CELLULAR AUTOMATON MODEL OF COUPLED MASS TRANSPORT AND CHEMICAL REACTIONS

T. Karapiperis

Laboratory for Waste Management

Abstract

Mass transport, coupled with chemical reactions, is modelled as a cellular automaton in which solute molecules perform a random walk on a lattice and react according to a local probabilistic rule. Assuming molecular chaos and a smooth density function, we obtain the standard reaction-transport equations in the continuum limit. The model is applied to the reactions $a+b \leftrightarrow c$ and $a+b \rightarrow c$, where we observe interesting macroscopic effects resulting from microscopic fluctuations and spatial correlations between molecules. We also simulate autocatalytic reaction schemes displaying spontaneous formation of spatial concentration patterns. Finally, we propose and discuss the limitations of a simple model for mineral-solute interaction.

1 Introduction

The need to account for the coupling between transport and chemical reactions arises when modelling a wide variety of processes in geological media. These processes may be natural (e.g. diagenesis connected with ground-water flow), or they may be the result of engineered conditions (e.g. oil recovery processes or contaminant migration). In particular, the transport of radionuclides released from a nuclear waste repository will be coupled with a variety of physical and chemical processes: radioactive decay and chemical reactions among the dissolved nuclides; sorption on the surfaces of water-conducting fractures, matrix pores or other particles (e.g. colloids); changes in the porosity of the host rock by precipitation/dissolution. The effectiveness of the surrounding geological medium as a potential barrier to radionuclide migration depends crucially on such processes. Modelling coupled transport and reactions in their full complexity constitutes a formidable task. A key insight, in this case, is the realisation that the complex macroscopic behaviour that one wishes to describe arises from the simultaneous evolution of a vast number of tiny, spatially-distributed parts of the system of interest that hardly interact with each other (they do so only if they happen to be sufficiently close). The dynamics that determines (sic) the evolution of these subsystems (which, in an extreme case, might be individual molecules of a migrating solute) is much simpler than that of the full system. The subsystems evolve in parallel (i.e. simultaneously at different locations) and interact locally (i.e. with subsystems within a small spatial neighbourhood). Under these circumstances, it is natural to model the system of reacting particles as a cellular automaton (CA).

A CA is a dynamical system which consists of a *discrete-valued field*, defined on the sites of a regular spatial lattice, and a *local rule* that determines the time evolution of the field. The rule is applied in discrete iterative steps and each time determines the next value of the field on a given site in terms of the present field values at a small neighbourhood of sites [1]. Our CA describes a system of molecules in liquid and solid phase, belonging to various species and subject to arbitrary chemical reactions with each other. The state of the system is defined by a set of *occupation numbers*, i.e. an integer-valued field with a species label giving the number of molecules of the different species at each lattice site. The microscopic evolution rule consists of a 'transport' operation followed by a 'chemical reaction' operation. Each operation is applied synchronously to all the lattice sites. The transport algorithm prescribes the probability with which a molecule in the mobile phase will move to a neighbouring site; the probability may depend on species label, time, location on the lattice or spatial direction of the displacement; upon iteration of the transport operation each particle performs a random walk. If sufficient molecules of different species meet on a lattice site, they react with a probability that reflects a prescribed microscopic reactive scheme; the chemical reaction operation is defined once the algorithm for calculating the above probability is given. Sequential application of the two operations completes an evolution step. Starting from an appropriate distribution of molecules on the lattice (initial condition), we apply the evolution rule iteratively up to the desired time (subject to certain boundary conditions) and then extract the final values of the occupation numbers. To derive macroscopic quantities we average over an *ensemble* of independent copies of the system of interest that evolve in parallel. Thus *particle density* is defined by averaging over the occupation number. Alternatively, or in combination, local space and time averages may be performed.

The occupation numbers of our CA obey evolution equations that lead to discrete (in time and space) dynamical equations for the particle densities if certain physical assumptions hold (see discussion below). In the continuum limit (i.e. when we let the lattice spacing and the time step shrink to zero), these equations go over to a set of reaction-transport *partial differential equations (PDE's)*. The latter contain the standard advection and diffusion terms describing transport of the mobile

molecules, as well as reaction-rate terms that are, in general, non-linear functions of the densities. Conventionally, the transport and reactions of solutes in geological media are modelled by systems of PDE's. Apart from a few rare cases, the reaction-transport PDE's cannot be solved by analytical means. Thus, the continuum equations are approximated by discrete equations, such as *finite difference equations (FDE's)*, which can be implemented on general-purpose computers. Then a substantial amount of effort goes into testing the stability and convergence properties of the numerical algorithms. Instabilities arise when, for some choice of the physical and discretisation parameters, round-off errors accumulate exponentially in the course of the iteration and result in computer overflows. Although general stability criteria are available for the FDE's approximating linear PDE's (e.g. pure transport), establishing stability has to be tackled on a case by case basis for non-linear PDE's.

