

**Multiscale kinetic Monte Carlo algorithm for simulating epitaxial growth**

Jason P. DeVita\* and Leonard M. Sander†

*Physics Department, University of Michigan, Ann Arbor, Michigan 48109-1040, USA*

Peter Smereka‡

*Department of Mathematics, University of Michigan, Ann Arbor, Michigan 48109-1043, USA*

(Received 29 April 2005; published 15 November 2005)

We present a fast Monte Carlo algorithm for simulating epitaxial surface growth, based on the continuous-time Monte Carlo algorithm of Bortz, Kalos, and Lebowitz. When simulating realistic growth regimes, much computational time is consumed by the relatively fast dynamics of the adatoms. Continuum and continuum-discrete hybrid methods have been developed to approach this issue; however, in many situations, the density of adatoms is too low to efficiently and accurately simulate as a continuum. To solve the problem of fast adatom dynamics, we allow adatoms to take larger steps, effectively reducing the number of transitions required. We achieve nearly a factor of ten speed up, for growth at moderate temperatures and large  $D/F$ .

DOI: [10.1103/PhysRevB.72.205421](https://doi.org/10.1103/PhysRevB.72.205421)

PACS number(s): 68.55.Jk, 68.35.Fx, 81.15.Aa

**I. INTRODUCTION**

Molecular beam epitaxy (MBE) is a popular technique for growing materials, and is also an interesting example of a nonequilibrium statistical process. A simplified view of MBE, which captures many of the salient features, is the following. Atoms arrive at a crystal surface at a rate  $F$ . Once on the surface, the atoms diffuse with diffusion constant  $D$  until they attach to an existing island or nucleate a new island. Atoms detach from island edges with rate  $w_n = D \exp(-n\varepsilon)$  where  $n$  is the number of in-plane nearest neighbors and  $\varepsilon$  is the ratio of effective interatomic bond energy to the thermal energy ( $k_B T$ ). Over time, adatoms nucleate new islands, islands grow by absorbing adatoms, and the surface grows layer by layer.

A common technique for modeling epitaxial growth is kinetic Monte Carlo (KMC). In KMC, atoms are treated as individual particles which diffuse across flat substrates by taking nearest neighbor hops at a rate  $4D$ . Atoms that are attached to islands execute nearest neighbor hops with a reduced rate ( $w_n$ ). By representing each atom individually, KMC automatically incorporates internal noise. In many cases, KMC is a very effective tool for simulating epitaxial growth, as it retains its accuracy and flexibility through a wide range of growth environments. However, this dependable accuracy and flexibility sometimes comes at the expense of computational speed. For example, when the adatom density becomes large (such as at high temperature) or when the adatom diffusion rate is much greater than the island-atom detachment rate, a KMC simulation spends much of its computational effort computing adatom dynamics. To compensate for the adatom bottleneck, we devised a new algorithm for simulating epitaxial growth, which is faster than KMC, while retaining KMC's accuracy and flexibility. Our method is a KMC approach, but we modify the microscopic dynamics of the adatoms to increase efficiency, while retaining accurate mesoscopic and macroscopic dynamics. The essential step in the technique is to allow the adatoms to hop many

steps at once. We show that this modified dynamics reproduces the true dynamics in the regime of interest.

**II. KINETIC MONTE CARLO ALGORITHM**

The most basic KMC algorithm used to simulate epitaxial growth, which we will refer to as rejection KMC, goes as follows. Consider an  $L \times L$  grid, choose a grid point at random, and calculate the hopping rate for the atom at that site [ $w_n = \exp(-n\varepsilon)$ ], where  $n$  is the number of in-plane nearest neighbor bonds. Since adatoms have the fastest hopping rates, we let our time step be the time it takes for an adatom to hop. Thus  $\Delta t = 1$ , and the probability for any atom to hop in one time step is  $w_n$ . We then choose a random number between 0 and 1, and if that number is less than  $w_n$ , the atom executes a nearest neighbor hop. To complete one time step, this is repeated  $L^2$  times—enough so that on average every particle is given one chance to move.

