White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/79862/

Version: Accepted Version

Article:
Mullis, AM, Bollada, PC and Jimack, PK (2014) Towards a 3-dimensional phase-field model of non-isothermal alloy solidification. Materials Science Forum, 783-78. 2166 - 2171. ISSN: 0255-5476

https://doi.org/10.4028/www.scientific.net/MSF.783-786.2166

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Towards a 3-Dimensional Phase-Field Model of Non-Isothermal Alloy Solidification

Andrew M Mullis[1,a], Peter C Bollada[1,b] and Peter K Jimack[2,c]

[1]Institute for Materials Research, University of Leeds, Leeds, LS2-9JT, UK.

[2]School of Computing, University of Leeds, Leeds, LS2-9JT, UK.

[a]a.m.mullis@leeds.ac.uk, [b]p.c.bollada@leeds.ac.uk, [c]p.k.jimack@leeds.ac.uk

Keywords: Phase-field; Adaptive Mesh Refinement; Multigrid; High Performance Computing.

Abstract. We review the application of advanced numerical techniques such as adaptive mesh refinement, implicit time-stepping, multigrid solvers and massively parallel implementations as a route to obtaining solutions to the 3-dimensional phase-field problem for coupled heat and solute transport during non-isothermal alloy solidification. Using such techniques it is shown that such models are tractable for modest values of the Lewis number (ratio of thermal to solutal diffusivities). Solutions to the 3-dimensional problem are compared with existing solutions to the equivalent 2-dimensional problem.

Introduction

The modelling of dendritic growth is a subject of enduring interest within the metallurgical community, both because dendrites are a prime example of spontaneous pattern formation and because they have a pervasive influence on the engineering properties of metals. Remnants of these dendritic microstructures often survive subsequent processing operations, such as rolling and forging, and the length scales established by the dendrite can influence not only the final grain size but also micro- and hence macro-segregation patterns.

The technique which over the last few decades has perhaps received the most attention for simulating crystal structures from the melt is that of phase-field simulation [1, 2], in which a non-conserved order parameter ϕ, which encodes the phase state of the material, is defined over the whole domain. By assuming the interface between the solid and liquid to be diffuse, ϕ is rendered continuous, wherein standard techniques for partial differential equations (PDEs) may be used without the need to explicitly track the solid-liquid interface. The technique has been subjected to careful validation against experiment, specifically in relation to dendrite growth velocities and tip radii, with good agreement being shown against experiments performed in microgravity conditions as part of the US Isothermal Dendrite Growth Experiment [3]. As a result, phase-field modelling is now the technique of choice for simulating solidification microstructures.

However, phase-field modelling presents significant computational challenges in that the resulting set of PDE's is highly non-linear and generally the width of the diffuse interface must be much narrower than the smallest physical feature to be simulated. This results in very large computational meshes, particularly when the problem is solved in 3-dimensions. Consequently, most phase-field models employ a restrictive set of assumptions wherein, with only a very few exceptions [4, 5], the models are constructed within the idealized limit that solidification is controlled by the diffusion of one species only, either heat or solute. This is true for the solidification of pure elements (thermal diffusion only) and approximately true during the very slow solidification of alloys, wherein the system may be considered isothermal (solutal diffusion only). However, this severely restricts the application of phase-field modelling in the field of rapid solidification where, even within concentrated alloys, growth rates are such that the influence of the thermal field cannot be ignored.

The phase-field modelling of non-isothermal alloy solidification poses particular multi-scale challenges. Not only do we require that the diffuse interface be much finer than the scale of the

microstructure to be simulated, we must also resolve the diffusion fields for both heat and solute. The thermal and solutal diffusivities for most liquid metals are typically 10^{-5} m^2 s^{-1} and 10^{-9} m^2 s^{-1} respectively. The ratio of these two, referred to as the Lewis number ($Le = D/\alpha$), is therefore typically of the order 10,000 and gives a measure of the multi-scale nature of the problem, with the thermal field being Le times more extensive than the solute field, which is itself likely to be > 10 times larger than the diffuse phase-field interface.

In this paper we describe how advanced numerical techniques can be combined to solve complex multi-scale problems in materials science. Adaptive mesh refinement is an economical strategy for reducing the total mesh size in a discrete element simulation, but if used in isolation may not be an efficient means of tackling multiscale problems, as the time step will generally be limited by the size of the smallest element in the mesh. Here we describe how adaptive mesh refinement can be combined with implicit time-stepping to achieve economical meshing combined with a time step which is mesh independent. A multigrid solver is used to invert the large, but sparse, algebraic system that results at each time step, such that the computational resource required scales linearly with the number of degrees of freedom in the problem. In the parallel implementation of this scheme dynamic load balancing is used to ensure the even distribution of work between cores is retained as the mesh adapts.