Our CA model applies directly to its elementary constituents those properties of the microscopic dynamics that are essential to the macroscopic behaviour we wish to describe. Thus *microscopic fluctuations and correlations due to chemical reactions are naturally carried through all stages of our simulations*. By contrast, in order to arrive at the macroscopic reaction-transport equations: (i) we neglect correlations between molecules of different species (Boltzmann's 'molecular chaos' assumption), (ii) we define macroscopic quantities by averaging away fluctuations, and (iii) we assume that the spatial dependence of the microscopic quantities is sufficiently smooth to justify the equivalence between the definitions of macroscopic variables as either ensemble averages or local averages. Keeping the microscopic information is, however, sometimes essential in order to obtain the correct macroscopic behaviour, as we shall see later. While being closer to physical reality, our model does not waste resources by accounting for inessential details of the microscopic dynamics. The 'particles' of our model are mathematical abstractions of the actual molecules. The number of particles is large enough to make statistical concepts meaningful, but is still many orders of magnitude smaller than the real number of molecules involved. In this sense, our approach is to be seen as intermediate between the *macroscopic* approach of PDE's and the *microscopic* one of molecular dynamics.

An important property of our CA model is that the evolving field (i.e. the occupation number) is integer-valued and is exactly represented by computer words. *The absence of round-off errors guarantees the stability of our algorithms* in sharp contrast to the tedious procedure of establishing the stability of non-linear algorithms iterating real-valued fields (e.g. FDE for the particle density).

The parallel and local character of cellular automata maps naturally onto the architecture of massively parallel computers, in which a large number of processors run in parallel and communication is very fast between neighbouring processors (for example, one can envisage one processor being assigned to every lattice site). The massively parallel paradigm is bound to dominate computer architectures in the coming decades, as it offers the only way to surpass the physical limitations (e.g. speed of light) of serial and vector machines. Of course, when microscopic effects are uninteresting, other numerical approaches (e.g. FDE) can be implemented on massively parallel computers and the relative efficiency of different algorithms will depend on the specific applications.

Probabilistic CA models similar to ours have been applied to the treatment of various reaction-diffusion processes [2,3]. These models were inspired by the *lattice-gas automata (LGA)* that are used to simulate fluid motion [4]. By freeing itself of some of the standard features of LGA, our model is able to surpass certain serious limitations of prior models. Most importantly, *we see no reason to impose an a priori limit on the number of particles per site* [1] and, thus, unnecessarily constrain the statistics of our simulations (the number of particles per site still remains finite in any given simulation). Moreover, *advection arises naturally from our microscopic rule*, whereas models with an exclusion principle cannot model advection without introducing unwanted non-linearities. Finally the connection between our CA reaction rules with the prescribed molecular reaction scheme is more transparent than in reactive LGA and facilitates the modelling of complex chemical reactions.

In Section 2, we introduce formally the basic ingredients of our CA model and outline the derivation of the continuum equations; we then describe applications to various chemical reaction schemes, emphasising the effects of microscopic fluctuations and correlations on macroscopic behaviour (Section 3). The experience derived from the present work is summarised in Section 4 and future directions are indicated.

2 The Model

For simplicity we formulate the model on a one-dimensional lattice with lattice spacing λ. At time t there are $N_\alpha(x,t)$ particles of species α on the site with position coordinate x. One time step later, i.e. at time $t+\tau$, the value of N_α is given by

$$N_\alpha(x,t+\tau) = \mathcal{C} \circ \mathcal{T} N_\alpha(x,t), \qquad (1)$$

where $\mathcal{T}$ and $\mathcal{C}$ are the transport and chemical reaction operators, respectively. $\mathcal{T}$ lets each particle move to the right or left or remain stationary with probabilities p_α, q_α and $1 - p_\alpha - q_\alpha$, respectively. $\mathcal{C}$ adds (subtracts) the appropriate number of products (reactants) for each reaction that takes place. We write

$$\mathcal{C}N_\alpha(x,t) = N_\alpha(x,t) + \sum_{r=1}^{R} \left(\nu_{\alpha r}^{(f)} - \nu_{\alpha r}^{(i)}\right)\eta_{x,r}\,, \qquad (2)$$

where $\eta_{x,r}$ is a random Boolean variable and $\nu_{\alpha r}^{(i)}$ and $\nu_{\alpha r}^{(f)}$ are stoichiometric coefficients referring to initial and final species, respectively. The summation runs over all R one-way reactions obtained by separating, in the case of reversible reactions, the forward from the reverse process.

[1] A so-called 'exclusion principle' is useful for optimisation purposes in LGA simulations on special-purpose and vector computers.

Reaction r takes place at x when $\eta_{x,r} = 1$ and we choose this to happen with probability

$$P_r F_r\left(\left\{N_\beta(x,t): \beta \in S\right\}\right),$$

where the P_r is a real constant, F_r is a function of the occupation numbers and S is the set of all species. We shall see later that, in the continuum limit, P_r is related to the *rate constant* and F_r to the functional dependence of the reaction rate on the concentrations.

