Atomistic modelling of diffusional phase transformations with elastic strain

2004 J. Phys.: Condens. Matter 16 S2679

(http://iopscience.iop.org/0953-8984/16/27/007)

INSTITUTE OF PHYSICS PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **16** (2004) S2679–S2697 PII: S0953-8984(04)71492-7

Atomistic modelling of diffusional phase transformations with elastic strain

D R Mason[1], **R E Rudd**[2] **and A P Sutton**[1,3]

[1] Department of Materials, Oxford University, OX1 3PH, UK
[2] Lawrence Livermore National Laboratory, 7000 East Avenue, L-045, Livermore, CA 94550, USA
[3] Helsinki University of Technology, Laboratory of Computational Engineering, PO Box 9203, FIN 02015 HUT, Finland

E-mail: daniel.mason@materials.ox.ac.uk

Received 6 November 2003
Published 25 June 2004
Online at stacks.iop.org/JPhysCM/16/S2679
doi:10.1088/0953-8984/16/27/007

Abstract

Phase transformations in 2xxx series aluminium alloys (Al–Cu–Mg) are investigated with an off-lattice atomistic kinetic Monte Carlo simulation incorporating the effects of strain around misfitting atoms and vacancies. Atomic interactions are modelled by Finnis–Sinclair potentials constructed for these simulations. Vacancy diffusion is modelled by comparing the energies of trial states, where the system is partially relaxed for each trial state. No special requirements are made about the description of atomic interactions, making our approach suitable for more fundamentally based models such as tight binding if sufficient computational resources are available. Only a limited precision is required for the energy of each trial state, determined by the value of $k_B T$. Since the change in the relaxation displacement field caused by a vacancy hop decays as $1/r^3$, it is sufficient to determine the next move by relaxing only those atoms in a sphere of finite radius centred on the moving vacancy. However, once the next move has been selected, the entire system is relaxed.

Simulations of the early stages of phase separation in Al–Cu with elastic relaxation show an enhanced rate of clustering compared to those performed on the same system with a rigid lattice.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

The mechanical properties of commercial aluminium alloys are manipulated by the process of ageing the supersaturated solid solution. In the Al–4 wt% Cu system, copper atoms cluster together on $\{200\}_{Al}$ host planes to form metastable particles known as Guinier–Preston

(GP) zones. These particles are copper-enriched discs just a few atoms thick, which are coherent with the matrix, and up to 50 nm in diameter [1–3]. On further ageing the GP zones undergo a change of structure before transforming eventually into the thermodynamically stable θ phase $CuAl_2$, which is incoherent with the host. There is a corresponding peak in the flow stress, where dislocations switch from cutting coherent GP zones to forming loops around incoherent θ phase particles. In the 2xxx series of alloys, magnesium is added and the precipitation sequence changes. First coherent GPB zones form as $\langle 100 \rangle$ rods, which ripen into laths of the orthorhombic S phase Al_2CuMg [4–8]. The precipitates formed can be changed again with the addition of a fourth alloying element. The addition of Si leads to a finer precipitate dispersion, while adding Ag can produce Mg–Ag clusters in the early stages of decomposition leading to the Ω phase forming on $\{111\}_{Al}$ planes.

While the observed structures of GP zones are well established [9, 10], and the effects of the addition of certain trace elements have been studied experimentally [11–13], modelling the effect of trace elements on kinetics is a difficult problem. Their influence may stem from the formation of complexes with vacancies, which will reduce the rate of diffusion and of phase separation. It is also possible that large trace element atoms may be attracted to small solute atoms by the concomitant reduction in strain energy of the system, and in this way promote the formation of a co-cluster. Apart from reductions in the total strain energy of the system through such an association there may also be much shorter-range gains in bond energies. It follows that the influence of trace elements is inherently atomistic in nature, in that it stems from discrete atomic interactions with point defects such as vacancies and other solute atoms. However, misfitting atoms and vacancies interact over large distances compared with a bond length through their elastic fields. It is the requirement to treat both discrete atomic interactions and long-range elastic fields on an equal footing in dilute alloys that makes the modelling of the influence of trace elements so challenging. The purpose of this paper is to present a methodology to meet this challenge. Since our methodology treats atomic vacancies explicitly we are able to include the possibility that vacancies may become trapped at interfaces surrounding second-phase particles, which may lead to Brownian motion of particles as described in [14–16].

A vacancy is a centre of tensile dilation, and so it will be attracted to regions of compressive stress. The pattern of compressive stress will follow the underlying local composition of the system, enhancing the time spent by a vacancy in some areas, while reducing it in others. If the pattern of compressive stress is enhanced along particular crystallographic directions, owing to elastic anisotropy, the increased probability of finding a vacancy in these regions may be expected to influence the mobility of solute atoms there, which in turn may influence the morphology of the second-phase particle.

Eshelby's sharp interface construction [17, 18] may go some way to predicting equilibrium morphologies of the emerging phase from the elastic constants and lattice misfit, but a general analytic form may not be found for anything other than the simplest geometric shapes, requiring numerical solutions of the remaining boundary integrals [19, 20]. Khachaturyan's formulation [21–23] for the elastic energy has recently been extended to model elastic inhomogeneity as a perturbation expansion for the anisotropic elastic case [24], but it may still have difficulty addressing the motion of vacancies, and their localized interaction with very small concentrations of trace elements.

As we will show, finding the displacement field to sufficient accuracy around a vacancy in an inhomogeneous system necessitates going beyond a harmonic lattice model. This forces us to perform off-lattice Monte Carlo simulations, which is a much more expensive calculation. But there are also considerable benefits, such as the ability to model the breakdown of interfacial coherency, and changes in crystal structure of second-phase regions. To the best

of our knowledge this is the first presentation of an atomistic simulation of a bulk diffusional transformation which does not preserve a reference lattice.

There are three principal steps we have taken to enable us to model diffusional processes atomistically with elastic interactions. Firstly we have improved the kinetic Monte Carlo algorithm used to select each vacancy hop, by developing an efficient combination of the stochastic first and second order residence algorithms which outperforms the n-fold way for this problem. This is described in detail elsewhere [25]. Secondly we have developed an iterative scheme based on the Lanczos method to tridiagonalize a matrix, which appears particularly suited to relaxing the system following a vacancy–atom exchange. We believe that this is the first application of the Lanczos algorithm to an elastic relaxation problem. Finally, although the relaxation propagates to the boundaries of the system, we have found that the majority of the *change* of the elastic energy due to a vacancy–atom exchange can be recovered by relaxing only those atoms in a near field region.

The purpose of this paper is to demonstrate the importance of considering elastic stress effects when performing kinetic Monte Carlo simulations. We show that the harmonic lattice approximation is insufficient for accurately predicting the difference in energy between trial Monte Carlo states. However, an approximation based on relaxing those atoms within a small sphere centred on the vacancy is enough to determine the next move, provided that eventually the whole system is relaxed after each move. Simulations of diffusion performed in the aluminium–copper–magnesium ternary alloy system show that for a fixed number of vacancy–atom exchanges, larger clusters grow on a flexible lattice owing to the reduction in elastic strain energy this can provide. However, it is also found that since vacancies are preferentially attracted to compressive regions, they can become trapped within emerging magnesium-rich regions, and this can have the opposite effect of slowing the growth rate of second-phase particles.

2. Computational method

2.1. Kinetic Monte Carlo simulations

Kinetic Monte Carlo simulations [26–28] enable us to associate a meaningful time with simulated stochastic events. In our example of vacancy diffusion, adjacent atomic configurations are linked by a single vacancy–atom exchange event. If there are V vacancies in the system and the vacancy is in a region of crystal with a local coordination number of z (in fcc crystals $z = 12$), there are Vz accessible neighbouring configurations[4]. We have developed [25] a new algorithm that combines the n-fold way with a stochastic version of the second order residence time algorithm to move between configurations. This new algorithm is more efficient than either the n-fold way or the second order algorithm.

