

Adaptive mesh computation of polycrystalline pattern formation using a renormalization-group reduction of the phase-field crystal model

Badrinarayan P. Athreya,¹ Nigel Goldenfeld,² Jonathan A. Dantzig,¹ Michael Greenwood,³ and Nikolas Provatas³

¹*Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, 1206 W. Green Street, Urbana, Illinois 61801, USA*

²*Department of Physics, University of Illinois at Urbana-Champaign, 1110 W. Green Street, Urbana, Illinois 61801, USA*

³*Department of Materials Science and Engineering, McMaster University, 1280 Main Street West, Hamilton, Ontario, Canada L8S 4L7*

(Received 15 January 2007; published 16 November 2007)

We implement an adaptive mesh algorithm for calculating the space and time dependence of the atomic density field in microscopic material processes. Our numerical approach uses the systematic renormalization-group formulation of a phase-field crystal model of a pure material to provide the underlying equations for the complex amplitude of the atomic density field—a quantity that is spatially uniform except near topological defects, grain boundaries, and other lattice imperfections. Our algorithm employs a hybrid formulation of the amplitude equations, combining Cartesian and polar decompositions of the complex amplitude. We show that this approach leads to an acceleration by three orders of magnitude in model calculations of polycrystalline grain growth in two dimensions.

DOI: [10.1103/PhysRevE.76.056706](https://doi.org/10.1103/PhysRevE.76.056706)

PACS number(s): 05.10.Cc, 81.15.Aa, 81.16.Rf, 46.15.–x

I. INTRODUCTION

A fundamental theoretical and computational challenge in materials modeling is that of concurrently treating phenomena over a wide range of length and time scales. For example, in studying the mechanical response of polycrystalline materials, one must take into account the dynamics and interactions of vacancies, impurities, dislocations, and grain boundaries, on time scales ranging from atomic vibrations to system-wide diffusion times.

Numerous approaches to handling the wide range of length scales have been proposed [1], including quasicontinuum methods [2–5], the heterogeneous multiscale method [6,7], multiscale molecular dynamics [8–11], multigrid variants [12], and phase-field models [13–16]. In general one can classify different techniques as being either atomistic or continuum, and differentiate them further by their characteristic time scale: density functional theory, for a quantum mechanical description of processes at the atomic time scale; molecular dynamics or Monte Carlo methods, appropriate for collective dynamics at the atomic scale; and coarse-grained descriptions involving continuum fields at the mesoscale on diffusive time scales. The difficulty of merging descriptions at different length and time scales limits the effective application of most of these methods. Lack of a continuous transition between scales can induce artifacts, such as spurious reflections in a transition region between two levels [7,17]. Further, any method using molecular dynamics is typically restricted to subnanosecond time scales, whereas many interesting phenomena during materials processing, such as microstructural pattern formation, recrystallization, heat and solute diffusion, dislocation glide, etc., occur over time scales which are typically greater than 10^{-6} s.

One continuum approach that has been used successfully, especially in the multiscale modeling of solidification problems [18] is the *phase-field method* [13]. Through the effective use of matched asymptotic analysis [14] and adaptive mesh refinement [19,20], the phase-field method has been

used to quantitatively study phenomena spanning several orders of magnitude in length, from micrometers to centimeters, and occurring on time scales of seconds to minutes. Extensions of the method by Kobayashi and co-workers [21,22], and Warren [16] also make it possible to model polycrystalline systems. Special forms of the free energy that incorporate strain energy have been used to model the qualitative features of strain-induced phase transformations [23–28]. The phase-field method represents a coarse-graining in space to length scales much greater than those of the interfaces and defects of interest in this work. As a result, the kinetic coefficients that emerge in the final continuum equations are inherently phenomenological, and can be related to experimentally measurable parameters only after a suitable asymptotic matching of the phase-field equations with corresponding sharp-interface models [18,29,30]. As such, traditional phase-field models do not fundamentally embody the emergent kinetic and elastoplastic mechanisms that originate at the atomic scale. Perhaps the most important limitation of phase-field models is that, in general, they do not preserve any record of the underlying crystal lattice, so that *ad hoc* approaches must be used to model the variety of phenomena that result from lattice interactions.

The phase-field crystal (PFC) methodology [31,32] is a promising extension of the phase-field model approach, in which the equilibrium free energy is constructed to produce periodic atomic density states, rather than ones uniform in space. The conserved dynamics of the PFC model then naturally reproduce many of the nonequilibrium dynamics arising in real polycrystalline materials. The PFC methodology is founded on the insight that a free energy functional that is minimized by a periodic field necessarily includes elastic energy, anisotropy, and symmetry properties of that field. Thus the model naturally incorporates all properties of a crystal that are determined by symmetry, as well as vacancies, dislocations, and other defects. Moreover, the PFC model represents the evolution of the system over a time scale that is much longer than the vibrational period of atoms

[$O(10^{-15}$ s)], but much shorter than the time scale of diffusive processes in the system, such as the viscous glide of dislocations, which typically occur over a time scale of $O(10^{-6}$ s). The PFC model yields a relatively simple and well-behaved partial differential equation (PDE) for the evolution of the time-averaged density, giving it access to phenomena occurring on atomic length scales, but over diffusive time scales. The PFC method is thus able to incorporate atomic-scale elasticity and the interaction of topological defects on the same time scales that govern diffusive processes during phase transformations in pure materials [32–34] and alloys [35].

As with any model that resolves at the atomic scale, the PFC model is limited in its ability to model systems of realistic dimensions, because the computational grid must resolve the periodicity of the field. For grid-converged results, a minimum of nine grid points per period are typically required. In a physical system, the periodicity represents interatomic distance, $O(10^{-10}$ m). Thus, to simulate a system having a characteristic dimension of $1 \mu\text{m}$ would require about 10^9 degrees of freedom per spatial dimension on a uniform computational mesh. This would be a heroic computation in two dimensions (2D), and well beyond reach in 3D, even with the use of massive parallelization. Furthermore, the periodic lattice precludes the effective use of traditional adaptive mesh refinement (AMR) algorithms.

Athreya, Goldenfeld, and Dantzig have recently described an approach to overcome this difficulty [36,37], using the perturbative renormalization group (RG) method [38,39] to systematically coarse-grain the PFC equation [40]. The basic idea is to obtain renormalization group equations of motion for the complex amplitude of the periodic density field, a quantity whose modulus and phase are spatially uniform except near regions of lattice disruption, as at grain boundaries and at topological defects. From the complex amplitude, it is possible to reconstruct the atomic-scale density field at least within the one-mode approximation, and to compute non-trivial materials properties and dynamics to high accuracy (within 1%) [36,37]. This approach, which we will sometimes denote as the PFC-RG method, is much faster than solving the PFC equation directly, because the complex amplitude varies on much larger spatial length scales than the density itself, thus permitting the use of an adaptively generated coarse mesh over much of the computational domain [36]. It is important to appreciate that the equations of motion for the complex amplitude must be rotationally covariant, in order that a polycrystalline material or heterogeneous microstructure can be represented without any preferred orientations imposed; this is readily achieved using renormalization group methods [40].

In a practical numerical implementation of the PFC-RG method, the reciprocal lattice vectors of the equilibrium crystal structure are represented within a particular basis, and there is the potential for interference between the Fourier components of the atomic density field and this basis [36] (also see Appendix A). This interference gives rise to artifactual “fringes” or “beats” in the corresponding Fourier components. While the overall density does not, of course, exhibit these interference fringes, their presence in the individual Fourier components means that, to be properly

resolved, an adaptive mesh algorithm (which deals with the individual Fourier components) must generate grid refinement in their vicinity. As a result, efficient computation becomes compromised.

The purpose of this paper is to develop a computationally efficient formulation of the PFC-RG method and a corresponding numerical algorithm that enables the implementation of adaptive mesh refinement up to micro and meso length scales, without being deflected by artifacts arising from the choice of basis set. The approach is to use a hybrid representation of the complex amplitude, switching between Cartesian and polar coordinates as appropriate in a seamless fashion to avoid beating and coordinate singularities. The resultant description is fast, accurate, and provides mesh refinement and coarsening in the physically correct locations, without artifacts arising from choice of basis or other implementation-dependent details. As such, our work represents a first step toward providing a systematic description of materials processing using continuum fields across all relevant length scales.

The remainder of this paper is organized as follows: We introduce the PFC model and the complex amplitude equations (interchangeably called the RG equations) in Sec. II and use them to derive the polar formulation which addresses the problem of beats, but also exhibits coordinate singularities which makes it unwieldy for numerical solution. We next present a hybrid formulation in Sec. III, which is a procedure for solving the Cartesian equations concurrently with a reduced form of the polar equations in different parts of the computational domain, allowing scope for effective use of AMR. Section IV presents numerical simulations using AMR and results, including efficiency benchmarks, that clearly demonstrate the computational advantage of our AMR-RG approach. Section V concludes and presents directions for future work.

II. EVOLUTION EQUATIONS

A. The PFC model and complex amplitude equations

In the PFC model of a pure material, the evolution of the density ρ is given by

$$\frac{\partial \rho}{\partial t} = \Gamma \nabla^2 \left(\frac{\delta \mathcal{F}}{\delta \rho} \right) + \eta, \quad (1)$$

where \mathcal{F} is the free energy functional, which can be written as $\mathcal{F} = \int d\vec{r} [f(\rho, \nabla^2 \rho, \dots)]$, where f is the local free energy density, Γ is a constant, and η is a stochastic noise with zero mean and correlations $\langle \eta(\vec{r}, t) \eta(\vec{r}', t') \rangle = -\Gamma k_B T \nabla^2 \delta(\vec{r} - \vec{r}') \delta(t - t')$. The specific form of \mathcal{F} is chosen such that at high temperatures \mathcal{F} is minimized by a spatially uniform liquid state, and at low temperatures by a spatially periodic “crystalline” phase. Furthermore, f must be chosen such that \mathcal{F} is independent of crystal orientation. These constraints naturally incorporate both elastic and plastic effects during deformations.

