Coarse bifurcation analysis of kinetic Monte Carlo simulations: A lattice-gas model with lateral interactions

Alexei G. Makeev[a)]
Department of Chemical Engineering, Princeton University, Princeton, New Jersey 08544

Dimitrios Maroudas[b)]
Department of Chemical Engineering, University of California, Santa Barbara, Santa Barbara, California 93106

Athanassios Z. Panagiotopoulos and Ioannis G. Kevrekidis[c)]
Department of Chemical Engineering, Princeton University, Princeton, New Jersey 08544

(Received 8 July 2002; accepted 14 August 2002)

We present a computer-assisted study of ''coarse'' stability/bifurcation calculations for kinetic Monte Carlo simulators using the so-called coarse timestepper approach presented in A. G. Makeev, D. Maroudas, and I. G. Kevrekidis, J. Chem. Phys. **116**, 10083 (2002). Our illustrative example is a model of a heterogeneous catalytic surface reaction with repulsive adsorbate–adsorbate interactions and fast diffusion. Through numerical continuation and stability analysis, we construct one- and two-parameter coarse bifurcation diagrams. We also discuss several computational issues that arise in the process, the most important of which is the ''lifting'' of coarse, macroscopic initial conditions (moments of adsorbate distributions) to fine, microscopic initial conditions (distributions conditioned on these moments). © *2002 American Institute of Physics.*
[DOI: 10.1063/1.1512274]

I. INTRODUCTION

A persistent feature of complex reaction and transport systems is the emergence of macroscopic, coherent behavior from the interactions of microscopic agents—molecules, cells, individuals in a population—between themselves and with their environment. The implication is that macroscopic rules (description of behavior at a high level) can, in some cases, be deduced from microscopic ones (description of behavior at a much finer level). For some problems, like Newtonian fluid mechanics, or simple homogeneous chemical kinetics, successful macroscopic descriptions (the Navier–Stokes equations, mass action laws) predated their microscopic derivation from kinetic theory. In many current problems, however, ranging from chemistry to engineering, and from ecology to materials science, the physics are known at the microscopic/individual level, and the closures required to translate them to an accurate, high-level, macroscopic description are simply not explicitly available. Severe computational limitations arise in trying to directly bridge the enormous gap between the scale of the available description and the scale at which the questions of interest are asked and the answers are required.

In this paper we study surface reactions for which the available description is at the level of microscopic lattice-gas models with lateral interactions, and the desired information encompasses the location, stability, and bifurcation analysis of macroscopic stationary states (expected average coverages, expected stationary reaction rates). Over the last few years we have been working towards the development of a two-tier computational methodology to address problems for which closed macroscopic equations conceptually exist *but are not explicitly available*. The inner component of this methodology is the best available microscopic simulator of the process: for this paper, a kinetic Monte Carlo (KMC) lattice-gas simulator. The outer component consists of a combination of (a) system identification and (b) traditional numerical analysis. Consider the computation, through a Newton–Raphson iteration, of the steady states of the macroscopic, coarse closed equations for average surface concentrations (coverages). If these macroscopic equations are explicitly available, the Newton–Raphson procedure requires repeated evaluations of the right-hand side and of its Jacobian (the first partial derivatives of the rate expressions with respect to coverages). If these equations are *not* explicitly available, we will show below how to use short bursts of KMC simulation (appropriately initialized, evolved and averaged) to identify exactly these numbers (as opposed to just evaluating them through a simple one-line formula). The second component of our computational methodology will in this case be precisely the Newton–Raphson algorithm that we would use if the macroscopic equations had been explicitly available. The only difference is that the quantities processed by this algorithm to give us the next steady state guess will not come from explicit function evaluations, but from system identification based on short local bursts of KMC simulation data. In effect, we will use the KMC evolution code as an experiment, and by perturbing its parameters and initial conditions we will identify ''on demand'' (''just in time'' is another term used in the literature) the quantities we need to do macroscopic numerical analysis

[a)]Permanent address: Moscow State University, Faculty of Computational Mathematics and Cybernetics (BMK) Moscow, 119899, Russia.
[b)]Present address: Department of Chemical Engineering, University of Massachusetts, Amherst, MA 01003.
[c)]Author to whom correspondence should be addressed.

with. And while Newton–Raphson is the simplest illustrative example, there exists an immense computational arsenal of tools (numerical bifurcation theory, optimization, control) that has been erected by current mathematics to help get information from models easier, faster, better than simple simulation.[1–6] We will return to the implications and possible impact of such a computational methodology in the Discussion (Secs. V and VI).

In a previous publication,[1] we were able to construct what we termed "coarse" bifurcation diagrams for the expected values of KMC simulations; our particular examples were simple models of surface reaction problems. The KMC technique used was the so-called "stochastic simulation algorithm" for well-mixed reaction systems, proposed by Gillespie,[7] which constitutes a coarse graining of the corresponding master equation. For this algorithm, exact closed mesoscopic equations for the expected values of the adsorbate distribution averages exist, and can be used to test the numerical results. In this paper, we show how the computational methodology can be extended to apply to lattice-KMC models with lateral adsorbate interactions. For these problems such exact mesoscopic equations are not available in closed form, although several levels of approximation (mean field, quasichemical, etc.) can be derived. We apply our coarse-timestepper-based procedure and compute numerical bifurcation diagrams, which we compare both to those produced through approximate explicit closures, and to typical long-term KMC simulations.

The paper is organized as follows: In Sec. II we discuss our illustrative example (a simplified model of heterogeneous catalytic CO oxidation with lateral adsorbate interactions) and our KMC simulation protocols. In Sec. III we outline our procedure along the lines presented in Ref. 1. Numerical bifurcation results are presented in Sec. IV as various parameters, such as the gas phase reactant pressures, the adsorbate–adsorbate interaction energies, and the temperature are varied. In Sec. V we discuss various computational issues, including the effect of diffusion on these computations. One of the most important issues we discuss in this section is the nonunique, "one-to-many" lifting operator, a vital part of the procedure: the construction of (ensembles of) *microscopic* initial conditions consistent with (alternatively, distributions conditioned on) given *macroscopic* initial conditions. Finally, in Sec. VI, we summarize and provide connections with literature using comparable methods to enable different, non-KMC types of microscopic simulators, such as lattice-Boltzmann kinetic-theory inspired ones.

II. THE ILLUSTRATIVE MODEL: KINETIC MONTE CARLO SIMULATIONS

We consider a set of the elementary reactions occurring on a perfect lattice with $N=N_1\times N_2$ adsorption sites and periodic boundary conditions. These reactions may involve one or two adsorption sites,

$$A_i \rightarrow B_i\,,$$

$$A_i + B_j \rightarrow C_i + D_j\,.$$

Here the index j runs over all the nearest neighbors of the site labeled "i" ($i=1,\ldots,N$). Each lattice site can be occupied by an adsorbed particle (say, of type A) or it can be empty $[A=(*)]$. The list of possible kinetic events may include the elementary events of migration (diffusion),

$$A_i + (*)_j \rightarrow (*)_i + A_j\,,$$

which describe the individual jumps of adsorbed particles to their neighboring vacant sites. A reaction mechanism is given by the set of all of its possible elementary reactions. The time evolution of the reaction system can be described by the chemical master equation,

$$dP_x(t)/dt = \sum_{x'} \left(W_{x'\rightarrow x}P_{x'}(t) - W_{x\rightarrow x'}P_x(t)\right),$$

where $P_x(t)$ $[P_{x'}(t)]$ denotes the probability of finding the system in configuration x $[x']$ at time t, and $W_{x\rightarrow x'}$ are used to denote transition probabilities per unit time for various elementary events such as adsorption, desorption, migration, reaction, etc. We also take into account the lateral interactions between the adsorbed particles over an interaction range up to second-nearest-neighbor separations. Due to lateral interactions, the rates of elementary reactions depend on the local environment, and because of these interactions an ordered adsorbed layer (microstructure) may form on the lattice. In general, the master equation cannot be solved directly; therefore, one has to either use some uncontrolled approximations in order to derive macroscopic evolution equations, or to implement KMC simulations which can provide, in principle, the correct solution of the problem.

The basic steps of the KMC algorithm we have implemented are as follows:[7–10]

(1) The time Δt that the system spends in a current configuration is $\Delta t = -\ln(\xi)/W$, where ξ is a random number taken from a uniform distribution on (0, 1); $W = \Sigma_{x'}(W_{x\rightarrow x'})$ is the total transition probability per unit time. (In principle, transition probabilities can be computed accurately through *ab initio* calculations of energy surfaces for both reaction and diffusion events in conjunction with implementation of rate theory.) Accordingly, the time is updated by an increment Δt: $t = t + \Delta t$.
(2) One reaction is selected from the set of all possible kinetic events, with a probability proportional to its rate. For this purpose, one more random number uniformly distributed on (0, 1) is generated. The selected kinetic event is performed and the set of all possible reaction rates is updated to take the local compositional changes on a lattice into account.
(3) Repeat the previous steps if $t<\tau$.

