Efficient numerical algorithm for multiphase field simulations

Srikanth Vedantam[1] and B. S. V. Patnaik[2]
[1]*Department of Mechanical Engineering, Faculty of Engineering, National University of Singapore, 117576 Singapore*
[2]*Department of Computational Science, Faculty of Science, National University of Singapore, 117543 Singapore*
(Received 20 September 2005; published 19 January 2006)

Phase-field models have emerged as a successful class of models in a wide variety of applications in computational materials science. Multiphase field theories, as a subclass of phase-field theories, have been especially useful for studying nucleation and growth in *polycrystalline* materials. In theory, an infinite number of phase-field variables are required to represent grain orientations in a rotationally invariant free energy. However, limitations on available computational time and memory have restricted the number of phase-field variables used in the simulations. We present an approach by which the time and memory requirements are drastically reduced relative to standard algorithms. The proposed algorithm allows us the use of an unlimited number of phase-field variables to perform simulations without the associated burden on computational time or memory. We present the algorithm in the context of coalescence free grain growth.

DOI: 10.1103/PhysRevE.73.016703 PACS number(s): 81.10.Aj, 64.10.+h, 02.60.Cb

I. INTRODUCTION

Prediction of nucleation, grain growth, and concomitant microstructure in polycrystalline materials is of great technological importance. It is well known that grain size, shape, and topology significantly affect the mechanical properties. Much experimental, theoretical, and numerical effort has been focused on this problem. Computational methods are particularly exciting since they allow us to isolate and study the dominant controlling effects on the microstructural evolution. An easily understood computational approach would be to explicitly track all the grain boundaries and apply specific constitutive relations for their motion. Clearly, in problems involving complex microstructures, tracking each individual grain boundary is unfeasible. An alternative approach is the phase-field theory, first developed in the context of solidification, in which interfaces evolve naturally at each step of a numerical simulation. The latter can be achieved by the introduction of a single or many *phase-field variables*, whose *diffuse* values will differentiate the bulk from interfacial regions. Different groups have extended this class of diffuse interface models into a variety of applications with quite different spirits (see the recent reviews: [1–3], and references therein).

In the context of grain growth in polycrystalline materials, the phase-field theory comprises of several nonconserved *order parameters* $\eta_1(\boldsymbol{x},t), \eta_2(\boldsymbol{x},t), \ldots, \eta_Q(\boldsymbol{x},t)$ each of which notionally represents the volume fraction of grains of a particular orientation. Inside each grain, a single-order parameter representing the particular orientation, will remain nonzero (usually taking values ± 1) whereas all other order parameters will assume values close to zero. Across interfaces between two grains, order parameters corresponding to both orientations will take nonzero values, varying smoothly from one grain into the other, reflecting the diffuse nature of these interfaces. The evolution of these order parameters is assumed to follow standard relaxational dynamics governed by the Allen-Cahn (for nonconserved) or Cahn-Hilliard (for conserved) –type equations.

Several phase-field models have been proposed to simulate the grain growth kinetics of polycrystalline materials. To the best of our knowledge, one of the first successful models in which the grains of different crystallographic orientations are represented by a set of nonconserved order parameter fields is that of Chen and Yang [4]. This model has been extensively applied by Chen and co-workers for two-dimensional (2D) [5,6] and 3D systems [7]. Another class of multiphase field models was proposed by Steinbach *et al.* [8], with a constraint on the order parameters, $\Sigma_i \eta_i = 1$, i.e., the sum of all order parameters at a given point is unity. The physical interpretation of this constraint is that the order parameters represent the volume fraction of grains of different orientations. Warren and co-workers [9,10] have proposed a two-order parameter model, in which the crystalline order and predominant local orientation of the crystal are represented. Nevertheless, it appears that multiorder parameter models are more popular at this point. The effect of grain boundary energy and mobility anisotropy on grain growth has been studied [11,12] by extending the multiorder parameter model of Chen and Yang [4].

In the literature a variety of numerical approaches have been developed [6,13,14,16,15] to solve the governing evolution equations for the phase-field variables. One simple approach is to use a second-order finite-difference method on a uniform spatial grid with explicit time marching. To circumvent the stability restrictions posed by such schemes, fast Fourier transforms have also been widely applied [13,14]. These spectral methods are valid on uniform grids with periodic boundary conditions. Such restrictions can be overcome by using spatially adaptive grid designs and a finite element method [15,16]. All the existing algorithms employ only a small number of order parameters, since using more field variables readily translates into high computational time and memory costs.