The probability that any one of these events will be the next to occur is proportional to the rate at which the event occurs. The rate appropriate for the movement of a point defect is given by Vineyard's equations [29]. While the terms in this expression could be calculated exactly, in practice it is difficult to establish converged saddle point energies 'on the fly'. Instead we have used Flynn's approximation derived from dynamical theory [30] to determine a rate for the transition based on the difference in internal energy between states. The rate of migration from state i to state j passing through a saddle point s is given by

$$r_{i\to j} = \left(\frac{3}{5}\right)^{\frac{1}{2}} \nu_{\mathrm{D}} \exp\left(-\frac{\langle c\rangle \Omega \Delta^2}{k_{\mathrm{B}}T}\right) \exp\left(-\frac{E_j - E_i}{2k_{\mathrm{B}}T}\right) \tag{2.1}$$

[4] This is an upper limit as we exclude vacancy–vacancy exchanges.

where E_i is the energy of the system in state i when the system is fully relaxed, ν_D is the Debye frequency, and Ω the volume of the migrating atom. The parameter Δ is the displacement, expressed as a fraction of the bond length, at which the force on the migrating atom is a maximum, and is taken to be the constant value 0.316. $1/\langle c\rangle$ is an effective elastic compliance. For fcc crystals it is given by

$$\frac{1}{\langle c\rangle} \approx \frac{2}{15}\left(\frac{3}{c_{11}} + \frac{2}{c_{11}-c_{12}} + \frac{1}{c_{44}}\right). \tag{2.2}$$

Only $\frac{1}{2}\langle 110\rangle$ hops are considered. Saddle points for exchanges with the next-nearest neighbour shell of atoms are around 2 eV higher in energy and so are assumed so infrequent as not to affect the kinetics.

2.2. Interatomic potentials

We have constructed a set of relatively short-ranged Finnis–Sinclair potentials [31] for this study. They have been parameterized to reproduce the elastic constants, cohesive energy and lattice parameters of the pure metals, and we use a simple interpolation scheme for the parameters of the potentials for interactions between dissimilar atoms. Our potentials have been described fully in [25].

Our potentials will not reproduce correctly the subtle energy differences between intermetallic structures, and hence they will not reproduce the phase diagram, which is very complex. Nevertheless, they suit the present need by providing a means of estimating relaxation forces and energies for arbitrary alloy configurations, while reproducing the elastic properties of the pure metals. Monte Carlo simulations of annealing at zero external stress found that our potentials favour intermetallic phases at the compositions Al_3Cu, AlCu and Cu_3Al with crystal structures $L1_2$, $L1_0$ and $L1_2$, respectively.

2.3. The Lanczos method for relaxation

An efficient algorithm for finding relaxed states is required because there may be many millions of configurations generated during a kinetic Monte Carlo simulation of the early stages of diffusional phase separation by a vacancy mechanism. Consider first the harmonic lattice approximation. The energy may be expanded as a Taylor series about the current configuration which may be represented by a vector $\mathbf{x}_0$ in $3N$-dimensional space describing the positions of each of the N atoms in the system:

$$E(\mathbf{x}_0+\mathbf{u}) = E(\mathbf{x}_0) + \left.\frac{\partial E}{\partial \mathbf{x}}\right|_{\mathbf{x}_0}\cdot\mathbf{u} + \frac{1}{2}\mathbf{u}\cdot\left.\frac{\partial^2 E}{\partial\mathbf{x}\partial\mathbf{x}}\right|_{\mathbf{x}_0}\mathbf{u}. \tag{2.3}$$

The $3N$-dimensional vector $\mathbf{u}$ is a displacement from the current configuration $\mathbf{x}_0$, which is assumed to be small in magnitude. We identify the first derivative of the energy as minus the $3N$-dimensional force $\mathbf{f} = -\left.\frac{\partial E}{\partial \mathbf{x}}\right|_{\mathbf{x}_0}$, and the second as the matrix of force constants $\mathbf{D} = \left.\frac{\partial^2 E}{\partial\mathbf{x}\partial\mathbf{x}}\right|_{\mathbf{x}_0}$. These two derivatives are evaluated analytically at the current configuration $\mathbf{x}_0$ using the interatomic potentials.

Equation (2.3) has a minimum with respect to $\mathbf{u}$ given by

$$\begin{aligned}\mathbf{f} &= \mathbf{D}\mathbf{u}\\ \mathbf{u} &= \mathbf{D}^{-1}\mathbf{f}.\end{aligned} \tag{2.4}$$

The change in energy brought about by making this displacement from the current configuration is obtained by substituting equation (2.4) into (2.3):

$$\Delta E = -\tfrac{1}{2}\mathbf{f}\cdot\mathbf{D}^{-1}\mathbf{f}. \tag{2.5}$$

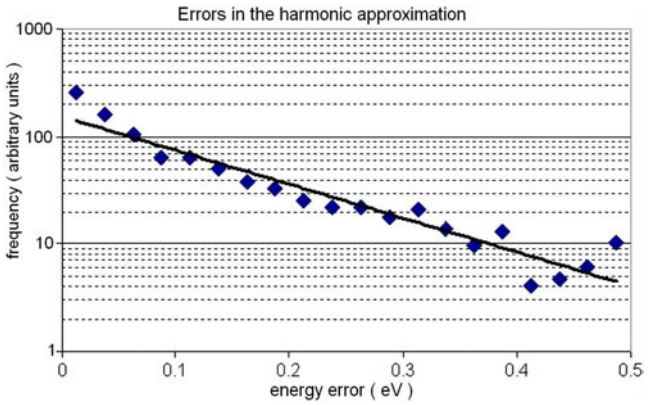

Figure 1. Errors in the estimation of the energy difference between neighbouring configurations using a harmonic lattice model. The vertical axis is measured in arbitrary units. The line shows a decaying exponential fit to the data and is present to guide the eye. A substantial fraction of the estimates of energy change (6%) is in error by more than 0.1 eV, and so of the same size as the total energy difference between states.

For the aluminium–copper alloys of interest to this study the displacements of atoms near a vacancy are relatively large and sensitive to the local atomic structure, and the harmonic approximation is unable to reproduce the energy differences between neighbouring states with sufficient accuracy. The error in the estimate of the energy difference between adjacent atomic configurations can be judged by comparing the energy change predicted by the harmonic approximation with the energy change if both before and after configurations are allowed to relax completely using an anharmonic relaxation algorithm. These errors are plotted in figure 1. 100 trial configurations were constructed for atomic concentrations of Cu in Al of 2% to 20% in 2% increments. The copper atoms were randomly positioned, and a single vacancy considered. Each of the twelve accessible vacancy–atom exchanges was performed at each configuration.

The difference in energy between harmonically relaxed states is in error by the order of 0.1 eV. Since this is the same order as the total energy difference between states, counting both chemical bond energy changes and elastic energy changes, the harmonic relaxation approximation is seen to be unsuitable for this study.

An anharmonic calculation could be done in a number of ways, for example using damped molecular dynamics or energy minimization with conjugate gradients. In this way atoms may be displaced considerable distances from their ideal lattice positions, and changes in the state of coherency of interfaces surrounding second-phase particles may be modelled, as may local changes in crystal structure, provided the relaxation algorithm explores the local configuration phase space sufficiently. However, it is considerably more expensive computationally than the simple harmonic approximation.

We have found that significant computational gains may be made if an algorithm derived from the Lanczos method of tridiagonalizing a matrix [32] is employed. The strategy is to find an approximate solution to equation (2.4) and displace atoms through this vector. We will denote an approximate value with a tilde ($\tilde{}$). The force and force constant matrix elements are then recalculated and further displacements constructed until the system energy has converged to an acceptable tolerance. Restarting the calculation in this way not only prevents error accumulation, allowing a machine precision solution if desired, but also sidesteps the known problems associated with keeping the Lanczos basis vectors orthonormal.

Equation (2.4) is notoriously difficult to solve for $\mathbf{u}$. The force constant matrix is sparse, but its inverse is not, and it may well be too large to store. However, the Lanczos method allows us to recast this equation in a basis where the matrix $\mathbf{D}$ may be represented as a tridiagonal matrix. The tridiagonal matrix is constructed through repeated operations involving a multiplication by the matrix $\mathbf{D}$.

