ELSEVIER

Computer Physics Communications 105 (1997) 139–150

Computer Physics Communications

Fuzzy clustering approach to hierarchical molecular dynamics simulation of multiscale materials phenomena

Aiichiro Nakano [1]

Department of Computer Science, Concurrent Computing Laboratory for Materials Simulations, Louisiana State University, Baton Rouge, LA 70803-4020, USA

Received 2 April 1997

Abstract

A new algorithm is developed for large-scale, long-time molecular dynamics simulations of multiscale materials phenomena by combining a hierarchy of subdynamics: (i) rigid-body cluster dynamics for global conformational changes; (ii) implicit integration of Newton's equations for the coalescence of the clusters; and (iii) normal-mode analysis of fast atomic oscillations. Fuzzy clustering is used to facilitate the seamless integration of the multiple levels of abstraction. The fuzzy-body/implicit-integration/normal-mode (FIN) scheme speeds up a simulation of nanocluster sintering by a factor of 28 over a conventional explicit integration scheme. A parallel implementation of the scheme achieves an efficiency of 0.94 for a 12.7 million-atom nanocrystalline solid on 64 nodes of an IBM SP2 computer. © 1997 Elsevier Science B.V.

Keywords: Molecular dynamics; Parallel computing; Multiscale algorithm; Fuzzy clustering; Rigid-body dynamics; Implicit integration; Normal-mode analysis

1. Introduction

A new concept is emerging in materials development. The "multiscale structure control" attains seemingly contradictory properties (such as superior porosity and mechanical stability for catalytic applications) by controlling structures extending over multiple length scales [1,2]. Fig. 1 shows a computer model [3] of cluster-assembled ceramics [4,5]. Strength of such materials must be optimized with respect to atomic bonding (at a length scale of 10^{-10} m), pore (10^{-9} m) and interface (10^{-8} m) morphologies, and the shape and distribution of clusters ($> 10^{-7}$ m).

Molecular dynamics (MD) [6] is a powerful tool for understanding how atomistic processes are related to macroscopic phenomena. Recent developments in parallel processing technology and multiresolution numerical algorithms [7] have established large-scale MD simulations as a new research mode for studying macroscopic materials phenomena such as fracture [3,8–10]. However, many important processes (such as sintering and sol-gel processes) are characterized by time scales that are many orders-of-magnitude larger than atomic time

[1] E-mail: nakano@bit.csc.lsu.edu; URL: http://www.cclms.lsu.edu

0010-4655/97/$17.00 © 1997 Elsevier Science B.V. All rights reserved.
PII S0010-4655(97)00044-1

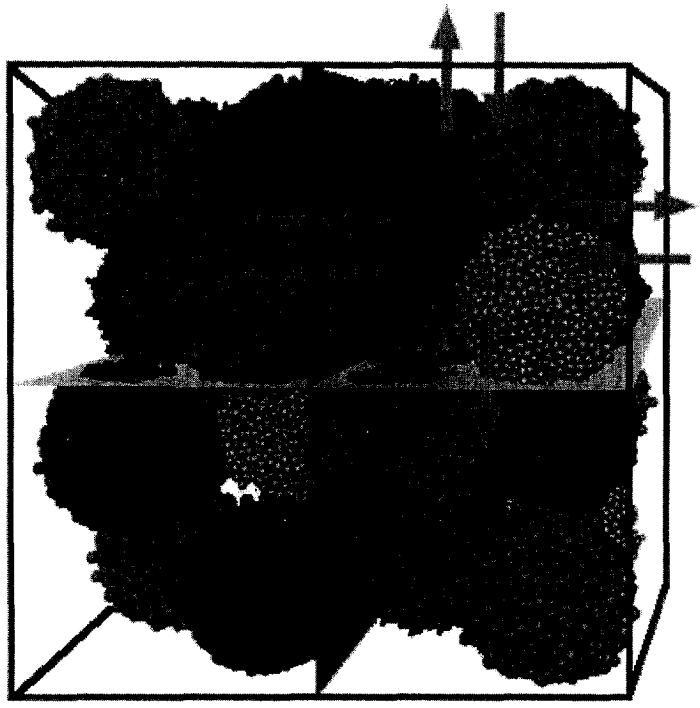

Fig. 1. MD configuration of nanocluster-assembled Si_3N_4. The system contains 108 clusters each consisting of 10052 atoms (in total of 1085616 atoms). Small spheres represent Si or N atoms. The figure also illustrates schematically the domain decomposition scheme used for parallel implementation. The walls represent partition boundaries to divide workloads. The arrows demonstrate the flow of atomic information during message passing.

scales [11]. The required large system size and long simulated time are beyond the scope of current simulation technologies.

Dynamic systems involving a wide range of time scales are called "stiff" [12]. Stiff systems are ubiquitous in many areas. An example is the conformational change of biomolecules [13]. In quantum chemistry, stiff harmonic action arising from the kinetic energy of quantum particles (characteristic time $\sim 10^{-17}$ sec) is inseparable from molecular motions on the time scale of 10^{-12} sec [14]. In mineralogical simulations, stiff atomic motions in structural units (such as SiO_4 tetrahedra) are coupled to slow rearrangement of these units [15]. In astrophysics, the billion-year evolution of the solar system is affected by planet motions on the time scale of a year [16].

Various approaches have been developed for long-time MD simulations of stiff systems. One approach introduces constraints to freeze high-frequency modes. This enables the use of a larger time step, assuming that the high-frequency modes are unimportant for global conformational changes [17,18]. In the subspace dynamics approach [19], low-frequency modes are selected by diagonalizing the dynamical matrix of the system. Other schemes use different time steps for different force components [20–22]. Ordinary differential equations can be integrated using a large time step by implicit-integration schemes [23]. For example, the implicit Euler integrator using time step Δt is a low-pass filter which selects motions with eigenfrequencies less than $1/\Delta t$; the frozen fast motions can be integrated separately by normal-mode analysis [13]. Recently, more accurate implicit integrators have been proposed which are symplectic [23]. Symplectic integrators preserve the phase-space volume, and this symplecticness is essential for the long-time stability of orbitals [24]. It is also possible to take advantage of the geometry of a particular system to speed up MD simulations. For example, an $O(N)$

algorithm has been developed for the constrained dynamics of linear-chain and tree-geometry molecules [25]. In addition, proper coordinate transformations often separate the degrees of freedom with different time scales and thereby speed up numerical integrations [14–16].

