



Coupling kinetic Monte-Carlo and continuum models with application to epitaxial growth

Tim P. Schulze ^{a,*}, Peter Smereka ^b, Weinan E ^c

^a Department of Mathematics, University of Tennessee, Knoxville, TN 37996-1300, USA

^b Department of Mathematics, University of Michigan, Ann Arbor, MI 48109-1109, USA

^c Department of Mathematics, Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, USA

Received 25 April 2002; received in revised form 17 December 2002; accepted 28 January 2003

Abstract

We present a hybrid method for simulating epitaxial growth that combines kinetic Monte-Carlo (KMC) simulations with the Burton–Cabrera–Frank model for crystal growth. This involves partitioning the computational domain into KMC regions and regions where we time-step a discretized diffusion equation. Computational speed and accuracy are discussed. We find that the method is significantly faster than KMC while accounting for stochastic fluctuations in a comparable way.

© 2003 Elsevier Science B.V. All rights reserved.

1. Introduction

This paper explores the relationship between two models of epitaxial film growth – kinetic Monte-Carlo (KMC) and Burton–Cabrera–Frank (BCF) models – and introduces a simulation technique that combines features of both in a hybrid scheme for film growth simulation. KMC is a stochastic model aimed at simulating film growth on an atom-by-atom basis using probabilistic rules to govern deposition, diffusion and other growth processes [1,2]. BCF models partition the discrete topography of the growing crystal into smooth planar surfaces (terraces) and unit-cell high step edges [3,4]. The steps themselves may be viewed as either continuous curves marking the boundary of film-surface level-sets (TS models) or as discrete curves with kinks that correspond to unit-cell width variations in the step boundaries (TSK models). Both TS and TSK models are referred to as BCF models, after Burton, Cabrera and Frank [3] who introduced these models and were the first to solve diffusion equations on the terraces that govern the quasi-equilibrium “adatom” distribution. This distribution is then coupled to conservation laws that govern the step-edge velocities. Much recent work has focused on solving the TS version of the BCF model using the level-set method [5–7].

* Corresponding author.

E-mail addresses: schulze@math.utk.edu (T.P. Schulze), psmerek@umich.edu (P. Smereka), weinan@math.princeton.edu (W. E).

BCF models can be viewed as coarse-grained KMC models. In a previous paper, Schulze and E [8] make this connection explicit by studying a quasi-equilibrium growth process governed by discrete rate-equations

$$\frac{\rho_i^{n+1} - \rho_i^n}{\Delta t} = k_{i-1}^+ \rho_{i-1}^n - (k_i^- + k_i^+) \rho_i^n + k_{i+1}^- \rho_{i+1}^n - \alpha_i \rho_i^n + F, \quad (1)$$

where ρ_i^n is the expected number of adatoms at site i and discrete time-step n . The various coefficients (k_i^\pm , α and F) are analogous to those used in KMC simulations. In particular, the hopping rates k_i^\pm are based on the local surface topography and evolve with it on a slow time scale. It was also shown that Eq. (1) could be used to simulate film growth, providing what amounts to a numerical solution of the BCF model.

By discretizing the model on the actual crystal lattice, one can potentially accommodate atomic-scale details like kinks, nucleation and surface chemistry which would otherwise be difficult to treat in a conventional BCF model formulated in terms of PDE's. Two drawbacks of this approach are that it is a deterministic model and that solving a diffusion equation discretized on atomic length scales is very slow. These shortcomings are shared by other numerical treatments of the BCF model. The first issue can be addressed by including noise in the model, but it may be difficult to do this in a way that is consistent with KMC. Both problems are addressed here through the use of a hybrid scheme that combines the BCF and KMC models in a way that is reminiscent of domain decomposition. Like other domain decomposition methods, the hybrid method is well-suited to parallel computation. The essential idea is to solve the BCF model using a coarse grid on terraces where the KMC simulation reduces to a homogeneous random walk and to couple this with KMC simulations in regions with more complicated topography. The result is a fast method that provides a good approximation to KMC. This method will be especially well suited to micron-scale calculations with widely separated steps that would overwhelm a standard KMC simulation.

The idea of combining continuum and atomistic simulations has recently received much attention. Examples include the quasi-continuum method [9], work combining Navier–Stokes with MC simulation of fluids [10], work combining molecular dynamics simulations with continuum models in crystal growth applications [11] and multi-scale computations of dendritic growth using MC and continuum methods [12]. E and Engquist discuss a general hybrid scheme for multi-scale problems [13]. While we discuss the present method in the context of epitaxial growth, it is clear that it can be applied to any mathematical model involving inhomogeneous random walks.

Below, we begin by reviewing the KMC and BCF models. In Section 4 we describe our method in detail, discussing the computational cost and errors in Section 5. In Section 6, we evaluate the method using a simple one-dimensional step-flow problem and compare its performance with an analogous KMC simulation. We conclude in Section 7.

2. KMC simulations

KMC simulations have emerged as one of the fundamental tools for studying the growth of thin films and help to illustrate the physics of the growth process. This technique was first adopted in the early 1970s [14,15] and has bifurcated in numerous directions. The thinking behind this approach, like most epitaxial models, is largely focused on what is called “solid-on-solid” or, more descriptively, “cube-on-cube” epitaxy, where the crystalline films grow with an ortho-rhombic structure. The bulk of the work focuses on single-species, mono-atomic crystals.