Initial conditions are given by the lattice configuration at $t=0$. This algorithm can be viewed as a "timestepper" that operates on the initial species coverage distribution on the lattice (input) for a given time period τ, and generates the final distribution (output). The number of all possible reactions on a lattice can be very large. Usually, there is no need

to recalculate all rates at each time step because of the local character of microscopic events occurring on a lattice.

Realistic description of diffusion is one of the main implementation problems for KMC simulations. In most practical cases, adspecies hop rates are many orders of magnitude greater than all other rates. In such a case, the adsorbed layer can be considered to be in a (quasi)equilibrium state at each time instant. Therefore, to accelerate the calculations, one can apply an additional procedure to produce an equilibrium distribution (for the current coverages), and remove migration events from the list of possible elementary events.[11] This separate relaxation procedure, which does not affect the average coverages, is executed after each time step. For this purpose, one can apply the classical Metropolis importance sampling algorithm for simulating species migration in the canonical ensemble: jumps of randomly selected adsorbed particles into the empty sites occur with a probability equal to 1 if $\Delta E \leqslant 0$, otherwise equal to $\exp(-\Delta E/RT)$; ΔE is the energy difference between the initial and final configurations, which is caused by the lateral interactions between adsorbed particles. Nearest neighbor (NN) jumps or long-range jumps can be considered. One Monte Carlo step (MCS) in the Metropolis algorithm corresponds, on average, to one attempted jump per lattice site. The number of these steps, N_{MCS} is a parameter of the algorithm. A summary of the various algorithm parameters and their meaning can be found in Appendix A 1.

In this paper we consider fast diffusion of adsorbates. The computed macroscopic steady-state solutions correspond to infinitely fast adspecies hop rates. In the simulations we used two KMC methods: for the first one the migration events are considered explicitly, i.e., the algorithm employs both diffusional transition probabilities and reaction probabilities; for the second one the Metropolis relaxation algorithm is used instead, along with reaction probabilities. The infinite mobility, in particular, allows us to suggest that the long-term, macroscopic system behavior can be described in terms of average coverages: higher order correlation functions are quickly slaved to (become functionals of) the coverages, that is, the zeroth moments of the species distributions on the lattice.

The lattice gas model: We consider the standard model of the $A+\frac{1}{2}B_2 \rightarrow AB$ reaction, which mimics the CO oxidation reaction occurring on a square lattice. Our lattice-gas (lattice-KMC) model contains 6 elementary steps involving two types of adsorbed species in the reaction mechanism:

(1) $\mathrm{CO_{gas}} + (*)_i \rightarrow \mathrm{CO_{ads},}_i$,

(2) $\mathrm{O_{2,gas}} + (*)_i + (*)_j \rightarrow \mathrm{O_{ads},}_i + \mathrm{O_{ads},}_j$,

(3) $\mathrm{CO_{ads},}_i \rightarrow \mathrm{CO_{gas}} + (*)_i$,

(4) $\mathrm{CO_{ads},}_i + \mathrm{O_{ads},}_j \rightarrow \mathrm{CO_{2,gas}} + (*)_i + (*)_j$,

(5) $\mathrm{CO_{ads},}_i + (*)_j \rightarrow (*)_i + \mathrm{CO_{ads},}_j$,

(6) $\mathrm{O_{ads},}_i + (*)_j \rightarrow (*)_i + \mathrm{O_{ads},}_j$.

Here, i and j are NN sites on a square lattice. Steps (5) and (6) describe the individual jumps of adsorbed particles to neighboring empty sites. In this paper, we consider only lateral interactions between $\mathrm{CO_{ads}}$ species at NN lattice sites, and assume no interactions for the activated complexes.

The standard mean-field (MF) equations for this model are

$$d\theta_1/dt = \alpha\theta_0 - \gamma\theta_1 \exp(4\theta_1\varepsilon/(RT)) - 4k_r\theta_1\theta_2 \exp(3\theta_1\varepsilon/(RT)), \tag{1a}$$

$$d\theta_2/dt = 4\beta\theta_0^2 - 4k_r\theta_1\theta_2 \exp(3\theta_1\varepsilon/(RT)), \tag{1b}$$

where $\theta_1 \equiv \theta_{\mathrm{CO}}$ ($\theta_2 \equiv \theta_{\mathrm{O}}$) is the CO (O) coverage; $\theta_0 = 1 - \theta_1 - \theta_2$ is the concentration of empty sites; ε is the energetic parameter of lateral interactions between NN $\mathrm{CO_{ads}}$ adspecies (negative for repulsive interactions); T is the absolute temperature, and R the ideal gas constant. The dimensionless constants α, β, γ, and k_r are associated with the rates of adsorption of CO, dissociative adsorption of oxygen, desorption of CO and the surface reaction, respectively. α and β can be varied through varying the gas phase pressures of the two gases, while γ and k_r may depend on temperature, since both CO desorption and surface reaction are activated processes. Oxygen desorption is not taken into account, consistently with experimental observations.

At the next level of modeling, constructing the equations of the quasichemical approximation (QCA), involves introducing the *pair probabilities* (the normalized dimensionless number of NN (ij) pairs on a lattice), g_{ij}. The kinetic (differential algebraic) system of equations of the QCA consists of

$$d\theta_1/dt = \alpha\theta_* - \gamma\theta_1(S_1)^4 - 4k_r g_{12}(S_1S_2)^3, \tag{2a}$$

$$d\theta_2/dt = 4\beta\theta_*^2 - 4k_r g_{12}(S_1S_2)^3, \tag{2b}$$

along with the quasichemical relations expressing local equilibrium,[12–14]

$$r_{11}g_{11}g_{00} = (g_{01})^2, \tag{3a}$$

$$r_{22}g_{22}g_{00} = (g_{02})^2, \tag{3b}$$

$$r_{12}g_{12}g_{00} = (g_{01}g_{02}), \tag{3c}$$

and the symmetry relations,

$$g_{ij} = g_{ji}; \quad (i,j = 0,1,2),$$

where

$$\sum_j g_{ij} = \theta_i \ \ (i,j = 0,1,2), \ \ \theta_0 = 1 - \theta_1 - \theta_2,$$

$$r_{ij} = r_{ji} = \exp[-\varepsilon_{ij}/(RT)],$$

$$S_i = (g_{i0} + r_{i1}g_{i1} + r_{i2}g_{i2})/\theta_i, \quad (i,j = 1,2);$$

and $\varepsilon_{ij} = \varepsilon_{ji}$ are the energetic parameters of lateral interactions between NN adsorbed particles ($\varepsilon_{ij} = 0$ for $i,j = 0$). Symmetry reduces the 9 g_{ij} unknowns to 6, and the conservation of total sites to 5. We thus have a set of 5 coupled differential-algebraic equations (DAEs). Conceptually, for an index-1 system, we use the three algebraic equations to eliminate three of the five remaining g_{ij}, and end up with a set of two coupled nonlinear differential equations for two independent variables (or for two of their independent combinations, for example the two coverages). If limited mobil-

ity of adsorbates is considered (slow diffusion limit), the time evolution of pair probabilities is given by the following ODEs:[14]

$$dg_{11}/dt = 2\alpha g_{10} - 2\gamma r_{11} g_{11} (S_1)^3 - 6k_r r_{11} g_{12} (g_{11}/\theta_1) \times (S_1)^2 (S_2)^3 + 6d_1 g_{10} [S_1 g_{10}/\theta_0 - r_{11} g_{11}/\theta_1] \times (S_1)^2, \quad (4a)$$

$$dg_{12}/dt = \alpha g_{20} - \gamma r_{12} g_{12} (S_1)^3 + 3\beta g_{10} g_{00}/\theta_0 - k_r g_{12} (S_1 S_2)^3 - 3k_r r_{12} (g_{12} S_1 S_2)^2 (S_2/\theta_1 - S_1/\theta_2) + 3d_1 g_{10} [S_1 g_{20}/\theta_0 - r_{12} g_{12}/\theta_1] \times (S_1)^2 + 3d_2 g_{20} [S_2 g_{10}/\theta_0 - r_{12} g_{12}/\theta_2] \times (S_2)^2, \quad (4b)$$

$$dg_{22}/dt = \beta g_{00} + 6\beta g_{20} g_{00}/\theta_0 - 6k_r r_{22} g_{12} (g_{22}/\theta_2) \times (S_1)^3 (S_2)^2 + 6d_2 g_{20} [S_2 g_{20}/\theta_0 - r_{22} g_{22}/\theta_2] \times (S_2)^2, \quad (4c)$$

where d_1 and d_2 are the migration rate constants for CO and O species, respectively. These three differential equations [Eqs. (4)] combined with the two differential equations for the coverages [Eqs. (2)] provide a set of five differential equations for five unknowns. If $d_1, d_2 \to \infty$ and all other rate constants are finite, these three ODEs reduce to the algebraic Eqs. (3). The relations (3) are, in effect, quasi-steady-state approximations that can be obtained in the limit of high mobility for all adsorbed species. The algebraic nature of Eqs. (3) is, therefore, associated with the fact that pair probabilities (higher order moments of the species distribution) have fast dynamics compared to the lower order moments (coverages) and they get quickly slaved to the coverages. Indeed, prescribing the two coverages leads, through the quasichemical and the symmetry relations, to reconstruction of all 9 pair probabilities.

