



A hybrid scheme for simulating epitaxial growth

Tim P. Schulze*

Department of Mathematics, University of Tennessee, 121 Ayres Hall, 1403 Circle Drive, Knoxville, TN 37996-1300, USA

Received 8 August 2003; accepted 18 November 2003

Communicated by G.B. McFadden

Abstract

This article continues the development of a hybrid scheme for simulating epitaxial growth that combines the Burton–Cabrera–Frank (BCF) model with kinetic Monte-Carlo (KMC) simulation. This is the first implementation of the scheme for “2+1” dimensional growth. Other improvements over an earlier version include the use of a more conventional KMC model and some refinement in the handling of the boundary condition between the KMC and continuum regions. The method is used to examine unstable step-flow with direct comparison to KMC simulations. The results are extremely good with respect to computational speed and reveal effects due to fluctuations to a much greater extent than the BCF model alone. This method will be especially useful in scenarios with widely separated steps and high adatom densities, as these are situations that cannot be easily simulated with KMC due to increased computational cost. The hybrid method is extremely flexible and can be coupled interchangeably to any KMC scheme.

© 2003 Published by Elsevier B.V.

PACS: 68.55.–a

Keywords: A3. Epitaxy; A3. Monte Carlo simulation

1. Introduction

The simulation of epitaxial growth by the kinetic Monte-Carlo (KMC) method is a popular and well-established way of studying the evolution of surface morphology [1]. Even with rapidly increasing computational resources, however, this tool is limited to small length scale simulations, on the order of a micron, and scenarios, like narrow step-spacing, that favor low adatom densities. As a result, the Burton–Cabrera–Frank (BCF), or

terrace–step, model [2] has gained some popularity for larger scale simulations when combined with level-set techniques [3–5]. Without the latter enhancement, the difficulty of the BCF free-boundary problem has limited this model to theoretical studies featuring uniform step trains, island stability or other idealized geometries. In principle, these three- (or “2+1”)- dimensional simulations should use a terrace–step–kink version of the BCF model and some recent progress in this direction has been made by accounting for kink density [6]. This paper continues the development of another alternative—using a hybrid KMC-BCF technique that was recently introduced by Schulze et al. (SSE) [7]. In a similar approach, described

*Corresponding author. Tel.: +1-865-974-4302; fax: +1-865-974-6576.

E-mail address: schulze@math.utk.edu (T.P. Schulze).

further in Section 2, Russo et al. [8] have also combined the BCF and KMC models, applying their scheme in the study of diffusion limited aggregation. For a discussion of a broad framework for developing hybrid schemes in the context of multi-scale simulation, see Ref. [9].

In this paper we extend the hybrid scheme of SSE from (1 + 1)- to 2 + 1-dimensional epitaxial growth and compare the performance of the new method to that of a typical KMC scheme. Making use of a simplified MC scheme, the earlier study SSE suggested that a hybrid scheme would be computationally efficient while retaining more of the stochastic details that dominate KMC simulations (and presumably real growth) than deterministic BCF simulations. The present paper will show that the computational advantages of the hybrid scheme are readily achieved, and even somewhat enhanced, for the more realistic growth model studied here. As in the previous study, we find that the greater challenge comes in accurately reproducing the results of KMC simulations.

To make detailed comparisons, we focus on a step-flow regime with widely separated steps and negligible nucleation. This regime is desirable from a theoretical viewpoint as it allows one to explore step-flow instabilities induced by large Schwoebel barriers [10]. This regime may also be desirable from a manufacturing viewpoint, as it allows one to select a unique, self-organized surface morphology. Previously, KMC-based studies of this phenomenon [11] have been limited to narrowly spaced steps for two reasons: (1) to avoid nucleation when terraces are wide requires relatively high rates of diffusion, producing a large number of hopping events on the terraces and (2) the width of the terrace itself further aggravates this situation, as more hops are needed for an adatom to reach a step. In contrast, these conditions are ideal for the hybrid scheme, which combines KMC near steps with the coarse-grained solution of an adatom diffusion equation on the terraces.