To obtain the particle density we average $N_\alpha(x,t)$ over several independent simulations:

$$\rho_\alpha(x,t) \equiv \left\langle N_\alpha(x,t)\right\rangle.$$

It follows that [5]:

$$\rho_\alpha(x,t+\tau) =$$

$$\left\langle \mathcal{J}N_\alpha(x,t)\right\rangle + \sum_{r=1}^{R}\left(\nu_{\alpha r}^{(f)} - \nu_{\alpha r}^{(i)}\right)P_r\left\langle F_r\left(\left\{\mathcal{J}N_\beta(x,t): \beta \in S\right\}\right)\right\rangle. \quad (3)$$

The first term on the RHS of (3) gives the result of the transport operation on the density:

$$\left\langle \mathcal{J}N_\alpha(x,t)\right\rangle =$$

$$p_\alpha\rho_\alpha(x-\lambda,t) + (1-p_\alpha-q_\alpha)\rho_\alpha(x,t) + q_\alpha\rho_\alpha(x+\lambda,t). \quad (4)$$

A chemical reaction rule that yields the standard law of mass action in the continuum limit is defined as follows:

$$F_r\left(\left\{N_\beta(x,t): \beta \in S\right\}\right) \equiv \prod_\beta \prod_{m=1}^{\nu_{\beta r}^{(i)}}\left(N_\beta(x,t) - m + 1\right). \quad (5)$$

Here the reaction probability is weighted by a product involving the number of particles of each species; as a result, the reaction becomes likelier the greater the number of particles present. For example, according to this rule, if r is the reaction $a + 2b \rightarrow c$, then it occurs with probability $P_r N_a N_b(N_b - 1)$. Of course, the reaction probability has to be ≤ 1 (probability conservation); since the occupation number is finite in a simulation, we can ensure probability conservation by choosing P_r to be sufficiently small (so that $P_r N_a N_b(N_b - 1) \leq 1$, $\forall x,t$, in the above example); since the physical reaction rate is related to P_r/τ, τ has to be chosen correspondingly small; the limit is set by the resulting increase in computation time due to the greater number of iterations.

The time evolution of the density is obtained by substituting (4) and (5) in (3). For the purpose of deriving the macroscopic equations, we have to discard some microscopic information by making two assumptions:

(i) There are no correlations between the occupation numbers of different species (*molecular chaos*).

(ii) N_α does not vary appreciably over several lattice spacings (*smoothness assumption*).

The first assumption allows us to break up $\langle F_r\rangle$ into a product of expectation values for different species, which can then be expressed in terms of the local density ρ_α, thanks to the second assumption. It should be clear that neither assumption is implied in our simulations and they only serve the purpose of deriving the macroscopic equations. This derivation is useful in order to delineate the differences between the CA and PDE approaches.

With the above assumptions we arrive at the discrete evolution equation [5]:

$$\rho_\alpha(x,t+\tau) = \rho_\alpha(x,t+\tau/2) +$$

$$\sum_{r=1}^{R}\left(\nu_{\alpha r}^{(f)} - \nu_{\alpha r}^{(i)}\right)P_r\prod_\beta \sum_{n=\nu_{\beta r}^{(i)}}^{\infty}\frac{\rho_\beta^n(x,t+\tau/2)}{n!}e^{-\rho_\beta(x,t+\tau/2)}, \quad (6)$$

where we have used the notation

$$\rho_\alpha(x,t+t/2) \equiv \left\langle \mathcal{J}N_\alpha(x,t)\right\rangle.$$

The continuum limit is defined by letting λ,τ and $p_\alpha - q_\alpha \rightarrow 0$, while keeping λ^2/τ and $(p_\alpha - q_\alpha)\lambda/\tau$ constant. We thus derive the standard reaction-transport PDE [5]:

$$\frac{\partial\rho_\alpha(x,t)}{\partial t} =$$

$$-V_\alpha\frac{\partial\rho_\alpha(x,t)}{\partial x} + D_\alpha\frac{\partial^2\rho_\alpha(x,t)}{\partial x^2} + \sum_{r=1}^{R}\left(\nu_{\alpha r}^{(f)} - \nu_{\alpha r}^{(i)}\right)k_r\prod_\beta \rho_\beta^{\nu_{\beta r}^{(i)}}(x,t), \quad (7)$$

where we define the *advection velocity* V_α, the *diffusion coefficient* D_α and the rate constant k_r as follows:

$$V_\alpha \equiv (p_\alpha - q_\alpha)\frac{\lambda}{\tau}, \quad D_\alpha \equiv (p_\alpha + q_\alpha)\frac{\lambda^2}{2\tau}, \quad k_r \equiv \frac{P_r}{\tau}. \quad (8)$$

We note that advection is the macroscopic manifestation of a bias in the displacement probabilities of the random walk ($p_\alpha \neq q_\alpha$). We treat species with different transport properties by choosing p_α and q_α individually for each species (e.g. $p_\alpha = q_\alpha = 0$ for immobile species). Inhomogeneities are accounted for by making p_α and q_α position dependent; then, if we define V_α and D_α using Eq. (8), we have to identify the advection velocity with

$$W_\alpha(x) \equiv V_\alpha(x) - dD_\alpha(x)/dx.$$

Similarly, the parameters in (8) can be made time-dependent.

The transition from particle densities to species *concentrations* is made by means of a scaling factor γ: $C_\alpha(x,t) \equiv \gamma\rho_\alpha(x,t)$. The C_α satisfy Eq. (7) with appropriately rescaled rate constants.