Because of the wide range of length and time scales involved in multiscale materials phenomena, the above long-time integrators alone are not sufficient for the simulation of these phenomena if they are used in the context of a single physical model. In fact, new hybrid approaches are emerging which introduce multiple levels of abstraction. For example, the quasi-continuum model uses direct atomistic calculations to provide inputs to the finite element analysis of continuum mechanics [26]. Kinetic Monte Carlo simulations of long-time semiconductor growth may be augmented by MD analysis of short-time local energy equilibration [27]. Merging multigrid methods with MD simulations has also been explored [28]. Efforts are underway to combine MD with first-principles electronic structure calculations to describe the breaking and formation of chemical bonds during macroscopic material processes [29,30].

In this paper, we develop a new long-time integrator for large-scale MD simulations of multiscale materials phenomena. The new algorithm combines a hierarchy of subdynamics: (i) quaternion-based rigid-body cluster dynamics [31] for global conformational changes; (ii) implicit integration of Newton's equations [23] for the coalescence of the clusters; and (iii) normal-mode analysis of fast atomic oscillations [13]. The greatest challenge, however, is to integrate these heterogeneous abstraction levels into a seamless, unified scheme. To facilitate such integration, we find it useful to introduce the concept of fuzzy clustering [32–37]. Fuzzy set theory [38] uses a membership function to describe the uncertainty of belonging. We use the fuzzy membership to define the clusters in the above hierarchical MD scheme.

The outline of this paper is as follows. Section 2 describes our fuzzy-body/implicit-integration/normal-model (FIN) algorithm. Numerical tests on the new scheme is presented in Section 3, and Section 4 deals with its implementation on parallel computers. Finally, Section 5 contains the summary.

2. Fuzzy-body/implicit-integration/normal-mode method

In this paper, sintering of nanometer-size clusters is simulated to test the new hierarchical MD scheme. Sintering is a process of densification by heating [11]. This process is used to consolidate cluster-assembled nanophase materials [3–5] such as shown in Fig. 1. The early stage of sintering involves three processes: (i) the relative rotation of clusters [39,40]; (ii) the formation and growth of necks between clusters via surface diffusion or viscous flow [11]; and (iii) local thermal motion of atoms which assists the first two processes (see Fig. 2). Processes (i) and (ii) are characterized by time scales ($> 10^{-9}$ sec) which are much longer than the characteristic time ($\sim 10^{-13}$ sec) of process (iii). A reliable long-time MD integrator must at least capture these processes.

In MD simulations, a system is represented by a set of atomic coordinates, $\{\boldsymbol{x}_i | i = 1, \ldots, N\}$, where N is the number of atoms. Time evolution of the system is governed by Newton's second law of motion [6],

$$m_i \frac{d^2 \boldsymbol{x}_i}{dt^2} = \boldsymbol{g}_i(\{\boldsymbol{x}_i\}) , \tag{1}$$

where m_i and $\boldsymbol{g}_i(\{\boldsymbol{x}_i\}) = -\partial V/\partial \boldsymbol{x}_i$ are the mass and force for the ith atom. The potential energy function $V(\{\boldsymbol{x}_i\})$ consists of a sum over atomic pairs and triples [41]. Given atomic coordinates $\boldsymbol{x}_i^n$ at time $t_n = n\Delta t$ (n = integer), Eq. (1) is numerically integrated for Δt using a reference system $\boldsymbol{r}_i$, The reference system represents two types of essential dynamics, i.e., relative rotation of clusters and fast atomic oscillations. The reference system is thus defined as a superposition, $\boldsymbol{r}_i = \boldsymbol{r}_{0i} + \boldsymbol{r}_{hi}$, where $\boldsymbol{r}_{0i}$ and $\boldsymbol{r}_{hi}$ are the collective and harmonic parts, respectively.

The collective part of the reference system represents the rigid-body motion of clusters, and it has $6M$ degrees of freedom for the translational and rotational motion of M clusters ($M \ll N$). Equations of motion for the

