11 the height of the FSDP in the bulk silica glass is 15% higher than that of the nanophase glasses. Additionally, the position of the FSDP in nanophase silica glasses shifts toward smaller  $k$  with respect to the bulk. We find that the Si–O partial static structure factor is largely responsible for the change in the height of the FSDP in the nanophase glasses at different densities. The behavior of the FSDP position




Fig. 11. Neutron-scattering static structure factor,  $S_n(k)$ , for bulk (solid line) and nanophase (dashed lines) amorphous  $\text{SiO}_2$ . Only the region around the FSDP is shown in order to illustrate the different behavior of the bulk and nanophase systems.

in the nanophase glasses is reflected in the Si–O and O–O partial static structure factors.

Analysis of the short-range order (SRO) in nanophase SiC reveals that while the interior of the clusters remains crystalline, the interfacial regions are more disordered having a structure similar to that of amorphous SiC. The peaks in pair-distribution functions and bond-angle distributions for atoms in interfacial regions of nanophase systems appear broadened when compared with the sharp peaks in the interior regions of the nanoparticles. From the Si–C pair distribution function, we find the average nearest-neighbor coordination for Si atoms reduces from 4 in the interior of particles to 3.5 in the interface regions, indicating that there are nearly the same number of three- fold and four-fold coordinated atoms in the interface regions. The pair-distribution functions, nearest-neighbor coordinations, and bond-angle distributions in the interfacial regions of n-SiC are found to be similar to those of bulk amorphous SiC. These results are consistent with neutron and X-ray diffraction measurements, and also with electron microscopy studies [49].

In both the nanophase silica glasses and nanophase SiC we find that the elastic moduli exhibit a power-law dependence on the density. This can be seen in Fig. 12, which shows a log–log plot of the Young's ( $E$ ) moduli with respect to the relative density for both nanophase systems. The nanophase and bulk silica glasses (solid squares) are given as a function of the bulk silica density; results for the nanophase SiC

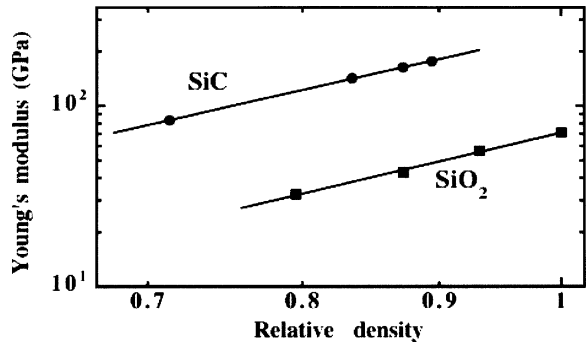


Fig. 12. Log–log plot of the Young's modulus for nanophase SiC (solid circles) and nanophase amorphous  $\text{SiO}_2$  (solid squares) as a function of the relative density. The nanophase SiC densities are relative to the bulk crystalline SiC density and the nanophase  $\text{SiO}_2$  densities are relative to the bulk amorphous  $\text{SiO}_2$  density. The solid lines are the best least-squares fits.

(solid circles) are shown as a function of the density relative to the bulk crystalline SiC. The solid lines are the best least-squares fits for each of the systems. The scaling exponents for the Young's modulus of nanophase silica and SiC are found to be  $3.5 \pm 0.2$  and  $3.34 \pm 0.03$ , respectively. Experimental measurements on high-density silica and carbon aerogels indicate very similar power-law dependence of elastic moduli on the density [50].

## 7. Conclusion

In summary, highly efficient  $O(N)$  multiresolution algorithms on parallel architectures have enabled us to perform multimillion atom molecular-dynamics simulations of various nanostructured materials. These systems include dynamics of oxidation of aluminum nanoclusters and properties and processes in nanostructured silicon carbide and nanostructured amorphous silica.