3 Applications

The CA presented above can describe a wide variety of systems subject to coupled mass transport and chemical reactions. Our aim in this section is to demonstrate the versatility of the model and to compare the results of the discrete simulations with those obtained from the corresponding PDE's. We shall see that microscopic information, such as density fluctuations and correlations between the molecules of different species, is neglected by the usual macroscopic approach at a certain cost. Thus, for a system of migrating particles subject to either of the reactions $a + b \leftrightarrow c$ and $a + b \rightarrow c$, the macroscopic long-time behaviour is dominated by microscopic effects that are naturally accounted for by the CA model but are left out by the standard reaction-diffusion PDE's. When such effects

are not significant, the CA and PDE approaches agree well with each other, as we shall see below. For details of the simulations, we refer the reader to Ref. [5].

3.1 $a + b \leftrightarrow c$

We begin with a homogeneous system of a-, b- and c-particles diffusing[2] and reacting according to the reversible reaction $a + b \leftrightarrow c$. In the macroscopic approach, all particle densities are spatially independent and the density of, say, the a-particles evolves in time according to the rate equation

$$\frac{d\rho_a(t)}{dt} = -k_1\rho_a(t)\rho_b(t) + k_2\rho_c(t), \qquad (9)$$

where k_1 and k_2 are the rate constants for $a + b \rightarrow c$ and $c \rightarrow a + b$, respectively. In the long-time limit, the system reaches equilibrium, with the densities satisfying the *law of mass action* [3]:

$$\frac{\rho_a(t)\rho_b(t)}{\rho_c(t)} = \frac{k_2}{k_1}, \qquad (10)$$

In our simulations we find that the asymptotic steady-state does not satisfy (10) if a single diffusion operation separates successive reaction operations. This is due to correlations at the microscopic level: the a- and b-particles originating from the decay of the same c-particle remain sufficiently close to each other for their probability of meeting again to be higher than in a well-mixed system. The conditions of molecular chaos are systematically approached by increasing the amount of diffusion (i.e. the number of diffusion operations per reaction operation) in our simulations (see also [6]).

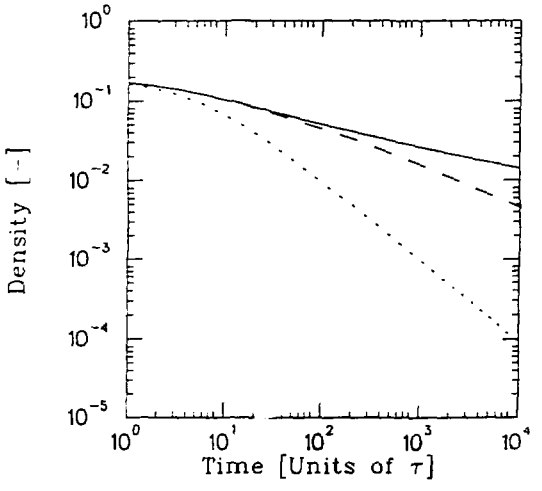

Fig. 1: a-density for 1-dimensional, homogeneous system reacting via $a + b \rightarrow c$, as function of time: (a) only diffusion (solid line), (b) diffusion+relative advection (dashed line), and (c) random redistribution (dotted line)

We consider next an inhomogeneous system of diffusing a-, b- and c-particles. Initially there are only a- and b-particles occupying the $x \leq 0$ and $x \geq 0$ halves of the lattice, respectively. The particles combine according to $a + b \rightarrow c$ to produce c-particles, which decay back to a and b according to $c \rightarrow a + b$. In our simulations, we find that *the width of the reaction zone* (i.e. of the region around $x = 0$ where the production rate of c is non-zero) increases asymptotically with time as $t^{1/2}$, as expected in the diffusion-limited regime prevailing for our choice of transport and reaction parameters. At the centre of the reaction zone, the densities of all species approach asymptotically stationary equilibrium values. If the time-scale for c-production is much shorter than that for c-decay, then the width increases as $t^{1/2}$ for $t >> k_2^{-1}$, as expected, but for $t << k_2^{-1}$ a different power-law appears to set in temporarily in one-dimensional simulations [7]; this is reminiscent of the asymptotic behaviour observed with the irreversible reaction, which we now proceed to describe in detail.

3.2 $a + b \rightarrow c$

When the reverse reaction is switched off, then a- and b-particles will deplete each other until there are only inert c-particles left over, unless there is a indefinite external supply of reactants. This is quite unlike the state of global or local chemical equilibrium attained with the reversible reaction. Consequently, questions of a qualitatively different nature have to be posed concerning the asymptotic behaviour of a system subject to the irreversible reaction $a + b \rightarrow c$.

We first consider a homogeneous system. Initially, there is a *uniform random distribution of equal numbers* (hence, also, densities) of a- and b-particles. Species c is inert and can be ignored. It is easily seen that the densities of the reactants remain equal at all times and the rate equation simplifies to:

$$d\rho_a(t)/dt = -k_1\rho_a^2(t).$$

The long-time prediction $\rho_a \propto t^{-1}$ follows trivially.

