# Hybrid deterministic/stochastic algorithm for large sets of rate equations

M. Gherardi<sup>a</sup>, T. Jourdan<sup>a,\*</sup>, S. Le Bourdiec<sup>b</sup>, G. Bencteux<sup>b</sup>

<sup>a</sup>CEA, DEN, Service de Recherches de Métallurgie Physique, F-91191 Gif-sur-Yvette, France <sup>b</sup>EDF R&D, F-92141 Clamart, France

## Abstract

We propose a hybrid algorithm for the time integration of large sets of rate equations coupled by a relatively small number of degrees of freedom. A subset containing fast degrees of freedom evolves deterministically, while the rest of the variables evolves stochastically. The emphasis is put on the coupling between the two subsets, in order to achieve both accuracy and efficiency. The algorithm is tested on the problem of nucleation, growth and coarsening of clusters of defects in iron, treated by the formalism of cluster dynamics. We show that it is possible to obtain results indistinguishable from fully deterministic and fully stochastic calculations, while speeding up significantly the computations with respect to these two cases.

Keywords: rate theory, cluster dynamics, Monte Carlo, stiffness, iron, helium, irradiation

# 1. Introduction

The use of rate equations covers a wide range of applications in fields such as chemistry, biology and physics. In material science it has proved to be successful to study long term evolution processes that affect the microstructure of materials. This approach, also known as "rate theory" or "cluster dynamics", enables one to study phenomena such as homogeneous precipitation under thermal ageing [1] or formation of voids and interstitial loops under irradiation [2–4]. The system is considered as a gas of clusters, which are defined by their solute and self-defect (vacancies, self-interstitials) content, and which can emit and absorb mobile species. The evolution of each cluster type is given by a rate equation, which also involves the number or concentration of other types of clusters.

In general, such a set of ordinary differential equations (ODE) is stiff, which roughly means that some modes have very small time constants with respect to others. In such cases, implicit integration methods are necessary, which implies the solving, at each time step, of a linear system whose size is given by the number of cluster types considered in the simulation. Since it is customary to simulate large clusters (containing several millions single species), the solving of such an ODE set can become challenging, from a point of view of both computation time and memory storage. It is all the more the case when multicomponent systems are treated: besides industrial alloys, some simple cases include the modelling of pure metals with solute impurities such as carbon, nitrogen, etc., and gas atoms, such as helium and hydrogen [5]. When self-defects and a single type of impurity are considered, some numerical approximations have proved useful to reduce the number of equations [6, 7]. However, alternatives must be envisaged once additional components are included in the simulations.

One interesting way to solve these equations is to use a Monte Carlo algorithm [8], as suggested by Gillespie in a general framework [9]. At variance with other fields, where stochastic effects must sometimes be taken into account to correctly model noisy systems [10], differences in the physical results are not expected in our case, but using Monte Carlo simulations leads to a tremendous decrease in memory storage. Since each reaction firing has to be processed, it can however be computationally demanding when reactions occur frequently; due to the stiffness of the underlying ODE set, this case is quite common, and approximations are needed. Some methods rely on the grouping of events [11], others are based on the partitioning between slow and fast reactions, which are treated differently [12]. Recently, a hybrid stochastic and deterministic algorithm has been devised for small biochemical systems, with a stochastically exact coupling [14]. We have adapted such an algorithm to the case of a large set of species, coupled together by a small number of species.

It should be emphasized here that the reason for the use of such an hybrid method in material science is fundamentally different from what motivates this kind of approach in biology. In general, for biochemical systems, the deterministic equations are introduced as a useful approximation of the chemical master equation [13], but can lead to erroneous results if the approximation is not carefully controlled. In our case, we aim to speed up simulation of systems which are already correctly described by deterministic rate equations, by introduction of a stochastic part.

In the following we describe the algorithm in a general way and present some simple test cases taken from material science, namely the formation kinetics of voids and interstitial loops in iron and in iron with helium under irradiation. Provided an approximation is made to the initial algorithm devised for small biochemical systems, hybrid simulations are found to be significantly faster than the purely deterministic and purely stochastic simulations.

<sup>\*</sup>Corresponding author. Tel.:+33 1 69 08 73 44. Fax: +33 1 69 08 68 67 Email address: thomas.jourdan@cea.fr (T. Jourdan)

Preprint submitted to Computer Physics Communications

#### 2. Background and notation

We consider a well-stirred system of species (labeled by indices in a finite set  $\Lambda$ ) interacting through reaction channels (labeled by indices in a finite set  $\mathcal{R}$ ) in a volume *V*. The dynamical state of the system at time *t* is specified by the set  $X_{\Lambda}(t) \equiv \{X_i(t)\}_{i \in \Lambda}$ , whose elements are the numbers of particles of type *i*. The probability, given the state at time *t*, that a reaction of type  $j \in \mathcal{R}$  will occur in the time interval [t, t + dt] can be written under general assumptions as [15]

$$P(\text{reaction } j \text{ during } [t, t + dt] | X_{\Lambda}(t)) = a_j(X_{\Lambda}(t)) dt, \quad (1)$$

where  $a_j$  is called the *propensity function* for the *j*-th reaction channel. The change in the number of molecules of type *i* produced by a reaction of type *j* is called the *state-change vector*  $v_{ij}$ . All propensity functions and state-change vectors completely specify the reaction channels. Notice that information about the space localization of the particles in the physical system is completely neglected in this framework.