A free energy form that satisfies these criteria naturally produces mobile regions of liquid-solid coexistence separated by free surfaces, i.e., phase transformations. Elastic en-

ergy and defects in the crystalline phase arise from the requirement that \mathcal{F} be minimized by a spatially periodic density field that is independent of crystal orientation. Elder *et al.* [31,32] demonstrated these properties of the model for a variety of applications, including studies of grain boundary energy, liquid phase epitaxial growth, and the yield strength of nanocrystalline materials. The particular model they used made the following choice for the function f :

$$f = \rho[\alpha\Delta T + \lambda(q_o^2 + \nabla^2)^2]\rho/2 + u\rho^4/4, \quad (2)$$

where α , λ , q_o , and u are model parameters that can be specified to match some specific material properties, such as Young's modulus and lattice spacing [31,32]. In order to discuss the dynamical behavior of the PFC model, it is useful to rewrite the free energy in dimensionless units: $\vec{x} \equiv \vec{r}q_o$, $\psi \equiv \rho\sqrt{u/\lambda}q_o^4$, $r \equiv \alpha\Delta T/\lambda q_o^4$, $\tau \equiv \Gamma\lambda q_o^6 t$, and $F \equiv \mathcal{F}u/\lambda^2 q_o^{8-d}$ so that

$$F = \int d\vec{x} \{ \psi [r + (1 + \nabla^2)^2] \psi / 2 + \psi^4 / 4 \}. \quad (3)$$

In these units the conservation law of Eq. (1) becomes

$$\frac{\partial \psi}{\partial t} = \nabla^2 \{ [r + (1 + \nabla^2)^2] \psi + \psi^3 \} + \zeta \quad (4)$$

where $\langle \zeta(\vec{r}_1, t_1) \zeta(\vec{r}_2, \tau_2) \rangle = \mathcal{E} \nabla^2 \delta(\vec{r}_1 - \vec{r}_2) \delta(\tau_1 - \tau_2)$ and $\mathcal{E} \equiv uk_B T q_o^{d-4} / \lambda^2$. Equation (4), introduced by Elder *et al.* [31,32], will be referred to as the PFC equation in what follows. This equation can be used in any dimension by simply introducing the appropriate form for the Laplacian operators.

The spatial density ψ can be approximated in terms of the complex amplitudes A_j as

$$\psi \approx \sum_{j=1}^3 A_j e^{i\mathbf{k}_j \cdot \mathbf{x}} + \sum_{j=1}^3 A_j^* e^{-i\mathbf{k}_j \cdot \mathbf{x}} + \bar{\psi}, \quad (5)$$

where

$$\begin{aligned} \mathbf{k}_1 &= k_0(-i\sqrt{3}/2 - \vec{j}/2), \\ \mathbf{k}_2 &= k_0\vec{j}, \\ \mathbf{k}_3 &= k_0(i\sqrt{3}/2 - \vec{j}/2) \end{aligned} \quad (6)$$

are the reciprocal lattice vectors of a crystal with hexagonal symmetry, and k_0 is the dominant wave number of the pattern. For all the calculations shown in this paper, length has been scaled such that $k_0=1$, which corresponds to an interatomic spacing of $a_0=2\pi/(\sqrt{3}/2)$. The complex amplitude equations, which constitute a coarse-grained approximation to the PFC equation, were shown in our earlier work [36,37] to be given by

$$\frac{\partial A_j}{\partial t} = \tilde{\mathcal{L}}_j A_j - 3A_j |A_j|^2 - 6A_j \sum_{k:k \neq j} |A_k|^2 - 6\bar{\psi} \prod_{k:k \neq j} A_k^* \quad (7)$$

where $j, k \in [1, 3]$ and

$$\tilde{\mathcal{L}}_j = (1 - \nabla^2 - 2i\mathbf{k}_j \cdot \nabla)(-r - 3\bar{\psi}^2 - \{\nabla^2 + 2i\mathbf{k}_j \cdot \nabla\}^2) \quad (8)$$

is a rotationally covariant operator. The asterisk denotes complex conjugation. The parameters r (≤ 0) and $\bar{\psi}$ (≥ 0) control the bifurcation from a uniform liquid phase to a crystalline phase with hexagonal symmetry. Specifically, r is proportional to the temperature difference relative to a critical temperature T_c , while $\bar{\psi}$ is the mean density in the system. We refer to this form as the *Cartesian* representation because the amplitudes are expressed along each coordinate direction. The rotational covariance of the operator $\tilde{\mathcal{L}}$ permits the incorporation of multiple crystal orientations using only the basis vectors in Eq. (6) (see Appendix A).

B. Limitations of the Cartesian representation of Eq. (7)

A straightforward approach to solving Eq. (7) is to determine the real and imaginary parts of the complex amplitudes A_j directly, using the Cartesian definition. This leads to six equations that can be evolved concurrently using a suitable time integration scheme. The second-order finite-difference spatial discretizations of the Laplacian and gradient operators are given in Appendix B. This approach leads to limited success with AMR techniques because of the beats (this phenomenon is explained further in Appendix A and [36]).

To illustrate this effect, we simulated heterogeneous nucleation and growth of a two-dimensional film, randomly placing 12 randomly oriented crystals of initial radius 8π in a square domain of side 256π with periodic boundary conditions. The largest misorientation angle between grains was $\theta = \pi/12$. The amplitude equations in Cartesian form were solved using an adaptively evolving mesh algorithm (described in detail below). The model parameters were $r = -0.25$ and $\bar{\psi} = 0.285$, the smallest mesh spacing was $\Delta x_{\min} = \pi/2$, while the largest mesh spacing at any given time was $\Delta x_{\max} = 2^4(\Delta x_{\min})$ corresponding to five levels of refinement. On a uniform grid, this simulation requires $1025 \times 1025 = 1\,050\,625$ nodes with the PFC equation, and $513 \times 513 = 263\,169$ nodes with the amplitude equations. A time step of $\Delta t = 0.04$ was used.

Figure 1 shows the crystal boundaries and grid structure at various times during the simulation. The field plotted is the average amplitude modulus, $\sum_{j=1}^3 |A_j|/3$. Although the grid starts out quite coarse ($t=0$ and 88), at several locations in the computational domain, because of the large liquid fraction, this advantage falls off dramatically once the crystals evolve, collide, and start to form grain boundaries. In particular, once all the liquid freezes, only a few grains that are favorably oriented with respect to \mathbf{k}_j show any kind of grid coarsening at all. Those that are greatly misorientated with respect to \mathbf{k}_j lead to frequency beating, causing the number of nodes in the adaptive grid to increase rather than decrease. The polycrystal mesh shown in Fig. 1(f) has 219 393 nodes, which is very near that on a uniform grid. Therefore the adaptive refinement algorithm applied to a Cartesian formulation of Eq. (7) gives at best a marginal improvement over a fixed grid implementation. The main purpose of this paper is to present a methodology for overcoming this problem.

C. Complex amplitude equations in a polar representation

We find that the computational benefits of AMR are potentially greater if, instead of solving for the real and imaginary components of A_j , we solve for the amplitude moduli $\Psi_j = |A_j|$, and the phase angles $\Phi_j = \arctan[\text{Im}(A_j)/\text{Re}(A_j)]$ [here we use $\text{Im}()$ to denote the imaginary and $\text{Re}()$ to denote the real component of a complex number], which are spatially uniform fields irrespective of crystal orientation. Together these two fields constitute a *polar representation* of A_j .

In this section we derive evolution equations for Ψ_j and Φ_j directly from Eq. (7), by applying Euler's formula for a complex number, i.e., $A_j = \Psi_j e^{i\Phi_j}$, and then by equating the corresponding real and imaginary parts on the left- and right-hand sides of the resulting equations. In this manner we get the coupled system of equations

$$\begin{aligned} \frac{\partial \Psi_j}{\partial t} = & (r + 3\bar{\psi}^2)[- \Psi_j + \mathcal{C}^{\text{Re}}(\Psi_j, \Phi_j)] - [\mathcal{C}^{\text{Re Re}}(\Psi_j, \Phi_j) \\ & - \mathcal{C}^{\text{Im Im}}(\Psi_j, \Phi_j)] + [\mathcal{C}^{\text{Re Re Re}}(\Psi_j, \Phi_j) \\ & - \mathcal{C}^{\text{Re Im Im}}(\Psi_j, \Phi_j) - \mathcal{C}^{\text{Im Im Re}}(\Psi_j, \Phi_j) \\ & - \mathcal{C}^{\text{Im Re Im}}(\Psi_j, \Phi_j)] - 3\Psi_j \left(\Psi_j^2 + 2 \sum_{k \neq j} \Psi_k^2 \right) \\ & - 6 \frac{\bar{\psi}}{\Psi_j} \left(\prod_k \Psi_k \right) \cos \left(\sum_k \Phi_k \right) \end{aligned} \quad (9)$$

and

$$\begin{aligned} \frac{\partial \Phi_j}{\partial t} = & \{ (r + 3\bar{\psi}^2) \mathcal{C}^{\text{Im}}(\Psi_j, \Phi_j) - [\mathcal{C}^{\text{Re Im}}(\Psi_j, \Phi_j) \\ & + \mathcal{C}^{\text{Im Re}}(\Psi_j, \Phi_j)] + [\mathcal{C}^{\text{Im Re Re}}(\Psi_j, \Phi_j) \\ & - \mathcal{C}^{\text{Im Im Im}}(\Psi_j, \Phi_j) + \mathcal{C}^{\text{Re Im Re}}(\Psi_j, \Phi_j) \\ & + \mathcal{C}^{\text{Re Re Im}}(\Psi_j, \Phi_j)] \} / \Psi_j + 6 \frac{\bar{\psi}}{\Psi_j^2} \left(\prod_k \Psi_k \right) \sin \left(\sum_k \Phi_k \right), \end{aligned} \quad (10)$$

where

$$\begin{aligned} \mathcal{C}^{\text{Re}}(\Psi_j, \Phi_j) &= \text{Re} \left\{ \frac{(\nabla^2 + 2i\mathbf{k}_j \cdot \nabla)(\Psi_j e^{i\Phi_j})}{e^{i\Phi_j}} \right\}, \\ \mathcal{C}^{\text{Im}}(\Psi_j, \Phi_j) &= \text{Im} \left\{ \frac{(\nabla^2 + 2i\mathbf{k}_j \cdot \nabla)(\Psi_j e^{i\Phi_j})}{e^{i\Phi_j}} \right\}, \\ \mathcal{C}^{\text{Re Re}}(\Psi_j, \Phi_j) &= \text{Re} \left\{ \frac{(\nabla^2 + 2i\mathbf{k}_j \cdot \nabla)[\mathcal{C}^{\text{Re}}(\Psi_j, \Phi_j) e^{i\Phi_j}]}{e^{i\Phi_j}} \right\}, \\ \mathcal{C}^{\text{Im Re}}(\Psi_j, \Phi_j) &= \text{Im} \left\{ \frac{(\nabla^2 + 2i\mathbf{k}_j \cdot \nabla)[\mathcal{C}^{\text{Re}}(\Psi_j, \Phi_j) e^{i\Phi_j}]}{e^{i\Phi_j}} \right\}, \end{aligned} \quad (11)$$

and so on for the remaining \mathcal{C} 's. From here on we refer to the evolution equations for Ψ_j and Φ_j as the phase/amplitude




FIG. 1. (Color online) Evolution of a polycrystalline film simulated with the complex amplitude equations, Eq. (7), on an adaptive grid. Note that the grid does not coarsen inside many of the grains (misoriented with respect to \mathbf{k}_j) because of the fine scale structure of the beats. $t =$ (a) 0, (b) 88, (c) 168, (d) 248, (e) 320, and (f) 552.

equations, whereas Eq. (7) will be referred to as the complex amplitude equation. Unfortunately, the phase/amplitude equations in Eqs. (9) and (10) turn out to be quite difficult to solve globally. The principal difficulties are summarized below.