The hybrid scheme, as implemented here, contains some features that were not present in SSE. In particular, the earlier paper used a simplified KMC model, featuring only two distinct rates—a homogeneous hopping rate and a detach-

ment rate. Thus, the model did not have a Schwoebel barrier and the only interaction between adatoms was mediated by the steps. The number of steps was fixed at the beginning of the simulation and remained constant—there was no possibility of nucleation, step bunching nor steps annihilating one another. In contrast, the KMC we use here, based on a model put forward by Smilauer and Vvdensky [12], is more typical of that found in the KMC literature. In addition to these essential changes, there are some modifications of the technique aimed at enhancing accuracy.

Ultimately, we find that the hybrid scheme can exhibit the step-flow instability under the same growth conditions that trigger the instability in pure KMC simulations, but that there are a number of potential pitfalls and a tendency to suppress the effects of stochastic fluctuations. In Section 2, we describe the hybrid scheme and discuss its implementation. In Section 3, we compare and discuss simulations produced via pure KMC and hybrid simulation. In the final section, we summarize and draw some conclusions about the likely benefits of using a hybrid scheme of the type discussed here.

2. The hybrid scheme

The BCF model [2] is based on the idea of using a number density of adatoms ρ , rather than precise adatom locations, to describe the surface configuration. One then models the evolution of the number density using a diffusion equation:

$$\rho_t = D\nabla^2\rho + F,$$

where ∇^2 is the surface Laplacian, F is the deposition flux and the diffusion constant D is related to the hopping rate for an isolated adatom performing a random walk on a flat surface. Additional terms, corresponding to evaporation for example, are often added to the model. Finally, boundary conditions are posed at the steps, typically modeled as continuous curves, that conserve mass and relate the flux of adatoms into the edge from either side to the local adatom density.

It is natural to make contact with the atomistic KMC models by discretizing the PDE on the crystal lattice. Scaling the horizontal and vertical coordinates by the lattice spacing(s) then gives:

$$\frac{\partial}{\partial t} \rho_{ij} = D(\rho_{i+1,j} + \rho_{i-1,j} + \rho_{i,j+1} + \rho_{i,j-1} - 4\rho_{ij}) + F_{ij}. \quad (1)$$

The difference equation allows for a great deal of flexibility in modeling atomic scale processes, providing, for example, a simple way of solving the terrace–step–kink version of the BCF free-boundary problem [13]. In this deterministic model, which Schulze and E refer to as an atomistic difference scheme (ADS), the discrete topography can be modeled by introducing time-dependent coefficients into (1) that are slaved to the surface height h_{ij} . These coefficients then evolve slowly with the topography by discrete shifts whenever the accumulated flux into “edge” sites passes through a multiple of the unit cell volume. A similar approach to triggering attachment events is used by Russo et al. [8], who also add a stochastic element to the model by attaching the atom at a random position on the boundary of an evolving island. This latter approach is principally aimed at the simulation of diffusion limited aggregation [14].

Ultimately, discretizing the BCF model on the crystal lattice is far too slow for terrace-scale simulation, and leads to the complicated question of how to properly introduce fluctuations into the model. In general, simulations based upon the BCF model can be made faster by discretizing on a larger length scale. Away from the steps, this can be done with a controlled loss of accuracy since the terraces are characterized by a homogeneous diffusion environment with very little interaction between adatoms. Unlike other implementations of the BCF model, the hybrid scheme described here retains a fully atomistic simulation in the vicinity of the steps, where the details of the film assembly are characterized by a number of widely varying rates and subject to complicated mechanisms like the step-flow instability that we examine in Section 3. In this way, we aim to produce

simulations that are faster than KMC while retaining much of the atomic level detail.

2.1. The domain decomposition

To this end, we use the step locations as a guide for partitioning the computational domain into two subdomains—one where we perform a conventional KMC simulation and a second where we solve the diffusion equation on the coarsened mesh (Fig. 1). We refer to each collection of sites in the coarse grid as a *cell*. Thus we reinterpret Eq. (1) with $\rho_{i,j}(t)$ representing the average adatom density per site in the cell with horizontal and vertical indices i and j , respectively

$$\frac{\rho_{ij}^{n+1} - \rho_{ij}^n}{\Delta t} = D \left(\frac{\rho_{i+1,j}^n - 2\rho_{ij}^n + \rho_{i-1,j}^n}{\Delta x} + \frac{\rho_{i,j+1}^n - 2\rho_{ij}^n + \rho_{i,j-1}^n}{\Delta y} \right) + F_{ij}, \quad (2)$$

where Δx and Δy are integers representing the cell-width and -height in lattice spacings. The time-step, Δt , is arbitrary provided the scheme remains stable. As explained in SSE [7], the explicit scheme is a natural choice for coupling with the