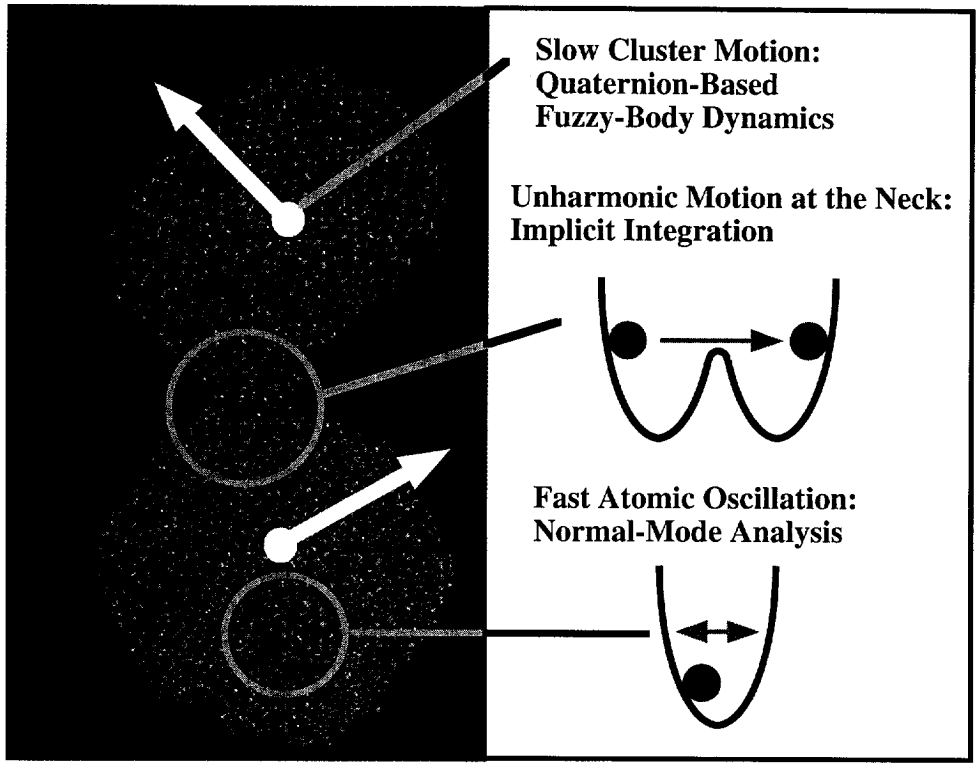

Fig. 2. Various physical processes involved in the sintering of nanoclusters. (i) Relative rotation of clusters is included by rigid-body dynamics with fuzzy clustering; (ii) unharmonic atomic motions lead to surface diffusion and the growth of the neck between clusters, and these motions are included by implicit integration of Newton's equations; (iii) thermal atomic motions assist the above diffusion process, and these high-frequency motions are dealt with normal-mode analysis.

clusters are

$$M_c \frac{d^2 \boldsymbol{R}_c}{dt^2} = \boldsymbol{g}_c \,, \tag{2}$$

$$\frac{d\boldsymbol{L}_c}{dt} = \boldsymbol{T}_c \,, \tag{3}$$

where M_c, $\boldsymbol{R}_c$, and $\boldsymbol{L}_c$ are the mass, center-of-mass, and angular momentum of the cth cluster ($c = 1, \ldots, M$). Here, $\boldsymbol{g}_c = \sum_i P(i \in c) \boldsymbol{g}_i^{\text{inter}}(\{\boldsymbol{r}_{0i}\})$ is the force acting on the cth cluster and $\boldsymbol{T}_c = \sum_i P(i \in c) m_i \boldsymbol{r}_{0i} \times \boldsymbol{g}_i^{\text{inter}}(\{\boldsymbol{r}_{0i}\})$ is the torque. These quantities are computed from atomic forces $\boldsymbol{g}_i^{\text{inter}}(\{\boldsymbol{r}_{0i}\}) = -\partial V_{\text{inter}} / \partial \boldsymbol{r}_{0i}$ arising from the intercluster potential energy V_{inter}. Only those atomic pairs and triples whose elements belong to different clusters contribute to V_{inter}. The membership function $P(i \in c)$ describes the degree of association between atom i and cluster c. Eqs. (2) and (3) are integrated for Δt using the quaternion formulation of rigid-body dynamics in order to avoid the numerical singularity associated with angular coordinates [31].

We use fuzzy clustering [32–37] to merge the cluster dynamics seamlessly with the other dynamics. Fuzzy clustering uses fuzzy membership functions [38] whose range is continuous, [0,1]. To determine the membership function, we apply the principle of maximum entropy [34,35]. First a cost function $E_c(i)$ is introduced such that

$$E_c(i) = \tfrac{1}{2} \sum_{j(\neq i)} P(j \in c) v_{ij}(|\boldsymbol{x}_i - \boldsymbol{x}_j|) \,, \tag{4}$$

where $v_{ij}(x)$ is the atomic pair potential used in the interatomic potential energy V [41]. The cost function $E_c(i)$ is the cost for assigning atom i to cluster c. The membership function is determined to maximize the entropy, $S_i = -\sum_c P(i \in c) \ln P(i \in c)$, subject to the constants (i) $\sum_c P(i \in c) = 1$; and (ii) $\sum_c P(i \in c) E_c(i) = \langle E_i \rangle =$ constant. This constrained minimization problem is solved by introducing the Lagrange multiplier $\beta = 1/k_B T$, to give the solution

$$P(i \in c) = \frac{\exp(-E_c(i)/k_B T)}{\sum_{c=1}^{M} \exp(-E_c(i)/k_B T)} . \tag{5}$$

Note that Eqs. (4) and (5) constitute an implicit equation system for $P(i \in c)$. This system is solved by fixed-point iterations [34] concurrently with the MD simulation.

The harmonic part $\boldsymbol{r}_{hi}$ of the reference system represents the fast oscillation of each atom around the local potential minimum, and it is defined as the solution to the linear system

$$m_i \frac{d^2 \boldsymbol{r}_{hi}}{dt^2} = -\boldsymbol{H}_{ii}(\boldsymbol{r}_{hi} - \boldsymbol{x}_i^*) , \tag{6}$$

where $\boldsymbol{x}_i^*$ is the local potential minimum near $\boldsymbol{x}_i^n$ and the Hessian matrix

$$\boldsymbol{H}_{ij} = \frac{\partial^2}{\partial \boldsymbol{x}_i \partial \boldsymbol{x}_j} V(\{\boldsymbol{x}_i\}) \tag{7}$$

is evaluated at $\boldsymbol{x}_i^*$. Eq. (6) describes independent harmonic motion of individual atoms, and it can be integrated analytically in terms of trigonometric functions [13].

The residual system is defined as $\boldsymbol{z}_i = \boldsymbol{x}_i - \boldsymbol{r}_i$, and it satisfies

$$m_i \frac{d^2 \boldsymbol{z}_i}{dt^2} = \boldsymbol{g}_i(\{\boldsymbol{z}_i + \boldsymbol{r}_i\}) - \boldsymbol{g}_i^{\text{inter}}(\{\boldsymbol{r}_{0i}\}) + \boldsymbol{H}_{ii}(\boldsymbol{r}_{hi} - \boldsymbol{x}_i^*) . \tag{8}$$

The residual system is expected to vary slowly, since the rapidly oscillating harmonic motions have been subtracted. Therefore, Eq. (8) is integrated by an implicit integration scheme using Δt which is much larger than atomic time scales. One implicit integration scheme [23] has the same algorithmic structure as the conventional velocity-Verlet integrator [6], except that the force $\boldsymbol{g}_i^{n+1}$ is obtained by solving a nonlinear equation system,

$$\boldsymbol{g}_i^{n+1} = \boldsymbol{g}_i\left(\left\{\boldsymbol{x}_i^{n+1} + \frac{\Delta t^2}{2m_i} \boldsymbol{g}_i^{n+1}\right\}\right) - \boldsymbol{g}_i^{\text{inter}}(\{\boldsymbol{r}_{0i}^{n+1}\}) + \boldsymbol{H}_{ii}(\boldsymbol{r}_{hi}^{n+1} - \boldsymbol{x}_i^*) . \tag{9}$$