Eq. (1) lies at the heart of the *Stochastic Simulation Algorithm* (SSA), introduced by Gillespie [9], which is based on the observation that the probability at a given time  $t_0$  that the next reaction is *j* and occurs in the time interval [t, t+dt] is  $p(j, t|t_0)dt$ , with

$$p(j,t|t_0) = a_j(X_{\Lambda}(t_0)) \exp\left[-(t-t_0)A_{\mathcal{R}}(X_{\Lambda}(t_0))\right], \quad (2)$$

where we have introduced the total propensity

$$A_{\mathcal{R}}(X_{\Lambda}(t)) = \sum_{j \in \mathcal{R}} a_j \left( X_{\Lambda}(t) \right).$$
(3)

Regardless of the reaction *j*, the probability that at a given time  $t_0$ , a reaction occurs in [t, t + dt] is  $p(t|t_0)dt$ , with

$$p(t|t_0) = A_{\mathcal{R}}(X_{\Lambda}(t_0)) \exp\left[-(t-t_0)A_{\mathcal{R}}(X_{\Lambda}(t_0))\right].$$
(4)

From Eq. (2) and (4), one infers that the probability that the reaction occurring in [t, t + dt] is of type *j* reads

$$\frac{p(j,t|t_0)}{p(t|t_0)} = \frac{a_j(X_\Lambda(t_0))}{A_\mathcal{R}(X_\Lambda(t_0))}.$$
(5)

The evolution of the system is simulated by computing realizations of the continuous-time Markov chain  $X_{\Lambda}(t)$ . By calling  $N_j(t)$  the stochastic process counting the number of times that a reaction of type *j* fires in the interval [0, *t*] (which is an inhomogeneous Poisson process of rate  $a_j(X_{\Lambda}(t))$  by Equation (1)), one has

$$dX_i(t) = \sum_{j \in \mathcal{R}} v_{ij} dN_j(t).$$
(6)

Operatively, at each time step  $t_0$  of the SSA the following operations are performed:

- 1. Compute all propensities  $a_j(X_{\Lambda}(t_0))$ , together with their sum  $A_{\mathcal{R}}(X_{\Lambda}(t_0))$ .
- 2. Select a reaction *j* with probability proportional to its propensity  $a_i(X_{\Lambda}(t_0))$ , following Equation (5).

3. Select a next-reaction time *t* by sampling the probability distribution in Equation (4), which means that

$$t - t_0 = \frac{\xi}{A_{\mathcal{R}}(X_{\Lambda}(t_0))},\tag{7}$$

where  $\xi$  is a random number exponentially distributed with parameter 1.

4. Update time  $(t_0 \mapsto t)$  and configuration  $(X_i \mapsto X_i + v_{ij}$  for each  $i \in \Lambda$ ).

A consequence of Equation (4) is that the larger the total propensity  $A_{\mathcal{R}}$  is, the slower the evolution. In fact, the typical time between two reactions is  $1/A_{\mathcal{R}}$ . This problem is most painful when the system is *stiff* — meaning that at least two very different time scales are present - and just a few reaction channels contribute most to the total propensity. In such a case the algorithm spends most of the time performing fast reactions and taking small time steps. Some solutions to speed up the simulations have been proposed, based on diverse ideas such as reuse of pseudo-random numbers [16], quasi-steady-state theory [17], grouping of reactions (the so-called *tau leaping* methods) [11]. In many applications where stochastic noise is negligible for all practical purposes, such as in cluster dynamics [4], a commonly followed approach consists instead in solving a deterministic system of equations describing the evolution of average quantities. More precisely, a well-known consequence of Equation (1) is the chemical master equation (see for instance [18]), which in turn implies the following rate equations, provided fluctuations are neglected:

$$\frac{\mathrm{d}X_i(t)}{\mathrm{d}t} = \sum_{j \in \mathcal{R}} v_{ij} a_j \left( X_\Lambda(t) \right). \tag{8}$$

In this equation,  $X_i$  are now real functions of time. The precise conditions under which such an evolution "on average" is a good approximation of the original jump process are described in Reference [18]. In our case populations  $X_i$  are large for realistic volumes V, so we will assume that these conditions are satisfied. The rate equations can then be integrated deterministically using various numerical methods adapted to ODE systems.

### 3. The algorithm

#### 3.1. The hybrid method

The deterministic approach is not well-suited to situations where the number of degrees of freedom is large, due to the poor scaling of known numerical integration methods. What we propose here is to adopt a strategy similar to the one already proposed by Alfonsi *et al.* [14] for applications in biochemistry, namely splitting the reactions into two sets, ideally according to their typical time scale, and modeling one of them stochastically through the SSA and the other one deterministically by using one among the well-known numerical methods for ODE systems. In this subsection we recall briefly the hybrid algorithm presented in Reference [14]. Let the set of reactions  $\mathcal{R}$  be partitioned into two disjoint subsets,  $\mathcal{R} = \mathcal{D} \cup \mathcal{S}$ . The evolution equation for  $X_i(t)$  will be given by

$$dX_i(t) = \sum_{j \in \mathcal{D}} v_{ij} a_j \left( X_\Lambda(t) \right) dt + \sum_{j \in \mathcal{S}} v_{ij} dN_j(t), \tag{9}$$

where  $X_i(t)$  are real variables. The idea is to evolve the system deterministically by using only the reactions in  $\mathcal{D}$ , and account for the second sum in the right-hand side by performing reactions in S stochastically. More precisely, we consider the reduced system whose right-hand side function is given by the sum of the reaction terms on the set  $\mathcal{D}$ 

$$\frac{\mathrm{d}X_i(t)}{\mathrm{d}t} = \sum_{j\in\mathcal{D}} v_{ij} a_j \left( X_\Lambda(t) \right),\tag{10}$$

then at any given time  $t_0$  we compute its numerical solution, with initial condition  $X_i(t_0)$ , up to the first-reaction time  $t > t_0$ , when we perform a stochastic reaction and update the configuration accordingly. This stochastic reaction is instantaneous, in that it causes a change in the variables  $(X_i \mapsto X_i + v_{ij} \text{ for } i \in \Lambda)$  without increasing time.