There are two immediate drawbacks to using a limited number of order parameters. Since each order parameter represents a particular grain orientation, intermediate values of orientations are excluded. This may be viewed as restricting the allowed grain orientations to a finite discrete set of angles

(with loss of rotational invariance of the free energy) or each order parameter as representing a range of orientations. In the context of solidification of a polycrystalline material, this prevents the potential nucleation of a new randomly oriented grain. Obviously, a larger number of order parameters is required for a more continuous description of the frequency distribution of grains of different orientations. Furthermore, another unfortunate consequence of limiting the number of order parameters is the coalesence of grains during grain growth. Coalesence is the situation in which two grains which have the same order parameters come into contact and instantaneously form a single large grain. This leads to incorrect growth rates and unphysical grain shapes. The likelihood of coalescence involving a given grain varies with the probability $p(Q)$ that at least one of its second-nearest-neighboring grains shares the same orientation and is given by [17]

$$p(Q) \simeq 1 - \left(1 - \frac{1}{Q}\right)^Z, \tag{1}$$

where Z represents the average number of second-nearest-neighbor grains in the microstructure. In the limit of large Q, $p(Q)$ is approximately equal to Z/Q. In 2D simulations, an average grain has about 6 sides and hence, there will be 12 nearest-neighbor grains [18]. With $Q=48$, the probability of grain coalesence is about $p(48)=0.223$. Approximately 110 order parameters are needed to effectively suppress $p(Q)$ to below 10% [7]. The problem is even more acute in 3D with each grain having about 14 sides on average ($Z \approx 28$), resulting in a probability of coalesence of 0.36 [18]. Therefore, more than 200 order parameters would have to be employed at each grid point in order to keep $p(Q)$ below 13% in 3D [7]. There have been some attempts to avoid this effect when using smaller Q values, via a dynamic reorientation algorithm [7]. However, to our knowledge, there are no multiorder parameter models with the framework to handle an unlimited number of phase-field variables.

In this paper, we propose an efficient new algorithm for solving the Allen-Cahn equations of the multiorder parameter method. For a fixed number of order parameters, the proposed method significantly reduces the CPU time and memory requirements. Further, the new method *does not engender any additional computational expense associated with increasing the number of order parameters used.* We apply this method to achieve microstructure devoid of coalesence and quantify its role in overall growth rates of grains in the standard algorithm. The key idea for the scheme stems from the following observation, which is gainfully exploited.

In the standard approach, all order parameter values at every grid point are stored and evolved at every iteration step. However, at any given point away from the grain boundaries, only one order parameter is *active* (has a nonzero value). And near a grain boundary a small number of order parameters corresponding to the adjacent grains are active. Even at the junction of the boundaries between several adjacent grains, the number of active order parameters is small. Therefore, instead of evolving all the phase-field variables at every grid point, we evolve only the active phase-field variables. This allows us to implicitly track the location of the grain boundaries since points which have more than one active order parameter may be considered to be on the grain boundaries. The computational advantage arises from the fact that fewer order parameters are stored at each point compared to the standard method and fewer evolution equations solved. An additional computational benefit arises from the fact that the evolution equations are solved only on grain boundaries which occupy a smaller fraction of the total area, especially towards later times when grains become larger.

The successful implementation of this strategy along with the computational advantages are explained in Sec. II. For convenience, we refer to the modified algorithm as the active parameter tracking (APT) algorithm. A comparison of accuracy and computational time for the standard algorithm and the APT, and coalesence free grain growth simulations are presented in Sec. III.

II. THEORY AND COMPUTATIONAL ALGORITHM

First, we briefly describe the phase-field theory of Chen and co-workers (see [1,5,6], and references therein for more details). Next, we describe the proposed APT algorithm and its implementation.

A. Theory

Basic to phase-field theories are continuous field variables $\eta_1(\boldsymbol{x},t), \eta_2(\boldsymbol{x},t), \cdots, \eta_Q(\boldsymbol{x},t)$ (referred to as order parameters) which are functions of material points $\boldsymbol{x}$ and time t. In the context of polycrystalline materials, the order parameters represent the volume fraction of grains of a particular orientation. The evolution of the order parameters is specified by the time dependent Ginzburg-Landau equations for each of the Q order parameters

$$\frac{\partial \eta_i}{\partial t} = -L_i \frac{\delta \mathcal{F}}{\delta \eta_i}, \quad i = 1,2,\ldots,Q, \tag{2}$$

where L_i are mobility coefficients and $\mathcal{F}$ is the free energy functional. The free energy functional is taken to be of the form

$$\mathcal{F} = \int_V \left\{ f(\eta_1, \eta_2, \ldots, \eta_Q) + \sum_{i=1}^{Q} \frac{\kappa_i}{2} |\nabla \eta_i|^2 \right\} dV, \tag{3}$$

in which κ_i are positive constants for an isotropic gradient energy and f is a local free energy density.