Fig. 1. A sketch of the film surface showing the coarse grid, a step and shading cells assigned to the KMC region.

nearest-neighbor KMC simulations, as one finds the CFL condition

$$\Delta t < \min\left(\frac{\Delta x^2}{4D}, \frac{\Delta y^2}{4D}\right), \quad (3)$$

is important not only for maintaining numerical stability in stepping the difference equation (2), but in maintaining the integrity of the overall simulation as the shape of the KMC region slowly evolves with the moving steps.

The partitioning scheme for the (2 + 1)-dimensional simulations has to be more flexible than in lower dimension, as steps can become curved, resulting in additional cells being added to the KMC domain. We are also allowing for the possibility of steps merging, nucleation within the KMC region and islands forming by portions of a step becoming detached by a necking mechanism. With these considerations in mind, it is convenient to keep track of cells containing steps rather than tracking the steps themselves. A cell is considered to be an “edge-cell” unless all of its sites are vicinal—have four lateral nearest neighbors sharing their height—or are adatoms—have a height one greater than their four lateral nearest neighbors. We also want to maintain a buffer of KMC cells around the steps, so that they can move seamlessly from one cell to another. For this reason, we include cells in the KMC region if they are edge-cells or they neighbor edge-cells. While this may seem complicated and does require some care in implementation, the machinery that is inherently part of any KMC simulation makes it computationally inexpensive to maintain this information. During the course of the simulation, it is both convenient and efficient to maintain two cross-indexed arrays for this purpose—a cell-list that is sorted by region (KMC or BCF) and an inverse cell-list that can, given the cell-coordinates, tell you which region (KMC or BCF) the cell is in and its location on the cell-list. This type of dual data structure is used for the KMC algorithm as well.

2.2. The KMC model

In the present implementation, the underlying KMC model assumes a cube-on-cube epitaxy

described by an integer height variable h_{ij} . The adatom concept that is central to the BCF model is not explicit in this description, but we imagine that a site with height exactly one more than its lateral nearest neighbors plays this role. The basic idea behind any KMC simulation is to enumerate a finite number of possible changes in the current configuration and prescribe a rate for each of those to occur. By modeling these events as being independent Poisson processes, one then uses the Bortz–Kalos–Leibowitz (BKL) [15] algorithm:

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

We indicate the total number of events as M , the sum of the individual rates r_m being R . After generating a random number $r \in [0, R)$, one must locate the corresponding event by identifying the interval where

$$\sum_{k=1}^{m-1} r_k \leq r < \sum_{k=1}^m r_k.$$

Many of the details in this subsection concern making the KMC algorithm more efficient irrespective of whether or not one implements the hybrid scheme. This is important for two reasons—first, the hybrid scheme will seamlessly reduce to a KMC simulation if the surface becomes extremely irregular, and second, we want to benchmark the performance of the hybrid scheme against the most efficient KMC simulation we can achieve.

In a typical KMC model there is one deposition and four hopping events possible for each site on the lattice. Thus, a 1000×1000 site lattice produces five million events. Normally, there are far fewer distinct rates, less than 100, and by organizing the lists of events by rate, one need only locate the appropriate list after which one can randomly select an event from it rather than search for the correct interval. This is only possible, however, if one maintains a second list, indexed by surface location, that tells you where each event’s rate is located in the rate lists [16]. Without this second list, a search will be required when you

update the rate lists after executing the local change in configuration. Alternately, one can use a single list and a binary search [17], but it is faster to maintain the cross-indexed lists.

For large simulations, memory can become limiting if one declares the event list for each rate to be long enough to hold all of the events, which is the easiest thing to do if memory is not an issue. If necessary, all of the events can be efficiently packed into a linear array of length M , however, provided one does a little extra shuffling after executing each change in configuration. To do this, partition the list as $\{e_{11}, e_{12}, \dots, e_{21}, \dots, e_{n_r, 1}, \dots\}$ into the appropriate n_r sublists which contain events that share common rates, keeping track of how many events correspond to each distinct rate. When a particular event—defined by a location and direction to hop in—must be moved from one rate category to another, move it to the end of its new category, fill the resulting gap by replacing it with the last event in the original category and repack the list by moving only the last event on all of the sublists that lie between the two affected categories from one end to the other. Though somewhat complicated, the cost of this procedure again scales with the number of distinct rates rather than with the (much larger) number of events in the simulation.