The field Ψ_j is nearly constant within the individual grains and varies sharply only near grain boundaries, rendering its equation ideally suited for solution on adaptive meshes. The field Φ_j , on the other hand, if computed naively as $\arctan[\text{Im}(A_j)/\text{Re}(A_j)]$, is a periodic and discontinuous function¹ bounded between the values $-\pi$ and π , with a frequency that increases with increasing grain misorientation. This poses a problem similar to that previously posed by the beats, with the grid this time having to resolve the fine-scale structure of Φ_j . Further, one may need to resort to shock-capturing methods in order to correctly evaluate higher-order derivatives, and resolve jumps where Φ_j

¹ $\Phi_j \approx \mathbf{q}(\theta) \cdot \mathbf{x}$, $\Phi_j \in [-\pi, \pi]$, where $\mathbf{q}(\theta)$ is the phase vector, constant for a particular orientation of the grain, and θ is the misorientation angle of the grain. Thus Φ_j , roughly speaking, has the structure of a sawtooth wave form.




FIG. 2. (Color online) (a), (b) Contours of Ψ_1 and Φ_1 respectively. (c), (d) Line-plots of Ψ_1 and Φ_1 along solid line in (a) and dashed line in (b) respectively. Field Ψ_1 is smooth everywhere except near interfaces and at defects. Φ_1 , computed naively as $\arctan(\text{Im}(A_1)/\text{Re}(A_1))$, is periodic and discontinuous. The chaotic fluctuations in Φ_1 regions outside the crystals correspond to the liquid phase where Φ_1 has no physical meaning. The rapid, but periodic, variations of Φ_1 in the left grain are due to its large misorientation angle of $\pi/6$. In contrast, the grain on the right is oriented along \mathbf{k}_j , causing Φ_1 to vary much more smoothly.

changes value from π to $-\pi$ and vice versa. Complications are also caused by Φ_j being undefined in the liquid phase, and the tendency for Ψ_j , which appears in the denominator on the right-hand side of Eq. (10), to approach zero at those locations. This calls for some type of robust regularization scheme² for the phase equations. These problems are clearly highlighted in Fig. 2, which shows the impingement of two misaligned crystals and the corresponding values of Ψ_1 and Φ_1 .

Ideally, one would like to reconstruct from the periodic Φ_j a continuous surface $\Phi_j + 2n\pi$ (where n is an integer), which would be devoid of jumps and therefore amenable to straightforward resolution on adaptive meshes. The implementation of such a reconstruction algorithm, however, even if possible, requires information about individual crystal orientations, and the precise location of solid-liquid interfaces, defects, and grain boundaries at every time step, making it very computationally intensive. Further, such an algorithm would be more appropriate in the framework of an interface-tracking approach such as the level set method [41], rather than our phase-field modeling approach.

²We have determined that simple tricks such as setting Ψ_j to some small nonzero value, or setting a heuristic upper bound on higher-order derivatives, have the effect of destroying defects and other topological features in the pattern.

Despite these issues with the polar (phase and amplitude) equations progress can be made, under certain noncritical approximations, by solving the phase and amplitude equations in the interior of crystalline regions, in conjunction with the Cartesian complex amplitude equations in regions closer to domain boundaries and topological defects.

D. Reduced equations and the frozen phase gradient approximation

The main idea that will be developed in this and subsequent sections is that of evolving the phase and amplitude and complex amplitude equations simultaneously in different parts of the domain, depending on where they can most appropriately be applied. The phase and amplitude formulation is solved in the crystal interior, away from defects, interfacial regions, and the liquid phase. The complex amplitude equations are solved everywhere else in the computational domain. This does away with the need for regularizing the phase equations where $\Psi_j \rightarrow 0$ (since $\Psi_j \gg 0$ in the crystal interior) as well as the issue of the phase being undefined in certain regions. We overcome the remaining issues with the phase equation, i.e., the difficulty of evaluating derivatives of the phase and the need to resolve its periodic variations via certain controlled approximations described next.

Let us examine the results of a fixed grid calculation performed using the complex amplitude equations, illustrated in Fig. 3, showing a sequence of line plots of the quantity $\Delta(\partial\Phi_1/\partial x) = \partial\Phi_1/\partial x|_{x=840} - \partial\Phi_1/\partial x|_r$. The quantity $\nabla\Phi_j$ inside the growing crystal is seen to be essentially time invariant. As the crystal on the left grows, it can be seen that $\Delta(\partial\Phi_1/\partial x)$ stays close to zero inside. We have verified that this is also true for the y component of $\nabla\Phi_1$, and both components of $\nabla\Phi_2$ and $\nabla\Phi_3$.

These results suggest that we may employ a *locally frozen phase gradient*. Note that the assumption of a frozen phase gradient does not mean that Φ_j itself cannot change. Φ_j can continue to evolve as per Eq. (13) under the constraint of a fixed $\nabla\Phi_j$, although the changes may actually be quite small. On the other hand, when similarly oriented crystals collide to form a small-angle grain boundary, it is energetically more favorable for grains to locally realign (i.e., for $\nabla\Phi_j$ to change close to grain boundaries) in order to reduce orientational mismatch [42–45], rather than to nucleate dislocations. Since such interaction effects originate at the grain boundary, where the full complex RG equations will be solved, we anticipate that our assumption will not lead to artificially “stiff” grains.

This approximation allows us to neglect third and higher order derivatives of Ψ_j and Φ_j ,³ which allows us to reduce Eqs. (9) and (10) to the following second-order PDEs:

$$\begin{aligned} \frac{\partial\Psi_j}{\partial t} &= (r + 3\bar{\psi}^2)[- \Psi_j + \mathcal{C}^{\text{Re}}(\Psi_j, \Phi_j)] - 3\Psi_j \left(\Psi_j^2 + 2 \sum_{k \neq j} \Psi_k^2 \right) \\ &\quad - 6 \frac{\bar{\psi}}{\Psi_j} \left(\prod_k \Psi_k \right) \cos \left(\sum_k \Phi_k \right), \end{aligned} \quad (12)$$

³To consistent order, we can also neglect second-order derivatives of Ψ_j .




FIG. 3. (Color online) Φ_1 and its time evolution for the pair of crystals shown in Fig. 2. (a) Contours of $\partial\Phi_1/\partial x$ at $t=520$. (b) Line plot of $\partial\Phi_1/\partial x$ along solid line in (a). (c) $\Delta(\partial\Phi_1/\partial x)$ along dashed line in (a). Just as with Ψ_j , the components of $\nabla\Phi_j$ are also practically constant inside the individual crystals. The spike in (b) corresponds to a defect on the grain boundary. As seen from the time series in (c) for $\partial\Phi_1/\partial x$, $\nabla\Phi_j$ hardly changes in the crystal bulk during its evolution.

$$\frac{\partial\Phi_j}{\partial t} = \frac{(r + 3\bar{\psi}^2)\mathcal{C}^{\text{Im}}(\Psi_j, \Phi_j)}{\Psi_j} + 6\frac{\bar{\psi}}{\Psi_j^2} \left(\prod_k \Psi_k \right) \sin\left(\sum_k \Phi_k \right), \quad (13)$$

where \mathcal{C}^{Re} and \mathcal{C}^{Im} contain only first- and second-order derivatives of Ψ_j and Φ_j . Equations (12) and (13) are referred to as the reduced phase and amplitude equations.

The task of evolving the phase and amplitude equations is now considerably simplified, as only derivatives up to sec-




FIG. 4. Sketch illustrating the idea of selectively evolving the complex amplitude and phase and amplitude equations in different regions of the computational domain. Ψ_j and Φ_j are evolved inside the shaded circles, which fall well inside the crystalline phase, while the real and imaginary components of A_j are evolved everywhere else.

ond order in Φ_j need to be computed. While the Laplacian and gradient of Ψ_j can be computed in a straightforward manner using Eqs. (B1) and (B5), respectively, the gradient of Φ_j needs to be computed with a little more care (in order to avoid performing derivative operations on a discontinuous function). The result is that

$$\nabla\Phi_j = \frac{\text{Re}(A_j) \nabla \text{Im}(A_j) - \text{Im}(A_j) \nabla \text{Re}(A_j)}{\Psi_j^2}. \quad (14)$$

Thus, the gradient operation on a discontinuous function Φ_j is now transformed into gradient operations on the smooth components of the complex amplitude A_j . Further, $\nabla^2\Phi_j$ is computed as $\nabla \cdot \nabla\Phi_j$, where the divergence operator is discretized using a simple second-order central difference scheme.

However, as can be seen from Eq. (14), $\nabla\Phi_j$ now depends on the gradients of the real and imaginary components of A_j , which may not be properly resolved in the crystal bulk as we intend to coarsen the mesh there. To address this point, we assume that $\nabla\Phi_j$ is frozen temporally in the crystal bulk. This assumption implies that, once $\nabla\Phi_j$ is accurately initialized in the crystal interior via Eq. (14), after ensuring adequate resolution of the components of A_j , it need not be computed again. For example, in simulations of crystal growth from seeds, we can start with a mesh that is initially completely refined inside the seeds, so that $\nabla\Phi_j$ is correctly computed. Once initial transients disappear and the crystals reach steady state evolution, the growth is monotonic in the outward direction. From this point on, $\nabla\Phi_j$ hardly changes inside the crystal bulk and the grid can be unrefined inside the grains while correctly preserving gradients in $\nabla\Phi_j$. Note that the apparent discontinuities in Φ_j no longer need be resolved by the grid.

III. A HYBRID FORMULATION

In order to implement our idea of evolving Eq. (7) and Eqs. (12) and (13) selectively within different regions, we begin by dividing the computational domain into two regions

TABLE I. Domain decomposition algorithm. The parameters γ , ϵ_1 , and ϵ_2 are heuristic.

```

Compute  $\Psi_j^{max}$ 
 $\Psi_j^{max} = \gamma \times \Psi_j^{max}$ 
{Split domain based on the magnitude of  $\Psi_j$  and  $|\nabla\Psi_j|$ }
FOR i=1 to maxnode DO {loop over all nodes}
  count=0
  FOR j=1 to 3 DO {loop over amplitude components}
    if  $\Psi_j \geq \Psi_j^{max}$  and  $|\nabla\Psi_j| \leq \epsilon_1$  THEN
      count++
    END IF
  END FOR
  IF count=3 THEN
    domain=P {passed test, solve phase and amplitude
equations}
  ELSE
    domain=X {failed test, solve complex equations}
  END IF
END FOR
{Split domain based on  $|\nabla(|\nabla\Phi_j|)|$ }
FOR i=1 to maxnode DO {loop over all nodes}
  count=0
  IF domain=P THEN {check only nodes that passed previous
test}
    FOR j=1 to 3 DO {loop over amplitude components}
      IF  $|\nabla(|\nabla\Phi_j|)| \geq \epsilon_2$  THEN
        count++
      END IF
    END FOR
    IF count  $\neq$  3 THEN
      domain=X {failed test, solve complex equations}
    END IF
  END IF
END FOR

```

where each set of equations may be evolved simultaneously in a stable fashion. The region where A_j is computed in terms of its real and imaginary parts is called X, and the region where Ψ_j and Φ_j are computed is called P. We ensure that subdomain P is well separated from locations with sharp gradients, such as interfaces and defects. Otherwise, errors resulting from our approximations may grow rapidly, causing X to invade P, which will in turn require us to solve the complex equations everywhere. We will further assume that the decomposition algorithm is implemented after a sufficient time, when initial transients have passed, and that the crystals are evolving steadily, which implies that Ψ_j inside the crystals has reached some maximum saturation value Ψ_j^{max} . The scenario we have in mind is sketched in Fig. 4, with P constituting the shaded regions and all other regions corresponding to X.