The Lanczos method is perhaps most familiar in solid state physics in the context of calculating local densities of electronic or vibrational states [33]. Our use of the method here is quite different: namely we are generating an approximate solution for $\mathbf{u}$ in equation (2.4) for which ΔE in equation (2.5) has converged. The method constructs a new orthonormal basis set of vectors. After m such basis vectors have been generated, the subspace they span is known as the Krylov subspace:

$$\mathcal{K}_m(\mathbf{D}, \mathbf{f}) = \{\mathbf{f}, \mathbf{D}\mathbf{f}, \ldots, \mathbf{D}^{m-1}\mathbf{f}\}. \tag{2.6}$$

The representation of $\mathbf{D}$ in this basis is the tridiagonal $m \times m$ matrix $\mathbf{T}_m$.

$$\text{define } \mathbf{T}_m: \qquad \mathbf{\Phi}_m^{\mathrm{T}}\mathbf{D}\mathbf{\Phi}_m \equiv \mathbf{T}_m = \begin{pmatrix} \alpha_0 & \beta_1 & & & \\ \beta_1 & \alpha_1 & \beta_2 & & \\ & \beta_2 & \alpha_2 & \ddots & \\ & & \ddots & \ddots & \beta_{m-1} \\ & & & \beta_{m-1} & \alpha_{m-1} \end{pmatrix} \tag{2.7}$$

where the columns of the $3N \times m$ matrix $\mathbf{\Phi}_m$ are the m generated basis vectors $\{\phi_0, \phi_1, \ldots, \phi_{m-1}\}$ each of dimension $3N$. This matrix spans the same subspace as $\mathcal{K}_m(\mathbf{D}, \mathbf{f})$.

The coefficients α and β and the basis vectors ϕ are found by the Lanczos algorithm, which is an iterative process [32]:

$$\begin{aligned} &\phi_{-1} = \mathbf{0} \\ &\phi_0 = \frac{1}{\|\mathbf{f}\|}\mathbf{f} \\ &\beta_{m+1}\phi_{m+1} = \mathbf{D}\phi_m - \alpha_m\phi_m - \beta_m\phi_{m-1} \qquad (2.8) \\ &\text{where} \qquad \alpha_m = \phi_m \cdot \mathbf{D}\phi_m \\ &\text{and} \qquad \beta_m = \begin{cases} \phi_{m-1} \cdot \mathbf{D}\phi_m, & m > 0, \\ \|\mathbf{f}\|, & m = 0. \end{cases} \qquad (2.9) \end{aligned}$$

Finding each successive pair of α and β requires a matrix–vector multiplication.

We may obtain approximate solutions for ΔE and $\mathbf{u}$ in the subspace $\mathcal{K}_m(\mathbf{D}, \mathbf{f})$. The predicted relaxation energy at this stage is

$$\Delta E_m = -\tfrac{1}{2}\|\mathbf{f}\|^2[\mathbf{T}_m^{-1}]_{00}. \tag{2.10}$$

$[\mathbf{T}_m^{-1}]_{00}$ is the element in the top left-hand corner of the inverse of the matrix $\mathbf{T}_m$. The predicted displacement at this level is

$$\mathbf{u}_m = \|\mathbf{f}\|\mathbf{\Phi}_m\mathbf{T}_m^{-1}\mathbf{e}_0 \tag{2.11}$$

where $\mathbf{e}_0$ is the m-dimensional unit vector $[1, 0, 0, \ldots, 0]^{\mathrm{T}}$. The iterative procedure of equation (2.8) can be halted when $\Delta\tilde{E}_m$ has converged, and then the displacement vector may be found.

We note that the displacement vector may be constructed directly from equation (2.8) as the basis vectors are constructed, in a manner analogous to that used by the method of conjugate gradients. The SYMMLQ algorithm [34] for instance computes the solution of $\mathbf{T}_m\mathbf{y}_m = \|\mathbf{f}\|\mathbf{e}_0$ progressively, using Gaussian elimination with partial pivoting. However, we

use an alternative method for calculating the displacement vector which does not rely on inverting the tridiagonal matrix, but instead is an iterative scheme using the change in energy at level m. Our method does require storing the basis vector set $\mathbf{\Phi}_m$, but we believe this extra cost to be offset by the improvement in convergence available by the method of section 2.4. We expand the displacement in the m generated basis vectors $\{\phi_0, \phi_1, \ldots, \phi_{m-1}\}$:

$$\mathbf{u}_m = \sum_{j=0}^{m-1} \gamma_j \phi_j. \tag{2.12}$$

The orthonormality condition of the basis vectors ϕ_m and equation (2.8) gives

$$\Delta E_m = -\tfrac{1}{2}\mathbf{u}_m \cdot \mathbf{f} = -\tfrac{1}{2}\beta_0\gamma_0 \qquad \text{so } \gamma_0 = -\frac{2\Delta E_m}{\beta_0} \tag{2.13}$$

$$\mathbf{u}_m \cdot \mathbf{D}\phi_0 = \mathbf{f} \cdot \phi_0 = \beta_0 = \sum_j \gamma_j \phi_j \cdot \mathbf{D}\phi_0 = \gamma_0\alpha_0 + \gamma_1\beta_1 \qquad \text{so } \gamma_1 = \frac{\beta_0 - \gamma_0\alpha_0}{\beta_1}. \tag{2.14}$$

The coefficients γ_k for $k > 1$ are found by the iterative procedure

$$k > 1: \qquad \mathbf{u}_m \cdot \mathbf{D}\phi_k = \mathbf{f} \cdot \phi_k = 0 = \sum_j \gamma_j \phi_j \cdot \mathbf{D}\phi_k = \gamma_k\alpha_k + \gamma_{k+1}\beta_{k+1} + \gamma_{k-1}\beta k$$

$$\text{so } \gamma_k = -\frac{\gamma_{k-2}\beta_{k-1} + \gamma_{k-1}\alpha_{k-1}}{\beta_k}. \tag{2.15}$$

As is seen in equation (2.8), the initial vector ϕ_0 is a unit vector in $3N$-dimensional space in the direction of the force $\mathbf{f}$. Each successive multiplication by $\mathbf{D}$ in equation (2.8) enables the relaxation to spread further from the centres where the forces in $\mathbf{f}$ are located. Since the change in the relaxation energy associated with a vacancy hop decays rapidly with distance from the vacancy, a rapidly convergent estimate of the relaxation energy may be obtained, together with a good approximation to $\mathbf{u}$.

Thus there is a sound physical basis for our choice of the Lanczos method to relax the system for each trial vacancy–atom exchange.

2.4. Improving the solution

We have found that the efficiency of the convergence of the Lanczos procedure can be improved by storing and reusing information generated during previous relaxations. We will show that the values of α and β found during the tridiagonalization process can improve the estimate of the energy change after only m levels have been performed. By substituting this improved energy into equations (2.13)–(2.15), the displacement vector is also improved.

The change in energy after m levels can be written as a continued fraction (see, for example, [33]). From equation (2.10) we write

$$\Delta E_m = \cfrac{-\frac{1}{2}\beta_0^2}{\alpha_0 - \cfrac{\beta_1^2}{\alpha_1 - \cfrac{\beta_2^2}{\vdots \cfrac{}{\alpha_{m-1} - \beta_m^2}}}}. \tag{2.16}$$

Going to another level $m + 1$ does not affect the previously found coefficients α and β, but can improve the energy estimate by adding to the end of the continued fraction. It would be possible to improve the estimate of the energy change if we could terminate the continued fraction at level m with something which better represents the tail of the fraction. It is found

numerically that averages a and b constructed from the coefficients α and β converge to constant values, but that even after thousands of levels the values for α and β rapidly oscillate about these averages. As a first choice, a quadratic terminator constructed from the averages could be used to replace the tail of the continued fraction at level m, by substituting β_m^2 in equation (2.16) with β_m^2/t, where

$$t = a - \frac{b^2}{t} = \frac{a}{2}\left(1 + \sqrt{1 - \left(\frac{2b}{a}\right)^2}\right) \tag{2.17}$$

$$a = \frac{1}{m}\sum_{j=0}^{m-1}\alpha_j, \qquad b = \frac{1}{m}\sum_{j=1}^{m}\beta_j. \tag{2.18}$$

The form for the quadratic terminator in equation (2.17) is valid because the discriminator is always positive.