The problem of how to compute this reaction time is nontrivial. Indeed, contrary to the standard SSA, propensities evolve between two reactions treated stochastically due to the deterministic evolution of the system. As shown by Alfonsi *et al.* [14], it is possible to find the first-reaction time by solving the following equation:

$$g(t|t_0) \equiv \int_{t_0}^t A_{\mathcal{S}}(X_{\Lambda}(\tau)) \, \mathrm{d}\tau = \xi, \tag{11}$$

where  $\xi$  is an exponentially distributed random number with parameter 1 and  $A_S$  is the *total stochastic propensity* 

$$A_{\mathcal{S}}(X_{\Lambda}(t)) = \sum_{j \in \mathcal{S}} a_j(X_{\Lambda}(t)).$$
(12)

One can see that Equation (11) is a straightforward generalization of Equation (7) to the case when propensities vary between two stochastic reactions.

#### 3.2. Partitioning the variables

With the foregoing discussion, the original problem has been split into a deterministic and a stochastic part, in such a way that the latter is faster than a full SSA algorithm, provided that the sum of all stochastic propensities (12) is less than the total propensity (3). On the other hand, inspection of Equation (9) shows that we have not gained much on the deterministic side, since the number of equations to be treated has not changed. In fact, reaction terms with  $j \in D$  are still involved in the equations for all the variables  $X_i$  with  $i \in \Lambda$ . In the type of hybridization proposed in the biochemical literature one can afford having the deterministic solver handle all species  $X_{\Lambda}$  because one is typically in the situation where the number of degrees of freedom is modest and a full SSA would be slower. In the present case instead we want to take advantage of both approaches and tackle systems with a large number of degrees of freedom. In order to address the problem in a general way, we propose to split the variables into two sets (in the same spirit as in [19]), by identifying a proper subset  $\Delta$  of  $\Lambda$  and treating variables with indices inside or outside  $\Delta$  differently. A partition of the reactions is induced by the choice of  $\Delta$ . We impose that a reaction is treated deterministically if and only if *every* species involved in it (both the reactants and the products) belongs to  $\Delta$ . With this definition we can rewrite Equation (9) as

$$\begin{cases} dX_{i\in\Delta}(t) = \sum_{j\in\mathcal{D}} v_{ij}a_j \left(X_{\Delta}(t)\right) dt &+ \sum_{j\in\mathcal{S}} v_{ij}dN_j(t) \\ dX_{i\notin\Delta}(t) = \sum_{j\in\mathcal{S}} v_{ij}dN_j(t) \end{cases}$$
(13)

where  $X_{\Delta}(t)$  is a shorthand for the set of all  $X_i$  with  $i \in \Delta$ . Species outside  $\Delta$  evolve only through stochastic terms. Moreover, the deterministic reactions in the other set of equations now only depend on the species inside  $\Delta$  so that the deterministic part of the method must deal with a reduced number of degrees of freedom. Namely, all one has to do in between two successive stochastic reactions is to solve the upper-left corner of the system (13), that is

$$\frac{\mathrm{d}}{\mathrm{d}t}X_{i\in\Delta}(t) = \sum_{j\in\mathcal{D}}v_{ij}a_j\left(X_{\Delta}(t)\right).$$
(14)

Stochastic reactions are then triggered by looking at the timechange function g (11). General knowledge about deterministic solvers and the SSA, together with the discussion at the beginning of this subsection, suggest that we should choose  $\Delta$  so that a few (fast) degrees of freedom are treated deterministically, while the bulk of them should be left to the stochastic part.

### 3.3. Decoupling and the final algorithm

Actual implementation of the hybrid method as it is described above brings out some technical difficulties. The main problem resides in the abrupt changes in  $X_i$  caused by the stochastic reactions, since such variations in the right-hand side function can considerably slow down the deterministic solving of the ODE set. In order to reduce the severity of this behavior we introduce an approximation.

We define another subset D of  $\Lambda$ , such that  $D \subseteq \Delta$ , and decouple the deterministic and stochastic dynamics inside D. What this means is that whatever stochastic reaction occurs, only the variables  $X_i$  with  $i \in \Delta \setminus D$  will be updated, while those inside D will be left unchanged. One of course has to compensate in some way for the modification imposed to the dynamics. We do this by adding an effective source term  $\rho_i$  to each equation in D, crafted in such a way as to approximately balance the discrepancy. Namely, we take  $\rho_i$  to be the mean over all possible stochastic reactions j — of the change that  $X_i$ would have undergone if it were not decoupled, weighted by the reaction's propensity:

$$\rho_i = \sum_{j \in \mathcal{S}} v_{ij} a_j \left( X_{\Lambda}(t) \right). \tag{15}$$

Since  $X_i$  is not updated when a stochastic reaction occurs if  $i \in D$ , the deterministic solving is much more efficient, provided

most stochastic reactions involving species in  $\Delta$  actually involve species in D and no species in  $\Delta \setminus D$ . A possible source of slowing down for the deterministic solver is the abrupt change of propensities in  $\rho_i$  when a stochastic reaction occurs. However, in general the firing of a single stochastic reaction does not change  $\rho_i$  appreciably, so the deterministic solver is virtually not affected. We are not going to give analytical estimates of the errors this approximation introduces to the dynamics; we justify it *a posteriori* by checking the numerical results against exact simulation methods. Let us just note here that inclusion of such effective source terms essentially restores the full right-hand side of Equation (8) for the species in D.

