Coarse-grained numerical bifurcation analysis of lattice Boltzmann models

Leemput, P. Van; Lust, K.W.A.; Kevrekidis, I.G.

[Link to publication in University of Groningen/UMCG research database](#)

ELSEVIER

Coarse-grained numerical bifurcation analysis of lattice Boltzmann models

P. Van Leemput[a,*], K.W.A. Lust[a,b], I.G. Kevrekidis[c]

[a] *Department of Computer Science, K.U.Leuven, B-3001 Heverlee, Belgium*
[b] *Institute for Mathematics and Computing Science, University of Groningen, 9700 AV Groningen, The Netherlands*
[c] *Department of Chemical Engineering and PACM, Princeton University, Princeton, NJ 08544, USA*

Abstract

In this paper we study the coarse-grained bifurcation analysis approach proposed by I.G. Kevrekidis and collaborators in PNAS [C. Theodoropoulos, Y.H. Qian, I.G. Kevrekidis, "Coarse" stability and bifurcation analysis using time-steppers: a reaction-diffusion example, Proc. Natl. Acad. Sci. 97 (18) (2000) 9840–9843]. We extend the results obtained in that paper for a one-dimensional FitzHugh–Nagumo lattice Boltzmann (LB) model in several ways. First, we extend the coarse-grained time stepper concept to enable the computation of periodic solutions and we use the more versatile Newton–Picard method rather than the Recursive Projection Method (RPM) for the numerical bifurcation analysis. Second, we compare the obtained bifurcation diagram with the bifurcation diagrams of the corresponding macroscopic PDE and of the lattice Boltzmann model. Most importantly, we perform an extensive study of the influence of the lifting or reconstruction step on the minimal successful time step of the coarse-grained time stepper and the accuracy of the results. It is shown experimentally that this time step must often be much larger than the time it takes for the higher-order moments to become slaved by the lowest-order moment, which somewhat contradicts earlier claims.
© 2005 Elsevier B.V. All rights reserved.

Keywords: Coarse-grained modeling; Lattice Boltzmann method; Newton–Picard method; Numerical bifurcation analysis; Reaction-diffusion systems

1. Introduction

Arguably the most common approach to study dynamical systems starts from the derivation of some sort of macroscopic description of the system, often in the form of a set of partial differential equations (PDEs). For a reaction-diffusion system with S species on a one-dimensional domain and with space-independent isotropic diffusion, these equations take the form

$$\rho_t^s = D^s \rho_{xx}^s + F^s(\rho^1, \ldots, \rho^S), \quad s = 1, \ldots, S, \quad (1)$$

* Corresponding author. Tel.: +32 16 327082; fax: +32 16 327996.
  E-mail address: pieter.vanleemput@cs.kuleuven.be
(P. Van Leemput).

where the subscripts t and x denote differentiation with respect to time and space, respectively. In many cases, an evolution equation for the macroscopic quantities (here, the densities or concentrations $\rho^s(x, t)$) is already known. There are however cases where an appropriate macroscopic description of a system is not yet known but a microscopic description is available. At the most detailed level one has molecular dynamics simulations that model all interactions between all individual particles (atoms or molecules). Kinetic Monte Carlo methods provide a higher level of abstraction by modeling the statistics of the various interactions between particles. Even more coarse-grained are lattice gas cellular automata (LGCA) and lattice Boltzmann methods (LBMs). These models do no longer model individual microscopic particles and are therefore often called *mesoscopic models*. Instead, they model the behavior of an idealized particle limited to move in certain directions with particular velocities only. While LGCA track the evolution of individual idealized particles, LBMs evolve the distributions of such particles characterized by position and speed.

In many applications one is not interested in the detailed microscopic behavior of a system but only in its macroscopic behavior, i.e., the evolution of macroscopic variables over a large domain and relatively long time interval. These variables are typically the first few moments of a microscopic distribution, e.g., the concentration of the various species. Kevrekidis et al. [1–3] proposed an approach to realize a macroscopic time step for a system for which only a microscopic description is available. The crucial assumption is that a closed macroscopic description in terms of those variables conceptually exists. A macroscopic time step is then performed by first constructing one or more microscopic initial states corresponding to the macroscopic initial condition, then evolving those microscopic initial states using the microscopic evolution laws and finally computing a new macroscopic state. The macroscopic initial state does not contain enough information to initialize the microscopic simulator and the missing information has to be filled in. Using a multiple time scales argument, Kevrekidis et al. argue that the effect of the errors from the initializations will disappear very fast (compared to the macroscopic time step) as the higher-order moments of the microscopic distribution quickly become slaved by (or equivalently, become functionals of) the lower-order ones (corresponding

to the macroscopic variables) i.e., it is assumed that the long-term dynamics of the microscopic model take place on or very near to a lower-dimensional manifold which can be parametrized adequately by the macroscopic variables, and that any orbit started away from this manifold is very quickly attracted to it (at a time scale much smaller than the typical macroscopic time scales). This lower-dimensional manifold is related to the slow manifold in multiple time scale systems and to inertial manifolds. It is expected that this coarse-grained time stepper can replace a time integrator for the (unknown) macroscopic equations in many applications such as bifurcation analysis and control.

Simulating the microscopic models over the whole domain and time interval of interest is often impossible. To cope with this problem, schemes that fully exploit the range of temporal and spatial scales are proposed in [2,4,5], such as the so-called "projective integrators", "gap-tooth scheme", "patch dynamics" and the heterogeneous multiscale method. However, as the number of variables in our lattice Boltzmann (LB) simulations remains limited, we will only use the most basic variant of the coarse-grained time integrator, i.e., we will simulate over the full physical space.

In this article, we will study the application of the coarse-grained time stepper for bifurcation analysis. Bifurcation theory studies the possible transitions between various stable and unstable static and dynamic equilibria in parameter-dependent systems as the parameters are varied. In numerical bifurcation analysis one computes a branch of solutions obtained by varying one parameter of the system and detects or computes points along the branch where the stability of the solutions changes (the bifurcation points). In such points, other branches of solutions often intersect or branches of solutions of a different type emerge or end. Numerical bifurcation analysis is well established for small systems with several software packages available. e.g., AUTO [6] and MATCONT [7]. More recently, several methods have been proposed for large-scale systems. Of particular interest for this paper are methods that operate on top of an existing time integration code such as the Recursive Projection Method (RPM) [8] or the Newton–Picard method [9,10]. We have chosen to use the latter since it is more robust and better suited to compute branches of periodic solutions.

The macroscopic model in this article is the FitzHugh–Nagumo PDE system, while the "microscopic" model is an equivalent lattice Boltzmann (LB) model [11] (which, strictly speaking, is a mesoscopic model), designed to reproduce the behavior of the PDE accurately. The steady states and periodic solutions of this LBM can be analyzed with numerical bifurcation analysis techniques for maps. As such, it is the ideal benchmark to compare the coarse-grained bifurcation results with. Indeed, as we shall argue later, if the microscopic model has a steady state, the best one can hope for is to compute the same steady state using the coarse-grained time stepper (and similarly for periodic solutions). These states may be slightly different from those of an equivalent macroscopic model since every macroscopic model is only an approximation and thus involves modeling errors. Our LBM is fully deterministic and has no chaotic dynamics. Our initialization of the LBM at each coarse-grained time step is also fully deterministic. Therefore, this model is an ideal example to study the errors caused by the imperfect initialization of the "microscopic" state in the coarse-grained time stepper.

The plan of the paper is as follows. We discuss the macroscopic PDE model, the LBM and the coarse-grained time integrator in Sections 2–4, respectively. The Newton–Picard scheme is discussed in Section 5. In Section 6, we compute the bifurcation diagrams for the PDE, the LBM and the coarse-grained time stepper. We also make a careful study of the effects of the initialization of the microscopic simulator on the results. In Section 7, we study the spectra for the different models. Finally, in Section 8 we summarize the main conclusions of this paper.