Following Chen and Yang [4] we choose the specific form

$$f(\eta_1, \eta_2, \ldots, \eta_Q) = \sum_{i=1}^{Q} \left(-\tfrac{1}{2}\alpha \eta_i^2 + \tfrac{1}{4}\beta \eta_i^4\right) + \gamma \sum_{i=1}^{Q} \sum_{j \neq i}^{Q} \eta_i^2 \eta_j^2 \tag{4}$$

such that the energy is independent of the orientation of the grains. If the coefficients α, β, γ are chosen to be positive constants with $\alpha=\beta$ and $\gamma > \alpha/2$, the energy then has $2Q$ wells of equal depth at $(\eta_1, \eta_2, \ldots, \eta_Q) = (\pm 1, 0, \ldots, 0)$,

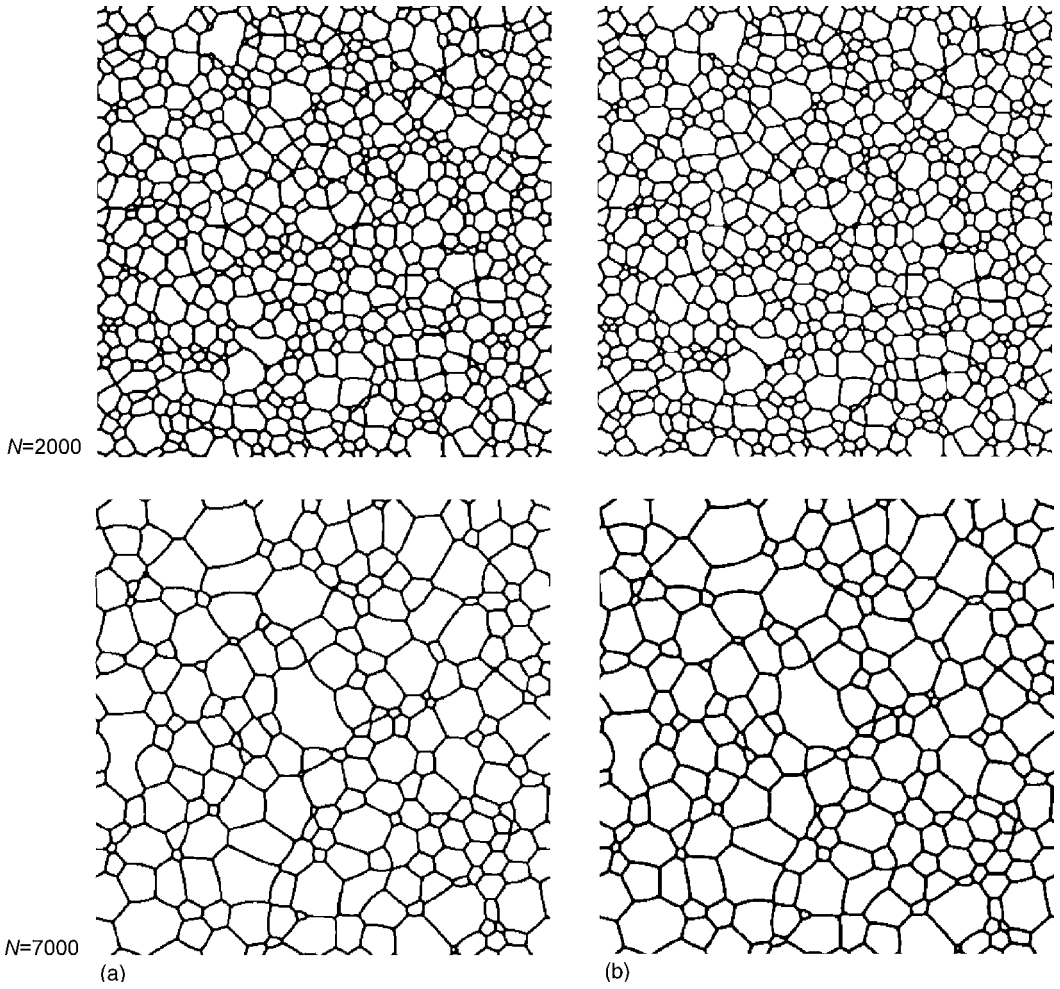

FIG. 1. Comparison of Fan and Chen's [6] algorithm against the present effective phase-field simulation. (a) Basic algorithm [6]; (b) present approach. Pixel by pixel comparison shows little difference between the results of the two methods.

$(0,\pm 1,\ldots,0),\ldots,(0,0,\ldots,\pm 1)$ representing $2Q$ possible orientations of the grains.

Substituting (4) and (3) in the evolution equations (2), we obtain the governing equations for grain growth

$$\frac{\partial \eta_i}{\partial t} = -L_i\left(-\alpha\eta_i + \beta\eta_i^3 + 2\gamma\eta_i\sum_{j\neq i}^{Q}\eta_j^2 - \kappa_i\nabla^2\eta_i\right), \quad i = 1,2,\ldots,Q, \tag{5}$$

which are a set of Q coupled, nonlinear, hyperbolic equations.

The governing equations (5) are solved using the finite difference method. Assuming a uniform spacing Δ and denoting $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$ as unit vectors in the x and y directions, respectively, the spatial derivatives are calculated using a two level discretization

$$\nabla^2\eta_i(\boldsymbol{x},t) = \frac{1}{2\Delta^2} \times\left\{\sum_{p=-1}^{1}\sum_{q=-1}^{1}\left(\frac{\eta_i(\boldsymbol{x}+p\Delta\boldsymbol{e}_1+q\Delta\boldsymbol{e}_2,t)-\eta(\boldsymbol{x},t)}{|p|+|q|}\right)\right\}, \quad p=q\neq 0. \tag{6}$$

The time derivatives are discretized using a forward-Euler difference

$$\frac{\partial \eta_i(\boldsymbol{x},t)}{\partial t} = \frac{\eta_i(\boldsymbol{x},t+\Delta t)-\eta_i(\boldsymbol{x},t)}{\Delta t}, \quad i=1,2,\ldots,Q \tag{7}$$

to obtain an explicit finite difference scheme. It should be pointed out that in explicit schemes, the stability of the scheme is governed by the time step chosen. Despite this drawback, we prefer this simple scheme to illustrate our proposed strategy of implicit boundary tracking via active parameter tracking algorithm.