Finally, the complete algorithm is as follows:

- 1. At time  $t_0$  generate random number  $\xi$  exponentially distributed with parameter 1.
- 2. Compute  $\rho_i$  for each  $i \in D$  as in (15).
- 3. Evolve the variables  $X_i(t)$  with  $i \in \Delta$  following

$$\begin{cases} \frac{\mathrm{d}}{\mathrm{d}t} X_{i \in D}(t) = \sum_{j \in \mathcal{D}} v_{ij} a_j \left( X_\Delta(t) \right) + \rho_i \\ \frac{\mathrm{d}}{\mathrm{d}t} X_{i \notin D}(t) = \sum_{j \in \mathcal{D}} v_{ij} a_j \left( X_\Delta(t) \right). \end{cases}$$
(16)

Keep track of the evolution of the total stochastic propensity (12) (which depends also on variables in  $\Delta$ ) and of the time-change function *g* (11).

- 4. When g(t|t<sub>0</sub>) = ξ, choose a stochastic reaction j ∈ S with probability proportional to its propensity and update X<sub>i</sub> → X<sub>i</sub> + v<sub>ij</sub> for each i ∈ Λ \ D
- 5. Update  $t_0$  and go back to step 1.

# 3.4. Remarks on the implementation

The method presented above leaves some open choices to the implementer. Apart from the parameters V,  $\Delta$  and D (the tweaking of which should be done on a case-by-case basis; see Section 4), some freedom is left as far as other details are concerned.

Let us consider the computation of the effective term  $\rho_i$  at step 2 of the algorithm. One could choose not to compute it at every stochastic step, but rather to update it just once every  $n_{\rho}$ reactions. Moreover, one could choose to sample and hold its value between two updates, thus neglecting its smooth dependence on the variables inside the deterministic region  $\Delta$  between two reactions. We tested these different approximations for the cases described in Section 4 and found no big differences in performance or precision for values of  $n_{\rho}$  between 1 and 10<sup>5</sup>.

An important design-related issue concerns the data structure used to store the variables outside  $\Delta$ , which should be the large majority. For the most frequently used reaction types [20], propensity functions are such that some of them vanish whenever a variable takes on the special value 0. The best choice is then to use a dynamical structure, so that slots are assigned to  $X_{i\in\Lambda\setminus\Delta}$ only when it gets above zero, and they are released whenever it goes back to zero, in such a way as to take full advantage of the shape of the distribution  $X_{\Lambda}$  and avoid the computation of useless propensities. This is the implementation we use in the next section.

In order to solve the stiff set of ordinary differential equations in the hybrid model (as well as in the fully deterministic simulations that will be used as reference), a multistep method with backward differentiation formulas is used, as implemented in CVODES from the SUNDIALS package [21]. This solver contains a root-finding procedure, which we use to detect time *t* when  $g(t|t_0) = \xi$ , i.e. when a stochastic reaction occurs. Since the method is implicit, it requires the solving of a nonlinear system, which is done using Newton iteration. The solving of linear systems in the Newton iteration is carried out by a dense direct solver. Although more sophisticated methods can be used for the solving of the linear system, this one is chosen for its simplicity of implementation.

### 4. Results

# 4.1. Cluster dynamics model

The model we use here, called cluster dynamics, describes the time evolution of defect clusters in materials, within a meanfield approach [4]. Clusters (or species) can absorb and emit some species that diffuse in an effective medium. Generally, the set of mobile species is limited to small clusters.

In the following we assume that clusters are made of two elements, namely self-defects and solute (S) atoms, but the method could be extended to more complicated clusters. Selfdefects can either be vacancies (V) or self-intertitial (I) atoms. Clusters are thus identified by  $i = (n, p) \in \mathbb{Z} \times \mathbb{Z}^+ \setminus \{(0, 0)\}$ , where |n| is the number of self-defects (n is positive for self-interstitials and negative for vacancies) and  $p \ge 0$  is the number of solute atoms in the cluster. When no solute is present, clusters are simply labeled by  $i = n \in \mathbb{Z} \setminus \{0\}$ . Extended defects, such as dislocations, grain boundaries and surfaces, are treated as sinks and sources for mobile clusters. The evolution equation for the concentration or number  $X_i$  of species of type i is

$$\frac{\mathrm{d}}{\mathrm{d}t}X_{i} = G_{i} + \sum_{l \in M_{i}} J_{i-l,i} - \sum_{l \in M} J_{i,i+l} - \sum_{l \in \Lambda} J_{l,i+l} - k_{i}D_{i}(X_{i} - X_{i}^{e}).$$
(17)

In Eq. (17),  $G_i$  is a creation rate of species *i* by irradiation,  $k_i$  is the sink strength due to dislocations, grain boundaries or surfaces,  $D_i$  is the diffusion coefficient and  $X_i^e$  is the concentration or number of species *i* in thermodynamical equilibrium with the sinks (only valid for self-defects). A is the set containing all species and *M* contains only mobile species. The set  $M_i$  is a subset of *M*, to ensure that the species involved in the flux  $J_{i-l,i}$  exist (the number of solute atoms of i - l must be non-negative). The flux  $J_{i,i+l}$  is defined by

$$J_{i,i+l} = \beta_{i,l} X_i X_l - \alpha_{i+l,l} X_{i+l}, \qquad (18)$$

where  $\beta_{i,l}$  and  $\alpha_{i+l,l}$  are absorption and emission rates respectively. The absorption rate is defined by solving the diffusion equation in stationary state around an isolated cluster, which leads to

$$\beta_{i,l} = r_{i,l} D_l / V, \tag{19}$$



Figure 1: (Color online) Partition of the set of clusters. Shaded areas correspond to the deterministic  $(\Delta)$ , decoupling (D) and mobile (M) regions. The arrows describe an emission reaction; the hollow point represents a cluster whose variable is not affected by the reaction, since it lies inside D.