The pseudocode shown in the algorithm in Table I presents a simple algorithm to achieve this decomposition. The algorithm first determines nodes with Ψ_j exceeding some




FIG. 5. (Color online) Filled contour plot showing the time evolution of three misoriented crystals. The field plotted is Ψ_3 . Superimposed on the plot as solid curves are the boundaries that separate domains X and P, with P being enclosed by the curves. $t=(a)$ 120, (b) 200, (c) 280, and (d) 360.

minimum value $\gamma\Psi_j^{max}$, and $|\nabla\Psi_j|$ beneath some limit ϵ_1 . The nodes satisfying these conditions constitute domain P, while those failing to constitute X. The P nodes are then checked again to see if the quantity $|\nabla(|\nabla\Phi_j|)|$ is under some limit ϵ_2 . Nodes in set P that fail to satisfy this condition are placed in set X. The parameters γ , ϵ_1 , and ϵ_2 are chosen to ensure the largest possible size of set P. A small problem is caused by the fields Ψ_j and $|\nabla\Phi_j|$ not being perfectly monotonic. As the limits ϵ_1 and ϵ_2 are sharp, several small islands (clusters of grid points) of X or P can be produced, which are detrimental to numerical stability. We have resolved this issue via a coarsening algorithm that eliminates very small clusters of X and P by locally averaging the two fields in a recursive fashion.

Figure 5 shows results from a uniform grid implementation of the algorithm in Table I. No islands are present, as the algorithm decomposes the domain in an unsupervised manner. It is noteworthy that domain boundaries are distorted in Figs. 5(c) and 5(d) in response to the formation of a grain boundary between the two crystals, after being roughly hexagonal at earlier times. The fact that the domain separatrices maintain a safe distance from the grain boundary ensures that the phase and amplitude equations are not evolved in regions containing sharp gradients in $\nabla\Phi_j$. Parameter values used were $\gamma=0.85$, $\epsilon_1=0.0005$, and $\epsilon_2=0.003$.

The remarkable feature of our numerical scheme is that solving different sets of equations in X and P does not require doing anything special near the domain boundaries, such as creating “ghost” nodes outside each domain, or constraining solutions to match at the boundaries. Both sets of variables $\{\Psi_j, \Phi_j\}$ and $\{A_j\}$ are maintained at all grid points irrespective of the domain they belong to, with one set al-




FIG. 6. (Color online) Evolution of a polycrystalline film simulated with Eq. (7), and Eqs. (12) and (13), using our adaptive mesh refinement algorithm. The conditions in this simulation are identical to those in Fig. 1 of Sec. II B. Note that the grid now coarsens inside grains that are misoriented with respect to \mathbf{k}_j , and beats are no longer a limitation. $t=(a) 0$, (b) 88, (c) 168, (d) 248, (e) 320, and (f) 552.

lowing easy computation of the other.⁴ Therefore the transition between the two domains is a continuous one in terms of field variables, which allows the finite-difference stencils in Eqs. (B1) and (B5) to be applied to the respective fields without any modification near domain boundaries.

IV. RESULTS AND COMPUTATIONAL EFFICIENCY

Using the scheme described in Sec. III we can now solve the phase and amplitude and complex equations simultaneously in different parts of our computational domain using adaptive mesh refinement. The interested reader is referred to the supplemental material [46] for both conceptual and implementation-related development of adaptive mesh refinement and associated algorithms.

We simulated the same problem (same initial and boundary conditions and problem parameters) that was solved adaptively in Sec. II B using only the complex amplitude equations. Figure 6 shows the crystal boundaries and grid structure at various times during the simulation. The transient time N_{tr} was chosen to be 3000 for this simulation. With $\Delta t=0.04$, this implies that this simulation is identical to the previous one until $t=N_{tr}\times\Delta t=120$. Thus, Figs. 6(a) and 6(b) are identical to Figs. 1(a) and 1(b). The advantage of the hybrid implementation starts to appear from Fig. 6(c), whenceforth, unlike in Fig. 1, even grains that are misoriented with respect to the basis \mathbf{k}_j show grid unrefinement


within. It is also noteworthy that the grid remains refined near solid-liquid interfaces, grain boundaries, and defects, ensuring that key topological features are correctly resolved.

We now compare solutions from the two simulations quantitatively. We find it more informative to make a point-wise comparison of the two solutions along cross sections of the domain, rather than comparing solution norms, as we believe that this is a more stringent test of our implementation. We choose two random cuts, one running parallel to the y axis at $x_{cut}=70\pi$, and the other parallel to the x axis at $y_{cut}=118\pi$. The solutions are compared along these cuts at two different times, $t=168$ and $t=552$ in Figs. 7 and 8, respectively. The solid curves in the figures (labeled “hybrid”) are variations in Ψ_1 and $\partial\Phi_1/\partial x$ along the entire length of the domain as computed with the current (“hybrid”) implementation, whereas the symbols (labeled “complex”) are variations in the same variables as computed using fully complex equations (Sec. II B). The agreement is essentially perfect, indicating that our simplifications based on approximations in the preceding sections work reasonably well.


Because the performance of our algorithm is sensitively tied to the type of problem that is being solved, it is difficult to come up with a simple metric that quantifies its computational efficiency. The difficulty lies in accounting for the change in CPU time per time step, which increases with the number of mesh points. For example, Fig. 9 shows the number of nodes in this simulation over time. Clearly, an adaptive grid implementation has a significant computational advantage over an equivalent fixed grid implementation at the early stages of the simulation.

One performance measure is the projected speed of our implementation compared to a uniform grid implementation

⁴For example, in domain X where $\{A_j\}$ is the field variable, $\Psi_j=|A_j|$ and $\Phi_j=\text{Im}(A_j)/\text{Re}(A_j)$, whereas in domain P where $\{\Psi_j, \Phi_j\}$ are the field variables, $\text{Re}(A_j)=\Psi_j \cos(\Phi_j)$ and $\text{Im}(A_j)=\Psi_j \sin(\Phi_j)$.



(a)




(b)

FIG. 7. Numerical solution along the line $x=70\pi$ in Fig. 6 compared to the results using the hybrid scheme. Some of the data points in the complex solution were omitted for clarity of presentation. $t=(a)$ 168 and (b) 552.


of the PFC equation. This speedup is estimated by the simple formula

$$S = \frac{N_{PFC}}{N_{RG-AG}} \times \frac{\Delta t_{RG-AG}}{\Delta t_{PFC}} \times \frac{1}{6} \times \beta, \quad (15)$$

where N_{PFC} is the number of grid points required to solve the PFC equation, N_{RG-AG} is the number of grid points required in a hybrid implementation of the amplitude and RG equations, Δt_{PFC} and Δt_{RG-AG} are the time steps used in the respective implementations, the factor $1/6$ comes from solving six RG equations in place of the (one) PFC equation directly, and $\beta \in [0, 1]$ is the overhead of the AMR algorithm. The difficulty lies in fixing N_{RG-AG} , which is constantly changing with time. One estimate for N_{RG-AG} is the number of nodes averaged over the entire simulation. This can be computed easily by dividing the area under the hybrid curve in Fig. 9 by the total number of time steps taken, which gives



(a)



(b)

FIG. 8. Numerical solution along the line $y=118\pi$ in Fig. 6 compared to the results using the hybrid scheme. Some of the data points in the complex solution were omitted for clarity of presentation. $t=(a)$ 168 and (b) 552.

$N_{RG-AG}=104\,747$. Further, based on heuristics collected while running our code, we conservatively estimate mesh refinement and coarsening to constitute about 3% of the CPU time, which gives $\beta=0.97$. Therefore, from Eq. (15) we have

$$S = \frac{1\,050\,625}{104\,747} \times \frac{0.04}{0.008} \times \frac{1}{6} \times 0.97 = 8.1. \quad (16)$$

We do recognize that for a more accurate estimate of S we would also need to consider overhead costs that may come from suboptimal cache and memory usage owing to the data structures used. Hence these numbers should only be considered as rough estimates of true speedup.

While a speedup factor of 8 may not seem to be a great improvement in computational efficiency, one should bear in mind that the number of nodes in the AMR algorithm scales (roughly) linearly with interface or grain boundary length, which is quite substantial in the system we just simulated. Thus, one should not expect to derive the maximum compu-




FIG. 9. Number of computational nodes in the grid as a function of time, for simulations in Figs. 1 and 6. The number of nodes reaches a constant value after all the liquid freezes. The number of nodes required by a uniform grid implementation of the complex amplitude equations for the same problem is also shown for comparison.

tational benefit when simulating small systems with large numbers of grains. On the other hand, with this method, we can now simulate the growth of a few crystals in a much larger system. We choose a square domain of side 4096π , which in physical dimensions translates to $0.722 \mu\text{m}$, if we assume an interatomic spacing of 4 \AA .⁵ We initiate three randomly oriented crystals, two a little closer together than the third, so that a grain boundary forms quickly. The crystals are shown at different times in Fig. 10. The simulation was terminated at $t=3960$ when memory requirements exceeded 1 Gbyte, after running on a dedicated 3.06 GHz Intel Xeon processor for about one week.

Let us calculate the speedup factor for this simulation as we did previously, after 70 000 time steps [$t=2800$, Fig. 10(f)]. Figure 11 shows that the number of nodes in the adaptive grid varies nearly linearly with the number of time steps, and we estimate the average number of nodes N_{AG-RG} to be 200 721. The same simulation on a uniform grid using the PFC equation would have required 268 435 456 nodes (not possible on our computers). We estimate $\beta=0.98$. In this case the speedup is about three orders of magnitude,

$$S = \frac{268\,435\,456}{200\,721} \times \frac{0.04}{0.008} \times \frac{1}{6} \times 0.98 = 1091. \quad (17)$$

Figure 12 shows vividly the range of length scales from nanometers to micrometers spanned by our grid in this simulation, highlighting its “multiscale” capability.

⁵This is the interatomic spacing in aluminum [50], which has a face-centered cubic lattice.