For some KMC schemes, including the hybrid scheme described here, some or all of the rates may not fall into discrete categories. In this case, it is possible to use a combination of the acceptance–rejection technique [18] and the cross-indexing described above. For the present application, the percentage of events falling outside the finite set of distinct rates is too small to bother with this, but for the sake of generality we describe the procedure. Acceptance–rejection works by finding an upper bound r^* , preferably sharp, on the set of rates. One then chooses a random number $r \in [0, M*r^*)$, with an interval of width r^* corresponding to each event. This allows one to identify the interval corresponding to a particular event without searching and either accept or reject the change in configuration, depending upon whether r falls within the correct portion of this interval. On average, the fraction of accepted events is R/Mr^* . For most KMC simulations this is highly

inefficient due to the wide range of rates—which is why the BKL algorithm is normally preferred for this type of simulation. The cross-indexed algorithm essentially uses the acceptance–rejection technique on sublists that contain only a single rate so that the efficiency becomes 100%. If the rates do not naturally fall into a small number of distinct categories, one can artificially sort them into categories based on order of magnitude and use acceptance–rejection on these sublists.

The hybrid scheme can be implemented using any of the numerous KMC models in the literature. In all of these schemes, the rates for the various nearest-neighbor processes are calculated from the Boltzmann factors

$$r_m = K \exp\left(-\frac{\Delta\phi_m}{k_B T}\right),$$

where K is a temperature-dependent attempt frequency, k_B is Boltzmann's constant, T is the substrate temperature and $\Delta\phi_m$ is the energy barrier that the m th event must overcome. It is the modeling of this energy landscape that accounts for the differences in various KMC models. The simulations presented here use a model due to Smilauer and Vvdensky [12] where the rate for each hopping event is determined by a bond-counting formula which is a function of the in-plane, lateral, nearest neighbors $\mu \in \{0, 1, 2, 3, 4\}$ and the difference between the remaining lateral, nearest neighbors (i.e. in the planes above and below the site being considered) before and after an event where an atom moves one lateral site $\nu \in \{0, 1, 2, \dots, 8\}$. As explained in Ref. [12], this latter number ν is set to zero if it is not positive in order to model the influence of the Schwoebel effect—the tendency for atoms to avoid sites near the top of steps:

$$\Delta\phi_m = E_0 + \mu E_e + \nu E_s,$$

where E_0 , E_e and E_s are the substrate, edge-barrier and Schwoebel barrier contributions to the energy, respectively. Of the 45 possible rates described by this formula, 10 may be discarded because there are no geometric configurations corresponding to those parameters (i.e. $\nu \in \{0, 1, \dots, 4 + \mu\}$). It is best to pre-compute and store these rates.




Fig. 2. Schematic diagram indicating the numerous KMC events (small hash-marks) that produce small random time intervals that accumulate to give the macro-time step Δt used in Eq. (2).

Note that each time a KMC event occurs, a small random time-step is taken. It is inefficient to step the discretized diffusion equation (2) by these small amounts. Instead, time is accumulated until a threshold is met, at which point the coarse-grained adatom density is stepped by the accumulated time Δt . Thus there are a large number of KMC events which occur over random sub-intervals of each macro-timestep, as illustrated in Fig. 2. At the beginning of a simulation that starts from a step-train devoid of adatoms, the global sum of rates R will be much less than normal. This can result in the individual KMC time steps exceeding the CFL limit during the early portion of the simulation and should be guarded against by checking for this and subdividing the macro-timestep if necessary. After a brief transient period, this situation should occur rarely, if at all.

2.3. The interface between KMC and continuum regions

The interface between the KMC and BCF domains presents a number of issues concerning how to transfer mass between the regions and how to update the partitioning of the cells.

The transfer of mass from the KMC into the continuum region is relatively straightforward. If an event results in an atom hopping into the continuum region, the storage location of the correct cell can be identified by using the inverse cell-list described above. The atom is then removed from the site it is hopping from, as usual, while its mass is added to the density of the cell it hops into rather than adding one to the height of the site

onto which it would otherwise hop. Note that the adatom density is updated immediately, rather than waiting until the end of the macro-timestep.