where  $r_{i,l}$  has the dimension of a distance and is related to the geometry of clusters *i* and *l*, and *V* is the total simulation volume if *X* is a number of species, or is equal to 1 if *X* is a concentration. The emission coefficient is defined by

$$\alpha_{i+l,l} = \frac{r_{i,l}D_l}{V_{\text{at}}} \exp\left(-\frac{F_{i+l,l}}{k_{\text{B}}T}\right).$$
(20)

In this equation  $V_{\text{at}}$  is the atomic volume,  $F_{i+l,l}$  is the binding free energy of species *l* to cluster i + l,  $k_{\text{B}}$  is the Boltzmann constant and *T* is the temperature.

The number of ODEs to solve can amount to several millions with only vacancies, self-interstitials and solute atoms. In addition the propensities can be different by several orders of magnitude. The resolution of such a stiff system must be done by an implicit method, which involves the resolution of a linear system based on the jacobian matrix. This can considerably slow down the computations and even render the simulations impractical, mainly for memory storage considerations.

However, the distribution of clusters is often sparse, which means some concentrations can be considered as zero with a good approximation. This is one of the motivations to use a stochastic method: as only data related to clusters that exist in the simulation box are stored in memory, and since the inversion of a linear system is no longer necessary, we can expect to solve the memory problems. Contrary to biochemical systems for example, where stochastic methods have proved necessary to obtain reliable results [10], we do not expect effects due to the stochastic nature of the processes, mainly because the elementary species that contribute most to the kinetics are present in large numbers.

The equations of cluster dynamics (17) are coupled to each other by mobile clusters, which are supposed to be small clusters. Each stochastic reaction will affect at least one cluster in M, modifying the corresponding variable. It is then natural to choose the decoupling zone D in such a way as to include at least all these small and active clusters, in order to smooth out part of their fluctuations. Let M be a rectangular region  $[-m_V, m_I] \times [0, m_S]$  as in Figure 1 (in general, if the set of mobile clusters is bounded it suffices to take the smallest rectangle that

contains it). Correspondingly, we define  $D = [-D_V, D_I] \times [0, D_S]$ for the decoupling region and  $\Delta = [-\Delta_V, \Delta_I] \times [0, \Delta_S]$  for the deterministic one (all these sets do not contain (0,0), but we will omit this in the notation for clarity). Consider a non-mobile cluster made of  $n > m_I$  interstitial atoms and  $p > m_S$  solute atoms, corresponding to  $X_{(n,p)}$ . If  $n \le \Delta_I - \max(m_V, m_I)$  and  $p \le \Delta_S - m_S$  then no stochastic reaction can affect  $X_{(n,p)}$  (the same reasoning holds symmetrically for vacancy-type clusters). Therefore, any choice of D such that  $m_I < D_I \le \Delta_I - \max(m_V, m_I)$ ,  $m_V < D_V \le \Delta_V - \max(m_V, m_I)$  and  $m_S < D_S \le \Delta_S$  will have the same effect on the algorithm. To highlight the advantages of such a choice, it is useful to specify more precisely the rate equations for clusters in D under the same form as Eq. (16):

$$\frac{\mathrm{d}}{\mathrm{d}t}X_i = \sum_{j\in\mathcal{D}} v_{ij}a_j \left(X_\Delta(t)\right) + \rho_i^+ \qquad i\in D\setminus M \qquad (21)$$

$$\frac{\mathrm{d}}{\mathrm{d}t}X_i = \sum_{j\in\mathcal{D}} v_{ij}a_j \left(X_\Delta(t)\right) + \rho_i^+ - \rho_i^- X_i \qquad i\in M.$$
(22)

For  $i \in D \setminus M$ ,  $\rho_i^+$  is simply equal to  $G_i$ ; although in previous sections, this term was formally included in the reaction terms, we write it here separately for more clarity, restricting reaction terms to fluxes between clusters. It should be noted that, as a consequence of the choice of D, Eq. (21) is the full evolution equation for clusters in  $D \setminus M$  and requires no additional stochastic contribution. This is not the case for  $i \in M$ , where stochastic contributions are in general very large since these clusters are coupled to all existing clusters. However, these stochastic contributions are all contained in the source term  $\rho_i$ , which can be written in terms of two quantities:

$$\rho_{i}^{+} = G_{i} + \sum_{l} \alpha_{l+i,i} X_{l+i} + k_{i} D_{i} X_{i}^{e}$$
(23)

$$\rho_i^- = \sum_l \beta_{l,i} X_l + k_i D_i, \tag{24}$$

where the summation is performed on all indices that are not taken into account through reactions in  $\mathcal{D}$ , i.e. on l such that  $l \in \Lambda \setminus \Delta$  and/or  $l + i \in \Lambda \setminus \Delta$ . In practice, these source terms are updated every  $n_{\rho} = 1000$  reactions. In order to avoid the entire recomputation of these source terms, it is useful to keep track of the changes induced by the stochastic reactions by updating a copy of  $\rho_i^+$  and  $\rho_i^-$  at every stochastic reaction.