2. The macroscopic model

Our macroscopic model in this article is the FitzHugh–Nagumo system of two reaction-diffusion equations

$$
\begin{cases}
\rho_t^{\mathrm{ac}} = \rho_{xx}^{\mathrm{ac}} + \rho^{\mathrm{ac}} - (\rho^{\mathrm{ac}})^3 - \rho^{\mathrm{in}}, \\
\rho_t^{\mathrm{in}} = \delta\rho_{xx}^{\mathrm{in}} + \varepsilon(\rho^{\mathrm{ac}} - a_1\rho^{\mathrm{in}} - a_0),
\end{cases}
\tag{2}
$$

with homogeneous Neumann boundary conditions on a one-dimensional domain of length $L = 20$. The variables $\rho^{\mathrm{ac}}(x, t)$ and $\rho^{\mathrm{in}}(x, t)$ denote the activator

and inhibitor "concentration", respectively. (Strictly speaking, these are no concentrations in the physical sense; the values can also be negative.) In all our computations, we fix $\delta = 4$, $a_0 = -0.03$ and $a_1 = 2$ and vary $\varepsilon \in [0, 1]$. For this choice of parameters, we computed a branch of steady states and a branch of periodic solutions. In our numerical experiments in Sections 6 and 7, (2) was discretized using a second-order spatial discretization at the midpoints of 200 grid intervals and the trapezoidal rule for time integration. The time step is fixed during the computation of a single trajectory but varied between 0.02 and 0.47 for different computations in the paper.

3. The lattice Boltzmann model

3.1. Model structure

Lattice Boltzmann models are inherently discrete in space and in time. They model the evolution of a distribution function for each species (activator and inhibitor in our case). The distribution function depends on space, time and velocity and is defined on a space-time lattice with grid spacing Δx in space and Δt in time. We use a D1Q3-type model, i.e., only three values are considered for the velocity:

$$
v_{-1} = -\frac{\Delta x}{\Delta t}, \quad v_0 = 0 \quad \text{and} \quad v_1 = \frac{\Delta x}{\Delta t}.
$$

Let $f_i^s(x_j, t_k)$ denote the value of the distribution function for species $s \in \{\mathrm{ac}, \mathrm{in}\}$ at grid point x_j and time t_k for particles with velocity v_i. The corresponding concentrations (the macroscopic variables in (2)) are defined as

$$
\rho^s(x_j, t_k) = \sum_{i=-1}^{1} f_i^s(x_j, t_k),
\tag{3}
$$

i.e., the zeroth-order moment of the distribution function. Lattice Boltzmann models are often specified using the dimensionless variables for space, time and velocity obtained by rescaling space and time in units of grid spacing Δx and Δt, respectively. To avoid confusion in our notation when moving between the mesoscopic and macroscopic space, we will only express the

higher-order moments in dimensionless form. The dimensionless first- and second-order velocity moments (up to the factor 1/2 for the second-order moment) are

$$\phi^s(x_j, t_k) = \sum_{i=-1}^{1} i\, f_i^s(x_j, t_k)$$

$$= f_1^s(x_j, t_k) - f_{-1}^s(x_j, t_k), \qquad (4)$$

$$\xi^s(x_j, t_k) = \frac{1}{2} \sum_{i=-1}^{1} i^2\, f_i^s(x_j, t_k)$$

$$= \frac{1}{2}(f_1^s(x_j, t_k) + f_{-1}^s(x_j, t_k)). \qquad (5)$$

We will refer to these moments as respectively the "momentum" ϕ^s and (kinetic) "energy" ξ^s (although these are non-conserved quantities in a diffusive system). The state of our one-dimensional LBM at time t_k is fully determined by specifying, for each species and at all lattice points, either the distribution functions $f_i^s(x_j, t_k)$, $i \in \{-1, 0, 1\}$ or the three moments $\rho^s(x_j, t_k)$, $\phi^s(x_j, t_k)$ and $\xi^s(x_j, t_k)$.

The evolution law for the distribution functions is

$$f_i^s(x_{j+i}, t_{k+1}) - f_i^s(x_j, t_k)$$

$$= \Omega_i^s(x_j, t_k) + R_i^s(x_j, t_k), i \in I := \{-1, 0, 1\}. \quad (6)$$

The *collision term* Ω_i^s models the diffusion while the *reaction term* R_i^s models the chemical reactions. A LB time step is usually executed in two phases. In the *collision phase*, the terms Ω_i^s and R_i^s are evaluated and added to $f_i^s(x_j, t_k)$. In the *propagation* or *streaming* phase, the distributions at a lattice site hop to a neighbouring site according to their velocity direction. Eq. (6) is augmented with no-flux boundary conditions which we implemented using the halfway bounce-back scheme [12,13]. This puts the lattice points at the same location as in our PDE discretization.

3.2. The collision operator

For the collision operator we use the Bhatnagar–Gross–Krook (BGK) approximation [14,15]

$$\Omega_i^s(x_j, t_k) = -\omega^s[f_i^s(x_j, t_k) - f_i^{s,\mathrm{eq}}(x_j, t_k)], i \in I \qquad (7)$$

which expresses relaxation to the local equilibrium $f_i^{s,\mathrm{eq}}(x_j, t_k)$. Since the macroscopic mean flow of the reactants in a reaction-diffusion system is zero, the

more general expression for the equilibrium distribution in [14,16,17] simplifies to

$$f_i^{s,\mathrm{eq}}(x_j, t_k) = \nu_i \rho^s(x_j, t_k), \quad i \in I \qquad (8)$$

with ν_i, $i \in I$, satisfying the constraints

$$\sum_{i=-1}^{1} \nu_i = 1 \quad \text{and} \quad \nu_{-1} = \nu_1.$$

This still leaves one degree of freedom for the choice of ν_i. In a reaction-diffusion system, all weights are usually chosen equal [11,18], i.e.,

$$\nu_i = \frac{1}{3}, \qquad (9)$$

which is also the choice we made in all our experiments. Notice that for this choice of weights,

$$\phi^{s,\mathrm{eq}} = f_1^{s,\mathrm{eq}} - f_{-1}^{s,\mathrm{eq}} = 0 \quad \text{and}$$

$$\xi^{s,\mathrm{eq}} = \frac{1}{2}(f_1^{s,\mathrm{eq}} + f_{-1}^{s,\mathrm{eq}}) = \frac{1}{3}\rho^s.$$

The BGK relaxation coefficient ω^s is related to the diffusion coefficient in (1). In [11] it is shown that

$$\omega^s = \frac{2}{1 + 3\, D^s \frac{\Delta t}{\Delta x^2}} \qquad (10)$$

for a one-dimensional model with rest particles and the choice (9) of weights in (8).

3.3. The reaction term

The reaction term is modeled according to [11,18]:

$$R_i^{\mathrm{ac}}(x_j, t_k) = \nu_i \Delta t(\rho^{\mathrm{ac}}(x_j, t_k) - (\rho^{\mathrm{ac}})^3(x_j, t_k)$$

$$- \rho^{\mathrm{in}}(x_j, t_k)),$$

$$R_i^{\mathrm{in}}(x_j, t_k) = \nu_i \Delta t \varepsilon(\rho^{\mathrm{ac}}(x_j, t_k)$$

$$- a_1 \rho^{\mathrm{in}}(x_j, t_k) - a_0), i \in I. \qquad (11)$$

Here it is assumed that the reactions occur at the local diffusive equilibrium [19]. Hence, the weights ν_i are the same as for the equilibrium distribution.

3.4. Extension to continuous time

In Section 6, we will compute a branch of periodic solutions of the LB model. Strictly speaking, a periodic

orbit is a phenomenon in a continuous time model while an invariant circle is the corresponding phenomenon in maps and thus in discrete time systems. However, since the (discrete time) LB model clearly models a continuous time system and since the LB time step is so small compared to the period of the limit cycle, it makes sense to consider a continuous time extension of the LB model and to use numerical techniques developed for such models.

To evaluate the LB state at an arbitrary time T, we use the same strategy as many time integration codes for ordinary differential equations (ODEs) (e.g., LSODE [20]). We first determine k such that $t_{k-1} < T \leq t_k$ and then use a linear interpolation between the results at time t_{k-1} and t_k. Because there is a linear transformation between the distribution functions and the moments, it does not matter whether we interpolate the distribution functions or corresponding moments.