B. Active parameter tracking algorithm

As mentioned earlier, at every grid point only a few phase-field variables are nonzero and contribute to the evolution of grains via boundary migration. At each point $O(\boldsymbol{x},t)$, the set of ordered pairs

$$\mathcal{P}(\boldsymbol{x},t) = \{(i,\eta_i): |\eta_i(\boldsymbol{x},t)| > \varepsilon\} \tag{8}$$

contains the list of active phase-field variables with magnitude greater than ε (chosen to be a small positive threshold

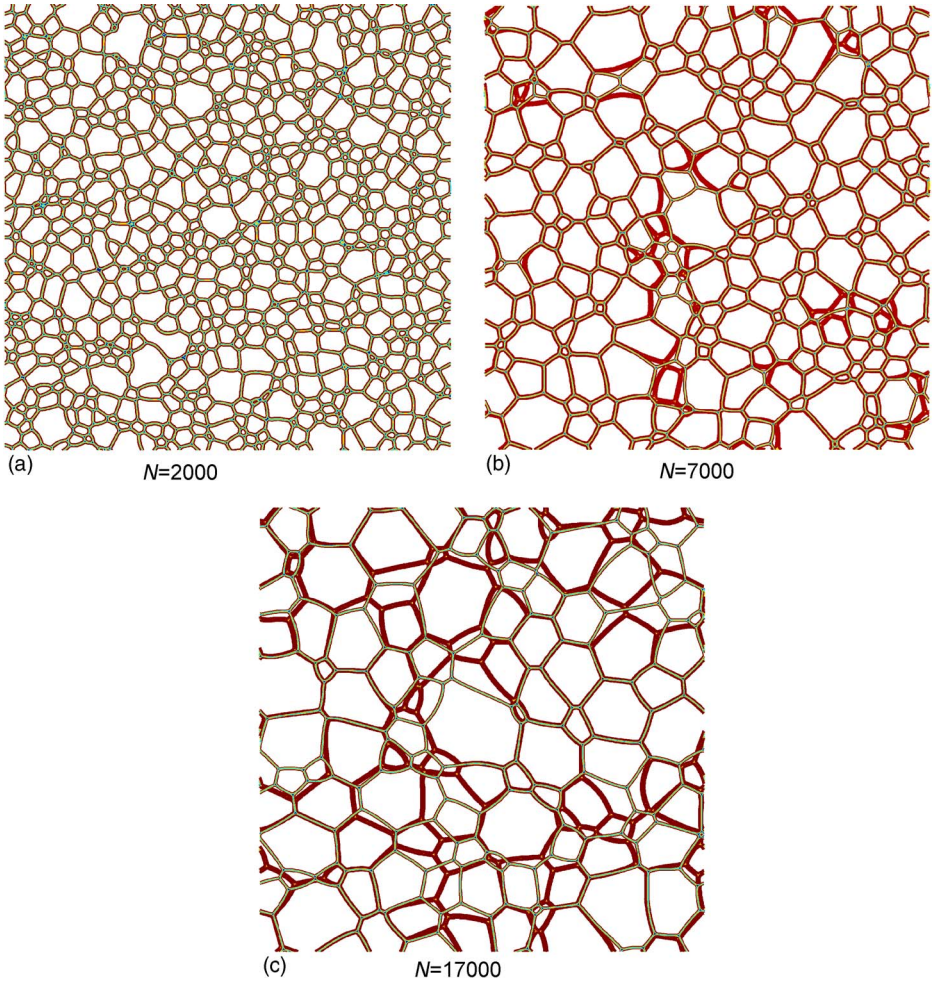

FIG. 2. (Color online) Overlay plot of microstructures at three different time levels of the simulation, N above refers to the number of integration steps: (a) 2000; (b) and 7000; and (c) 17 000. For the purpose of clarity, the microstructures from the present algorithm are represented with thin lines while those from the Fan and Chen [6] algorithm are presented with thick boundaries.

value).[1] In order to account for the possibility of a grain entering or leaving a point from neighboring points, we consider the *active parameter list*: $\mathcal{Q}(\boldsymbol{x},t)$ which is the union of the sets $\mathcal{P}$ at the two levels of nearest neighbors

$$\mathcal{Q}(\boldsymbol{x},t) = \bigcup_{p,q=-1}^{1} \mathcal{P}(\boldsymbol{x} + p\Delta\boldsymbol{e}_1 + q\Delta\boldsymbol{e}_2, t). \tag{9}$$

The number of active parameters at a point is $Q_a(\boldsymbol{x},t) = |\mathcal{Q}(\boldsymbol{x},t)|$ where $|\mathcal{S}|$ denotes the cardinality of set $\mathcal{S}$.