Mathematical Model

The phase-field model used here is a hybrid of the 3-dimensional solute only model used by Echebarria *et al*. [6] and the 2-dimensional coupled thermo-solutal model due to Ramirez and Beckermann [4]. The formulation follows the 'thin-interface' methodology [7], whereby quantitatively valid results can be obtained in a manner independent of the interface width. Following non-dimensionalization against characteristic length and time scales, W_0 and τ_0, the phase equation is given by

$$A^2\left[\frac{1}{Le}+Mc_\infty\left(1+(1-k_E)U\right)\right]\frac{\partial\phi}{\partial t}=\nabla\cdot(A^2\nabla\phi)+\phi(1-\phi^2)+\lambda(1-\phi^2)^2(\theta+Mc_\infty U)$$
$$+\frac{\partial}{\partial x}\left[A\left(-\frac{\partial A}{\partial\vartheta}\frac{\phi_y|\nabla\phi|^2}{\phi_x^2+\phi_y^2}+\frac{\partial A}{\partial\psi}\frac{\phi_z\phi_x}{\sqrt{\phi_x^2+\phi_y^2}}\right)\right]+\frac{\partial}{\partial y}\left[A\left(\frac{\partial A}{\partial\vartheta}\frac{\phi_y|\nabla\phi|^2}{\phi_x^2+\phi_y^2}+\frac{\partial A}{\partial\psi}\frac{\phi_z\phi_y}{\sqrt{\phi_x^2+\phi_y^2}}\right)\right] \quad , \tag{1}$$
$$-\frac{\partial}{\partial z}\left[A\left(\frac{\partial A}{\partial\psi}\sqrt{\phi_x^2+\phi_y^2}\right)\right]$$

with an anisotropy, $A = A(\vartheta,\psi)$, given in terms of the standard spherical angles, ϑ and ψ, by

$$A(\vartheta,\psi)=A_0\left\{1+\varepsilon\left[\cos^4\psi+\sin^4\psi\left(1-2\sin^2\vartheta\cos^2\vartheta\right)\right]\right\} \quad . \tag{2}$$

This corresponds to a preferred growth direction aligned with the Cartesian coordinate axis. The small parameter ε governs the strength of the anisotropy, M is the scaled magnitude of the liquidus slope (the dimensional form of which is m), c_∞ is the solute concentration far from the interface, D is the dimensionless solute diffusivity and θ is a scaled temperature given by

$$\theta=\frac{\Delta T-mc_\infty}{L/c_p} \quad , \tag{3}$$

where L and c_p are the latent heat of fusion and the specific heat capacity respectively. λ is a coupling parameter which determines the width of the diffuse interface, W_0, via the relations

$$\lambda = \frac{a_1 W_0}{d_0} = \frac{D}{a_2}\ , \quad a_1 = \frac{5\sqrt{2}}{8}, \quad a_2 = 0.6267 \quad . \tag{4}$$

The evolution equation for the dimensionless concentration field, U, is given by

$$\left(\frac{1+k}{2} - \frac{1-k}{2}\phi\right)\frac{\partial U}{\partial t} = \nabla\cdot\left\{D\frac{1-\phi}{2}\nabla U + \frac{1}{2\sqrt{2}}\left[1+(1-k_E)U\right]\frac{\partial\phi}{\partial t}\frac{\nabla\phi}{|\nabla\phi|}\right\} + \frac{1}{2}\left[1+(1-k_E)U\right]\frac{\partial\phi}{\partial t} \tag{5}$$

where k_E is the equilibrium partition coefficient. Here, U is related to the concentration, c, via

$$U = \frac{1}{1-k_E}\left(\left(\frac{2c/c_\infty}{1+k_E-(1-k_E)\phi}\right)-1\right) \quad . \tag{6}$$

Finally, the evolution of the temperature field is given by

$$\frac{\partial\theta}{\partial t} = \alpha\nabla^2\theta + \frac{1}{2}\frac{\partial\phi}{\partial t} \quad .$$

Numerical Methods

The implementation of explicit solution methods for systems of PDE's is straight forward, but they suffer from a time-step restriction in order to ensure the stability of the discretization scheme, which is of the form $\Delta t \le C.h^2$ where h is the minimum element size within the mesh. Consequently, if explicit methods are used with adaptive meshing the time step to be applied is determined by the smallest element in the mesh. This can severely limit the application of such techniques if fine spatial resolution is required. Moreover, it can be shown that for the complex equations considered here $C = C(\lambda, Le)$, with C varying from ≈ 0.3 at $Le = 1$ to $C \le 0.001$ at $Le = 500$ [5] at which point the solution becomes intractable using the explicit Euler method, despite $Le = 500$ still being significantly short of values typical of liquid metals. However, if utilising an implicit time discretization scheme the appropriate selection of Δt shifts from being a question of numerical stability to one of accuracy. We use the 2nd Order Backward Differences (BDF2) method with a variable time step of ratio, where the choice of appropriate time steps is now based upon a set of local error estimators. This method can be shown to be A-stable [8].

