Modelling Simul. Mater. Sci. Eng. **15** (2007) 711–738 doi:10.1088/0965-0393/15/7/002

A three-dimensional microstructural model of reactions and transport in aqueous mineral systems

Jeffrey W Bullard

Materials and Construction Research Division, National Institute of Standards and Technology, Gaithersburg, MD, USA

E-mail: jeffrey.bullard@nist.gov

Received 6 December 2006, in final form 8 August 2007
Published 27 September 2007
Online at stacks.iop.org/MSMSE/15/711

Abstract

A stochastic three-dimensional microstructure model is introduced for simulating spatial and temporal variations in aqueous mineral systems. Dissolution, nucleation, precipitation and solute transport are governed by local probabilistic rules applied on a regular computational lattice. The model is shown to accurately simulate ion diffusion in a dilute electrolyte. The reaction algorithms faithfully reproduce kinetics expected from standard rate equations, and reversible reactions are shown to converge to the correct equilibrium state determined by detailed balance of forward and reverse reaction rates, or the law of mass action. Accounting for the exponential temperature dependence of the reaction rate constants is shown to provide accurate predictions of the influences of temperature on both the kinetics and equilibrium of reactions. A simulation of the hydration of a generic metal oxide in water demonstrates the important relationships between microstructure development and the mechanisms of nucleation, growth and solute diffusion.

1. Introduction

This paper describes a computer model for predicting the rate and extent of chemical and structural changes in three-dimensional (3D) aqueous mineral systems, which are relevant to a number of natural and commercial processes. When minerals come into contact with water, a number of chemical and physical changes can occur that are capable of changing the properties of the system. In general terms, each mineral approaches equilibrium with the adjacent liquid by dissolution and precipitation reactions at its surface, and solute species generated at the surface may then be transported through the solution by diffusion or advection. At any location where the solution becomes supersaturated with respect to formation of a solid, that solid can potentially nucleate and grow until local equilibrium is achieved. These chemically driven processes can all affect the composition of the solution, as well as both the amount and the spatial distribution of solids.

Dissolution, transport and precipitation of materials play a fundamental role in areas of materials processing such as synthesis of colloidal particles, sol–gel thin film deposition and hydration of cementitious materials and gypsum slurries. A wide range of other commercial and natural processes are also governed by these same rate phenomena, including corrosion of metals, diagenesis of porous sedimentary rocks, weathering of rocks and the formation of mineral deposits in water pipes and tanks.

Thermodynamic computer models such as PHREEQC [1] are available for predicting equilibrium states of complex geochemical systems, and therefore can serve as a guide for the long-time changes in phase composition in such systems. But the instantaneous rate of a chemical reaction is generally described by a nonlinear partial differential equation and, when multiple reactions can occur, the kinetic equations are inextricably coupled to each other. Furthermore, dissolution and precipitation reactions occur at surfaces that change their shape and position as the reactions proceed. That is, microstructure development in these systems is basically a 3D moving boundary problem. Therefore, linking the rates of coupled reactions and transport to the corresponding changes in realistic microstructures requires a 3D numerical model that can track not only the solution composition and solid phase volume fractions but also the spatial distribution of the phases as a function of time. As stated earlier, this paper presents a computer model which has been developed for just these tasks.

In developing the model, the main objectives have been (1) to base it on accepted principles of reaction kinetics, thermodynamics and solute transport; (2) to provide a direct link between the coupled reactions and the development of 3D microstructure and (3) to make the model sufficiently general that it can be applied as a research tool to study a range of aqueous mineral systems, including cement paste, gypsum/water systems, porous rocks and mineral deposits in supersaturated water environments.

The purpose of this paper is to introduce the main features of the model and to illustrate its utility through a series of examples. The model is described in section 2. Many of the basic principles and underlying assumptions of the model are based on earlier work by Karapiperis and Blankleider [2] for ideal diffusion and homogeneous reactions, and that work is reviewed briefly. But this model also extends that earlier work in significant ways by including the kinetics of nucleation of solids in supersaturated solutions, the ability to model heterogeneous reactions and the effects of temperature on both kinetics and equilibrium states.

Section 3 provides examples for verifying the model. The examples are divided into three sets: (1) diffusion of solute species, (2) homogeneous reactions between solute species in a solvent and (3) heterogeneous dissolution/precipitation reactions at mineral–water interfaces.

2. Model description

2.1. Assumptions and principles

The numerical methods used by the model are based closely on a cellular automaton (CA) model described first in a paper by Karapiperis and Blankleider [2] and in a later paper by Karapiperis [3]. Hereafter, this model is referred to as the KB model. In these papers, the authors not only provide clear details of the algorithms used, but they also demonstrate, both mathematically and with numerous examples, the convergence of their numerical approach to the standard rate equations for diffusion and chemical reactions. Because those papers provided such clear detail and evidence for the validity of the numerical approach, the description here is brief.

As with any CA algorithm, a stochastic approach is used to model transport and reactions. The material system is discretized on a regular cubic lattice with lattice spacing λ. At each

lattice site, a number of automata, or independent agents, may be located. This paper refers to these agents as *cells*. Each cell represents a particular material (solid, liquid, solute) with specific chemical composition and physical properties. The number of cells of a given type present at a lattice site, called the *occupation number* at that site, determines the local concentration (for solutes) or volume fraction (for solids) of the material that the cell represents.

A simulation proceeds according to local rules for diffusion and reactions that are applied synchronously over all lattice sites at regular time intervals τ. The time interval is supposed to encompass a two-step sequence: a transport step (with no reactions) and a reaction step (with no transport). Each rule i is executed at each lattice site by assigning a probability, $p(i) \in [0, 1]$ that it occurs. A random number $q \in [0, 1]$ is selected from a standard uniform distribution and compared with the probability. The event i is executed at the site if and only if $q \leqslant p(i)$. Therefore, the assignment of probabilities for the various possible events provide the link between the model implementation and the actual chemical and physical properties of the system being modeled.

In the following description, lattice sites are denoted by their vector position $\boldsymbol{x}$ relative to an arbitrarily chosen origin. Chemical species generally are indexed by a lower case Greek letter (α, β, etc). Chemical reactions, and sometimes lattice sites, are indexed by a lower case Latin letter (i, j, etc).

To model transport by diffusion of solute species, each cell is allowed to execute a random walk, taking at most one step from its current location to any one of its six nearest-neighbor sites in each time interval τ. To model diffusion in the continuum limit, where $\lambda \to 0$ and $\tau \to 0$, the probability $p(\boldsymbol{x}, \boldsymbol{x} + \delta\boldsymbol{x}, \alpha)$ that a cell of solute species α is displaced from a lattice site at vector position $\boldsymbol{x}$ to a nearest-neighbor site located at $\boldsymbol{x} + \delta\boldsymbol{x}$ is given by

$$p(\boldsymbol{x}, \boldsymbol{x} + \delta\boldsymbol{x}, \alpha) = \frac{\tau D_\alpha(\boldsymbol{x}, \boldsymbol{x} + \delta\boldsymbol{x})}{\lambda^2}, \tag{1}$$

where $D\alpha(\boldsymbol{x}, \boldsymbol{x} + \delta\boldsymbol{x})$ is the effective diffusion coefficient governing the local rate of diffusion of the solute from vector position $\boldsymbol{x}$ to $\boldsymbol{x} + \delta\boldsymbol{x}$. It can be noted here that transport by advection can be modeled by a similar approach [2], although this paper will not consider advection.

Tabulated values are available for the diffusion coefficient at infinite dilution, D_α°, of many solute species in water at 298 K. The actual values D_α can differ from D_α° and depend on factors such as temperature, the concentrations of each solute species [4] and the tortuosity [5] of the diffusion path between sites.

Equation (1) is general and captures the basic scaling of diffusion phenomena. The details of the physics and chemistry of a given diffusion process are contained in the calculation of the local diffusion coefficients. As long as care is taken to calculate a meaningful local value for D_α, the equation may be applied to unbiased diffusion of neutral solutes, to diffusion of ions and to diffusion in either dilute or more concentrated solutions. An example will be given in a later section to illustrate how (1) can be applied to nonideal electrolytes.

In the reaction step, each chemical reaction is represented as

$$\sum_\alpha \nu_\alpha^{(r)} s_\alpha \xrightarrow{k} \sum_\alpha \nu_\alpha^{(p)} s_\alpha, \tag{2}$$

where the sum is taken over all the available chemical species, and $\nu_\alpha^{(r)}$ and $\nu_\alpha^{(p)}$ are the molar stoichiometric coefficients for species s_α as a reactant and product, respectively. Reactions that can proceed at appreciable rates in both the forward and reverse directions,

$$\sum_\alpha \nu_\alpha^{(r)} s_\alpha \underset{k_-}{\overset{k_+}{\rightleftharpoons}} \sum_\alpha \nu_\alpha^{(p)} s_\alpha \tag{3}$$

are modeled as two one-way reactions of the form given in the previous equation.

Designating $N_\alpha(\boldsymbol{x}, t)$ as the number of cells, or occupation number, of species s_α at location $\boldsymbol{x}$ and time t, these occupation numbers are updated according to a sum over all R of the possible reactions:

$$N_\alpha(\boldsymbol{x}, t+\tau) = N_\alpha(\boldsymbol{x}, t) + \sum_{j=1}^{R} \left(\nu_{\alpha,j}^{(p)} - \nu_{\alpha,j}^{(r)}\right) \eta_{\boldsymbol{x},j}, \tag{4}$$

where $\eta_{\boldsymbol{x},j} = 1$ if the reaction j occurs at vector position $\boldsymbol{x}$, and $\eta_{\boldsymbol{x},j} = 0$ if not. The probability $p(\boldsymbol{x}, j)$ that $\eta_{\boldsymbol{x},j} = 1$ depends on the reaction rate constant k_j and on the occupation numbers at $\boldsymbol{x}$. For homogeneous reactions, that is, reactions that occur in the bulk of a phase, the derivation of $p(\boldsymbol{x}, j)$ is given in [2]; the result in modified form is

$$p(\boldsymbol{x}, j) = k_j \xi^{\left(\sum_\alpha \nu_{\alpha,j}^{(r)}\right)-1} \tau \prod_\alpha \max\left[0, \prod_{m=1}^{\nu_{\alpha,j}^{(r)}} \aleph_\alpha(\boldsymbol{x}, t, m)\right], \tag{5}$$

where ξ is a constant parameter of the model that relates the local occupation number of species s_α, N_α, to the corresponding molar concentration of s_α. $\aleph_\alpha$ is a function of the occupation number for species α given by

$$\aleph_\alpha(\boldsymbol{x}, t, m) = \begin{cases} (N_\alpha(\boldsymbol{x}, t) - m + 1) & \text{if } s_\alpha \text{ is solute,} \\ 0 & \text{if } s_\alpha \text{ is condensed phase and } N_\alpha < \nu_{\alpha,j}^{(r)}, \\ 1 & \text{if } s_\alpha \text{ is condensed phase and } N_\alpha \geqslant \nu_{\alpha,j}^{(r)}, \end{cases} \tag{6}$$

where m is the index used in equation (5).