### 4.2. Iron under irradiation

As a first application we consider the cluster dynamics description of pure  $\alpha$ -iron under irradiation by neutrons. The set  $\Lambda$ labeling the species is the set of non-zero integer numbers (no solute is present), and the deterministic and decoupling regions are punctured intervals in  $\mathbb{Z} \setminus \{0\}$ . The dynamical variables  $X_n$  will represent the number of clusters of type n, but will nonetheless be treated as continuous quantities in the deterministic sector  $\Delta$ . Vacancy clusters are mobile up to size  $m_V = 4$  and interstitial clusters up to  $m_I = 3$ . The creation rate is adapted to describe cascades and is non zero for monomers ( $n = \pm 1$ ), clusters of 4 interstitial atoms and of 8 vacancies. The total irradiation time reached is around 10<sup>6</sup> s ( $\approx 11.6$  days) and the damage



Figure 2: (Color online) Cluster distribution at  $t = 4 \cdot 10^5$  s for pure iron at 300°C (simulation volume is  $V = 10^{-13}$  cm<sup>3</sup>). Red crosses (a) correspond to the hybrid algorithm; pink squares (b) are histograms of width 1000, obtained from the same data-set; the dashed blue line (c) is obtained by a purely deterministic method. The inset shows a close-up for small clusters.

rate is  $3.9 \times 10^{-8}$  dpa/s (dpa: displacements per atom). Other parameters of the model are taken from Reference [22].

The choice of  $\Delta$  is found to have no influence on the physical result, so it should only be driven by performance considerations. A minimal requirement for  $\Delta$  is that it should include all species produced by cascades, since irradiation is a potentially fast reaction channel. Apart from that, the choice of  $\Delta_V$  and  $\Delta_I$  has to be adapted to the characteristics of the system. In the  $\alpha$ -iron case we expect that the distribution of  $X_n$  becomes hollow in the vacancy sector at large times, at least for sufficiently high temperatures due to thermal dissociation of clusters, with one peak at small sizes (due to the creation by irradiation) and one at large negative values of *n*. The best choice for  $\Delta_V$  then probably lies in the middle. In the following we will report performance results for two different choices of the  $\Delta$  region.

Figure 2 shows the distribution of clusters  $C_n \equiv X_n/V$  at  $t = 4 \cdot 10^5$  s and  $T = 300^{\circ}$ C. Results obtained both by our hybrid algorithm in a single run and by a purely deterministic method are presented. The simulation volume is  $V = 10^{-13}$  cm<sup>3</sup>; clearly, all points in the stochastic region outside  $\Delta$  lie on or above the threshold 1/V. In order to obtain points more easily comparable with the exact solution, we consider binned data (see figure caption). The shape of the distribution is correctly reproduced, both in the stochastic and deterministic regions, the sole deviations being close to the interface. The distribution of small clusters is indistinguishable from the expected one on the scale of the figure: the effective corrections (15) are successful, even if their values are not updated at every iteration.

Figure 3 shows the CPU time as a function of the physical time reached (simulations were carried out on an Intel Xeon X5650 2.66 GHz). We report the performances for a purely deterministic method, a standard SSA algorithm and our hybrid algorithm. The SSA simulation was done with a volume  $V = 10^{-13}$  cm<sup>3</sup>. The performances obtained by the hybrid algorithm are shown for two different simulation volumes  $(V = 10^{-14} \text{ cm}^3, 10^{-13} \text{ cm}^3)$  and two choices of the determinis-



Figure 3: (Color online) CPU time versus physical time (pure iron at 300°C) for different methods and several choices of the parameters ( $\Delta_1 = [-300, 300]$ ,  $\Delta_2 = [-450, 600]$ ,  $V_1 = 10^{-14}$  cm<sup>3</sup>,  $V_2 = 10^{-13}$  cm<sup>3</sup>). Pink triangles correspond to the purely deterministic method; the dot-dashed green line to the SSA; the remaining four lines are obtained by the hybrid algorithm, the two dashed blue lines corresponding to the wider  $\Delta$  and the two solid red lines to the narrower.

tic interval ( $\Delta = [-300, 300]$ , [-450, 600]). The deterministic curve was obtained by first running a hybrid simulation up to  $t = 10^6$  s and noting the size of the largest populated interstitial and vacancy clusters (let us call them I(t) and V(t) respectively) at several intermediate times  $t_i$ . Afterwards, purely deterministic solutions were calculated at said times, for the system constituted of the  $N(t_i) = I(t_i) + V(t_i)$  equations describing interstitial clusters not larger than  $I(t_i)$  and vacancy clusters not larger than  $V(t_i)$ . One of the advantages of our hybrid approach is that the number of stochastic species is not fixed before the simulation, and the number of degrees of freedom involved in the computation automatically follows the growing number of populated species. As the plot reveals, the hybrid algorithm outperforms both the purely deterministic solving and the standard SSA solving, and has the additional advantage of being highly tunable.

Changing V and  $\Delta$  has a big impact on the algorithm's behavior. Increasing V has the direct consequence of decreasing the threshold 1/V and thus increasing the number of clusters participating in stochastic reactions: the total stochastic propensity  $A_S$  increases approximately linearly with V (see Figure 4); as a consequence, the CPU time it takes to reach a certain physical time increases with the volume. Decreasing the width of the deterministic interval  $\Delta$  has a positive effect in the short term, since it decreases the strain on the deterministic solver, but it affects negatively the long-time behavior, for the total propensity stays larger. The mutual influence of  $\Delta$  and V on the performance is a model-dependent aspect that has to be considered carefully when trying to find the best trade-off between speed and precision.