4. The coarse-grained time integrator

4.1. Performing a single coarse-grained time step

In [1–3], a procedure is proposed to perform a macroscopic-level (or coarse-grained) time step for an unknown macroscopic equation using only a microscopic simulator. It is important to first select an appropriate set of macroscopic variables. For the procedure to work well, a macroscopic description must conceptually exist and close using only those variables, i.e., in principle it should be possible to write down an evolution equation for those variables in which all terms can be expressed in function of those variables only, though the precise relationships may not yet be known. This also implies that the chosen macroscopic variables should be a suitable parametrization of the "slow" manifold to which all initial states are quickly attracted (see the discussion in Section 1 of the paper). When using too few variables, the procedure would very likely fail or produce wrong results. On the other hand, using too many moments of the microscopic distribution might reveal too much of the microscopic behavior, and our "macroscopic" model might no longer exhibit the steady states or periodic solutions that we expect to find, but much more complicated dynamics. This will not happen in our case since the

LBM has well-defined steady states and periodic solutions. The existence of a macroscopic PDE in our case confirms that a macroscopic description using only the zeroth-order moment ρ^s of the LB distribution is possible.

A time step of length ΔT with the coarse-grained time stepper consists of three substeps. First, one needs to construct an initial condition for the microscopic simulator which corresponds to the macroscopic state. This step is called the *lifting* (in [2]) or *reconstruction* step (in [5]). Since the microscopic simulator needs more information than provided by the macroscopic variables, the missing information has to be reconstructed. Since a macroscopic model is really a description for the dynamics on the lower-dimensional "slow" manifold mentioned above, it is clear that the best initial condition for the microscopic simulator is the point *on the manifold* corresponding to the particular initial values of the macroscopic variables. However, in [1–3], Kevrekidis et al. argue that the errors caused by the initialization of the missing higher-order moments away from this manifold disappear as the higher-order moments get slaved. According to this argument, the initialization should not matter too much. However, in Section 6, we will show that the fast slaving of the higher-order moments does not imply that all influences of the deviation of the reconstructed initial condition from the corresponding correctly slaved initial condition, disappear quickly, and that a good reconstruction scheme is important.

In our particular case, we need to initialize the missing momentum (4) and energy (5). Given the macroscopic initial condition $\rho^s(x_j, 0)$, we distribute, as in [3], the particles at each grid point over the three velocities using

$$f_i^s(x_j, 0) = w_i \rho^s(x_j, 0), i \in I, \tag{12}$$

where the only constraint on the weights w_i needed for correspondence of the microscopic state with the macroscopic state is

$$\sum_{i=-1}^{1} w_i = 1.$$

This leaves two degrees of freedom. These should be used to initialize the LBM as close as possible to the

"slow" manifold. A very reasonable first choice are the weights of the local diffusive equilibrium distribution, i.e.,

$$w_i = \frac{1}{3}, \tag{13}$$

which corresponds to $\phi^s(x_j, 0) = 0$ and $\xi^s(x_j, 0) = (1/3)\rho^s(x_j, 0)$.

In the second step, the microscopic initial condition is evolved over the macroscopic time step ΔT using the microscopic simulator (the LB model in our case). Finally, the macroscopic variables at the end of the time step are computed from the final microscopic state. This step is called the *restriction step* in [2]. For our LB model, this is done using (3).

When the underlying microscopic model is a stochastic model or when the lifting scheme is stochastic, the result of a coarse-grained time step is again a stochastic variable, characterized by an average, a variance, etc. To get a sufficiently small variance, which is needed for our numerical bifurcation techniques but also for other methods such as the projective integration mentioned in the introduction, one should run many microscopic simulations from the same (if the simulator itself is stochastic) or equivalent initial conditions. In the restriction step, the result for all the simulations must then be averaged. A similar problem occurs for deterministic microscopic models with chaotic dynamics. In this case, a lot of nearby initial conditions have to be used. These difficulties do not occur in our LBM. It is sufficient to evolve a single initial condition once, interpolating as explained before between two LB time steps at the end if ΔT is not a multiple of the LB time step Δt.

The last difficulty is the choice of the macroscopic time step ΔT. According to [2,21], this time step should be larger than the time it takes for the higher-order moments to become slaved by the lower-order ones, i.e., the time that the solution takes to approach the lower-dimensional "slow" manifold mentioned before. The latter time interval is also called the *healing time* in [22] and is typically small compared to the relevant macroscopic time scales. On the other hand, as is clearly demonstrated in [21], ΔT should not be too large either, in particular in the case of a stochastic or chaotic microscopic model because the various realizations might diffuse irreparably over a large part of the attractor, losing phase

information. The latter is not a problem for our LB model.

In this paper, we will show that the claim made in previous papers that the initialization of the higher-order moments does not matter much, is not always right. Our numerical experiments will demonstrate that if the initialization of the higher-order moments is too far from the unknown correctly slaved state, the effects of this deviation might decay very slowly (on a much longer time scale than the healing time, in fact, on a time scale comparable to the slowest macroscopic time scales) and sufficiently accurate results might only be obtained using a coarse-grained time step which is several orders of magnitude larger than the healing time.

4.2. Extension to continuous time

In applications, one often has to integrate over a time interval T much larger than the maximal allowable macroscopic time step ΔT. The above scheme is then repeated until the end time T is reached. In this procedure, the restriction followed immediately by lifting from the end point to generate microscopic initial conditions for the next integrations is essential. Since we remove state information at every restriction step and add slightly different information again to the system in the following reconstruction step, performing k coarse-grained time steps with time step ΔT is not equivalent to performing a single coarse-grained time step with time step $k\Delta T$, in particular when there is an upper limit to ΔT. Though it is not needed in our case to use multiple coarse-grained time steps since there is no maximum to the allowable time step contrary to the test cases in [21,22], we will still use this procedure to experiment with small macroscopic time steps and to test the concept. We believe that many of the conclusions we draw from this experiment will carry over to systems where the maximal successful macroscopic time step is limited.

If T is not a multiple of ΔT, we have two choices. If ΔT is small compared to the dominant time scales of the macroscopic dynamics, i.e., if ΔT is comparable to what a time step would be in a typical numerical integrator for the macroscopic model if the latter were known explicitly, we can use the same approach as for the LBM and interpolate between the two last states. If ΔT is larger, we need another approach. Let

$\Delta T_{\min}$ and $\Delta T_{\max}$ be the minimum and maximum allowed macroscopic time step, estimated according to the criteria specified in Section 4.1. Varying ΔT between these limits should have very little influence on the result at time T. Therefore, we change ΔT such that T is an integer multiple of ΔT. This approach may fail however if T is not sufficiently larger than $\Delta T_{\max}$ and the time step limits are too close to each other. We used the second approach in our experiments.

5. Numerical bifurcation analysis

To perform a numerical bifurcation analysis of a system of autonomous PDEs, the PDEs are first space-discretized. In many cases, this leads to a large system of ordinary differential equations (ODEs)

$$\frac{\mathrm{d}u}{\mathrm{d}t} = f(u, \gamma), \quad f : \mathbb{R}^N \times \mathbb{R}^\Gamma \mapsto \mathbb{R}^N, \tag{14}$$

though in some cases, a system of differential-algebraic equations is obtained. In (14), γ denotes the parameters of the system. For most discretization schemes, the Jacobian matrix $\partial f(u, \gamma)/\partial u$ is a large but very sparse matrix. In bifurcation analysis, steady states are usually computed by applying some version of Newton's method to $f(u, \gamma) = 0$. The stability of the resulting steady state is determined by the eigenvalues λ_l of $\partial f(u, \gamma)/\partial u$. A steady state is asymptotically stable if

$$Re(\lambda_l) < 0, \quad l = 1, \dots, N.$$

For a discretized PDE, only the rightmost eigenvalues will be good approximations to the true eigenvalues of the continuous PDE problem, but these are precisely the eigenvalues which determine stability. Bifurcations occur when one or more eigenvalues cross the imaginary axis as the parameter is changed.

Assume $\varphi_T(u(0), \gamma)$ is the solution $u(T)$ at time T of (14) with initial condition $u(0)$ at the parameter values γ. A steady state of (14) is also a fixed point of the map

$$u \mapsto \varphi_T(u, \gamma) \tag{15}$$

for any value of T. A periodic solution of (14) is a fixed point of (15) only when T is a multiple of the (unknown) period. Solutions of (14) can therefore be

studied by analyzing fixed points of (15) instead. The stability of a fixed point of (15) is determined by the eigenvalues μ_l of

$$M := \frac{\partial \varphi_T(u, \gamma)}{\partial u}.$$

The fixed point is stable if all eigenvalues have modulus smaller than one. If u is a steady state of (14) then

$$M = \frac{\partial \varphi_T(u, \gamma)}{\partial u} = \exp\left(T \frac{\partial f(u, \gamma)}{\partial u}\right)$$

and hence

$$\mu_l = \exp(\lambda_l T). \tag{16}$$

Consequently, the stability information obtained with both approaches is equivalent. In practice, we have to use a numerical time integrator. Most classical time integration schemes preserve steady states of (14). However, (16) will only be satisfied approximately. If the step size is sufficiently small, the dominant eigenvalues of the numerical time integrator will be very good approximations to the eigenvalues μ_l of the exact time integrator and can be used to judge the stability of the computed fixed points. One can then still use (16) to compute approximations to the eigenvalues λ_l.