The equations of the MF approximation [Eqs. (1)] follow from those of QCA [Eqs. (2)] if one ignores the spatial correlations, i.e., set $g_{ij} = \theta_i \theta_j$, and also, concomitantly, set $S_i = \exp[-(\theta_1 \varepsilon_{i1} + \theta_2 \varepsilon_{i2})/(RT)]$. A summary of the various model parameters can be found in Appendix A 2.

III. THE COARSE TIMESTEPPER, AND COARSE-KMC STABILITY/BIFURCATION ANALYSIS

Both in the mean field and the quasichemical approximation, the dynamics of the average coverages (zeroth moments of the species lattice distribution) are captured through macroscopic differential equations that close at the level of coverages alone. Mean field does not take spatial correlations into account; QCA goes so far as to take into account pair probabilities, and uses a separation of time scales argument to close the coverage equations by slaving pair probabilities to coverages. In our work we will assume that the dynamics of the average coverages (in a neighborhood of the macroscopic steady-state solutions of the lattice-gas model) can indeed be described by a system of coarse ODEs that close in terms of the average coverages themselves,

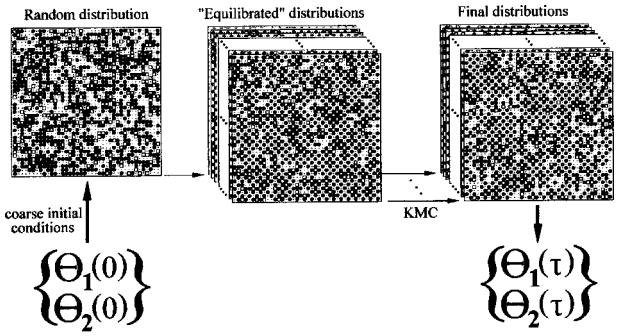

FIG. 1. Schematic of the coarse-KMC (C-KMC) timestepper. For the given macroscopic initial conditions this procedure produces coarse final results via a sequence of microscopic simulations.

$$d\Theta/dt = \mathbf{F}(\Theta; p), \quad (5)$$

where $\mathbf{F}$ exists, but it is unavailable in closed form; p denotes the dependence of the kinetics on parameters, such as gas phase pressures or temperature. The idea is that not only pair probabilities, but all higher order correlations get quickly (after short initial transients) slaved to coverages. We do not, however, know how to construct the corresponding set of differential-algebraic equations for the dynamics of the (in principle infinite) hierarchy of higher moments of the species distributions: we do not have an explicit closure of the coverage equations. Our computational procedure circumvents the derivation of an explicit closure by identifying, through direct ''on demand'' KMC simulation,[15] quantities that would be simply evaluated if $\mathbf{F}$ were explicitly available. This ''closure on demand'' nature of our computational procedure will be revisited in the Discussion. Most elements of this coarse-KMC procedure have been introduced in Ref. 1; while we refer the reader to that article, we revisit here, for completeness, some of the salient algorithmic features and the rationale behind them.

Figure 1 shows a schematic of our coarse timestepper computational procedure, which assumes that a dynamic equation describing the evolution of adsorbate coverages exists and closes at the level of coverages (zeroth moments of the adlayer particle distributions). To be more precise, this equation is what one might call a ''mesoscopic'' equation for the expected values of the coverages; ''mesoscopic'' here implies that these dynamic equations are valid over a time scale that is relatively short compared to the time that it would take for a typical KMC simulation to sample the entire probability distribution function (PDF) of the system (see Ref. 1). The coarse timestepper consists of the following conceptual steps:

(a) We start with an initial condition for the coverages, the zeroth moments of the adsorbate distribution. Since we have assumed that a deterministic mesoscopic equation exists and closes at this level, we will refer to the coverages as the determining moments of the distribution. The illustrative terminology used here is borrowed from the theory of inertial manifolds for dissipative partial differential equations.[16] In that problem, very qualitatively, due to a separation in time scales, fast, higher-order components (modes) of the solution of the problem become quickly slaved to (become functionals

of) a few slower, master modes. By analogy, in our case higher-order spatial moments of the adsorbate distribution become quickly slaved to (functionals of) the master or determining moments: the coverages. We will expand slightly upon this point in the final discussion (Sec. VI).

(b) Coarse, macroscopic initial conditions (coverages) are then lifted to *consistent* microscopic initial conditions, i.e., lattice configurations with adsorbate distributions conditioned on these coverages. This lifting operator is *not unique*—many distributions exist that have the same zeroth moment. We could loosely characterize some of these distributions as ''mature'' (or ''bred'') in the sense that for them the higher-order moments are already slaved to the determining ones (coverages). In principle, if the assumption that an equation exists and closes at the coverage level is correct, *it should not matter* which of these distributions we choose for our procedure; furthermore, it should not matter if the initial distribution(s) we choose are indeed mature: if they are not, they will become so very quickly (compared to the reporting horizon of our timestepper). An ensemble of such consistent initial conditions—distributions conditioned on their low-order moments—is necessary because of variance reduction purposes, as will be further discussed below.

(c) Each member of this ensemble is evolved (using the KMC algorithm) over the timestepper reporting horizon, τ. As we have discussed in detail in Ref. 1, this reporting horizon should be long compared to the ''healing'' or ''slaving'' time (the time that it takes for higher-order moments to get slaved to the governing ones, namely the time it takes for a fresh initial condition to mature, or for the errors we made in lifting to ''heal''). In addition, τ should be short compared to the time it takes for the KMC simulations to diffuse over the long-term PDF of the problem because of the noise. That such a plateau of acceptable timestepper reporting horizons exists is intimately related to the assumption that a meaningful mesoscopic equation for expected coverages exists and closes.

(d) The spatial moments of these final adsorbate distributions are then computed. We can think of this as the application of a restriction operator from the microscopic (distribution) phase space to the macroscopic (moments) one. This operator *is* one-to-one, and the restriction of *any* lifting of a coverage should be that same coverage. We also average over all realizations of the timestepper for variance reduction purposes. This procedure (lifting, evolving microscopically and restricting) combined with ensemble averaging,[15] is the *coarse timestepper for coverages*.

The pseudocode computational construction of the coarse timestepper is then:

(1) [Given macroscopic initial conditions, Θ_0] →
(2) [Create *consistent* microscopic initial conditions (lattice distributions)] →
(3) [Evolve these in time using the KMC timestepper for the time interval τ; perform N_{run} statistically independent runs] →
(4) [Calculate the average $\Theta=\Theta_\tau\equiv\Phi_\tau(\Theta_0;p)$].

The *lifting operator* [step (2)], which constructs microscopic initial conditions (initial lattices having equilibrium distributions of adsorbed species) consistent with a given set of macroscopic initial conditions (coverages), is performed by means of the Metropolis relaxation procedure.

Clearly, a steady state of Eq. (5) is a fixed point of the iteration,

$$\Theta\rightarrow\Phi_\tau(\Theta)\rightarrow\Phi_\tau(\Phi_\tau(\Theta))\equiv\Phi_{2\tau}(\Theta)$$

and to find it we formulate the algebraic system,

$$\mathbf{G}(\Theta;\tau;p)\equiv\Theta-\Phi_\tau(\Theta,p)=\mathbf{0}, \tag{6}$$

which is then solved by means of the Newton–Raphson (NR) iterative method. Steady states of Eq. (5) in general are fixed points of Eq. (6) *for arbitrary* τ. [Periodic solutions of Eq. (5) are also fixed points of Eq. (6) *but only for particular values of* τ. We will also see below cases where spurious apparent fixed points of Eq. (6) arise, again for particular values of τ, that are *not* true steady states of Eq. (5).] In this way, we can calculate both stable and unstable stationary solutions of the unavailable Eq. (5) as fixed points of the coarse timestepper. To estimate the partial derivatives required in the NR iteration, we use numerical differentiation of the coarse timestepper itself; in this paper, we used centered differences in both the coarse variables and the parameters. For noisy problems, however, variance reduction is crucial in estimating derivatives. While in this paper variance reduction was mostly achieved through a large number of samples, filtering and maximum likelihood estimation should (and will) be explored as additional paths to precise derivative estimation.

In order to compute coarse bifurcation diagrams we apply pseudo-arclength continuation with p as the bifurcation parameter. In this case, we also have to estimate the derivatives $\partial\Phi/\partial p$ (coarse derivatives, derivatives of the coarse timestepper with respect to the bifurcation parameter). The eigenvalues λ of the matrix $\partial\mathbf{F}(\Theta,p)/\partial\Theta$ evaluated at these steady-state solutions (the coarse eigenvalues, the eigenvalues of the coarse Jacobian) are expressed via the eigenvalues μ of the matrix $\partial\Phi(\Theta,\tau,p)/\partial\Theta$ (the eigenvalues of the linearization of the coarse timestepper at its fixed point). The matrix $\partial\Phi(\Theta,\tau,p)/\partial\Theta$ is a form of state transition matrix, and $\mu_l=\exp(\lambda_l\tau)$.