To allow mass to flow in the other direction, events are added to the KMC event list that correspond to atoms hopping from continuum cells into KMC cells. These are examples of events, discussed above, with rates that do not fall into the discrete set of rate categories. To handle these events, particularly for $(2 + 1)$ -dimensional simulations, it is necessary to maintain a list of boundary segments that divide the KMC and BCF regions. This list is in addition to the cell-lists described above. Each segment on this list represents an additional event in the KMC process. The rates for these extra events are proportional to the diffusion constant and the adatom density in the cell from which the adatom is hopping.

$$r_{\text{hop}} = \alpha D \rho. \quad (4)$$

In Ref. [7], the constant α was set to one. Here, we have found that it is much better to adjust this constant to approximate the correct equilibrium adatom density at the boundary of the continuum region. In the calculations below, α was adjusted so that the average adatom density was close to zero, mimicking the homogeneous Dirichlet condition supplied at the steps in some implementations of the BCF model. One could do somewhat better by estimating the density $1\frac{1}{2}$ cell-spacings away from the step by using the BCF model for a uniform step-train. The value of α needed to achieve this density can then be determined iteratively as the simulation proceeds or via a few trial runs that are terminated as soon as the density at the boundary has settled into a meaningful average. We refer to the constant α as the sampling rate and discuss the undesirable consequences of it being too high or too low in the next section.

When one of the boundary events is selected, one atom's mass is removed from the appropriate continuum cell and deposited at a randomly chosen site along the adjoining KMC boundary. Here, one must address the issue that the KMC simulation is discrete while the continuum is not. Thus, it may be the case that these events are

selected when the adatom density is less than one. In fact, this is nearly always the case unless one artificially prevents atoms from leaving a cell until a threshold density is reached. In the earlier paper [7], the preferred approach to solving this problem was to make up the mass deficit by searching neighboring cells. This slows the simulation down, however, and is much more complicated for $(2 + 1)$ -dimensional simulations. A second approach was to only allow this type of event to occur if the adatom density was at least one. For most operating conditions, this has the undesirable effect of retaining more atoms in the terrace region than would be the case in a normal KMC simulation. In the present implementation, we adopt a third approach that seems to improve upon these earlier ideas. Here, we remove the mass locally with no search of neighboring cells, but allow the adatom density to become negative rather than establishing a threshold that avoids this possibility. When this happens, the hopping rate for that cell is set to zero until the cell is naturally replenished by the deposition and diffusion process. In effect, the threshold for an active cell is now zero, rather than one adatom.

When the accumulated time steps for the various KMC events approach the CFL constraint (3), adatom densities are updated using (2) along with a homogeneous Neumann boundary condition,

$$\hat{\mathbf{n}} \cdot \nabla \rho = 0,$$

that prevents further mass transfer between KMC and BCF regions. This is followed by a reevaluation of which region each of the cells should be in. As KMC events occur, a variable that counts the number of “interesting” sites in each cell is maintained so that one can easily identify edge-cells and determine which cells need to change their affiliation. This is done on the slower time scale of the diffusion process, which dictates the maximum rate at which steps can flow from one cell into the next. If a cell is to be moved into the diffusion region, the process is straightforward: remove the nonuniformities in height (i.e. the adatoms) and all of the events corresponding to that cell from the KMC event lists, add the appropriate amount of mass to the adatom density

for that cell and update the cell- and boundary-list variables.

The opposite conversion, from a continuum to a KMC cell, is again complicated by the fact that the KMC simulation is discrete. If there should be enough mass for one or more adatoms, these are deposited randomly in the new KMC cell. Next, events are added to the KMC event-lists corresponding to all of the sites in the cell. Usually there will not be enough mass to generate adatoms, but there will be a fractional amount of mass left over. Here, we follow the practice of the earlier paper, which was to leave this residual mass in the cell for the next time the cell comes into a continuum region. While this conserves mass, it can lead to anomalous transfer of mass across KMC regions and, hence, over steps. In the present implementation, this problem is avoided by a secondary benefit of choosing the rate constant α so that the adatom density near the cell boundaries is zero on average.