In the continuum limit, the reaction rule described by (4) and (5) simulates reaction kinetics governed by the standard rate equation [2]. Consider the generic homogeneous reaction given by (2). Equations (4) and (5) can be shown to converge [2], as $\lambda \to 0$ and $\tau \to 0$, to the standard equation for the rate of production of species β given by

$$\frac{\partial c_\beta}{\partial t} = \left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) k \prod_\alpha C_\alpha^{\nu_\alpha^{(r)}}, \tag{7}$$

where the product is taken over all the species that are defined for the system. This rate equation holds when each reactant is dilute enough that its activity is well approximated by its concentration. Although not considered in the KB model, the concentrations in (7) must be replaced by activities if the reactants are more concentrated [6], so that the rate equation for species β becomes

$$\begin{aligned} \frac{\partial c_\beta}{\partial t} &= \left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) k \prod_\alpha a_\alpha^{\nu_\alpha^{(r)}}, \\ &= \left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) k \prod_\alpha (y_\alpha c_\alpha)^{\nu_\alpha^{(r)}}, \end{aligned} \tag{8}$$

where y_α is the molar activity coefficient for species s_α defined by

$$y_\alpha \equiv \frac{a_\alpha}{c_\alpha}.$$

The activity coefficients can be associated with the rate constant, so that a form of (7) is recovered

$$\begin{aligned} \frac{\partial c_\beta}{\partial t} &= \left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) \left(k \prod_\alpha y_\alpha^{\nu_\alpha^{(r)}}\right) \prod_\alpha c_\alpha^{\nu_\alpha^{(r)}}, \\ &\equiv \left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) \bar{k} \prod_\alpha c_\alpha^{\nu_\alpha^{(r)}}. \end{aligned} \tag{9}$$

Here, and throughout the rest of the paper, the symbol $\overline{k}$ is used to represent the rate constant, k, multiplied by the product of the molar activity coefficients of the reactants raised to their respective molar stoichiometric coefficient.

Heterogeneous reactions, that is, reactions that occur at an interface between two condensed phases, require a different expression for the reaction probability. Unlike homogeneous reactions, the absolute rate of a heterogeneous reaction is proportional to the area S of the reacting surface. This important class of reactions was not considered explicitly in the KB model. Instead, such reactions were essentially modeled as homogeneous reactions by coarse-graining the system to the point that the reactions could be considered to occur within a uniformly saturated rock with prescribed and constant porosity, such that the solid–liquid surface area remained fixed and proportional to the volume of the system [3]. While this approach is helpful at large length scales, it cannot be used to model the detailed growth of precipitates or the changing surface/volume ratio and morphology of the porosity at the microstructural level.

To model the kinetics of heterogeneous reactions, one may extend the ideas of the KB model for homogeneous reactions. Without loss of generality, assume that species s_β is one of the condensed phases involved in a heterogeneous reaction. Then the generic rate law for the production of s_β in a volume V is

$$\frac{\partial c_\beta}{\partial t} = \frac{S}{V}\left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) k \prod_\alpha a_\alpha^{\nu_\alpha^{(r)}},$$

where c_β, which has units of concentration, is the number of moles of condensed phase s_β per unit volume, S and V are the reacting surface area and the volume associated with the lattice site, respectively and k is the rate constant. Within a cubic volume element of dimension λ intersecting the reacting surface, S and V scale as λ^2 and λ^3, respectively, and so the rate equation for production of s_β can be rewritten as

$$\frac{\partial c_\beta}{\partial t} = \frac{1}{\lambda}\left(\nu_\beta^{(p)} - \nu_\beta^{(r)}\right) \overline{k} \prod_\alpha c_\alpha^{\nu_\alpha^{(r)}}.$$

Therefore, the rate of a heterogeneous reaction takes the same form as that of a homogeneous reaction, (9), with the exception of the spatial scaling factor $1/\lambda$. This scaling factor therefore can be incorporated in (5) to determine the probability of the jth heterogeneous reaction occurring at the lattice site at vector position $\boldsymbol{x}$:

$$p(\boldsymbol{x}, j) = k_j \xi^{\left(\sum_\alpha \nu_{\alpha,j}^{(r)}\right)-1} \frac{\tau}{\lambda} \prod_\alpha \max\left[0, \prod_{m=1}^{\nu_{\alpha,j}^{(r)}} \aleph_\alpha\,(x, t, m)\right]. \tag{10}$$

In the rate equation (9), the exponent corresponding to the concentration of a given reactant is equal to the corresponding stoichiometric coefficient in the chemical reaction. This equality is guaranteed only if the reaction is *elementary*, meaning that only one elementary step is involved. For more complex, multi-step reactions, the exponents may differ significantly from the corresponding stoichiometric coefficients, and the values of the exponents usually can be determined only by experimental observation [6]. This should not be construed as an inherent limitation of the formulation given in (4)–(10). If the values of the exponents in the rate equation are known to differ from the stoichiometric coefficients, (5) or (10) can be modified such that the maximum value of the index m is equal to that exponent instead of $\nu_{\alpha,j}^{(r)}$. However, for the purposes of this paper, the equality between the exponents and their corresponding stoichiometric coefficients is assumed to be a good approximation.

With reaction models for both homogeneous and heterogeneous reactions in hand, the remaining issue is the nucleation of a new solid phase when the solution becomes supersaturated with respect to it. In the KB model, the formation and growth of solid phases are allowed anywhere that solid already exists or where the solution is saturated with respect to it [3]. However, such a model cannot discriminate between homogeneous and heterogeneous nucleation phenomena, and also cannot account for the significant degrees of supersaturation that often are observed before a thermodynamically favored solid phase actually nucleates and grows in solution.

The calculation of the rate of nucleation per unit volume differs fundamentally from that for the rates of other types of reactions, and the model must use a different procedure to calculate nucleation rates. Following a review by Kaschiev and van Rosmalen [7], consider the reaction given by

$$\sum_{\alpha} \nu_{\alpha}^{(r)} s_{\alpha,(aq)} \rightarrow M_{(s)} \tag{11}$$

in which a number of aqueous species react to form a single solid product M. If solid M is not already available to act as a template for further growth at a proposed reaction site, then the formation of a stable, or supercritical, nucleus is required if the reaction is to occur. The number of supercritical nuclei formed per unit volume per unit time is [7]

$$I = A\mathbb{K}\mathrm{e}^{-W^*/k_{\mathrm{B}}T}, \tag{12}$$

where A is a prefactor that represents the frequency of attempts to form a supercritical nucleus, W^* is the work required to form a single supercritical nucleus, k_{B} is the Boltzmann constant, T is the absolute temperature and $\mathbb{K}$ is the saturation index defined by

$$\mathbb{K} \equiv \prod_{\alpha} \left(\frac{a_{\alpha}}{a_{\alpha,\mathrm{eq}}} \right)^{\nu_{\alpha}^{(r)}} = K_{\mathrm{eq}} \prod_{\alpha} a_{\alpha}^{\nu_{\alpha}^{(r)}} \tag{13}$$

with a_{α} and $a_{\alpha,eq}$ being the actual and equilibrium activity, respectively, of s_{α}. k_{eq} is the equilibrium constant for the reaction in (11). Neither A nor W^* are constants but depend on the molecular volume, surface free energy, temperature and $\mathbb{K}$ as described in [7].

The nucleation rate I can be transformed to the rate of change in molar density n_{M}, in moles per unit volume, of nucleated material M as follows:

$$\frac{\mathrm{d}n_{\mathrm{M}}}{\mathrm{d}t} = \frac{IV^*}{\Omega}, \tag{14}$$

where V^* is the volume of a supercritical nucleus. Substituting (13) and (14) into (12) yields

$$\frac{\mathrm{d}n_{\mathrm{M}}}{\mathrm{d}t} = A'\mathrm{e}^{-W^*/k_{\mathrm{B}}T} \prod_{\alpha} c_{\alpha}^{\nu_{\alpha}^{(r)}}. \tag{15}$$

The leading coefficient A' is given by

$$A' = \frac{AV^*}{\Omega} K_{\mathrm{eq}} \prod_{\alpha} y_{\alpha}^{\nu_{\alpha}^{(r)}}, \tag{16}$$

where again y_{α} is the molar activity coefficient for species s_{α}.

In terms of the model formulation, (15) is recast as a finite difference equation, with concentrations rewritten as cell occupation numbers at lattice site located at $\boldsymbol{x}$, to produce

$$\Delta N_{\mathrm{M}}(\boldsymbol{x}, t + \tau) = \eta_{\boldsymbol{x}}, \tag{17}$$

where N_{M} is the cell occupation number of solid M, and $\eta_{\boldsymbol{x}} = 1$ if nucleation occurs but is zero otherwise. Equations (14)–(16) can be used to determine the probability $p(\boldsymbol{x})$ that $\eta_{\boldsymbol{x}} = 1$:

$$p(\boldsymbol{x}) = \overline{k}_{\mathrm{nuc}} \xi^{\left(\sum_\alpha \nu_\alpha^{(r)}\right)-1} \tau \prod_\alpha \max \left[0, \prod_{m=1}^{\nu_\alpha^{(r)}} \left(N_\alpha(\boldsymbol{x}, t) - m + 1 \right) \right]. \tag{18}$$

Equation (18) has been intentionally made to resemble the corresponding (5) for other reactions by collecting several terms into an effective nucleation rate constant, $\overline{k}_{\mathrm{nuc}}$:

$$\begin{aligned} \overline{k}_{\mathrm{nuc}} &\equiv \overline{k}_0 \, \mathrm{e}^{-W^*/k_{\mathrm{B}}T}, \\ &= \overline{k}_0 \, \exp\left[\frac{-W^*}{T^3 \, (\ln \mathbb{K})^2} \right], \end{aligned} \tag{19}$$

where $W^* = (T \ln \mathbb{K})^2 W^*/k_{\mathrm{B}}$ can be assumed constant, to a good approximation, for a given nucleation process [7]. By comparing (18) and (19) to the general equation for homogeneous reaction probabilities, (5), it is evident that the same general reaction algorithm can be used for nucleation phenomena as for other types of reactions.

2.2. *Implementation*

As stated already, the model proceeds by repeating a cycle that consists of two steps which are assumed to be separable: a transport step followed by a reaction step. Each two-step cycle corresponds to a time increment of τ. The maximum time increment $\tau_{\max}$ is fixed by requiring during each cycle that the probability of every possible event does not exceed unity. Thus for the transport step, (1) can be used to estimate

$$\tau_{t,\max} \leqslant \frac{\lambda^2}{2d \max(D)} \tag{20}$$

with d being the dimensionality of the computational lattice and $\max(D)$ being the maximum expected local diffusion coefficient for any species at any position. The diffusion coefficient for any species should be constant only in ideal solutions in the absence of variable potential fields. In more realistic situations, the value of the diffusion coefficient, even for a particular species, can vary substantially from point to point within the system. In such situations, one approach is to make a conservative estimate of the maximum possible value of the diffusion coefficient over all species to determine a suitable value for $\tau_{t,\max}$. Alternatively, adaptive time increments can be implemented by making a sweep through the lattice to calculate the actual maximum value of the diffusion coefficient prior to calculating $\tau_{t,\max}$.

The same principle applies for the reaction step: one can estimate a maximum time increment $\tau_{r,\max}$ by appealing to (5), (10) and (18) and requiring that the probability of any event cannot exceed one. Then, for the entire two-step cycle, the maximum time increment is given by

$$\tau_{\max} = \min\left(\tau_{t,\max}, \tau_{r,\max}\right).$$

As long as the time increment of any given cycle does not exceed $\tau_{\max}$, then the simulation is guaranteed to be numerically stable, regardless of the possibly large differences in rate constants among the various processes that may take place. In fact, the maximum time increment can even be increased somewhat, with acceptably small error, in the following way. Suppose that one allows $\tau_{r,\max}$ to be n times greater than the value estimated from the expressions for reaction/nucleation probabilities. Then one can expect that the calculated probability of some reaction j at some lattice site located at $\boldsymbol{x}$ will exceed one. When this occurs, the given reaction

is executed m times, where m is the greatest integer less than $p(\boldsymbol{x}, j)$. A random number $q \in [0, 1]$ is selected from a standard uniform distribution, and one additional occurrence of the reaction is executed at the lattice site if and only if $q \leqslant p(\boldsymbol{x}, j) - m$. The error in the predicted reaction kinetics caused by this approximation increases as the multiplication factor n increases. However, when $\xi < 1$ it has been found that choosing $n \leqslant 1/\xi$ produces negligible errors in the kinetics.