In Figure 1, we show the time evolution of ρ_a as obtained from one-dimensional simulations using three different kinds of transport operation: (a) both species diffuse with the same diffusion coefficient, (b) the same as before, but one of the species is subject to an additional drift (advection), and (c) particles are randomly redistributed between reaction steps. The asymptotic time dependence in the three cases listed above is $t^{-1/4}$, $t^{-1/2}$ and t^{-1}, respectively. It is clear that the prediction of the macroscopic rate equation corresponds to the well-mixed third case. In the other two cases, the long-time behaviour of the system is dominated by long-range density fluctuations which are present in the initial state and which diffusion or advection are not able to destroy by the time of interest; after the bulk of particles has been depleted, the remaining excess particles within such fluctuations form segregated regions of a- and b-concentrations; the reaction then proceeds only along the boundaries of these regions

[2] Unless otherwise stated, the diffusion coefficients of all species are assumed equal.

[3] The law of mass action in this simple form is valid only for infinite dilution. For higher concentrations, interactions (e.g. of electrostatic nature) complicate the situation and the relation holds properly for *activities* rather than concentrations.

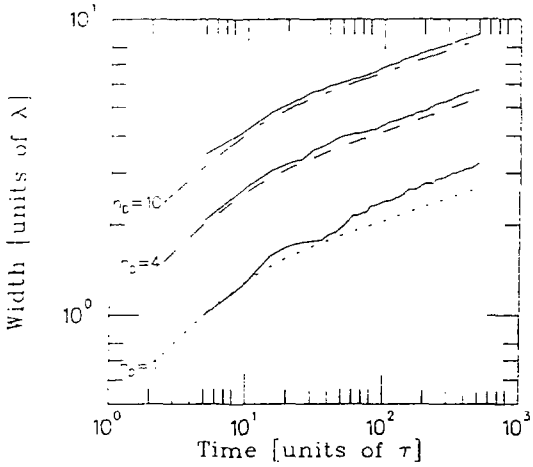

Fig. 2: Width of reaction zone for a two-dimensional, initially separated system reacting via $a + b \rightarrow c$, as a function of time: solution of reaction-diffusion equations (dotted, dashed and dot-dashed curves for 1, 4 and 10 transport operations per reaction operation, respectively) compared with result of simulations (solid curves).

and is slower than under the well-mixed conditions assumed by the rate equation [8]. Interestingly, the advection of one species relative to the other mixes the species faster than diffusion and speeds up the reaction compared to pure diffusion [9].

Moving away from the homogeneous system, we take up again the initially separated system of a- and b-particles considered in the previous subsection. This time there is no reverse reaction and the c-particles can be left out of the discussion. It can be shown from the reaction-diffusion equations that, under the diffusion-controlled conditions relevant here, the width of the reaction zone grows asymptotically with time as $t^{1/6}$ [10].

In Figure 2 we compare the width of the reaction zone obtained from two-dimensional simulations with the result of the reaction-diffusion equations. The differential equation results display asymptotically the expected $t^{1/6}$ behaviour. Figure 2 shows, however, that the width grows faster in the simulations. The discrepancy is substantial for n_D = 1, but it diminishes with enhanced diffusion (n_D = 4, 10). It appears that, if diffusion does not mix sufficiently the two species, a- and b-particles deplete each other less vigorously; as a result, each species penetrates deeper into the region of the other and the width of their overlap is enhanced. The observation that the relative enhancement diminishes with more diffusion suggests that the slowing down of the reaction can probably be traced back to density fluctuations, as in the homogeneous case above.

3.3 Autocatalytic reactions

Our aim in formulating a general CA model of reaction-transport phenomena has certainly been to go beyond the relatively simple reactions presented so far. As a first excursion into complex chemical systems, we model non-linear reaction schemes with *autocatalytic reactions*, which require the presence of a certain species in order to produce more of it. When systems of species subject to such reactions are driven far from chemical equilibrium, they can undergo phase transitions to new stable states, where the concentrations of the constituents exhibit different kinds of striking behaviour (oscillations, non-linear travelling waves, steady spatial patterns). Since life processes may depend crucially on the emergence of ordered states, autocatalytic reactions are of extreme importance in the chemical and biological sciences [11,12].

The formation of steady concentration patterns was predicted by A. Turing in 1952 [13] and experimentally observed recently [14]. With current computers, pattern formation can be more efficiently studied using standard FDE methods than with our microscopic simulation, but future massively parallel architectures may render the difference immaterial. To illustrate the versatility of our model we show the results of two simulations of autocatalytic systems.

The Brusselator [15] is a relatively simple model of a non-linear chemical system that displays a lot of the remarkable behaviour mentioned above; it is defined by the reaction scheme:

$$A \xrightarrow{k_1} X \quad , \quad B + X \xrightarrow{k_2} Y + D \quad ,$$

$$2X + Y \xrightarrow{k_3} 3X \quad , \quad X \xrightarrow{k_4} E \quad ,$$

where the concentrations of A and B are kept constant (e.g. by external supply of reactants), whereas D and E are continuously removed. Figure 3 shows the result of a simulation on a one-dimensional lattice of length L. The interior of the lattice is initially filled with uniform random distributions of X and Y particles, whereas the boundary concentrations of X and Y are fixed to the values of the homogeneous steady state. We choose the parameter values of our simulation so that an interplay between chemical

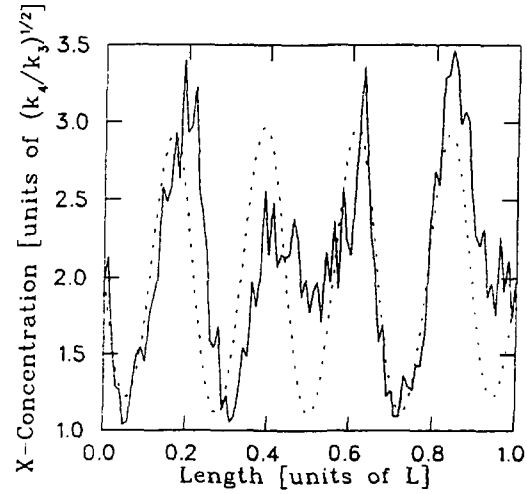

Fig. 3: Simulation of one-dimensional Brusselator (solid curve) compared with solution of reaction-diffusion equations (dotted curve).