3. Step flow instability

It is a well-known fact that the attachment of adatoms to the steps occurs asymmetrically and that normally attachment from the lower terrace is preferred. When the effect is pronounced, there are at least two instabilities of the evolving film surface that can occur. One of these, not examined here, features the nucleation of new islands, one on top of the other, due too the increased adatom density that is the result of adatoms being trapped on the highest exposed layers of growth. This is often referred to as the “wedding cake” instability. The second instability that can be caused by a large Schwoebel barrier leads to the formation of Taylor–Saffman-like ripples along a step that can subsequently evolve into crevices, pinch off to form islands or even lead to what has been called a “surface dendrite”. The mechanism behind this instability is easily explained in the context of the BCF model, with perturbations of the steps resulting in an enhanced flux of atoms to portions that protrude into the lower terrace as a result of the locally enhanced surface-to-step-length ratio on the preferred side of the step. If attachment

from the top, rather than the bottom, of a step is preferred, there can be a step-bunching instability. In this case, a perturbation that leaves one terrace narrower than its neighbors will tend to narrow further as the step in front of it suffers from a decreased supply of adatoms. The study of such step-flow phenomena is an ideal task for the hybrid scheme, as the topic is technologically relevant but often difficult to simulate using other techniques.

To illustrate the hybrid scheme, we focus on the step-flow instability that is the result of preferred attachment from the lower terrace. This instability has been observed experimentally and analyzed extensively using the BCF model beginning with the work of Bales and Zangwill [19]. In contrast, the simulation of this instability by KMC has been limited by the extreme computational cost associated with widely separated steps. In one example of work in this area, Rost et al. [11] present the results of KMC simulations that reveal some of the rich nonlinear phenomena seen in experiments, with the necking of ripples and subsequent formation of islands. To complete the simulations in a reasonable amount of time, the authors limited themselves to an average step spacing around five sites, greatly reducing the equilibrium adatom density and the number of hops that each atom must take to reach a step. This study illustrates the potential benefits of using a faster method, like the hybrid scheme presented here, to explore the strongly nonlinear behavior that can occur during growth by step-flow. In particular, this is a situation where the nano-scale surface structures one would like to study are probably too small for continuum models of the BCF type.

The rest of this section uses the step-flow scenario to demonstrate the hybrid scheme. The aim is to make direct comparison with a corresponding KMC simulation rather than present a detailed exploration of step-flow using the hybrid scheme. The calculations focus on 300×300 site sections of film with periodic boundary conditions for the adatom density. The steps themselves adhere to periodic boundary conditions in one direction and form an “Escher-”staircase in the other direction. In contrast to the earlier study of Rost, et al. [11] we focus on widely spaced steps.

In Figs. 3–5, we show a surface where three steps separated on average by 100 sites have evolved through 20 layers of growth. Fig. 3 was generated by a KMC simulation and required about 300 h (over 12 days) of CPU time on a 2 GHz Pentium IV processor. Fig. 4 was generated using the hybrid scheme in 56 h using the same computer hardware. This simulation used cells which were 10 sites on each side. The result shown in Fig. 5 used a cell width of 15 and took about 70 h. Thus, the first hybrid simulation was just over five times faster than KMC and the second was a little over four times faster. As explained in Ref. [7] the computational cost of the hybrid scheme is dominated by the KMC portion of the domain, so that the improvement in performance scales with the reduction in the KMC portion of