A LBM defines a map. Since a LBM is deterministic, one can use the same techniques to analyze this map as for the time integrator map. If the dynamics of the LBM are equivalent to those of a PDE, the dominant eigenvalues μ_l computed from the LBM and the rightmost eigenvalues λ_l computed from the equivalent PDE will also approximately satisfy (16). The same framework can also be used in combination with the coarse-grained time integrator to compute steady states of the unknown but assumed to exist macroscopic description and to analyze their stability.

In numerical bifurcation analysis, one typically computes a discrete set of points on a branch of steady states or fixed points obtained by varying one parameter of the ODE system (14) or map (15) while monitoring the stability-determining eigenvalues along the branch to detect bifurcation points. The tool for this is a continuation method. Given the already computed points on a branch, a prediction is made for the position of the

next point and that point is then computed by solving the nonlinear system

$$\begin{cases} \varphi_T(u, \gamma) - u = 0, \\ n(u, \gamma; \eta) = 0. \end{cases} \tag{17}$$

The vector u and one of the components of the parameter vector γ are the unknowns. The last equation, $n(u, \gamma; \eta) = 0$ is a scalar equation that determines the position of the point along the branch through a reparametrization with parameter η. In our experiments we used pseudo-arclength continuation [23].

When computing a branch of isolated periodic solutions, the period T is also unknown. We use single shooting to compute a point on the branch. This point is computed by adding a phase condition $p(u, T, \gamma) = 0$ to (17) which fixes the starting point for integration along the periodic orbit, i.e., the phase condition effectively defines a Poincaré section. The resulting system

$$\begin{cases} \varphi_T(u, \gamma) - u = 0, \\ p(u, T, \gamma) = 0, \\ n(u, T, \gamma; \eta) = 0, \end{cases} \tag{18}$$

is solved for the vector u, T and one component of γ. The stability is determined by the eigenvalues (now called *Floquet multipliers*) of M evaluated at the solution. One of the multipliers will be one and should not be taken into account. The others determine the stability of the periodic orbit. Instead of a numerical time integrator for (14), we can also use the coarse-grained time integrator or the LBM (with the extension to a continuous time variable). This approach to bifurcation analysis, using the time integrator map, can also be extended to compute periodic solutions of a Hamiltonian system or to compute traveling waves, etc. The problems that arise in those computations are similar to those that arise in other continuation methods.

In [8], Shroff and Keller proposed the Recursive Projection Method to compute solutions of (17). This procedure is essentially a stabilization and acceleration procedure for iterating with the map. This method was used in [3]. Though it is possible to extend RPM to compute periodic solutions also, we used the Newton–Picard method [9,10]. This method tends to be more robust and was originally developed for bifurcation analysis of periodic solutions of large systems. Both

methods are based on the assumption that M has only few eigenvalues close to or outside the unit circle. The Newton–Picard method starts from Newton's method, but computes only an approximate solution to each linearized system. At each step, the low-dimensional dominant subspace $\mathcal{U}$ of M (i.e., the subspace of all eigenvalues larger than some user-determined threshold θ, with $0 < \theta < 1$) is determined using orthogonal subspace iteration [24], the linearized system is projected on this space and its orthogonal complement $\mathcal{U}^\perp$, and the resulting system is solved approximately by combining fixed-point (or Picard) iterations in $\mathcal{U}^\perp$ with a direct solver in $\mathcal{U}$. The dominant eigenvalues are easily computed from the projection of M onto $\mathcal{U}$.

When using time stepper based bifurcation analysis to compute steady state solutions, T can be chosen freely. However, T should not be too large since the nonlinear behavior of the map (15) may become more pronounced, causing trouble when solving the nonlinear system (17). This is especially the case when computing unstable steady states and is essentially the same problem as encountered in single shooting methods for computing periodic solutions [25,26]. As can be seen from (16), T also influences the eigenvalues of M and thus the convergence of the Newton–Picard method. Assuming the threshold θ is squared so that the dimension of the subspace $\mathcal{U}$ remains the same, doubling T will roughly halve the number of orthogonal subspace iteration and Picard iteration steps, but each step will be twice as expensive. The overall computational cost remains roughly the same. The value of T is not critical for the Newton–Picard method, although we did observe some problems with the subspace convergence criterion we used if T becomes too small.

Both RPM and the Newton–Picard method need matrix–vector products with M. These can be computed using finite difference methods. These matrix–vector products must be accurate enough, otherwise it will be impossible to compute a basis for $\mathcal{U}$ with enough accuracy. Note that the subspace $\mathcal{U}$ may be better conditioned than individual eigenvalues associated with that subspace. Therefore, in order to compute the eigenvalues with enough accuracy to reliably determine the stability and bifurcation points, even higher accuracy of the matrix–vector products may be required. Problems are possible with the coarse-grained

time stepper if the underlying microscopic model is stochastic or has chaotic dynamics. Since differentiation is by itself an ill-conditioned operation, we will need a very low variance of the result of the coarse-grained time stepper to compute the matrix–vector products accurately enough. This may require a lot of microscopic simulations at each time step as was experienced in [21] for a system with an ODE-like one- to three-dimensional macroscopic state. If too many simulations are needed to obtain a low enough variance, the numerical bifurcation analysis using the coarse-grained time stepper may be too expensive and thus unfeasible. However, in our test case, the "microscopic" model has clear steady states and periodic solutions.

We use only a single LB simulation at each coarse-grained time step, and the initial conditions are prescribed in a deterministic way. Therefore, we have no problems computing the matrix–vector products. A stochastic microscopic simulator would be a better choice to study this aspect of coarse-grained bifurcation analysis. However, by avoiding this problem in our test case, we are in a much better position to study the influence of the initialization of the microscopic simulator and the macroscopic time step ΔT which is the subject of Section 6.2.

6. Study of the bifurcation diagrams

We will first compare the bifurcation diagrams for the PDE model, the LBM and the coarse-grained LB time integrator, using the set of weights (13) in the reconstruction step and a fairly large macroscopic time step. Next, we will study how the minimal macroscopic time step needed for accurate results depends on the weights used in the reconstruction step. This will show that the minimal macroscopic time step is often much larger than the healing time.

6.1. The reference bifurcation diagram

Fig. 1 presents the steady state bifurcation diagram for the LBM, the coarse-grained time integrator and the PDE model. On the vertical axis, we show $\int_0^L \rho^{ac}(x)\, dx$, computed using the midpoint quadrature rule. This quantity should be essentially the same for all three models and is also, up to discretization errors, independent of the grid size. For all computations, we used a grid with 200 grid cells ($\Delta x = 0.1$). All branches, also the branch for the PDE model, were computed using the time stepper based bifurcation approach explained in Section 5, using the publicly available package PDEcont [27]. We set $T = 5$. For this value of T, we had
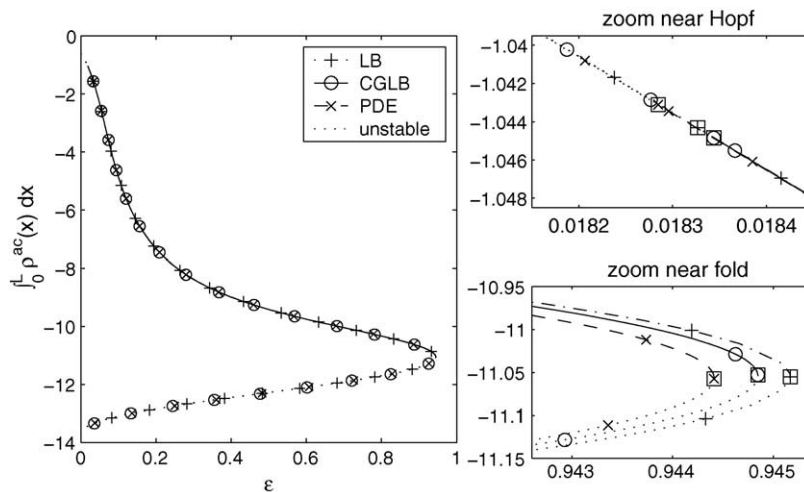

