Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Computational Physics 211 (2006) 596–615
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Temporal acceleration of spatially distributed kinetic Monte Carlo simulations

Abhijit Chatterjee, Dionisios G. Vlachos *

Department of Chemical Engineering and Center for Catalytic Science and Technology (CCST), University of Delaware, Newark, DE 19716-3110, United States

Received 12 May 2005; received in revised form 31 May 2005; accepted 13 June 2005
Available online 26 July 2005

Abstract

The computational intensity of kinetic Monte Carlo (KMC) simulation is a major impediment in simulating large length and time scales. In recent work, an approximate method for KMC simulation of spatially uniform systems, termed the binomial τ-leap method, was introduced [A. Chatterjee, D.G. Vlachos, M.A. Katsoulakis, Binomial distribution based τ-leap accelerated stochastic simulation, J. Chem. Phys. 122 (2005) 024112], where molecular bundles instead of individual processes are executed over coarse-grained time increments. This temporal coarse-graining can lead to significant computational savings but its generalization to spatially lattice KMC simulation has not been realized yet. Here we extend the binomial τ-leap method to lattice KMC simulations by combining it with spatially adaptive coarse-graining. Absolute stability and computational speed-up analyses for spatial systems along with simulations provide insights into the conditions where accuracy and substantial acceleration of the new spatio-temporal coarse-graining method are ensured. Model systems demonstrate that the r-time increment criterion of Chatterjee et al. obeys the absolute stability limit for values of r up to near 1.

1. Introduction

Kinetic Monte Carlo (KMC) on a lattice has extensively been employed to study dynamic and equilibrium phenomena in diverse areas ranging from catalysis to crystal growth to surface diffusion and phase transitions on single crystals to transport in microporous materials to cell membrane receptor dynamics [2–7]. KMC provides the exact solution of an underlying master equation along with information about correlations, fluctuations, spatial distributions, etc. These traits make KMC a powerful tool for modeling

* Corresponding author. Tel.: +1 302 831 2830; fax: +1 302 831 1048.
E-mail address: vlachos@che.udel.edu (D.G. Vlachos).

0021-9991/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2005.06.004

molecular scale and noise-dependent phenomena that are crucial, for example, in nucleation and growth, phase transitions, and pattern formation. However, KMC simulation is computationally intensive, and as a result, despite the increase in computational power, is restricted to small lattices and short time scales.

An appealing multiscale paradigm builds upon coarse-graining of the same simulation tool resulting in multiscale molecular methods. Different strategies for arriving at coarse-grained or mesoscopic models are currently being explored. In the *top-down approach*, one postulates the rules of the coarse-grained simulation tool (e.g., coarse potential of interactions, transition probabilities, etc.) by attempting to match an outcome of the molecular and coarse-grained simulators. A common outcome used in off lattice simulations is the radial distribution function. Reviews describing methods, limitations of the top-down approach of molecular simulation, and examples are given in [8–10].

In the *bottom-up* approach, one starts from the microscopic description and through non-equilibrium statistical mechanics attempts to derive the coarse-grained rules. It is this latter strategy that is being discussed here. The coarse-grained Monte Carlo (CGMC) method [11–13] is such a bottom-up spatial coarse-graining method that groups lattice sites into coarse cells to study lattice dynamics and equilibrium. The idea of grouping in CGMC is reminiscent of renormalization group theory [14,15]. However, CGMC is conceptually different from the latter that focuses on determining critical exponents of equilibrium systems typically with first nearest neighbor interactions. The adaptive CGMC (ACGMC) [16,17] is an extension of CGMC that employs adaptive spatial coarse-graining to resolve with accuracy and low computational cost large gradients in solutions. An alternative approach is the wavelet accelerated Monte Carlo (WAMC) method that employs wavelet transformations of the Hamiltonian to determine thermodynamic properties [18,19].

Temporally CGMC techniques have also been developed. In particular, the τ-leap methods that coarse-grain time were developed for studying dynamics and equilibrium in well-mixed (spatially homogeneous) reacting systems. Different implementations, such as the explicit, implicit and trapezoidal τ-leap methods, using different distributions for generating random deviates, have been proposed [1,20–23].

Methods that simultaneously coarse grain space and time for spatially distributed systems are currently lacking. In this paper, a new coarse-grained method called the τ-leap adaptive coarse-grained Monte Carlo (τ-leap ACGMC) method is introduced that simultaneously coarse-grains space and time of lattice KMC. The method combines the binomial τ-leap method of [1] and the ACGMC of [16] for time and space, respectively. These methods are first outlined and the new multiscale Monte Carlo method is subsequently introduced. The issues of accuracy, computational requirements and absolute stability of the τ-leap ACGMC are analyzed and numerically assessed in the grand-canonical, canonical, and combined ensembles using prototype systems. Finally, conclusions are drawn.

2. Stochastic simulations of spatially homogeneous systems

The KMC, also referred to as exact stochastic simulation algorithm or dynamic Monte Carlo, of [24,25] follows the dynamics at a discrete level of description as a Markov process. Consider a well mixed system with N_{sp} species S_i, $i = 1, 2, \ldots, N_{sp}$, among M reactions R_j, $j = 1, \ldots, M$. A reaction event is selected based on the transition probability per unit time (often referred to as the propensity) a_j of all chemical reactions at time t (for details refer to [24,25]). This procedure is repeated and populations are tracked in time.

The original KMC method is extremely slow, especially when large populations of species are encountered. To improve the efficiency of the technique, an approximate stochastic simulation method, termed the τ-leap method, was introduced by Gillespie [20]. Instead of executing one reaction event at a time, in τ-leap methods a 'bundle' of events are allowed to trigger during a time interval $[t, t + \tau)$ and the time is advanced by a mesoscopic amount τ. A central assumption in the τ-leap method, termed as the leap condition [20], requires τ to be sufficiently small so that the change in species population X_i, $i = 1, \ldots, N_{sp}$, is small. In practice one is interested in choosing τ as large as possible in order to reach long times.

The original Poisson distribution based τ-leap method developed by Gillespie samples the reaction bundle size k_j for reaction R_j, i.e., the number of times R_j is executed during τ, using a Poisson distribution. The Poisson distribution random variable is unbounded. As a result, negative populations are encountered [1,26], when k_j exceeds the reactant population size. This is especially the case for small populations or long simulations. Recently, two binomial distribution based τ-leap methods were proposed [1,23] to overcome the negative population issue of the original algorithm. There are subtle differences between the two methods that may potentially lead to different computational behavior.

In the binomial τ-leap method [1], the number of events k_j of each reaction R_j in a time increment τ is sampled using the binomial distribution

$$P_{\mathrm{BD}}(k_j; p_j, k^j_{\max}) = \frac{k^j_{\max}!}{k_j!(k^j_{\max} - k_j)!} p_j^{k_j} (1 - p_j)^{k^j_{\max} - k_j}. \tag{1}$$

Here $k^j_{\max}$ is the limiting reactant population size for each reaction, and $p_j = \min(a_j\tau/k^j_{\max}, 1)$ is the probability associated with a successful event among $k^j_{\max}$ possible events. Eventually, once k_j for all M reactions are determined, the species populations are updated according to

$$X_i(t+\tau) = X_i(t) + \sum_{j=1}^{M} v_{ij} k_j, \quad i = 1, 2, \ldots, N_{\mathrm{sp}} \tag{2}$$

and the time is incremented to $t + \tau$.

There are various methods introduced in [1,20,27] to choose the time increment τ. In this paper, τ is determined using a simple, inexpensive criterion [1] termed the *r-criterion*

$$\tau = r \min_i \left(X_i \Big/ \sum_{\substack{j=1 \\ v_{ij}<0}}^{M} |v_{ij}| a_j \right), \tag{3}$$

where $0 < r < 1$ is a user specified temporal coarse-graining factor. Later it will be shown for prototype examples that the value of $r \sim 1$ is closely related to the onset of numerical instabilities in the algorithm.

Numerical simulations for simple reaction networks [1] and complex biological networks [28] demonstrated that the method accurately captures the probability density function (pdf) of the time-dependent species populations for small τ. Furthermore, fluctuations are usually more accurately captured using the binomial τ-leap method than the Poisson τ-leap method.

3. Stochastic simulations of spatially distributed lattice systems

3.1. The microscopic lattice-gas model for variable-ranged interactions

The physical system is represented by a microscopic lattice $\mathscr{L}$ of N lattice sites in contact with a fluid phase. Each lattice site represents a local minimum of the potential energy surface, and can be either occupied by an adsorbed atom/molecule or vacant. For a single adsorbed species the lattice-gas model [29] is analogous to the Ising model, and the two possible states of a lattice site $p \in \mathscr{L}$, $p = 1, \ldots, N$, are given by its occupation function

$$\sigma_p = \begin{cases} 0, & \text{unoccupied site } p, \\ 1, & \text{occupied site } p. \end{cases} \tag{4}$$

The strength of lateral adsorbate–adsorbate interactions is described by a two-body intermolecular potential $J(r)$ where r is the distance between site p and any other site. The convention used is that $J(r) > 0$ (<0) for attractive (repulsive) interactions. Piecewise continuous interactions of the form

$$J(r) = \begin{cases} J_0/n, & 1 \leqslant r \leqslant L, \\ 0, & r > L \end{cases} \tag{5}$$

are assumed in this paper, where L is the length of intermolecular potential, n is the number of interacting lattice sites, $\beta J_0 = \beta \sum_{1 \leqslant r \leqslant L} J(r)$ is the dimensionless zeroth moment of the potential, $\beta = (k_B T)^{-1}$, k_B is the Boltzmann constant, and T is the absolute temperature.