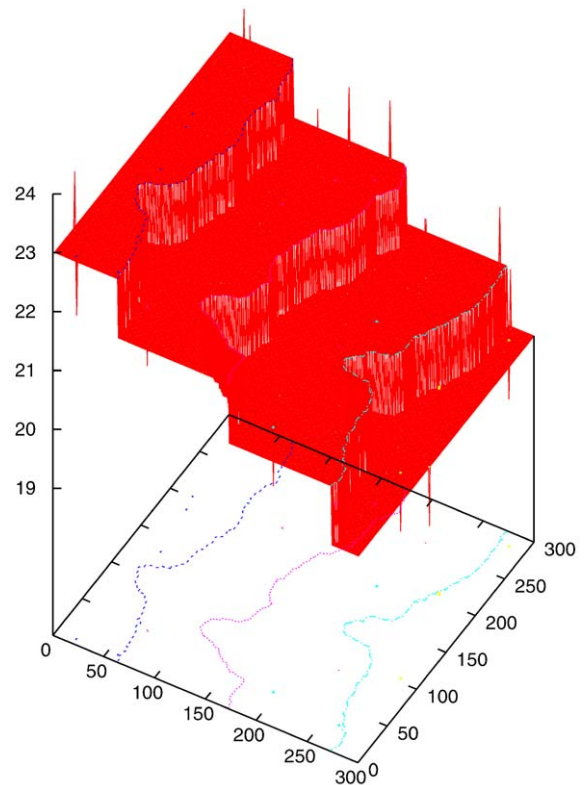


Fig. 3. A KMC-generated surface showing a 300×300 site region with three steps and the cyclic boundary conditions described in the text. After about twenty layers of growth this simulation reveals a fingering instability with highly irregular steps that appear to be forming a deep crevice.




Fig. 4. This image was generated by the hybrid scheme using parameters identical to those of the simulation shown in Fig. 3 along with 10×10 cells in the continuum region. While this simulation took far less time than the KMC simulation and clearly exhibits a step-flow instability, the step has a smoother appearance and has not formed the crevices that appear along the steps in Fig. 3.




Fig. 5. This figure shows the result of a hybrid simulation using continuum cells containing 15 sites on each side. While somewhat slower than the simulation shown in Fig. 4, it does a better job at capturing the level of noise in the corresponding KMC simulation.

the domain. As a result, the computational speed would see further improvements with smaller cell widths or with wider terraces.

To be useful for exploring surface phenomena, the hybrid scheme must not only be faster than KMC, but it must retain a significant amount of the atomistic detail that KMC enjoys over simulations based on the BCF model. The verdict here is somewhat mixed and further improvements of the scheme may be necessary for some applications. While it is significant that a simulation that uses only parameters that are drawn directly from the corresponding KMC model is able to reproduce the step-flow instability, it is also clear that the instability develops somewhat slower in the hybrid scheme and that the level of noise is

reduced. In particular, notice that the faster of the two hybrid simulations (shown in Fig. 4) has significantly smoother steps than the KMC simulation. While better, Fig. 5 also suffers from this defect to some extent. It is important to keep in mind that the BCF model alone produces deterministic variations in the interface shape that are perfectly smooth and uniformly spaced, requiring the addition of noise terms with further modeling and parameter estimation to do better. These enhancements are available within the context of the hybrid scheme as well, where it would presumably be easier to insert a single parameter correction within the continuum region to enhance the fluctuations in the rate at which adatoms are delivered to the KMC-regions.

It is worth noting that it is not especially easy to locate parameter values—choices for the the step- and Schwoebel-barriers along with the step-spacing—that exhibit the step-flow instability while avoiding the island mode of growth. Indeed, from the simulations that were performed, it would seem likely that islands might form with the values used here— $E_0 = 1.6$, $E_e = 0.3$, $E_s = 0.2$ and $L = 100$ —if the simulation were continued long enough. The long-time behavior is not necessarily relevant, however, since manufacturing processes ultimately aim for films of a certain prescribed thickness. Fortunately, the hybrid scheme could be used to search for these windows of parameter space without performing the corresponding KMC simulations.

We conclude this section with a discussion of the role of the sampling rate α in Eq. (4). If one arbitrarily sets $\alpha = 1$ as was done in Ref. [7], one finds that the adatom density at the boundary becomes overly depleted. If one considers ρ to be an estimate of the probability that an atom is present in the cell, then this can be understood by realizing that when the estimate is used repeatedly on successive KMC time-steps, the correlations between these successive samplings are neglected. This is very significant, because if an atom was not found in the cell at a given time, it is most likely that it would not be there a short time later. Thus the sampling rate must be reduced to compensate for this tendency. Failure to do so will result in too low of an adatom density in the continuum region and an artificial barrier to atoms entering the region from adjacent KMC cells. In the presence of a strong Schwoebel barrier, one way this error manifests itself is nucleation of islands near the top side of a step as atoms are trapped between the Schwoebel barrier and this artificial barrier introduced by too large a sampling rate. Conversely, if the sampling rate is too low, the adatom density will be artificially high. Recall that the fractional portion of the adatom density is being retained in the cell upon its conversion to a KMC-cell. This can combine with a low α value to shift the dominant source of adatoms supplying a step from the lower to the upper side—a shift that tends to favor step-bunching. In diagnosing whether α is too large or too small, it is important to realize

that the nucleation of islands near the top of the step also leads to bunched steps. The two situations can be hard to distinguish after these occasional backward-moving steps—the result of nucleation—merge with a forward-facing step. Fortunately, both of these problems are easily avoided by adopting the condition, described above, of adjusting the sampling rate so that the adatom density is zero or near-zero at the boundary between the domains.