Fig. 1. Bifurcation diagram for the steady state solutions. "PDE" denotes the solution branch of the PDE system (2), "LB" of the LBM and "CGLB" of the coarse-grained LB time stepper (using $T = \Delta T = 5$). Unstable solutions are plotted using a dotted line. The markers represent only a subset of the points computed by the continuation code. The two figures on the right zoom in on the Hopf and the fold point. The bifurcation points are marked with a square.

no difficulties computing the unstable solutions. Also, for this T, the eigenvalues are well enough separated for the robust operation of our Newton–Picard implementation. Note also that the computed steady states do not depend on the value of T. The LB time step Δt was set equal to 0.001. The time step for the trapezoidal rule was 0.02. Smaller values of Δt only led to minor changes (on the order of the spatial discretization error) in the bifurcation diagram, while the bifurcation diagram changed considerably for larger values of Δt. In other words, the LB bifurcation diagram has converged for $\Delta t = 0.001$. Also, the correspondence with the PDE bifurcation diagram is very good. To compute the fixed points of the coarse-grained integrator, we set $\Delta T = 5$ and used the weights (13) in the reconstruction operator. Note that the results do depend on the choice of ΔT, see Section 6.2.4. With this choice of parameters, the difference with the PDE results is also of the order of the discretization error. We will study the influence of these parameters in more detail in Section 6.2. There is a fold point at $\varepsilon \approx 0.945$ and a supercritical Hopf bifurcation at $\varepsilon \approx 0.0183$ giving rise to stable periodic orbits at smaller parameter values. The location of the bifurcation points corresponds very well for all three approaches.

In Fig. 2, we show the branch of periodic solutions emanating from the Hopf point. Since the phase condition was not the same for all simulations, it is impossible to compare the computed point on the orbit

for all three models. Therefore, we now plot the period T on the vertical axis. To compute periodic solutions using the coarse-grained integrator, we set $\Delta T \approx 5$, cf., the second approach in Section 4.2. For the trapezoidal rule, we used 1000 time steps per period. There is a fold point at $\varepsilon \approx 0.00087$. Solutions on the unstable part of the branch have (at least initially) almost the same parameter-period dependence as on the stable part, but the periodic orbit is slightly different. The amplitude of the temporal oscillations of the inhibitor is larger for the unstable periodic solution than for the stable one at the same parameter value. We did not succeed in computing the unstable branches far past the fold point. This demonstrates the lack of robustness of single shooting methods, in particular when computing unstable solutions [25,26]. The Hopf points for the three methods are slightly different from those in Fig. 1 because they are computed from a different eigenvalue problem. The results in Fig. 2 are less accurate since the location of the Hopf point is determined by detecting an ill-conditioned algebraically double eigenvalue at 1.

We can draw two conclusions from this section. First, the steady states and periodic solutions of the LBM correspond very well to those of the macroscopic PDE. Second, computing a bifurcation diagram for the coarse-grained time integrator does work and produces the expected results (at least for this choice of ΔT and reconstruction scheme).


Fig. 2. Bifurcation diagram for the periodic solutions. The labels and markers are the same as in Fig. 1.

6.2. Influence of the reconstruction step and the macroscopic time step

6.2.1. The slaved state

In Fig. 3, we study the slaving of momentum and energy at the stable steady state of the LB model at $\varepsilon = 0.05$. For both the activator and the inhibitor (not shown), the momentum is small compared to the concentration. When investigating the solution more carefully, one notes that the momentum is in fact proportional to the gradient of the concentration, i.e.,

$$\phi^s \approx -d^s \rho_x^s. \tag{19}$$

We computed $d^s = \|\phi^s\|/\|\rho_x^s\|$, where $\|\cdot\|$ denotes the two-norm of the discrete state vector. The ratio d^s is essentially constant along the solution branches with $d^{ac} \approx 0.04338$ and $d^{in} \approx 0.07334$. Fig. 3 also clearly shows that the (dimensionless) kinetic energy ξ^s is almost perfectly one third of the concentration.

These relationships can be proven quite easily for a diffusion problem. In the latter case, the LB variables can be written up to first order terms as [28]

$$f_i^s = f_i^{s,\text{eq}} + f_i^{s,\text{neq}} = v_i \rho^s - \frac{v_i i \Delta x}{\omega^s} \rho_x^s,$$

where $f_i^{s,\text{neq}}$ is the deviation from the equilibrium distribution (the "non-equilibrium part"). The corresponding momentum (4) and energy (5) are

$$\phi^s = 2 f_1^{s,\text{neq}} = -\frac{2 v_1 \Delta x}{\omega^s} \rho_x^s,$$
$$\xi^s = f_1^{s,\text{eq}} = v_1 \rho^s. \tag{20}$$

Substituting the parameter values of our problem into (20), we obtain $\phi^{ac} = -0.04333 \rho_x^{ac}$, $\phi^{in} = -0.07333 \rho_x^{in}$ and $\xi^s = (1/3)\rho^s$. So far, we have no proof of these relations for reaction-diffusion problems, but they appear to hold at least for our example.

Eq. (20) enables us to develop an almost perfect reconstruction scheme by first computing the slaved (i.e., on-manifold) values of $\phi^s(x_j, 0)$ and $\xi^s(x_j, 0)$ from $\rho^s(x_j, 0)$ using (20) and numerically approximating ρ_x^s by finite differences. The corresponding distributions $f_i^s(x_j, 0)$ can then be obtained from definitions (3)–(5). However, in the remainder of this paper we will focus on the behavior of the coarse-grained integrator using initializations away from the "slow" manifold.

6.2.2. Initialization of the LB model in the coarse-grained time stepper

We studied several reconstruction schemes in our experiments. Three schemes are based on (12) but use different sets of weights. Initialization with the local diffusive equilibrium distribution, i.e.,

$$w_{-1} = w_0 = w_1 = \frac{1}{3} \tag{21}$$

in (12) is a straightforward choice. For this choice, the kinetic energy is almost perfectly slaved. The

Fig. 3. Slaving of the activator higher-order moments for the stable steady state at $\varepsilon = 0.05$. Note that we plot $25\phi^{ac}$ rather than ϕ^{ac} in the top figure.

momentum is identically zero. Though the momentum is small in the correctly slaved state also, this reconstruction scheme – as all others based on (12) – does not satisfy (19). For the second reconstruction scheme, we chose the symmetric set of weights

$$w_{-1} = w_1 = 0.01, \quad w_0 = 0.98. \tag{22}$$

The momentum is still zero and thus "close" to the correctly slaved state, but the kinetic energy (the second-order velocity moment) is very different from the correctly slaved state. Our third scheme uses the asymmetric weights

$$w_{-1} = 0.75, \quad w_0 = 0.24, \quad w_1 = 0.01. \tag{23}$$

For this choice, both the first-order and second-order velocity moments, i.e., both momentum and kinetic energy, differ significantly from the correctly slaved state.

Though we can initialize the kinetic energy very well with reconstruction scheme (12), the momentum cannot be initialized correctly unless a more complicated scheme such as (20) is used. Since this relation is as yet unproven for our class of problems, and to avoid any small error resulting from the approximations made in the derivation of (20) and the computation of ρ_x^s, we also performed experiments with a coarse-grained time stepper using both concentration and momentum as the macroscopic variables. This leaves only one degree of freedom per species and per lattice point for the initial state of the LB model. In this case, the reconstruction scheme is

$$f_{-1}^s = \frac{1}{2}(1 - w_0^*)\rho^s - \frac{1}{2}\phi^s, \quad f_0^s = w_0^*\rho^s \quad \text{and}$$

$$f_1^s = \frac{1}{2}(1 - w_0^*)\rho^s + \frac{1}{2}\phi^s \tag{24}$$

where w_0^* can be chosen freely. For this initialization,

$$\xi^s = \frac{1}{2}(1 - w_0^*)\rho^s.$$

We considered two choices for the parameter w_0^*. For the first choice,

$$w_0^* = 0.98, \tag{25}$$

the energy differs significantly from the correctly slaved state. The second choice,

$$w_0^* = \frac{1}{3}, \tag{26}$$

provides an (almost) perfect initialization of all velocity moments.