Processes such as adsorption, desorption, surface diffusion and reaction result in exchange of atoms and/or molecules between the lattice sites and the fluid phase. Transition probabilities per unit time for these processes are employed in a lattice KMC simulation to study the evolution of the configuration space $\underline{\sigma}$, i.e., the vector of occupation functions. In this paper, the role of interactions in the transition probabilities per unit time is modeled using Arrhenius dynamics [5], i.e., the activation energy for desorption and hopping from one site to a neighboring empty site depends only on the strength of adsorption from the departing site and accounts for adsorbate–adsorbate interactions [30].

3.2. Spatially adaptive coarse-grained model

A brief outline of adaptive coarse-grained model is given here. For details see [16,17]. In ACGMC, microscopic sites of $\mathscr{L}$ are grouped into m coarse cells (each cell denoted as C_k, $1 \leqslant k \leqslant m$) of a coarse lattice $\mathscr{L}_c$. The number of lattice sites q_k in C_k is variable, i.e., $\mathscr{L}_c$ is non-uniform. Note that q_k and m are integers and $\sum_{k=1}^{m} q_k = N$. Furthermore, the original coarse-grained MC (CGMC) [11–13] is simply a special case of ACGMC when q_k is the same for all cells. The coarse-grained occupancy function at C_k is $\eta_k = \sum_{p \in C_k} \sigma_p$, such that $0 \leqslant \eta_k \leqslant q_k$.

The coarse-grained transition probabilities, interaction energy and Hamiltonians are derived by projecting the corresponding microscopic entities of $\mathscr{L}$ onto $\mathscr{L}_c$ [16], which are then used to study the evolution of the vector of coarse-grained occupation functions $\underline{\eta}$ using KMC. First, a Haar wavelet basis is employed to coarse-grain the interaction potential. One assumes local mean-field inside a coarse cell. The coarse-grained Hamiltonian and the coarse-grained interaction energy are expressed in terms of the coarse intercell interactions $\bar{J}_{kl}$ between atoms in cells C_k and C_l and the intracell interactions $\bar{J}_{kk}$ of atoms in C_k and are given by

$$\bar{H}(\underline{\eta}) = -\frac{1}{2}\left[\sum_{k \in \mathscr{L}_c} \sum_{l \in \mathscr{L}_c} \bar{J}_{kl}\eta_k\eta_l + \sum_{k \in \mathscr{L}_c} \bar{J}_{kk}\eta_k(\eta_k - 1)\right] + \sum_{k \in \mathscr{L}_c} \bar{h}\eta_k, \tag{6}$$

and

$$\bar{U}_k = \bar{J}_{kk}(\eta_k - 1) + \sum_{\substack{l \in \mathscr{L}_c \\ l \neq k}} \bar{J}_{kl}\eta_l - \bar{h}, \tag{7}$$

respectively. Here $\bar{h}$ is the external field, e.g., chemical potential. Overbars are used to denote coarse-grained variables.

The expressions of transition probabilities per unit time are postulated so that they reduce to the microscopic ones when $q_k = 1$, $k = 1, \ldots, m$. Furthermore, the unknown pre-factors in these expressions are computed so that detailed balance is obeyed and the microscopic and coarse processes result in the same continuum mesoscopic models upon 'complete' coarse-graining. Coarse-grained transition probabilities derived for typical surface processes are tabulated in Table 1. The derivations are given in [16,17].

ACGMC accurately captures the dynamics, equilibrium, and spatial fluctuations [16,17]. In particular, when the potential of interactions is long-ranged, considerable coarse-graining is allowed without

Table 1
List of coarse-grained transition probabilities per unit time for surface processes at coarse cells C_i and C_j [16,17]

Process	Coarse-grained transition probabilities per unit time
Adsorption	$\bar{a}_a(i) = K_a P(q - \eta_i)$
Desorption	$\bar{a}_d(i) = K_d \eta_i e^{-\beta \bar{U}_i}$
Diffusion (nearest cell jump is assumed)	$\bar{a}_m(i \to j) = \frac{\Gamma_m}{q_i q_j (q_i + q_j)} \eta_i (q_j - \eta_j) e^{-\beta \bar{U}_i}$
Unimolecular reaction	$\bar{a}_r(i) = K_r \eta_i$

significant errors. The CPU requirements to reach the same real time in grand-canonical ensemble (adsorption–desorption) simulations using a uniform coarse lattice of cell size q are reduced by q and q^2 for short and long potentials, respectively [12]. Likewise, in the canonical ensemble (diffusion) the corresponding reduction is q^3 and q^4 times for short and long potentials, respectively [11]. The error in the ACGMC solution can be estimated using information loss theory [12]. Recently, systematic adaptive grid generation techniques have been developed that use such information loss theory based error estimates to refine the mesh wherever errors are large [17,31].

4. The proposed τ-leap ACGMC method

Despite the time acceleration of ACGMC, the execution of the algorithm still follows the "one process at a time" approach of microscopic lattice KMC. Thus far, the τ-leap methods have exclusively been used to model spatially uniform, well-mixed systems. The τ-leap methods cannot be directly implemented for a microscopic lattice since any surface process would result in a population change (in the occupation function) that violates the leap condition. On the other hand, if the cell sizes q_i, $i = 1, \ldots, m$, of a coarse lattice are sufficiently large, then the τ-leap method can be combined with the ACGMC method, without violating the leap condition given that the time increment τ is not as large.

Here the binomial τ-leap method of [1] is combined with the ACGMC method to study the evolution of the coarse-grained occupation state vector. The robustness of the binomial τ-leap method is crucial in preventing negative populations and accurately describing spatio-temporal phenomena. The implementation of the new method, termed as (*binomial*) τ-*leap ACGMC method*, is straightforward. For simplicity consider a single component system. There are two species, namely, adsorbed species and vacant sites. The adsorbed and vacant site populations in C_k are given by η_k and $\gamma_k = q_k - \eta_k$, respectively. The total number of species is twice the number of cells m, $N_{sp} = 2m$, and the total number of coarse-grained surface processes is $M = mp$. Here p is the number of surface processes at a microscopic site. For instance, $p = 2$ if only adsorption and desorption are present, $p = 2$ for diffusion on a 1D lattice (jumps to the left and right), and $p = 4$ if all adsorption, desorption, and 1D surface diffusion occur simultaneously. The aforementioned counting implies periodic boundary conditions. Note that M is mesh dependent, and boundaries in boundary value problems could involve different processes affecting the precise values of M and N_{sp}.

The transition probabilities per unit time of each coarse-grained surface process are functions of η_k and γ_k, $k = 1, \ldots, m$, as shown in Table 1 and the coarse-grained surface processes in each coarse cell are analogous to reactions in well-mixed systems. For example, adsorption in C_k is written as a unimolecular reaction

$$A + S_k \to AS_k, \tag{8}$$

desorption as a unimolecular reaction

$$AS_k \to A + S_k, \tag{9}$$

and diffusion as a bimolecular reaction

$$AS_k + S_j \rightarrow AS_j + S_k, \tag{10}$$

where A is an atom/molecule in the fluid phase, S_k is a vacant site in C_k, S_j is a vacant site in a neighboring cell C_j and AS_k denotes an adsorbed species in C_k. The former two reactions are reminiscent of the reversible isomerization reaction that has been studied extensively in a well-mixed (spatially homogeneous) batch reactor. Note that these lattice processes are termed unimolecular or bimolecular reactions because, even though other species (neighboring adsorbed molecules) affect the transition probabilities through the interaction energy, only the populations of one or two species are affected by the particular process. The definition of η_k and γ_k ensures mass conservation at the coarse cell level, i.e., $\eta_k + \gamma_k = q_k$, $k = 1,\ldots,m$. A significant difference from the well-mixed case in a batch reactor is that the corresponding connectivity or reaction matrix of various species is sparse, as compared to the relatively dense matrix of a reaction network, but the transition probabilities may depend on the occupation function of multiple cells in the case of long-ranged interactions. The population-dependent rate constants are reminiscent of variable volume simulations used in cell biology studies (see discussion in [32]).

Below the accuracy and computational requirements of the τ-leap ACGMC are assessed using prototype examples and compared to those of ACGMC where one event per time is executed. Long-ranged interactions are used throughout this paper in order for the local mean-field assumption within each cell to be valid. The validity of the ACGMC in comparison to the microscopic KMC can be found in [11,12,31]. An analysis is also carried out to gain insight into the absolute stability of the algorithm and the computational requirements in terms of the lattice size and the spatial and temporal coarse-graining factors, q and r, using first a uniform mesh. A non-uniform mesh simulation with combined processes is presented later. For all cases, the desorption and migration parameters $K_d = 1\ s^{-1}$ and $\Gamma_m = 3000\ s^{-1}$ are fixed (see Table 1 for their exact meaning).