This integration scheme is symplectic, and it is stable for an arbitrarily large Δt [23]. Solving Eq. (9) is equivalent to the minimization of the dynamics function [13]

$$\Phi(\{\boldsymbol{x}_i\}) = \sum_{i=1}^{N} \frac{m_i}{\Delta t^2}\left[\boldsymbol{x}_i - \boldsymbol{x}_i^{n+1} + \frac{\Delta t^2}{2m_i}(\boldsymbol{g}_i(\{\boldsymbol{r}_{ci}^{n+1}\}) - \boldsymbol{H}_{ii}(\boldsymbol{r}_{hi}^{n+1} - \boldsymbol{x}_i^*)\right]^2 + V(\{\boldsymbol{x}_i\}) . \tag{10}$$

In the variational formulation, Eq. (10) is minimized with respect to $\{\boldsymbol{x}_i\}$. (We use the truncated Newton method for this minimization [13].) The solution $\boldsymbol{x}_i^*$ to the minimization problem is then used to compute the self-consistent force, $\boldsymbol{g}_i^{n+1} = (2m_i/\Delta t^2)(\boldsymbol{x}_i^* - \boldsymbol{x}_i^{n+1})$, which in turn is used to accelerate the velocity for the residual system, $\boldsymbol{z}_i$.

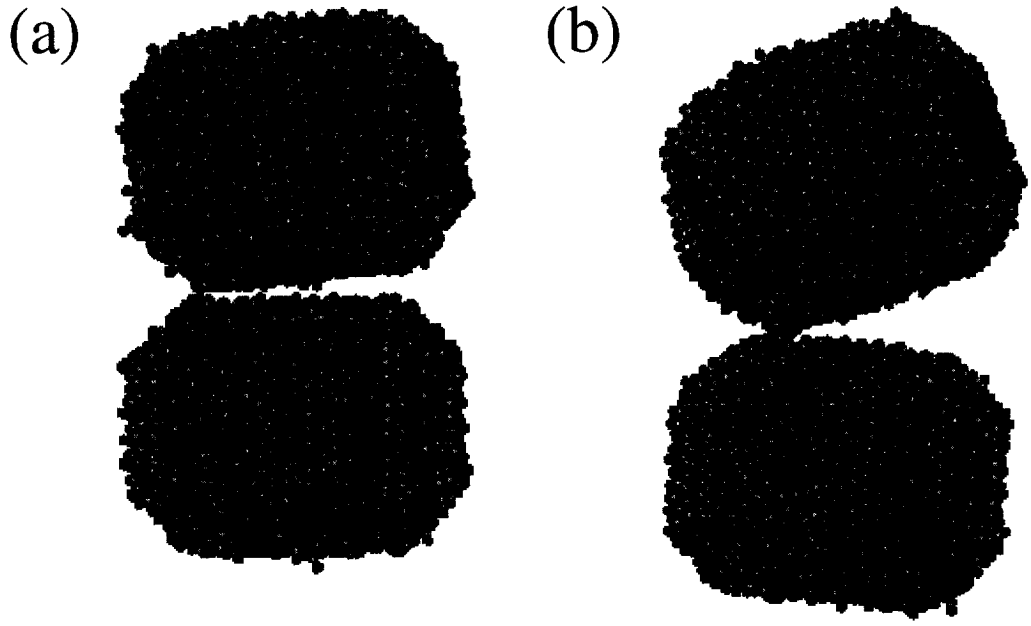

Fig. 3. (a) Starting MD configuration of two 20335-atom Si_3N_4 clusters. (b) The same system after 2×10^{-11} s. The FIN scheme is used for the integration of MD equations.

3. Numerical experiments

The fuzzy-body/implicit-integration/normal-mode (FIN) scheme is applied to the simulation of two silicon nitride (Si_3N_4) nanoclusters each consisting of 20335 atoms. We compare four integration schemes for the Newton's equations:

(i) Explicit integration with the velocity Verlet scheme [6] using $\Delta t = 2 \times 10^{-15}$ s.
(ii) Rigid-body (RB) cluster dynamics by freezing the other $3N-6M$ degrees of freedom ($M = 2$) completely. In this scheme, crisp membership is used instead of the fuzzy membership, i.e., $P(i \in c)$ is a Boolean function whose value is either 0 or 1.
(iii) The rigid-body/implicit-integration/normal-mode (RIN) scheme which is equivalent to the FIN scheme except that the crisp membership is used for $P(i \in c)$.
(iv) The FIN scheme using the fuzzy membership function, Eq. (5).

In the RB, RIN, and FIN schemes, $\Delta t = 10^{-12}$ s is used as a time step.

A reliable long-time integrator must first reproduce the long-time global motion of the clusters correctly. Figs. 3a shows the starting configuration of the clusters, and Fig. 3b is the final configuration at time 2×10^{-10} s obtained by the FIN integration scheme. The translational and rotational motions of the clusters in the FIN scheme are indistinguishable from the results of the explicit integration scheme. The FIN scheme thus describes the global motion of the clusters accurately. Significance of such rotation for the sintering of nanophase materials has been emphasised by Tsuruta et al. [39].

Sintering progresses through the formation of a "neck" between two clusters (see Fig. 2). Fig. 4a shows a close-up of the neck region at 1.5×10^{-10} s in the FIN scheme. The membership value of each atom to the top cluster ($c = 1$) is represented in gray scale; the brighter atoms have the larger values of $P(i \in 1)$. In Fig. 4b, $P(i \in 1)$ for the atoms are plotted as a function of their z coordinates. The atoms at the neck region are characterized by membership values which are intermediate between 0 and 1.