It should be emphasized that NR iteration and pseudo-arclength continuation allows the calculation of both stable *and* unstable steady states. More generally, numerical bifurcation algorithms can be now ''wrapped around'' the coarse timestepper and allow the direct calculation/continuation of turning points, other co-dimension one, and higher co-dimension, bifurcation points, as well as of coarse limit cycles and their bifurcation points. Of course, the stable coarse steady states can be calculated as a time average along the stochastic trajectory in phase space by means of the KMC algorithm if the time over which one averages is long

enough. But this computation becomes very time consuming and difficult to interpret close to marginal stability boundaries in parameter space.

We can think of the coarse timestepper for the coverages as simply the result of numerical integration of an equation (like the MF or QCA equations) *that we do not have in closed form*. The typical way of using an integrator subroutine computationally is to run for a step, take the result, feed it back in as an initial condition, and repeat the process. It is also clear, however, and we have discussed this in our previous work, that we can use the same integrator subroutine *not in the usual way* (integrate again and again and again), but differently. Indeed, we can call the integrator for nearby initial conditions, as well as nearby parameter values, to estimate partial derivatives with respect to both the states and to the parameters, which can then be used to perform contraction mappings (like the Newton–Raphson method) to find steady states, or to perform continuation/bifurcation tasks. There is a sequence of references where such numerical enabling technologies for timesteppers are discussed and applied, for both spatially lumped and spatially distributed processes (i.e., for coarse ODEs and for coarse PDEs). Let us also mention that there is an emerging set of techniques (for which the term ''projective integrators'' is used) that try to expand the time horizons over which a microscopic simulator can be used to help estimate the evolution of the *coarse* problem. In the case studied in this article the coarse problem is lumped in space (the coarse equations in question are ordinary differential equations for coverages). Therefore, we use numerical derivatives estimated through the coarse timestepper to perform Newton–Raphson contraction mappings (to find coarse steady states); to augment the system so that one can follow steady state branches through arclength continuation; and to augment the system so that one can converge to co-dimension one bifurcation points (such as, for our example, turning points). When the describing macroscopic equations are partial differential equations (i.e., when the macroscopic problem is spatially distributed) then the coarse timestepper must be combined with *matrix free* (matrix-vector product based) iterative techniques (like the RPM method of Shroff and Keller,[17] or Newton–Picard methods). Instead of numerically estimating partial derivatives, we estimate the action of the coarse Jacobian on selected vectors, and use that to build contraction mappings to compute coarse steady states.

IV. RESULTS

Figures 2, 3, and 4 show bifurcation diagrams with respect to the parameter β, whose variation corresponds to variation of the gas-phase pressure of oxygen. The mean-field approximation (MFA), quasichemical approximation (QCA) as well as the coarse-kinetic Monte Carlo (C-KMC) bifurcation diagrams are overlayed in the figures, the latter in the form of computed points joined by interpolation curves. The MFA diagram is obtained through Eq. (1), while the QCA diagram through the set of Eqs. (2) and (3). In each case (MFA, QCA, and KMC) the stability of the solution branches is indicated, the leading coarse eigenvalue is included in the inset, and the exchange of stability at turning

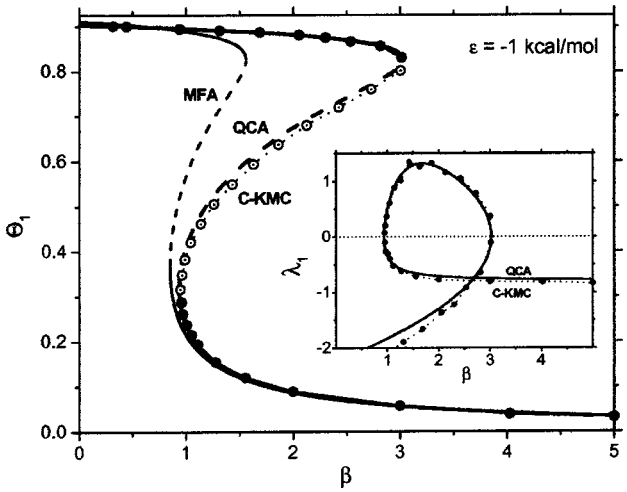

FIG. 2. Bifurcation diagram with respect to the parameter β. Weak interactions: $\varepsilon=-1$ kcal/mol, $\gamma=0.004$, $D=0$, $\tau=0.001$, $N_{\mathrm{MCS}}=100$, 100×100 lattice, $N_{\mathrm{run}}=10^5$. The inset shows the maximal eigenvalues for QCA and for C-KMC.

points (coarse eigenvalue crossing zero) can be seen clearly. Realizations of lattice configurations representative of coarse states marked on the diagrams are also included in Fig. 5; these lattice realizations provide a simple qualitative feeling of the corresponding microscopic states.

It is clear that the QCA does a good job in representing the coarse system dynamics and stability in the case of Fig. 2 and even that of Fig. 3 ($\varepsilon=-1$ and -1.5 kcal/mol, respectively), while the MFA is seen to be inadequate already in Fig. 2. What is remarkable, however, is the bifurcation diagram obtained in Fig. 4 for $\varepsilon=-2$ kcal/mol ($-\varepsilon/RT\approx2$). A long-range-ordered $c(2\times2)$ phase (corresponding to an Ising antiferromagnet) is known to form above the critical value of $|\varepsilon|/RT=2\ln(2^{1/2}+1)\approx1.76$, given by the Onsager exact solution for $\theta_1=0.5$. Therefore, it is expected that large deviations of the MFA and QCA models from the MC results are seen when the lateral interactions between adsorbed par-

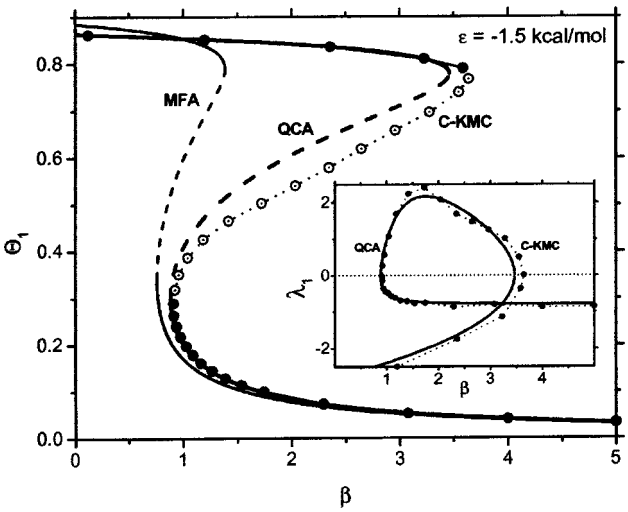

FIG. 3. Intermediate interactions: $\varepsilon=-1.5$ kcal/mol, $D=0$, $\tau=0.001$, $N_{\mathrm{MCS}}=100$, 100×100 lattice, $N_{\mathrm{run}}=10^5$. The inset shows the maximal eigenvalues for QCA and C-KMC.

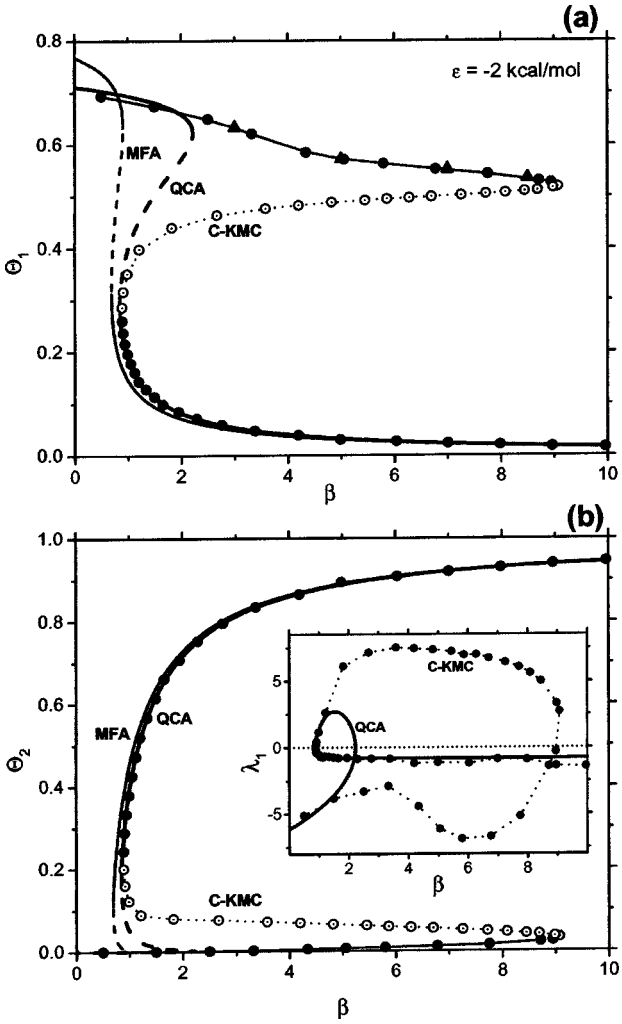

FIG. 4. Strong interactions: $\varepsilon=-2$ kcal/mol, $D=0$, $\tau=0.001$, $N_{\mathrm{MCS}}=100$, 100×100 lattice, $N_{\mathrm{run}}=10^5$. The frame (a) shows CO coverage; (b) O coverage. The inset of frame (b) shows the maximal eigenvalues for QCA and C-KMC. Triangles on the frame (a) give the long time average CO coverage of the usual KMC algorithm with fast diffusion on a 200×200 lattice.