In some scenarios, it may be desirable to simulate a system in which the liquid phase is being mixed in some way. Although slow or incomplete mixing is difficult to simulate, the limiting case of instantaneous, complete mixing is more straightforward. In the latter case, one can simply distribute each solute species homogeneously throughout the liquid phase during each transport step without regard to displacement. When this algorithm is adopted, the value of $\tau_{t,\max}$ is essentially infinite, so only the reaction step limits the maximum time step in any given cycle.

A detailed description of the organization of data and algorithms are beyond the scope of this paper, but a few words may be helpful to describe the general design of the model. Written in C++ and using an object-oriented approach, the overall design is composed of a hierarchy of classes. The basic building block of all materials is considered to be the 'ion' class. A generic material is represented as a 'material' base class that stores data about the material, such as its composition, molar volume, density, heat capacity and enthalpy of formation. This class also includes basic functions to change its properties in response to changes in environmental conditions, to calculate its local activity coefficient and other functions. All materials in the model are actually specialized classes derived from the generic class, which enables each class of material (solute, solid, liquid, etc) to have its own specialized functions that are appropriate to only that kind of material. For example, the activity of a solute can be calculated with a function that multiplies its local concentration by an appropriate activity coefficient, but the activity of a pure liquid or solid is generally taken to be one. All materials are stored in a material database class.

Different kinds of reactions (dissolution, precipitation or homogeneous reactions) are represented as classes that store a list of the reactants, products, stoichiometric coefficients, enthalpy of reaction, activation energy and reference reaction rate constant at 298 K. Again, each different reaction class has specialized functions that are characteristic of its behavior. As an example, only precipitation reactions include functions to check for the necessity of nucleation and therefore decide whether (10) or (18) should be used for calculating the probability.

Structurally, the basic building block of the computational system is the lattice site, represented as a class. This class stores data about the cell occupation numbers of each material, the internal porosity, temperature and the heat generated in any cycle. It also includes functions for calculating the local concentration of solute species, the volume fraction of solids and liquids, the transport factor for diffusion to a neighbor and for determining which reactions are possible from the list of materials it contains. Finally, the entire computational lattice is represented as a class that stores a list of all its sites as well as parameters such as the site spacing λ, the elapsed time and current time increment and the total heat released or absorbed by the system.

3. Results

In this section, eight examples are given to demonstrate various applications of the model and to provide validation of the algorithms.

3.1. Nonideal diffusion of NaCl in water

Detailed analysis of the diffusion model is the subject of a separate paper [8], in which the model is applied to both ideal solutions and to dilute nonideal solutions. In that paper, transition state theory is used to derive the form of the diffusion rate constants that should be used in these applications. The stability of the model and its numerical convergence to the generalized diffusion equation also are demonstrated in that paper. Therefore, this paper will briefly examine only nonsteady-state diffusion in a simple dilute electrolyte.

Ignoring off-diagonal effects such as thermal diffusion, the diffusion of an ion α in an electrolyte is driven by gradients in its electrochemical potential, $\tilde{\mu}_\alpha$, and the linear phenomenological rate law for the material flux $\boldsymbol{j}_\alpha$ is [9]

$$\boldsymbol{j}_\alpha = -L_\alpha \nabla \tilde{\mu}_\alpha, \tag{21}$$

where L_α is the Onsager coefficient for diffusion of α, which is related to the more familiar diffusion coefficient D_α by

$$D_\alpha = \left(\frac{\partial \tilde{\mu}_\alpha}{\partial c_\alpha} \right) L_\alpha.$$

In [8], the author has shown a derivation of the following expression for the rate constant of diffusion from a site $\boldsymbol{x}$ to an adjacent site $\boldsymbol{x}+\delta\boldsymbol{x}$, $k_{D,\alpha}(\boldsymbol{x}, \boldsymbol{x}+\delta\boldsymbol{x})$, of a species s_α in a nonideal solution when a potential field may bias the diffusion trajectory:

$$k_{D,\alpha}(\boldsymbol{x}, \boldsymbol{x}+\delta\boldsymbol{x}) = \frac{D_\alpha^\circ}{\lambda^2} \exp\left(-\frac{c_\alpha(\boldsymbol{x}+\delta\boldsymbol{x}) - c_\alpha(\boldsymbol{x})}{2} \frac{\mathrm{d} \ln y_\alpha}{\mathrm{d}c_\alpha} - \frac{\psi_\alpha(\boldsymbol{x}+\delta\boldsymbol{x}) - \psi_\alpha(\boldsymbol{x})}{2RT} \right), \tag{22}$$

where D_α° is the diffusion coefficient at infinite dilution, values of which are tabulated for many ions [10], y_α is the molar activity coefficient and ψ_α is the potential per mole of component s_α, given by

$$\psi_\alpha \equiv -q_\alpha \psi \tag{23}$$

with q_α being the electric charge (C mol^{-1}) of species s_α and ψ being the electrostatic potential, in J C^{-1}.

The molar activity coefficients for both Na^+ and Cl^- were estimated using a modified Davies equation, which has been shown to be reasonably accurate up to ionic strengths of about 0.1 mol L^{-1} [11]. The diffusion coefficients at infinite dilution at 298 K are $D^\circ_{Na^+} = 1334\,\mu m^2\,s^{-1}$ and $D^\circ_{Cl^-} = 2033\,\mu m^2\,s^{-1}$, respectively [10]. Therefore, (22) can be used to calculate the specific rate constants at 298 K:

$$k_{\boldsymbol{x},\boldsymbol{x}+\delta\boldsymbol{x},Na^+} = \frac{1334}{\lambda^2} \sqrt{\frac{y_{Na^+}(\boldsymbol{x})}{y_{Na^+}(\boldsymbol{x}+\delta\boldsymbol{x})}} \exp\left[\frac{-\mathcal{F}\left(\phi(\boldsymbol{x}+\delta\boldsymbol{x}) - \phi(\boldsymbol{x})\right)}{2RT} \right], \tag{24}$$

$$k_{\boldsymbol{x},\boldsymbol{x}+\delta\boldsymbol{x},Cl^-} = \frac{2033}{\lambda^2} \sqrt{\frac{y_{Cl^-}(\boldsymbol{x})}{y_{Cl^-}(\boldsymbol{x}+\delta\boldsymbol{x})}} \exp\left[\frac{-\mathcal{F}\left(\phi(\boldsymbol{x}+\delta\boldsymbol{x}) - \phi(\boldsymbol{x})\right)}{2RT} \right], \tag{25}$$

where $\mathcal{F}$ is Faraday's constant.

The electrostatic potential at each lattice site is estimated using a screened sum of charges over a neighborhood of the $M_{2\lambda}$ lattice sites separated from the given site by a distance of 2λ or less:

$$\phi(\boldsymbol{x}) = \frac{\kappa}{4\pi\epsilon_w} \sum_j^{M_{2\lambda}} \frac{\xi\lambda^3}{r(\boldsymbol{x}, j)} \sum_k N_k q_k, \tag{26}$$

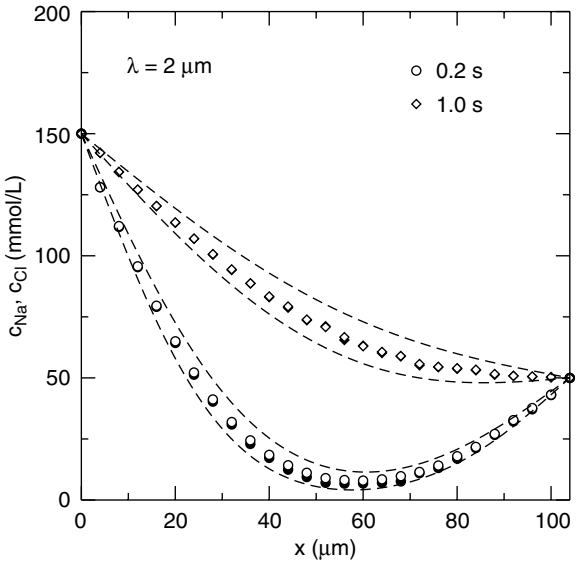

Figure 1. Diffusion of NaCl. The open symbols indicate the concentration profiles of Cl^- and the closed symbols indicate the concentration profiles of Na^+. The upper and lower dashed curves surrounding each profile correspond to the ideal concentration profiles of Cl^- and Na^+, respectively, based on their self-diffusion coefficients at infinite dilution.

where κ is a dimensionless factor used to simulate the Coulombic screening of charge by water, set at 6×10^{-8} in this simulation, ϵ_w is the dielectric permittivity of water, and $r(\boldsymbol{x}, j)$ is the distance between lattice site at $\boldsymbol{x}$ and its jth neighbor, in units of λ. The conversion factor ξ again is set to $2 \times 10^{-5}\,\mathrm{mol\,L^{-1}}$.

To test this approach, a small orthorhombic diffusion cell was modeled, having dimensions $X = 100\,\mu\mathrm{m}$ and $Y = Z = 10\,\mu\mathrm{m}$ and a lattice spacing $\lambda = 2\,\mu\mathrm{m}$. At time $t = 0$, the cell is filled with pure water, with the following boundary conditions: $c_{\mathrm{Na}} = c_{\mathrm{Cl}} = 150\,\mathrm{mmol\,L^{-1}}$ at $x = 0\,\mu\mathrm{m}$, and $c_{\mathrm{Na}} = c_{\mathrm{Cl}} = 50\,\mathrm{mmol\,L^{-1}}$ at $x = 100\,\mu\mathrm{m}$ for all $t > 0$. Figure 1 shows predicted concentration profiles at times of 0.2 s and 1.0 s. Open symbols correspond to Cl^- and the closed symbols correspond to Na^+ concentrations. The upper and lower dashed curves surrounding each profile are the hypothetical concentration profiles of Cl^- and Na^+, respectively, if each behaved as an ideal neutral component.

Figure 1 shows that the diffusion rate of Cl^- is decreased significantly, and that of Na^+ is increased modestly, from the infinite-dilution limit. This is due to the influence of the local electrostatic potential field that arises from transient charge separation. In most regions of the plot, the open symbol completely obscures the underlying filled symbol, and it is only near the minima in the profiles that differences are noticeable. Microscopically, the field produced should be exactly that required to cause both ions to diffuse with the same velocity [4], which maintains local charge neutrality. The figure suggests that the rough estimation of the electrostatic potential in (26) is sufficient to approximately enforce charge neutrality across the diffusion cell.

Another way of testing the diffusion simulation is to compare the nonsteady-state concentration profiles to those predicted by the 1D diffusion equation:

$$\frac{\partial c_{\mathrm{Na}}}{\partial t} = \frac{\partial}{\partial x}\left(D_{\mathrm{NaCl}} \frac{c_{\mathrm{Na}}}{\partial x}\right), \tag{27}$$

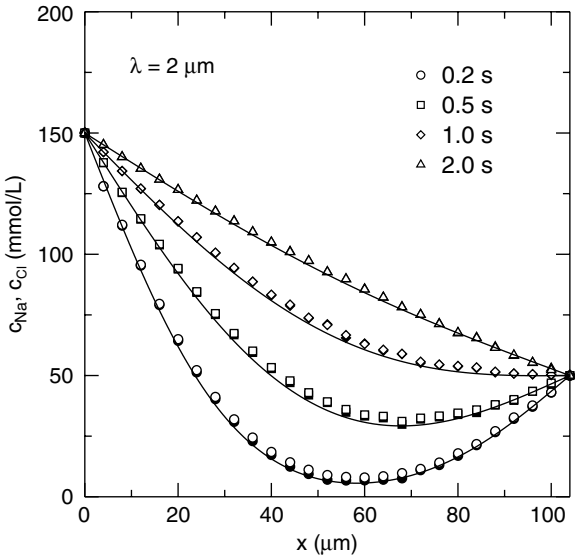

Figure 2. Comparison of NaCl concentration profiles predicted by the stochastic model (symbols) with those predicted using experimentally observed effective diffusion coefficients for NaCl in [4], (solid curves). The open symbols indicate the concentration profiles of Cl^- and the closed symbols indicate the concentration profiles of Na^+.

where the *effective* diffusion coefficient of NaCl in water, D_{NaCl}, is dependent on solute concentration and temperature. Values for D_{NaCl} have been measured experimentally with high accuracy over a range of concentrations [4]. For the concentration range relevant to this example, a reasonable quadratic fit to these measurements is

$$D_{NaCl} \approx 1556.4 - 1.178\,42\,c_{NaCl} + 0.003\,81\,c_{NaCl}^2 \quad (\mu m^2\,s^{-1}). \tag{28}$$

Numerical solution of (27) using a finite difference approach [12] thus provides 'experimental' concentration profiles with which the stochastic model predictions can be compared.