This algorithm is simple enough to be useful for a wide variety of problems in epitaxial growth, including heteroepitaxy. However, it can be quite slow. Since the hopping rates differ exponentially, many moves are outright rejected. To avoid this bottleneck, we use the Bortz, Kalos, and Lebowitz (BKL) algorithm,<sup>1</sup> which samples according to the appropriate rates without rejections. The system we are simulating has a very simple rate structure. There are only five possible transitions in basic homoepitaxy, corresponding to the 0 to 4 possible in-plane bonds. Thus we can use the BKL algorithm—often referred to as continuous-time KMC. We keep lists of atoms with each bond count—a list of adatoms, a list of singly bonded atoms, and list of doubly bonded atoms, and so on. One loop of the algorithm involves exactly one transition. The hopping rate for each type of atom is  $w_n = \exp(-n\varepsilon)$ . The total transition rate for the system is then  $W = \sum_{j=0}^4 w_j N_j$  where  $N_0$  is the number of adatoms,  $N_1$  the number of singly bonded atoms, and so on. The mean waiting time for a transition to occur is  $1/W$ . So each loop of the algorithm advances time by  $1/W$  (which varies, since  $N_j$  var-

ies). One of the five types of atoms is then chosen with probability

$$P_j = \frac{w_j N_j}{W} \quad (1)$$

and a random member of that set of atoms is hopped. In a system with a few fast moving atoms and many slow moving ones, BKL runs much faster than rejection KMC, and so has become the standard for homoepitaxy simulations. It is worth noting that the ideas presented in this paper are also applicable to rejection KMC, even though we focus on BKL.

### III. ADATOM DIFFUSION

As mentioned above, there are five rates in our problem. The fastest is adatom hopping ( $w_0=1$ ), and the next fastest is hopping of singly bonded atoms [exponentially slower at  $w_1=\exp(-\varepsilon)$ ]. So if  $\varepsilon=5$  (about 600 K for Cu for example), then about 148 adatom diffusion events occur for every singly bonded edge-atom hop—and about 22 000 adatom events for every doubly bonded edge atom hop. Thus, in regimes where adatom diffusion plays a significant role, much of the computation is spent calculating the adatom trajectories.

Our goal is to reduce the disparity in the transition rates. We do this by altering what we consider to be an adatom event. Normally, an adatom transition involves one nearest neighbor hop; instead, we allow an adatom to diffuse for  $n_d$  steps. In other words, when we choose an adatom, we allow it to execute  $n_d$  random nearest neighbor hops before moving on to the next atom. This new  $n_d$ -step adatom transition must occur at a rate  $w'_0=1/n_d$ . The total transition rate is the sum of the rates for all the possible transitions

$$W = \sum_{j=0}^4 N_j w_j = N_0 + \sum_{j=1}^4 N_j w_j \quad (2)$$

which becomes

$$W' = \frac{N_0}{n_d} + \sum_{j=1}^4 N_j w_j. \quad (3)$$

Since the  $w_j$ s are Boltzmann factors, the total rate is dominated by the adatom term (when the temperature is low enough). Thus  $W'$  is often much smaller than  $W$ .

Though the total transition rate has been greatly reduced, the included events have become more complex and time consuming (an  $n_d$ -step random walk versus a one-step nearest neighbor hop). Some increase in efficiency is realized through this rate reduction technique, but there is more to be gained. For the growth conditions considered within this paper (high  $D/F$ ), the step density is often quite low, and many adatoms are far from an step. Given an adatom on an open terrace (no step edges or other adatoms around) the subsequent motion is known *a priori* (in a probabilistic sense). If it is known that an adatom is far enough from any obstacles, then the adatom can “complete” an  $n$ -step random walk in one step. The difficulty lies in knowing how much room the adatom has. A KMC method that includes long steps has

been developed in one dimension (1D)<sup>2</sup>; however, tracking distances in two dimensions is significantly more challenging.

What is needed is a way of knowing the distance to the nearest attachment site (a site with a free in-plane bond). Several methods have been developed to solve this issue in various related problems, such as diffusion limited aggregation (DLA).<sup>3</sup> A fast method for doing this that has been suggested<sup>4</sup> is to use a hierarchy of coarse grained grids, where the depth of the hierarchy is variable, and determined by the distance from an edge. This works nicely in the context of DLA, where the growth is irreversible, there is no nucleation, and there is only one random walker at a time. But when there are multiple random walkers, detachment, edge diffusion, and nucleation/breakup of dimers, the map hierarchy becomes too cumbersome to be efficient. Methods have also been developed to directly compute the distance to an edge.<sup>5-7</sup> Such methods are very fast when computing the distance globally; they are too slow, however, when only a local update is needed.