While it is possible to ensure convergence in these averages by going to enough levels, we are looking for a displacement vector estimated from only the first few basis vectors. Too many levels of the procedure gives a well-converged solution only to the harmonic problem of equation (2.4). We find that $\Delta\tilde{E}_m$ has typically converged to 0.1% when $\mathbf{T}_m$ is constructed to level $m = 6$–12. Note that the numerical errors at this level remain small. This is why we have no need for renormalization of the basis set $\mathbf{\Phi}_m$. We are trying to find a displacement vector which will minimize the system energy on a weakly anharmonic potential energy surface. To converge to the true minimum we calculate the force and matrix elements at the current atom displacement, call the Lanczos procedure of (2.8) to generate a better solution, and move the atoms. Starting from an initial position $\mathbf{x}_0$, this produces a succession of estimates for the local energy minimum $\mathbf{x}_0$, $\mathbf{x}_0 + \tilde{\mathbf{u}}^{(0)}$, $\mathbf{x}_0 + \tilde{\mathbf{u}}^{(0)} + \tilde{\mathbf{u}}^{(1)}$, ... The displacement vectors are generated by

$$\begin{aligned} \mathbf{D}|_{\mathbf{x}_0}\tilde{\mathbf{u}}^{(0)} &\approx \mathbf{f}|_{\mathbf{x}_0} \\ \mathbf{D}|_{\mathbf{x}_0+\tilde{\mathbf{u}}^{(0)}}\tilde{\mathbf{u}}^{(1)} &\approx \mathbf{f}|_{\mathbf{x}_0+\tilde{\mathbf{u}}^{(0)}} \\ &\cdots. \end{aligned} \tag{2.19}$$

We are using the Newton–Raphson method to approach the solution, so the input forces $\mathbf{f}|_{\mathbf{x}_0}, \mathbf{f}|_{\mathbf{x}_0+\tilde{\mathbf{u}}^{(0)}}, \ldots$ diminish in magnitude quadratically. However, the input matrices $\mathbf{D}|_{\mathbf{x}_0}, \mathbf{D}|_{\mathbf{x}_0+\tilde{\mathbf{u}}^{(0)}}, \ldots$ change only slightly as the position of the atoms is updated, reflecting the weakly anharmonic nature of the potentials near the minimum position.

As a function of the number of levels m, information about the extreme eigenvalues converges faster than information about the eigenvalues in the middle of the spectrum [33]. Therefore, it is possible to estimate the limits of the eigenvalue spectrum from a finite number of values of α and β. Long-range elastic interactions are determined by the properties of the lower limit of the spectrum. The short-range displacement field around a vacancy is determined primarily by the first few levels of the continued fraction. By attaching a terminator to the continued fraction not only are we able to describe long-range elastic interactions but we also embed a finite cluster centred on the vacancy in an infinite medium characterized by average elastic properties. This embedding also influences the short-range displacement field around the vacancy.

The values of the coefficients α and β to level m' are stored, and they are overwritten by new values as they are generated until the criterion for convergence of the relaxation energy is satisfied. Subsequently we wish to calculate the change in the relaxation energy when a vacancy is exchanged with a neighbour. Convergence of the change in the elastic relaxation energy is attained with a smaller number, m, of exact levels. To calculate the change in the elastic relaxation we use the following continued fraction, where we note that the first m pairs of coefficients α and β are those computed for the displaced vacancy, the remaining m'–m

pairs of coefficients α and β are those stored and not overwritten, and t is the terminator of equation (2.17), averaged over all m' coefficients:

$$\Delta\tilde{E}_m = \cfrac{-\frac{1}{2}\beta_0^2}{\alpha_0 - \cfrac{\beta_1^2}{\cfrac{\vdots}{\alpha_{m-1} - \cfrac{\beta_m^2}{\alpha_m - \cfrac{\beta_{m+1}^2}{\cfrac{\vdots}{\alpha_{m'-1} - \cfrac{\beta_{m'}^2}{t}}}}}}}. \tag{2.20}$$

The improved estimate for the energy change calculated using (2.20) is substituted in equation (2.13) to get an improved estimate for the displacement vector. Note that since the values of α and β are exact to level m, the projection of this new estimate on the true solution to equation (2.4) is greater than that available previously, even though it is within the same Krylov subspace $\mathcal{K}_m(\mathbf{D}, \mathbf{f})$.

This algorithm can be efficiently coded in parallel. A domain decomposition is used to construct the elements of $\mathbf{D}$. The matrix elements do not then need to be broadcast to perform the matrix–vector multiply in equation (2.8).

To summarize our relaxation scheme:

- Calculate the force $\mathbf{f}$ and force-constant matrix $\mathbf{D}$.
- Use equation (2.8) to construct a tridiagonal representation of $\mathbf{D}$ to the mth level.
- Write over the stored values of α and β to the mth level.
- Use equation (2.20) to calculate the energy $\boldsymbol{\Delta}\tilde{\mathbf{E}}_m$ at level m, and to decide if this has converged satisfactorily.
- Use equations (2.13)–(2.15) to construct a displacement vector $\tilde{\mathbf{u}}$.
- As with all Newton–Raphson methods, ensure that the displacement $\tilde{\mathbf{u}}$ is genuinely a reduction in system energy, and perform a 1D line search for the energy minimum along the direction of $\tilde{\mathbf{u}}$ if not.
- Repeat these steps until the true local energy minimum has been found.

2.5. The short-range relaxation region

The displacement field around a point defect scales in magnitude as $1/r^2$, where r is the distance from the defect. Starting from a fully relaxed system of atoms, a single atom–vacancy exchange is performed to construct a new trial configuration. The additional displacement field introduced by this exchange will be dipolar, and so scale as $1/r^3$. This displacement is short-ranged, and so the greatest contribution to the elastic energy change due to an exchange will come from the minor corrections of the positions of the atoms close to the exchange event.

As we are undertaking a Monte Carlo simulation, it is necessary that the energies be accurate only to within a fraction of the thermal energy. At 300 K, this energy is 0.025 eV. If we ensure that the energy differences are accurate to about 10% of this value, the decision of the next state to visit will be made incorrectly about 1% of the time. We need an approximation method which can reliably make estimates of the energy difference between states to within about 1 meV.

A good approximation to the elastic energy change is found by pinning all atoms more than a fixed radius R distant from the original vacancy location, and only relaxing those within.

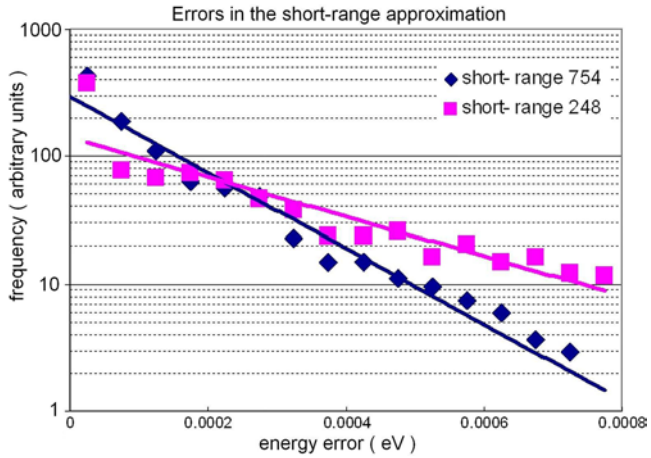

Figure 2. Errors introduced using the short-range relaxation region, with cut-off regions encapsulating 248 and 754 atom sites centred on the original vacancy position. The vertical axis shows the frequency of errors. The tests were performed over the same configurations as figure 1, so the frequency is at the same scale. Here it is seen that the probability of the difference in energy between relaxed states being in error by more than 0.001 eV becomes very small with the larger cut-off region.

If this is done, after relaxation only those atoms immediately exterior to the sphere will still have a finite net force acting on them. The error in the elastic energy will be due to the failure of the embedding external medium to respond to the change in shape of the interior and vice versa. This energy error will be positive, and scale as $1/R^3$. However, if a second virtual state is considered, it too will have an error of roughly the same size. All virtual states will therefore have a systematic component to the error due to the non-propagation of the force dipole to infinite range. An additional error component will be due to the coupling of the inhomogeneities in the far field and the force dipole. As larger near field regions are assumed, this error component will tend to zero. As this error must be kept small compared to the thermal energy, we have determined our cut-off range R to be three times the range of the potential, covering a sphere of 754 lattice sites plus the original vacancy position.

The systematic error component may be removed without affecting the statistical likelihood of selecting a given state. This error will appear in the free energy as a constant additional term, and so in the rate as a constant multiplier. If after the move the system is allowed to relax fully, then a comparison may be made between exact and approximate energies. This difference can be ascribed to the systematic error, and so the rates may be corrected. If the system is relaxed after every move, the energy of the previously visited configuration will always be stored exactly, so that the rates may be corrected even before the exchange is made.