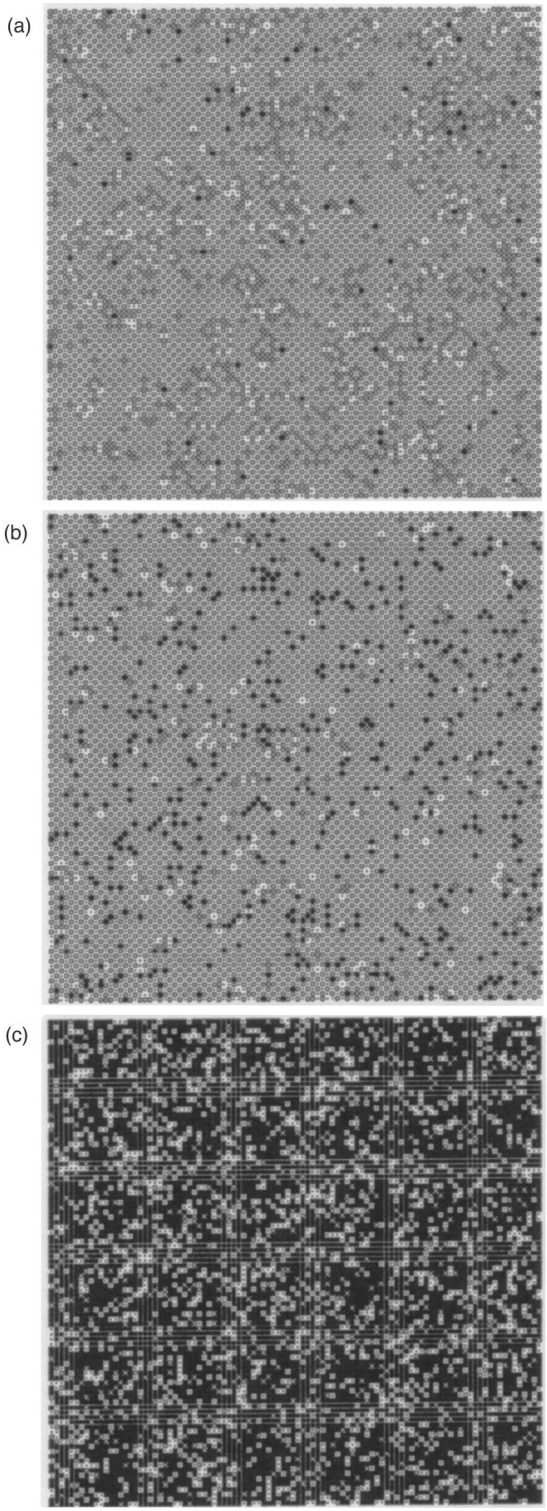

FIG. 5. Snapshots of the 100×100 lattice representative of stationary solutions (results correspond to Fig. 4). Adsorbed CO, empty circles; adsorbed O, filled squares; empty sites, dots. Top frame, $\beta\approx5.8$, $\theta_{\mathrm{CO}}\approx0.563$, $\theta_{\mathrm{O}}\approx0.008$, stable steady state; middle frame $\beta\approx7$, $\theta_{\mathrm{CO}}\approx0.5$, $\theta_{\mathrm{O}}\approx0.055$, unstable steady state; bottom frame, $\beta\approx1.87$, $\theta_{\mathrm{CO}}\approx0.089$, $\theta_{\mathrm{O}}\approx0.7$, stable steady state.

ticles are strong enough and the CO coverage is in the range $0.35<\theta_{\mathrm{CO}}<0.65$. We do not observe any significant effect on the shape of the steady state branch close to the point at which one might expect long-range $c(2\times2)$ order of the CO–adsorbate layer to appear. This is consistent with KMC studies of adsorption isotherms of single species lattice-gas models with repulsive interactions.[18]

It appears, however, that the coarse timestepper procedure is capable of capturing the true hysteresis region much better than either the MFA or the QCA (remember that the coarse NR is capable of converging on both stable and unstable coarse steady states). A comparison with long-term traditional KMC runs, indicated by triangles in Fig. 4(a), shows that the agreement is almost quantitative. These results underscore the potential that timestepper based methods have in capturing dynamic transitions in the coarse behavior of microscopic simulators.

In the case of supercritical interactions ($\varepsilon=-2$ kcal/mol), we performed a large set of C-KMC simulations studying the influence of different parameters of the algorithm. Essentially the same results (norm differences within the NR tolerance) were obtained for various values of: time horizon, τ, ranging from 0.0002 to 0.01; lattice sizes, ranging from 40×40 to 1000×1000; migration rate, D,

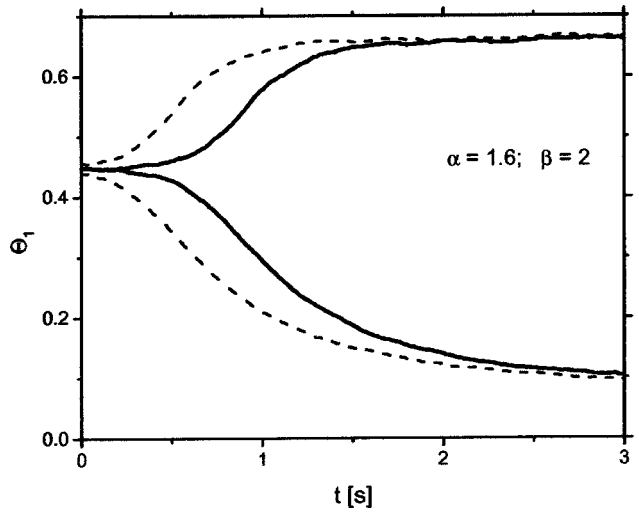

FIG. 6. Evolution of the lattice-gas model starting from the unstable steady state at $\beta=2$. Shown are the typical single runs of KMC algorithm with infinite migration rate on a 200×200 lattice. Initial conditions: $\theta_1(0)=0.448$, $\theta_2(0)=0.077$ for solid curves, while for dashed curves this initial value of $\theta_1(0)$ was perturbed by $\pm7\times10^{-3}$ keeping the same $\theta_2(0)$; the initial equilibrated distribution was produced by the Metropolis algorithm with $N_{\mathrm{MCS}}=200$.

ranging from 0 to 100; $N\times N_{\mathrm{run}}$, ranging from 10^8 to 4×10^9; N_{MCS}, ranging from 30 to 1000.

Of course, the changes of N and N_{run} influence the accuracy of calculations. Comparing several independent runs of the C-KMC algorithm for $N\times N_{\mathrm{run}}=10^9$ (for Figs. 2–4), we find that typically it converges with an absolute accuracy in evaluation of steady state coverages equal to $\approx10^{-3}$. The accuracy in the computation of the eigenvalues was much lower (which is normal, since they are based on numerically estimated coarse derivatives). For the results shown in Fig. 4, the time-horizon of the C-KMC procedure was constant ($\tau=0.001$), the amount of microscopic events during a single run of KMC timestepper was about 30 on a 100×100 lattice (accordingly, 750 on a 500×500 lattice). Typically, the Newton–Raphson procedure converges after 3 or less iterations. However, near the high-θ_{CO} saddle-node point the iterations did not converge so fast. Thus, the step size of the pseudo-arclength continuation was decreased automatically, providing a better initial guess and thus restoring convergence. With sufficient variance reduction, as it was for the data shown in Fig. 4, the continuation algorithm is able to produce the complete one-parameter bifurcation diagram, going around both turning points, during a single simulation run.

V. COMPUTATIONAL ISSUES

Having presented what we consider a successful application of our computer-assisted approach to a nontrivial problem, we now proceed to discuss several issues that arise in its computational implementation. In particular, we discuss those computational issues that we perceive as linked to the physics of the problem and of its microscopic simulator.

Figure 6 demonstrates why the procedure succeeds in locating *unstable* steady states. Indeed, a forward-in-time integrator, especially a noisy one, will move away from an unstable steady state even if it is initialized close to it. What is shown in Fig. 6, however is that it takes some time for this deviation to occur; the corresponding run was initialized very close to the coarse unstable steady state at $\beta=2$. Indeed, since at a steady state the rate of change is zero, change is very slow close to it. So, even in deviating away from the unstable steady state, short-time integration can be used to detect that one is *close* to a steady state. Starting at nearby initial conditions, all of them close to the unstable steady state, short-term integration also can be used to estimate its linearization, and thus to assist in a Newton–Raphson based contraction mapping that will converge to the unstable steady state. It is also worth observing that, if the coarse initial conditions are very close to the unstable solution, the system will sometimes evolve to the stable CO-covered state, and sometimes to stable O-covered state. It takes a significant perturbation of the coarse unstable steady state for the system to almost always evolve to only one of the two coarse stable states.