It is important to note that coping with the growth of N(t) is the main weakness of purely deterministic methods. If the number of equations is fixed, then a deterministic solution of the equations is asymptotically faster than the SSA. Indeed, for our problems, increasingly large time steps are taken by the solver in the fully deterministic scheme, whereas in the hybrid



Figure 4: (Color online) Evolution of the total stochastic propensity (in s<sup>-1</sup>) normalized by the volume (in cm<sup>3</sup>) for pure iron at 300°C. The two red curves on top correspond to  $\Delta_1 = [-300, 300]$ , the two blue curves underneath to  $\Delta_2 = [-450, 600]$ . In each couple, the fluctuating dashed curves correspond to  $V = 10^{-15}$  cm<sup>3</sup> and the smoother ones to  $V = 10^{-13}$  cm<sup>3</sup> (The data for  $V = 10^{-14}$  cm<sup>3</sup> is not shown here for clarity, since points collapse very tightly onto the curves corresponding to the higher volume).

one the time step is constrained by the root finding procedure, i.e. by the firing of stochastic reactions. However, since in the Newton iteration the matrix to be inverted is considered as dense and a direct linear solver is used, the CPU time scales as  $N^3$ , and the memory storage as  $N^2$ . The memory storage issue can even render calculations unfeasible. In the hybrid case, N should be replaced by the size of the deterministic region  $(\Delta_I + \Delta_V + 1) \times (\Delta_S + 1)$ , which is far smaller than  $N^2$ . This is why if N(t) is an increasing function of time, the hybrid algorithm becomes better in the long run from the point of view of CPU time and memory storage. In the present case, N(t) grows sublinearly in t. A fit of the form  $N(t) \propto t^{\alpha}$  gives  $\alpha \approx 0.35$ . Systems where N(t) grows faster than this will benefit even more from the hybrid approach (see Section 4.3).

As we already mentioned, raising the temperature yields peaked distributions. The dynamic implementation of the stochastic part of our algorithm (see Section 3.4) permits to take advantage of these situations by neglecting the degrees of freedom corresponding to the hollow regions. An example of such a case is shown in Figure 5, where a second peak is present, corresponding to very large vacancy-type clusters. Again, the position and shape of the peak (as well as the distribution at small sizes) are well reproduced.

### 4.3. Iron/helium under irradiation

The case of helium in iron is a particulary interesting one, from both physical and numerical point of views. From the physical side, helium is created by transmutation during irradiation with neutrons. As a noble gas, its interaction with iron atoms is repulsive and it thus tends to agglomerate in vacancies and vacancy clusters, forming helium bubbles [23]. Due to the helium pressure inside these bubbles, the emission rate of vacancies is reduced as the helium-to-vacancy ratio increases. At high temperature, only bubbles containing enough helium remain stable, which can lead to sparsely populated distributions. This is a case where the SSA and hybrid simulations should perform



Figure 5: (Color online) Cluster distribution at  $t = 4 \cdot 10^5$  s for pure iron at 400°C (simulation volume is  $V = 10^{-13}$  cm<sup>3</sup>). See the caption to Figure 2 for an explanation of the symbols.

particularly well. Very recently, SSA calculations have been performed in iron containing both helium and hydrogen [24].

For clusters containing no helium, the parametrization is the same as in the previous section, except the source term which corresponds to ion irradiation [22]. The damage rate is  $10^{-4}$  dpa/s and the helium-to-dpa ratio is  $10^{-4}$ . We intend to model irradiation conditions close to those obtained by dual ion beam experiments, where one beam is used to create damage (vacancies and self-interstitials), while the other one injects helium into the sample. Compared to irradiations performed with neutrons, where helium is created by transmutations, the helium content can be significantly higher in dual beam irradiations, so a role of helium can be more clearly highlighted.

Other input data of our model, specific to helium, include the diffusion coefficient of helium, deduced from first principle calculations [25], and binding energies of helium atoms, vacancies and interstitials to bubbles. These data are taken from first principle calculations for small clusters [26] and extrapolated from molecular dynamics calculations for larger sizes [27]. For the sake of simplicity, interstitial clusters are assumed to contain no helium atoms. In the following we fix the volume to  $V = 10^{-13}$  cm<sup>3</sup> and the deterministic region to  $\Delta = [-10, 10] \times [0, 2]$ . The temperature is set to 500°C.

Adding a solute atom to the picture heavily affects the computational cost of a completely deterministic approach, since the space of cluster types is now two-dimensional. This has an impact both on the speed of the simulations and on the memory requirements. The latter are especially important in the present case, where 1.3 GB of RAM are needed in order to reach t = 3 s in a fully deterministic computation. Clearly, reaching much longer times with such an approach is unfeasible. On the other hand, the memory requirements of the hybrid algorithm are much less demanding (less than 10 MB for all the times considered), since the deterministic region is fixed and the number of variables in the stochastic region follows the shape of the distribution, as explained above.

This distribution is shown at t = 1000 s in Fig. 6, where



Figure 6: (Color online) Distribution of clusters at t = 1000 s obtained by the hybrid algorithm in the iron/helium case, at 500°C. Data have been averaged on 50 simulations and smoothed by taking histograms: each point in the mesh represents the average on a region of size  $1000 \times 20$  around it. The inset shows the distribution of interstitial clusters, containing no helium, obtained by taking histograms of size 2000. Concentrations  $C_{(n,p)}$  are in cm<sup>-3</sup>.

data have been averaged on 50 simulations and grouped into histograms. Only a small part of the set of clusters is populated, which supports the use of a dynamical structure to store concentrations, neglecting non-populated classes. If such an approach were not considered, the memory needed at t = 1000 s would amount to around 3 GB instead of the few MB previously mentioned.