For simplicity, the model will be described in a two- (i.e., “1+1”) dimensional setting. The ideas in the three- (i.e., “2+1”) dimensional case are essentially the same. The basic notion is that individual “adatoms” rain down upon a film/substrate, where they are allowed to occupy positions on a lattice. They can move from one site to an adjacent one according to some probabilistic hopping rate k_i^\pm which may vary from one site to the next:




Fig. 1. Some steps, an adatom and an illustration of the potential ϕ associated with the topography. On the left step edge we illustrate only a step edge barrier; on the right we have a Schwoebel barrier as well.

$$k_i^\pm = K(T) \exp \frac{\Delta\phi}{k_b T},$$

where i is an integer indicating the lattice site, a $+$ or $-$ superscript indicates the direction of the hop, $K(T)$ is a temperature-dependent attempt frequency, k_b is the Boltzmann constant and $\Delta\phi$ represents the potential barrier to be overcome in the transition from one site to the next. Typically, this barrier, which is illustrated in Fig. 1, depends on a count of the nearest in- and out-of-plane neighbors. Note the deep well at the lower end of each step, due to the significant attraction of the in-plane nearest-neighbor. It is also common to have a step-edge barrier at the top end of steps. This is a manifestation of what is known as the Schwoebel effect [16] and is due to the fact that the last site on each step has one fewer out-of-plane neighbor than the other sites on the terrace.

Once the atoms are on the surface of the film, they perform a random-walk before becoming incorporated into the crystal or evaporating from the surface. The incorporation process is most likely to take place at step edges and, in the three-dimensional case, at kinks in the step edges where the atomic potential favors permanent attachment.

KMC algorithms are based on the method of Bortz, Kalos and Lebowitz (BKL) [17]. The BKL algorithm is built on the assumption that the model features M independent Poisson processes with rates r_m that sum to give an overall rate R which can be used to

1. Decide which event to execute.
2. Randomly select the time it takes for that event to occur from a Poisson distribution.

After generating a random number $r \in [0, R)$, a binary search can locate the event in $O(\log(M))$ operations [18]. Alternately, for models based on nearest neighbors which feature relatively few distinct rates, the event can be located without searching [19].

3. Continuum model

Using the KMC picture, one can motivate a formal semi-discrete model where one divides the surface of the film into terraces and steps, solving continuum equations on the former with discrete boundary conditions on the latter. Historically, this approach was adopted first [3], as KMC is completely impractical without the use of high-speed computing – the simulation of crystals on the micron scale being prohibitively expensive even with today's machines.

For simplicity, we continue to look at the 1 + 1-dimensional scenario, where a simple version takes the form:

$$\rho_t = D\rho_{xx} + F, \quad x \in (x_i, x_{i+1}). \quad (2)$$

Here $\rho(x)$ is a continuous adatom density on each terrace, but is expected to be discontinuous at step locations x_i . This piecewise continuous density may be thought of as the expected value of the discrete ρ_i in a KMC model at lattice sites sufficiently far from the step edges and in a limit where there are a large number of such sites per terrace. The diffusion constant D is, of course, directly related to the uniform hopping rate k_s that one expects to apply to all sites with no in-plane nearest neighbors. Included on the right-hand side of the equation is a source term F , which is often taken to be constant.

The locations of the steps must be determined as part of the solution, so that a total of three boundary conditions are needed at each step:

$$D\rho_x|^{+} = \alpha^{+}(\rho|^{+} - \rho_e), \quad -D\rho_x|^{-} = \alpha^{-}(\rho|^{-} - \rho_e), \quad v \frac{dx_i}{dt} = \left[D\rho_x + \frac{dx_i}{dt} \rho \right]_{-}^{+}, \quad (3)$$

where the $+$ and $-$ superscripts indicate that the quantity is to be evaluated as $x \rightarrow x_i$ from the left- and right-hand sides, respectively. The coefficients α^{\pm} , the equilibrium adatom concentration ρ_e and the volume of an adatom v are, in principle, related to the various atomistic parameters in the KMC description. In particular, the α^{\pm} represent the tendency for adatoms to attach to the upper and lower sides of a terrace.

4. The hybrid scheme

It is evident that fluctuations play a large role in KMC simulations and that statistical details cannot be well-approximated by a deterministic version of the BCF model. The aim of the hybrid scheme is to simulate film growth faster than KMC while addressing this issue. We argue that it is events which alter the homogeneous landscape which contribute the most to any differences. This would include attachment, detachment, vacancy formation, nucleation and other interactions between adatoms. One approach to eliminating these differences would be to include stochastic terms in the BCF model to compensate for the missing details. Such simulations can be calibrated to perform well in certain regimes, but a general theory that does not rely on fitting parameters has not been developed. Here, we adopt a more direct approach of using KMC to simulate growth in inhomogeneous regions while using a diffusion equation to track the adatom density on terraces. Since the vast majority of events in a KMC simulation consist of hopping on the terrace, this approach can significantly reduce simulation time.