Here, we take a different approach to examining an adatom's local environment. We search locally. An  $n$ -step random walk on square grid potentially covers an area of  $2n(n+1)+1$  (a diamond of diagonal length  $2n+1$ ). In other words, an adatom executing a 25-step random walk can end up, or pass through, any of 1301 sites (all the sites 25 or fewer hops away). To simulate a 25-step random walk, we can compute (*a priori*) the adatom probability density over the 1301 site diamond, and then choose the final position of the adatom by sampling randomly from the probability distribution. This method is only accurate if the 1301 site diamond is void of obstructions (steps, attachment sites, and other adatoms). Otherwise there would be a sink for the adatom probability density, and the distribution would differ from the *a priori* computed distribution. So to replace a 25-step random walk with a single hop requires searching 1301 sites for obstructions, which is computationally cumbersome.

A simple approximation can reduce the number of sites to be searched. Consider that the probability density is very low at the perimeter of the 1301-site diamond. In fact, for an  $n$ -step random walk, approximately 98% of the probability density is contained within a circle of radius  $2\sqrt{n}$ . As a first approximation, we can simulate the 25-step random walk by searching an area of about 314 sites, and then sample from the truncated probability distribution. That would be faster than searching all 1301 sites, but still too slow. We could continue to truncate the distribution closer and closer to the origin, but the approximation becomes less and less accurate.

Ideally, we would sample from a distribution that is as small as possible, yet one that provides for accurate dynamics. We start by recognizing that the probability of finding a particle at  $(i, j)$  after a 1-step random walk is

$$P_{i,j}^{t+1} = \frac{1}{4}(P_{i-1,j}^t + P_{i+1,j}^t + P_{i,j+1}^t + P_{i,j-1}^t) \quad (4)$$

which becomes the diffusion equation

$$\frac{\partial p}{\partial t} = \frac{1}{4} \frac{\partial^2 p}{\partial x^2} \quad (5)$$

in the continuum limit. Given the approximation

$$\frac{\partial^2 p}{\partial x^2} \approx \frac{1}{m^2} (p_{i+m,j} - 2p_{i,j} + p_{i-m,j}), \quad (6)$$

we can replace Eq. (4) with

$$p_{i,j}^{t+m^2} = \frac{1}{4} (p_{i-m,j}^t + p_{i+m,j}^t + p_{i,j+m}^t + p_{i,j-m}^t). \quad (7)$$

and retain the same continuum limit. Essentially, we are sampling from a distribution that contains only four points  $(0, \pm m), (\pm m, 0)$ , by allowing an adatom to hop a distance  $m$  in one of four directions. To make sure that the long-hopping adatom does not skip over anything, we need to search the square that extends out  $(m-1)$  spaces  $-(2m-1)^2$  sites. To simulate a 25-step random walk, we search a square containing 81 sites, and then hop 5 spaces in one of four directions. Since each hop of a random walk requires the generation of a pseudorandom number, an 81-site search can be done faster than a 25-step random walk.

Our algorithm is like the standard continuous time algorithm, but with the following changes. A step distance ( $m$ ) is chosen, so that the associated time step ( $m^2$ ) does not exceed the ratio of the adatom hopping rate to the next fastest process (usually singly-bonded atoms). All the rates  $w_1 \rightarrow w_4$  are unchanged, but the adatom transition rate  $w_0$  becomes  $1/m^2$ . Then, if an adatom is chosen, we search the box around the adatom to find out how far it can hop without hitting something (up to a maximum of  $m$ ). If it has room to hop a distance of  $m$ , then it does, and we move on to the next transition. If it has room only for a hop of distance  $s < m$ , then it does so, but still has  $m^2 - s^2$  time left. So it repeats the process, this time with a maximum step size of  $\sqrt{m^2 - s^2}$ , until it uses up the full  $m^2$  time-step. That way all the adatoms diffuse for the same length of time, but the ones near the edges (or other adatoms) take several smaller steps instead of one large one.