As there can be considerable variability in the speedup factor depending on the number of crystals and system size, we now propose a simple expression that relates speedup and these quantities. Assuming that the advantage of using larger time steps in an explicit scheme with the RG method offsets the disadvantage of solving six PDEs, the speedup is simply

$$S \approx \frac{N_{PFC}}{N_{RG-AG}}. \quad (18)$$

While $N_{PFC} \propto l^d$ in d dimensions, where l is the characteristic size of the system, $N_{RG-AG} \propto l_{int}$ [47], where l_{int} denotes the solid-liquid interface length. We can conservatively estimate l_{int} by assuming that the crystals grow, roughly, as circles (spheres in 3D), so that $l_{int} \propto n_{cryst} R_{cryst}$, where n_{cryst} and R_{cryst} are the number of crystals and the average crystal size, respectively. If the crystals occupy the entire volume (this would correspond to the largest interface length), $n_{cryst} R_{cryst}^d \propto l^d$, and therefore $R_{cryst} \propto n_{cryst}^{-1/d} l$. This implies that

$$S \propto \frac{l^{d-1}}{n_{cryst}^{(d-1)/d}}. \quad (19)$$

Clearly, the RG method, with AMR, would demonstrate a significant computational advantage when simulating large systems. The constant of proportionality in the above equation would depend on the algorithmic details of the specific AMR implementation.

V. CONCLUDING REMARKS

In this paper, we have presented an efficient hybrid numerical implementation that combines Cartesian and polar representations of the complex amplitude with adaptive mesh refinement, and allows the modeling capabilities of the PFC equation to be extended to microscopic length scales. Depending on the choice of application, we have shown that our scheme can be anywhere from one to three orders of magnitude times faster than an equivalent uniform grid implementation of the PFC equation, on a single processor machine, and, more generally, the speedup scales as $l^{d-1}/n_{cryst}^{(d-1)/d}$. We anticipate that this advantage will be preserved when both implementations are migrated to a parallel computer, which is an important next step required to give the RG extension of the PFC model full access to micro- and meso-scale phenomena.

In conclusion, we have shown that multiscale modeling of complex polycrystalline material microstructure is possible using a combination of continuum modeling at the nanoscale using the PFC model, RG and related techniques from spatially extended dynamical systems theory, and adaptive mesh refinement.

We regard this work as only a first step for our modeling approach with the RG extension of the PFC to be successfully applied for studying important engineering and materials science applications. We have identified a few issues that require immediate attention. The first, although an implementation issue, is critical, and has to do with using amplitude equations for applications involving externally applied loads and displacements to a polycrystal that has been

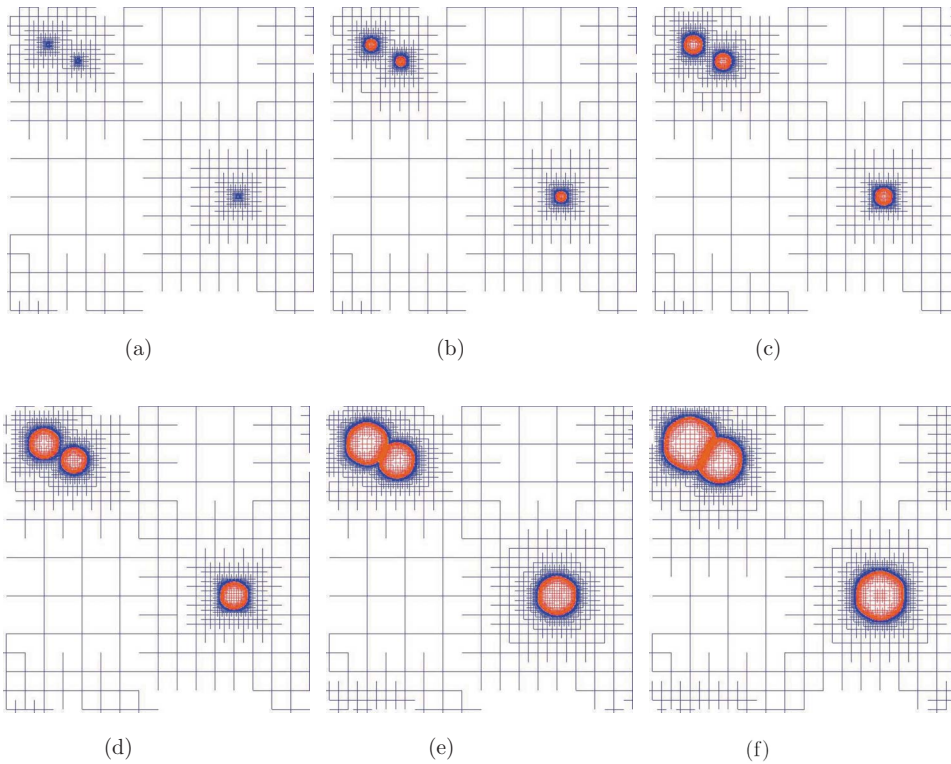


FIG. 10. (Color online) Micro-scale simulation of two-dimensional crystal growth with amplitude equations using the hybrid scheme and AMR. (a) $t=0$, (b) 840, (c) 1120, (d) 1620, (e) 2080, and (f) 2800.

evolved with our equations. Simple applications could be subjecting the polycrystal to shear, uniaxial, or biaxial loading states [32,33]. Such boundary conditions are difficult enough to apply to the scalar field ψ in the PFC equation [34]. Meaningful translation to equivalent boundary conditions on the amplitudes and phases of ψ can be a very difficult task, requiring the solution of systems of nonlinearly coupled equations at the boundaries. We have not yet investigated this issue in any detail.

Our derivation of the amplitude equations [40] was based on a one-mode approximation to the triangular lattice, and as we always chose parameters fairly close to the boundary between the triangular phase and coexisting triangular and constant phases, i.e., $|r+3\bar{\psi}^2| \ll 1$, the amplitude equations we derived were within their domain of validity and our results were quite accurate. It is almost certain that a one-mode approximation will not give similarly accurate results when $|r+3\bar{\psi}^2| \sim O(1)$ (although it would be interesting to see how large the error actually is). It is not clear if this in any way precludes certain phenomena from being studied with our equations, as we can always choose parameters to stay in the regime where the one-mode approximation is valid, but if it does, amplitude equations for dominant higher modes need to be systematically developed.

An important assumption made in the derivation of our so-called “hybrid” formulation of the complex amplitude equations is that of locally freezing the phase gradient vector $\nabla\Phi_j$. In fact, it is this assumption that allows us to effectively unrefine the interior of grains and gain significant speedup over the PFC equation. If for example, the problem we are studying involves the application of a large external shear strain that could change $\nabla\Phi_j$ in the grain interior via grain rotation, it is uncertain whether our algorithm would

continue to maintain its computational efficiency over the PFC. This is again a matter worth investigating.

ACKNOWLEDGMENTS

We thank Ken Elder and Nicholas Guttenberg for several useful discussions. This work was partially supported by the National Science Foundation through Grant No. NSF-DMR-01-21695. One of the authors (N.P.) wishes to acknowledge




FIG. 11. Number of computational nodes in the grid as a function of time for the $1 \times 1 \mu\text{m}^2$ domain. The growth is almost linear.




FIG. 12. (Color online) The above grid spans roughly three orders of magnitude in length scales, from a nanometer up to a micrometer. The leftmost box resolves the entire computational domain whereas the rightmost resolves dislocations at the atomic scale.

support from the National Science and Engineering Research Council of Canada.

APPENDIX A: THE PROBLEM OF BEATS

Consider a density field ψ defined by Eq. (5) with triangular lattice basis vectors $\mathbf{k}_j(\theta)$ {where $|\mathbf{k}_j(\theta)|=1$ } that are rotated by an angle θ from the basis vectors \mathbf{k}_j in Eq. (6), i.e.,

$$\psi(\theta) = \sum_{j=1}^3 A_j e^{i\mathbf{k}_j(\theta)\cdot\mathbf{x}} + \sum_{j=1}^3 A_j^* e^{-i\mathbf{k}_j(\theta)\cdot\mathbf{x}} + \bar{\psi}. \quad (\text{A1})$$

Equation (A1) describes the density field of a grain misoriented with respect to the basis vectors. Writing the basis vectors as $\mathbf{k}_j(\theta) = \mathbf{k}_j + \delta\mathbf{k}_j(\theta)$, where the vector $\delta\mathbf{k}_j(\theta)$ measures the rotation of each lattice vector, we obtain

$$\psi(\theta) = \sum_{j=1}^3 A_j e^{i\delta\mathbf{k}_j(\theta)\cdot\mathbf{x}} e^{i\mathbf{k}_j\cdot\mathbf{x}} + \sum_{j=1}^3 A_j^* e^{-i\delta\mathbf{k}_j(\theta)\cdot\mathbf{x}} e^{-i\mathbf{k}_j\cdot\mathbf{x}} + \bar{\psi} \quad (\text{A2})$$

or

$$\psi(\theta) = \sum_{j=1}^3 A_j^\theta e^{i\mathbf{k}_j\cdot\mathbf{x}} + \sum_{j=1}^3 A_j^{\theta*} e^{-i\mathbf{k}_j\cdot\mathbf{x}} + \bar{\psi}, \quad (\text{A3})$$

where

$$A_j^\theta = A_j e^{i\delta\mathbf{k}_j(\theta)\cdot\mathbf{x}}. \quad (\text{A4})$$

Thus grains arbitrarily misoriented from the global basis \mathbf{k}_j can still be described in terms of \mathbf{k}_j by suitably representing the complex amplitude A_j in polar form according to Eq. (A4). A straightforward way to include differently oriented grains in the system is to specify an initial condition via Eq. (A3). By making the amplitude a nonuniform complex function with a periodic structure, multiple grain orientations are automatically included. Figure 13 illustrates this idea. Figure




FIG. 13. (Color online) (a) Real component of the complex amplitude A_1 . As the grain in the bottom left corner is aligned with the basis \mathbf{k}_j in Eq. (6) its amplitude is constant, while amplitudes of the remaining misoriented grains have beats. (b) Density field ψ reconstructed using Eq. (A3). Clockwise from the lower left corner, $\theta=0, \pi/24$, and $\pi/6$.

13(a) shows the real component of one of the three complex amplitude functions A_j , specified by Eq. (A4), and Fig. 13(b) shows the corresponding density field constructed using Eq. (A3). Since Eq. (7) is rotationally covariant, it allows these beat structures in the amplitudes (and therefore the corresponding orientation of the grain) to be preserved as the system evolves, thereby enabling the representation of polycrystalline systems with a single set of basis vectors.

APPENDIX B: DISCRETIZATION OF OPERATORS

1. Laplacian

The Laplacian of a function $f(x,y)$ is discretized at point $(x_i, y_j) = (i\Delta x, j\Delta x)$ using a nine-point finite-difference stencil as shown below, where Δx is the mesh spacing,

$$\begin{aligned} \nabla^2 f|_{i,j} = & \frac{f_{i+1,j} + f_{i-1,j} + f_{i,j+1} + f_{i,j-1}}{2\Delta x^2} \\ & + \frac{f_{i+1,j+1} + f_{i-1,j-1} + f_{i-1,j+1} + f_{i+1,j-1}}{4\Delta x^2} - \frac{3f_{i,j}}{\Delta x^2} \\ & + O(\Delta x^2). \end{aligned} \quad (\text{B1})$$

A Fourier transform of this *isotropic* discretization, described by Tomita in [48], is shown to very nearly follow the $-k^2$ isocontours.