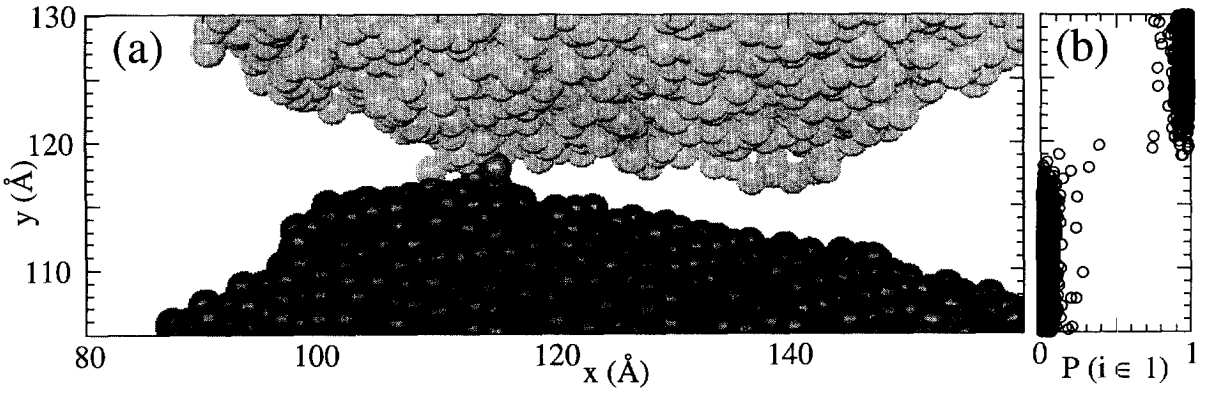

Fig. 4. (a) A close-up of the neck region at 1.5×10^{-10} s in the FIN scheme. The membership of an atom to the top cluster ($c = 1$) is represented in gray scale; the brighter atom has the larger value of $P(i \in 1)$. (b) The membership function $P(i \in 1)$ for the atoms as a function of their z coordinates.

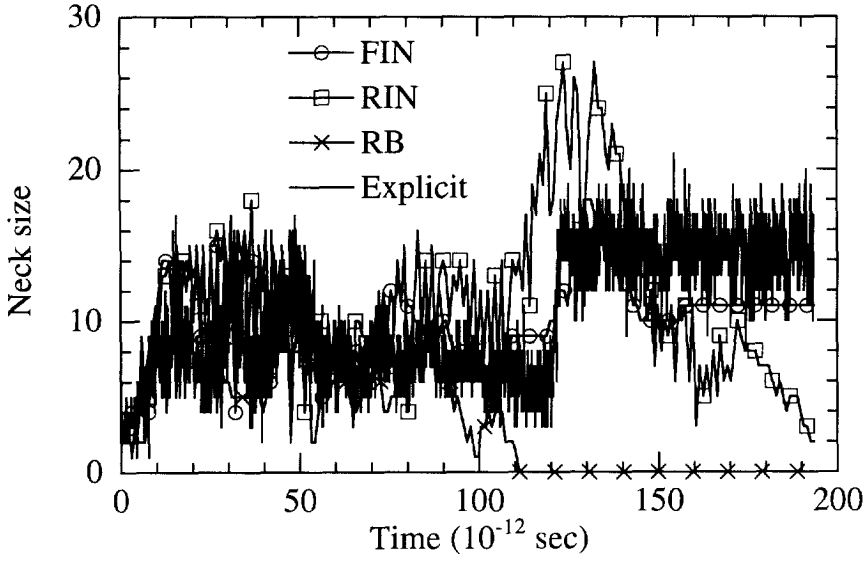

Fig. 5. Neck size between two Si_3N_4 clusters as a function of time. The circles, squares, and crosses represent the results with the FIN, RIN, and RB schemes, respectively. They are compared with the result of the explicit integration scheme (solid curve).

Fig. 5 shows the time evolution of the neck size which is defined as the number of atomic pairs whose elements belong to different clusters but are within a cut-off length (2 Å). We note that the rigid-body dynamics cannot describe the neck formation. In fact, the two clusters bounce back and become apart in this scheme. Accordingly, the neck size in Fig. 5 becomes 0 after 1.0×10^{-10} s with the RB scheme. Both the RIN and FIN schemes reproduce the sudden neck formation at 1.2×10^{-10} s which is observed in the explicit-integration scheme. However, the RIN scheme does not represent the step-functionlike increase in the neck size correctly. Rather, the neck size overshoots and then decreases to almost 0. Fuzzy clustering enables the FIN scheme to reproduce the sudden increase at 1.2×10^{-10} sec, as well as the subsequent plateau in the neck size until the end of the simulation at 2×10^{-10} s.

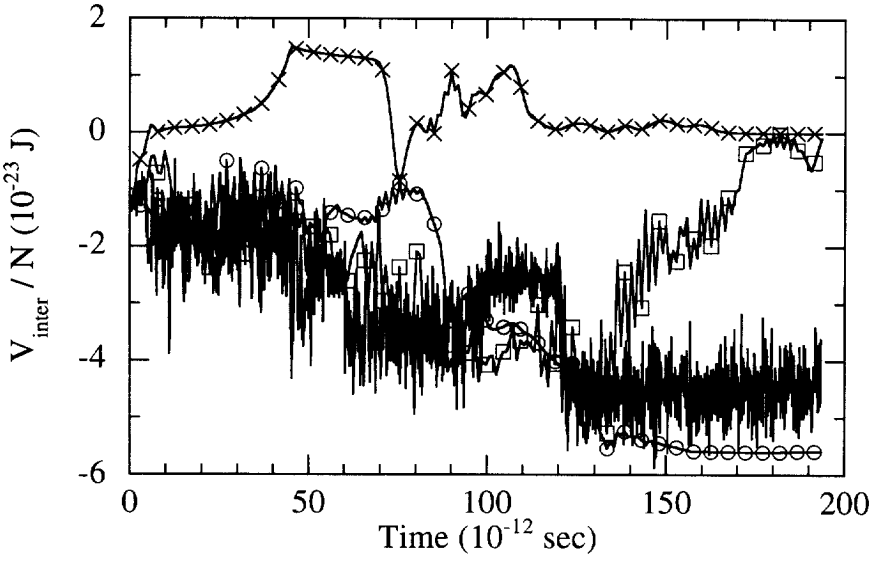

Fig. 6. Intercluster potential energy, V_{inter}, per atom as a function of time. The FIN, RIN, and RB schemes (represented by circles, squares, and crosses) are compared with the result of the explicit integration scheme (solid curve).