5. Grand-canonical ensemble (adsorption and desorption on a lattice)

5.1. Accuracy

First adsorption and activated desorption on an initially empty 1D lattice (grand-canonical ensemble) is studied using periodic boundary conditions for two parameter sets using the ACGMC and τ-leap ACGMC methods. The lattice contains $N = 20{,}000$ microscopic sites. Two uniform coarse lattices with $m = 2000$ and 200 coarse cells (spatial coarse-graining of $q = 10$ and 100, respectively) were employed. This level of coarse-graining corresponds roughly to a small 3×3 coarse cell and a 10×10 coarse cell of a 2D lattice.

Fig. 1 compares the spatially averaged coverages and Hamiltonians of the two methods. Results are shown only for the finest mesh of $q = 10$. Similar results are obtained for coarser lattices as well. The τ-leap ACGMC method gives correct dynamics and equilibrium for temporal coarse-graining factors up to $r \leqslant 0.2$, even though each coarse cell contains roughly 3–5 adsorbed molecules. The excellent agreement observed in the Hamiltonians in Fig. 1(b) indicates that the τ-leap ACGMC accurately captures spatial correlations. Error (shown only in Fig. 1(a)) is observed for a larger temporal coarse-graining factor ($r = 0.3$) during the transient where the cells are basically empty. Similar error in transients is also found for $q = 100$ when $r > 0.3$ (not shown). The maximum value of r that provides accurate solutions is consistent with studies of reaction networks in a well-mixed reactor [1,28] and is rationalized below.

Besides the expected value of the coverage, the standard deviation is important when modeling fluctuation-driven phenomena, such as phase transitions and pattern formation. The fluctuations are measured at equilibrium in terms of the standard deviation in the average coverage θ in Fig. 2 as the equilibrium constant $K_a P/K_d$ varies. Here θ is defined as the spatial average of the fraction of occupied sites, i.e.,

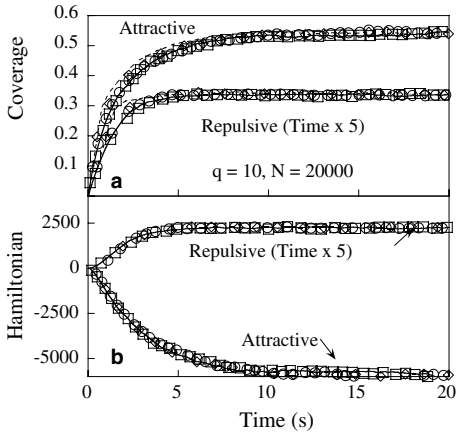

Fig. 1. (a) Spatial average coverage and (b) Hamiltonian for adsorption–desorption starting from an initially empty lattice using ACGMC (solid lines) and τ-leap ACGMC ($r = 0.01$, squares; $r = 0.1$, circles; $r = 0.2$, diamond; $r = 0.3$, dashed lines). Attractive ($\beta J_0 = 2$) and repulsive ($\beta J_0 = -2$) interactions of length $L = 30$ microscopic units are employed, with $K_a P/K_d = 0.4$ and 1, respectively. The ACGMC and τ-leap ACGMC are in excellent agreement except when the temporal coarse-graining factor is $r > 0.3$ (observed only in (a)).

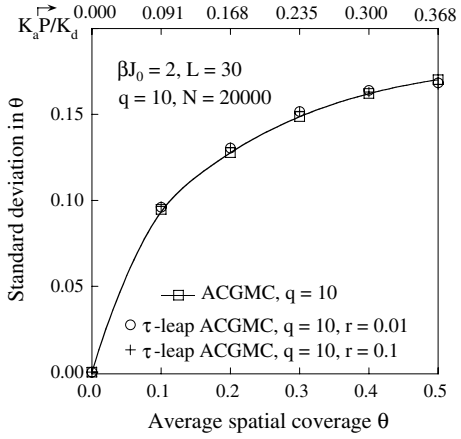

Fig. 2. Fluctuations from ACGMC and τ-leap ACGMC adsorption–desorption simulations for different equilibrium constants indicated. Excellent agreement between the two methods is observed.

$$\theta = \frac{1}{m}\sum_{k=1}^{m}\langle \eta_k \rangle / q_k, \tag{11}$$

$\langle\,\rangle$ denotes ensemble average. The standard deviation is symmetric about $\theta = 0.5$. The τ-leap ACGMC method provides accurate fluctuations for sufficiently low values of r.

The corresponding average bundles per cell in the τ-leap ACGMC simulations are shown in Fig. 3. At equilibrium, nearly 10 times larger bundles adsorb/desorb for $q = 100$, as compared to $q = 10$. The average bundle size per cell in the τ-leap ACGMC method ($\sim$0.4 for $q = 10$ and $\sim$4 for $q = 100$) is substantially larger than the corresponding value of a single event per cell of ACGMC (5×10^{-4} for $q = 10$ and 5×10^{-3} for $q = 100$). The probability distributions of adsorption and desorption bundle sizes per cell at

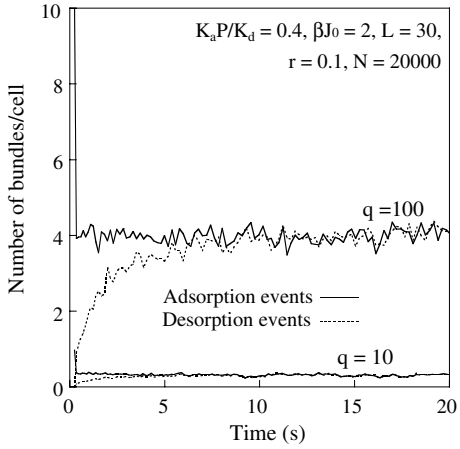

Fig. 3. Molecular bundles per cell vs. time for parameters indicated.

equilibrium are shown in Fig. 4. Both the spatial coarse-graining factor q and the temporal coarse-graining factor r affect the bundle size. However, as shown below the bundle size is of little importance for the τ-leap ACGMC computational speed-up and in fact, it is the size of the time leap τ that is more important. The most important observation so far is that even small spatial coarse-graining, e.g., $q = 10$, enables temporal coarse-graining and provides good accuracy.

5.2. Computational requirements and potential speed-up

The average time increments of the ACGMC (i.e., $\Delta t_{\mathrm{ACGMC}} = 1/\sum_j \bar{a}_j$, summation occurs over all processes j) and of the τ-leap ACGMC (i.e., τ given by Eq. (3)) are plotted in Fig. 5. τ is approximately 10^3 times larger than $\Delta t_{\mathrm{ACGMC}}$ for $r = 0.1$ and appears to be independent of q. In order to understand the time scales and when is the new method really advantageous, we focus on the simpler case of equilibrium. Taking expected values,

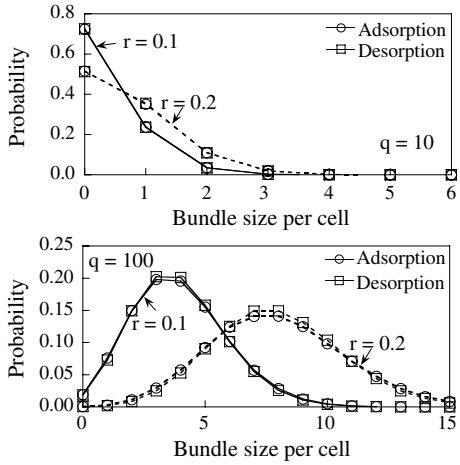

Fig. 4. Probability of bundle size at equilibrium with uniform meshes of $q = 10$ (top) and 100 (bottom) in the grand-canonical ensemble for conditions of Fig. 3. Both the spatial and temporal coarse-graining factors affect the bundle size.