Essentially we have replaced the typical adatom transition (one nearest neighbor hop) with a rescaled adatom transition ( $m^2$  hops), so that the adatoms execute transitions with a rate nearer to that of singly bonded atoms. Furthermore, we can do the same to the one-bond atoms, since they make transitions at a rate much faster than the doubly bonded atoms. The general idea is to have each type of atom operate on its own time scale, so that the rates become equalized. Of course, for this specific problem, there is little reason to go beyond singly bonded atoms, so we stop there. It is worth noting that if we were to have enhanced edge diffusion, this multiscale approach would provide additional speed up by also rescaling edge diffusion.

#### IV. THE ALGORITHM

We now assemble the ideas presented above into an algorithm for growing on a substrate. As in continuous time KMC, each loop of the algorithm is responsible for making

one transition. The time is then updated according to the total transition rate, and the loop is run until we reach the finishing time.

(1) Compute the rates associated with each transition. Let the number of atoms with  $k$  bonds be  $N_k$ ; then the total transition rate is

$$W = \frac{N_0}{m^2} + \frac{N_1}{q} e^{-\varepsilon} + \sum_{k=2}^4 N_k e^{-k\varepsilon} + F, \quad (8)$$

where  $m$  is the hopping distance for an adatom (and  $m^2$  the rescaled time step) and  $q$  is the rescaled time step for singly bonded atoms.

(2) Advance time by  $1/W$ .

(3) Choose a type of atom to move. Do this by generating a pseudorandom number  $r$  between 0 and  $W$ . If  $r < F$  then we add flux (one randomly placed atom); if  $F \leq r < N_0 w_0$  do an adatom event, and so on. The six possible events are the five atom hops (adatom, singly bonded, etc.) and deposition.

(4) Update the lists of each type of atom, and then restart the loop.

Four of the six possible events are identical to continuous-time KMC (adding flux, and hoping atoms with two, three, or four bonds). The algorithm for an adatom event is as follows:

(a) Choose uniformly from the list of adatoms.

(b) Scan a square region of side length  $2m-1$ . Throughout the simulation we keep track of the number of bonds available to attach to at a given site. Any site with more than zero attachment bonds is a potential attachment site. So, we start at the innermost square surrounding the adatom, and sum up the attachment bonds in all the sites in that ring. If that sum is zero, we move on to the next concentric square. We keep doing that until either the sum is nonzero or we reach the  $(m-1)$ th ring. Then we know how long our step can be, say  $s$ .

(c) We then take a step of size  $s$  in a randomly chosen direction. If  $s=m$ , then we are done; otherwise, we go back to step two, but with a maximum step of  $\sqrt{m^2 - s^2}$ . If at some point the adatom attaches, we then choose another adatom to finish out the time-step.

The singly bonded atoms act like the adatoms, except that they only take nearest neighbor hops, and only for  $q$  steps.

#### V. RESULTS

To arrive at our multiscale kinetic Monte Carlo (MSKMC) model, we have made some approximations regarding the adatom dynamics. While the approximations are physically reasonable, it is important to verify that the results produced are accurate. Since we have only modified the dynamics, any equilibrium results should be unaffected. We can confirm this by computing equilibrium island shapes. The shape of an island in equilibrium with the surrounding adatoms is known exactly by mapping onto the 2D Ising model.<sup>8</sup> We took as an initial condition a square island on a periodic grid, and allowed the island to come to equilibrium with the surrounding adatoms. The resulting island is shown in Fig. 1




FIG. 1. The equilibrium island shape for a large island (approximately 500 000 atoms) closely matches the exact (Wulff) solution (solid line). Because of the large size of the island, ensemble averaging is unnecessary. The temperature is  $E/kT=4$ .

with the exact (Wulff) solution superimposed. By using a large island ( $\sim 5 \times 10^5$  atoms, or about 250 nm for copper), we were able to achieve good agreement with the exact solution, without ensemble averaging. This demonstrates that our method preserves the original energetics.