Petaflop computers expected to be build in ten years will enable trillion-atom simulations. This will allow the study of the effects of microstructures spanning diverse length scales above mesoscales. We also expect that interatomic potential models used in molecular-dynamics simulations will be further refined with inputs from first-principles calculations [51,52]. These atomistic simulations will be combined with continuum schemes based on finite-element methods to

model truly macroscopic materials properties [53,54]. With these developments, materials simulations will soon be able to predict properties of materials reliably in advance of fabrication.

## Acknowledgement

This work is partially supported by DOE (Grant No. DE-FG02-96ER45570), NSF (Grant No. DMR-9711903 and ASC-9701504), AFOSR (Grant No. F 49620-99-1-0250 and F 49620-98-1-0086), USC-LSU MURI (Grant No. F 49620-95-1-0452), NASA (Grant No. NCC 2-5320 and NAG2-1318), ARO (Grant No. DAAH04-96-1-0393), and Louisiana Education Quality Support Fund. Large-scale simulations were performed using parallel computers at Department of Defense's Major Shared Resource Centers under a DoD Challenge Applications Award.

## References

- [1] R.W. Siegel, *J. Phys. Chem. Solids* 55 (1994) 1097.
- [2] H. Gleiter, *Nanostructured Materials* 1 (1992) 1.
- [3] R.W. Siegel, H. Hahn, in: *Current Trends in the Physics of Materials*, M. Yussouff (Ed.) (World Scientific, Singapore, 1987) p. 403.
- [4] R.W. Siegel and et al., *J. Mater. Res.* 3 (1988) 1367.
- [5] R.W. Siegel, J.A. Eastman, in: *Proc. Mater. Res. Soc. Symp.*, Vol. 132, 1989, p. 3.
- [6] R.C. Cammarata, *Thin Solid Films* 240 (1994) 82.
- [7] R.W. Siegel, in: *Materials Interfaces: Atomic-Level Structure and Properties*, D. Wolf, S. Yip (Eds.) (Chapman Hall, London, 1992) p. 431.
- [8] R.W. Siegel, G.J. Thomas, *Ultramicros* 40 (1992) 376.
- [9] R.W. Siegel, in: *Physics of New Materials*, F.E. Fujita (Ed.) (Springer, Heidelberg, 1994) p. 65.
- [10] H. Gleiter, *Prog. Mat. Sci.* 33 (1989) 223.
- [11] H. Gleiter, in: *Prog. Mat. Sci.*, Vol. 33, J. Christian, P. Haasen, T. B. Massalski (Eds.) (Pergamon Press, NY, 1990) p. 223.
- [12] H. Gleiter, in: *Deformation of Polycrystals: Mechanisms and Microstructures*, N. Hansen et al. (Eds.) (Risø National Laboratory, Roskilde, 1981) p. 15.
- [13] J.C. Sánchez-López, A. Fernández, C.F. Conde et al., *Nanostruct. Mater.* 7 (1996) 813.
- [14] J.C. Sánchez-López, A.R. González-Elipse, A. Fernández, *J. Mater. Res.* 13 (1998) 703.
- [15] S. Sako, K. Ohshima, T. Fujita, *J. Phys. Soc. Jpn.* 59 (1990) 662.
- [16] C.E. Aumann, G.L. Skofronick, J.A. Martin, *J. Vac. Sci. Technol. B* 13 (1995) 1178.
- [17] T.G. Nieh, P. Luo, W. Nellis et al., *Acta Mater.* 44 (1996) 3781.
- [18] B.H. Suits, P. Apte, D.E. Wilken et al., *Nanostruct. Mater.* 6 (1995) 609.
- [19] N.F. Mott, *Trans. Faraday Soc.* 35 (1939) 1175.
- [20] N. Cabrera, N.F. Mott, *Rep. Prog. Phys.* 12 (1948–1949) 163.