Figure 2 shows such a comparison of concentration profiles at 0.2 s, 0.5 s, 1.0 s and 2.0 s. Once again, the open and closed symbols correspond to predicted concentrations of Cl^- and Na^+, respectively. The solid curves correspond to the numerical solution of (27) using (28) for D_{NaCl}.

Based on figures 1 and 2, the model appears to capture some basic electrolyte diffusion phenomena quite well. A more in-depth discussion of its strengths and limitations regarding other electrolyte phenomena can be found in [8].

3.2. Homogeneous reactions

The validity of the reaction algorithms are now verified by comparing the model predictions of the rates of irreversible reactions against the standard rate law. The reactions considered in this section are homogeneous in the sense that they occur uniformly in space. Heterogeneous reactions that occur at the interface between condensed phases are considered in the following section.

For simplicity, it is assumed that (1) the reactions are elementary, in the sense discussed in section 2.1, (2) the activity of any species in solution is equal to its molar concentration, i.e. $\overline{k} \equiv k$ and (3) the activity of any pure condensed phase is equal to one. The impact of solution nonideality on reaction kinetics is taken up in the following section.

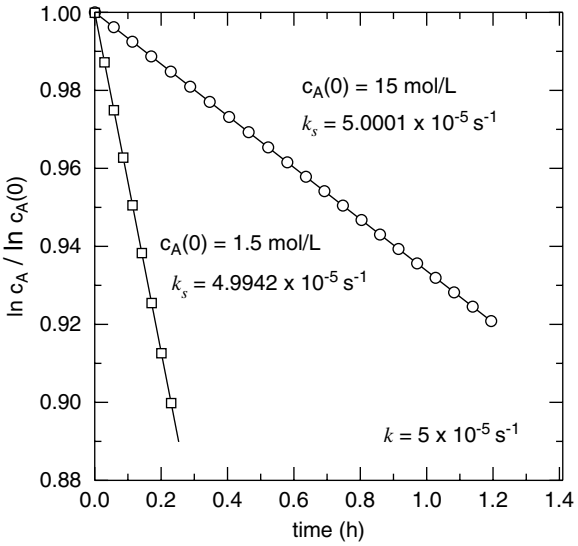

Figure 3. Time dependence of reactant concentration for the irreversible homogeneous unimolecular reaction A $\longrightarrow$ B with rate constant $k = 5 \times 10^{-5}\,\mathrm{s}^{-1}$. Open symbols are the simulation results and the solid lines are fit to the simulation data by linear regression. For clarity, only 1/100 or 1/200 of the simulation points are shown for the upper and lower lines, respectively.

For each of the next four examples, the 'reactor' is simulated as a cube with an edge length of 30 μm. The concentration parameter ξ is $0.025\,\mathrm{mol\,m}^{-3}$. Unless stated otherwise, the lattice spacing $\lambda = 3\,\mu$m. The solution is assumed to be well mixed, which means that after each computational cycle, every component in solution is distributed uniformly throughout the liquid.

Example 1. Homogeneous unimolecular reaction

The simplest possible reaction is a homogeneous one by which an individual molecule transforms into another:

$$A \xrightarrow{k} M.$$

The standard rate law for this reaction, subject to the stated assumptions in this section, is

$$\frac{\mathrm{d}c_\mathrm{A}}{\mathrm{d}t} = -kc_\mathrm{A}.$$

Solving this differential equation yields

$$\frac{\ln c_\mathrm{A}(t)}{\ln c_\mathrm{A}(0)} = -\frac{k}{\ln c_\mathrm{A}(0)}t + 1. \tag{29}$$

Therefore, a plot of the left side of this equation versus t should be a straight line with slope equal to $-k/\ln c_\mathrm{A}(0)$. Figure 3 shows such a plot generated by the model when the rate constant is $5 \times 10^{-5}\,\mathrm{s}^{-1}$ and the initial concentration is $15\,\mathrm{mol\,L}^{-1}$ or $1.5\,\mathrm{mol\,L}^{-1}$. The symbols show the simulation results, and the solid lines are the best-fit lines determined by linear regression. As shown in the figure, the simulation predicts a linear relation at both concentrations. The apparent rate constant determined from the simulation, k_s, differs from the 'true' rate constant by 0.0002% or 0.5% when the initial concentration is $15\,\mathrm{mol\,L}^{-1}$ or $1.5\,\mathrm{mol\,L}^{-1}$, respectively.

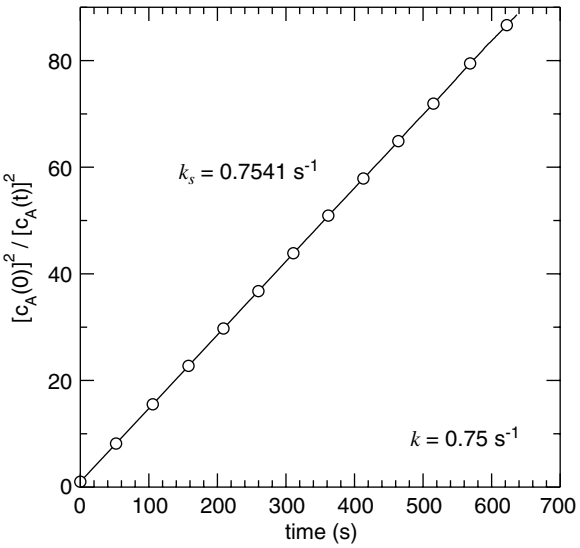

Figure 4. Time dependence of reactant concentration for the irreversible homogeneous trimolecular reaction A + 2B $\longrightarrow$ M with rate constant $k = 0.75\,\mathrm{s}^{-1}$. Open symbols are the simulation results and the solid line is a fit to the simulation data by linear regression. The best-fit slope to the simulation results is $0.7541\,\mathrm{s}^{-1}$, which differs from the target rate constant by 0.5%. For clarity, only 1/50 of the simulation points are shown.

Although the error is small in each case, the relatively poorer accuracy at the lower initial concentration is due to the lower average occupation numbers at each lattice site. Recall that the relationship in the model between concentration c and occupation number N is $c = \xi N$. Because $\xi = 0.025\,\mathrm{mol\,m}^{-3}$ for both simulations, the average initial occupation number at a lattice site is 6×10^5 when the initial concentration is $15\,\mathrm{mol\,L}^{-1}$, and is 6×10^4 when the initial concentration is $1.5\,\mathrm{mol\,L}^{-1}$. The same accuracy could be achieved in both simulations by adjusting ξ to obtain equal values of the average initial occupation number.

Example 2. Homogeneous trimolecular reaction

As a step up in complexity, consider the trimolecular homogeneous reaction,

$$A + 2\mathrm{B} \xrightarrow{k} \mathrm{M}.$$

The standard rate law is

$$\frac{\mathrm{d}c_\mathrm{A}}{\mathrm{d}t} = -kc_\mathrm{A}c_\mathrm{B}^2.$$

Solving for c_A yields

$$\left(\frac{c_\mathrm{A}(0)}{c_\mathrm{A}(t)}\right)^2 = 8c_\mathrm{A}^2(0)kt + 1. \tag{30}$$

In this case, a plot of the left side of this equation versus t should be a straight line with slope equal to $8c_\mathrm{A}^2(0)k$. Figure 4 shows such a plot generated by the model when the rate constant is $5 \times 10^{-5}\,\mathrm{s}^{-1}$. Again, the symbols show the simulation results, and the solid lines are fits to the data using linear regression. As in the previous example, the simulation correctly predicts a linear relation. The apparent rate constant determined from the simulation, k_s, differs from the 'true' rate constant by 0.5%.

3.3. Heterogeneous reactions

Heterogeneous reactions occur between phases that are separated by an interface. In this model, if one lattice site contains a condensed phase and a nearest-neighbor site does not, then a surface implicitly exists between those sites, and a heterogeneous reaction may be possible at both of them. However, this situation presents a potential problem. Reactions in the model are executed at each lattice site using whatever collection of reactants is available at that site. If a surface is discontinuous—meaning that one site is completely occupied by one condensed phase and the adjacent site is completely occupied by the other—then any reaction between them will not happen because the reactants in the two sites are effectively separated. To detect and remedy this problem with discontinuous surfaces, at least two different approaches are possible. First, each site could poll its nearest neighbors and borrow whatever reactants may be available to make a reaction happen. This approach has the advantage that it can be applied uniformly at all sites for all heterogeneous reactions. However, the act of borrowing species from nearest neighbors makes it difficult to calculate the innermost product in (5) or (10) in a consistent way without doing violence to the rate equation they are supposed to approximate.

A second approach to handling discontinuous surfaces is to smooth out the discontinuity before attempting a reaction there. Smoothing can be accomplished by exchanging a fraction ε of the contents of each lattice site comprising the discontinuous surface. The fraction ε will hereafter be called the 'smoothing parameter'. Once smoothing occurs, then any surface reaction can proceed at both of the sites that undergo the exchange of material. This approach has the advantages of being simple to execute and of only being required in relatively rare instances where a discontinuity occurs. And although it is a numerical way to circumvent discontinuities, the idea of smoothing may have some plausibility from a physical perspective as well. In Gibbs' original theory of the thermodynamics of surfaces, he considered a surface to be a continuous transition in properties from one phase to another, with a two-dimensional dividing surface being defined as a mathematical convenience [13]. The Cahn–Hilliard theory of inhomogeneous fluids [14], as well as the phase-field models that are based on that theory [15], explicitly model a surface as a gradient of an order parameter from one condensed phase to another. In the Cahn–Hilliard theory, this gradient contributes to the free energy of the system, that is, excess interface free energies are included [14, 15]. However, in the present model, the smoothing at discontinuous interfaces is a numerical convenience and does not have an interpretation analogous to a gradient of an order parameter. Because of this, the smoothing used here is assumed to have no effect on the thermodynamics of the system.

In the following example, the smoothing idea will be used to remove discontinuous surfaces, and the impact on kinetics of the smoothing parameter ε will be evaluated.