Since we have modified the dynamics of the adatoms, it is important to verify that the new microscopic dynamics yield the correct mesoscopic and macroscopic results. A common difficulty among continuum and quasicontinuum methods is predicting the rate and location for nucleation of new islands. In models that treat that adatoms as a continuum field, the formation of new islands must be computed from a mean-field representation of the adatom density. In situations where fluctuations play a large role, a mean field treatment can give poor results. This is especially noticeable at early stages of growth, when nucleation, and subsequent breakup, occur very rapidly, and correlations in the adatom density can build up. MSKMC avoids these problems. To demonstrate this, we compare it to ordinary continuous time Monte Carlo. Both models are allowed to grow starting from a flat substrate, and the island sizes are tallied at various coverages. The results are shown in Fig. 2, where the island size histograms are plotted for coverages of 10%, 15%, and 20%. There is good agreement between the two models. This lends confidence in the accuracy of our new adatom dynamics. For this case, the MSKMC simulation ran 6.5 times faster than the standard BKL simulation.

We now test the long-time kinetics, by observing mounding. In the presence of an Ehrlich-Schwoebel (ES) barrier, a growing surface roughens and forms mounds. The ES barrier is a step-edge barrier that reduces the rate at which atoms ascend and descend steps. This tends to increase the adatom density on the island tops, thus increasing the nucleation rate on top of existing islands. This gives rise to a mounding

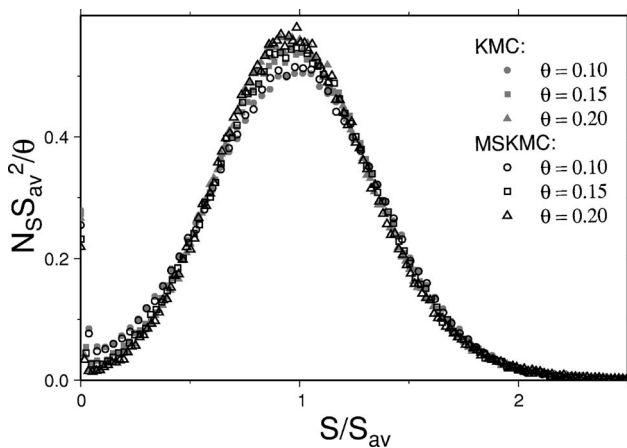


FIG. 2. The island size distribution for several coverages, showing good agreement between KMC and MSKMC.

instability, which has been observed in both simulation and experiment.<sup>9–11</sup> The degree of mounding and mound coarsening is sensitively dependent upon the nucleation on the island tops,<sup>12,13</sup> especially when the adatom density is low. This is one of the critical challenges of continuum based modeling; mean-field derived nucleation rates are often inadequate for predicting island-top nucleation.<sup>13,14</sup> The sensitivity of mounding on island-top nucleation makes this a challenging test for a method which modifies the adatom dynamics.

Rejections are used to implement the ES barrier. The ES barrier can be handled using the BKL framework; however, doing so greatly increases the number of lists required, and thus the complexity of the sorting. Since the ES barrier is typically small and only arises in certain configurations, the number of rejections is low and so rejection can be an efficient alternative to sorting lists. All the other barriers in the simulation are handled using BKL, and when an ES barrier is encountered, rejection is used. We grew, using both regular BKL and MSKMC, 1000 monolayers on a 1024 by 1024




FIG. 3. When a small Ehrlich-Schwoebel barrier is included, the oscillations in the roughness diminish over time as the total roughness increases. For this case,  $E/kT=5$ , the grid is 1024 by 1024,  $D/F=10^6$ , and the ES barrier is  $1kT$ . The difference between the two lines is a result of noise.




FIG. 4. Surface images at 500 and 1000 monolayers for growth with a small Ehrlich-Schwoebel barrier.

grid. Figure 3 shows the surface roughness as a function of time. The roughness starts off oscillating. Eventually the oscillations die down, and the overall roughness increases due to the mounding instability. The inefficiency of the BKL computation makes ensemble averaging over multiple realizations prohibitive; however, the difference between the two results is consistent with observed fluctuations. Snapshots of the surface at 500 and 1000 monolayers are shown in Fig. 4, for both KMC and MSKMC. Our MSKMC method ran seven times faster than the standard BKL method. Our method took about 24 hours to run the simulation on a 2.8 GHz Pentium desktop computer.