Below, we discuss simulation details for the continuum and KMC regions, the coupling between them and issues associated with the free boundary, including cases where nucleation or other events require converting portions of the domain from one type to the other.

4.1. Continuum regions

Following the BCF tradition, the fundamental currency in the continuum region is an adatom density $\rho(x)$ which measures the expected number of adatoms per site. From the KMC point of view, a single adatom hopping on a terrace sufficiently far from steps performs a homogeneous discrete-space, continuous time random walk. For the nearest-neighbor model described above, the equation

$$\frac{\partial \rho_i}{\partial t} = \frac{D}{\Delta x^2} (\rho_{i-1} - 2\rho_i + \rho_{i+1}) \quad (4)$$

exactly describes the evolution of the adatom's probability distribution ρ_i .

For two or more adatoms performing independent random walks, we could evolve separate equations analogous to (4) for each adatom, so that given information about their initial locations, we could exactly track their subsequent distributions. The linearity of (4) implies that the sum of these distributions satisfies

the same equation and it is this number density of adatoms that we use in our computations. There is no error in using this approximation until there are interactions between the continuum and KMC regions. This “coupling” error is discussed along with other errors in Section 5. For now, it suffices to say that the number-density approximation is valid if there are a large number of random walkers per cell. While this is not the case for epitaxial growth if we discretize on the lattice scale, the coarse-grained description that we use to speed up the simulation does have this property.

The algorithm employs a uniform coarse mesh, with M sites per cell. This number is chosen to be commensurate with the total domain size of N sites. We refer to M as the cell-width and the corresponding regions as cells. From now on, we will denote the number-density of adatoms per cell as ρ_j and use a j -subscript when speaking of cell-based quantities. The cells are partitioned into a set Ω_C corresponding to continuum regions and a set Ω_K corresponding to KMC regions. A cell is placed in the KMC region if it contains an edge or it is next to a cell that contains an edge. This partitioning must be updated as the simulation proceeds and is designed to keep the step edges, which produce an inhomogeneous diffusion environment, separated from the continuum region. These features of the method are illustrated for the 2+1-dimensional case in Fig. 2.

Like the diffusion, stochastic effects due to deposition can be treated to good approximation within the continuum description. Thus, after introducing a discrete time-step Δt , the algorithm within the continuum region can be written as

$$\rho_j^{n+1} = \rho_j^n + D \frac{\Delta t}{\Delta x^2} (\rho_{j-1}^n - 2\rho_j^n + \rho_{j+1}^n) + F\Delta t, \quad (5)$$

where $\Delta x = aM$ is dimensional cell-width in terms of the lattice spacing a and number of sites per cell M . Evaporative effects could also be included.

One might try to further accelerate the simulation process by adopting an implicit version of (5), but our investigations show that this is of limited use because the time-step restriction implied by the CFL condition,

$$\Delta t < \frac{\Delta x^2}{2D}, \quad (6)$$




Fig. 2. A plan view of a 2+1-dimensional film surface showing the coarse grid, an edge and shading cells assigned to the KMC region.

is no more severe than other constraints on the system. The explanation for this lies in understanding that the origin of the CFL numerical instability is the nearest-neighbor accounting practice in an explicit finite difference scheme. Nearest neighbor KMC models suffer from an analogous problem – information can only propagate through the system at a finite rate. If you take too large of a time-step you will find, for example, that step edges travel through more than one cell in a single time-step. This will rapidly lead to numerical problems. Also, note that the larger values of Δx provided by a coarse spatial grid tend to ease the CFL restriction.

4.2. KMC regions

The adatom concept does not appear in traditional KMC, which tracks the height of the lattice at each site h_i . While it is possible to use this conventional framework in the present scheme, it is more convenient to adapt the KMC algorithm so that adatoms and height are tracked separately. Thus, we define an adatom as a location on the lattice which has a height h_i which is one more than any of its lateral nearest neighbors, all of which share the same height. By itself, this does not actually change the KMC model, only the data structure. One replaces the old h_i with a new height variable and a variable ρ_i which contains either zero or one adatom, as defined above. This facilitates counting the number of adatoms in a KMC cell, information which is needed for communicating with neighboring continuum cells and is useful for presenting data. The diffusion equation does not, by itself, allow for interactions between random walkers nor the formation of vacancies. As discussed below, these elements can be introduced into the algorithm, but in the present implementation we have limited the set of possible events within the KMC region to those corresponding to adatom motion, attachment and detachment. Thus, the adatoms do not “see” one another when moving on an otherwise smooth terrace, there is no possibility for nucleation, vacancy formation or step bunching and there can be more than one adatom per site. The edges are put in at the beginning of the simulation and their number remains fixed throughout the simulation. These simplifications seem appropriate for 1+1-dimensional simulations where adatoms cannot go around one another. To do otherwise would overemphasize adatom interaction. Finally, there are no step-edge-barriers, so that just three rates describe the KMC simulation: a deposition rate F , a diffusion constant D and a detachment rate β . Attachment occurs when an adatom hops to the site in front of an edge; hence it occurs with the same rate D as any other hop on the homogeneous terrace. Fig. 3 illustrates this reduced KMC model. These simplifications are not necessary to make the hybrid scheme work, but are introduced to simplify this initial exploration of the method.