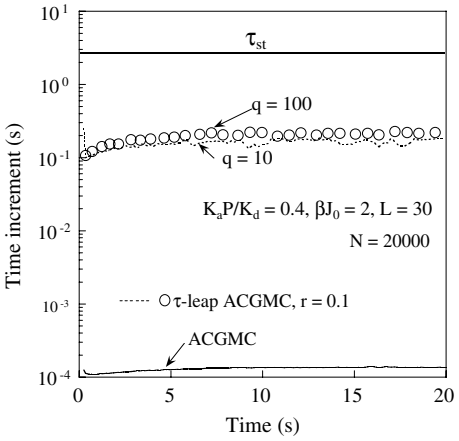

Fig. 5. Time increments of the τ-leap ACGMC and ACGMC simulations corresponding to simulations of Fig. 3. The maximum time increment τ_{st} from the absolute stability criterion is also plotted.

$$\langle \Delta t_{\mathrm{ACGMC}} \rangle = \left[N(K_a P(1-\theta) + K_d \theta e^{-\beta \langle \bar{U} \rangle}) \right]^{-1} \tag{12}$$

and

$$\langle \tau \rangle = r \min_{i=1,\ldots,m} \left[(K_a P)^{-1}, (K_d e^{-\beta \langle \bar{U} \rangle})^{-1} \right], \tag{13}$$

where $\langle\,\rangle$ denotes ensemble average. Here we have omitted spatial correlations and taken $\langle \bar{U} \rangle$ to be $\frac{1}{m}\sum_{k=1}^{m} \langle \bar{U}_k \rangle$. This mean-field type of analysis has been found before to give analytical estimates of CPU in very good agreement with computational times [33]. Eqs. (12) and (13) indicate $\langle \tau \rangle / \langle \Delta t_{\mathrm{ACGMC}} \rangle$ is independent of the cell size q as also evidenced numerically in Fig. 5. This q-independence of Eqs. (12) and (13) can be explained physically. Since adsorption and desorption at a site are independent Poisson processes, the average time increment between any two events is $(K_a P(1-\theta) + K_d \theta e^{-\beta \langle \bar{U} \rangle})^{-1}$. For the entire lattice, the average time interval between any two events is scaled down by N as suggested by Eq. (12) and is independent of the mesh used. In the τ-leap ACGMC, O(q) bundles trigger in each of the m cells in a time interval τ. As the lattice gets further coarse-grained for a fixed lattice size N, the bundle size increases (see Figs. 3 and 4) but the number of coarse cells decreases, leading to a nearly fixed number of adsorption and desorption bundles per time step. This analysis remains unchanged for a non-uniform mesh.

The larger time step of the τ-leap ACGMC method comes at the expense of larger computational requirements per step compared to ACGMC. The computational requirements of the two methods are assessed by performing an operation count for reaching the same real time t. At equilibrium, the CPU time for any algorithm is

$$t_{\mathrm{CPU}} = \Omega \hat{t},$$

where Ω is the number of time steps (leaps for τ-leap ACGMC or MC events for ACGMC) and $\hat{t}$ is the average CPU time required per step. Note that $\Omega_{\mathrm{ACGMC}} = t/\langle \Delta t_{\mathrm{ACGMC}} \rangle$ and $\Omega_{\tau\text{-ACGMC}} = t/\langle \tau \rangle$. Thus, the ratio of CPU times is

$$\frac{t_{\mathrm{CPU,ACGMC}}}{t_{\mathrm{CPU},\tau\text{-ACGMC}}} = \frac{\langle \tau \rangle}{\langle \Delta t_{\mathrm{ACGMC}} \rangle} \cdot \frac{\hat{t}_{\mathrm{ACGMC}}}{\hat{t}_{\tau\text{-ACGMC}}}. \tag{14}$$

In ACGMC, $\sum_{k=1}^{m}(n_k+1)$ multiplications are required for computing all $\bar{U}_i$, $i=1,\ldots,m$. Here n_k is the number of interacting neighbors of cell C_k. Additionally, $3m+1$ multiplications and 1 exponentiation are required in evaluating the transition probabilities and selecting a process. For a uniform mesh of size q, $n_k \approx 2L/q$. Thus

$$\hat{t}_{\text{ACGMC}} = [m(2L/q+1)+3m+1]\hat{t}_{\text{mult/division}} + m\hat{t}_{\text{exp}} + 2\hat{t}_{\text{rand}}. \tag{15}$$

The computational requirement is assumed to be the same for a multiplication and a division, and is denoted as $\hat{t}_{\text{mult/division}}$. The CPU time for exponentiation and generation of a uniform random deviate are $\hat{t}_{\text{exp}}$ and $\hat{t}_{\text{rand}}$, respectively.

In the τ-leap ACGMC method, after computing the transition probabilities of all processes, an additional $4m$ multiplications/divisions are required to determine the time leap using Eq. (3), and $4m$ multiplications and $2m$ binomial random numbers are needed to select the bundle sizes. $\hat{t}$ can be written as

$$\hat{t}_{\tau\text{-ACGMC}} = [m(2L/q+1)+3m+1]\hat{t}_{\text{mult/division}} + m\hat{t}_{\text{exp}} + 8m\hat{t}_{\text{mult/division}} + 2m\hat{t}_{\text{binomial}} \tag{16}$$

and includes contributions from both spatial coarse-graining and the binomial τ-leaping. Finally, one gets

$$\frac{t_{\text{CPU,ACGMC}}}{t_{\text{CPU},\tau\text{-ACGMC}}} = rN\phi\frac{(4+2L/q)\hat{t}_{\text{mult/division}} + \hat{t}_{\text{exp}} + 2\hat{t}_{\text{rand}}/m}{(12+2L/q)\hat{t}_{\text{mult/division}} + \hat{t}_{\text{exp}} + 2\hat{t}_{\text{binomial}}}, \tag{17}$$

where ϕ is a constant

$$\phi = \min\{(K_{\text{a}}P)^{-1}, K_{\text{d}}^{-1}\text{e}^{\beta\langle\bar{U}\rangle}\}(K_{\text{a}}P(1-\theta) + K_{\text{d}}\theta\text{e}^{-\beta\langle\bar{U}\rangle}).$$

Using the equilibrium relation $K_{\text{a}}P(1-\theta) = K_{\text{d}}\theta\text{e}^{-\beta\langle\bar{U}\rangle}$ one gets

$$\phi = 2(1-\theta)\min(1, \theta/(1-\theta)). \tag{18}$$

Typically, the cost of generating a random binomial deviate is much higher than that of computing other terms in Eq. (16), i.e., $\hat{t}_{\tau\text{-ACGMC}} \approx 2m\hat{t}_{\text{binomial}} \gg \hat{t}_{\text{ACGMC}}$ (exceptions to this include extremely long-ranged interactions). Therefore, based on Eq. (14) the τ-leap ACGMC method has a computational advantage only when $\langle\tau\rangle/\langle\Delta t_{\text{ACGMC}}\rangle$ is large. Eq. (17) indicates that computational savings are $O(rN)$ when $L \sim q$ or $L \gg q$. This dependence implies that τ-leap ACGMC is useful especially for large size systems and when the temporal coarse-graining factor r can be large.

The actual CPU at equilibrium vs. the temporal coarse-graining factor r is shown in Fig. 6. The CPU comparison is performed for coarse-graining factors that give accurate solutions. Spatial coarse-graining

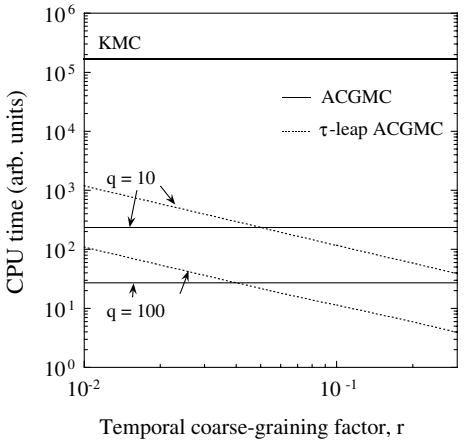

Fig. 6. Computational requirements of ACGMC and τ-leap ACGMC methods for the simulation depicted in Fig. 3.

lowers the computational requirements of ACGMC with respect to the microscopic KMC. Specifically, ACGMC simulations require about 50 and 700 times less CPU time for $q = 10$ and $q = 100$, respectively, than the microscopic KMC method to reach the same real time. The τ-leap ACGMC method provides an additional speed-up for sufficiently large values of r. The linear dependence of τ on r gives a linear dependence of the τ-leap ACGMC CPU on r, in agreement with Eq. (17). The smaller speed-up obtained in Fig. 6 compared to the ratio $\langle\tau\rangle/\langle\Delta t_{\mathrm{ACGMC}}\rangle$ shown in Fig. 5 can be rationalized with Eq. (14). Below a certain value of r, the τ-leap ACGMC method is computationally more expensive than the ACGMC method. The same conclusion holds true for a well-mixed system [1].

The maximum speed-up depicted in Fig. 6 is moderate (just one order of magnitude). However, as evident from Table 2 the speed-up is significant for larger lattice sizes. With $N = 2\times10^5$ (2×10^6) and $r = 0.1$, the τ-leap ACGMC simulation is 10^2 (10^3) times faster than the corresponding ACGMC simulation. This linear N-dependence is in agreement with Eq. (17). Thus, where the real advantage of the τ-leap method lies is in simulation of large systems. As an example, for a 2 mm × 1 mm lattice, one has $N = 2\times10^{12}$ sites (assuming a microscopic lattice spacing of 1 nm). With such a system size, the τ-leap ACGMC method outperforms ACGMC by a factor of 10^9(roughly 30 years of an ACGMC calculation compared to 1 s of the τ-leap ACGMC method!).

Finally, the factor ϕ incorporates the effect of coverage on speed-up and can be simplified as

$$\phi = \begin{cases} 2\theta, & \theta \leqslant 1/2, \\ 2(1-\theta), & \theta > 1/2 \end{cases} \tag{19}$$

for the grand-canonical ensemble. As shown in Table 2, for very low or very high coverages the efficiency of the τ-leap ACGMC method decreases because the corresponding bundle sizes are small. This analysis suggests that τ-leaping is more advantageous for intermediate coverages.