The difference in energy between elastically relaxed states made using the short-range approximation is shown in figure 2. It can be seen that if a sufficiently large relaxation region is chosen the magnitude of the energy error can indeed be made sufficiently small. Note that our relaxation region is considerably larger than that used previously [35–37], and that we can achieve very small errors in the energy difference by ensuring that the whole system is relaxed before each move is performed.

We have chosen to use a relaxation region of 754 atoms about the original vacancy location, which produces errors in the energy of the order of 1 meV.

Table 1. Aluminium alloy compositions.

	Aluminium			Copper			Magnesium		
Alloy	at.%	(wt%)	n_{Al}	at.%	(wt%)	n_{Cu}	at.%	(wt%)	n_{Mg}
1	98	(95.4)	10 755	2	(4.6)	220			
2	97.8	(95.2)	10 733	2	(4.6)	220	0.2	(0.2)	22
3	97.5	(95.0)	10 700	2	(4.6)	220	0.5	(0.4)	55
4	96	(91.1)	10 536	4	(8.9)	439			
5	95.8	(90.9)	10 514	4	(8.9)	439	0.2	(0.2)	22
6	95.5	(90.6)	10 481	4	(8.9)	439	0.5	(0.4)	55
7	94.0	(86.9)	10 316	6	(13.1)	659			

3. Clustering in the Al–Cu and Al–Cu–Mg systems

The alloy compositions we have studied are listed in table 1. We have chosen a supercell of 10 976 atoms ($14 \times 14 \times 14$ fcc unit cells) run for at least one hundred thousand vacancy exchanges.

We have modelled natural ageing, that is to say ageing at room temperature (300 K) as opposed to artificial ageing at elevated temperatures. This is to provide a large driving force for clustering. We have also chosen to use a single vacancy in these diffusion simulations. The equilibrium vacancy concentration ρ at a given temperature T is determined by the free energy of formation ΔG_f of a vacancy, given by $\rho = \exp(-\Delta G_f / k_B T)$. The constant vacancy concentration we have chosen, being nearly 1×10^{-5}, is in equilibrium at somewhat higher solutionizing temperatures. We are assuming that the system has been held at this higher temperature until a random solid solution of atoms with this vacancy concentration has formed. The system is then rapidly quenched to our simulation temperature, here being 300 K.

Note that, as we are performing off-lattice Monte Carlo simulation, the removal of vacancies by dislocation climb or absorption at grain boundaries is allowed in the simulation procedure, even though these cases have not been explicitly considered here.

The aim of this set of computer simulations is to answer the following questions:

(i) Does incorporating elastic strain effects into the determination of the diffusion path of a vacancy have any real effect on the kinetics of phase separation?
(ii) What difference, if any, is made by adding a small concentration of magnesium atoms to an Al–Cu alloy?

3.1. The effect of elastic strain on vacancy diffusion

The kinetic Monte Carlo algorithm [25] returns a real time, which is a stochastic variable approximating the actual elapsed time during the transition drawn from gamma distributions. In the Flynn model of equation (2.1), the time taken is determined from the average elastic compliance of the system, the atomic volume of the migrating atom and the difference between the energies of the initial and final states for each jump. To complete the parameterization of the Flynn model, we need a value for the Debye frequency. The results we present assume a constant Debye frequency equal to that of pure aluminium, taken to be 1.19×10^{13} Hz.

Only one type of transition, a nearest neighbour vacancy–atom exchange, is permitted in this simulation. So we can identify the rate of hopping as proportional to the mobility of the vacancy. The time taken for the same number of vacancy moves with and without atomic relaxation is shown in figure 3. In the case of the Al–Cu alloys, we find that running the

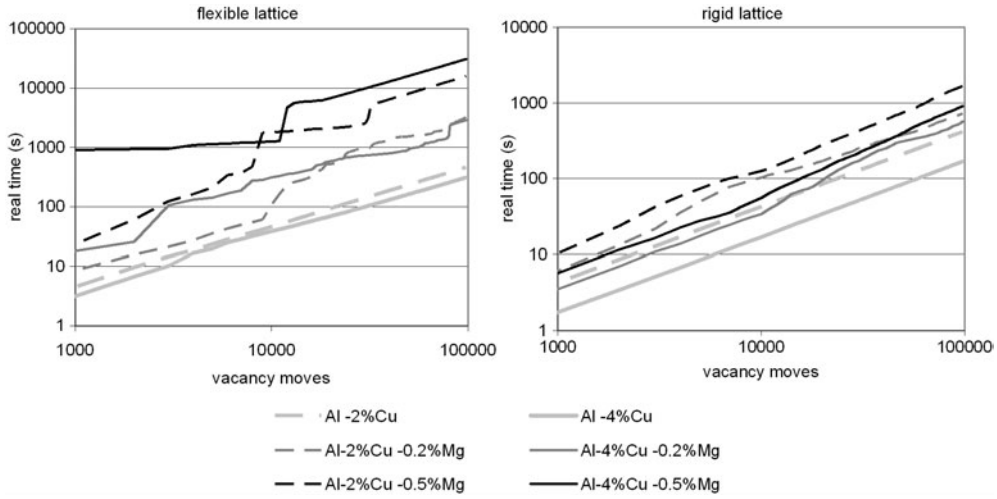

Figure 3. The effect of performing the simulation on a flexible lattice on the real time simulated. Times shown are the simulated real time (in seconds) for a sequence of 100 000 vacancy moves.

simulation on a flexible lattice increases the real time taken for a given number of exchanges. On a flexible lattice, the real time simulated is roughly 10%–20% greater over 100 000 exchanges, reflecting a commensurate decrease in the mean vacancy hopping rate. In more familiar terms, the *mobility* of the vacancy is reduced on a flexible lattice. The origin of this reduction must be that the relaxations in a flexible lattice produce a greater variability in the configurational energy of the system as the vacancy migrates. In a flexible lattice the vacancy is attracted to certain sites, where it spends longer periods of time. We are confident that the variations in the configurational energies we calculate are real in the sense that they are not a consequence of inadequately converged relaxations, as discussed in section 2.5.

In the case of Al–Cu–Mg alloys, the time differences in the flexible and rigid lattices are even more pronounced, with the time simulated being five times greater with 0.2% Mg (alloys 2, 5), and around ten times greater with 0.5% Mg (alloys 3, 6). Note that steep steps in the plot are seen in the flexible lattice calculation when Mg is present, indicating that the vacancy is trapped temporarily. In the rigid lattice calculation these steps are absent in these early stages, although steps do appear later when there is significant clustering [25].

3.2. The effect of elastic strain on clustering

3.2.1. Al–Cu. We now consider the short-range ordering of copper atoms. We know from separate Monte Carlo simulations using the same potentials that $L1_0$- and $L1_2$-type intermetallic structures, rather than the θ phase, are preferentially generated. We can define a short-range order parameter to measure the degree of ordering of copper atoms into pairs separated by a $\langle 100 \rangle$ vector. There are n_{Cu} copper atoms in the system. At any time, there will be a certain number of such pairs of atoms $n_{\langle 100 \rangle}$. A periodically repeated block of $L1_2$ structure Al_3Cu with the same number n_{Cu} of copper atoms would have every neighbouring copper pair linked by a $\langle 100 \rangle$-type vector. An illustration of this type of clustering is shown in figure 4. There would be $\frac{1}{2} \times 6 \times n_{\mathrm{Cu}}$ such pairs, where the factor of a half stops us from double counting each pair. Our short-range order parameter, $\mathrm{sro}_{\mathrm{Cu}}$, measures how close we are getting to this