When using implicit time discretisation methods it is necessary to solve a large, but sparse, system of non-linear algebraic equations at each time-step. Multigrid methods are among the fastest available solvers for large sparse systems of equations and in this work we apply the non-linear generalization known as Brandt's full approximation scheme (FAS) [9] directly to the algebraic system that arises at each time step. For a large sparse system of equations in N unknowns such methods scale as N, rather than as $N^{3/2}$ for conjugate gradient and successive over-relaxation methods or N^2 for Gauss elimination. Multigrid is therefore the key ingredient that allows the practical implementation of implicit time stepping and is based upon two principles: the coarse grid principle and the smoothing principle. For the coarse grid correction one has to define grid transfer operators to transfer the solution and the residual from the fine grid to the coarse grid and the solution from the coarse to the fine grid. Bilinear interpolation is used for the coarse to fine grid transfer and injection is used for the fine to coarse transfer. For smoothing the error we use a basic pointwise nonlinear weighted Gauss-Seidel method. The number of pre- and post-smoothing operations required for convergence has been investigated and it has been found that a V-cycle iteration with 2 pre- and 2 post- smoothing operations is optimal.

In 2-dimensions we use a quadrilateral mesh structure which is non-conforming, in the sense that we allow hanging nodes, with adaptive refinement based upon an elementwise gradient criterion. The serial implementation of this is reported in detail elsewhere [5] and has recently been made available by us as part of the open source package PhAIM-2d (Phase-field by Adaptive Implicit Multigrid in 2-dimensions) [10]. However, in 3-dimesions parallel implementation is required, wherein dynamic load balancing during adaptive refinement and coarsening of the mesh becomes a significant issue if efficient computation is to be achieved. In order to control the three-dimensional mesh refinement and de-refinement we use the open source library, PARAMESH [11]. This library provides functions to generate hexahederal mesh blocks in an oct-tree structure. Starting with a base block (of 8×8×8 cubic cells for example) it is possible to refine this into up to 8 child blocks (with each block always being of the same dimension as the base block) and then to refine any of these child blocks successively. Functions are also provided to undo regions of this local refinement (i.e. de-refinement) and to interpolate or restrict solution fields between meshes. A further capability of PARAMESH is that it is able to undertake this meshing in parallel in a manner that is hidden from the user – each block is simply treated as independent of its neighbours and PARAMESH takes care of which process owns each block, using its own dynamic load balancing scheme. A price that has to be paid for this simplicity is that every block is required to store guard cells in each dimension regardless of whether its neighbouring blocks are actually owned by a different process: PARAMESH's guard cell update routines then take care of all of the transfer of data between neighbouring blocks, regardless of their location in memory.

The use of PARAMESH imposes a number of constraints upon our choice of finite difference stencil. Specifically, we avoid the use of any points around cell (i, j, k) that are not of the form (i±1, j±1, k±1) as this ensures that our parallel implementation needs only a single layer of guard cells – which reduces the memory and communication overhead significantly. For the results reported here a compact 27-point stencil is used, which is found to significantly reduce mesh induced anisotropy relative to the standard 2^{nd} order 7-point stencil in 3-dimensions.

The local refinement and de-refinement capability provided by PARAMESH is essential for this work since our phase-field models require very fine meshes around the solid-liquid interface in order to ensure that the interface is resolved with sufficient accuracy. The nondimensionalization used to derive the systems introduced in section 2 is such that the interface width is $O(1)$ and so our mesh spacing cannot be greater than $h = 1$ around the interface. Hence the finest grid resolution needs to be at least this size (for a domain of dimension (0,800)×(0,800)×(0,800) at least ten levels of refinement are required, wherein $800/2^{10}$ gives $h = 0.78125$ – though an eleventh level is necessary if we wish to ensure that the interface is even moderately well resolved in its normal direction). Without the use of local mesh refinement and de-refinement there would need to be an excessive number of cells, creating a computational load that would be unmanageable without the very largest supercomputing resources, a uniform mesh with comparable resolution to our level 11 mesh here having ≈ 8.6 billion elements.