Example 3. Dissolution at a plane surface

The unimolecular dissolution reaction,

$$\mathrm{B}_{(\mathrm{s})} \xrightarrow{k} \mathrm{B}_{(\mathrm{aq})}$$

has an associated standard rate law given by

$$\frac{V}{S}\frac{\mathrm{d}c_{\mathrm{B}_{(\mathrm{s})}}}{\mathrm{d}t} = -\bar{k}c_{\mathrm{B}_{(\mathrm{s})}}, \tag{31}$$

where now V is the system volume, S is the total area of the solid–liquid interface and c_{B} is the number of moles of solid phase B per unit volume. If the reacting interface is large and

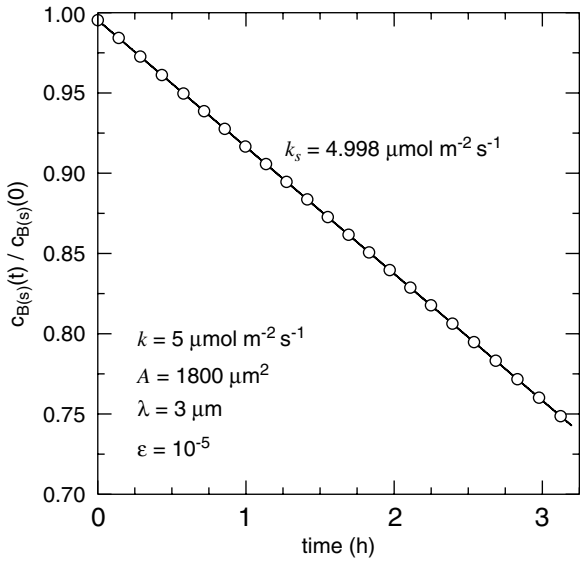

Figure 5. Time dependence of reactant concentration for the irreversible dissolution reaction $B_{(s)} \longrightarrow B_{(aq)}$ with rate constant 5×10^{-6} mol m^{-2} s^{-1}. Open symbols are the simulation results and the solid line is a fit to the simulation data by linear regression. The best-fit slope to the simulation results implies an apparent rate constant of 4.998×10^{-6} mol m^{-2} s^{-1}, a -0.04% difference from the true value. For clarity, only 1/500 of the simulation points are shown.

planar, then S remains constant with time, and the analytic solution to the rate law is

$$\frac{c_B(t)}{c_B(0)} = -\frac{\overline{k}}{c_B(0)} St + 1, \tag{32}$$

where it also has been assumed that the activity of a pure solid is equal to one.

In this example, the system is again modeled as a cube with an edge length of 30 μm. The lattice resolution $\lambda = 3\,\mu$m and the concentration parameter $\xi = 0.025$ mol m^{-3}. Initially, a plane sheet of the solid $B_{(s)}$, 15 μm thick, is inserted into the system, with its surface normal oriented parallel to one of the lattice axes. The remainder of the system is filled with initially pure water. Periodic boundary conditions are used in each direction, so the computational domain itself can be thought of as a unit cell of a system composed of an infinite number of uniformly spaced, parallel thin plates. With this setup, the volume of the unit cell is 2.7×10^{-8} cm^3 and the initial volume of solid in the unit cell is 1.35×10^{-8} cm^3 or, equivalently, 4.1×10^{-10} mol if one assumes that the molar volume is that of $Ca(OH)_2$, 3.31×10^{-5} m^3 mol^{-1}. The total surface area per unit cell is $S = 1800\,\mu$m^2, which remains constant throughout the simulation.

Figure 5 shows a plot of the left side of (32) versus time. The rate constant is chosen to be 5 μmol m^{-2} s^{-1}, which is approximately the value for dissolution of $Ca(OH)_2$ at 298 K. Again, the symbols are the simulation results, and the solid line is a linear regression fit to the simulation. Evidently, the simulation correctly captures the linear relation predicted by (32). Data regression yields a slope of -8.997×10^{-15} mol s^{-1}, implying an apparent rate constant of 4.998 μmol m^{-1} s^{-1}. This value differs from the target rate constant by -0.04%.

In figure 5, the smoothing parameter ε was chosen to be 10^{-5}. This means that, prior to executing the reaction, any given 'solid' lattice site on one side of the surface donates 10^{-5} of its solid volume to the 'water' site adjacent to it, and simultaneously accepts 10^{-5} of the water

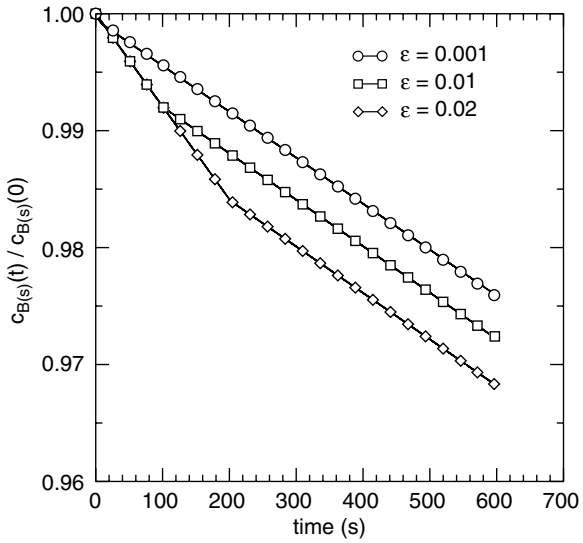

Figure 6. Influence of the value of the smoothing parameter, ε, on the kinetics of the irreversible dissolution reaction in figure 5. For clarity, only 1/1500 of the simulation points are shown on each curve.

volume from that site. To investigate how the smoothing parameter might influence the kinetics of the reaction, the same simulation was executed three more times using $\varepsilon = 0.001$, $\varepsilon = 0.01$ and $\varepsilon = 0.02$. Figure 6 shows the resulting plots for all three simulations. For each value of ε, the figure indicates two linear regions with an abrupt transition between them, instead of a single line. The slope of lines at later times all are equal to the slope in figure 5, but the slope at early times is twice that value in each case. The reason for the change in slope is that the smoothing process, by donating a small amount of solid to adjacent water sites, doubles the number of sites that are participating in the dissolution reaction at early times. That is, the effective surface area is doubled by the smoothing process at early times. However, when the donated solid has dissolved completely, the sites that initially accepted that solid no longer participate in dissolution and the effective surface area is halved. Thus, it is not the apparent rate constant, but the effective surface area that is impacted by smoothing.

Figure 6 also shows that this surface area artifact can be made arbitrarily short-lived by choosing a small enough value of the smoothing parameter ε. In fact, the time at which the change in slope occurs is linearly proportional to ε. In figure 5, where $\varepsilon = 10^{-5}$, the increased surface area at early times is present for such a brief period that it cannot be discerned in the plot.

Example 4. Dissolution of a sphere

When a spherical particle of a solid material B dissolves, both its volume V_B and its area S_B decrease continuously as material dissolves:

$$dV_B = \Omega_B dn_B,$$

$$dS_B = \left(\frac{32\pi\Omega_B^2}{3n_B}\right) dn_B,$$

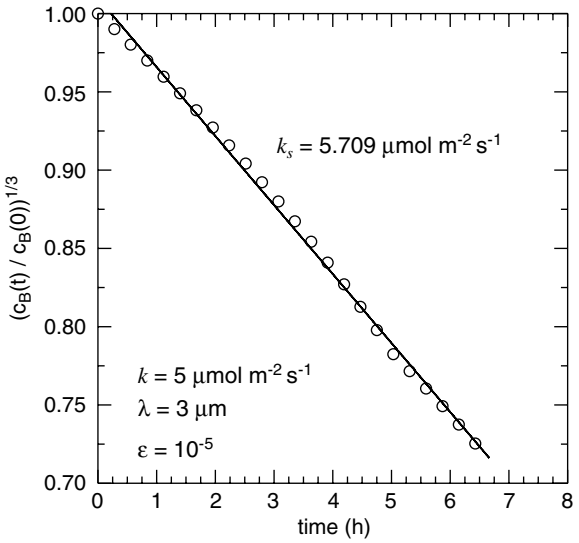

Figure 7. Time dependence of reactant concentration for the irreversible dissolution of a 21 μm diameter sphere by the reaction $B_{(s)} \longrightarrow B_{(aq)}$ with rate constant 5 μmol m^{-2} s^{-1}. Open symbols are the simulation results and the solid line is a fit to the simulation data by linear regression. The best-fit slope to the simulation results implies an apparent rate constant of 5.709 μmol m^{-2} s^{-1}, a 14.2% difference from the true value. For clarity, only 1/1000 of the simulation points are shown.

where Ω_B and n_B are the molecular volume and number of molecules of the solid, respectively. With these relationships, the rate law in (31) integrates to

$$\left(\frac{c_B(t)}{c_B(0)}\right)^{1/3} = -\left(\frac{4\pi\Omega_B^2}{3c_B(0)}\right)^{1/3} \bar{k}t + 1. \tag{33}$$

A plot of the left side of this equation against time, therefore, should be a straight line with negative slope proportional to the rate constant. Figure 7 shows such a plot for a simulation using a solid sphere with an initial diameter of 21 μm and molar volume $\Omega_m = 3.31 \times 10^{-5}$ m^3 mol^{-1}, which is the value for $Ca(OH)_2$ at 298 K. The initial sphere was approximated as a digitized collection of cubic volume elements, in which a given volume element with centroid at (x, y, z) belongs to the sphere if

$$\lambda\sqrt{(x - x_c)^2 + (y - y_c)^2 + (z - z_c)^2} \leqslant 10.5,$$

where (x_c, y_c, z_c) is the center of the sphere. The other parameters of the simulation are the same as in the previous example.

Figure 7 shows that the data are indeed nearly linear. The solid line, which was determined by linear regression, has a coefficient of determination $R^2 = 0.9988$. The effective rate constant calculated from the slope of the line is 5.709 μmol m^{-2} s^{-1}, which is 14.2% greater than the value used for the reaction. The discrepancy can be attributed primarily to two artifacts of the microstructure representation, which partially offset each other. First, a digitized sphere always has a greater surface-to-volume ratio S_B/V_B than that of a continuum sphere of the same diameter, the difference being due to the terraced surface of the digitized sphere [16]. In this example, a continuum sphere with a diameter of 21 μm has $S_B/V_B = 0.2857\ \mu$m^{-1}, while its digitized counterpart has $S_B/V_B = 0.5551\ \mu$m^{-1} after the smoothing procedure has been performed, as described at the beginning of this section. This is a difference of 94%, so the instantaneous dissolution rate at $t = 0$ is 94% greater than expected for a real sphere.

However, as the small quantity of material originally transferred during the smoothing process dissolves, the value of S_B/V_B in the microstructure decreases gradually towards a value of approximately $0.1738\,\mu m^{-1}$, which is 39% *less* than the value for the analogous continuum sphere. The net effect over the course of the simulation is an overestimation of the rate by about 14.2%. Much of this error could be eliminated by starting off with a more accurate representation of the spherical particle. By allowing sites at the surface to be partially occupied with solid, one could generate an initial sphere that is much smoother than the binary digitized sphere used in this example. Nevertheless, this example at least illustrates that the linear integrated rate law for a sphere in (33) can be captured by the simulation.

3.4. Reversible reactions

Consider the following generic dissociation reaction:

$$A \underset{k_-}{\overset{k^+}{\rightleftharpoons}} \nu_B B + \nu_C C.$$

The standard rate equations for the forward and reverse reactions are

$$\frac{dc_A}{dt} = -k_+ a_A \qquad \text{(forward)}, \tag{34}$$

$$\frac{dc_A}{dt} = k_- a_B^{\nu_B} a_C^{\nu_C} \qquad \text{(reverse)}. \tag{35}$$

At equilibrium, the net rate of change of c_A is zero by definition. Setting the sum of the forward and reverse rates to zero gives

$$\frac{k_+}{k_-} = \frac{a_B^{\nu_B} a_C^{\nu_C}}{a_A} \equiv K_{eq}. \tag{36}$$

Therefore, to simulate any elementary reaction that can approach equilibrium, one needs only to account for both the forward and the reverse reactions. The absolute rates of these reactions will continually change as equilibrium is approached, until a dynamic equilibrium is established, such that $k_+/k_- = K_{eq}$.