6.2.3. The healing process

In Fig. 4, we study the healing process. Diagram (a) shows the difference between the momentum and the scaled concentration gradient while diagram (b) shows the difference between the kinetic energy and one third of the concentration for the activator at the lattice point $x = 9.95$ in the first few LB time steps using the three reconstructions based on (12), i.e., both diagrams show the deviation from the best available approximation of the correctly slaved state. The macroscopic initial state is the stable steady state of the LB model at $\varepsilon = 0.05$. From this state, we generated microscopic initial conditions using (12) with weights (21)–(23). The figures show that both the momentum and kinetic energy become slaved in about 10 to 15 LB time steps of $\Delta t = 0.001$. This claim is further verified by the experiment in Fig. 5. If the higher-order moments are correctly slaved for a given state, the evolution from then on could be described by a macroscopic model with the lower-order moments as the unknowns. Since our LB model is designed to correspond to the (macroscopic) FitzHugh–Nagumo PDE model (2), it is clear what the macroscopic model should be in this case. Hence we started time integration of the PDE model from the LB trajectory obtained for initialization (12) with weights (23), at $t = 0.002 = 2\Delta t$ (i.e., before slaving is obtained) and at $t = 0.02 = 20\Delta t$ (after obtaining slaving). The PDE trajectory started from the LB trajectory at $t = 0.002$ differs significantly from the LB trajectory, while the PDE trajectory started at $t = 0.02$ follows the LB trajectory more closely. Note that the LB and PDE trajectories converge to a slightly different steady state as discussed in Section 6.1. This experiment confirms that after 20 time steps, the higher-order moments are slaved by the lower-order ones, while this is not yet the case after 2 time steps.

6.2.4. The bifurcation diagram

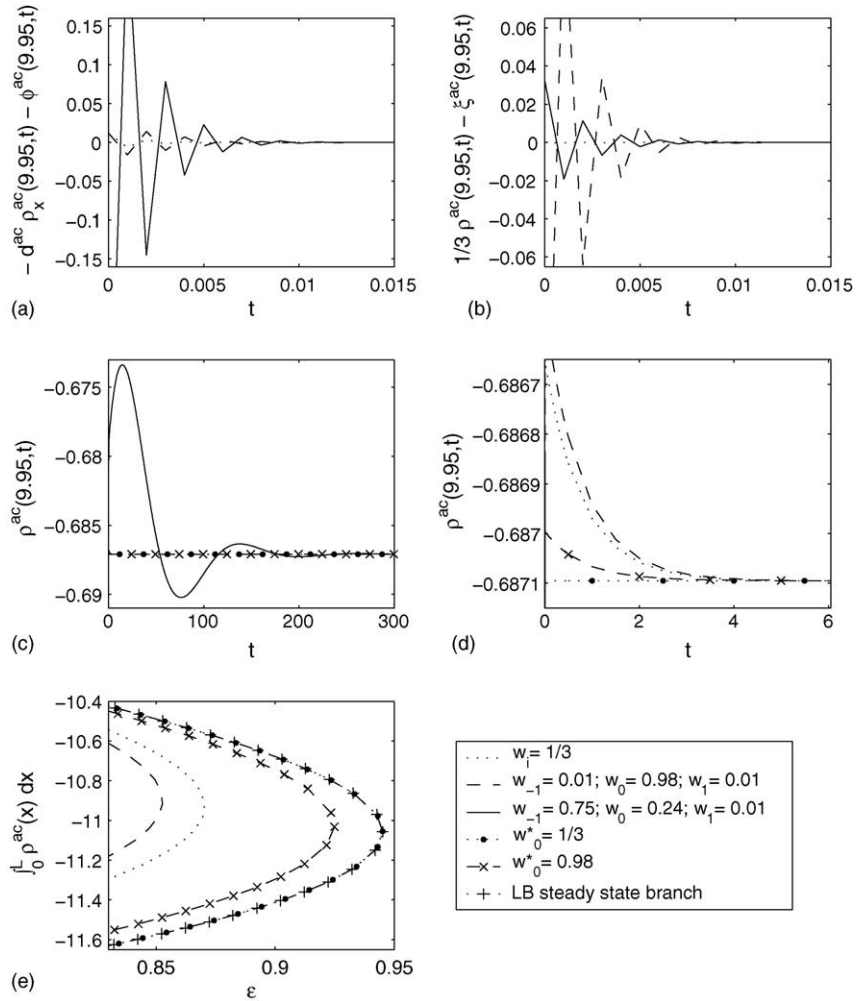The fact that slaving is obtained so quickly, suggests that a coarse-grained time step $\Delta T = 20\Delta t = 0.02$

Fig. 4. Healing of the lifting error. Evolution of (a) $-d^{ac}\rho_x^{ac}(x, t) - \phi^{ac}(x, t)$, (b) $(1/3)\rho^{ac}(x, t) - \xi^{ac}(x, t)$ and (c) concentration $\rho^{ac}(x, t)$ at lattice site $x = 9.95$ for the LB trajectories started from the stable steady state at $\varepsilon = 0.05$ using the reconstruction schemes (12) and (24). (d) is a close-up of (c). (e) Steady state bifurcation diagrams for the coarse-grained integrator using $\Delta T = 0.02$. The steady state bifurcation diagram for the LBM from Fig. 1 is also shown for comparison.

would be sufficient to compute the bifurcation diagram accurately. Diagram (e) in Fig. 4 shows the bifurcation diagrams near the fold point computed with this time step. The results for all reconstructions except (24) with (26) have an unacceptably large error. In fact, the line for (12) with weights (23) even falls off the figure. The reason for this can be seen in Fig. 4, diagram (c) and (d). Though slaving is obtained quickly, the LB simulation does not follow the intended trajectory, i.e., the trajectory that would be followed by a macroscopic

model using the same macroscopic initial condition. (In this experiment, we initialized from a steady state, so the correct trajectory is constant.) In the healing process, the lower-order moments change also and even at a fairly fast time scale. At the end of the healing process, these lower-order moments are different from what they would have been for a "perfect" initialization, and so the trajectories differ. Since we are in the neighborhood of a stable steady state, all trajectories ultimately converge to the steady state. However, it takes

Fig. 5. Trajectory of $\rho^s(9.95, t)$ for a LB simulation started from the initial condition obtained by reconstruction using the asymmetric reconstruction scheme, and for two PDE simulations started from the LB trajectory at $t = 0.002$ and $t = 0.02$.

about five time units (5000 LB time steps) with most reconstruction schemes to return to the steady state, while with the reconstruction scheme (12) with the asymmetric weight choice (23) (i.e., both momentum and kinetic energy are badly initialized), it even takes on the order of 300 time units. Therefore, the coarse-grained time step ΔT must be much larger than the healing time unless an accurate reconstruction scheme is used.

In Fig. 6 we show the bifurcation diagram obtained with the initialization (12) with weights (21) (the equilibrium distribution). The left panel shows the bifurcation diagram near the fold point for different values of ΔT. In the right diagram, we plot the estimated discretization error for the stable steady state at $\varepsilon = 0.93$, close to the fold point. The estimate was obtained by
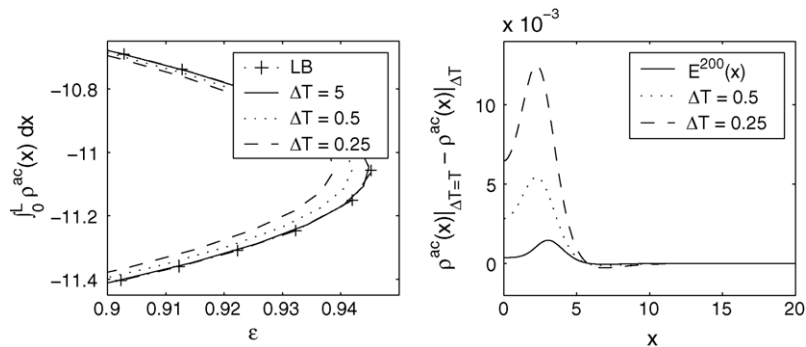
comparing with a LB steady state on a lattice with 16,200 lattice points. As ΔT increases, the computed equilibria and bifurcation diagram become more accurate. For $\Delta T = 0.5$, the error of the activator concentration is on the order of four times the discretization error and the bifurcation diagram is also acceptable. For $\Delta T = 5$, the bifurcation diagram is virtually the same as for the LB model. This agrees with Fig. 4(d), where it took about five time units for the LB simulator to converge to the correct trajectory. However, near an unstable solution, the trajectories diverge and one would expect that the results would only get worse as ΔT is increased. This is true when plotting the trajectories, but when computing fixed points, we still notice an improvement as ΔT is increased.