2. Gradient

The gradient of a function $f(x,y)$ is discretized at point $(x_i, y_j) = (i\Delta x, j\Delta x)$ using a nine-point second-order finite-difference stencil as shown below, where Δx is the mesh spacing. The stencil is designed to minimize effects of grid

anisotropy which can introduce artifacts in the solution, especially on adaptive grids. We have

$$\begin{aligned} \nabla f|_{i,j} &= \tilde{\nabla}_{\oplus} f|_{i,j} + O(\Delta x^2) = \left(\frac{f_{i+1,j} - f_{i-1,j}}{2\Delta x} \right) \vec{i} \\ &+ \left(\frac{f_{i,j+1} - f_{i,j-1}}{2\Delta x} \right) \vec{j} + O(\Delta x^2). \end{aligned} \quad (\text{B2})$$

But

$$\nabla f = \left(\frac{f_x + f_y}{\sqrt{2}} \right) \left(\frac{\vec{i} + \vec{j}}{\sqrt{2}} \right) + \left(\frac{-f_x + f_y}{\sqrt{2}} \right) \left(\frac{-\vec{i} + \vec{j}}{\sqrt{2}} \right), \quad (\text{B3})$$

and hence we also have

$$\begin{aligned} \nabla f|_{i,j} &= \tilde{\nabla}_{\otimes} f|_{i,j} + O(\Delta x^2) = \left(\frac{f_{i+1,j+1} - f_{i-1,j-1}}{2\sqrt{2}\Delta x} \right) \left(\frac{\vec{i} + \vec{j}}{\sqrt{2}} \right) \\ &+ \left(\frac{f_{i-1,j+1} - f_{i+1,j-1}}{2\sqrt{2}\Delta x} \right) \left(\frac{-\vec{i} + \vec{j}}{\sqrt{2}} \right) + O(\Delta x^2) \\ &= \left(\frac{f_{i+1,j+1} - f_{i-1,j-1} - f_{i-1,j+1} + f_{i+1,j-1}}{4\Delta x} \right) \vec{i} \\ &+ \left(\frac{f_{i+1,j+1} - f_{i-1,j-1} + f_{i-1,j+1} - f_{i+1,j-1}}{4\Delta x} \right) \vec{j} + O(\Delta x^2). \end{aligned} \quad (\text{B4})$$

Using the discrete forms for the gradient in Eqs. (B2) and (B4) we can write the isotropic second-order discretization as

$$\nabla f|_{i,j} = \frac{1}{2} (\tilde{\nabla}_{\oplus} f|_{i,j} + \tilde{\nabla}_{\otimes} f|_{i,j}) + O(\Delta x^2). \quad (\text{B5})$$

A discretization scheme similar to Eq. (B5) is given by Set-
hian and Strain [49].

-
- [1] R. Phillips, *Crystals, Defects and Microstructures: Modeling across Scales* (Cambridge University Press, Cambridge, U.K., 2001).
- [2] E. B. Tadmor, M. Ortiz, and R. Phillips, *Philos. Mag. A* **73**, 1529 (1996).
- [3] V. B. Shenoy, R. Miller, E. B. Tadmor, R. Phillips, and M. Ortiz, *Phys. Rev. Lett.* **80**, 742 (1998).
- [4] J. Knap and M. Ortiz, *J. Mech. Phys. Solids* **49**, 1899 (2001).
- [5] R. E. Miller and E. B. Tadmor, *J. Comput.-Aided Mater. Des.* **9**, 203 (2002).
- [6] W. E, B. Enquist, and Z. Huang, *Phys. Rev. B* **67**, 092101 (2003).
- [7] W. E and Z. Huang, *Phys. Rev. Lett.* **87**, 135501 (2001).
- [8] R. E. Rudd and J. Q. Broughton, *Phys. Rev. B* **58**, R5893 (1998).
- [9] J. Q. Broughton, F. F. Abraham, N. Bernstein, and E. Kaxiras, *Phys. Rev. B* **60**, 2391 (1999).
- [10] C. Denniston and M. O. Robbins, *Phys. Rev. E* **69**, 021505 (2004).
- [11] S. Curtarolo and G. Ceder, *Phys. Rev. Lett.* **88**, 255504 (2002).
- [12] J. Fish and W. Chen, *Comput. Methods Appl. Mech. Eng.* **193**, 1693 (2004).
- [13] J. S. Langer, in *Directions in Condensed Matter Physics*, edited by G. Grinstein and G. Mazenko (World Scientific, Singapore, 1986), Vol. 1, p. 165.
- [14] A. Karma and W. J. Rappel, *Phys. Rev. E* **57**, 4323 (1998).
- [15] C. Beckermann, H.-J. Diepers, I. Steinbach, A. Karma, and X. Tong, *J. Comput. Phys.* **154**, 468 (1999).
- [16] J. A. Warren, R. Kobayashi, A. E. Lobkovsky, and W. C. Carter, *Acta Mater.* **51**, 6035 (2003).
- [17] D. D. Vvedensky, *J. Phys.: Condens. Matter* **16**, R1537 (2004).
- [18] N. Provatas, M. Greenwood, B. P. Athreya, N. Goldenfeld, and J. A. Dantzig, *Int. J. Mod. Phys. B* **19**, 4525 (2005).
- [19] N. Provatas, N. Goldenfeld, and J. Dantzig, *Phys. Rev. Lett.* **80**, 3308 (1998).
- [20] J. Jeong, N. Goldenfeld, and J. Dantzig, *Phys. Rev. E* **64**, 041602 (2001).
- [21] R. Kobayashi, J. A. Warren, and W. C. Carter, *Physica D* **119**, 415 (1998).
- [22] R. Kobayashi, J. A. Warren, and W. C. Carter, *Physica D* **140**, 141 (2000).
- [23] A. Onuki, *J. Phys. Soc. Jpn.* **58**, 3065 (1989).
- [24] A. Onuki, *J. Phys. Soc. Jpn.* **58**, 3069 (1989).
- [25] J. Muller and M. Grant, *Phys. Rev. Lett.* **82**, 1736 (1999).
- [26] K. Kassner, C. Misbah, J. Muller, J. Kappey, and P. Kohlert, *Phys. Rev. E* **63**, 036117 (2001).
- [27] A. Karma, D. A. Kessler, and H. Levine, *Phys. Rev. Lett.* **87**, 045501 (2001).
- [28] M. Haataja, J. Mahon, N. Provatas, and F. Léonard, *Appl. Phys. Lett.* **87**, 251901 (2005).
- [29] A. Karma, *Phys. Rev. Lett.* **87**, 115701 (2001).
- [30] B. Echebarria, R. Folch, A. Karma, and M. Plapp, *Phys. Rev. E* **70**, 061604 (2004).
- [31] K. R. Elder, M. Katakowski, M. Haataja, and M. Grant, *Phys. Rev. Lett.* **88**, 245701 (2002).
- [32] K. R. Elder and M. Grant, *Phys. Rev. E* **70**, 051605 (2004).
- [33] J. Berry, M. Grant, and K. R. Elder, *Phys. Rev. E* **73**, 031609 (2006).
- [34] P. Stefanovic, M. Haataja, and N. Provatas, *Phys. Rev. Lett.* **96**, 225504 (2006).
- [35] K. Elder, N. Provatas, J. Berry, P. Stefanovic, and M. Grant, *Phys. Rev. B* **75**, 064107 (2007).
- [36] N. Goldenfeld, B. P. Athreya, and J. A. Dantzig, *Phys. Rev. E* **72**, 020601(R) (2005).
- [37] N. Goldenfeld, B. P. Athreya, and J. A. Dantzig, *J. Stat. Phys.* **125**, 1015 (2006).
- [38] L. Y. Chen, N. Goldenfeld, and Y. Oono, *Phys. Rev. E* **54**, 376 (1996).

- [39] K. Nozaki and Y. Oono, *Phys. Rev. E* **63**, 046101 (2001).
- [40] B. P. Athreya, N. Goldenfeld, and J. A. Dantzig, *Phys. Rev. E* **74**, 011601 (2006).
- [41] J. A. Sethian, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 1591 (1996).
- [42] K. E. Harris, V. V. Singh, and A. H. King, *Acta Mater.* **46**, 2623 (1998).
- [43] D. Moldovan, V. Yamakov, D. Wolf, and S. R. Phillpot, *Phys. Rev. Lett.* **89**, 206101 (2002).
- [44] D. Moldovan, D. Wolf, S. R. Phillpot, and A. J. Haslam, *Acta Mater.* **50**, 3397 (2002).
- [45] D. Moldovan, D. Wolf, S. R. Phillpot, and A. J. Haslam, *Philos. Mag. A* **82**, 1271 (2002).
- [46] See EPAPS Document No. E-PLLEE8-76-165710 for supplemental material for both conceptual and implementation-related development of adaptive mesh refinement and associated algorithms. For more information on EPAPS, see <http://www.aip.org/pubservs/epaps/html>.
- [47] N. Provatas, J. Dantzig, and N. Goldenfeld, *J. Comput. Phys.* **148**, 265 (1999).
- [48] H. Tomita, *Prog. Theor. Phys.* **85**, 47 (1991).
- [49] J. A. Sethian and J. Strain, *J. Comput. Phys.* **98**, 231 (1992).
- [50] W. D. Callister, *Materials Science and Engineering* (Wiley, New York, 1997).

Adaptive mesh computation of polycrystalline pattern formation using a renormalization-group reduction of the phase-field crystal model - Supplementary Material

Badrinarayan P. Athreya¹, Nigel Goldenfeld², Jonathan A. Dantzig¹, Michael Greenwood³, and Nikolas Provatas³

¹*Department of Mechanical Science and Engineering,
University of Illinois at Urbana-Champaign,
1206 W. Green Street, Urbana, IL 61801, USA*

²*Department of Physics, University of Illinois at Urbana-Champaign,
1110 W. Green Street, Urbana, IL 61801, USA*

³*Department of Materials Science and Engineering, McMaster University,
1280 Main Street West, Hamilton, Ontario, L8S 4L7, Canada*

PACS numbers: 81.15.Aa, 81.16.Rf, 46.15.-x, 05.10.Cc

I. SOLVING THE RG EQUATIONS WITH ADAPTIVE MESH REFINEMENT

As highlighted in earlier work [1–4] the use of dynamic adaptive mesh refinement alters the numerical mesh resolution dynamically such as to place high resolution near phase boundaries and a very low resolution in bulk regions where there is little activity. This dramatically reduces computer memory requirements, allowing larger systems to be simulated. It also significantly reduces overall simulation times. The determining factor guiding the use or otherwise of an AMR technique to solve a particular problem is the simple criterion

$$\frac{\text{Interface length (or Area)}}{\text{Domain area (or Volume)}} \ll 1 \quad (1)$$

The phase/amplitude RG equations discussed in the main article [5] are precisely in the class of problems that can benefit from and is amenable to adaptive mesh refinement.