Example 5. Reversible dissociation of calcium hydroxide in water

For this example, we assume a one-step reaction for the dissolution/growth of calcium hydroxide in water:

$$Ca(OH)_2 \underset{k_-}{\overset{k_+}{\rightleftharpoons}} Ca^{2+} + 2OH^-. \tag{37}$$

The system size and geometry used for this example is the same as in example 3. The equilibrium constant for dissolution of $Ca(OH)_2$ is 6.59×10^{-6} [1]. Based on measurements of the rate of growth of $Ca(OH)_2$ in highly supersaturated solutions [17], the rate constant for precipitation at 298 K is estimated as $k_- = 0.1\,mol\,m^{-2}\,s^{-1}$. The ratio of the rate constants should be equal to the equilibrium constant, which implies that $k_+ = 0.66\,\mu mol\,m^{-2}\,s^{-1}$. In this example, one of two initial conditions for the ion concentrations was used. The first initial condition was

$$c_{OH} = c_{Ca} = 0,$$

which simulates dissolution into initially pure water. The other initial condition was

$$c_{Ca} = 0.5 c_{OH} = 23.6\,mmol\,L^{-1}.$$

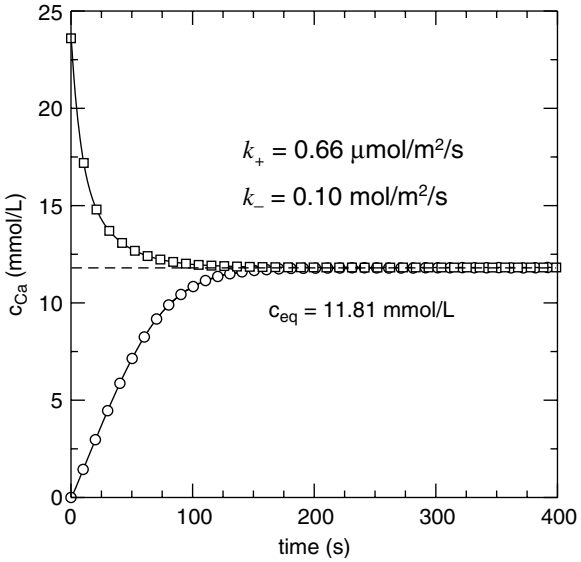

Figure 8. Time dependence of calcium concentration for the reversible reaction between $Ca(OH)_2$ and water. The circles are the results for dissociation of $Ca(OH)_2$ into initially pure water, and the squares are the results for growth of $Ca(OH)_2$ from an initially supersaturated solution. For clarity, only 1/3 of the simulation points are shown, although the solid curve follows the other points, not shown.

This value is twice that predicted by solving the equilibrium constant for concentration, so it represents a supersaturated solution that should promote growth of $Ca(OH)_2$ as the solution approaches equilibrium. Once again, molar concentration is used to approximate activity, primarily because this makes it easy to predict equilibrium concentrations in terms of the equilibrium constant.

Figure 8 shows the results of both simulations. The concentration of Ca^{2+} is plotted versus time. Evidently, in both simulations the concentration of Ca^{2+} approaches a stationary value of 11.81 mmol L^{-1}. Based on K_{eq}, the equilibrium concentration of Ca^2+ is 11.81 mmol L^{-1}, so the stationary state adopted in the simulation is quite accurate relative to the equilibrium state predicted by K_{eq}.

From a kinetic standpoint, the initial rate of dissolution, given by the initial slope of the lower curve in figure 8, is less than the initial rate of growth indicated by the initial slope of the upper curve. This difference in initial rates is expected because the rate constant for the dissolution reaction, k_+, is less than the rate constant for growth, k_-. In addition, the net rates of both dissolution and growth decrease substantially as equilibrium is approached because, although the rate constants are fixed, the thermodynamic driving force for reaction decreases exponentially as equilibrium is approached from either direction.

The concentration in figure 8 converges to the value implied by K_{eq} if it is assumed that the molar concentration of each solute species is equal to its activity. However, the experimentally observed concentration of Ca^{2+} in a saturated solution of $Ca(OH)_2$ at 298 K is 22 mmol L^{-1}, almost twice as high as that predicted in figure 8. The agreement with experiment can be improved significantly by estimating the activities of the solute species instead of using their concentrations. When the molar activity coefficients of Ca^{2+} and OH^- are estimated using the modified Davies equation [11], then the simulations predict an equilibrium Ca^{2+} concentration of 19.1 mmol L^{-1}. Although this value still underestimates the experimentally

measured value by 13%, it is a considerable improvement over the result when the activity coefficients are neglected. Further improvements probably could be made by using even more refined estimates of the activity coefficients, and also by taking account of ion complexation reactions in solution.

3.5. Temperature effects

The effect of temperature on absolute reaction rates and equilibrium states can be modeled by incorporating the temperature dependence of the rate constants. Consider again the generic reversible reaction in (3). Let the reactants on the left side have a collective Gibbs free energy of $G^{(r)}$, and let the products of the reaction on the right side have a collective Gibbs free energy of $G^{(p)}$. The free energies are all relative to some common reference state. According to transition state theory [18], both the forward and the reverse reactions involve passage through a state of higher free energy, G^*, called the activated complex, so that the temperature dependence of the forward rate constant is [18]:

$$\begin{aligned} k_+(T) &= \frac{k_B T}{h} \exp\left[\frac{-\left(G^* - G^{(r)}\right)}{k_B T}\right], \\ &= \frac{k_B T}{h} e^{\left(s^* - s^{(r)}\right)/k_B} \exp\left[\frac{-\left(H^* - H^{(r)}\right)}{k_B T}\right], \\ &= \frac{k_B T}{h} e^{\Delta S_+^*/k_B} \exp\left[\frac{-\Delta H_+^*}{k_B T}\right]. \end{aligned} \tag{38}$$

In equation (38), h is Planck's constant, and ΔS_+^* and ΔH_+^* are the entropy of activation and the enthalpy of activation, respectively, for the forward reaction at temperature T. An analogous expression holds for the rate constant in the reverse direction:

$$k_-(T) = \frac{k_B T}{h} e^{\Delta S_-^*/k_B} \exp\left[\frac{-\Delta H_-^*}{k_B T}\right]. \tag{39}$$

In practice, one generally assesses the values of the rate constants at some reference temperature T_0. If the entropy of activation is assumed to be independent of temperature, then the rate constants at temperature T can be written in terms of their values at T_0:

$$k_+(T) = k_+(T_0) \exp\left[\frac{-\Delta H_+^*}{k_B}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right], \tag{40}$$

$$\begin{aligned} k_-(T) &= k_-(T_0) \exp\left[\frac{-\Delta H_-^*}{k_B}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right], \\ &= k_-(T_0) \exp\left[\frac{-\left(\Delta H_+^* + \Delta H_{rx}\right)}{k_B}\left(\frac{1}{T} - \frac{1}{T_0}\right)\right], \end{aligned} \tag{41}$$

where ΔH_{rx} is the enthalpy of the net reaction. For small temperature excursions away from T_0, say 10 K–20 K, one can assume without grave error that the enthalpy values do not change from their values at T_0.

As one would expect, the expressions for the relative rate constants in (40) and (41) indicate that both the forward and the reverse rates of reactions should increase with increasing

temperature. In addition, because the ratio k_f/k_r is the equilibrium constant for the overall reaction, one may write

$$K_{\mathrm{eq}}(T) \equiv \frac{k_+(T)}{k_-(T)} = \frac{k_+(T_0)}{k_-(T_0)} \exp\left[\frac{H^{(p)} - H(r)}{k_{\mathrm{B}}}\left(\frac{1}{T_0} - \frac{1}{T}\right)\right],$$

$$= K_{\mathrm{eq}}(T_0) \exp\left[\frac{\Delta H_{rx}}{k_{\mathrm{B}}}\left(\frac{1}{T_0} - \frac{1}{T}\right)\right]. \tag{42}$$

The latter result is the well-known van't Hoff equation. Thus, the temperature dependence of equilibrium states can be captured dynamically through the temperature dependence of the balance between the forward and the reverse processes.

Example 6. Dissolution of calcium hydroxide

$Ca(OH)_2$ is one of the relatively few minerals for which dissolution is exothermic [1], with $\Delta H_{rx} = -17.88\,\mathrm{kJ\,mol^{-1}}$ at 298 K for the forward reaction in equation (37). This means that the solubility *decreases* with increasing temperature. As in example 3, the equilibrium constant at 298 K is taken to be 6.59×10^{-6}, and the rate constant for growth at 298 K is assumed to be $k_- = 0.1\,\mathrm{mol\,m^{-2}\,s^{-1}}$. According to equation (42), the solubility product at 308 K should be 5.22×10^{-6}. Furthermore, a study of the rate of growth of $Ca(OH)_2$ from supersaturated solutions, between 288 K and 313 K, has reported an activation enthalpy of $75.36\,\mathrm{kJ\,mol^{-1}}$ for growth [17]. This reported value is used for ΔH_-^*, and since $\Delta H_{rx} = -17.88\,\mathrm{kJ\,mol^{-1}}$, the activation barrier for dissolution $\Delta H_+^* = 57.48\,\mathrm{kJ\,mol^{-1}}$.

Figure 9 shows plots at 298 K and 308 K of the predicted concentration of Ca^{2+} as a function of time as $Ca(OH)_2$ dissolves in initially pure water. As in the previous example, activities were assumed equal to concentrations for simplicity in calculating the equilibrium concentrations from the equilibrium constant. As the figure indicates, the initial rate of dissolution increases with increasing temperature as it should for a thermally activated process. In addition, each curve converges to the equilibrium concentrations of $11.81\,\mathrm{mmol\,L^{-1}}$ and $10.93\,\mathrm{mmol\,L^{-1}}$ (shown by the horizonal dashed lines) at 298 K and 308 K, respectively, which are the values predicted by the van't Hoff equation. Thus, the model reproduces the temperature dependence of the equilibrium state required by the van't Hoff equation. This happens even though the equilibrium state is achieved in the model by a dynamic balance of forward and reverse reactions, each of which have thermally activated rate constants.

Example 7. Nucleation and microstructure: metal oxide hydration

The examples up to this point have focused on kinetics and equilibrium on a global scale, without much attention paid to microstructure development. The examples involved simple solutions or solution/solid interfaces—to validate kinetic and thermodynamic predictions—so the neglect of microstructure was excusable. For more complex systems, however, in which multiple minerals are present and several dissolution and precipitation processes may be occurring simultaneously, one can expect significant interactions between microstructure development and chemical kinetics.

This final example illustrates nucleation phenomena as well as interactions between microstructure and kinetics. Consider a generic hydration reaction, in which a divalent metal oxide, MO, is transformed into a sparingly soluble metal hydroxide by a through-solution mechanism. If the dissolution of the oxide is assumed irreversible, then the coupled reactions

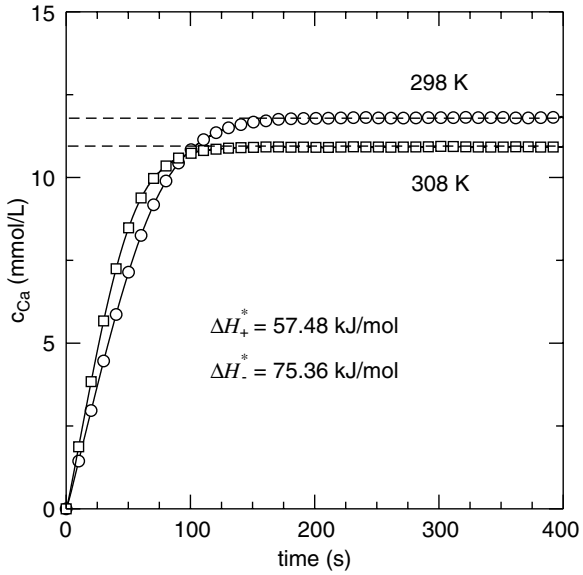

Figure 9. Simulated dissolution of $Ca(OH)_2$ in initially pure water at 298 K (circles) and 308 K (squares). For clarity, only 1/20 of the simulation points are shown, although the solid curve follows the points in between.

may be written as follows:

$$MO + H_2O \xrightarrow{k_{1,+}} M^{2+} + 2OH^-,$$

$$M^{2+} + 2OH^- \underset{k_{2,-}}{\overset{k_{2,+}}{\rightleftharpoons}} M(OH)_2.$$

The concentration of ions in solution increases by the first reaction until the system is saturated with respect to $M(OH)_2$. At even higher concentrations, there is a thermodynamic driving force for precipitation of $M(OH)_2$ but, assuming that none is present initially, nucleation is necessary before growth can occur. Homogeneous nucleation in the bulk solution is possible but, in many cases, heterogeneous nucleation on an existing solid surface has a lower value of W^*. Therefore, heterogeneous nucleation often commences at lower driving forces than those required for homogeneous nucleation.