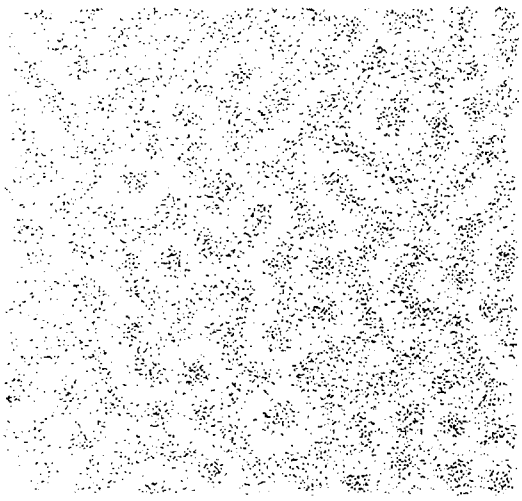

Fig. 4: Spatial concentration pattern obtained from two-dimensional simulation of the Schnackenberg model.

reactions and diffusion drives the system towards a structured stable state. A crucial condition for the emergence of Turing structures is a difference in the diffusion coefficients of X and Y. Linear stability analysis shows that, for our choice of parameters, the ratio D_Y/D_X has to exceed 3.05 for the homogeneous steady state to become unstable (in the simulation we take $D_Y/D_X = 3.75$). Unfortunately, the simulation is not yet efficient enough to allow, within reasonable computational time, the good statistics needed for a quantitative comparison with the result of the macroscopic equations.

The reaction scheme

$$A \leftrightarrow X \quad , \quad 2X + Y \rightarrow 3X \quad , \quad B \rightarrow Y$$

was introduced by Schnackenberg [16] as the simplest, but chemically plausible, two-species reaction mechanism that admits periodic solutions. The model has been studied extensively using the macroscopic reaction-diffusion equations and is known to exhibit a Turing instability over a wide domain of control parameters [12,17,18]. Figure 4 shows a two-dimensional concentration pattern obtained with our CA model. Due to time limitations the simulation was not run up to a time when the pattern becomes stationary.

3.4 Precipitation/dissolution

So far we have considered chemical reactions between species in solution. Aqueous solutions react, however, with the host rock in which they are transported. Thus, mineral dissolution and precipitation alter rock properties, e.g. the porosity, which have a direct impact on mass transport through the rock. Modelling the intricate coupling between hydrological and chemical processes constitutes a formidable task since the flow characteristics vary, in general, in space and time. We believe the CA modelling approach to offer significant advantages in this connection. We have seen in Section 2 that transport with spatially varying parameters can be treated with no additional effort. Moreover, the fact that we model at a microscopic level allows us to readjust the characteristics of the flow as they vary in the process of the simulation.

We consider here the simple reversible reaction

$$m(s) \leftrightarrow a + b,$$

where a and b are mobile solute species and m is a stationary mineral. The microscopic rule for this reaction can be stated as follows: If m-particles are present at a site, one of them dissociates into an a- and a b-particle with probability P_1. When an a- and a b-particle meet at a site, they react with probability P_2 if there is at least one m-particle at the same site or if the product of the a- and b-densities at that site is not less than the *solubility product* $\chi \equiv P_1/P_2$ (we do not allow supersaturation in our simple model). Thus, precipitation is always allowed to occur if there is already solid present, while otherwise the solution must be saturated before solid begins to form. The solute densities needed in order to decide if the solution is saturated are calculated by averaging over an ensemble of macroscopically identical systems.

As an example, we assume that at $t = 0$ a quantity of solid is uniformly distributed in the left half of a one-dimensional lattice. A fluid flows with constant velocity across the lattice, from left to right. We assume the amount of solid to be small, so that variations in its quantity do not appreciably affect the flow. The solid dissolves gradually and the dissolved particles are carried down the lattice by the fluid. We choose a lattice of 100 sites and a fluid velocity of $0.1\ \lambda/\tau$. Initially, there are on the average 10 solid particles per site in the left half of the lattice. The dissolved particles are transported with the above advection velocity and diffuse with diffusion constants $D_a = D_b = \lambda^2/2\tau$. The probabilities for the forward and backward reactions are 0.04 and 0.4, respectively. The density profile of the solid particles after 5000 iterations is shown in Figure 5. The solid, dotted and dashed curves correspond to ensemble averages over 250, 1000 and 4000 systems, respectively. It is interesting that there is precipitation of solid downstream from the original solid edge. This is entirely caused by statistical fluctuations in the solute densities. Due to the fluctuations the product of the densities occasionally exceeds the solubility product and precipitation occurs, an effect which is absent in the macroscopic approach. As expected, Figure 5 shows that averaging over a larger ensemble reduces the fluctuations and the resulting precipitation.