A major property of the multigrid method is h-independent convergence, which means that the convergence rate does not depend upon the element size. This behaviour is vital in respect of being able to solve the extreme multi-scale problem arising from coupled thermo-solutal phase-field simulations. The extension of the PARAMESH capability to include nonlinear multigrid is explained in [12]. The essential ingredients are the extension of the restriction and prolongation operators for the FAS scheme and for the use of the multi-step BDF formula (requiring data from previous time steps to be used at each multigrid level).

Results

Extensive numerical validation of both the 2- and 3-dimensional models has been undertaken. For the 2-dimensional model validation has been undertaken by comparing with published results obtained using the same equation set [4], albeit solved with an explicit computational scheme. Consequently, validation against directly comparable data set is only possible over a relatively

limited range of Lewis numbers, typically Le ≈ 40. Validation of the 3-dimensional model is rather more difficult as there are no published results for coupled thermo-solutal phase-field simulations in 3-dimensions. Validation has therefore been undertaken according to the following protocol:

i) The concentration, Mc_∞, is set to zero and Le is set to 1 in order to recover the pure thermal model, which can then be compared against the published data in 3-dimensions [13].

ii) The thermal field is everywhere fixed at a constant value and $1/Le$ is set to zero, wherein the pure solutal model is recovered. The solution for equiaxed dendritic growth has then been compared against a set of simulations runs specifically for this purpose using the 3-dimensional adaptive, explicit code due to Dantzig *et al*. [14].

iii) The 3-D model is reduced to 2-D, a straightforward process in PARAMESH, wherein the reduced PARAMESH model can be compared both against published data for 2-dimensional coupled thermo-solutal growth and against the results for our PhAIM2-d model.

Validation was by comparison of the velocity, V, and radius of curvature, R, at the dendrite tip. The results of these test, all of which agreed with published values to within 5% or better, are given in [5] for the 2-dimensional model and in [15] for the 3-dimensional model.

In 2-dimensions solutions for Lewis number of $O(10{,}000)$, which is physically realistic for most liquid metals, can be routinely undertaken using serial execution [16]. In 3-dimensions the simulations are significantly more challenging and typically parallel execution on 256-1024 cores is required. An example 3-dimensional dendrite morphology calculated for an alloy with $Mc_\infty = 0.05$, $k_E = 0.3$, $\gamma = 0.02$, $D = 1.2533$ ($\lambda = 2$) and $Le = 40$ ($\alpha = 50.132$) is shown in Fig. 1. The simulation is run on a $[0{:}800]^3$ domain with 11 levels of refinement, wherein the minimum mesh spacing is $h = 0.39$. Also shown in Figure 1 is the evolution, as a function of dimensionless simulation time, τ, of the (dimensionless) tip radius and velocity, for both the 3-dimensional dendrite, and the equivalent 2-dimensional case.

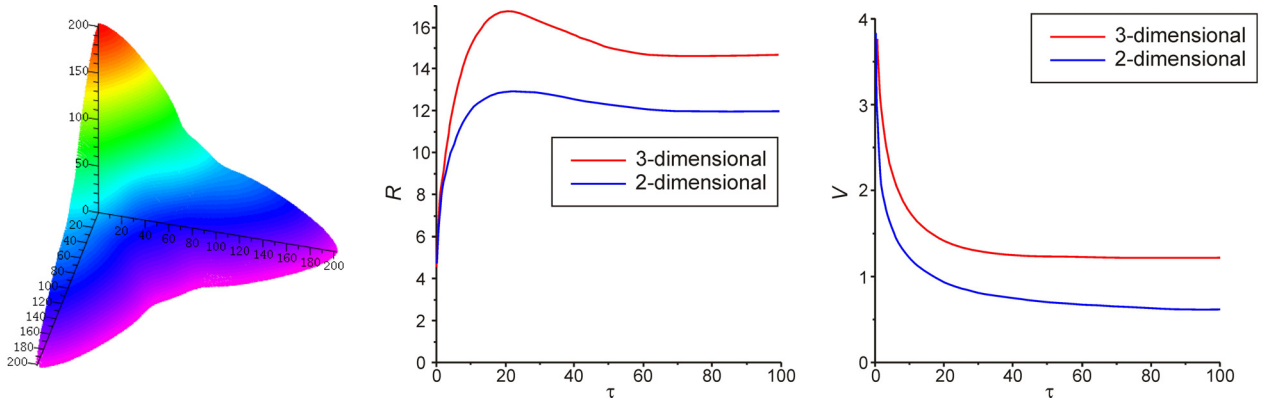

Figure 1. 3D rendering of the $\phi = 0$ isosurface of a non-isothermal alloy dendrite (left) and comparison of the tip radius and velocity of the 3-D dendrite and equivalent 2-D dendrite.