In Fig. 7 we show the bifurcation diagram obtained using (12) with the symmetric weight set (22) and the asymmetric one (23). With the symmetric weight set, we again obtain sufficiently accurate results for $\Delta T = 5$. For the asymmetric weight set, the results also get better as ΔT increases, but only become acceptable when $\Delta T = 75$. Though we can compute fairly accurate solutions even for this bad initialization, the coarse-grained time step ΔT and hence the time integration interval T for the Newton–Picard method becomes so large that it is hard to compute the unstable solutions as we already pointed out in Section 5. The computations break down at $\varepsilon \approx 0.636$ for $\Delta T = 25$ and at $\varepsilon \approx 0.928$ for $\Delta T = 75$.

Clearly, obtaining a correctly slaved state by the end of the microscopic integration in the coarse-grained

Fig. 6. Left: Steady state solution branches for the coarse-grained integrator using different values of ΔT. The reconstruction scheme (12) with the equilibrium distribution is used. Right: The estimated discretization error $E^{200}(x)$ for the coarse-grained integrator ($\Delta T = T = 5$) with 200 discretization points for a steady state at $\varepsilon = 0.93$ compared with the difference between the coarse-grained steady state using $\Delta T = 5$ and the corresponding states using $\Delta T = 0.5$ or $\Delta T = 0.25$.

Fig. 7. The steady state bifurcation diagram for the coarse-grained integrator with both a symmetric ($w_{-1} = 0.01$, $w_0 = 0.98$ and $w_1 = 0.01$) and asymmetric ($w_{-1} = 0.75$, $w_0 = 0.24$ and $w_1 = 0.01$) set of reconstruction weights. The diagram for the coarse-grained integrator with asymmetric reconstruction scheme is computed for three different macroscopic time steps. The steady state branch obtained with the LBM from Fig. 1 is also shown. The right figure zooms in on the area near the fold point.

time stepper is not sufficient to obtain accurate results. If the reconstruction is not very good, the microscopic simulator must be run over a much larger time interval ΔT. In this test case, there are no upper limits on this time interval other than those imposed by the Newton–Picard procedure. However, as we pointed out in Section 4.1, other microscopic models, and in particular stochastic simulations, may impose a more severe upper bound on ΔT. In these cases it may be impossible to compute an accurate bifurcation diagram unless a very good reconstruction scheme is used. We expect that the quality of the reconstruction will be even more important when using more advanced simulation schemes such as the projective integration and gap-tooth schemes suggested in [2,4]. In fact, unless the higher-order moments are initialized near-perfectly, it may be impossible to simulate trajectories accurately near unstable equilibria with those techniques. Correctly initializing the microscopic simulations is clearly an important area of further research. As shown in [29,30], ideas from approximate inertial manifolds may be useful here.

7. The spectra

7.1. Stability analysis

In Section 6.1 we noticed that the bifurcation diagrams for the PDE model, LB model and coarse-grained integrator are virtually the same, including the location of the bifurcation points. The latter fact indicates that the dominant, stability-determining eigenvalues will match very well also. We will now study this in more detail.

The dominant eigenvalues are computed in the Newton–Picard code by performing a number of additional orthogonal subspace iteration steps after the computation of the fixed point. In Table 1, we list the dominant eigenvalues for the unstable steady state at $\varepsilon = 0.01$ on the upper part of the branch in Fig. 1. We used $T = \Delta T = 5$ and transformed the eigenvalues to exponent form using (16). Table 1 also lists the trivial Floquet multiplier and the most dominant non-trivial multiplier for the stable periodic solution at the same parameter value. The existence of a trivial multiplier

Table 1

Dominant eigenvalues for the unstable steady state on the upper part of the branch and stable periodic solution at $\varepsilon = 0.01$ (using $\Delta T \approx 5$ in the CGLB integrator)

|	Steady state		Periodic solution	
	$\lambda_{1,2}$	λ_3	trivial μ_1	μ_2
LB	$0.002010 \pm 0.039461i$	-0.124867	1.000000	0.514888
CGLB	$0.002012 \pm 0.039463i$	-0.124863	1.000000	0.514452
PDE	$0.001999 \pm 0.039446i$	-0.124861	1.000000	0.516712

Fig. 8. Left: The full spectrum for the LB and discretized PDE model for the stable steady state at $\varepsilon = 0.05$. Right: Close-up of the most dominant eigenvalues.

at one is a general property of autonomous systems. Its computational accuracy is independent of the spatial discretization error. The remarkable precision of the computed value indicates that the time integration and eigenvalue computation are very accurate. Clearly, the eigenvalues (and also the corresponding eigenvectors) correspond very well for all three models.

7.2. Slaving and the spectrum of the lattice Boltzmann model

To conclude, we study the full spectrum of the LBM and the discretized PDE model. We computed the Jacobian matrix analytically for both cases. To compare

with the eigenvalues of the LBM, the eigenvalues λ_l obtained for the PDE were transformed to multiplier form using (16) with $T = \Delta t = 0.001$, the LB time step. The results for the stable steady state at $\varepsilon = 0.05$ are shown in Fig. 8. The LBM has 400 eigenvalues in the same zone along the real axis as the discretized PDE. However, only the dominant eigenvalues correspond well. This is not surprising. The less dominant eigenvalues depend very much on the discretization and have little relationship with the true eigenvalues of the continuous problem.

At first, one would expect to recognize slaving of the first- and second-order moment to the zeroth-order moment of the eigenvectors of those 400 LB eigenvalues, while in the other eigenvectors there would clearly


Fig. 9. Slaving of the activator momentum and energy of the (real part of the) full LB eigenvectors for the largest complex pair of eigenvalues and the first real eigenvalue from Fig. 8. Left: Difference between the eigenvector's momentum and its appropriately scaled concentration gradient. Right: Difference between the eigenvector's energy and one third of the concentration.

be no slaving. However, we only observed slaving in the eigenvectors for the most dominant eigenvalues that correspond very well to those of the discretized PDE. Fig. 9 illustrates this slaving for the real part of the eigenvectors corresponding to the rightmost complex pair of eigenvalues and for the eigenvector corresponding to the largest real eigenvalue. The eigenvector's momentum is small compared to its concentration and proportional to its concentration gradient, and its second-order moment, the energy, is very nearly one third of the concentration, so we note the same slaving relationships as for the state in Section 6.2. The discovery of such relationships between the higher-order and lower-order moments could be a step towards the development of some kind of constitutive equation or closure relation.

8. Conclusions

In this paper, we have studied the coarse-grained bifurcation analysis procedure proposed in [3], using the same test case, a FitzHugh–Nagumo lattice Boltzmann model. We have extended the work of [3] in several ways. We compared the results of a numerical bifurcation analysis using the coarse-grained time integrator not only with results for an equivalent PDE, but also with the bifurcation diagram for the (deterministic) LB model used in the coarse-grained time stepper. The results for all three approaches corresponded very well. We have also extended the coarse-grained integrator to produce results at an (almost) arbitrary time *T*. This enabled the computation of periodic solutions. Instead of the Recursive Projection Method used in [3], we used the Newton–Picard method [9,10]. Though the coarse-grained time stepper is not really needed to perform a numerical bifurcation analysis of the LB model, this test case did enable us to thoroughly study the effects of the reconstruction scheme. This led to the most important conclusion of this paper. Contrary to the claim in [3] that the quality of the reconstruction step does not really matter, we have shown that this step can be crucial to the success of the method. Though slaving is quickly obtained irrespective of the reconstruction scheme, with a bad initialization the trajectory of the microscopic simulator may be quite different from the intended one (the one which corresponds to the trajectory of a macroscopic model,

if such a model would be known explicitly, started from the same initial macroscopic state). Hence, good reconstruction schemes are clearly problem-dependent and are an interesting area of further research, see e.g., [29,30].