5.3. Absolute stability analysis of expected value

The accuracy and computational speed-up demonstrate that beyond a certain r, inaccuracies are encountered due to violation of the leap condition. A mathematical understanding of the maximum time increment is obtained by testing the absolute stability of the algorithm in analogy to deterministic integration methods. The concept of absolute stability was recently introduced in [27] for the Poisson based τ-leap methods using the reversible isomerization reaction in a well-mixed reactor as a test-bed problem. A stability criterion was derived for the maximum time increment τ_{st} by writing an evolution equation for average species populations using the τ-leap algorithm. Performing the analysis with uniform time increments, one identifies τ_{st} where the propagation coefficient grows unboundedly for long times. Here a stability criterion

Table 2
Computational time required for adsorption/desorption simulations on different uniform meshes indicating the effects of lattice size N, spatial coarse graining q, and rate constants (via the change in θ)

N	q	m	θ	ϕ[a]	$\left(\frac{t_{\mathrm{CPU,ACGMC}}}{t_{\mathrm{CPU},\tau-\mathrm{ACGMC}}}\right)_{\mathrm{observed}}$	$\mathrm{O}\left(\frac{t_{\mathrm{CPU,ACGMC}}}{t_{\mathrm{CPU},\tau-\mathrm{ACGMC}}}\right)_{\mathrm{expected}}$[b]
2×10^6	10^3	2×10^3	0.54	0.92	3700	10^3
2×10^6	10^4	200	0.54	0.92	1077	10^3
2×10^5	10^3	200	0.54	0.92	264	10^2
2×10^6	10^4	200	0.05	0.1	255	10^2
2×10^6	10^4	200	0.95	0.1	227	10^2

Here θ is the average coverage. The temporal coarse-graining factor is $r = 0.1$, the adsorption constant varies in some cases in order to change ϕ, and the desorption constant is $K_{\mathrm{d}} = 1\ \mathrm{s}^{-1}$.

[a] Computed using Eq. (18).

[b] O denotes order of magnitude, computed using Eq. (17).

for the mean of the populations is derived for the spatial τ-leap ACGMC method in the grand-canonical ensemble.

The evolution equation at coarse cell C_k for the $(n+1)$th time increment is given as

$$\eta_k^{n+1} = \eta_k^n + k_{a,k}^n - k_{d,k}^n, \tag{20}$$

where $k_{a,k}$ and $k_{d,k}$ denote the adsorption and desorption bundle size at C_k. Since the bundle size of the τ-leap ACGMC method belongs to a binomial distribution, the average number of firings of the jth coarse process R_j is $\langle a_j\tau\rangle$. Taking expected values, one has (see also Table 1 for the probabilities)

$$\langle \eta_k^{n+1}\rangle = \langle \eta_k^n\rangle + K_a P\tau(q - \langle \eta_k^n\rangle) - K_d\tau\left\langle \eta_k^n e^{-\beta\bar{U}_k^n}\right\rangle. \tag{21}$$

For long-ranged, weak interactions this reduces (see [16] for justification)

$$\left\langle \eta_k^n e^{-\beta\bar{U}_k^n}\right\rangle \approx \langle \eta_k^n\rangle e^{-\beta\langle\bar{U}_k^n\rangle}. \tag{22}$$

In this limit, Eq. (21) is written as

$$\langle \eta_k^{n+1}\rangle = \langle \eta_k^n\rangle + K_a P\tau(q - \langle \eta_k^n\rangle) - K_d\tau\langle \eta_k^n\rangle e^{-\beta\langle\bar{U}_k^n\rangle}. \tag{23}$$

If initial conditions are in the vicinity of the equilibrium occupancy $\eta_{eq} = q\theta$, a Taylor series expansion around η_{eq} for the nonlinear term gives

$$\langle \eta_k^n\rangle e^{-\beta\langle\bar{U}_k^n\rangle} = \eta_{eq}e^{-\beta\bar{U}_{eq}} + e^{-\beta\bar{U}_{eq}}(1 - \beta\bar{J}_{kk}\eta_{eq})\langle\Delta\eta_k^n\rangle - e^{-\beta\bar{U}_{eq}}\sum_{\substack{l\\ l\neq k}}\beta\bar{J}_{kl}\eta_{eq}\langle\Delta\eta_l^n\rangle + \cdots, \tag{24}$$

where $\langle\Delta\eta_k^n\rangle = \langle\eta_k\rangle - \eta_{eq}$ gives the separation from equilibrium and $\bar{U}_{eq} = \bar{J}_{kk}(\eta_{eq} - 1) + \sum_{\substack{l\\ l\neq k}}\bar{J}_{kl}\eta_{eq}$. In the large coarse-graining limit one approaches the mean-field limit such that $\eta_{eq} \gg 1$, $\bar{J}_{kk} = J_0/q$ and $\bar{J}_{kl} \approx 0$. Therefore, $\bar{U}_{eq} = J_0\eta_{eq}/q = J_0\theta$ and Eq. (24) reduces to

$$\langle \eta_k^n\rangle e^{-\beta\langle\bar{U}_k^n\rangle} = \eta_{eq}e^{-\beta J_0\theta} + e^{-\beta J_0\theta}(1 - \beta J_0\theta)\Delta\eta_k. \tag{25}$$

Plugging Eq. (25) into Eq. (23) one gets

$$\langle \eta_k^{n+1}\rangle = \langle \eta_k^n\rangle - (K_a P + K_d(1 - \beta J_0\theta)e^{-\beta J_0})\tau\Delta\eta_k,$$

or

$$\langle \eta_k^{n+1}\rangle = \langle \eta_k^n\rangle[1 - \lambda\tau] + \lambda\tau\eta_{eq}, \tag{26}$$

where λ is identical with the eigenvalue of the Jacobian of the corresponding mean-field ordinary differential equation (ODE), namely,

$$\frac{dc}{dt} = f(c) = K_a P(1 - c) - K_d c e^{-\beta J_0 c}. \tag{27}$$

In Eq. (27), $c(t)$ is the time-dependent mean-field coverage. The eigenvalue is evaluated at the equilibrium surface coverage of the adsorbed species and its magnitude is given by

$$\lambda = \left|\frac{df}{dc}\right|_\theta = K_a P + K_d(1 - \beta J_0\theta)e^{-\beta J_0\theta}. \tag{28}$$

Returning to Eq. (26) and by writing it as

$$\langle \eta_k^{n+1}\rangle = \langle \eta_k^0\rangle[1 - \lambda\tau]^{n+1} + \lambda\tau\eta_{eq}[1 + (1 - \lambda\tau) + (1 - \lambda\tau)^2 + \cdots + (1 - \lambda\tau)^n], \tag{29}$$

one concludes that when $n \to \infty$ the expected value converges to

$$\langle \eta_k^{\infty} \rangle = \eta_{\mathrm{eq}} = q\theta, \tag{30}$$

provided that the propagation coefficient $|1 - \lambda\tau| < 1$, i.e., $\tau < \tau_{\mathrm{st}}$, where

$$\tau_{\mathrm{st}} = \frac{2}{\lambda} = \frac{2}{K_{\mathrm{a}}P + K_{\mathrm{d}}(1 - \beta J_0\theta)\mathrm{e}^{-\beta J_0\theta}}. \tag{31}$$

The result $\tau_{\mathrm{st}} = 2/\lambda$ is similar to that obtained for stability analysis of the explicit Euler scheme for the mean-field equation (Eq. (27)).

Unlike the reversible isomerization reaction in a well-mixed reactor [27], the population state affects the maximum time increment τ_{st} through the interaction energy. Similar to Eqs. (12) and (13), the maximum time increment in Eq. (31) is independent of the spatial coarse-graining q. The maximum time increment τ_{st} of the τ-leap ACGMC method at equilibrium is also plotted in Fig. 5, and is roughly an order of magnitude larger than $\Delta t_{\mathrm{ACGMC}}$ obtained using Eq. (3) with $r = 0.1$ for which accurate results are obtained.

Figs. 7(a) and (b) show τ-leap ACGMC simulations around equilibrium for fixed time increments $\tau = \omega\tau_{\mathrm{st}}$. Note that the value of intermolecular potential chosen here does not satisfy the weak interaction approximation assumed in the aforementioned derivation and serves to assess the validity of Eq. (31) for strong interactions. For small ω, the τ-leap ACGMC simulation is similar to the ACGMC simulation. As ω approaches 1, entire populations are either adsorbed into coarse cells or desorbed from coarse cells and numerical instabilities are evidenced when $\omega > 1$. Far from equilibrium trajectories starting with an initially empty lattice are depicted in Fig. 7(c) for a spatial coarse-graining of just $q = 2$. The transients with $r = 0.1$ and $\omega = 0.1$ are qualitatively similar and errors arise when ω approaches 1.