We solved the RG equations in [5] using a new C++ adaptive mesh refinement (AMR) algorithm that uses finite differences (FD) to resolve spatial gradients [6]. While it is typical to use the finite element method (FEM) in situations involving non-structured meshes, adaption using finite differences schemes allows approximately a 5-10 fold increase in simulation speed (measured as CPU time per node) over previous (FEM) formulations involving traditional phase field models [1, 2]. This improvement increases further still in cases where model equations contain spatial

gradients of order higher than two, such as in the case of Eqs. (7), (12) and (13) in [5]. The basic reason for the difference in speeds is that FEM formulations generally have more overhead due to their reliance on local matrix multiplication at multiple Gauss quadrature points. This overhead time becomes even more pronounced when using elements of order higher than two, as is required if an FEM formulation is to be used to resolve the spatial derivatives involved with the RG equations in this work.

A. AMR Algorithm

At the heart of our algorithm is a routine that creates a non-uniform mesh that increases nodal density in specific regions according to a local error estimator. Nodes are grouped into pseudo-elements, managed by dynamic tree data structures, as used in a finite element formulation by Provatas et al [2]. The quad-tree structure illustrated in Figure 1 is a hierarchy of elements where every level deeper in the tree results in elements of higher refinement. Every element has associated with it 4 corner ‘nodes’ and 5 ‘ghost’ nodes: 4 in the center of each edge and 1 in the center of the element. The ghost nodes facilitate interpolation when neighboring elements are at different levels of refinement. Each node is a structure containing field values, such as phase, amplitude modulus and the real and complex components of the complex amplitude A_j . A node also contains information about its nearest neighbors. Edge ghost nodes can serve as field nodes if a resolution mismatch occurs across neighbor elements. This can be seen in the schematic in figure 1. The field equations are not solved at ghost nodes, but are instead interpolated linearly from the nodal values of the element (or edge) to which they belong. The inclusion of ghosts nodes simplifies the calculation of local derivatives.

An adapter algorithm refines/unrefines the tree structure by using a user-defined error estimator computed for each element. Refinement is done by bisection as shown in Figure 1, and unrefinement is done by fusing four child elements into their parent element. Once refinement/unrefinement of all elements is complete, the adapter produces an array of nodes, each of which is the center of a local grid. The equation solver accepts this array as input and then solves the equations at these nodal points. This method of node organization modularizes our algorithm and allows the solver and adapter to be separately parallelized.

The mesh data structure contains nodes with detailed knowledge of their local neighbors, each of which exist at the center of a 5×5 uniform mesh (See Fig. 2). During adaptation, the data structure applies rules that either increase (if higher accuracy is needed) or decrease (to decrease

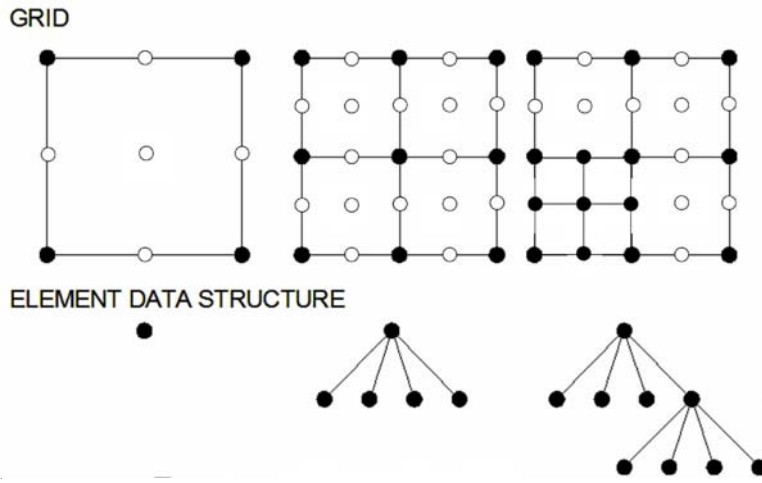


FIG. 1: Schematic of a quad-tree data structure used in adaptive meshing. An element splits and creates 'children' beneath it in the tree structure. Nodes are created at the corner of each element and 'ghost' nodes are placed in the center of elements and along edges that have no real nodes, to accommodate resolution mismatches. Ghost nodes are approximated by interpolation from the element to which they belong.

memory requirements) the size of this local nodal mesh. Also during adaptation, the data structure and its associated elements and nodes respect the following six rules (3 applied to each element and 3 to each node). Rule 1 ensures mesh cohesion and maintains accuracy in the solution of the PDEs.

RULE 1: Neighboring elements can vary by at most one level

RULE 2: All elements contain 9 nodes, real or ghost

RULE 3: Element neighbors are all at the same level (Therefore elements may have *NULL* as a neighbor)

RULE 4: Node neighbors are all at the same level (Nodes will never have *NULL* as a neighbor, but may instead have a ghost as a neighbor)

RULE 5: Each node is at the center of what is defined as a uniform mini-mesh

RULE 6: Each node is assigned the resolution (Δx) of the most refined element attached to it.

The adaptive process is controlled primarily at the tree level, but invokes function calls inside of the element and node structures. Its basic flow is illustrated in Algorithm 1. This process allows an element-by-element examination using recursion to maintain the rules above. Once the process

of adaptation is complete the adapter creates an array of nodes each with the index of its neighbors in the array, which is then used to solve the equations before adapting again. We also note that element ‘leaves’ are stored by their resolution level in an array of element lists. Element resolution is restricted to vary by at most one level compared to the element neighbors (Rule 1).

Algorithm 1 Regridding algorithm

```

Delete the Ghost Pointer List
while Some elements have split or unsplit do
  Check all elements for unsplitting criteria
  Check all elements for splitting criteria
end while
Update node neighbors
Build Ghost List
Set Ghost Averaging Information

```

Element splitting (refinement) is the dominant process in the algorithm, taking precedence over element coarsening (i.e. fusing four children elements into their parent). Elements are searched one refinement level at a time, starting from the second highest level of resolution. Each element is considered for splitting using an error criterion computed for that element. If splitting is required, an element data structure it is pushed onto a stack, where its neighbors are subsequently checked against Rule 1. If splitting will violate Rule 1, the neighbors are recursively split until all refined elements satisfy Rule 1. When an element is split, it and its updated neighbor elements generate new real and ghost nodes, as well as information about their neighbors. The splitting algorithm is illustrated in Algorithm 2.

The unrefinement algorithm starts at the lowest level of refinement. Again, Rule 1 above must be imposed. The elements are examined to see if a parent requires splitting. If it does not, the parent has its four child elements eliminated, assuming their relationship with the neighbors allows it (i.e. Rule 1). Recursive unsplitting of elements is not allowed. As in the case of refinement, during unrefinement, element neighbors and node neighbors are updated. The flow of unsplitting is shown in Algorithm 3.

As discussed above, nodes follow rules 4, 5 and 6. This is enforced by the introduction of ghost nodes in elements and by maintaining rule 1. Rule 5 is maintained by creating a list of local nodes (ghost or real) and by maintaining rule 1 during splitting. Rule 6 determines which node neighbors are chosen and the local node spacing (Δx). A real node will contain the minimesh on which is applied the governing equations; a ghost node provides instructions on how to interpolate the

Algorithm 2 Splitting Algorithm

```
for ElementLevel=maxLevel-1 to 0 do  
  while Level is not empty do  
    if Element Splitting criteria is met then  
      PUSH Element onto stack  
      while Stack is not empty do  
        if Top of Stack's neighbors need splitting then  
          Push Needed neighbors onto stack  
        else  
          Pop the Element from the Stack  
          Create 4 new children elements  
          Push children onto ElementLevel+1  
          Add New Nodes to the Node List  
          Remove former Ghosts from Ghost List  
        end if  
      end while  
    end if  
  end while  
end for
```

values needed for real node calculations.

B. Handling of Ghost Nodes in The Hybrid Formulation

A very useful feature of the AMR implementation outlined in the previous section is that each node sits at the center of a uniform 5×5 mini-grid, illustrated in Fig. 2. Let us focus on the node represented by the lightly shaded circle (labeled F). The data structure ensures that node F has access to all other nodes on the wireframe. The advantage of this construction is that it allows us to use the uniform grid finite difference stencils for the Laplacian and gradient operators in Eqs. (B1) and (B5) respectively in [5], instead of modifying them node-wise to accommodate variations in grid spacing. This requires the introduction of ghost nodes, shown as open circles in Figure 2.

The scheme used to interpolate values at the ghost nodes is a potential source of error in the numerical solution, and must be chosen carefully. Ψ_j , $\partial\Phi_j/\partial x$, $\partial\Phi_j/\partial y$, $\Re(A_j)$ and $\Im(A_j)$, where Ψ_j and Φ_j are the modulus and phase respectively of the complex amplitude A_j (see [5] for details) are very smoothly varying functions, and therefore we *linearly* interpolate their values to

Algorithm 3 UnSplitting Algorithm

```
for ElementLevel=maxLevel-1 to 0 do  
  while Level is not empty do  
    if Element Splitting criteria is NOT met then  
      if Element's Parent's Splitting criteria is NOT met then  
        if Parent's Children Have NO Children then  
          if Parent's neighbor's are Same level or One level Lower then  
            ADD Parent to ElementLevel+1  
            REMOVE Parent's Children from ElementLevel  
            REMOVE Unneeded nodes from node list  
            REMOVE Ghosts from Ghost list  
            Center Node Becomes a Ghost  
            Edge Nodes Become Ghosts depending on neighbor levels  
            DELETE Children  
          end if  
        end if  
      end if  
    end while  
  end for
```

the ghost nodes. Values at ghosts residing on element [9] edges, for example node M in Fig. 2, are obtained by averaging values of the two end nodes Q and S, whereas values at ghosts residing at the center of an element, node N for example, are obtained by averaging the values at the four corner nodes, P, Q, R, and S. We have found this interpolation scheme to be quite stable. We note however that, given the near-periodic variations in $\Re(A_j)$ and $\Im(A_j)$, especially in misoriented grains, higher order interpolation functions (such as cubic splines) could improve solution accuracy, while strongly enforcing continuity of fields across elements. This issue will be examined in future work.