With this identification of the active phase-field variables at a point, only $Q_a(\boldsymbol{x},t)$ governing equations (5) are required to be solved

$$\frac{\partial \eta_i}{\partial t} = -L_i\left(-\alpha\eta_i + \beta\eta_i^3 + 2\gamma\eta_i \sum_{\substack{j=1 \\ \xi_j \neq i}}^{Q_a} \eta_{\xi_j}^2 - \kappa_i \nabla^2 \eta_i\right),$$
$$\forall (i,\eta_i) \in \mathcal{Q}(\boldsymbol{x},t)$$
$$\text{and } (\xi_j, \eta_{\xi_j}) \in \mathcal{Q}(\boldsymbol{x},t), j = 1, \ldots, Q_a. \tag{10}$$

[1]We use the notation $\mathcal{S}(\boldsymbol{x},t)$ to emphasize that the set $\mathcal{S}$ is to be constructed for each point $\boldsymbol{x}$ and at each time step t.

Specifically, the algorithm consists of the following steps repeated for each iteration (for each grid point in the domain).

(1) The list of all active phase-field variables from the first- and second-nearest neighbors $\mathcal{Q}$ is assembled from stored values of $\mathcal{P}$.

(2) Q_a governing equations (10) are solved for all $(i,\eta_i) \in \mathcal{Q}$. The Laplacian term is evaluated from neighboring active parameter sets $\mathcal{P}(\boldsymbol{x}+p\Delta\boldsymbol{e}_1+q\Delta\boldsymbol{e}_2,t), p,q=-1,0,1$.

(3) A new set $\mathcal{P}(\boldsymbol{x},t+\Delta t)$ is assembled from order parameters satisfying $|\eta_i(\boldsymbol{x},t+\Delta t)| > \varepsilon$ obtained from (2) and stored.

In more generic terms the steps (1) and (3) are akin to the mesh refinement and/or derefinement strategy often used by computational physicists in a variety of other settings (shocks, vortex structures, etc). However, in the present context they serve the purpose of implicitly tracking the grain boundaries using minimal data storage. As the grain boundary moves away from a given lattice point, the phase-field variable representing that grain is automatically not tracked any more, since it is no longer part of the active set of any of the neighbors.

We find that there is usually only one effective grain orientation the grid points away from grain boundaries, while two or more variables are likely at the interface between adjoining grains. As can be imagined, at initial times of the

simulation when triple or higher junctions are possible for small grains, a larger number of active order parameters exist at several points. With increasing grain size, there remain smaller numbers of active order parameters at each point. We found that the number of effective grain orientations never crossed 10 in our 2D simulations. It can be easily seen that increasing Q has no effect on the CPU time or memory requirements since the number of active order parameters at each point remains unchanged. This provides significant savings in CPU time and memory. Furthermore, when the number of active order parameters is one at each grid point and for all its neighbors (which is fairly typical inside a grain), a simple test condition allows us to determine if a given grid point is interior to a grain or not. If the grid point is completely interior to a grain, we do not solve the corresponding evolution equation for the order parameter at that point. At later times in the simulations when the grain boundary area is a small fraction relative to the overall domain, an additional computational savings is thus obtained.

We end our description of the APT algorithm with a remark on the development of the initial microstrucure from a liquid state. Chen and co-workers start with $\eta_i, i=1,2,\ldots,Q$ taking small random values at each grid point. We perform simulations using two methods: In the first approach, we perform simulations using the standard algorithm for a few hundred time steps using as large a Q as computationally feasible until grains form. Then we reassign unique order parameters to all the grains and switch to the APT algorithm. Alternatively, we start by assigning small random values to a single unique order parameter in small overlapping domains, with the aim of simulating coalesence free grain growth from the initial time $t=0$.

III. RESULTS AND DISCUSSION

In this section we examine four different aspects of the APT algorithm in relation to the standard algorithm: (i) accuracy, (ii) robustness, (iii) savings in CPU time and memory requirements, and (iv) effect of coalescence on grain growth.

To begin with, the standard algorithm is implemented following Fan and Chen [6] and these results are used for benchmarking the simulations of the APT algorithm. For the ease of simulation, we use the same parameter values provided by Fan and Chen [6]

$$\kappa_i = 2.0, \quad L_i = 1.0, \quad \alpha = \beta = \gamma = 1, \quad i = 1,2,\ldots,Q. \tag{11}$$

The spatial and time steps are chosen to be, respectively, $\Delta=2.0, \Delta t=0.25$. Periodic boundary conditions are applied in both the $x-$ and $y-$ directions. We choose an active parameter threshold $\varepsilon=10^{-6}$.

The following simulations are performed.

(1) Simulation using the standard algorithm of Fan and Chen [6] with $Q=48$ phase-field variables over the entire simulation time.

(2) Simulation using the standard algorithm up to $N=1000$ time steps, and continuing the simulation using the APT algorithm for subsequent time steps.