Mounding can also occur in the absence of an ES barrier, however, it can take a long time to happen. Monte Carlo simulations have shown this effect,<sup>15</sup> but only when there is enhanced edge diffusion. The increased efficiency of the MSKMC method allows us to simulate larger grid sizes and longer times. In Fig. 5 we show the result of growing 1000 monolayers on a 2048 by 2048 grid, with no ES barrier and no enhanced edge diffusion. We are able to observe decay in the step-density oscillations, and also observe the formation of interesting large scale features. This simulation took a little less than two days on a 2.8 GHz Pentium desktop computer.

Epitaxial growth on a vicinal surface is another problem of recent interest. In the presence of an ES barrier, step-flow becomes unstable and the steps can meander considerably.<sup>16</sup> This effect has been demonstrated computationally and experimentally.<sup>17–21</sup> Some systems exhibit meandering wavelengths which are quite large. For example, Hibino *et al.*<sup>20</sup> show experiments involving step-flow on Si, where the meandering wavelength is several microns. Simulating such a





FIG. 5. Even in the absence of an ES barrier, the surface begins to roughen and interesting morphology can arise.

system requires a very large computational grid. Coupled with the low temperature and large  $D/F$  typical of many of these experiments, KMC can be quite slow. Our MSKMC method is well-suited to computations involving large scale features, low  $T$ , and high  $D/F$ .


Figure 6 shows an example of meandering step flow. Both KMC and MSKMC are shown for comparison. The computation was done on a 1000 by 1000 grid, with an initial step spacing of 20 atoms. The pattern that forms is very similar to patterns seen for growth on vicinal Cu.<sup>21</sup> The MSKMC method ran nine times faster than the standard BKL method. For smaller miscut (large step spacing) and lower temperature, the speed advantage would be even greater.

## VI. DISCUSSION


There have been several distinct attempts to increase the computational efficiency of epitaxy simulation. Here we will discuss a couple of these other approaches.

### A. Absorbing Markov chains

There are some similarities between our MSKMC method and Novotny's Monte Carlo with absorbing Markov chains (MCAMC) method.<sup>22</sup> Both methods desire to improve computational efficiency by rescaling the fast dynamics; how-



(a) KMC



(b) MSKMC

FIG. 6. Starting from a uniform train of steps, a step-edge barrier causes the steps to meander. The step spacing is 20 atoms,  $E/kT=5$ ,  $E_{ES}/kT=3$ , and  $D/F=10^6$ .

ever, there are significant differences. MCAMC follows the exact dynamics in a manner that waits for significant changes in the current system state before executing a transition. It requires knowledge of the transition rates for not only the current state, but also for a number of accessible states. There are some situations for which such knowledge is efficiently obtained, and as Novotny demonstrates, significant speed-up can be seen. Unfortunately, for a general epitaxy simulation, it is simply too expensive to compute the transition rates for nearby states. It is therefore necessary to carefully approximate the microscopic kinetics in a way that maintains the correct mesoscopic and macroscopic dynamics, which is what we have done.

### B. Continuum models

Continuum models, which treat the adatoms as a continuous fluid, have been suggested as an alternative to KMC.<sup>23,24</sup>

Since the adatoms are represented by a continuous density, the computational speed is independent of the number of adatoms involved. Continuum models with deterministic attachment and detachment have had substantial success, but their usefulness is limited to growth situations where fluctuations are unimportant. One key source of fluctuations is shot noise. Shot noise arises from the fact that atoms attach one at a time, as whole particles. Since diffusion controlled growth is inherently unstable, noise can play a vital role in determining the shape of islands. An extreme example is the diffusion limited aggregation (DLA) model of Witten and Sander.<sup>3</sup> DLA is a diffusive growth model with no smoothing processes (such as edge diffusion or detachment), which creates branching fractal structures. Continuum models lack sufficient noise to generate DLA-like clusters.