Concerning the computation time, the hybrid method is shown to be faster than standard SSA by more than one order of magnitude for small times and roughly three times faster at t = 500 s. Deterministic calculations could only be performed until 3 s of physical time due to memory requirements and are considerably slower than SSA and hybrid calculations. The number of equations N(t) corresponding to all species inside the rectangle containing the set of clusters grows as  $N(t) \sim t^{\alpha}$ , with  $\alpha = 1.6$ , which is much faster than the strongly sublinear growth observed for pure iron. This is why the deterministic solving fails so early. However, it should be noted that using a dense solver is clearly not optimal in the present case, since the jacobian matrix is sparse. Using a sparse direct solver, a deterministic algorithm could prove to be convenient. Since any improvement on the linear solver will be reflected in the hybrid algorithm, in general the hybrid algorithm should perform at least as well as the deterministic one, while possibly relieving technical difficulties associated with the solving of large matrices.

## 5. Conclusions

We proposed a new hybrid method for dealing with large sets of stiff rate equations, coupled by a small set of degrees of freedom, as those appearing in cluster dynamics. In this approach deterministic and stochastic strategies work in parallel across the time evolution. Essentially, a fixed small number of fast degrees of freedom is treated deterministically, while a non-constant (and possibly very high) number of the remaining variables is treated stochastically. The coupling between the



Figure 7: (Color online) CPU time versus physical time (iron/helium case). Pink triangles correspond to the purely deterministic method, the dashed green line to the purely stochastic (SSA) method, and the blue solid line to the hybrid algorithm.

two sectors is realized by the introduction of effective source terms for a small number of equations, where fluctuations due to the stochastic reactions are averaged out. This approximation does not yield appreciable errors in the results, and is helpful in reducing the time needed for the simulations.

We tested the hybrid method in the cluster dynamics description of two physical phenomena: the formation of voids and interstitial loops in pure iron under irradiation, and the nucleation and growth of helium bubbles in the same system. In the two cases, a significant decrease in computation time is observed with respect to both deterministic solving and standard SSA algorithm. Since the number of variables treated deterministically in the hybrid algorithm is far lower than in a fully deterministic solving, the memory storage is also orders of magnitude lower, and comparable to SSA. We emphasize that we have used a dense solver for the deterministic and hybrid methods; using a sparse solver would improve the performance of these two methods with respect to SSA. In particular, a larger number of deterministic variables could be used in the hybrid method, to keep a balance between the solving of the deterministic and stochastic parts.

### Acknowledgments

This work was partially funded by the French National Research Agency (ANR) through the ParMat Projet ANR-06-CIS6-006-03. We are much indebted to V. V. Bulatov for suggesting the use of SSA to solve the physical problems mentioned in the present article.

#### References

- [1] T. Jourdan, F. Soisson, E. Clouet, A. Barbu, Acta Mater. 58 (2010) 3400.
- [2] M. Hayns, J. Nucl. Mater. 56 (1975) 267–274.
- [3] N. Ghoniem, D. Cho, Phys. Status Solidi A 54 (1979) 171.
- [4] A. Hardouin-Duparc, C. Moingeon, N. Smetniansky-de-Grande, A. Barbu, J. Nucl. Mater. 302 (2002) 143.

- [5] C. J. Ortiz, M. J. Caturla, C. C. Fu, F. Willaime, Phys. Rev. B 80 (2009) 134109.
- [6] N. Ghoniem, Radiat. Eff. Defects Solids 148 (1999) 269.
- [7] S. Golubov, R. Stoller, S. Zinkle, A. Ovcharenko, J. Nucl. Mater. 261 (2007) 149.
- [8] M. Baskes, R. Fastenau, P. Penning, L. Caspers, A. van Veen, J. Nucl. Mater. 102 (1981) 235.
- [9] D. T. Gillespie, J. Comput. Phys. 22 (1976) 403-434.
- [10] R. Srivastava, L. You, J. Summers, J. Yin, J. Theor. Biol. 218 (2002) 309.
- [11] D. T. Gillespie, J. Chem. Phys. 115 (2001) 1716.
- [12] E. Haseltine, J. Rawlings, J. Chem. Phys. 117 (2002) 6959.
- [13] T. Jahnke, Multiscale Modeling & Simulation 9 (2011) 1646–1676.
- [14] A. Alfonsi, E. Cancès, G. Turinici, B. D. Ventura, W. Huisinga, ESAIM: Proc. 14 (2005) 1–13.
- [15] D. T. Gillespie, Physica A 188 (1992) 404-425.
- [16] M. A. Gibson, J. Bruck, J. Phys. Chem. 104 (2000) 1876–1889.
- [17] C. V. Rao, A. P. Arkin, J. Phys. Chem. 118 (2003) 4999-5010.
- [18] D. T. Gillespie, J. Chem. Phys. 113 (2000) 297-306.
- [19] M. Surh, J. Sturgeon, W. Wolfer, J. Nucl. Mater. 378 (2008) 86.
- [20] D. T. Gillespie, J. Phys. Chem. 81 (1977) 2340–2361.
- [21] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, C. S. Woodward, ACM Trans. Math. Softw. 31 (2005) 363–396.
- [22] E. Meslin, A. Barbu, L. Boulanger, B. Radiguet, P. Pareige, K. Arakawa, C. C. Fu, J. Nucl. Mater. 382 (2008) 190.
- [23] H. Trinkaus, B. N. Singh, J. Nucl. Mater. 323 (2003) 229.
- [24] J. Marian, V. V. Bulatov, J. Nucl. Mater. 415 (2011) 84.
- [25] C. C. Fu, F. Willaime, Phys. Rev. B 72 (2005) 064117.
- [26] C. C. Fu, F. Willaime, J. Nucl. Mater. 367-370 (2007) 244.
- [27] G. Lucas, R. Schäublin, J. Nucl. Mater. 386-388 (2009) 360.