The above simulations were performed with a very small amount of solid compared to the volume of liquid present. In realistic porous media, the solid phase occupies most of the rock volume and our code requires several hours of CPU time in order to simulate the dissolution of a macroscopic amount of solid. The difference in scales between the numbers of mineral and solute particles, and the requirement of reasonable statistics for the latter, render a straightforward simulation of realistic mineral-solute systems very inefficient with present-day computers. Solving this problem, while maintaining the valuable microscopic aspects of the model, constitutes a major challenge to our approach. It is a challenge worth facing,

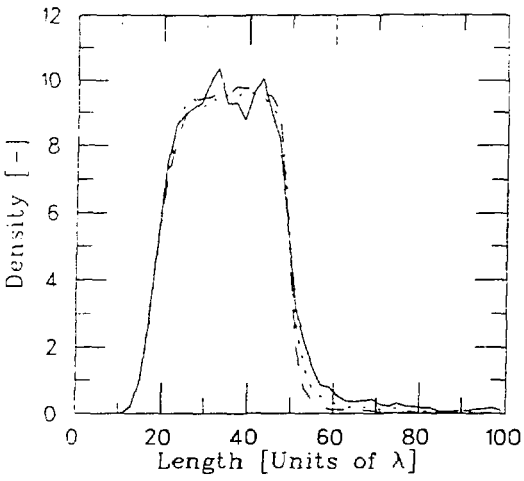

Fig. 5: Density profile of solid m-particles subject to the precipitation/dissolution process m(s) ↔ a + b.

since a microscopic approach is indispensable for the solution of transient problems in which the transport of solutes is continually modified by changes in the host rock caused by the same solutes.

4 Conclusions

The cellular automaton presented in this work models the transport of solutes coupled with chemical reactions. Solute particles perform a random walk on a regular lattice and react chemically among themselves or with stationary mineral particles according to a set of local probabilistic rules. At the macroscopic level, the random walk appears as a combination of diffusion and advection, the latter following from a directional bias of the random walk. Processes are modelled at an elementary level, albeit without the full detail of molecular dynamics. Due to this simplicity, arbitrary chemical reactions and boundary conditions can be modelled in a physically transparent way. The model is general enough to describe a wide variety of reaction-transport processes.

The evolution equations of the discrete model go over to the standard macroscopic reaction-transport equations in the continuum limit, if we assume molecular chaos and a smooth spatial dependence of the particle density. Since no such conditions are required by our model, our simulations account for microscopic effects that are left out by the macroscopic approach. We have seen this in particular in our simulations of homogeneous systems subject to the reactions a + b ↔ c and a + b → c, when the latter proceed on a much smaller time scale than transport. In the reversible case, we found out that correlations between the reacting particles can influence the macroscopic properties of the equilibrium state reached after sufficient time; this is an example of a breakdown of the molecular chaos assumption. In the case of the irreversible reaction, the long-time decline in the density of the reactants is dominated by residual fluctuations in the particle density that result in separate spatial domains of a- and b-particles; this gives rise to a slower annihilation mechanism between the two species. By neglecting microscopic fluctuations, the macroscopic approach misses completely the dominant reaction mechanism when it comes to the long-time limit and predicts the wrong power-law behaviour. In our simulations of the inhomogeneous a + b → c system we found that a- and b-particles penetrate deeper into the regions of each other as their reaction becomes weaker, due probably to a similar fluctuation effect. In microscopic simulations of autocatalytic reactions, density fluctuations result in transient pattern formation below the threshold predicted by linear stability analysis based on the macroscopic equations; this has been observed by other authors [19], but has also been confirmed by our simulations. Finally, density fluctuations can cause precipitation when the macroscopic densities lie just below the precipitation threshold. In conclusion, *we have demonstrated that our model is able to approximate systematically the results of the macroscopic equations, if microscopic properties have no significant consequences at the macroscopic level. If, on the contrary, effects of microscopic origin are not negligible, then our model, unlike the macroscopic approach, accounts naturally for them.*

The ability of our model to describe various reaction-transport systems in a physically transparent way is certainly reassuring. Equally important is the guaranteed stability of our algorithm, which is a valuable asset when modelling highly non-linear systems, as discussed in the Introduction. Another important criterion for the usefulness of the model as a practical tool is its computational efficiency. The possibility to improve the computation-time requirements of our simulations by utilising the freedom we enjoy in the selection of the microscopic reaction rule has been discussed in Ref. [5]. But the main computational gains that will make our model competitive with alternative methods of numerical solution of the macroscopic equations (to the extent that explicit microscopic effects are negligible or uninteresting) will most likely be attained by ingenious programming on an appropriate massively parallel computer. The performance of CA simulations is usually measured in site updates per second (u.p.s.). The efficiency of our simulations on a VAX-9000 in scalar mode ranges from 10^6 u.p.s. for pure transport of one species to 2×10^4 u.p.s. for the Brusselator. At the latter rate, the result displayed in Figure 3 required about 11 CPU hours, whereas the calculation illustrated in Figure 4 took about 30 CPU hours. On an early model of the massively parallel Connection Machine (CM-2a), we achieved a gain in the simulation of the a + b ↔ c reaction by a factor of about 2 compared to the performance on the VAX; an additional factor of 2 was attained thanks to the architectural improvements implemented in a later model (CM-200). The simulations can be made faster by utilising a low-level instruction set available on the Connection Machine, but we would rather maintain the present transparent form of our code. We point out that the above performance figures