As can be seen, based on either velocity or tip radius, the simulation appears to have reached steady-state by a simulation time of $\tau = 100$, this taking approximately 64 hours on 512 cores (32 × 16-core AMD Opteron 2.3GHz Interlagos processors with 1Gb memory per core). As might be expected the 3-dimensional case grows faster, which is related to the easier diffusion of heat and solute from a 3-dimensional needle dendrite, as opposed to one in 2-dimesions. For the case shown here the stead-state value of the 3-dimensional velocity is larger by a factor of approximately 2. The tip radius is also larger in the 3-dimensional case, although by the rather smaller factor of 1.22.

It is useful to assess these results against analytical models for dendritic growth, such as that due to Lipton, Kurz & Trivedi (LKT, [17]), although caution must be exercised in taking the results of such comparison too literally, as such models are known to have a number of deficiencies. In particular, they do not consider the role of anisotropy upon growth, making use of a stability parameter, σ^*, which is used as an adjustable parameter within the model but which in fact should

be an outcome of the model, σ^* being a Peclet number dependent function of the anisotropy strength. Consequently an exact comparison is not possible, although we consider the general behaviour of the LKT model (and its 2-dimensional analogue, which is obtained by replacing the 3-dimensional Ivantsov solution for the growth of a parabolic needle crystal with the equivalent 2-dimensional solution for the growth of a parabolic plate). Over the plausible range of values for σ^* we find that, as with the phase-field model, the growth velocity is higher in the 3-dimensional case, typically by a rather larger factor than is the case for the phase-field model, but that the tip radius is smaller in the 3-dimesnional case than in 2-dimensions, contrary to the phase-field model. Further phase-field simulations are being conducted to confirm this result.

Summary and Conclusions

- Using a range of advanced numerical techniques including mesh adaptivity, implicit time-stepping, a non-linear multigrid solver and parallel execution we have been able to demonstrate, for the first time, the ability to solve the phase-field equations for coupled thermo-solutal growth of a binary alloy in 3-dimesnions.
- Realistic dendrite geometries growing under stead-state conditions are typically obtained with 50-100 hours on 512 cores using HECToR, the UK National Supercomputer.
- 3-dimensional growth is significantly faster (typically by a factor of 2) when compared with 2-dimensional growth for the same parameter set. The tip radius is also marginally larger in 3-dimensions compared to 2-dimensional growth, counter to expectation from analytical growth models such as LKT [17]. Reasons for these discrepancies are being investigated.

References

[1] J.S. Langer, in G. Grinstein and G. Mazenko (Eds) Directions in condensed matter physics, World Scientific Publishing, Singapore, 1986 pp. 164.

[2] O. Penrose and P.C. Fife, Physica D 43 (1990) 44.

[3] J.C. LaCombe, M.B. Koss, and M.E. Glicksman, Phys. Rev. Lett. 83 (1999) 2997.

[4] J.C. Ramirez and C. Beckermann, Acta Mater. 53 (2005) 1721.

[5] J. Rosam, P.K. Jimack and A.M. Mullis, Acta Mater. 56 (2008) 4559.

[6] B. Echebarria, R. Folch, A. Karma and M. Plapp, Phys. Rev. E 70 (2004) 061604.

[7] A. Karma and W.-J. Rappel, Phys. Rev. E 53 (1996) R3017.

[8] W. Hundsdorfer and J.G. Verwer, Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations, Springer, Verlag, 2003.

[9] A. Brandt, Math. Comput. 31 (1977) 333.

[10] PhAIM2-d is available from www.digital.leeds.ac.uk/software as an OpenSource download.

[11] K. Olson, in A. Deane et al. (Eds) Parallel Computational Fluid Dynamics 2005: Theory and Applications, Elsevier, 2006.

[12] J. R. Green, P.K. Jimack, A.M. Mullis and J. Rosam, Partial Differential Eq, 27 (2011) 106.

[13] A. Karma and W.-J. Rappel, Phys. Rev. Lett., 77 (1996) 4050.

[14] J.H. Jeong, N. Goldenfeld and J. Dantzig, Phys. Rev. E 64 (2001) 041602.

[15] C. Goodyer, P.K. Jimack, A.M. Mullis, H. Dong, X. Yu, Adv. Appl. Math. Mech. 4 (2012) 665.

[16] J. Rosam, P. K. Jimack and A. M. Mullis, Phys. Rev. E 79 030601 (2009).

[17] J. Lipton, W. Kurz and R. Trivedi, Acta Metall. 35 (1987) 957.