4.3. Interface between KMC and continuum regions

We now turn to the exchange of mass between continuum and KMC regions. For the 1+1-dimensional implementation of the hybrid method, we maintain an edge list E_k which contains the site index i of the current edge positions. From this, we can also determine the cell index $j = \text{Int}[(i - 1)/M + 1]$. In a 2+1-dimensional simulation, one would want to maintain a list B_k of cells on the boundary of the KMC region,




Fig. 3. An illustration of two KMC cells containing four adatoms, an edge and having a cell-width $M = 8$. The edge will adsorb an adatom and move forward if an adatom lands in front of it; it will emit an adatom and move backward if a desorption event occurs.

but here we will assume that the steps remain separated, so that the cells to the left and right of a cell containing an edge are on the boundary.

From the KMC perspective, we add to the list of events in the interior of the KMC region events corresponding to atoms hopping into and out of the continuum region. Hops from $\Omega_K \rightarrow \Omega_C$ occur at the rate D for a hop without attachment; those from $\Omega_C \rightarrow \Omega_K$ occur at this rate multiplied by the expected number of atoms per site within the continuum cell, $D\rho_j$. On the continuum side of these exchanges, we immediately update the value of ρ_j on the cell boundary, setting

$$\rho_j^{n+v} = \rho_j^n \pm 1/\Delta x, \quad (7)$$

where the v in the first superscript indicates that the KMC time-step is an arbitrary fraction of the full time-step Δt . In general, there will be many KMC events in one time-step Δt of the continuum region. In these transitions, the adatom mass is immediately spread over or removed from the entire cell. If a hop from $\Omega_C \rightarrow \Omega_K$ would leave the adatom density negative in the continuum region, a remedy must be prescribed. In the present implementation we have used two approaches: either removing such events from the event list so that there must be a full adatom's worth of mass in a cell or if an adatom jumps out of the cell when the expected density is less than one, the mass is recovered from neighboring continuum cells. The former method is simpler and somewhat faster, but the second is more accurate and allows one to extend the methods applicability to somewhat lower adatom densities.

As the simulation proceeds, a check is made to determine if any of the edges have moved a full cell-width M in either direction. If this has happened, a shorter continuum time-step is taken on that particular iteration. If the cell-width is adjusted appropriately, however, this rarely happens and the KMC portion of the simulation normally continues for the full continuum time-step Δt . Eq. (5) is then used to update the continuum-region adatom density, applying homogeneous Neumann conditions at the boundaries with KMC regions to prevent any additional transfer of mass beyond that prescribed by (7).

Before continuing with the KMC portion of the simulation a check must be made to see if any of the edges have changed cells. If an edge has moved forward into a new cell, the cell in front of its new position must be converted from $\Omega_C \rightarrow \Omega_K$ and the one behind its old position must undergo the opposite conversion. The first of these is accomplished by randomly positioning $\text{Int}(M\rho_j)$ adatoms in the new KMC cell and adding the new events to the event-list L_k . There will normally be a fraction of an atom left in the cell; in the present implementation of the hybrid scheme we have chosen to leave this, so that it will remain in that cell until the next time it becomes a continuum cell. The second conversion is accomplished by locating the adatoms in the cell and removing them and the corresponding events from the event-list. To do this efficiently without searching requires an inverse list for the adatom positions, which is slightly complicated with the possibility of multiple adatoms per site. This feature of the present simulations can be avoided by allowing only one adatom per cell. If an edge moves backward, the analogous procedures are performed. Finally, note that for all cells, it is convenient to keep track of the number of adatoms per site using the variable ρ_j .

4.4. Nucleation and vacancy formation

Nucleation within the KMC region can obviously be accomplished by relaxing some of the simplifying assumptions being used here. This is better done in a 2+1-dimensional setting, where the coverage of the surface by adatoms will be more realistic. Within the continuum region, nucleation can be modeled by adding nucleation events to the KMC event list with rates proportional to $(\rho_j)^2$. A similar approach has proven effective in the context of level-set simulations of the BCF model [5,6]. Since these events will be rarely executed, they will not add appreciably to the KMC execution time. If one of the events is executed, simply convert the cell in which it occurs and the neighboring cells into KMC cells via the mechanism described above. Note that this creates a pair of edges in 1+1-dimensional simulations. In a similar way, a pair of edges could annihilate one another, in which case one would convert cells into continuum cells.

It is possible that there may ultimately be so much nucleation in 2+1 growth that the hybrid scheme will revert to a full KMC simulation. Fortunately, a full implementation of the method as described here will automatically detect this situation and add little overhead cost to a conventional KMC simulation. In contrast, the hybrid scheme will be at its best in situations with widely separated steps – a situation that KMC codes handle very poorly. In an island-dynamics simulation, for example, it could be that in the early stages of growth the simulation is essentially a KMC simulation, but that as coarsening occurs and grain sizes become large, the computational advantages of the hybrid scheme would be realized.