hold for a general code that we have applied to the different systems by simply modifying the input. Substantial improvements in performance are possible by means of application-specific optimisation. Thus, in Ref. [2], the reaction a + b → c was simulated on a CRAY-YMP at a rate of 2×10^7 u.p.s., while an improvement by a factor of more than 4 is possible on CM-2a [20]. It is certainly desirable to combine the advantages of the present general (but slow) and other fast (but specific) models in order to obtain a useful and efficient algorithm.

The applications of the model presented here have to be selected so as to maximise its advantages: (i) the ability to model transport with inhomogeneous and time-dependent parameters, as well as systems which are not necessarily in chemical equilibrium, (ii) the guaranteed numerical stability, and (iii) the capacity to treat arbitrary boundary conditions. Such potential applications include the interaction of solutes with mineral surfaces and the coupling of porosity changes with the transport of reactive solutes. When alternative modelling approaches are available, relative computational efficiency will be the decisive criterion. We expect, however, the model to realise its full potential with natural systems whose complexity makes the microscopic point of view indispensable.

Acknowledgements

A substantial part of the work reported here was performed in collaboration with B. Blankleider. The author has benefited a great deal from discussions with B. Chopard, M. Droz, L. Frachebourg (Université de Genève), J. Hadermann (PSI), J.-P. Boon and D. Dab (Université Libre de Bruxelles). Partial financial support by NAGRA is gratefully acknowledged.

References

[1] *Theory and Applications of Cellular Automata* (S. Wolfram, ed.), Advanced Series on Complex Systems, World Scientific, Singapore (1986).

[2] CHOPARD B. AND DROZ M., "Microscopic study of the properties of the reaction front in an A+B→C reaction-diffusion process."*Europhys. Lett.* **15** (1991) 459.

[3] LAWNICZAK A. ET AL., "Reactive lattice gas automata." *Physica D* **47** (1991) 132.

[4] *Lattice Gas Methods for Partial Differential Equations*, (G. Doolen, et al., eds), Santa Fe Institute Studies in the Science of Complexity, Addison-Wesley, (1990).

[5] KARAPIPERIS T. AND BLANKLEIDER B., "Cellular automaton model of mass transport with chemical reactions." PSI Report 93-06 (1993).

[6] DAB D. AND BOON J.-P., "Cellular automata approach to reaction-diffusion systems." in *Cellular Automata and Modeling of Complex Physical Systems* (P. Manneville, et al., eds), Springer, Berlin (1989), 257-273.

[7] CHOPARD B., ET AL., "Properties of the reaction front in a reversible $A+B\underset{k_2}{\overset{k_1}{\leftrightarrow}}C$ reaction-diffusion process." *Phys. Rev. E* **47** (1993) 40.

[8] TOUSSAINT D. AND WILCZEK F., "Particle-antiparticle annihilation in diffusive motion." *J. Chem. Phys.* **78** (1983) 2642.

[9] KANG K. AND REDNER S., "Fluctuation-dominated kinetics in diffusion-controlled reactions." *Phys. Rev. A* **32** (1985) 435.

[10] GÁLFI L. AND RÁCZ Z., "Properties of the reaction front in an A+B→C type reaction-diffusion process." *Phys. Rev. A* **38** (1988) 3151.

[11] NICOLIS G. AND PRIGOGINE I., *Self-Organization in Nonequilibrium Systems*, Wiley, New York (1977).

[12] MURRAY J. D., *Mathematical Biology*, Springer, New York (1989).

[13] TURING A. M., "The chemical basis of morphogenesis." *Phil. Trans. Roy. Soc. London B* **237** (1952) 37.

[14] CASTETS V., ET AL., "Experimental evidence of a sustained standing Turing-type nonequilibrium chemical pattern." *Phys. Rev. Lett.* **64** (1990) 2953.

[15] PRIGOGINE I. AND LEFEVER R., "Symmetry breaking instabilites in dissipative systems." *J. Chem. Phys.* **48** (1968) 1695.

[16] SCHNACKENBERG J., "Simple chemical reaction systems with limit cycle behaviour." J. *Theor. Biol.* **81** (1979) 389.

[17] DUFIET V. AND BOISSONADE J., "Conventional and unconventional Turing patterns." *J. Chem. Phys.* **96** (1992) 664.

[18] DROZ M. AND FRACHEBOURG L., "Turing structures in cellular automata models of reaction-diffusion systems." *Helv. Phys. Acta* **66** (1993) 97.

[19] KAPRAL R., LAWNICZAK A. AND MASIAR P., "Oscillations and waves in a reactive lattice-gas aotomaton." *Phys. Rev. Lett.* **66** (1991) 2539.

[20] CHOPARD B., personal communication.