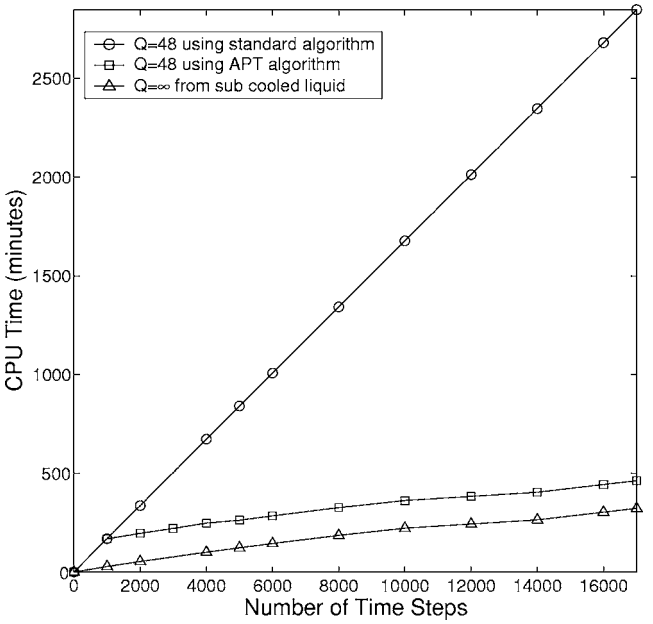

FIG. 3. CPU time against the number of integration steps of the algorithm. Note that the $Q=\infty$ from subcooled liquid simulation is equivalent to running the Fan and Chen algorithm with $Q \approx 25\,000$ phase-field variables. Yet, the time taken per iteration step is unchanged with respect to the $Q=48$ simulation.

(3) Simulation using the standard algorithm up to $N=1000$ time steps and reassigning unique order parameters to all the grains before continuing simulations using the APT algorithm.

(4) Simulation from an initial subcooled liquid state using small random values for unique order parameters in small overlapping subdomains and the APT algorithm.

From an initial subcooled liquid state (simulated by the random initial conditions), until the initial grain formation, the growth kinetics is governed by the basic algorithm. We use simulation (2) to verify the accuracy and robustness of our strategy compared with the basic algorithm results from simulation (1). Simulation (3) is used to study the role played by coalescence. We use simulation (4) to achieve the growth kinetics completely untainted by coalescence.

To visualize the microstructural evolution using the phase-field variables, we define the function

$$\psi(\boldsymbol{x},t) = \sum_{i=1}^{|\mathcal{P}|} \eta_{\xi_i}^2(\boldsymbol{x},t), \quad (\xi_i, \eta_{\xi_i}) \in \mathcal{P}, \tag{12}$$

which takes values of 1.0 within the grain and a significantly smaller positive values across the grain boundaries. Contours of $\psi(\boldsymbol{x},t)$ are shown in Fig. 1, using simulations (1) and (2), after 2000 and 7000 simulation time steps. The growth kinetics is exactly the same and the pixellated differences between the two visuals are fairly trivial. The smallest grains, where the most number of order parameters may be expected to be active compare very well in the two simulations. The results appear to vindicate our strategy of implicitly evolving the phase boundaries, as against all the variables within the domain in so far as the accuracy is concerned.