There has been some effort to include shot noise fluctuations in a continuum model.<sup>25–29</sup> Most of these methods utilize a hybrid approach where adatoms are modeled as a continuous density, but islands are composed of discrete atoms. Once such approach is the quasicontinuum Monte Carlo (QCMC) approach.<sup>28,29</sup> In the QCMC approach, atoms are added to island edges as discrete particles, with location determined (probabilistically) by the local adatom flux into the island edge. This incorporates shot noise fluctuations into the continuum formalism. This approach has had many successes, including the ability to grow DLA-like clusters and reproduce correct thermal roughening results,<sup>28,29</sup> and to some extent model multilayer growth with mounding.<sup>29</sup>

However, these hybrid models still have some shortcomings. For one, they're not necessarily faster than KMC. In a continuum model, the adatom density is defined over the entire grid, even if there are only a few adatoms. If the adatom density is low, then a continuum method must take very long time steps when solving the diffusion equation in order to be computationally efficient. This can introduce error into the computation. Thus most continuum-based models are inefficient at low adatom densities. It is possible for such methods to be efficient at high densities; however, even that is challenging. Since the nucleation (and subsequent breakup) of new islands occurs at a higher rate as the adatom density increases, the allowable time-step becomes shorter as the density goes up. It is a very delicate matter to take a sufficiently long time step, while still obtaining acceptable results for nucleation.

Continuum-based methods bring difficulties beyond efficiency. Since the adatoms are replaced by a density, new islands must be explicitly created. Often that is done using a mean-field estimate of the nucleation rate as a function of density. There is a significant literature on the subject of nucleation in epitaxy (see Ref. 30 for a review) much of which uses a rate equation approach to calculate the dimer formation rate as a function of density, temperature, and flux. The results of such calculations are used by many continuum-based models. This appears to work reasonably well in the absence of any step-edge barriers. However, it has been shown<sup>13,14</sup> that when there is a step-edge barrier, the rate equation approach fails to predict the nucleation rate on island tops. Moreover, growth in the presence of a step-edge barrier can lead to a mounding instability that is very sensitively dependent on the island top nucleation rate.

Our aim in developing the MSKMC method is to efficiently model adatom processes (and other fast processes) without the problems associated with continuum methods. Continuum-based methods are very useful in situations where fluctuations are less important, and when one wants to closely parallel analytical approaches; however, when noise and discrete effects become important, a fully discrete method is often better suited.

## VII. CONCLUSIONS

We have presented an algorithm for simulating epitaxial growth, which is between five and ten times faster than continuous-time kinetic Monte Carlo for some interesting physical regimes. Our multiscale kinetic Monte Carlo (MSKMC) algorithm is a continuous-time method, which rescales the adatom dynamics in both space and time. For many realistic growth conditions, the adatoms consume most of the computational time. By allowing adatoms to take long steps, we reduce the transition rate of the dominant process.

It is worth noting that the relative efficiency of MSKMC (as compared to BKL) is dependent on the random number generator. Our method replaces random processes (random walks) by deterministic ones (local searching). Due to the computational complexity involved in generating a pseudo-

random number, an efficient search is faster than a random walk. To be fair, we chose to use the fastest reasonable quality random number generator that we were able to obtain (a variation of a lagged Fibonacci generator),<sup>31</sup> which is much faster than many compiler default generators. Were we to use a slower generator (perhaps to ensure better randomness), then our method would be even faster than we have stated, relative to BKL.

To verify the validity of MSKMC, we have computed equilibrium island shapes and early-growth island size distributions, both of which match expected results, analytical and computational. We were able to reproduce correct equilibrium island shapes without resorting to ensemble averaging. Our method also produces a step-edge barrier induced mounding that is consistent with standard KMC. Due to our increased efficiency, we can carry out simulations on larger grids and/or for longer times, and perhaps observe phenomena previously unseen.

## ACKNOWLEDGMENTS

P.S. acknowledges useful discussions with Michael Falk. This work is supported by the National Science Foundation. J.P.D. and P.S. were supported by NSF Grant Nos. DMS-0207402 and DMS-0509124. L.M.S. and P.S. were supported by NSF Grant No. DMS-0244419.