While vacancy formation is not a common event and could be ignored to good approximation for many applications, it can be included in a similar way. In this case, simply add the vacancy-formation events within the continuum region back to the KMC event list. If a list-based algorithm [19] is being employed this will not add significantly to the cost of the KMC portion of the simulation. If a vacancy formation event occurs, convert the cell and its neighbors into KMC cells. More importantly from a performance perspective, one would want to monitor the KMC cells for situations where such surface inhomogeneities are eliminated, so that these cells could be converted to continuum cells.

5. Computational cost and accuracy

In this section we set out to identify the dominant sources of error and gain insight into their scaling properties. The results will show that there are complicated trade offs between accuracy and computational speed. The situation is further complicated by the need to obey the time-step restriction dictated by the CFL condition (6). We note at the outset that the main point of the method is that the overall error is significantly less than that of the BCF model while the computational speed is significantly faster than KMC.

5.1. Computational cost

We start by examining the scaling for the cost of the simulation. For a purely KMC simulation of an N -site region, there must ultimately be N net attachment events per layer of growth. Each of these may be accompanied by a number of attachment and detachment events, the latter being more costly as they in turn generate additional diffusion events. The dominant cost associated with each net attachment, however, is typically the initial diffusion of an adatom to an edge. Thus, if the typical terrace width is L sites, one expects the total cost for the KMC simulation to be $O(NL^2)$ events per layer.

For the hybrid scheme, we can divide the total cost into two pieces:

$$\text{total cost} = \text{KMC cost} + \text{continuum cost}.$$

The KMC cost still involves N net attachment events per layer, but now is only associated with diffusion events within the KMC region around each edge:

$$\text{KMC cost} = O(NM^2), \tag{8}$$

where M is the cell-width. This turns out to be the dominant cost and the fact that the KMC region is actually $3M$ wide weighs heavily in the actual performance of the algorithm. This suggests a potential area for further improvements in computation time, but these would come at the cost of further errors.

To see that (8) is dominant we must estimate the cost for the continuum region. We will measure this cost in terms of the $O(N/M)$ calculations that must be performed each time a continuum time-step is taken. The number of these time-steps per layer of growth can be estimated as

$$(\text{events/layer})(\text{time/event})(1/\Delta t),$$

where the number of KMC events per layer is given by (8) and the physical time associated with each event can be estimated as the reciprocal of the number of events per step M times the number of steps per layer N/L times the diffusion constant D , for a continuum region cost estimate of

$$O\left(\frac{NL}{\Delta t}\right).$$

If we further take the time-step Δt as scaling according to the CFL condition (6) we get

$$\text{Total cost} = O(NM^2) + O\left(\frac{NL}{M^2}\right),$$

which is dominated by the KMC cost for sufficiently large cell-width. When the latter term is negligible, the improvement in the overall computation time scales directly proportional to the area reduction in the region that is being simulated by KMC.

5.2. Accuracy

For purposes of discussing accuracy, we shall take the point of view that KMC is the exact process we wish to simulate. In this regard, one possible source of error are processes which are accounted for in the KMC domain, but not in the continuum region. For example, if we had a conventional KMC simulation where adatoms interact and, in particular, can nucleate, we would have no hope of achieving nucleation by merely time-stepping Eq. (5). This type of error is non-systematic and can only be eliminated by adjusting the two models so that they are analogous to one another. For example, we formulated an adapted version of KMC in Section 4.2 so that nucleation would play no role in either region. More elaborate simulations could account for nucleation in both regions. Beyond obvious incompatibility of the models in the two regions, we identify two additional sources of error: error due to coarse graining and error due to the number density approximation.

The coarse graining approximation leads to two effects. The first of these, discretization error, is the same error encountered in the standard solution of the diffusion equation by the finite difference algorithm. This suggests an $O(\Delta x^2) + O(\Delta t)$ error, so that this source of error is reduced by choosing smaller cell sizes and taking smaller time-steps. Recall that $\Delta x = aM$ where a is the lattice spacing and M is the number of sites per coarse-grained cell. In the continuum limit $a \rightarrow 0$, the resulting diffusion equation has smooth solutions and discretization is the only source of error. The lattice-based version of the diffusion Eq. (5) suffers an analogous error when we take $M > 1$. The stochastic nature of the simulation gives rise to a second coarse-graining issue in that the variance of the cell-based adatom density will be lower than that based on a non-coarse-grained density. This error also vanishes when $M = 1$.

The second source of error, due to the number density approximation, manifests itself in the coupling of the KMC and continuum regions. It has the opposite trend in that it gets worse as you decrease the cell-size. Eq. (4), which is discretized on the lattice scale, is an exact equation for the evolution of the number density as long as there are no interactions between the continuum and KMC regions. Interactions lead to errors because precise information about the probability distribution of the individual adatoms is lost in using the number density approximation.