4. Conclusion

This paper has continued the development of a hybrid scheme for simulating epitaxial growth by combining features of the BCF model and KMC simulations. This is the first implementation of the hybrid scheme for (2+1)-dimensional growth. Other improvements over an earlier version [7] include the use of a more conventional KMC model and some refinement in the handling of the boundary condition between the KMC and continuum regions. The code was applied to an unstable step-flow scenario with negligible nucleation effects and the results were compared to simulations that used only KMC. The results were extremely encouraging with respect to computational speed and revealed effects due to fluctuations along the steps to a much greater extent than the BCF model. Further improvements aimed at enhancing the accuracy of the technique with respect to pure KMC simulations may require modeling the various sources of noise in the spirit of work that has done in the BCF context. This should be easier in the case of the hybrid scheme because different sources of noise—from deposition, nucleation, detachment, inhomogeneities in the hopping rates, etc.—are easily isolated and treated separately.

The present method, with or without further enhancement, should be especially useful in scenarios with widely separated steps and high adatom densities—situations that can not be easily simulated with KMC due to increased computational cost. The method is extremely flexible and can be coupled interchangeably to any underlying KMC scheme so that complicated multi-species

interactions can be included. A nice feature of the method is that it will seamlessly revert to a KMC simulation if the surface becomes extremely complicated. In addition to the step-flow problem considered here, other ideal applications for this method include the late-stage coarsening of island mode growth, the deposition of film onto a continuously supplied substrate and diffusion limited aggregation.

Acknowledgements

The author would like to thank Weinan E and Peter Smereka for helpful discussions and acknowledge support from the National Science Foundation through grant DMS 0103825.

References

- [1] M. Kotrla, *Comput. Phys. Commun.* 97 (1996) 82.
- [2] W.K. Burton, N. Cabrera, F.C. Frank, *Philos. Trans. R. Soc. London.* 243A (1951) 299.
- [3] R.E. Caflisch, M.F. Gyure, B. Merriman, S.J. Osher, C. Ratsch, D.D. Vvedensky, J.J. Zinck, *Appl. Math. Lett.* 12 (1999) 13.
- [4] M.F. Gyure, C. Ratsch, B. Merriman, R.E. Caflisch, S.J. Osher, J.J. Zinck, D.D. Vvedensky, *Phys. Rev. E* 58 (1999) 6927.
- [5] P. Smereka, *Physica D* 138 (2000) 282.
- [6] R.E. Caflisch, W. E, M.F. Gyure, B. Merriman, C. Ratsch, *Phys. Rev. E* 59 (1999) 6879.
- [7] T.P. Schulze, P. Smereka, W. E, *J. Comput. Phys.* 189 (2003) 197.
- [8] G. Russo, L.M. Sander, P. Smereka, A Hybrid Monte Carlo Method for Surface Growth Simulations, pre-print.
- [9] W. E, B. Engquist, *Commun. Math. Sci.* 1 (2003) 87.
- [10] R.L. Schwoebel, *J. Appl. Phys.* 40 (1969) 614.
- [11] M. Rost, P. Smilauer, J. Krug, *Surf. Sci.* 369 (1996) 393.
- [12] P. Smilauer, D.D. Vvedensky, *Phys. Rev. B* 52 (1995) 14263.
- [13] T.P. Schulze, W. E, A continuum model for the growth of epitaxial thin films, *J. Crystal. Growth* 222 (2000) 414.
- [14] T.A. Witten, L.M. Sander, *Phys. Rev. Lett.* 47 (1981) 1400.
- [15] A.B. Bortz, M.H. Kalos, J.L. Lebowitz, *J. Comput. Phys.* 17 (1975) 10.
- [16] T. P. Schulze, *Phys. Rev. E* 65 (2002) 36704.
- [17] J.L. Blue, I. Beichl, F. Sullivan, *Phys. Rev. E* 51 (1995) 867.
- [18] W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Cambridge, 1986.
- [19] G.S. Bales, A. Zangwell, *Phys. Rev. B* 41 (1990) 5500.