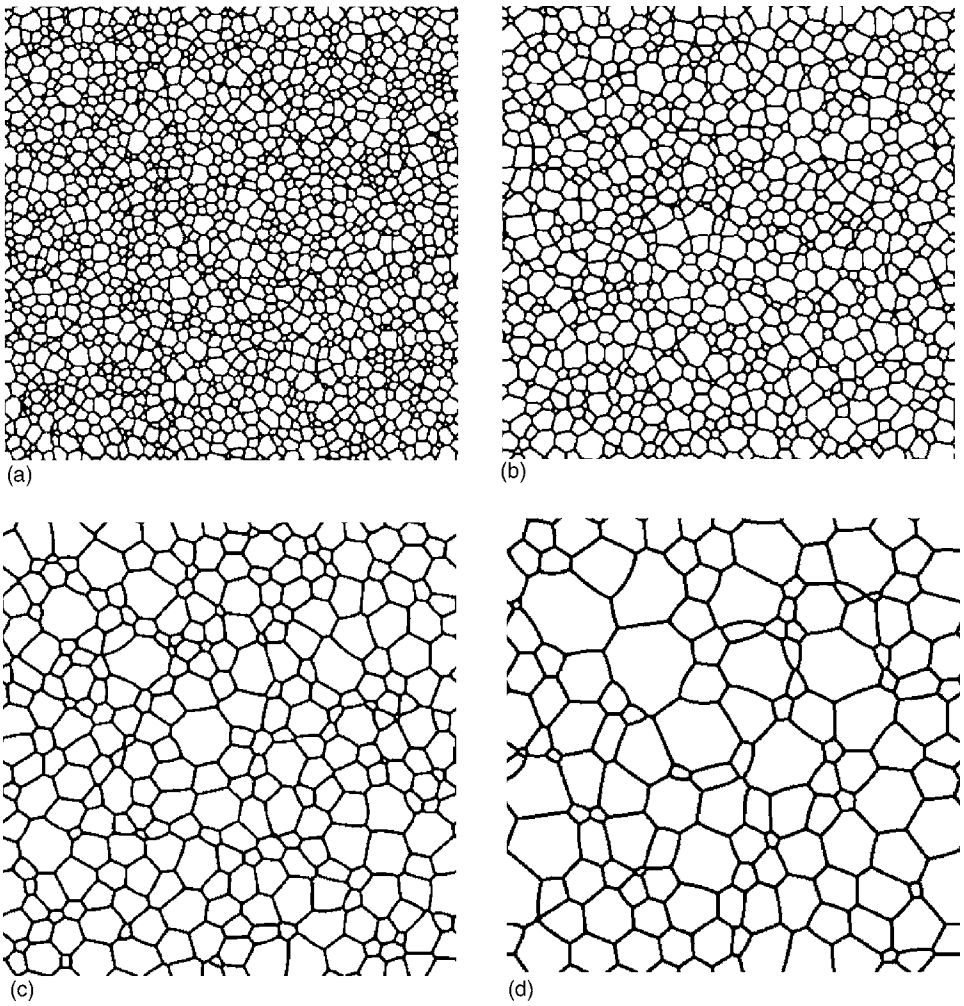

FIG. 4. Temporal evolution of coalescence free grain growth using unique grain orientation at each lattice point $N=$(a) 1000, (b) 2000, (c) 7000, and (d) 17 000.

The computational time taken for the basic algorithm and APT algorithm versus the number of time steps is shown in Fig. 3. The savings in computational time can clearly be seen to be very significant. While the computational time required for the standard algorithm increases linearly with the number of iterations as expected, the time required for the APT algorithm is slightly nonlinear. This is because the new method allows us to implicitly track the grain boundaries and perform evolution calculations only on grain boundaries. At later times, the grain boundaries occupy a smaller fraction of the overall area, leading to a fewer number of computations. Noting that the maximum number of the active order parameter Q_a at any point does not exceed 10, the computational memory requirements for the APT algorithm is equivalent to a standard algorithm simulation with $Q=10$. With dynamic memory allocation (which we have not implemented), the memory requirements at later simulation times can potentially be much smaller.

Next we focus our attention on simulation (3) which aims to achieve coalescence free simulations starting from $N=1000$ time steps. Each grain at 1000 steps is reassigned to a unique grain orientation (order parameter) and the simulations are continued thereafter. These results are compared to the simulations using basic algorithm [simulation (1)]. Figure 2 illustrates the role played by coalescence. The microstructure obtained from the present algorithm is overlayed with thin lines, while thick lines are used to distinguish the grains obtained using the basic algorithm. Around the central regions and in several other locations, we can visualize the grains affected by coalescence. Apparently there is no coalescence between $N=1000$ and $N=2000$, which is reflected in Fig. 2(a)) with a perfect match. However, when grains of different orientations existing between grains of like orientations are annihilated, it leads to an unphysical grain growth as depicted in Figs. 2(b) and 2(c) for $N=7000$ and $N=17000$, respectively.

In simulation (3), we have used the standard algorithm up to $N=1000$ time steps and continued using the unique order parameter reassignment and the APT algorithm for subsequent simulations. Using the basic algorithm from the $N=0$ has at least two limitations: (i) startup memory requirement is as large as the basic algorithm (approximately $32\times S\times S\times Q$ bytes assuming 16 byte double precision at two time levels, S being the grid size and Q the total number of phase-field variables). (ii) Between $N=0$ to $N=1000$, the small initial grains after nucleation will undergo significant coalescing (see Fig. 3). To circumvent these limitations, we have designed a scheme in which each lattice point is assigned a unique grain orientation, and the phase-field equations are solved. In the initial stages, every grid point has nine distinct

grain orientations for the possible evolution. The phase-field values above a certain threshold $|\eta_i| > 10^{-6}$ are carried over to the subsequent iterations. Thus using a Q value with as many as the number of grid points in the systems, a completely coalescence free growth is presented in Fig. 4. The growth rates of mean grain size is presented in Fig. 5. The mean grain size $\langle D \rangle$ is obtained by dividing the total domain area to an equivalent number of circular grains. The mean grain size $\langle D \rangle$ in the three simulations are matched at $N = 1000$ for a clear comparison of the growth rates for $Q = 48, 96$ and $Q = \infty$.[2]. For the latter, two grid sizes were used to check for the consistency of the present simulations. The role played by the coalescence even for $Q = 96$ is quite significant, which is completely avoided using the present scheme.

IV. CONCLUSIONS

In this paper, we describe a computationally efficient APT algorithm for multiorder parameter phase-field theories. In the proposed method we track the *active order parameters* at each grid point. Only the active order parameters are involved in the evolution equations and this provides an implicit means of tracking the grain boundaries in the phase-field method. This approach reduces computational effort for simulations involving a large number of order parameters and poses no additional computational burden in increasing the number of order parameters used without bound. We studied this approach in the context of polycrystalline grain growth. We have found that coalescence plays a role in the overall grain growth rates and we have compared these simulations in a coalescence free setting.