Neck formation is also monitored by the intercluster potential energy V_{inter}, as shown in Fig. 6. The FIN scheme reproduces the long-time behavior of V_{inter}, including the gradual decrease of V_{inter} before 1.2×10^{-10} sec and the plateau after that. The RB scheme does not describe the lowering of V_{inter} by neck formation at all. The RIN scheme reproduces the gradual decrease of V_{inter} well, but not the sudden drop of V_{inter} at 1.2×10^{-10} s.

In summary, the FIN scheme successfully describes the essential long-time dynamics during sintering. First, its rigid-body reference system $\boldsymbol{r}_{ci}$ captures the global translational and rotational motion of clusters. Second, the implicit integration of the residual system z_i accurately describes the neck formation. Thermal atomic motion represented by $\boldsymbol{r}_{hi}$ assists the neck formation by adding random fluctuations to the initial configuration for the implicit-integration step. Finally, the fuzzy clustering enables the description of sudden neck formation, which is not captured with the hard clustering using crisp membership.

Simulation of a 40 670-atom Si_3N_4 system for 10^{-10} sec in the FIN scheme takes 5.4 hrs on a single processor Digital Alpha 2100/275 (clock cycle = 275 MHz). It takes 150 hrs to simulate the same system using the standard velocity-Verlet integrator. The FIN scheme thus speeds up a conventional MD integration scheme by a factor of 28. It is only possible by a hybrid computation scheme such as FIN to simulate multimillion-atom systems for technologically important time scales.

4. Parallel implementation

The FIN algorithm is also implemented on distributed-memory parallel computers using a domain decomposition scheme [7]. The physical system is divided into subdomains of equal volume according to the uniform mesh topology, and each subdomain is assigned to a processor (see Fig. 1). The MPI (message passing interface) standard [42,43] is used to implement the message passing for exchanging atomic information among processors.

The execution time $T(p, N)$ of the parallel FIN algorithm for an N-atom systems on p processors is written as a sum of the computation and communication times: $T(p, N) = T_{comp}(p, N) + T_{comm}(p, N)$. The computation time is modeled as

$$T_{\mathrm{comp}}(p, N) = c_1 N/p + c_2 \log_2 p\,, \tag{11}$$

where the first term represents the computation of the local subdomain in each processor. The second term is the overhead for computing cluster properties globally. (An additional constant cost for performing the rigid-body cluster dynamics is negligible). The communication time consists of two terms: (i) the message-passing cost for performing global summation of cluster properties; (ii) the time required for copying the boundary-atom information among processors,

$$T_{\mathrm{comm}}(p, N) = c_3 \log_2 p + c_4 (N/p)^{2/3}\,. \tag{12}$$

The latter is necessary to compute interatomic potentials with a finite cut-off r_c. This cost is proportional to the number of "surface" atoms which are located within r_c from the partition boundaries.

We define the speed of the program as the number of atoms times the number of time steps calculated per unit time. The speedup is the speed of the program on p processors divided by the single-processor speed. The parallel efficiency E is the speedup for p processors divided by p [44,45]. If we use the memory-bounded scaling [44,45] such that the problem size N is proportional to the number of processors p, we obtain

$$E = \frac{T_{\mathrm{comp}}(1, N/p)}{T(p, N)} = \left[1 + \frac{c_2 + c_3}{c_1} p \log_2 p/N + \frac{c_4}{c_1} (p/N)^{1/3}\right]^{-1}. \tag{13}$$

In Eq. (13), the term proportional to $p \log_2 p/N$ arises from the overhead for performing global summations. The term proportional to the surface-to-volume ratio $(p/N)^{1/3}$ of each subdomain is due to the message-passing of boundary atoms.

The performance of the parallel FIN algorithm is tested on the IBM SP2 computer at the Maui High Performance Computing Center. The simulated system is aggregates of Si_3N_4 clusters each consisting of 49 728 atoms (cluster radius is 50 Å). We use the memory-bounded scaling, $N = 198\,912p$ ($1 \leq p \leq 64$, p is a multiple of 2). The largest system contains 12 730 368 atoms on 64 processors. The processors are logically arranged in a three-dimensional cube of size $p_x \times p_y \times p_z$ such that $\max(p_x, p_y, p_z) \leq 2 \times \min(p_x, p_y, p_z)$. Fig. 7 shows the total execution time and the time spent for communication as a function of the number of atoms N. The execution time is for integrating one FIN step ($\Delta t = 10^{-12}$ sec) involving one of Newton's iterations. The time increases only slightly as N increases, implying a nearly perfect parallel efficiency. In our numerical experiments, the granularity, $N/p \sim 2 \times 10^5$, is large, and accordingly both of the overheads in Eq. (13) are small.

Fig. 8 shows the parallel efficiency E and the communication overhead, $T_{\mathrm{comm}}(p, N)/T(p, N)$, as a function of the number of processors p. For the largest system (12.7 million atoms on 64 processors), the parallel efficiency $E = 0.94$ and communication overhead is only 2% of the total execution time.

5. Summary

A new algorithm is developed for long-time MD simulations of cluster-assembled nanophase materials. The algorithm combines: (i) rigid-body cluster dynamics; (ii) implicit integration of Newton's equations; and (iii) normal-mode analysis of fast atomic oscillations. Fuzzy clustering is used to integrate the multiple levels of abstraction seamlessly. The new scheme using a time step of 10^{-12} s speeds up a conventional explicit integration scheme using a time step of 2×10^{-15} sec by a factor of 28. A parallel implementation of the scheme achieves an efficiency of 0.94 for a 12.7-million-atom Si_3N_4 nanocrystalline solid on 64 nodes of an IBM SP2 computer.