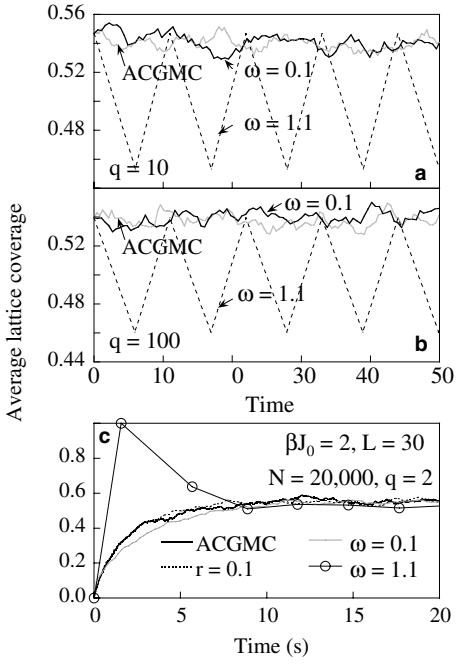

Fig. 7. (a,b) Test of the absolute stability limit of the τ-leap ACGMC method in the grand-canonical ensemble at equilibrium. Uniform time increments of size $\omega\tau_{\mathrm{st}}$ are employed, where τ_{st} is the time increment dictated from absolute stability. As ω approaches 1, the τ-leap ACGMC method becomes unstable. (c) Transients in the spatially averaged coverage far from equilibrium (starting from an initially empty lattice) for very small spatial coarse-graining. The other parameters are those of Fig. 3.

5.4. Relation between the r-criterion and the absolute stability limit

By comparing Eqs. (13) and (31) one can relate r with the loss of stability (for simplicity we omit the term $\beta J_0\theta$ in Eq. (31), i.e., this is the limit of weak interactions or low coverage)

$$\langle\tau\rangle/\tau_{\mathrm{st}} = \frac{r}{2}\min_{i=1,\ldots,m}\left[(K_{\mathrm{a}}P)^{-1},(K_{\mathrm{d}}\mathrm{e}^{-\beta\langle\bar{U}\rangle})^{-1}\right]\left(K_{\mathrm{a}}P + K_{\mathrm{d}}\mathrm{e}^{-\beta J_0\theta}\right), \tag{32}$$

which reduces to

$$\langle\tau\rangle/\tau_{\mathrm{st}} = \begin{cases} r/2, & K_{\mathrm{a}}P \ll K_{\mathrm{d}}\mathrm{e}^{-\beta J_0\theta}, \\ r, & K_{\mathrm{a}}P = K_{\mathrm{d}}\mathrm{e}^{-\beta J_0\theta}, \\ r/2, & K_{\mathrm{a}}P \gg K_{\mathrm{d}}\mathrm{e}^{-\beta J_0\theta} \end{cases} \tag{33}$$

for three limits indicated. Eq. (33) indicates that Eq. (3) with $r \leqslant 1$ obeys the absolute stability limit and rationalizes the fact that values of r up to $\sim$0.2 give reasonably accurate results.

Summarizing, it is demonstrated that the τ-leap ACGMC method gives correct transients, steady state and noise provided the time increments are not large to create numerical instabilities. The maximum time increment at equilibrium for stable behavior in the grand-canonical ensemble is determined from the absolute stability criterion given by Eq. (31). By using a time increment smaller than τ_{st} by about one order of magnitude, stability and accuracy are ensured. The τ-selection procedure following Eq. (3) with $r \leqslant 0.1$ meets this requirement.

6. Canonical ensemble (diffusion on a lattice)

6.1. Accuracy

Diffusion in the presence of long attractive interactions ($\beta J_0 = 2, L = 30$) on a uniform lattice containing $m = 200$ coarse cells and $N = 20{,}000$ microscopic lattice sites is studied ($q = N/m = 100$). The initial distribution of atoms consists of a cluster of a fully covered domain on an otherwise empty lattice. The initial coverage is shown in the inset of Fig. 8. Fig. 8 shows the Hamiltonian vs. time for two different coarse-graining factors r. Excellent agreement between the two methods is observed.

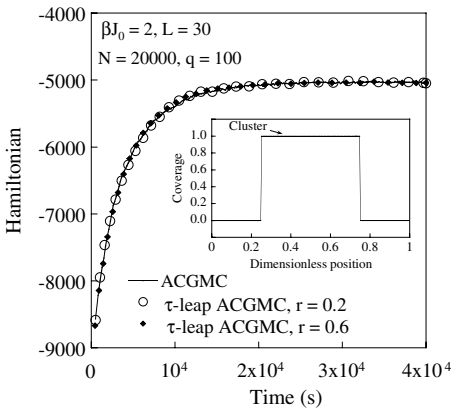

Fig. 8. Hamiltonian as a function of time for diffusion on a uniform lattice starting from a cluster (shown in the inset) using the ACGMC and τ-leap ACGMC methods. Attractive interactions with $\beta J_0 = 2$, $L = 30$ are employed.

6.2. Computational requirements and potential speed-up

The above τ-leap ACGMC simulations with $r = 0.2$ and $r = 0.6$ require roughly 11 and 34 times less computational time, respectively, than the corresponding ACGMC simulation for the specific parameters. In order to understand the acceleration of the τ-leap method, an analysis similar to that in the grand-canonical ensemble is carried out. The CPU for diffusion on a uniform mesh of cell size q is analyzed at equilibrium. The average time increments of the ACGMC and the τ-leap ACGMC methods are

$$\langle \Delta t_{\mathrm{ACGMC}} \rangle = q^2 \Big/ \left[N \Gamma_m \mathrm{e}^{-\beta \langle \bar{U} \rangle} \theta (1-\theta) \right] \tag{34}$$

and

$$\langle \tau \rangle = r q^2 \Gamma_m^{-1} \mathrm{e}^{\beta \langle \bar{U} \rangle} \min[\theta^{-1}, (1-\theta)^{-1}], \tag{35}$$

respectively. The q^2 dependence of $\Delta t_{\mathrm{ACGMC}}$ is explained in terms of larger jumps possible for diffusion on a coarse mesh [11]. Like adsorption and desorption, $\mathrm{O}(q)$ processes are triggered for each of the m cells; hence, $\langle \tau \rangle / \langle \Delta t_{\mathrm{ACGMC}} \rangle \approx \mathrm{O}(N)$. An operation count gives

$$\frac{t_{\mathrm{CPU,ACGMC}}}{t_{\mathrm{CPU},\tau\text{-ACGMC}}} = r N \phi' \frac{(7 + 2L/q)\hat{t}_m + 2\hat{t}_{\mathrm{exp}} + 2\hat{t}_{\mathrm{rand}}/m}{(15 + 2L/q)\hat{t}_m + 2\hat{t}_{\mathrm{exp}} + 2\hat{t}_{\mathrm{binomial}}}, \tag{36}$$

where ϕ' is a constant, $\phi' = \theta\,(1-\theta)\min(\theta^{-1},(1-\theta)^{-1})$ or

$$\phi' = \begin{cases} \theta, & \theta \leqslant 1/2, \\ 1-\theta, & \theta > 1/2. \end{cases} \tag{37}$$

Table 3 shows the effects of lattice size, spatial coarse-graining, and coverage on CPU. The ratios of CPU times for these simulations in column 6 are in agreement with the order of magnitude speed-up predicted by Eq. (36) given in column 7. Like the grand-canonical ensemble, the τ-leap ACGMC method is useful for diffusion on large lattices. Furthermore, τ-leaping is faster for intermediate coverages and a large temporal coarse-graining factor r.

Next we turn to non-uniform meshes. The time increment in Eq. (3) is determined by the smallest populations with the largest transition probabilities. Due to the nature of the diffusion–transition probabilities, this implies that for a non-uniform mesh, $\langle \tau \rangle / \langle \Delta t_{\mathrm{ACGMC}} \rangle$ is limited by the smallest cells of the mesh. For example, τ for a mesh of one small cell of size q_{small} and $(m-1)$ large cells of size q_{large} is

$$\langle \tau \rangle = q_{\mathrm{small}}(q_{\mathrm{small}} + q_{\mathrm{large}}) \min(\theta^{-1}, (1-\theta)^{-1})(2\Gamma_m \mathrm{e}^{-\beta \langle U \rangle})^{-1}. \tag{38}$$

Therefore, unlike the grand-canonical ensemble, τ-leaping in the canonical ensemble needs large spatial coarse-graining for significant acceleration.

Table 3
Computational time required for diffusion (canonical ensemble) on different uniform meshes

N	q	m	θ	ϕ'[a]	$\left(\frac{t_{\mathrm{CPU,ACGMC}}}{t_{\mathrm{CPU},\tau\text{-ACGMC}}}\right)_{\mathrm{observed}}$	$\mathrm{O}\left(\frac{t_{\mathrm{CPU,ACGMC}}}{t_{\mathrm{CPU},\tau-\mathrm{ACGMC}}}\right)_{\mathrm{expected}}$[b]
2×10^6	10^3	2×10^3	0.5	0.5	1045	10^3
2×10^6	10^4	200	0.5	0.5	3073	10^3
2×10^5	10^3	200	0.5	0.5	227	10^2
2×10^6	10^4	200	0.05	0.05	489	10^2
2×10^6	10^4	200	0.95	0.05	298	10^2

Here θ is the average coverage that is fixed. The temporal coarse-graining factor is $r = 0.2$ and the hopping frequency is $\Gamma_{\mathrm{m}} = 3000\ \mathrm{s}^{-1}$.
[a] Computed using Eq. (37).
[b] O denotes order of magnitude, computed using Eq. (36).