Several wide ranging applications exist in which a large number of grain orientations play an important role: for example, soap froth evolution, solidification of liquid into a polycrystalline solid, martensitic phase transformations in polycrystalline materials, polycrystalline ferroelectric transformations, etc. We believe that this method of implicitly tracking active order parameters at each point with a minimal additional computational burden will allow several more applications in computational physics to be accessible to solution using the phase field method.

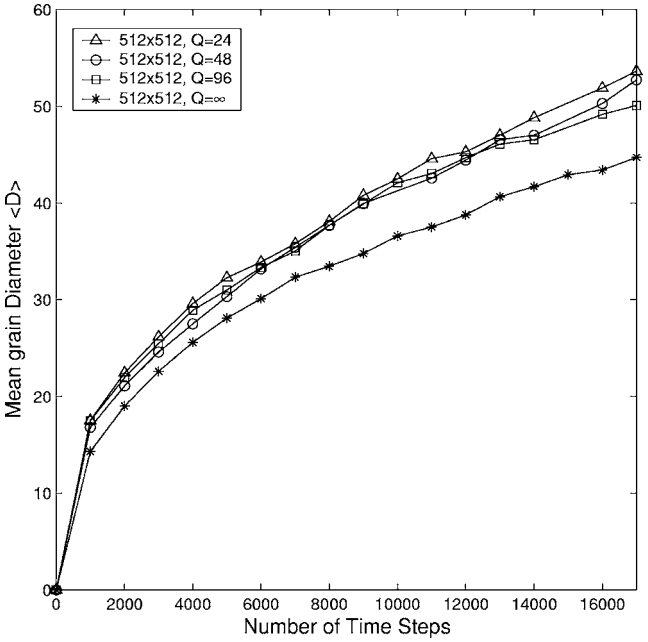

FIG. 5. Temporal evolution of average grain size for different Q values. The higher growth rates for $Q = 24$, 48, and 96 may be attributed to coalescence. The $Q = \infty$ simulation is run by seeding random values for unique phase-field variables in small overlapping domains.

[2]We use $Q = \infty$ as shorthand to denote unique order parameters for each grain in the simulation. For example, in simulation (3), reassigning unique order parameters for each grain at $N = 1000$ translates to running the simulation with $Q = 1180$ order parameters. In simulation (4), unique order parameters in small subdomains from $N = 0$ translates to running the simulation with $Q \approx 25\,000$ order parameters.

ACKNOWLEDGMENTS

The authors acknowledge the use of computational resources provided by the Department of Computational Science at the National University of Singapore.

[1] L. Q. Chen, Annu. Rev. Mater. Res. **32**, 113 (2002).
[2] W. J. Boettinger, J. A. Warren, C. Beckermann, and A. Karma, Annu. Rev. Mater. Res. **32**, 163 (2002).
[3] L. Granasy, T. Pusztai, and J. A. Warren, J. Phys.: Condens. Matter **16**, R1205 (2004).
[4] L. Q. Chen and W. Yang, Phys. Rev. B **50**, 15752 (1994).
[5] V. Tikare, E. A. Helm, D. Fan, and L. Q. Chen, Acta Mater. **47**, 363 (1999).
[6] D. Fan and L. Q. Chen, Acta Mater. **45**, 611 (1997).
[7] C. E. Krill and L. Q. Chen, Acta Mater. **50**, 3057 (2002).
[8] I. Steinbach, F. Pezzolla, B. Nestler, M. Sesselberg, R. Prieler, G. J. Schmitz, and J. L. L. Rezende, Physica D **94**, 135 (1996).
[9] R. Kobayashi, J. A. Warren, and W. C. Carter, Physica D **140**, 141 (2000).
[10] A. E. Lobkovsky and J. A. Warren, Phys. Rev. E **63**, 051605 (2001).
[11] A. Kazaryan, Y. Wang, S. A. Dregia, and B. R. Patton, Phys. Rev. B **61**, 14275 (2000).
[12] A. Kazaryan, Y. Wang, S. A. Dregia, and B. R. Patton, Phys. Rev. B **63**, 184102 (2001).
[13] L. Q. Chen and J. Shen, Comput. Phys. Commun. **108**, 147 (1998).
[14] J. Zhu, L. Q. Chen, J. Shen, and V. Tikare, Phys. Rev. E **60**, 3564 (1999).

[15] N. Provatas, N. Goldenfeld, and J. Dantzig, Phys. Rev. Lett. **80**, 3308 (1998).

[16] J. H. Jeong, N. Goldenfeld, and J. A. Dantzig, Phys. Rev. E **64**, 041602 (2001).

[17] M. P. Anderson, D. J. Srolovitz, G. S. Crest, and P. S. Sahni, Acta Metall. **32**, 783 (1984).

[18] B. Radhakrishnan and T. Zacharia, Metall. Mater. Trans. A **26A**, 167 (1995).