Next we discuss the further promise of numerical enabling approaches such as the one we presented. Within the context of numerical bifurcation theory, one can construct augmented systems whose solutions are precisely bifurcation points, i.e., points of neutral stability for the original systems. This allows one to circumvent the extensive computational effort required to accurately locate bifurcation points through integration, or through continuation of steady state branches. The idea is to focus precisely on what one wants to find (here a turning point) and write a new, augmented system to locate it. The new system is constructed by augmenting the original physical model with conditions that specify the mathematical properties of the object we want to find (here a zero eigenvalue of the vectorfield linearization, or an eigenvalue of the linearization of the timestepper at 1). This is the realm of traditional bifurcation theory and its numerical implementations; the only additional component is that the coarse timestepper makes this technology applicable in cases where the equation describing the problem *exists but is not available in closed form*.

Figure 7 shows a two-parameter continuation of the turning points of the bifurcation diagram of Fig. 4. For these calculations, significant variance reduction is required: here, one needs to evaluate rather precisely the coarse Jacobian matrix, so as to be able to approximate the necessary second-order coarse derivatives. Such calculations become rather time-consuming. For the data shown in Fig. 7 we used a 500×500 lattice and $N_{\mathrm{run}}=10^4$. The lower saddle-node branch was obtained as a result of an arclength turning point continuation procedure in the (α, β) parameter plane, while for the data points on the upper branch we used several fixed values of α. Upon convergence of the Newton–Raphson method for the augmented system, the estimated absolute value of the original system minimal eigenvalue (which should be zero at the exact turning point) was $<10^{-2}$.

The results of Fig. 8 are used to simply reinforce the notion that exactly the same computational superstructure can be exploited to perform coarse continuation/bifurcation calculations with respect to many other parameters (not just

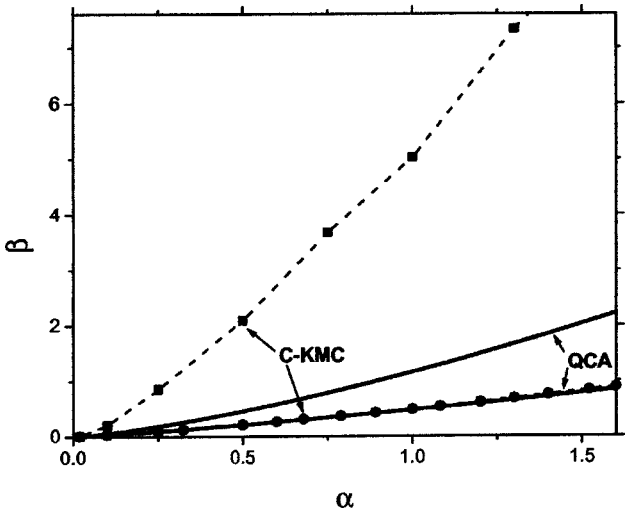

FIG. 7. Two-parameter bifurcation diagram. Lines correspond to the QCA; symbols to C-KMC with $D=0$, 500×500 lattice, $N_{\text{run}}=10^4$. The value of τ was adapted automatically during the continuation procedure.

gas phase pressures, but here also temperature). This brings up an important point. Microscopic evolution laws have parameters that are macroscopic (like the gas phase pressures) or microscopic (like the form of particle interaction potentials). In macroscopic, empirically closed equations one finds both clean macroscopic parameters (such as gas-phase pressures) as well as more complex apparent macroscopic parameters that are affected by the microscopic ones (such as apparent activation energies). Changing the microscopic parameters will simultaneously affect several of the apparent macroscopic ones in non-obvious ways. In complex materials problems, for example, changing one of the macroscopic constitutive law apparent parameters (holding the remaining ones constant) will not, in general, correspond to continuations in *any single* microscopic parameter and vice versa.

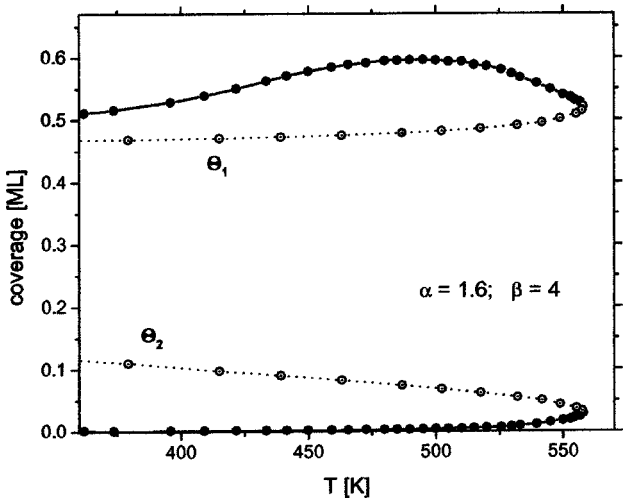

FIG. 8. Continuation with T as bifurcation parameter for fixed $\alpha=1.6$, $\beta=4$. Parameters of the C-KMC algorithm: $D=0$, $\tau=0.001$, $N_{\text{MCS}}=100$, 100×100 lattice, $N_{\text{run}}=4\times10^4$. Shown are the CO and O steady state coverages. There is an additional branch of stable steady states (not shown) for which the coverages are almost constant ($\theta_1(0)\approx0.04, \theta_2(0)\approx0.86$) in the temperature range shown here.

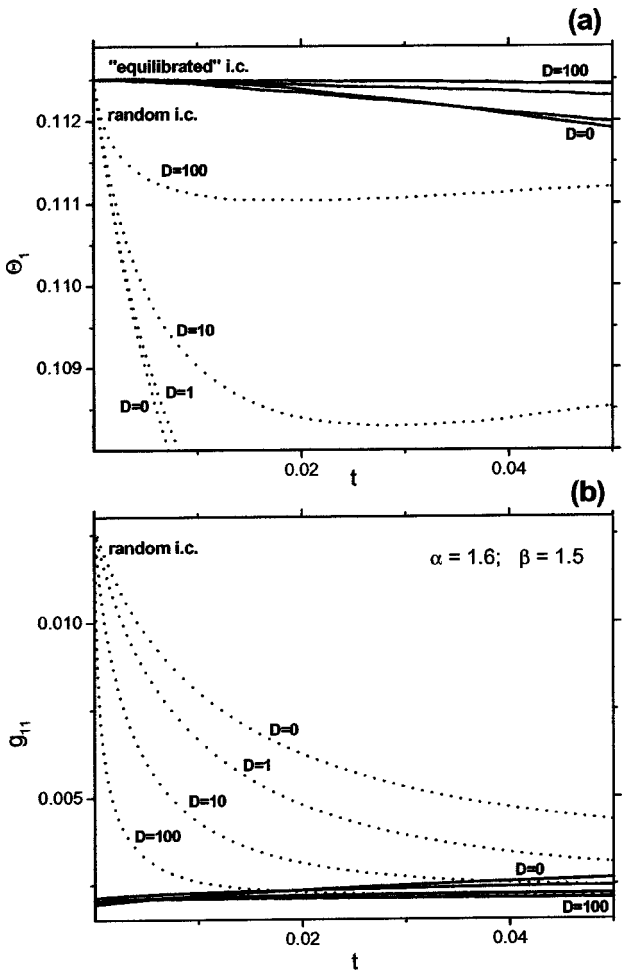

FIG. 9. The average of 400 independent runs of the KMC algorithm on a 500×500 lattice ($\alpha=1.6, \beta=1.5$). CO coverage (a) and pair probability g_{11} (b) are shown. The initial macroscopic conditions are $\theta_1(0)=0.1125$, $\theta_2(0)=0.6178$, and correspond to a stable steady state. The initial microscopic conditions: equilibrated adlayer (solid curves); random distribution (dotted curves). Results for various choices of D are compared.

One of the advantages of procedures like the one we are describing is that it allows naturally continuations with respect to *true microscopic* parameters (such as strengths in interaction potentials) without having to worry about how these enter in the various traditional apparent parameters of macroscopic laws (such as the ones appearing in rheology, transport, materials constitutive relations, etc.).

Lifting issues: Lifting the coarse initial conditions to consistent microscopic distributions is one of the most important elements of the computational procedure described in this article; as we discussed, it is clearly not a one-to-one mapping. More importantly, the microscopic realizations of distributions conditioned on their lower-order moments may be ''mature'' (i.e., they can have the higher moments already slaved) or ''fresh,'' in which case simulation is required to make this slaving take effect. While a strong separation of time scales is an assumption (and, really, a prerequisite) for the procedure, it would be clearly useful to have consistent initial distributions as mature as possible.