\*Electronic address: jdevita@umich.edu

†Electronic address: lsander@umich.edu

‡Electronic address: psmerka@umich.edu

- <sup>1</sup>A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, *J. Cell. Comp. Physiol.* **17**, 10 (1975).
- <sup>2</sup>C.-C. Chou and M. L. Falk (unpublished).
- <sup>3</sup>T. A. Witten and L. M. Sander, *Phys. Rev. Lett.* **47**, 1400 (1981).
- <sup>4</sup>R. C. Ball and R. M. Brady, *J. Phys. A* **18**, L809 (1985).
- <sup>5</sup>J. A. Sethian, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 1591 (1996).
- <sup>6</sup>D. Adalsteinsson and J. A. Sethian, *J. Comp. Physiol.* **118**, 269 (1995).
- <sup>7</sup>Y. H. R. Tsai, *J. Comp. Physiol.* **178**, 175 (2002).
- <sup>8</sup>C. Rottman and M. Wortis, *Phys. Rev. B* **24**, 6274 (1981).
- <sup>9</sup>M. D. Johnson, C. Orme, A. W. Hunt, D. Graff, J. Sudijono, L. M. Sander, and B. G. Orr, *Phys. Rev. Lett.* **72**, 116 (1994).
- <sup>10</sup>S. Clarke and D. D. Vvedensky, *J. Appl. Phys.* **63**, 2272 (1989).
- <sup>11</sup>J. Amar and F. Family, *Surf. Rev. Lett.* **5**, 851 (1998).
- <sup>12</sup>J. Krug and P. Kuhn, in *Atomistic Aspects of Epitaxial Growth* (Kluwer, MA, 2001), pp. 145–163.
- <sup>13</sup>J. Krug, P. Politi, and T. Michely, *Phys. Rev. B* **61**, 14037 (2000).
- <sup>14</sup>C. Castellano and P. Politi, *Phys. Rev. Lett.* **87**, 056102 (2001).
- <sup>15</sup>C. Ratsch, M. C. Wheeler, and M. F. Gyure, *Phys. Rev. B* **62**, 12636 (2000).
- <sup>16</sup>G. S. Bales and A. Zangwill, *Phys. Rev. B* **41**, 5500 (1989).
- <sup>17</sup>O. Pierre-Louis, *C. R. Phys.* **6**, 11 (2005).

- <sup>18</sup>J. Kallunki, J. Krug, and M. Kotrla, *Europhys. Lett.* **66**, 749 (2004).
- <sup>19</sup>R. Kato, M. Uwaha, and Y. Saito, *Surf. Sci.* **550**, 149 (2003).
- <sup>20</sup>H. Hibino, Y. Homma, M. Uwaha, and T. Ogino, *Surf. Sci.* **527**, L222 (2002).
- <sup>21</sup>N. Neel, T. Maroutian, L. Douillard, and H.-J. Ernst, *J. Phys.: Condens. Matter* **15**, S3227 (2003).
- <sup>22</sup>M. A. Novotny, *Phys. Rev. Lett.* **74**, 1 (1995).
- <sup>23</sup>M. Petersen, C. Ratsch, R. E. Caflisch, and A. Zangwill, *Phys. Rev. E* **64**, 061602 (2001).
- <sup>24</sup>C. Ratsch, M. F. Gyure, R. E. Caflisch, F. Gibou, M. Petersen, M. Kang, J. Garcia, and D. D. Vvedensky, *Phys. Rev. B* **65**, 195403 (2002).
- <sup>25</sup>L. Mandreoli, J. Neugebauer, R. Kunert, and E. Schöll, *Phys. Rev. B* **68**, 155429 (2003).
- <sup>26</sup>T. P. Schulze and W. E., *J. Cryst. Growth* **222**, 414 (2001).
- <sup>27</sup>T. P. Schulze, P. Smereka and W. E., *J. Comp. Physiol.* **189**, 197 (2003).
- <sup>28</sup>G. Russo, L. M. Sander, and P. Smereka, *Phys. Rev. B* **69**, 121406(R) (2004).
- <sup>29</sup>J. P. DeVita, L. M. Sander, and P. Smereka, in *Multiscale Modelling in Epitaxial Growth* (Oberwolfach, Germany 2004).
- <sup>30</sup>C. Ratsch and J. A. Venables, *J. Vac. Sci. Technol.* **21**, S96 (2003).
- <sup>31</sup>M. E. J. Newman (private communication).