When the number of adatoms per cell is small, there are several mechanisms that produce the latter type of errors. For example, when we remove an atom from the continuum region, we are unable to identify the correct distribution for the particular particle that we remove. If we had complete information about all of the particles, we could choose one particle to remove based on its relative probability of being in the boundary cell and then “collapse” just that component of the total distribution. Instead, in the version of the method that requires a threshold before an atom can hop out of the continuum region, we merely




Fig. 4. A sketch of a 1+1-dimensional continuum domain on the cell scale showing the adatom density ρ_j partitioned into eight regions (which should all have the same area in a perfect sketch) associated with individual adatoms. When adatoms are removed from the region, the mass is removed from cells on the boundary (shaded region) and this can be a significant source of error if the adatom density is low.

remove one adatom's worth of mass from the cell closest to the boundary (see Fig. 4). Similarly, by repeatedly approximating the rate for hopping from $\Omega_C \rightarrow \Omega_K$ as $D\rho_j$ we are not allowing for correlations in the probability that an adatom is in the boundary cell on subsequent hopping attempts. As the cell-size is decreased, severe coupling errors associated with the threshold-version of the method can be encountered. In general, the relative size of coupling errors can be reduced by making sure that there are several adatoms (the more the better) in each cell, at which point the number-density approximation is valid. A precise analysis is difficult due to the complicated nature of the interactions, but the important point is that, in contrast to conventional finite difference schemes, this prevents the arbitrary reduction in total error by reducing the cell size Δx .

In our present implementation, the size of the KMC region surrounding each step does not scale with the cell-width in that it is always three cells wide. As a result, the error is actually reduced to zero for sufficiently large cell sizes, as the simulation reduces to a pure KMC simulation. Thus, when the adatom density is low and coupling errors dominate, the total error will actually decrease monotonically with increasing cell-width and one must judge how much accuracy to give up in return for increased simulation speed. When the adatom density is larger and the coupling errors are not as dominant, there will be a local minimum in the error for an intermediate value of the cell-size. In this case, there is extra incentive to sacrifice accuracy so as to operate in this locally optimized regime.

6. Comparing hybrid and KMC simulations

Our aim in this section is to demonstrate the feasibility of the hybrid scheme by showing that it is indeed faster than KMC and that the error can be reduced to tolerable levels. The results also support the conclusions of the previous section regarding the behavior of the error as the cell-width is adjusted. In one set of calculations, we concentrate on the highly simplified situation of a single step growing with periodic boundary conditions so as to highlight these issues. For comparing the performance of the KMC and hybrid simulations, we use simulations that feature two steps, so that we can examine both the fluctuations in the number of adatoms on the surface and in the separation between the steps. Note that all of these fluctuations would be absent in a deterministic BCF model.

All of the calculations use periodic boundary conditions for the adatom density and/or local KMC environment. The steps descend to the right, with the lowest step reconnecting to the highest to form an "Escher" staircase. To give some sense of the overall behavior of the simulations, in Figs. 5(a)–(c) we present a series of adatom density profiles at three separate times for the same simulation. Here, the adatom density $\rho_j(t)$ recorded in the KMC region is obtained by counting the adatoms in the KMC cells. Note that fluctuations in $\rho_j(t)$ are much lower in the continuum region, where the parabolic profile that would result