As seen below, the details of where and when nucleation occurs can be decisive in the development of the microstructure. In addition, a net increase in total solid volume typically occurs during hydration, because the solid hydroxide has a greater molar volume than the corresponding metal oxide. In some cases, like the conversion of CaO to $Ca(OH)_2$, the volume change can be significant.

For this simulation, the initial system is a collection of spheres of MO, each 20 μm in diameter, suspended in pure water. The initial volume fraction of MO is 0.051. The computational system is a cube with 50 lattice sites in each dimension. The site spacing $\lambda = 4\,\mu$m. The temperature is maintained at 298 K, and the rate constants are $k_{1,+} = 0.01\,\text{mol}\,\text{m}^{-2}\,\text{s}^{-1}$, $k_{2,+} = 0.1\,\text{mol}\,\text{m}^{-2}\,\text{s}^{-1}$ and $k_{2,-} = 0.66\,\mu\text{mol}\,\text{m}^{-2}\,\text{s}^{-1}$. The molar volumes of MO and $M(OH)_2$ are $16.9\,\text{cm}^3\,\text{mol}^{-1}$ and $33.1\,\text{cm}^3\,\text{mol}^{-1}$, respectively. These values for the rate constants and molar volumes were chosen to be consistent with CaO/$Ca(OH)_2$ [17].

To illustrate the effects of nucleation on microstructure development, two different simulations were performed. In the first simulation, nucleation of $M(OH)_2$ occurs on a single

small seed in the solution, with a rate constant of $2.0 \times 10^{11}\,\mathrm{s}^{-1}$ and $w^* = 2.0 \times 10^9\,\mathrm{K}^3$ in (19). The seed is assumed to be inert, beyond its ability to catalyze nucleation. The nucleation barriers for homogeneous nucleation and for nucleation on the surfaces of the MO particles are assumed to be sufficiently large that nucleation can occur only on the seed particle during the entire simulation. In the second simulation, the seed is absent and nucleation occurs on the surfaces of the dissolving MO particles, with the same values of the nucleation rate constant and nucleation barrier as before.

Figure 10 shows for both simulations the initial microstructure and the predicted microstructure at 30 s, 60 s and 600 s. The seeded microstructures are shown on the left, and the unseeded microstructures undergoing surface nucleation are shown on the right. In the figures, the MO particles are shown in dark gray and the $M(OH)_2$ phase is rendered in light gray. The water between particles is not rendered. The figure demonstrates that the location and number of nucleation sites has a marked effect on the microstructure development. In both cases, the microstructure evolves rapidly at early times as the solution becomes supersaturated with respect to $M(OH)_2$. For the seeded microstructure, dissolution of MO particles, which is assumed to be irreversible, occurs rapidly and goes to completion by about 60 s. In contrast, the MO particles in the unseeded microstructure become almost completely encapsulated by precipitating $M(OH)_2$ by about 60 s. In the latter case, the $M(OH)_2$ provides an impermeable barrier that inhibits further dissolution of MO at longer times.

The changes in the microstructures can be quantified by plotting the volume fractions of MO and $M(OH)_2$ as a function of time, as shown in figure 11. The seeded microstructure exhibits an increase in $M(OH)_2$ that follows a smooth sigmoidal curve that is characteristic of phase transformations by classical nucleation and growth [19, 20]. Although not apparent on the plot, the first nucleation event in the seeded microstructure occurs by 2 s when $c_M = 151\,\mathrm{mmol\,L}^{-1}$. The rate of growth of $M(OH)_2$ then increases rapidly between 2 s and 120 s due both to the increasing surface area of $M(OH)_2$ and the increasing solution concentration. After 120 s, the rate of growth of $M(OH)_2$ decreases as the solution is depleted of M^{2+} and OH^-, until equilibrium between the solution and $M(OH)_2$ is achieved by about 900 s. The dashed line in the lower plot indicates the equilibrium concentration of M^{2+} for the parameters used in the simulation.

The situation is quite different for the unseeded microstructure, which hydrates by nucleation at the MO surfaces. The first nucleation event occurs at 0.1 s when $c_M = 9.2\,\mathrm{mmol\,L}^{-1}$. This concentration is *less than* the equilibrium concentration of $11.81\,\mathrm{mmol\,L}^{-1}$, but c_M is reported as an average for the whole solution, which is what would be measured experimentally. The local concentration near the surfaces of MO particles at 0.1 s, where nucleation of $M(OH)_2$ happens, is much higher, about $125\,\mathrm{mmol\,L}^{-1}$, at the same time. The average concentration of M^{2+} continues to rise for the first 23 s to a maximum value of $102\,\mathrm{mmol\,L}^{-1}$. This maximum concentration is considerably lower, and the subsequent decrease in concentration much more rapid in the unseeded microstructure than in the seeded one. There are two reasons for this, the first being that the volume fraction of $M(OH)_2$ is greater in the unseeded microstructure during the first minute (see upper plot in figure 11), due to the larger number of nucleation sites. The greater volume fraction corresponds to greater surface area for continued growth. The second reason for the lower maximum concentration in the unseeded microstructure is that the MO surfaces become partially covered by a protective layer of $M(OH)_2$, giving them less exposed surface area for continued dissolution.

At intermediate times in the unseeded microstructure, between 60 s and 600 s, figure 11 indicates that the changes in volume fraction of both solid phases are approximately linear in time (filled symbols in the upper plot), while the average composition of the solution remains approximately constant (filled symbols in the lower plot). Such behavior indicates that the

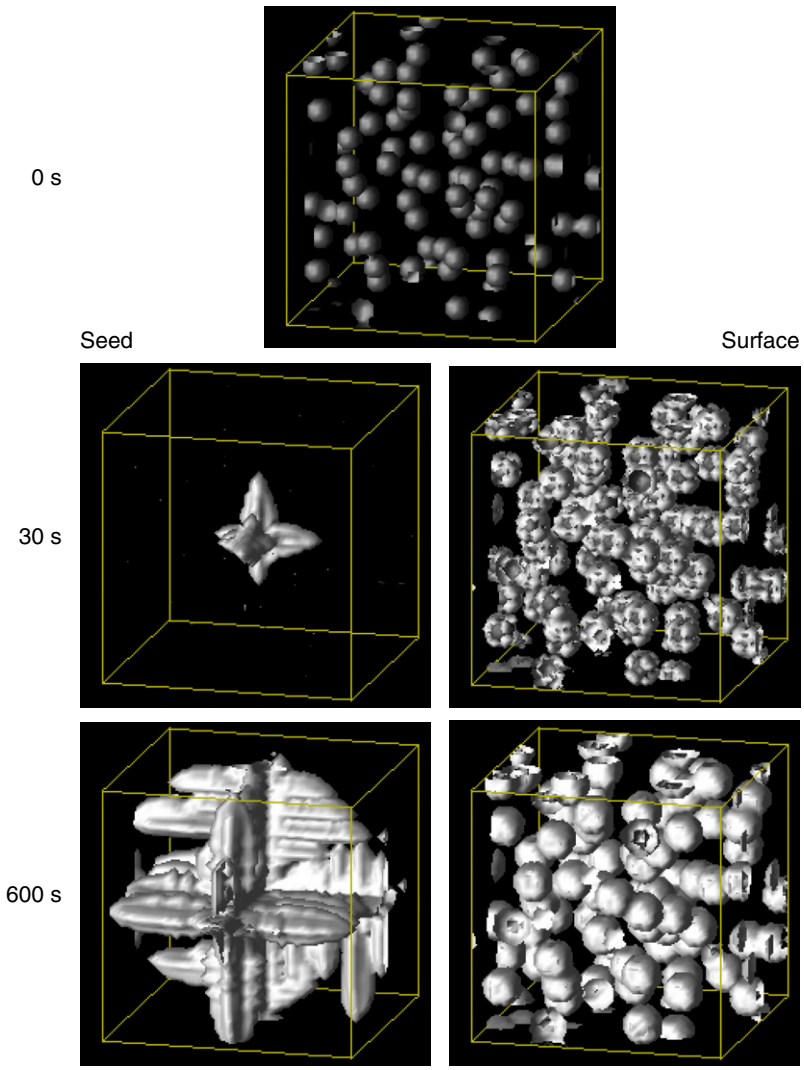

Figure 10. Predicted microstructure development as MO particles dissolve in water and $M(OH)_2$ nucleates either on a single inert seed (left) or on the surfaces of the dissolving particles (right). The temperature is kept constant at 298 K. Other than the difference in nucleation sites, all model parameters are identical in both simulations.

nutrients for continued growth of $M(OH)_2$ are being drawn, not from the bulk solution, but from freshly dissociated material immediately adjacent to the solid MO surfaces. Thus, although growth occurs by a 'through-solution' mechanism, it is almost indistinguishable from a direct topochemical conversion of MO to $M(OH)_2$. Only after 600 s does the bulk solution become further depleted, as the MO surfaces become almost completely covered.

4. Discussion

The model presented here uses algorithms for simulating both diffusion in non-electrolytes, (1), and the law of mass action for homogeneous reactions, (4) and (5) that are basically the same as

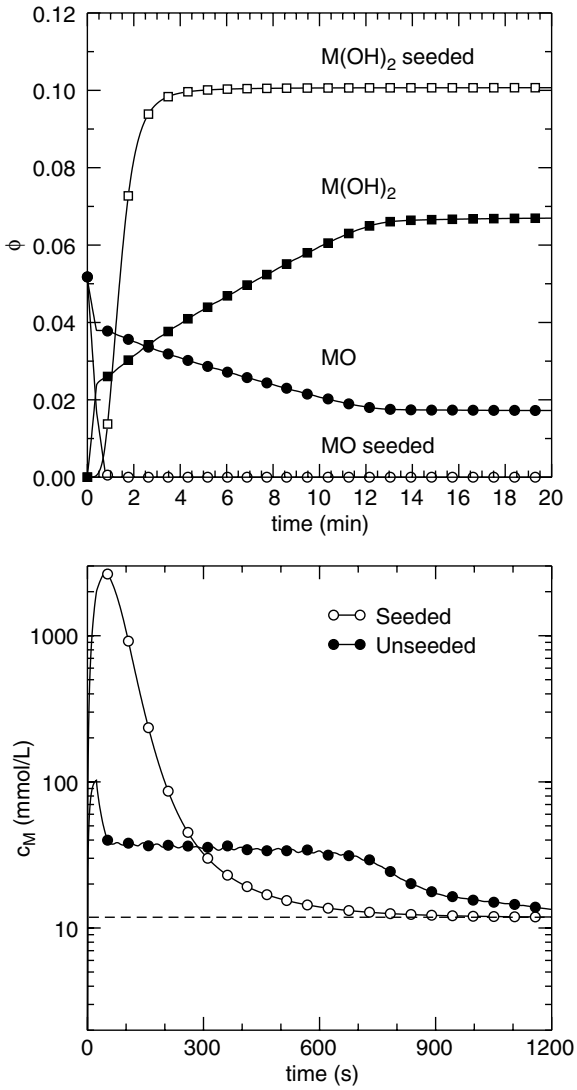

Figure 11. Volume fractions ϕ of MO and $M(OH)_2$ (upper plot) and concentration of M^{2+} (lower plot) as a function of time for the seeded and unseeded simulations shown in figure 10. For clarity, only 1/200 of the points are shown, although the solid curves follow the points in between.

those given earlier by Karapiperis and Blankleider [2]. But this model also incorporates several important modifications that enable the modeling of more realistic chemistry and physics in aqueous systems. Among these are (1) the ability to simulate diffusion in nonideal electrolyte solutions, (2) the modeling of heterogeneous reactions with correct scaling behavior (3) the estimation of activity coefficients for dissolved species, (4) explicit modeling of the kinetics of nucleation and (5) the temperature dependence of the relevant rate processes.