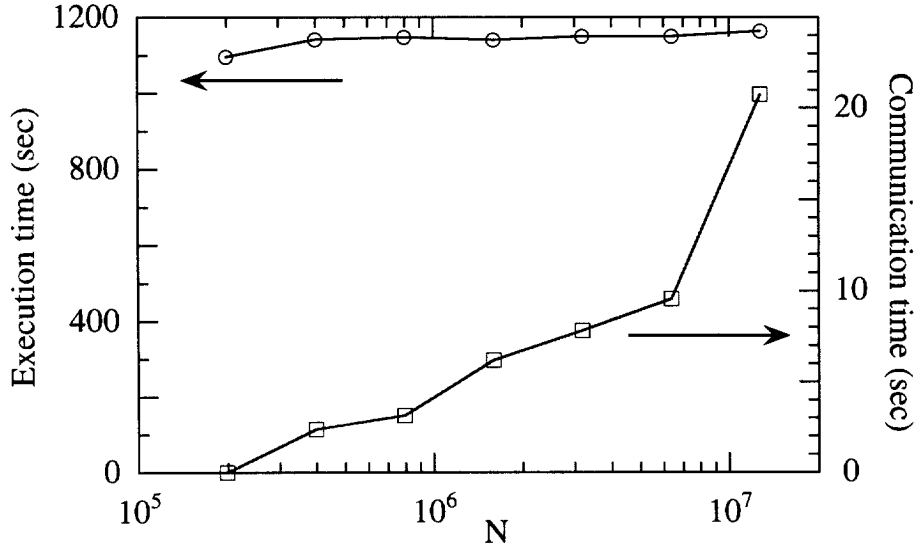

Fig. 7. Total execution time (circles) and the time spent for communication (squares) of the parallel FIN program on an IBM SP2 computer as a function of the number of atoms N. The system consists of α-Si_3N_4 nanoclusters each containing 49724 atoms. The number of atoms scales with the number of processors p as $N = 198912p$.

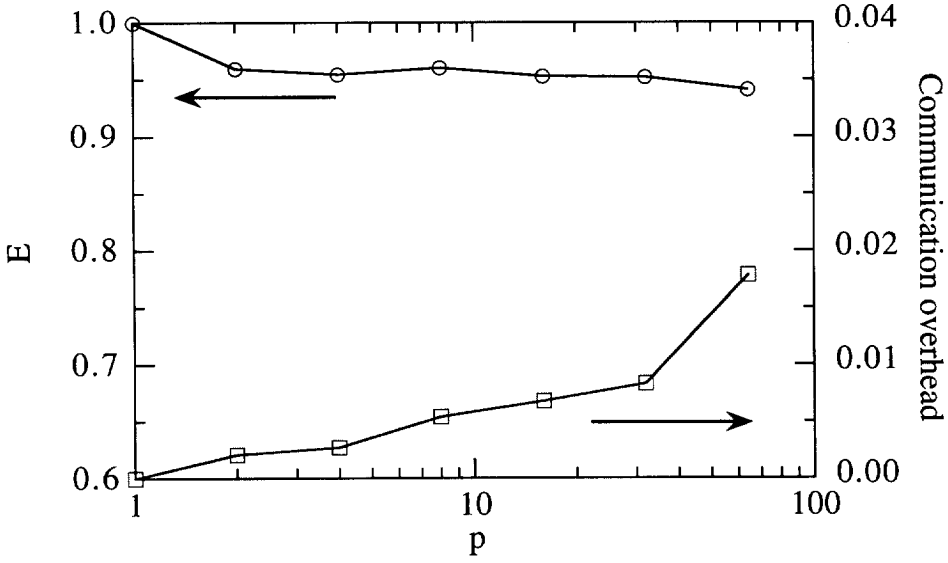

Fig. 8. Parallel efficiency (circles) and communication overhead (squares) of the parallel FIN program on an IBM SP2 computer for the same system as is shown in Fig. 7.

Acknowledgements

This work was supported by Army Research Office, Grant No. DAAH04-96-1-0393, National Science Foundation CAREER Program, Grant No. ASC-9701504, Louisiana Education Quality Support Fund (LEQSF), Grant No. LEQSF(96-99)-RD-A-10 and DoD/LEQSF(96-99)-03, and the Petroleum Research Fund, Grant No. 31659-AC9. Numerical experiments were performed on the IBM SP2 computer at the Maui High Perfor-

mance Computing Center. The computations were also performed on parallel architectures at the Concurrent Computing Laboratory for Materials Simulations (CCLMS) at Louisiana State University. The facilities in the CCLMS were acquired with the Equipment Enhancement Grants from the LEQSF. The author wishes to thank Drs. Priya Vashishta and Rajiv K. Kalia for fruitful discussions. The author also wishes to thank Dr. Tsuruta for his help in preparing Figs. 1 and 2, and for providing the initial configuration for the MD simulations.