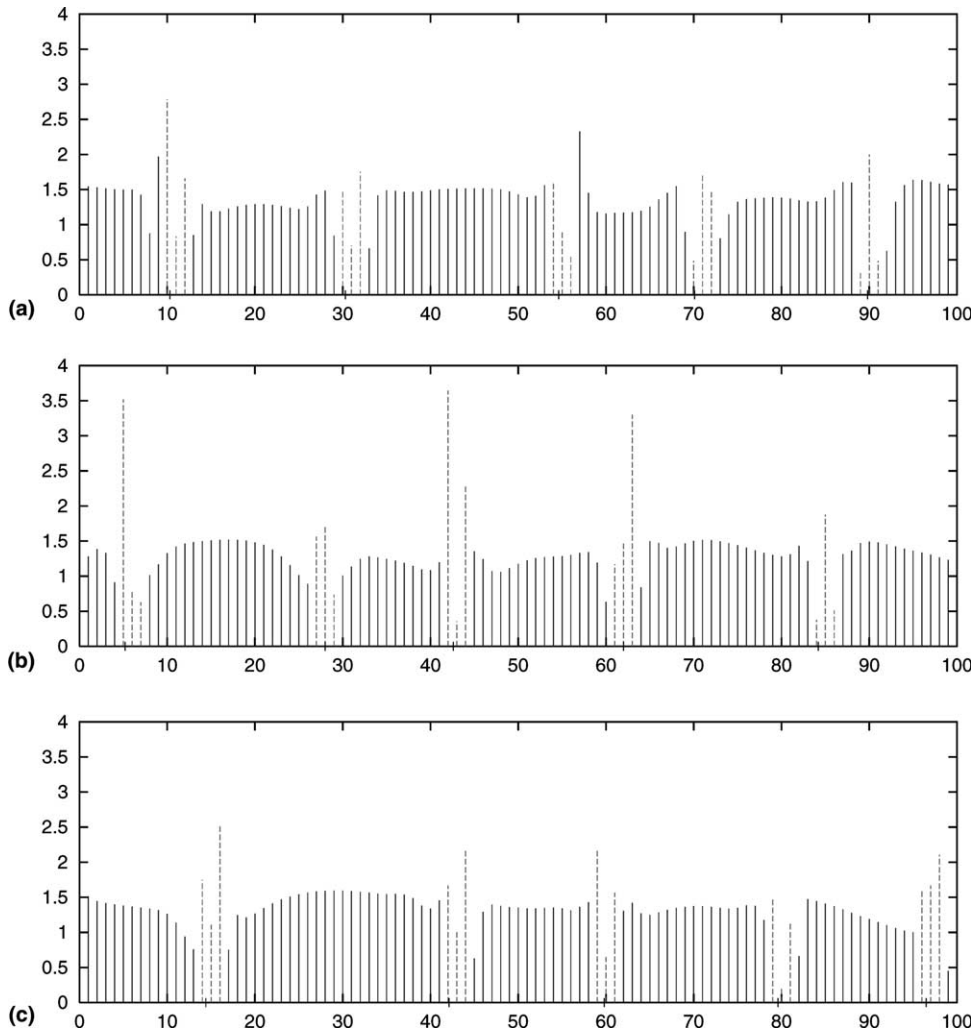


Fig. 5. Time series of adatom density on cell scale after approximately 10, 20, and 30 layers of growth. There are 100 cells with a cell-width of 10.

from a deterministic BCF model can be recognized, particularly toward the center of continuum regions. This effect becomes more pronounced with higher adatom densities.

The data that we collect in Tables 1 and 2 come from simulations with an $N = 200$ -site domain and a single step. The deposition rate was set to $F = .1$ adatoms per site per second and the homogeneous hopping rate to $D \approx 2 \times 10^5$ hops per second. While these choices were guided by previous simulations in the KMC literature, we do not claim any physical significance for them and do not study the behavior of the system as we change these parameters. For our present purposes, it is more relevant to understand how our method performs as the detachment rate varies relative to the hopping rate. To this end, we performed two sets of simulations: one with a low detachment rate of $\beta \approx 20$ and one with a higher rate of $\beta \approx 50$. The data were averaged over a period of time that allowed 100 layers of growth. Since the cell-width must be chosen commensurate with the domain width, we are limited to collecting data from simulations using $\Delta x = \{8, 10, 20, 25, 40\}$ and present this information in Tables 1 and 2.

Table 1
Data from simulations with low detachment

Δx	$\langle M \sum \rho_j \rangle$	$\langle (M \sum \rho_j)^2 - \langle M \sum \rho_j \rangle^2 \rangle^{1/2}$	CPU time
8	30.4	2.40	141
10	27.1	2.76	185
20	22.8	3.88	354
25	22.5	4.16	481
40	23.4	4.51	770
KMC	22.3	4.53	828

Columns correspond to cell-width, time average of adatoms on surface, variance of adatoms on surface and CPU time in seconds. The KMC data are shown in the last line.

Table 2
Data from simulations with high detachment

Δx	$\langle M \sum \rho_j \rangle$	$\langle (M \sum \rho_j)^2 - \langle M \sum \rho_j \rangle^2 \rangle^{1/2}$	CPU time
8	57.5	4.93	414
10	55.9	6.02	524
20	55.4	6.60	991
25	55.8	6.94	1364
40	55.9	7.25	2027
KMC	54.6	6.80	2154

Columns correspond to cell-width, time average of adatoms on surface, variance of adatoms on surface and CPU time in seconds. The KMC data are shown in the last line.

The first table is for a case where the error continues to decrease, as explained in Section 5 until it asymptotes to a full KMC simulation. The second table, featuring the increased detachment rate and a higher average adatom density, suggests an error structure with a local minimum. Again, this is consistent with the discussion in the previous section. In both cases there are values of Δx for which the data agree quite well with the corresponding KMC simulation and the computation time is significantly reduced, suggesting that the method will indeed be useful for improved large scale simulations of epitaxial growth. In comparing computation times between the two methods, we note that the simplified KMC used for these simulations is substantially faster than that typically used.

A more extensive set of simulations was performed to compare the results of the hybrid scheme to those of the KMC simulations. These simulations are summarized in Figs. 6–8 and feature a 1000-site periodic domain with two steps. The simulations use an edge detachment rate that is one tenth the diffusion constant and a range of D/F ratios. The calculations using the hybrid scheme were repeated for several values of the cell-width. The addition of a second step allows comparison of fluctuations in the step-width as well as fluctuations in adatom density. The standard deviation in these quantities was averaged in time from the beginning of the simulation until $t = 5000$. A larger simulation time would further reduce the fluctuations in the averages, but the cost of the KMC simulations discouraged this. As it stands, each data point for the KMC simulations took on the order of a day to obtain. The hybrid scheme was up to 10 times faster with the difference being the greatest for the smallest values of D/F – i.e., when the number of adatoms was the greatest. The quantitative comparison between the two methods was also best in this regime, as expected. For all values of D/F the hybrid scheme underestimates the level of noise in the system, though it does much better than a deterministic BCF model, which has no fluctuations. The agreement becomes quite good for low values of D/F . In principle, one could add additional parameters to better capture fluctuations. In the BCF context – particularly in 2+1 growth – this must be done in a hopelessly ad hoc fashion. For the hybrid scheme, one could model the dominant source of error and systematically correct




Fig. 6. The time and space averaged surface adatom density for the KMC simulations (diamonds) and the hybrid scheme (crosses) using cell-widths $M = \{20, 25, 40\}$.