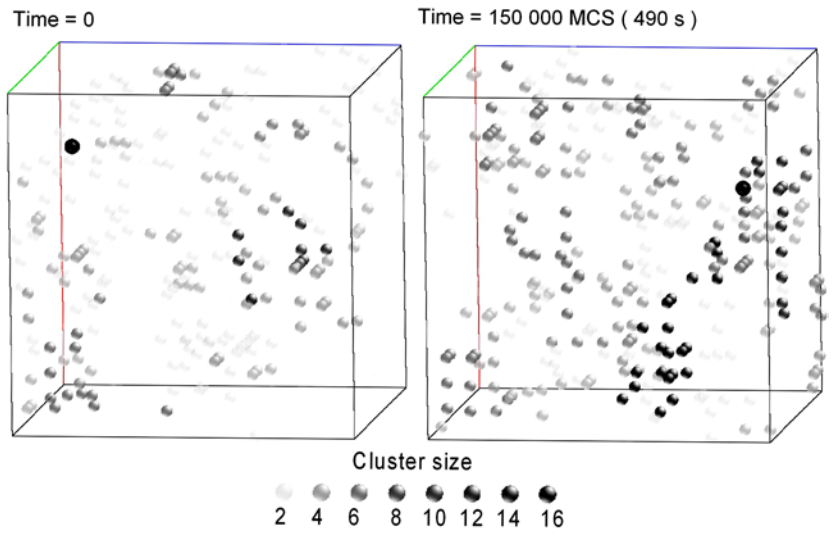

Figure 4. Snapshots from the Al–4 at.% Cu system with 10 976 atoms on a flexible lattice. The vacancy location is indicated with a black sphere. The copper atoms are shaded according to the number in the cluster, with single copper atoms and aluminium atoms not shown. Two copper atoms belong to the same cluster if they are nearest or next nearest neighbours. After 150 000 vacancy moves (a real time of 490 s) the copper atoms have clustered together significantly.

low energy intermetallic phase:

$$\mathrm{sro}_{\mathrm{Cu}} = \frac{n_{\langle 100\rangle}}{\frac{1}{2} \times 6 \times n_{\mathrm{cu}}}. \tag{3.1}$$

We can investigate how the clustering is affected by the lattice being flexible by plotting this short-range order parameter as a function of the real time simulated. In figure 5 we plot the evolution of the short-range ordering of copper atoms in Al–Cu on a flexible and on a rigid lattice. The traces corresponding to the rigid lattice calculations are much less computationally expensive than those on a flexible lattice, and so have been generated by averaging over five independent simulation runs. The order parameter is found as the average over fixed time intervals, evenly spaced on the logarithmic real timescale.

The traces for the flexible lattice calculations each took about one month of computation on four Sun UltraSPARC III 900MHx Processors. For this reason these simulations were done once only.

It is clear that at all three alloy compositions simulated there is a marked *increase* in the rate of ordering of copper atoms on a flexible lattice, despite the reduction of the vacancy mobility seen in figure 3. The speed of diffusing copper atoms is determined by the product of the vacancy mobility and the force driving their motion. Therefore there must be an additional driving force for clustering on a flexible lattice as compared with a rigid lattice. Once clustering has initiated somewhere in the system it will bias the motion of other copper atoms towards the cluster along lobes of compressive stress surrounding the cluster. The additional driving force must more than compensate the reduction of the vacancy mobility to produce the increased rate of ordering of copper atoms on a flexible lattice.

3.2.2. Al–Cu + 0.2 at.% Mg. Copper atoms are expected to be centres of tension, and magnesium atoms centres of compression in an aluminium matrix. Therefore, there will be

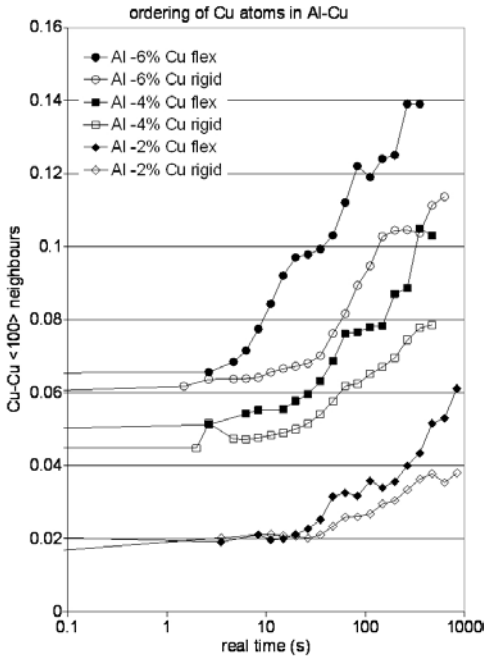

Figure 5. Ordering in the Al–Cu system. These simulations were performed with a single vacancy in a periodic supercell of 10 976 lattice points, with zero external stress, at a simulated temperature of 300 K. The simulations comprised approximately 200 000 vacancy moves for alloys 1 and 7 and 250 000 moves for alloy 4. The order parameter is defined in equation (3.1). The lines are to guide the eye.

an attractive elastic interaction between Cu and Mg atoms on a flexible Al lattice. We might therefore expect to see Cu atoms segregating to Mg-rich regions. However, the rate at which such a process occurs may be reduced by vacancies becoming trapped by Mg atoms. This trapping may be expected to occur more effectively on a flexible lattice where the motion of vacancies is biased by the elastic fields of Mg atoms, and more stable vacancy–Mg atom complexes may form through elastic relaxation.

In figure 6 we examine the effect of adding microalloying quantities of Mg atoms to Al–Cu alloys on the clustering of Cu and Mg atoms. The degree of clustering of magnesium atoms may be followed by defining an order parameter for the evolution of close packed clusters of magnesium atoms:

$$\mathrm{sro}_{\mathrm{Mg}} = \frac{n_{\langle \frac{1}{2}\frac{1}{2}0\rangle}}{\frac{1}{2} \times 12 \times n_{\mathrm{mg}}} \tag{3.2}$$

where $n_{\langle \frac{1}{2}\frac{1}{2}0\rangle}$ is the number of $(\mathrm{Mg–Mg})_{\langle \frac{1}{2}\frac{1}{2}0\rangle}$ pairs, and n_{Mg} is the number of magnesium atoms present in the system. It can be seen that the rate of ordering of copper atoms is still enhanced on a flexible lattice in the presence of 0.2 at.% Mg. However, the rate of clustering of magnesium atoms is reduced on a flexible lattice until later times.

3.2.3. Al–Cu + 0.5 at.% Mg. In figure 7 the effect of adding a higher concentration of Mg is simulated. With the higher Mg content it is no longer the case that the rate of ordering of copper atoms on a flexible lattice is enhanced.

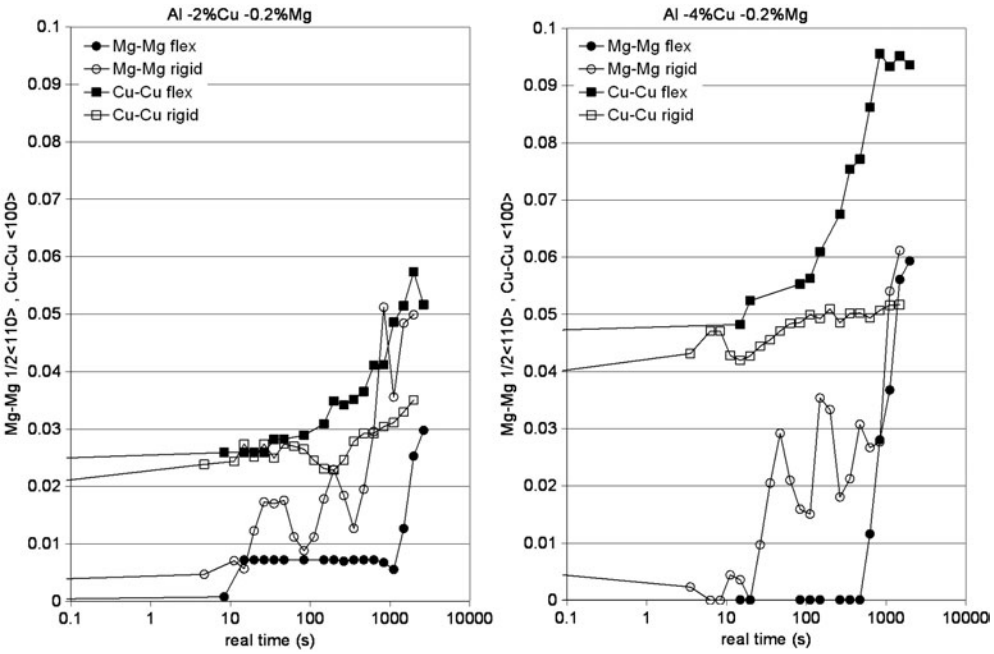

Figure 6. Ordering in the Al–Cu–0.2% Mg system. These simulations used the same parameters as those in figure 5, and the traces were generated in the same way. The degree of clustering of Mg atoms is measured by the short-range order parameter defined in equation (3.2). We see that there is some clustering of Mg atoms, and ordering of Cu atoms, present on both rigid and flexible lattices. The rate of ordering of copper atoms is significantly increased on the flexible lattices. The clustering of magnesium atoms seems to occur at later times on a flexible lattice as compared with a rigid lattice.