References

[1] W.E. King et al., Mat. Sci. Eng. A 191 (1995) 1.
[2] Research and Development of Synergy Ceramics (New Energy and Industrial Technology Development Organization, Tokyo, 1994).
[3] R.K. Kalia, A. Nakano, K. Tsuruta and P. Vashishta, Phys. Rev. Lett. 78 (1997) 689;
R.K. Kalia, A. Nakano, A. Omeltchenko, K. Tsuruta and P. Vashishta, Phys. Rev. Lett. 78 (1997) 2144.
[4] R.W. Siegel, in: Atomic-Level Properties of Interface Materials, D. Wolf and S. Yip, eds. (Chapman and Hall, London, 1992).
[5] H. Gleiter, Nanostructured Mat. 1 (1992) 1.
[6] M.P. Allen and D.J. Tildesley, Computer Simulation of Liquids (Oxford Univ. Press, Oxford, 1987).
[7] R.K. Kalia, S.W. de Leeuw, A. Nakano and P. Vashishta, Comput. Phys. Commun. 74 (1993) 316;
A. Nakano, P. Vashishta and R.K. Kalia, Comput. Phys. Commun. 77 (1993) 303;
A. Nakano, R.K. Kalia and P. Vashishta, Comput. Phys. Commun. 83 (1994) 197.
[8] F.F. Abraham, D. Brodbeck, R.A. Rafey and W.E. Rudge, Phys. Rev. Lett. 73 (1994) 272;
F.F. Abraham, Phys. Rev. Lett. 77 (1996) 869.
[9] S.J. Zhou, D.M. Beazley, P.S. Lomdahl and B.L. Holian, Phys. Rev. Lett. 78 (1997) 479.
[10] A. Nakano, R.K. Kalia and P. Vashishta, Phys. Rev. Lett. 73 (1994) 2336; 75 (1995) 3138;
A. Omeltchenko, J. Yu, R.K. Kalia and P. Vashishta, Phys. Rev. Lett. 78 (1997) 2148.
[11] G.W. Scherer, Sol-Gel Science (Academic Press, San Diego, 1990).
[12] A.C. Hindmarsh and L.R. Petzold, Comput. Phys. 9(2) (1995) 148.
[13] G. Zhang and T. Schlick, J. Comput. Chem. 14 (1993) 1212.
[14] M.E. Tuckeman, D. Marx, M.L. Klein and M. Parrinello, J. Chem. Phys. 104 (1996) 5579.
[15] K.D. Hammonds, M.T. Dove, A.P. Giddy and V. Heine, Am. Mineralog. 79 (1994) 1207.
[16] J. Wisdom and M. Holman, Astron. J. 102 (1991) 1528.
[17] J.-P. Ryckaert, G. Ciccotti and J.C. Berendsen, J. Comput. Phys. 23 (1977) 327.
[18] H.C. Andersen, J. Comput. Phys. 52 (1983) 24.
[19] B. Space, H. Rabitz and A. Askar, J. Chem. Phys. 99 (1993) 9070.
[20] W.B. Streett, D.J. Tildesley and G. Saville, Mol. Phys. 35 (1978) 639.
[21] M.E. Tuckeman, B.J. Berne and G.J. Martyna, J. Chem. Phys. 94 (1991) 6811.
[22] H. Grubmuller, H. Heller, A. Windemuth and K. Schulten, Mol. Sim. 6 (1991) 121.
[23] R.D. Skeel, G. Zhang and T. Schlick, SIAM J. Sci. Comput. 18 (1997) 203.
[24] V.I. Arnold, Mathematical Methods of Classical Mechanics (Springer-Verlag, Berlin, 1989).
[25] A. Jain, N. Vaidehi and G. Rodriguez, J. Comput. Phys. 106 (1993) 258;
M. Mathiowetz, A. Jain, N. Karasawa and W.A. Goddard, Proteins 20 (1994) 227.
[26] E.B. Tadmor, R. Phillips and M. Ortiz, Langmuir 12 (1996) 4529;
M.I. Baskes, Langmuir 12 (1996) 4535.
[27] A. Madhukar, W. Yu, R. Viswanathan and P. Chen, Mat. Res. Soc. Symp. Proc. 408 (1996) 413.
[28] T. Schlick and A. Brandt, IEEE Comput. Sci. & Eng. 3(3) (1996) 78;
T. Schlick and A. Brandt, Multigrid Tutorials with Applications to Molecular Dynamics, A. Brandt and T. Schlick, eds. (Oct. 10–12, 1995, Rehovot, Israel).
[29] The quantum-mechanical/molecular-mechanical (QMMM) approach embeds a quantum-mechanical description of charge-transfer and chemical reactions in a classical description of environmental atoms, see M.J. Field, P.A. Bash and M. Karplus, J. Comput. Chem. 11 (1990) 700;
C.S. Carmer, B. Weiner and M. Frenklach, J. Chem. Phys. 99 (1993) 1356;
F. Maseras and K. Morokuma, J. Comput. Chem. 16 (1995) 1170. The QMMM method is included in software packages such as GAMESS [M.W. Schmidt et al., J. Comput. Chem. 14 (1993) 1347].
[30] Macro-atomistic-ab initio-dynamics (MAAD) coalition has been formed by F. Abraham, J. Broughton, H. Gao, E. Kaxiras, R. Phillips and X. Xu, see F.F. Abraham, Phys. Today 50(2) (1997) 15.
[31] D.J. Evans and S. Murad, Mol. Phys. 34 (1977) 327.
[32] J.C. Dunn, J. Cybern. 3 (1974) 32.

[33] J.C. Bezdek, Pattern Recognition with Fuzzy Objective Function Algorithms (Plenum, New York, 1981);
J.C. Bezdek and R.J. Hathaway, IEEE Trans. Neural Network 3 (1992) 787.
[34] K. Rose, E. Gurewitz and G.C. Fox, Phys. Rev. Lett. 65 (1990) 945.
[35] E. Trauwaert, L. Kaufman and P. Rousseeuw, Fuzzy Sets Sys. 42 (1991) 42;
P. Rousseeuw, L. Kaufman and E. Trauwaert, Comput. Stat. Data Anal. 23 (1996) 135.
[36] M.-S. Yang, Math. Comput. Model. 18 (1993) 1.
[37] Recent works on fuzzy clustering can be found in: Proc. 5th Int. Conf. Fuzzy Sys., F.E. Petry and D.H. Kraft, eds. (IEEE, Piscataway, 1996).
[38] L.A. Zadeh, Inform. Contr. 8 (1965) 338.
[39] K. Tsuruta, A. Omeltchenko, R.K. Kalia and P. Vashishta, Europhys. Lett. 33 (1996) 441.
[40] H. Zhu and R.S. Averback, Mater. Sci. & Eng. A 204 (1995) 96.
[41] P. Vashishta, R.K. Kalia, J.P. Rino and I. Ebbsjö, Phys. Rev. B 41 (1990) 12197.
[42] W. Gropp, E. Lusk and A. Skjellum, Using MPI (MIT Press, Cambridge, 1994).
[43] M. Snir. S. Otto, S. Huss-Lederman, D. Walker and J. Dongarra, MPI: The Complete Reference (MIT Press, Cambridge, 1996).
[44] X.-H. Sun and J. Gustafson, Parall. Comput. 17 (1991) 1093.
[45] V. Kumar, A. Grama, A. Gupta and G. Karypis, Introduction to Parallel Computing (Benjamin/Cummings, Redwood City, 1994).