Fig. 7. The standard deviation (in time) of the surface averaged adatom density as a function of the ratio D/F . The solid curve (diamonds) are from the KMC simulations and the remaining curves are for the hybrid scheme with cell-widths $M = \{20, 25, 40\}$ sites per cell.




Fig. 8. The standard deviation (in time) of the the step width as a function of the ratio D/F . The solid curve (diamonds) are from the KMC simulations and the remaining curves are for the hybrid scheme with cell-widths $M = \{20, 25, 40\}$ sites per cell.

for it. However, we feel the better selling point of the hybrid scheme is that it is reasonably faithful to KMC without introducing new parameters and complications.

7. Conclusion

In this paper we have presented a new method for simulating epitaxial film growth that combines nano-scale KMC simulations with numerical solutions of the diffusion equation suitable to a scale of microns. The larger length scale of the continuum equation will allow future research to perform simulations and study phenomena that involve the flow of multiple steps, being particularly useful when the step separation is large. Similarly, this method will be useful for island dynamics simulations featuring many exposed layers, particularly in the later, highly coarsened stage of growth. Previously, this sort of study could only have been done with simulations that either neglected the effects of fluctuations or accounted for them in a more limited way. We believe this initial study of the method suggests that it will adequately account for stochastic effects and facilitate the study of their impact on large scale features of film growth. Finally, we note that the scheme presented here is not limited to epitaxial growth; it is well-suited to any model based on random walks where large length scales must be addressed.

Acknowledgements

The authors would like to acknowledge support through ONR Grant N00014-010100674 (W.E.). Additional support was provided by NSF Grant DMS 0103825 (T.P.S.) and DMS 0207402 (P.S.). We would also like to thank Eric Vanden-Eijnden for helpful discussions.

References

- [1] M. Kotrla, Numerical simulations in the theory of crystal growth, *Comp. Phys. Comm.* 97 (1996) 82.
- [2] P. Smilauer, D.D. Vvedensky, Coarsening and slope evolution during unstable epitaxial growth, *Phys. Rev. B* 52 (1995) 14263.
- [3] W.K. Burton, N. Cabrera, F.C. Frank, The growth of crystals and the equilibrium structure of their surfaces, *Philos. Trans. Roy. Soc. Lond.* 243A (1951) 299.
- [4] R. Ghez, S.S. Iyer, The kinetics of fast steps on crystal surfaces and its application to the molecular beam epitaxy of silicon, *IBM J. Res. Dev.* 32 (1988) 804.
- [5] R.E. Caflisch, M.F. Gyure, B. Merriman, S.J. Osher, C. Ratsch, D.D. Vvedensky, J.J. Zinck, Island dynamics and the level set method for epitaxial growth, *Appl. Math. Lett.* 12 (1999) 13.
- [6] M.F. Gyure, C. Ratsch, B. Merriman, R.E. Caflisch, S.J. Osher, J.J. Zinck, D.D. Vvedensky, Level-set methods for the simulation of epitaxial phenomena, *Phys. Rev. E* 58 (1999) 6927.
- [7] P. Smereka, Spiral crystal growth, *Physica D* 138 (2000) 282.
- [8] T.P. Schulze, W. E, A continuum model for the growth of epitaxial thin films, *J. Cryst. Growth* 222 (2000) 414.
- [9] E.B. Tadmor, M. Ortiz, R. Phillips, Quasicontinuum analysis of defects in crystals, *Philos. Mag.* A 73 (1996) 1529.
- [10] A.L. Garcia, J.B. Bell, W.Y. Crutchfield, B.J. Alder, Adaptive mesh and algorithm refinement using direct simulation Monte Carlo, *J. Comp. Phys.* 154 (1999) 134.
- [11] W. E, Z. Huang, Matching conditions in atomistic-continuum modeling of materials, *Phys. Rev. Lett.* 87 (12) (2002) 135501.
- [12] M. Plapp, A. Karma, Multiscale finite-difference-diffusion Monte-Carlo method for simulating dendritic solidification, *J. Comp. Phys.* 165 (2000) 592.
- [13] W. E, B. Engquist, Heterogeneous multi-scale method, preprint (2002).
- [14] F.F. Abraham, G.W. White, Computer simulation of vapor deposition on two-dimensional lattices, *J. Appl. Phys.* 41 (1970) 1841.
- [15] G.H. Gilmer, P. Bennema, Simulation of crystal growth with surface diffusion, *J. Appl. Phys.* 43 (1972) 1347.
- [16] R.L. Schwoebel, Step motion on crystal surfaces 2, *J. Appl. Phys.* 40 (1969) 614.

- [17] A.B. Bortz, M.H. Kalos, J.L. Lebowitz, New algorithm for Monte-Carlo simulation of Ising spin systems, *J. Comput. Phys.* 17 (1975) 10.
- [18] J.L. Blue, I. Beichl, F. Sullivan, Faster Monte-Carlo simulations, *Phys. Rev. E* 51 (1995) 867.
- [19] T.P. Schulze, Kinetic Monte-Carlo simulations with minimal searching, *Phys. Rev. E* 65 (2002) 36704.