Figure 9 shows the effect, on the computation, of using mature (or, as we will refer to them from now on,

equilibrated—with respect to diffusion—) microscopic initial conditions consistent with given coverages. The solid lines show the result of simulations for the evolution of one coverage and one pair probability (θ_1 and g_{11}) starting with equilibrated initial conditions, while the broken lines are for non-equilibrated, random ones. The top panel [Fig. 9(a)] shows CO coverage, while the bottom panel [Fig. 9(b)] shows one of the pair probabilities. It is clear by inspection of Fig. 9(b) that the equilibration time for the pair probabilities can be (depending on the diffusion coefficient) quite long. A practical question then arises: to locate infinite diffusion coarse steady states, we need simulations with very large diffusion rates; on the other hand, these become prohibitively expensive as the diffusion becomes faster. Figure 9(a) shows a practical solution: if the coarse initial condition is equilibrated, then—even with zero diffusion, and certainly for $D=10$ or $D=100$—the infinite diffusion steady state does not change for relatively long simulation periods (roughly $t=0.02$, a typical reporting horizon for our coarse timestepper). An idea would then be to select a coarse initial condition, create a random consistent microscopic distribution, equilibrate that with respect to diffusion, and only then turn on the reaction (with a very small, or even zero diffusion) to compute the timestepper, and through it, ultimately, the coarse steady state. All approaches have been tried for representative points in our diagram: very fast diffusion with no pre-equilibration as well as very fast, fast or no diffusion with pre-equilibration. The results (assuming that one rationally takes care of time-reporting horizons, etc.) have been comparable.

It is worth noting here a promising research direction that might provide an alternative to pre-equilibration through the Metropolis relaxation algorithm we used here. We started using this approach in a Brownian dynamics rheological problem with some success.[19] As we discussed in Ref. 1, if slaving of some higher moments becomes slow, these moments may need to be included in the coarse model as independent variables. In such a case, coarse initial conditions should be taken not only in coverages, but also in some additional moments (e.g., coverages *and* pair probabilities). These initial conditions should be therefore lifted to distributions conditioned *not only on coverages*, but also on the additional moments, (on coverages and on pair probabilities). This is one of the strong points of our procedure: when closure at some level fails, and more moments have to be included in a coarse model, the computational procedure does not have to change. The same exact computational superstructure can be used to analyze the system; the only difference is that coarse initial conditions should be taken in more moments than before (and lifted to consistent microscopic configurations). The KMC evolution part (the heart of the procedure) and the restriction back to governing moments remains the same. It will therefore become important to construct efficient algorithms to initialize distributions conditioned on several lower moments.

The last important issue has to do with a ''computational sanity check,'' the fixed points we find should not depend on (should be converged with respect to) the time-reporting horizon of the timestepper. This, however, should be true if the timestepper reporting horizon is *long enough* compared to the characteristic healing time—the time that it takes for errors made during lifting in the higher order moments of the problem to die, and for the higher order moments to become slaved to the governing ones. If this is not true, then erroneous (spurious) fixed points will result from the simulation. Figure 10 shows the dependence of the fixed points computed through the algorithm on the reporting horizon, on the diffusion rate, and on whether the lifting was to equilibrated microscopic distributions or not. For fast diffusion, whether with or without initial equilibration, one gets the correct fixed point. However, erroneous apparent fixed points can be computed, such as those shown for $D=100$, random initial conditions and short timestepper reporting horizons. Figure 11 explains how this can happen. It is possible that one can find (after time 0.001 for random initial conditions and $D=100$) that the coverage is equal to the initial coverage to within 10^{-7}. But the coverage has not remained constant for all this time—it has been changing, and it will continue to change because [see Fig. 11(b)] the pair probabilities have not yet been slaved to coverages, and they continue to evolve. The reason for the apparent fixed point then, is because we only report differences of the governing moments (the coverages) and because one of our most important assumptions (that the reporting horizon is long enough compared to the healing time) does not hold. The problem can be remedied easily through a number of rational checks: that the fixed point does not change with τ; that the next-highest moment (pair probability) does not change with τ; that the diffusion is fast enough for both above tests to be satisfied at the working τ. This example has been included in order to point out that such computational sanity checks, confirming that simulation parameters do not affect the numerical results, should be also implemented around a working algorithm, and be performed as we march on a solution branch

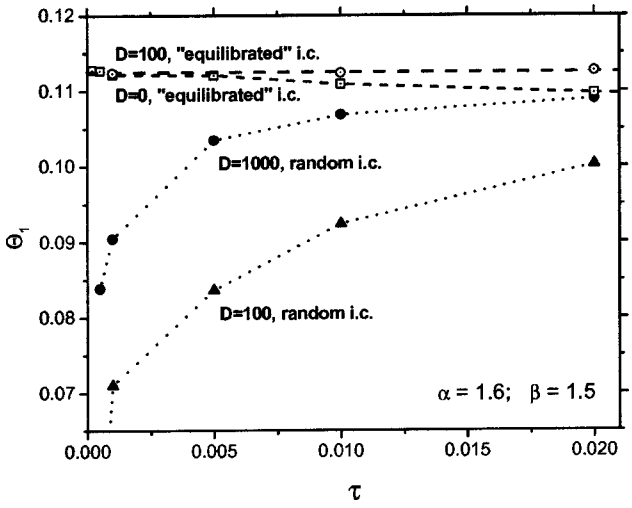

FIG. 10. Solutions found by Newton–Raphson iteration based on the C-KMC timestepper using 500×500 lattice in dependency upon time horizon τ ($\alpha=1.6,\beta=1.5,N_{\mathrm{run}}=4000$). Initial ''macroscopic'' conditions for the first iteration: $\theta_1(0)=0.05$, $\theta_2(0)=0.75$. The initial ''microscopic'' conditions (used at each iteration): ''equilibrated'' adlayer (open symbols); random distribution (filled symbols). Results for various choices of D are compared.

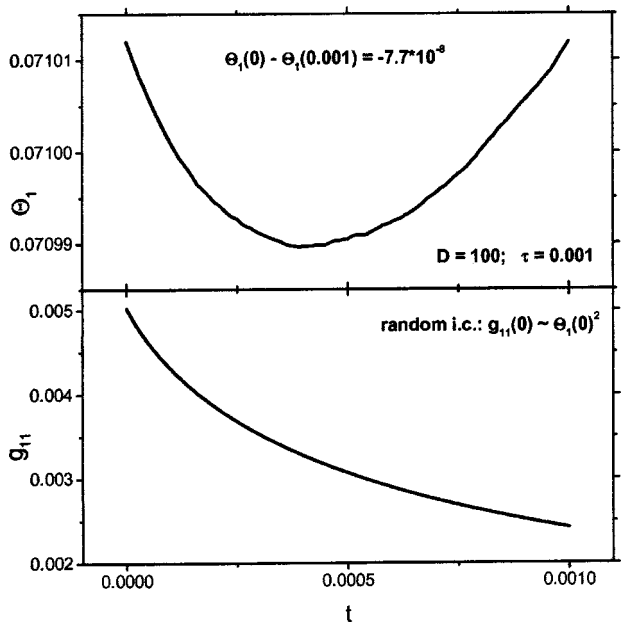

FIG. 11. Spurious apparent steady state solution found through the Newton–Raphson iteration (corresponds to the left filled triangle in Fig. 10), $\tau=0.001$, $D=100$, random initial distribution on a lattice. Shown are the average over 4000 runs CO coverage, θ_1, and pair probability, g_{11}, during the last iteration.

from time to time in order to ensure the validity of the results.

VI. SUMMARY AND DISCUSSION

We have presented a coarse timestepper based computational methodology that enables the continuation, stability and bifurcation analysis of equations for the (expected values of) moments of distributions evolved through microscopic simulators. The advantage of the method lies in that these equations *need not be available in closed form*. In this paper the moments in question were surface concentrations (coverages) of adspecies during surface reactions, and the microscopic simulator was a kinetic Monte Carlo (KMC) algorithm. Analogous timestepper based methodologies can be used for different types of microscopic simulators, e.g., kinetic theory based lattice Boltzmann simulators,[2,3] but also stochastic differential equations (Brownian dynamics, Ref. 19), molecular dynamics (MD), agent based models, etc. Furthermore, the approach is applicable beyond macroscopically ''lumped'' problems, modeled by coarse ODEs as in this paper, to macroscopically distributed problems, modeled by coarse PDEs.[2,3] Timestepper based approaches can thus be used to analyze traveling waves on surfaces at the so-called ''hydrodynamic limit'' (see, e.g., Ref. 20) without explicitly deriving such approximate hydrodynamic equations. Tasks beyond coarse bifurcation analysis (coarse integration, coarse control, coarse optimization, see discussion in Ref. 3) also become accessible.

The approach relies on the fact that, when deterministic macroscopic equations exist and close at some level, i.e., involving a finite number of moments of an evolving distribution, this intrinsically implies that the higher moments of the distribution become quickly slaved to (quickly evolve to functionals of) the slow governing moments. If this were not true, it would not be possible to have deterministic equations at this level of closure (with only so few independent variables). Traditional approaches to the derivation of macroscopic equations involve *closures*; modeling of the effect of the slaved, high modes, on the slow, governing ones. Our approach exploits the implicit separation of time scales, and uses short bursts of true microscopic simulation to estimate just in time the closure.[15] This is not done in the form of a closed formula; it is rather done in the form of the on demand estimation of the coarse timestepper, the result of integrating the unavailable equation for an appropriate time interval. As we have discussed in a previous paper,[1] if, as microscopic parameters vary, the system stops being ''closeable'' at a certain level, the same procedure/computational approach, but lifting now with *more* moments as independent variables, can be used for the coarse study of the problem. Algorithms that will routinely create distributions conditioned on several of their moments become, therefore, important. Variance reduction, and the accurate estimation of coarse derivatives with respect to either variables or parameters also becomes important in such tasks.