The interpolation of the phase Φ_j at the ghost nodes is a little more delicate. Since Φ_j is a discontinuous function [5], a simple average of the values at the neighboring real nodes may not always give the correct answer, especially because the grid does not resolve the discontinuities. Even if it did, a simple average could lead to the wrong result. As an example, consider the two real nodes Q and S in Fig. 2, with values $\Phi_1^Q = \pi - \delta_1$ and $\Phi_1^S = -\pi + \delta_2$ where δ_1 and δ_2 are very small but positive real numbers, on either side of a discontinuity in Φ_1 . We wish to determine the




FIG. 2: Schematic showing a portion of the adaptive grid where the refinement level changes. Filled circles (and node F) are real nodes where the fields are computed, whereas the open circles are non-computational ghost nodes where the fields are interpolated.

value at the ghost node M that lies between Q and S. Although the values of Φ_1 at Q and S are essentially equivalent in phase space, differing in magnitude by approximately 2π , a simple average gives $\Phi_1^M = (\delta_2 - \delta_1)/2 \approx 0$, which is quite wrong.

In order to interpolate correctly we need to make use of $\nabla\Phi_j$. For example in the above case, the total change in the phase from Q to S is obtained by integrating the directional derivative of Φ_1 along the edge QS, i.e.

$$\Delta\Phi_1^{QS} = \int_Q^S \nabla\Phi_1 \cdot d\mathbf{r} = \int_{y=y^Q}^{y^S} \frac{\partial\Phi_1}{\partial y} dy. \quad (2)$$

Eq. (2) can be evaluated numerically, and the accuracy of the result depends on how well $\partial\Phi_1/\partial y$ is approximated. Consistent with our earlier assumptions, we approximate $\partial\Phi_1/\partial y$ as piecewise constant where

$$\frac{\partial\Phi_1}{\partial y} = \frac{1}{2} \left(\left. \frac{\partial\Phi_1}{\partial y} \right|_S + \left. \frac{\partial\Phi_1}{\partial y} \right|_Q \right) \quad (3)$$

which leads to

$$\Delta\Phi_1^{QS} = \frac{1}{2} \left(\left. \frac{\partial\Phi_1}{\partial y} \right|_S + \left. \frac{\partial\Phi_1}{\partial y} \right|_Q \right) (y^S - y^Q). \quad (4)$$

Since $\partial\Phi_1/\partial y$ is constant along the edge QS, Φ_1 must vary linearly along QS. Hence at node M,

$$\Phi_1^M = \Phi_1^S + \frac{1}{2}\Delta\Phi_1^{QS} : \Phi_1^M \in [-\pi, \pi]. \quad (5)$$

Interpolation of Φ_j at element center ghost nodes, such as N, is done in a similar manner by interpolating linearly from ghost nodes at the centers of opposite element edges. Once again, this scheme might be improved by choosing higher order polynomials to approximate $\nabla\Phi_j$ inside elements.

C. Refinement Criteria in the Hybrid Formulation

Traditionally, AMR algorithms rely on some kind of local error estimation procedure to provide a criterion for grid refinement. Zienkiewicz and Zhu [7] developed a simple scheme for finite element discretization of elliptic and parabolic PDEs by computing the error in the gradients of the fields using higher order interpolation functions. Berger and Olinger [8] on the other hand estimated the local truncation error of their finite difference discretization of hyperbolic PDEs via Richardson extrapolation. Depending on the equations being solved and the numerical methods being used, one scheme may be more effective than another. We use a very simple and computationally inexpensive refinement criterion that works nicely for the equations in [5], based purely on gradients in the various fields.

The outline of the algorithm used to decide whether or not to split an element is given in Algorithm 4. The algorithm initially computes absolute changes in the real and imaginary parts of the complex amplitude A_j , and the x and y components of the phase gradient $\nabla\Phi_j$ in each element. We use absolute differences in place of derivatives in order for the refinement criterion to be independent of element size. We note that prior to implementing this algorithm, the domain decomposition algorithm described in [5] (see Algorithm 1) needs to be called first in order to split the computational domain into subdomains X and P.

The process begins by examining an element flag to see if the element lies on the separatrix between X and P, or in k layers from the boundary, within the P subdomain. If so this element is split. This ensures that the fields are always resolved on the interface between X and P, and just within the boundary on the P side. The latter is required because of the higher order derivative operations that need to be performed while evolving the complex amplitude equations in X.

If the element does not split and belongs to X, the variations in the real and imaginary parts of A_j are checked to see if they exceed a certain bound ϵ_1 . If any one of them does, this element is split. If, on the other hand, the element belongs to P where the phase/amplitude equations are solved, variations in the x and y components of $\nabla\Phi_j$ are checked to see if they exceed another limit ϵ_2 . If they do, this element is split. If none of the above criteria are satisfied, the element is not

Algorithm 4 Criteria for element splitting

{N1, N2, N3 and N4 are the element nodes in clockwise manner}

for i = 1 to 3 **do** {loop over amplitude components}

{Change in real part of A_i over element}

$$DR_i = |N1 \rightarrow \Re(A_i) - N2 \rightarrow \Re(A_i)| + |N2 \rightarrow \Re(A_i) - N3 \rightarrow \Re(A_i)|$$
$$+ |N3 \rightarrow \Re(A_i) - N4 \rightarrow \Re(A_i)| + |N4 \rightarrow \Re(A_i) - N1 \rightarrow \Re(A_i)|$$

{Change in imaginary part of A_i over element}

$$DI_i = |N1 \rightarrow \Im(A_i) - N2 \rightarrow \Im(A_i)| + |N2 \rightarrow \Im(A_i) - N3 \rightarrow \Im(A_i)|$$
$$+ |N3 \rightarrow \Im(A_i) - N4 \rightarrow \Im(A_i)| + |N4 \rightarrow \Im(A_i) - N1 \rightarrow \Im(A_i)|$$

{Change in x component of $\nabla\Phi_i$ over element}

$$DG PX_i = |N1 \rightarrow \partial\Phi_1/\partial x - N2 \rightarrow \partial\Phi_1/\partial x| + |N2 \rightarrow \partial\Phi_1/\partial x - N3 \rightarrow \partial\Phi_1/\partial x|$$
$$+ |N3 \rightarrow \partial\Phi_1/\partial x - N4 \rightarrow \partial\Phi_1/\partial x| + |N4 \rightarrow \partial\Phi_1/\partial x - N1 \rightarrow \partial\Phi_1/\partial x|$$

{Change in y component of $\nabla\Phi_i$ over element}

$$DG PY_i = |N1 \rightarrow \partial\Phi_1/\partial y - N2 \rightarrow \partial\Phi_1/\partial y| + |N2 \rightarrow \partial\Phi_1/\partial y - N3 \rightarrow \partial\Phi_1/\partial y|$$
$$+ |N3 \rightarrow \partial\Phi_1/\partial y - N4 \rightarrow \partial\Phi_1/\partial y| + |N4 \rightarrow \partial\Phi_1/\partial y - N1 \rightarrow \partial\Phi_1/\partial y|$$

end for

if element on X/P boundary OR k layers inside P **then**

Split element and exit

else if element inside X **then**

count=0

for i = 1 to 3 **do** {loop over amplitude components}

if $DR_i \geq \epsilon_1$ OR $DI_i \geq \epsilon_1$ **then**

count++

end if

end for

if count $\neq 0$ **then**

Split element and exit

end if

else {element is inside P}

count = 0

for i = 1 to 3 **do** {loop over amplitude components}

if $DG PX_i \geq \epsilon_2$ OR $DG PY_i \geq \epsilon_2$ **then**

count++

end if

end for

if count $\neq 0$ **then**

Split element and exit

end if

end if

split and is placed in the list of elements to be checked for coarsening. Since refinement criteria are recursively applied to the quadtree, the finest elements are automatically placed around domain separatrices, solid/liquid interfaces, and defects.

Algorithm 5 shows the flow of control in the main routine. The complex amplitude equations in [5] (Eq. (7)) are initially evolved everywhere until time N_{tr} , when initial transients have dissipated, and the crystals evolve steadily outward. The domain is then split into subdomains X and P, following which the reduced phase/amplitude equations in [5] (Eqs. (12) and (13)), are evolved using a forward Euler time stepping scheme in subdomain P. The grid is refined after a predetermined number of time steps N_{adapt} , which is chosen heuristically. We note that the current implementation can handle only periodic boundary conditions. Work is currently underway to enable handling of more general boundary conditions.

-
- [1] N. Provatas, J. Dantzig, and N. Goldenfeld, *Phys. Rev. Lett.* **80**, 3308 (1998).
 - [2] N. Provatas, J. Dantzig, and N. Goldenfeld, *J. Comp. Phys.* **148**, 265 (1999).
 - [3] J. Jeong, J. A. Dantzig, and N. Goldenfeld, *Met. Trans. A* **34**, 459 (2003).
 - [4] J. Jeong, N. Goldenfeld, and J. Dantzig, *Phys. Rev. E* **64**, 041602:1 (2001).
 - [5] B. P. Athreya, N. Goldenfeld, J. A. Dantzig, M. Greenwood, and N. Provatas (2007), in review.
 - [6] M. H. J. Fan, M. Greenwood and N. Provatas, *Phys. Rev. E* **74**, 031602 (2006).
 - [7] O. C. Zienkiewicz and J. Z. Zhu, *Int. J. Num. Meth. Eng.* **24**, 337 (1987).
 - [8] M. J. Berger and J. E. Oliger, *J. Comp. Phys.* **53**, 484 (1984).
 - [9] Here, we define an element as a square with real corner nodes.

Algorithm 5 Flow of control

```
InitVar() {Initialize program variables and parameters}
InitGrid() {decide where to initially refine/coarsen based on initial condition}
UpdateGhostsFV() {interpolate  $A_j$ ,  $\Psi_j$ ,  $\Phi_j$  at ghost nodes}
ComputePhaseGradients(1) {compute  $\nabla\Phi_j$  everywhere}
UpdateGhostsPG() {interpolate  $\nabla\Phi_j$  at ghost nodes}
for  $i = 1$  to  $N_{tr} - 1$  do {evolve until initial transients subside}
if  $i \bmod N_{adapt} = 0$  OR  $i = 1$  then
  AdaptGrid() {uses Algorithm 4}
end if
  EvolveComplexAmp() {evolve Eq. (7) [5] everywhere}
  UpdateAllFields() {compute  $\Psi_j$  and  $\Phi_j$  from  $A_j$ }
  UpdateGhostsFV()
  ComputePhaseGradients(1)
  UpdateGhostsPG()
end for
  DivideDomain() {call Algorithm 1 [5] to split domain into X and P}
for  $i = N_{tr}$  to  $N_{end}$  do {evolve after transients subside}
if  $i \bmod N_{adapt} = 0$  then
  DivideDomain()
  AdaptGrid()
end if
  EvolveComplexAmp() {evolve Eq. (7) [5] in X}
  EvolvePhaseAmp() {evolve Eqs. (12) and (13) [5] in P}
  UpdateAllFields() {evaluate  $\Psi_j$  and  $\Phi_j$  in X, evaluate  $A_j$  in P}
  UpdateGhostsFV()
  ComputePhaseGradients(i) {compute  $\nabla\Phi_j$  in X only, frozen gradient approx.}
  UpdateGhostsPG()
end for
```