We have also demonstrated that the techniques developed for time stepper based numerical bifurcation analysis of PDEs can be used for bifurcation analysis of steady states and periodic solutions of LB models using either the coarse-grained integrator or the LB model itself as the time stepper. As shown in [31], the amount of work when using the Newton–Picard method is roughly the same for both approaches, since this is mostly determined by the dominant eigenvalues. Since the state vector is lower-dimensional for the coarse-grained time stepper, the memory requirements will be less. However, this approach is much more complicated than bifurcation analysis using the LB model itself as the time stepper, since a good choice of the macroscopic variables must be made and a good reconstruction is needed for accurate results.

We have also studied the spectrum of the LB model and showed that the higher-order moments of the full eigenvectors are slaved in the same way as those of the corresponding LB solution.

This paper does not claim that numerical bifurcation analysis based on the coarse-grained time stepper of [2,3] will always work. Indeed, microscopic or mesoscopic simulations based on stochastic models or models with chaotic behavior, may (and likely will) pose additional numerical problems that cannot be studied with this simple test case. Further research is needed in this area. However, it does show that the idea of initializing microscopic simulators from a macroscopic state can produce valid macroscopic data already after a short time interval, provided the reconstruction of the microscopic state is done properly.

Another possible extension of this work is the combination with more efficient simulation techniques such as the schemes in [2,4,5].

References

[1] C.W. Gear, I.G. Kevrekidis, C. Theodoropoulos, 'Coarse' integration/bifurcation analysis via microscopic simulators: micro-Galerkin methods, Comput. Chem. Eng. 26 (7–8) (2002) 941–963.

[2] I.G. Kevrekidis, C.W. Gear, J.M. Hyman, P.G. Kevrekidis, O. Runborg, C. Theodoropoulos, Equation-free, coarse-grained multiscale computation: enabling microscopic simulators to perform system-level analysis, Commun. Math. Sci. 1 (4) (2003) 715–762.

[3] C. Theodoropoulos, Y.H. Qian, I.G. Kevrekidis, "Coarse" stability and bifurcation analysis using time-steppers: a reaction-diffusion example, Proc. Natl. Acad. Sci. 97 (18) (2000) 9840–9843.

[4] G. Samaey, D. Roose, I.G. Kevrekidis, The gap-tooth scheme for homogenization problems, Multiscale Model. Simul. 4 (1) (2005) 278–306.

[5] W. E, B. Engquist, The heterogeneous multiscale methods, Commun. Math. Sci. 1 (1) (2003) 87–133.

[6] E.J. Doedel, R.C. Paffenroth, A.R. Champneys, T.F. Fairgrieve, Y.A. Kuznetsov, B.E. Oldeman, B. Sandstede, X. Wang, AUTO2000: continuation and bifurcation software for ordinary differential equations (with HomCont), Tech. rep., Concordia University, 2002.

[7] A. Dhooge, W. Govaerts, Y.A. Kuznetsov, MATCONT: A MATLAB package for numerical bifurcation analysis of ODEs, ACM Trans. Math. Software 29 (2) (2003) 141–164.

[8] G.M. Shroff, H.B. Keller, Stabilization of unstable procedures: the Recursive Projection Method, SIAM J. Numerical Anal. 30 (4) (1993) 1099–1120.

[9] K. Lust, D. Roose, Computation and bifurcation analysis of periodic solutions of large-scale systems, in: E. Doedel, L. Tuckerman (Eds.), Numerical Methods for Bifurcation Problems and Large-Scale Dynamical Systems, The IMA Volumes in Mathematics and its Applications, vol. 119, Springer-Verlag, New York, 2000, pp. 265–302.

[10] K. Lust, D. Roose, A. Spence, A. Champneys, An adaptive Newton–Picard algorithm with subspace iteration for computing periodic solutions, SIAM J. Sci. Comput. 19 (4) (1998) 1188–1209.

[11] Y.H. Qian, S.A. Orszag, Scalings in diffusion-driven reaction $A + B \rightarrow C$: Numerical simulations by Lattice BGK Models, J. Stat. Phys. 81 (1–2) (1995) 237–253.

[12] I. Ginzbourg, P.M. Adler, Boundary flow condition analysis for the three-dimensional lattice Boltzmann model, J. Phys. II (France) 4 (1994) 191–214.

[13] X. He, Q. Zou, L.-S. Luo, M. Dembo, Analytic solutions of simple flows and analysis of nonslip boundary conditions for the lattice Boltzmann BGK model, J. Stat. Phys. 87 (1/2) (1997) 115–136.

[14] H. Chen, S. Chen, W.H. Matthaeus, Recovery of the Navier-Stokes equations using a lattice-gas Boltzmann method, Phys. Rev. A 45 (8) (1992) R5339–R5342.

[15] Y.H. Qian, D. D'Humières, P. Lallemand, Lattice BGK models for Navier–Stokes equation, Europhys. Lett. 17 (6) (1992) 479–484.

[16] B. Chopard, A. Dupuis, A. Masselot, P. Luthi, Cellular automata and lattice Boltzmann techniques: an approach to model and simulate complex systems, Adv. Complex Syst. 5 (2/3) (2002) 103–246.

[17] S. Succi, The Lattice Boltzmann Equation for Fluid Dynamics and Beyond, Numerical Mathematics and Scientific Computation, Oxford University Press, 2001.

[18] S.P. Dawson, S. Chen, G.D. Doolen, Lattice Boltzmann computations for reaction-diffusion equations, J. Chem. Phys. 98 (2) (1993) 1514–1523.

[19] D. Dab, J.-P. Boon, Y.-X. Li, Lattice-Gas Automata for coupled reaction-diffusion equations, Phys. Rev. Lett. 66 (19) (1991) 2535–2538.

[20] A.C. Hindmarsch, ODEPACK, A systematized collection of ODE solvers, in: R. Stepleman, M. Carver, R. Peskin, W. Ames, R. Vichnevetsky (Eds.), Scientific Computing, IMACS Transactions on Scientific Computing, vol. 1, North-Holland, Amsterdam, 1983, pp. 55–64.

[21] A.G. Makeev, D. Maroudas, I.G. Kevrekidis, "Coarse" stability and bifurcation analysis using stochastic simulators: Kinetic Monte Carlo examples, J. Chem. Phys. 116 (23) (2002) 10083–10091.

[22] A.G. Makeev, D. Maroudas, A.Z. Panagiotopoulos, I.G. Kevrekidis, Coarse bifurcation analysis of kinetic Monte Carlo simulations: A lattice-gas model with lateral interactions, J. Chem. Phys. 117 (18) (2002) 8229–8240.

[23] H.B. Keller, Numerical solution of bifurcation and nonlinear eigenvalue problems, in: P.H. Rabinowitz (Ed.), Applications of Bifurcation Theory, Academic Press, New York, 1977, pp. 359–384.

[24] Y. Saad, Numerical methods for large eigenvalue problems, Algorithms and architectures for advanced scientific computing, Manchester University Press, Manchester, 1992.

[25] H. Keller, Numerical methods for two-point boundary value problems, Blaisdell, New-York, 1968.

[26] U. Ascher, R. Mattheij, R. Russell, Numerical solution of boundary value problems for ordinary differential equations, Prentice-Hall, Englewood Cliffs, NJ, 1988.

[27] K. Lust, PDEcont, URL: http://www.math.rug.nl/~kurt/r_PDEcont.html.

[28] R.G.M. van der Sman, M.H. Ernst, Convection-diffusion lattice Boltzmann scheme for irregular lattices, J. Comput. Phys. 160 (2) (2000) 766–782.

[29] C.W. Gear, I.G. Kevrekidis, Constraint-defined manifolds: a legacy code approach to low-dimensional computation, Tech. Rep. physics/0312094, arXiv e-Print archive, 2003.

[30] C.W. Gear, T.J. Kaper, I.G. Kevrekidis, A. Zagaris, Projecting to a slow manifold: singularly perturbed systems and legacy codes, SIAM J. Appl. Dyn. Syst., 2005, in press.

[31] P. Van Leemput, K. Lust, Numerical bifurcation analysis of lattice Boltzmann models: a reaction-diffusion example, in: M. Bubak, G.D. van Albada, P.M. Sloot, J. Dongarra (Eds.), Computational Science—ICCS 2004, Lecture Notes in Computer Science, vol. 3039, Springer-Verlag, 2004, pp. 572–579.