What we have discussed here is a *computational enabling technology*: a set of subroutines that are wrapped around the best microscopic timestepper we have available for a process. Through the lift-evolve-restrict procedure and exploiting system identification techniques for noisy systems, these subroutines create a bridge between microscopic simulation and traditional macroscopic numerical analysis. By circumventing the derivation of macroscopic closed equations, this computational methodology (this set of subroutines, if you want) has the potential to extract large scale, expected, system level information easier, faster, better than current direct simulation practice. This is particularly true in regimes where simple simulation is very slow (e.g., when one tries to locate expected marginal stability conditions, e.g., see Ref. 21, or saddle-type, unstable solutions important in the detection of rare events). Should such work be successful, the extraction of relevant information from microscopic computer models may be accelerated, and our ability to analyze and even design complex self-organizing systems through computer modeling may be enhanced.

ACKNOWLEDGMENTS

This work was partially supported through AFOSR (Dynamics and Control), an NSF ITR grant, and a Humboldt Forschungspreis to I.G.K. Informative discussions with Professor J. Evans, Professor J. Li, Professor C. W. Gear, Professor S. Shvartsman, Professor P. G. Kevrekidis, Professor M. Katsoulakis, and Professor R. Kapral are also acknowledged.

APPENDIX: ALGORITHM PARAMETERS AND MODEL PARAMETERS

1. Parameters of the C-KMC algorithm

$N=N_1\times N_2$: Number of perfect square lattice sites; N_1 and N_2 set the cell size in each of the two surface dimensions.

N_{run}: The number of statistically independent KMC runs used for averaging. The same number of runs is utilized for calculations of $\Phi_\tau(\Theta)$ and the estimation of its derivatives.

τ: Reporting time horizon of the C-KMC timestepper. As we have discussed in Ref. 1 and briefly above, this time should be long enough for lifting errors to ''heal'' (that is, for higher order correlations to become slaved to coverages). It should also, however, be short compared to the time that it takes to sample the entire PDF of the problem; remember that the nonlinear equations for coverages we are estimating are mesoscopic equations. We will not repeat here the discussion of the two senses of infinite time (''infinite'' for these mesoscopic equations versus infinite for the system PDF) that can be found in Ref. 1. The characteristic times of the (macroscopic as well as the microscopic) system can and do often depend strongly on the bifurcation parameter changes. It may therefore be necessary to change τ during the arclength continuation. In our implementation, τ is controlled by the amount of microscopic events during a single run. Strictly speaking, τ is a simulation parameter, and we periodically check that the simulations are converged with respect to it, that is, that the increase/decrease of τ does not essentially change the results.

ϕ: User-prescribed tolerance for convergence of Newton–Raphson iterations; it is assumed that the iterative scheme has converged and the solution has been found when $|G_i| < \phi \ \forall i$, that is, that the norm of the residual is effectively zero. This parameter generally depends upon $N \times N_{\mathrm{run}}$ and τ; typically, ϕ was about 2×10^{-6}.

δ: Perturbation, which is used in order to estimate coarse numerical derivatives (with $\Theta_i \pm \delta$ as coarse initial conditions); typically, it was taken equal to 5×10^{-3} for a 500×500 lattice, or 2×10^{-2} for a 100×100 lattice. Once more, variance reduction through maximum likelihood estimators can and should be combined with finite difference formulas for best derivative estimation.

N_{MCS}: The number of successive jump attempts per site (on average) in the Metropolis relaxation algorithm which was used to initialize (equilibrate) the lattice, starting from the lattice obtained on a previous continuation step, and after the appropriate number of adsorbed particles is added/subtracted randomly. To initialize the lattice starting from the random distribution at the beginning of the continuation procedure we used $5 \times N_{\mathrm{MCS}}$ jump attempts (we consider jumps to NN empty sites). In addition, before each run of KMC timestepper (in the set of N_{run}) the Metropolis algorithm with one attempt per site (on average) was executed and the initial lattice was updated each time. The same values of N_{MCS} were utilized to prepare the initial lattice for calculations of both Φ and its derivatives. Once more, N_{MCS} is a simulation parameter, and we check that the results are equilibrated with respect to it (to the extent that we can estimate, total interaction energy on the lattice is indeed at a minimum).

2. Parameters of the model

α, β, γ, k_r, ε, T: The first four parameters are transition probabilities per second for corresponding microscopic events. The ''base set'' of parameters is $\alpha = 1.6\ (\mathrm{s}^{-1})$, $\gamma = 0.001\ (\mathrm{s}^{-1})$, $k_r = 1\ (\mathrm{s}^{-1})$, $\varepsilon = -2\ (\mathrm{kcal/mol})$, $T = 500\ (\mathrm{K})$. Unless otherwise stated the parameters of the model have these values.

$D = D_{\mathrm{CO}} = D_{\mathrm{O}}$: Migration rates at low coverages; note that at high coverages the rate of CO migration is strongly influenced by the lateral interactions. To simulate the migration of adspecies we adopted the independent saddle point scheme,[22] where the jump probability is determined only by its local environment in the initial state. (Note that this is not the case for the Metropolis relaxation algorithm where both initial and final states determine the jump probability.) The rates of CO desorption, CO_2 production, and CO migration depend on the local environment: the presence of NN CO_{ads} particles increases these rates. All steady-state solutions presented in this paper correspond to infinite mobility of adsorbed particles (see discussion in the text for the effect of finite or even zero mobility on the simulations).

[1] A. G. Makeev, D. Maroudas, and I. G. Kevrekidis, J. Chem. Phys. **116**, 10083 (2002).
[2] C. Theodoropoulos, Y.-H. Qian, and I. G. Kevrekidis, Proc. Natl. Acad. Sci. U.S.A. **97**, 9840 (2000).
[3] C. W. Gear, I. G. Kevrekidis, and C. Theodoropoulos, Comput. Chem. Eng. **26**, 941 (2002), see also http://www.neci.nj.nec.com/homepages/cwg/UCLA90.pdf
[4] H. B. Keller, in *Applications of Bifurcation Theory*, edited by P. H. Rabinowitz (Academic, New York, 1997), pp. 359–384.
[5] E. J. Doedel, H. B. Keller, and J.-P. Kernevez, Int. J. Bifurcation Chaos Appl. Sci. Eng. **1**, 493 (1991).
[6] L. S. Tuckerman and D. Barkley, in *IMA Volumes in Mathematics and its Applications*, edited by E. J. Doedel and L. S. Tuckerman (Springer, New York, 1999), pp. 453–466.
[7] D. T. Gillespie, J. Comput. Phys. **22**, 403 (1976); J. Phys. Chem. **81**, 2340 (1977); J. Chem. Phys. **115**, 1716 (2001).
[8] K. Binder, in *Monte Carlo Methods in Statistical Physics*, Topics in Current Physics (Springer, Berlin, 1978), Vol. 7.
[9] K. A. Fichthorn and W. H. Weinberg, J. Chem. Phys. **95**, 1090 (1991).
[10] J. J. Lukkien, J. P. L. Segers, P. A. J. Hilbers, R. J. Gelten, and A. P. J. Jansen, Phys. Rev. E **58**, 2598 (1998).
[11] B. Meng and W. H. Weinberg, J. Chem. Phys. **100**, 5280 (1994).
[12] R. Fowler and E. A. Guggenheim, *Statistical Thermodynamics* (Cambridge University Press, Cambridge, 1952).
[13] V. P. Zhdanov, Surf. Sci. **102**, L35 (1981); **111**, L662 (1981); **137**, 515 (1984).
[14] S. Sundaresan and K. R. Kaza, Surf. Sci. **160**, 103 (1985).
[15] G. Cybenko, in *Identification, Adaptation, Learning: The Science of Learning Models from Data*, NATO ASI Series F (Springer-Verlag, New York, 1996), Vol. 153, pp. 423–434.
[16] P. Constantin, C. Foias, B. Nicoleanko, and R. Temam, *Integral Manifolds and Inertial Manifolds for Dissipative Partial Differential Equations* (Springer Verlag, New York, 1988).
[17] G. M. Shroff and H. B. Keller, SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **30**, 1099 (1993).
[18] K. Binder and D. P. Landau, Phys. Rev. B **21**, 1941 (1980).
[19] C. Siettos, M. D. Graham, and I. G. Kevrekidis (unpublished).
[20] J. W. Evans, D.-J. Liu, and M. Tammaro, Chaos **12**, 131 (2002).
[21] J. Machta, Y. S. Choi, A. Lucke, T. Schweizer, and L. V. Chayes, Phys. Rev. Lett. **75**, 2792 (1995).
[22] C. Uebing and R. Gomer, J. Chem. Phys. **95**, 7626 (1991).