6.3. Absolute stability analysis of expected value

Unlike the adsorption–desorption, for diffusion a system of nonlinear equations is obtained for a lattice of N sites due to the bimolecular nature of these events (see Eq. (10)). A simpler model consisting of diffusion on a periodic 1D uniform lattice of three coarse cells of size q with the same initial lattice coverage θ for all three cells (see Fig. 9(a)) is considered to assess the numerical stability. The evolution equation at coarse cell C_k for the $(n+1)$th time increment from a single trajectory is given as

$$\eta_k^{n+1} = \eta_k^n + k_{m,k+1\to k}^n - k_{m,k\to k+1}^n + k_{m,k-1\to k}^n - k_{m,k\to k-1}^n. \tag{39}$$

Here $k_{m,k\to k+1}^n$ denotes the bundle size for diffusion from coarse cell C_k to cell C_{k+1}. Using the transition probability per unit time from Table 1, one has

$$k_{m,k\to j}^n = \frac{\Gamma_m \tau}{2q^3}\eta_k^n(q-\eta_j^n)\mathrm{e}^{-\beta\bar{U}_k^n}. \tag{40}$$

Taking expected values, one has

$$\left\langle \eta_k^{n+1}\right\rangle = \left\langle \eta_k^n\right\rangle + \left\langle k_{m,k+1\to k}^n\right\rangle - \left\langle k_{m,k\to k+1}^n\right\rangle + \left\langle k_{m,k-1\to k}^n\right\rangle - \left\langle k_{m,k\to k-1}^n\right\rangle. \tag{41}$$

For long-ranged, weak interactions

$$\left\langle \eta_k^n(q-\eta_j^n)\mathrm{e}^{-\beta\bar{U}_k^n}\right\rangle \approx \left\langle \eta_k^n\right\rangle(q-\left\langle \eta_j^n\right\rangle)\mathrm{e}^{-\beta\left\langle \bar{U}_k^n\right\rangle}. \tag{42}$$

By inserting Eq. (42) into Eq. (41) and subsequently taking a Taylor expansion of $\langle\eta_k^n\rangle(q-\langle\eta_j^n\rangle)\mathrm{e}^{-\beta\langle\bar{U}_k^n\rangle}$ around the equilibrium occupancy $\eta_{\mathrm{eq}} = q\theta$, one gets

$$\left\langle \eta_k^{n+1}\right\rangle = \left\langle \eta_k^n\right\rangle + \frac{\Gamma_m\tau}{2q^3}\left\langle \Delta\eta_{k+1}^n - 2\Delta\eta_k^n + \Delta\eta_{k-1}^n\right\rangle[q-\beta\eta_{\mathrm{eq}}(q-\eta_{\mathrm{eq}})(\bar{J}_{kk}+\bar{J}_{kl})]\mathrm{e}^{-\beta\langle\bar{U}_{\mathrm{eq}}\rangle}, \tag{43}$$

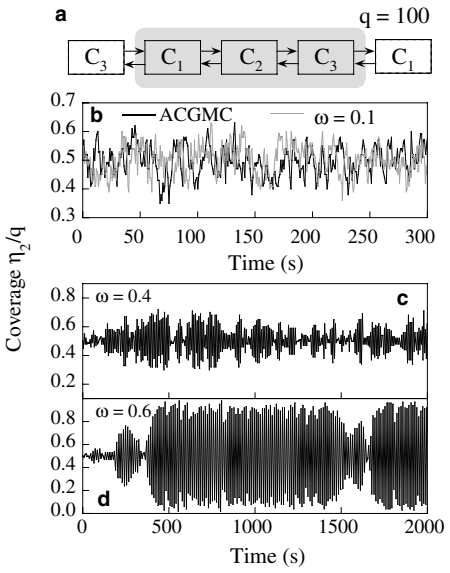

Fig. 9. (a) Schematic of a lattice of three cells with periodic boundary conditions. (b–d) Test of numerical stability in the canonical ensemble using the τ-leap ACGMC method for strong interactions with different uniform time increments of size $\omega\tau_{\mathrm{st}}$. τ_{st} is the time increment of absolute stability. As ω approaches 1, the τ-leap ACGMC method gives incorrect behavior.

where $\langle \Delta\eta_k^n \rangle = \langle \eta_k^n \rangle - \eta_{eq}$ gives the separation from equilibrium and $\bar{U}_{eq} = \bar{J}_{kk}(\eta_{eq} - 1) + \sum_{\substack{l \\ l \neq k}} \bar{J}_{kl}\eta_{eq}$. For large spatial coarse-graining, one approaches the mean-field limit such that $\eta_{eq} \gg 1$, $\bar{J}_{kk} = J_0/q$, $\bar{J}_{kl} \approx 0$ and $\bar{U}_{eq} = J_0\eta_{eq}/q = J_0\theta$. Using the material balance $\eta_1^n + \eta_2^n + \eta_3^n = 3\eta_{eq}$ in Eq. (43), one gets

$$\langle \eta_k^{n+1} \rangle = \langle \eta_k^n \rangle [1 - \lambda\tau] + \lambda\tau\eta_{eq}, \tag{44}$$

where

$$\lambda = \frac{3\Gamma_m}{2q^2}[1 - \beta J_0\theta(1-\theta)]e^{-\beta J_0\theta} = \frac{3D_{eff}}{q^2a^2}e^{-\beta J_0\theta}. \tag{45}$$

Here

$$D_{eff} = D(1 - \beta J_0\theta(1-\theta)) \tag{46}$$

is an effective diffusivity [16], D is the diffusion coefficient related to the hopping frequency in 1D as $D = \frac{\Gamma_m a^2}{2}$, and a is the microscopic hopping distance. As $n \to \infty$, the expected value of cell occupation converges to

$$\langle \eta_k^{\infty} \rangle = \eta_{eq} = q\theta \tag{47}$$

provided that $|1 - \lambda\tau| < 1$, i.e., $\tau < \tau_{st}$, where

$$\tau_{st} = \frac{4q^2 e^{\beta J_0\theta}}{3\Gamma_m[1 - \beta J_0\theta(1-\theta)]}. \tag{48}$$

Similar to Eqs. (34) and (35), the maximum time increment τ_{st} depends on q^2. It can easily be shown that the same eigenvalue as given by Eq. (45) is obtained from the corresponding mean-field system for the explicit Euler scheme.

τ-leap ACGMC simulations are performed on a lattice of three cells with a coverage of $\theta = 0.5$. Fig. 9 shows the coverage in the second cell. Uniform time increments are employed of size $\tau = \omega\,\tau_{st}$. Note again that Eq. (48) is assessed when interactions are strong. For small ω, the τ-leap ACGMC simulation is similar to the ACGMC simulation. For $\omega = 0.4$ enhanced noise is observed over certain time intervals. When $\omega = 0.6$, entire cell populations jump to neighboring cells resulting in huge oscillations. Overall, the stability of the deterministic Euler scheme provides an excellent starting point for determining the maximum time step regarding the expected value of the stochastic process as was also concluded for well-mixed systems in [27].

6.4. Relation between the r-criterion and the absolute stability limit

By comparing Eqs. (35) and (48) in the limit of very weak interactions one can relate the temporal coarse-graining factor r with the loss of stability

$$\langle \tau \rangle / \tau_{st} = 3r \min_{i=1,\ldots,m}[\theta^{-1}, (1-\theta)^{-1}]/4 \tag{49}$$

which reduces to

$$\langle \tau \rangle / \tau_{st} = \begin{cases} 3r/4, & \theta \to 0 \text{ or } 1, \\ 3r/2, & \theta \to 0.5 \end{cases} \tag{50}$$

for the particular limits indicated. Once again values of $r \leqslant 0.1$ fall within the stability region and provide accurate results, consistent with simulations employing the r-criterion shown in Fig. 8.

7. Simulations of multiple processes on non-uniform lattices

Fig. 10 shows a simulation with simultaneous adsorption, desorption and diffusion on a 1D non-uniform lattice of $N = 388{,}800$ microscopic sites using the ACGMC and the τ-leap ACGMC ($r = 0.2$ and 0.4) methods. Interactions between adsorbed particles are absent. The non-uniform mesh used for this calculation contains cells of size 100–1000, as shown in the inset of Fig. 10(a). Non-uniform adsorption rates are set over the lattice, and the ratio of adsorption to desorption rate constants is plotted in Fig. 10(a). Initially the surface is empty.

Fig. 10(b) shows the evolution of the spatially averaged coverage. In the presence of diffusion and spatially non-uniform pressure field, a cluster with diffuse boundary is formed at long times. The steady state coverage profile and the corresponding noise are plotted in Figs. 10(c) and (d), respectively. Excellent agreement between the two methods is observed. This example demonstrates that the τ-leap ACGMC method can be used for accurate adaptive simulations. The τ-leap ACGMC simulation requires roughly 24 and 48 times less computational time with $r = 0.2$ and $r = 0.4$, respectively, than the corresponding ACGMC simulation to reach the same time (comparison is done at steady state).