The modified Davies equation [11] was used here to estimate the activity coefficients of ionic solute species. Although valid only at fairly dilute concentrations, between 0.01 mol L^{-1} and 0.1 mol L^{-1}, this equation yields remarkably better predictions of equilibrium states than those produced by equating activity with molar concentration (see example 5). The modified

Davies equation also has the advantage that the activity coefficient of a particular ion is dependent only on temperature, ionic strength and properties of the ion itself. At higher concentrations, better approximations of activity coefficients of electrolytes can be obtained using the ion-interaction approach of Pitzer [21], although its implementation is considerably more computationally expensive. But if the goal of a modeling investigation is to obtain accurate predictions of equilibrium solution composition, including ion complexation effects, then thermodynamic modeling packages that implement approaches like that of Pitzer are preferable to the model described here. The primary strength of the present model is its ability to simulate spatial and temporal variations in microstructure at a local scale that are due to transport and any number of coupled homogeneous and heterogeneous reactions.

By including rate constants for nucleation of solids, the present model can simulate the unstable supersaturated states and finite induction times in solutions prior to precipitation of a solid phase [7]. These phenomena are extremely important in modeling the early-time kinetics of a wide variety of phase transformations, including evaporation-assisted precipitation of salts [22], production of slaked lime [23] and the hydration of gypsum, portland cement, and other cementitious materials [23–25]. In fact, this work was motivated largely by its intended application to model microstructural changes in these kinds of systems.

By accounting for the dependence of kinetics and equilibrium states on temperature, this model is able to simulate not only isothermal systems at different temperatures, as illustrated in example 6, but also adiabatic or semi-adiabatic systems in which heat generated (absorbed) by exothermic (endothermic) reactions causes temperature changes within the system. These temperature changes can be calculated from the volume fractions and heat capacities of the phases in the system, and so will be a function of the reaction history and the instantaneous state of the microstructure.

Because the model is based on principles of transition state theory, a complete description of the kinetics of any *elementary* reaction over a modest range of temperatures requires the input of values of k_+ and k_- at some reference temperature (or only one rate constant when K_{eq} is known), together with the activation enthalpy for either the forward or reverse direction and the enthalpy of reaction. For a very large number of reactions, the equilibrium constant and enthalpy of reaction have been catalogued [1]. When a reaction is elementary, the rate constant for, say, the forward reaction can be experimentally obtained by measuring the instantaneous rate of formation of a product of the reaction when there are no products of that reaction in the system that can participate in the reverse reaction. In addition, the enthalpy of activation in the forward direction can be measured by performing the same experiment at several different temperatures and then plotting the logarithm of the instantaneous rate as a function of $1/T$. Again, provided the reaction is elementary over this temperature range, such a plot should be linear with slope $-\Delta H_+^*/k_{\mathrm{B}}$. Therefore, the acquisition of the necessary parameters for elementary reactions is straightforward, at least in principle.

Regrettably, very few of the important reactions in materials science or environmental geochemistry are actually elementary. The distinguishing feature of a nonelementary rate process is that the concentration exponents on the right-hand side of the rate equation are not equal to the corresponding molar stoichiometric coefficients of the reactants. However, even when it appears to be nonelementary, a reaction must still be composed of a sequence of reaction steps, each of which is itself elementary. Therefore, within the context of this model, one of two approaches can be taken to acquire the necessary input data for nonelementary reactions. The first approach is to undertake the determination of the sequence of elementary reactions that are involved. In many cases, only one of these elementary reactions will be rate-controlling. When one reaction in the sequence is rate-controlling, then all the reactions that precede it must be near equilibrium. In such cases, the forward rate constant and activation enthalpy can be

measured for the rate-controlling reaction by the kinds of experiments described already. The difficulty, which often renders this approach impractical, is that the entire sequence of elementary reactions must be identified to use this approach.

When a determination of the sequence of elementary steps in a nonelementary reaction is impractical, an acceptable compromise may be to apply the standard rate equation as if the overall reaction were elementary, i.e. equating the exponents in the rate equation to the molar stoichiometric coefficients. The nonelementary aspects of the reaction could then be approximated by forcing the rate 'constant' to be some function of the concentrations of other species that are experimentally determined to influence the apparent rate of the overall reaction. Such an approach could be useful for modeling the effect of a catalyst or a poisoning agent on the rate of a highly irreversible reaction. The main objection to this second approach is that, for reversible reactions that can approach equilibrium, the simulation will generally predict a dynamic equilibrium state that differs from that required by the equilibrium constant (see equation (36)).

Finally, some comments are warranted about the morphology of the precipitates shown in example 7. In the seeded microstructure especially, the shape of the growing particle is reminiscent of the dendritic structures that are often observed during solidification of metals. However, dendrite formation during solidification is caused by a Mullins–Sekerka instability [26], which is related to the coupling between the local shape of the solid–liquid interface and the dissipation of the heat of solidification. In contrast, heat dissipation is not incorporated in example 7. Instead, the dendritic shape of the particle is an artifact both of the geometry of the computational lattice and of the smoothing algorithm, discussed in section 3.3, which shares material only between nearest neighbors. Both of these influences cause greater growth rates in the directions of the six nearest neighbors than in other directions. However, once the dendritic branching begins, it is exacerbated by its own geometry because the branch tips have greater specific surface area at which growth can occur. More isotropic growth behavior could be achieved by modifying the smoothing process so that 2nd-nearest and 3rd-nearest neighbors are included.

Many minerals exhibit well-defined, anisometric shapes that are due to their growth rate being a function of' crystallographic orientation. For example, portlandite ($Ca(OH)_2$) often grows in aqueous solutions as large, thin hexagonal plates, whereas gypsum ($CaSO_4.2H_2O$) grows as long needles. The model presented here does not account for anisotropic growth, but it could do so in a limited way by assigning a crystal orientation vector to newly formed nuclei and making the local growth rate constant in (10) depend upon the angle between this vector and the growth direction.

5. Conclusions

At the outset of this paper, three objectives were given for the development of a 3D microstructure model for aqueous mineral systems. First, *it should be based on principles of reaction kinetics and thermodynamics*. The stochastic model presented here uses transition state theory, the law of mass action, the classical theory of nucleation and the microscopic principles of diffusion. The examples given in the previous sections demonstrate that the model is stable, that it converges to the correct diffusion rate laws and that it accurately predicts a broad range of kinetic and equilibrium phenomena. Second, the model *should capture the relationships between reactions, transport and microstructure development*. As example 7 demonstrates, the model is able to simulate the interdependence of nucleation and growth mechanisms, solute transport and microstructure development. Finally, the model *should be as general as possible*. This paper has shown a broad range of reaction and

diffusion phenomena that can be simulated within a unified framework. Future work will expand its capabilities to include approximations of anisotropic crystal growth and reactions involving nonstoichiometric materials, but it already represents a research tool that can be tailored to particular materials systems to investigate relationships between their processing, microstructure and properties.

References

[1] Parkhurst D L 1995 User's guide to PHREEQC—a computer program for speciation reaction-path, advective-transport, and geochemical calculations *Water-Resources Investigations Report* 95-4227 US Geological Survey

[2] Karapiperis T and Blankleider B 1994 Cellular automaton model of reaction-transport processes *Physica D* **78** 30–64

[3] Karapiperis T 1995 Cellular automaton model of precipitation/dissolution coupled with solute transport *J. Stat. Phys.* **81** 165–80

[4] Harned H S and Owen B B 1958 *The Physical Chemistry of Electrolytic Solutions* (New York: Reinhold Publishing) 3rd edn

[5] Bear J 1972 *Dynamics of Fluids in Porous Media* (New York: Elsevier)

[6] Darken L S and Gurry R W 1953 *Physical Chemistry of Metals (Metallurgy and Metallurgical Engineering Series)* (New York: McGraw-Hill)

[7] Kaschiev D and van Rosmalen G M 2003 Review: nucleation in solutions revisited *Cryst. Res. Technol.* **38** 555–74

[8] Bullard J W 2007 Approximate rate constants for nonideal diffusion and their application in a stochastic model *J. Phys. Chem.* A **111** 2084–92

[9] de Groot S R and Mazur P 1984 *Non-Equilibrium Thermodynamics* (New York: Dover)

[10] Mills R and Lobo V M M 1989 *Self-Diffusion in Electrolyte Solutions* (Amsterdam: Elsevier)

[11] Samson E, Marchand J and Beaudoin J J 2000 Modeling the influence of chemical reactions on the mechanisms of ionic transport in porous materials: an overview *Cem. Concr. Res.* **30** 1895–902

[12] Press W H, Teukolsky S A, Vetterling W T and Flannery B P 1997 *Numerical Recipes in C* (London: Cambridge University Press) 2nd edn

[13] Gibbs J W 1928 *The Collected Works of J. Willard Gibbs, Thermodynamics* vol 1 (New Haven, CT: Yale University Press)

[14] Cahn J W and Hilliard J E 1958 Free energy of a nonuniform system: I. Interfacial free energy *J. Chem. Phys.* **28** 258–67

[15] Wang S L, Sekerka R F, Wheeler A A, Murray B T, Coriell S R, Braun R J and McFadden G B 1993 Thermodynamically-consistent phase-field models for solidification *Physica D* **69** 189–200

[16] Garboczi E J, Bentz D P and Martys N S 1999 Digital images and computer modelling *Methods in the Physics of Porous Media* vol 35 ed P Z Wong (San Diego, CA: Academic) pp 1–41

[17] Tadros M E, Skalny J and Kalyoncu R S 1976 Kinetics of calcium hydroxide crystal growth from solution *J. Colloid Interface Sci.* **55** 20–4

[18] Glasstone S, Laidler K J and Eyring H 1941 *The Theory of Rate Processes* (New York: McGraw-Hill)

[19] Avrami M 1941 Granulation, phase change, and microstruture: kinetics of phase change. III *J. Chem. Phys.* **9** 177–84

[20] Burke J E and Turnbull D 1952 Recrystallization and grain growth *Prog. Met. Phys.* **3** 220–92

[21] Pitzer K S 1991 Ion interaction approach: theory and data correlation *Activity Coefficients in Electrolyte Solutions* ed K S Pitzer (Boca Raton, FL: CRC Press) 2nd edn

[22] Franks F and Murase N 1992 Nucleation and crystallization in aqueous systems during drying *Pure Appl. Chem.* **64** 1667–72

[23] Taylor H F W 1997 *Cement Chemistry* (London: Thomas Telford) 2nd edn

[24] Garrault S and Nonat A 2001 Hydrated layer formation on tricalcium and dicalcium silicate surfaces: experimental study and numerical simulations *Langmuir* **17** 8131–8

[25] Gartner E M and Gaidis J M 1989 Hydration mechanisms: I. *Materials Science of Concrete* vol 1 ed J Skalny (Westerville, OH: American Ceramic Society) pp 95–125

[26] Mullins W W and Sekerka R F 1964 Stability of a planar interface during solidification of a dilute binary alloy *J. Appl. Phys.* **35** 444–51