It is seen that again the Mg atoms are clustering at a later time when the simulation is run with a flexible lattice. However, when it occurs, the degree of clustering of Mg is enhanced at this higher Mg concentration.

3.3. The association of vacancies with Mg atoms

The number of magnesium atoms surrounding each vacancy is plotted as a function of time for rigid and flexible lattices in figure 8. On a rigid lattice we see that the number of magnesium atoms around a vacancy increases with time, but that this is not the case on a flexible lattice, at least for real times of up to 1000 s. On the rigid lattice, magnesium clusters are forming, and the vacancies are associated with these clusters.

4. Summary of results and discussion

4.1. Summary

We may summarize the principal results of our simulations as follows.

(i) The real time taken for the first 100 000 vacancy moves in a flexible lattice is greater for all alloys we have considered than in a rigid lattice. This is evident in figure 3. Therefore, the mobility of the vacancy is reduced on a flexible lattice.

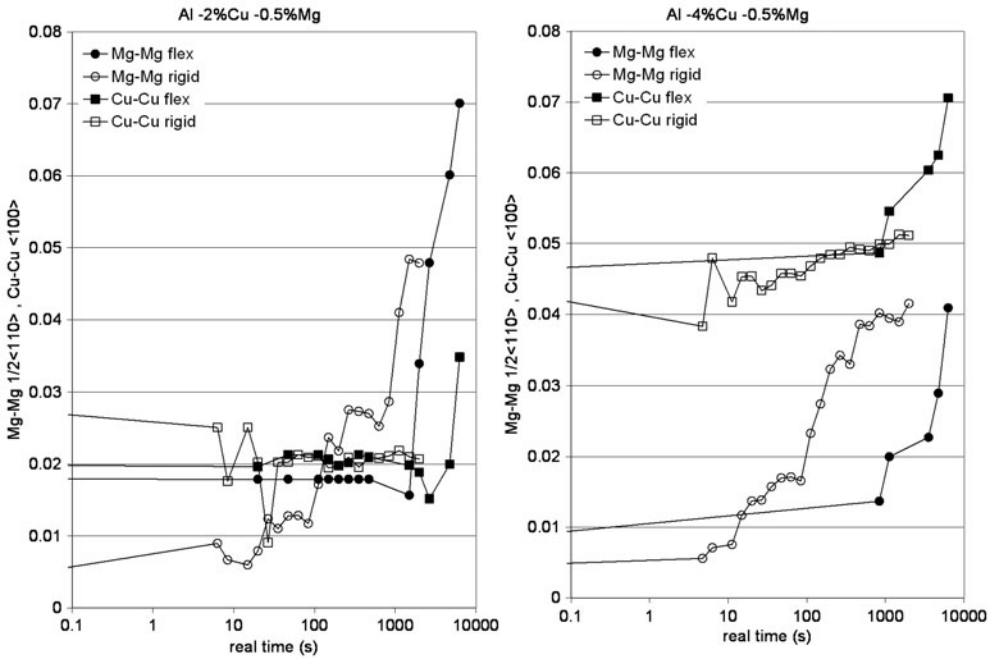

Figure 7. Ordering in the Al–Cu–0.5% Mg system. These simulations used the same parameters as those in figure 6, and the traces were generated in the same way. Again the magnesium atoms are seen to cluster at later times on a flexible lattice than were seen on the rigid lattice. However, here the clustering of Mg atoms is more extensive, and the clustering of copper atoms reduced.

(ii) In figure 3 we also saw that on a flexible lattice the presence of Mg in the alloy has the effect of increasing the time taken for 100 000 vacancy moves, and vacancy trapping is more evident.

(iii) Figure 5 shows that, in the absence of Mg, Cu orders more rapidly on a flexible lattice than on a rigid lattice. In view of the reduced mobility of vacancies in a flexible lattice, this implies that the driving force for ordering of Cu is greater in a flexible lattice. This is expected because on a flexible lattice the relaxation enables Al–Cu bonds to adjust closer to their ideal length.

(iv) At 0.2 at.% Mg, copper orders more rapidly on a flexible lattice than on a rigid lattice (see figure 6), but at 0.5 at.% Mg (see figure 7), this difference is much less pronounced, except possibly at the longest real times simulated. The driving force for ordering of Cu atoms is unlikely to be affected by such small additions of Mg. We suggest, therefore, that at the higher Mg concentration the vacancy is more frequently trapped by Mg atoms.

(v) Figures 6 and 7 show that Mg does not initially cluster as rapidly on a flexible lattice as on a rigid lattice. We attribute this behaviour to the increased trapping of the vacancy in the presence of Mg on a flexible lattice. We note that in figures 6 and 7 the real time is plotted on a logarithmic scale, and therefore most of the vacancy moves take place at the far right of these plots.

(vi) Figure 8 shows that at a given time there are on average more Mg atoms around each vacancy on a rigid lattice than on a flexible lattice. Moreover, this difference increases with time. This implies that clusters of Mg atoms are forming on a rigid lattice and that vacancies are associated with them.

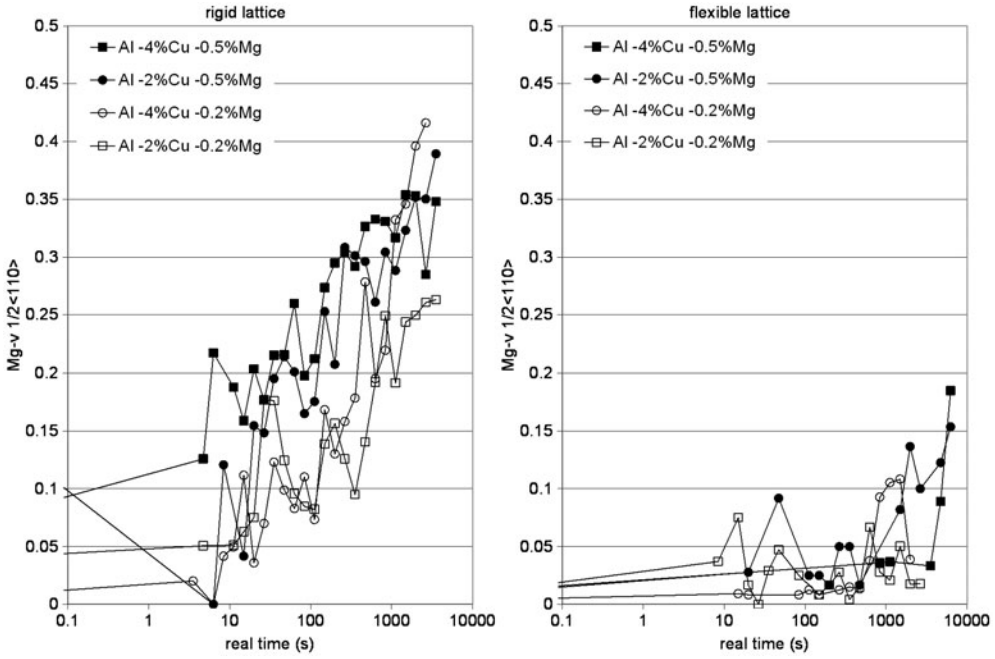

Figure 8. The proportion of atoms surrounding the vacancy which are magnesium. On a rigid lattice (left) the number of magnesium atoms surrounding a vacancy increases with time. This, together with the clustering of magnesium atoms seen in figures 6 and 7, shows that on a rigid lattice magnesium clusters are forming, and the vacancies are associated with them. On a flexible lattice (right) the number of magnesium atoms around a vacancy is not increasing.

4.2. Conclusions

We have presented a method to rapidly evaluate on-the-fly transition rates for possible vacancy moves incorporating atomistic elastic relaxation, by approximating the saddle points as a function of the internal energy of the initial and final states. We found that the harmonic lattice approximation was insufficient for measuring the tiny difference in the elastic relaxation energy between possible vacancy moves. However, a scheme based on relaxing only the local region around the vacancy was capable of reducing errors in the system energy to an acceptable fraction of the thermal energy. This short-range scheme was found to be successful for evaluating the transition energies for possible vacancy moves, although a full atomistic relaxation was required after the move was decided.