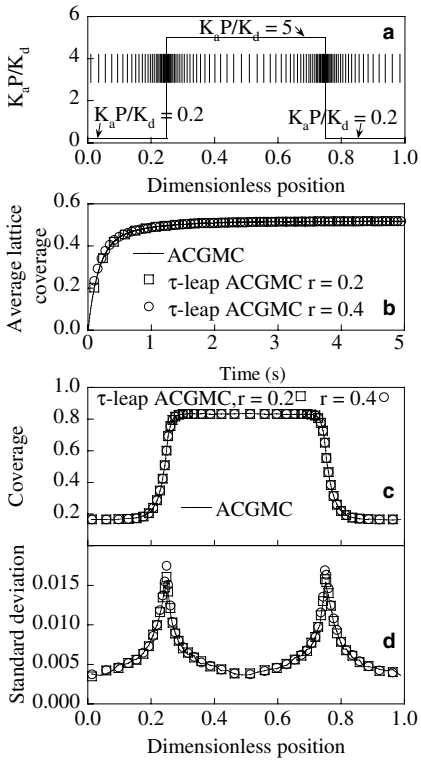

Fig. 10. τ-Leap ACGMC simulation on a non-uniform mesh (shown in the inset of (a)) with a discontinuous equilibrium constant K_aP/K_d. Transient evolution of spatially average coverage (b), steady state coverage profile (c), and noise in the steady state coverage (d) using the τ-leap ACGMC and ACGMC methods.

8. Conclusions

The τ-leap ACGMC method was introduced that combines two stochastic multiscale simulation techniques, namely, the adaptive coarse-grained Monte Carlo (ACGMC), which is a spatially coarse-grained mesoscopic tool, with the binomial τ-leap method, which is a temporally coarse-grained tool. Numerical examples in the grand-canonical, canonical, and combined ensembles demonstrate the ability of the τ-leap ACGMC method to accurately capture spatio-temporal phenomena and fluctuations in distributed lattice systems.

Absolute stability for the expected value and computational speed-up analyses were also performed. It was shown that with sufficient temporal coarse-graining, the τ-leap ACGMC method results in lower computational requirements than the microscopic KMC as well as the ACGMC methods without compromising the accuracy of the solution. The computational savings of the τ-leap ACGMC method over the ACGMC method can be substantial, especially for large size systems, and increase linearly with increasing lattice size and temporal coarse-graining factor r. The time increment τ herein was chosen using the r-based criterion. Model systems demonstrated that the maximum value of r is directly linked to the absolute stability of the expected value of the τ-leap ACGMC algorithm. The absolute stability analysis of the expected value can be well captured by that of the corresponding explicit Euler scheme of the corresponding deterministic model.

The τ-leap ACGMC method is the only discrete stochastic mesoscopic tool with simultaneous spatial and temporal coarse-graining. The above traits enable the τ-leap ACGMC method to reach large length and time scales and study nonlinear and fluctuation-driven phenomena, such as pattern formation, phase transitions, etc. Future work will elucidate the application of these spatio-temporally coarse-grained methods to complex systems.

Acknowledgment

The research was partially supported by the NSF through ITR-0219211 and CTS-0312117.

References

[1] A. Chatterjee, D.G. Vlachos, M.A. Katsoulakis, Binomial distribution based τ-leap accelerated stochastic simulation, J. Chem. Phys. 122 (2005) 024112.
[2] G.H. Gilmer, H.C. Huang, T.D. de la Rubia, J. Dalla Torre, F. Baumann, Lattice Monte Carlo models of thin film deposition, Thin Solid Films 365 (2000) 189.
[3] R.M. Ziff, E. Gulari, Y. Barshad, Kinetic phase transitions in an irreversible surface-reaction model, Phys. Rev. Lett. 56 (1986) 2553.
[4] T.S. Shimizu, S.V. Aksenov, D. Bray, A spatially extended stochastic model of the bacterial chemotaxis signalling pathway, J. Mol. Biol. 329 (2003) 291.
[5] R. Gomer, Diffusion of adsorbates on metal surfaces, Rep. Prog. Phys. 53 (1990) 917.
[6] D.P. Landau, K. Binder, A Guide to Monte Carlo Simulations in Statistical Physics, Cambridge University Press, Cambridge, 2000.
[7] K. Binder (Ed.), Monte Carlo Methods in Statistical Physics, vol. 7, Springer, Berlin, 1986.
[8] F. Müller-Plathe, Coarse-graining in polymer simulation: from the atomistic to the mesoscopic scale and back, ChemPhysChem 3 (2002) 754.
[9] F. Muller-Plathe, Scale-hopping in computer simulations of polymers, Soft Mater. 1 (2003) 1.
[10] S.O. Nielsen, C.F. Lopez, G. Srinivas, M.L. Klein, Coarse grain models and the computer simulation of soft materials, J. Phys. Condens. Matter 16 (2004) R481.
[11] M.A. Katsoulakis, D.G. Vlachos, Coarse-grained stochastic processes and kinetic Monte Carlo simulators for the diffusion of interacting particles, J. Chem. Phys. 119 (2003) 9412.

[12] M.A. Katsoulakis, A.J. Majda, D.G. Vlachos, Coarse-grained stochastic processes and Monte Carlo simulations in lattice systems, J. Comput. Phys. 186 (2003) 250.
[13] M. Katsoulakis, A.J. Majda, D.G. Vlachos, Coarse-grained stochastic processes for microscopic lattice systems, Proc. Natl. Acad. Sci. USA 100 (2003) 782.
[14] N. GoldenfeldLectures on Phase Transitions and the Renormalization Group, vol. 85, Addison-Wesley, New York, 1992.
[15] L.P. Kadanoff, Scaling laws for Ising models near T_c, Physics 2 (1966) 263.
[16] A. Chatterjee, D.G. Vlachos, M.A. Katsoulakis, Spatially adaptive lattice coarse-grained Monte Carlo simulations for diffusion of interacting molecules, J. Chem. Phys. 121 (2004) 11420.
[17] A. Chatterjee, M.A. Katsoulakis, D.G. Vlachos, Spatially adaptive grand canonical Monte Carlo simulations, Phys. Rev. E 71 (2005) 026072.
[18] A.E. Ismail, G. Stephanopoulos, G.C. Rutledge, Multiresolution analysis in statistical mechanics. II. The wavelet transform as a basis for Monte Carlo simulations on lattices, J. Chem. Phys. 118 (2003) 4424.
[19] A.E. Ismail, G.C. Rutledge, G. Stephanopoulos, Multiresolution analysis in statistical mechanics. I. Using wavelets to calculate thermodynamic properties, J. Chem. Phys. 118 (2003) 4414.
[20] D.T. Gillespie, Approximate accelerated stochastic simulation of chemically reacting systems, J. Chem. Phys. 115 (2001) 1716.
[21] D.T. Gillespie, L.R. Petzold, Improved leap-size selection for accelerated stochastic simulation, J. Chem. Phys. 119 (2003) 8229.
[22] M. Rathinam, L.R. Petzold, Y. Cao, D.T. Gillespie, Stiffness in stochastically reacting systems: the implict Tau-leaping method, J. Chem. Phys. 119 (2003) 12784.
[23] T. Tian, K. Burrage, Binomial leap methods for simulating stochastic chemical kinetics, J. Chem. Phys. 121 (2004) 10356.
[24] D.T. Gillespie, A general method for numerically simulating the stochastic evolution of coupled chemical reactions, J. Comput. Phys. 22 (1976) 403.
[25] D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions, J. Phys. Chem. 81 (1977) 2340.
[26] K. Burrage, T.H. Tian, P. Burrage, A multi-scaled approach for simulating chemical reaction systems, Prog. Biophys. Mol. Biol. 85 (2004) 217.
[27] Y. Cao, L.R. Petzold, M. Rathinam, D.T. Gillespie, The numerical stability of leaping methods for stochastic simulation of chemically reacting systems, J. Chem. Phys. 121 (2004) 12169.
[28] A. Chatterjee, K. Mayawala, J.S. Edwards, D.G. Vlachos, Time accelerated Monte Carlo simulations using the binomial τ-leap method, Bioinformatics, 21 (2005) 2136.
[29] T.L. Hill, An Introduction to Statistical Thermodynamics, Dover, New York, 1986.
[30] M.A. Katsoulakis, D.G. Vlachos, From microscopic interactions to macroscopic laws of cluster evolution, Phys. Rev. Lett. 84 (2000) 1511.
[31] A. Chatterjee, D.G. Vlachos, M. Katsoulakis, Numerical assessment of theoretical error estimates in coarse-grained kinetic Monte Carlo simulations: application to surface diffusion, Int. J. Multiscale Comput. Eng. 3 (2005) 59.
[32] M.A. Gibson, J. Bruck, Efficient exact stochastic simulation of chemical systems with many species and many channels, J. Phys. Chem. A 104 (2000) 1876.
[33] J.S. Reese, S. Raimondeau, D.G. Vlachos, Monte Carlo algorithms for complex surface reaction mechanisms: efficiency and accuracy, J. Comput. Phys. 173 (2001) 302.