- [21] L. Greengard, V. Rokhlin, *J. Comput. Phys.* 73 (1987) 325.
- [22] L. Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems* (MIT Press, Cambridge, MA, 1987).
- [23] M. Tuckerman, B.J. Berne, G.L. Martyna, *J. Chem. Phys.* 97 (1992) 1990.
- [24] P. Vashishta, R.K. Kalia, J.P. Rino et al., *Phys. Rev. B* 41 (1990) 12 197.
- [25] P. Vashishta, R.K. Kalia, A. Nakano et al., in: *Amorphous Insulators and Semiconductors*, M.F. Thorpe, M.I. Mitkova (Eds.) (Kluwer Academic, Netherlands, 1997) p. 151.
- [26] W. Jin, P. Vashishta, R.K. Kalia et al., *Phys. Rev. B* 48 (1993) 9359.
- [27] F.H. Streitz, J.W. Mintmire, *Phys. Rev. B* 50 (1994) 11 996.
- [28] R.P. Iczkowski, J.L. Margrave, *J. Amer. Chem. Soc.* 83 (1961) 3547.
- [29] A.K. Rappé, W.A. Goddard, *J. Phys. Chem.* 95 (1991) 3358.
- [30] W.J. Mortier, K. Van Genechten, J. Gasteiger, *J. Amer. Chem. Soc.* 107 (1985) 829.
- [31] G.H. Golub, C.F. Van Loan, *Matrix Computations* (Johns Hopkins Univ. Press, Baltimore, 1996).
- [32] A. Nakano, *Comput. Phys. Commun.* 104 (1997) 59.
- [33] S.W. Rick, S.J. Stuart, B.J. Berne, *J. Chem. Phys.* 101 (1994) 6141.
- [34] D.C. Rapaport, *Comput. Phys. Commun.* 62 (1991) 217.
- [35] M. Snir, S.W. Otto, S. Huss-Lederman et al., *MPI: The Complete Reference* (MIT Press, Cambridge, MA, 1996).
- [36] A. Nakano, T. Campbell, *Parallel Comput.* 23 (1997) 1461.
- [37] A. Nakano, *Concurrency: Practice and Experience* 11 (1999) 343.
- [38] J.R. Pilkington, S.B. Baden, *IEEE Trans. Parallel Distrib. Syst.* 7 (1996) 288.
- [39] I. Levin, D. Brandon, *J. Amer. Ceram. Soc.* 81 (1998) 1995.
- [40] A.J. Bourdillon, S.M. El-Mashri, A.J. Forty, *Philos. Mag. A* 49 (1984) 341.
- [41] A bond between Si and C atoms from two different clusters is said to exist if the separation between those atoms is less than or equal to the nearest-neighbor separation, 2.3 Å.
- [42] R. Sempéré, D. Bourret, T. Woignier et al., *Phys. Rev. Lett.* 71 (1993) 3307.
- [43] X. Zhu, R. Birringer, U. Herr et al., *Phys. Rev. B* 35 (1987) 9085.
- [44] J. Löffler, J. Weissmüller, H. Gleiter, *Nanostructured Materials* 6 (1995) 567.
- [45] G.J. Thomas, R.W. Siegel, J.A. Eastman, *Scripta Metall. Mater.* 24 (1990) 201.
- [46] D. Chen, *Nanostruct. Materials* 4 (1994) 753.
- [47] S.C. Moss, D.L. Price, in: *Physics of Disordered Materials*, D. Adler, H. Fritzsche, S. R. Ovshinsky (Eds.) (Plenum, New York, 1985) p. 77.
- [48] S.R. Elliott, *Phys. Rev. Lett.* 67 (1991) 711.
- [49] E.A. Stern, R.W. Siegel, M. Newville et al., *Phys. Rev. Lett.* 75 (1995) 3874.

- [50] J. Groß, J. Fricke, *NanoStruct. Mat.* 6 (1995) 905.
- [51] R. Car, M. Parrinello, *Phys. Rev. Lett.* 55 (1985) 2471.
- [52] M.L. Cohen, J.R. Chelikowsky, *Electronic Structure and Optical Properties of Semiconductors* (Springer, Berlin, 1989).
- [53] F.F. Abraham, J.Q. Broughton, E. Kaxiras, *Comput. Phys.* 12 (1998) 538.
- [54] E.B. Tadmor, R. Phillips, M. Ortiz, *Langmuir* 12 (1996) 4529.