The Lanczos method is particularly suitable for atomic relaxation after a single vacancy move, as each iteration effectively spreads the relaxation region further from the atoms experiencing a force around the vacancy. Although the relaxations remain time-consuming, our use of this relaxation scheme has enabled us to model hundreds of thousands of vacancy moves while keeping the system atomistically relaxed.

The rate at which phase separation can occur in alloys is proportional to the mobility of vacancies and to the driving force for diffusion. Both of these may be affected by performing the simulation on a rigid lattice rather than on a flexible lattice. We have seen in figure 3 that the vacancy mobility is quite different on flexible and rigid lattices. Vacancies may be attracted to impurities through long-range elastic fields, and may form complexes that are more stable on a flexible lattice than on a rigid lattice. As a result, the vacancy mobility may be reduced

significantly on a flexible lattice. The exclusion of elastic relaxation in a rigid lattice has the effect of smoothing the configurational potential energy surface, especially in those regions of configuration space where vacancies would have become trapped if elastic interactions had been included.

However, we also saw in figure 5 that the rate of ordering of copper atoms is lower in a rigid lattice than in a flexible lattice. The apparent paradox is resolved by noting that there is a reduced driving force for ordering of copper atoms on a rigid lattice, where bonds are constrained to remain longer than they would be in a relaxed lattice, and this more than outweighs the greater vacancy mobility in a rigid lattice.

A small concentration (0.2–0.5 at.%) of magnesium atoms added to the system are rapidly precipitated out of solution in a rigid lattice, owing to the very high energy compressed bonds around each Mg atom. In a flexible lattice the solubility of Mg is enhanced somewhat, and instead their principal action is to reduce severely the vacancy mobility. Mg atoms are especially trapping in a flexible lattice as the compression in the lattice around each atom is alleviated by the formation of an Mg–vacancy pair.

At present, we have applied the coupling of the Lanczos relaxation algorithm and stochastic second order kinetic Monte Carlo method only to the investigation of phase nucleation in aluminium–copper–magnesium alloys. We believe that the method has more general applicability. With an estimation for the free energy barriers to atom movement into divacancies and larger vacancy clusters, it would be possible without further refinement of the code to model the formation of pores in irradiated materials under elastically relaxed conditions.

Acknowledgments

DRM was supported by an EPSRC grant. The computer simulations were performed on Oswell, the Oxford University Supercomputer. This work was performed in part under the auspices of the US Department of Energy by the University of California, Lawrence Livermore National Laboratory, under Contract No. W-7405-Eng-48. DRM gratefully acknowledges the fellowship in the LLNL Summer Institute on Computational Materials Science and Chemistry through which part of this work was done.

References

[1] Guinier A 1938 Structure of age-hardened Al–Cu alloys *Nature* **142** 569–70
[2] Preston G D 1938 The diffraction of x-rays by age-hardening aluminium copper alloys *Proc. R. Soc.* A **167** 526–38
[3] Silcock J M, Heal T J and Hardy H K 1954 Structural aging characteristics of binary aluminium–copper alloys *J. Inst. Met.* **82** 239
[4] Veitz J T and Polmear I J 1966 The influence of silver on ageing of Al–Cu–Mg alloys *J. Inst. Met.* **94** 410–9
[5] Muddle B C and Polmear I J 1989 The precipitate omega phase in Al–Cu–Mg–Ag alloys *Acta Metall.* **37** 777–89
[6] Ringer S P, Sakwai T and Polmear I J 1997 Origins of hardening in aged Al–Cu–Mg (Ag) alloys *Acta. Metall.* **45** 3731
[7] Hardy H K 1954 The ageing characteristics of some ternary Al–Cu–Mg alloys with Cu:Mg weight ratios of 7:1 and 2.2:1 *J. Inst. Met.* **84** 17–34
[8] Mukhopadhyay A K 2002 Coprecipitation of Ω and σ phases in Al–Cu–Mg–Mn alloys containing Ag and Si *Metall. Mater. Trans.* A **33** 3635–48
[9] Hono K 2002 Nanoscale microstructral analysis of metallic materials by atom probe field ion microscopy *Prog. Mater. Sci.* **47** 621–729
[10] Gerold V 1989 On the structure of GP zones in Al–Cu alloys *Scr. Met.* **22** 927–32
[11] Ringer S P, Hono K and Sakurai T 1995 The effect of trace additions of Sn on precipitation in Al–Cu alloys: an atom probe field ion microscopy study *Metall. Mater. Trans.* A **26** 2207–17

[12] Silcock J M, Heal T J and Hardy H K 1955 The structural ageing characteristics of ternary Al–Cu alloys with Cd, In, Sn *J. Inst. Met.* **84** 23–32
[13] Silcock J M and Flower H M 2002 Comments on a comparison of early and recent work on the effect of trace additions of Cd, In or Sn on nucleation and growth of theta prime in Al–Cu alloys *Scr. Mater.* **46** 389–94
[14] Rautiainen T T and Sutton A P 1999 Influence of the atomic diffusion mechanism on morphologies, kinetics, and the mechanisms of coarsening during phase separation *Phys. Rev.* B **59** 13681–92
[15] Rautiainen T T 1999 Modelling microstructural evolution in binary alloys *DPhil Thesis* University of Oxford
[16] Weinkamer R and Fratzl P 2003 *Europhys. Lett.* **61** 261–7
[17] Eshelby J D 1957 The determination of the elastic field of an ellipsoidal inclusion, and related problems *Proc. R. Soc.* A **241** 376–96
[18] Eshelby J D 1959 The elastic field outside an ellipsoidal inclusion *Proc. R. Soc.* A **252** 561–9
[19] Schmidt I, Mueller R and Gross D 1998 The effect of elastic inhomogeneity on equilibrium and stability of a two particle morphology *Mech. Mater.* **30** 181–96
[20] Schmidt I and Gross D 1997 The equilibrium shape of an elastically inhomogeneous inclusion *J. Mech. Phys. Solids* **45** 1521–49
[21] Khachaturyan A G 1983 *Theory of Structural Transformations in Solids* (New York: Wiley)
[22] Khachaturyan A G 1967 Some questions concerning the theory of phase transitions in solids *Sov. Phys.—Solid State* **8** 2163–8
[23] Rubin G and Khachaturyan A G 1999 Three-dimensional model of precipitation of ordered intermetallics *Acta Mater.* **47** 1995–2002
[24] Khachaturyan A G, Semenovskaya S and Tsakalakos T 1995 Elastic strain energy of inhomogeneous solids *Phys. Rev.* B **52** 15909–19
[25] Mason D R, Rudd R E and Sutton A P 2003 Stochastic kinetic Monte Carlo algorithms for long-range Hamiltonians, submitted
[26] Fichthorn K A and Weinberg W H 1991 Theoretical foundations of dynamical Monte Carlo simulations *J. Chem. Phys.* **95** 1090–6
[27] Bortz A B, Kalos M H and Lebowitz J L 1975 A new algorithm for Monte Carlo simulation of Ising spin systems *J. Comput. Phys.* **17** 10–8
[28] Novotny M A 1995 *Comput. Phys.* **9** 46
[29] Vineyard G H 1957 Frequency factors and isotope effects in solid state rate processes *J. Phys. Chem. Solids* **3** 121–7
[30] Flynn C P 1972 *Point Defects and Diffusion* (Oxford: Clarendon)
[31] Finnis M W and Sinclair J E 1984 A simple empirical *n*-body potential for transition metals *Phil. Mag.* A **50** 45–55
[32] Saad Y 1996 *Iterative Methods for Sparse Linear Systems* (Boston, MA: PWS) http://www-users.cs.umn.edu/~saad/books.html
[33] Heine V, Bullett D W, Haydock R and Kelly M J 1980 *Solid State Physics* vol 35 (New York: Academic)
[34] Paige C C and Saunders M A 1972 Solution of sparse indefinite systems of linear equations *SIAM J. Numer. Anal.* **12** 617–24
[35] Ikeda H and Matsuda H 1992 Computer simulations of cluster formation processes with relaxation of lattice distortion *Mater. Trans. JIM* **33** 466–71
[36] Ikeda H and Matsuda H 1993 Effects of difference in elastic moduli between coefficients on spinodal decomposition processes *Mater. Trans. JIM* **32** 651–7
[37] Ikeda H and Matsuda H 1996 Computer simulation of phase decomposition process generating precipitates harder than matrix *Mater. Trans. JIM* **37** 1413–21