Parallel Programming and Computing (CSCI-6360)

Hybrid Potts-Phase field model to simulate microstructure evolution

Scott Peters, Sagar Bhatt

Contents

List of Figures

List of Tables

1. Introduction

One of the fundamental ideas in Materials Science is that the structure of a material is highly important for predicting its properties, whether mechanical, electronic, or otherwise. An example of this is in the grain size of metals, which tends to follow the Hall-Petch equation:

$$\sigma_y = \sigma_0 + \frac{k_y}{\sqrt{d}}$$

Within this equation, σ_y is the yield strength of the material, and d is the average size of a grain in the material (by diameter). Therefore, as the grain size decreases, the yield strength of the material will increase, relative to the two constants σ_0 and k_y. This relationship does not hold for very small grains (less than 10 nm), but matches well with experimental data for larger grains.

Therefore, to accurately predict the strength of materials, it is very important to first be able to predict what the grain size of the particles will be

after a certain length of time being processed. This is what grain growth models are designed to do. By modeling the diffusion of atoms between grains (using both bulk and interfacial energy terms), it is possible to model the flow of atoms from one grain to another, and to show how certain grains will grow, and others will shrink until they disappear.

In following the model we use from the literature, we combine this diffusive process with a Monte Carlo algorithm, in order to more efficiently model this process at the cost of determinism. In the bulk of this report, we describe the exact methods we use to simulate this task, show our results of both outputs and strong scaling, and discuss the results.

2. Background and Methodology

2.1. MMSP

Our project is built around the Mesoscale Microstructure Simulation Project, or MMSP. This project has been developed in-house by previous members of our respective research groups, and has found use in simulating various Materials Science concepts, ranging from phase field problems (Cahn-Hilliard model, dendritic growth, solidification, etc.) to statistical grain-growth problems (e.g. In a 3-D polycrystalline material, how many neighbors, on average, does each grain have?). The implementation of this project is written in C++, and may be downloaded from its Github repository (https://github.com/mesoscale/mmsp) freely for use by anyone.

The primary utility of this software (that we take advantage of for this project) is MMSPs Grid class. This object stores an array of grid points (in 1-3 dimensions), where each point contains some information about the material represented within the grid. This could include the order parameter (whether the material is a solid or a liquid at a given point), the concentration (how much of 1 type of element is present at that location), and/or the orientation (in a crystalline material, what direction the crystal is pointing). This list is not all-inclusive, and may include many other parameters that play a role in whatever simulation is being conducted.

In addition to holding all the relevant simulation data within a single object, this class also includes utilities that allow it to be used very effectively in parallel systems. When an MMSP-based program is run with multiple ranks, the grid will automatically be divided amongst the ranks. Additionally, MMSP implements a "`ghostswap`" functionality for the Grid class, where neighboring cells across ranks are made available, in case gradients of the parameters stored within the grid must be computed. Because most material simulations do depend on these gradients, parallel programs must balance the computational load of computing the rate equations for various Grid parameters on every grid point, with the bandwidth required to communicate between nodes to resolve these ghost rows/columns. Therefore, programs that utilize this feature, like ours, are not embarrassingly parallel, despite seeming so due to the convenience MMSP provides for writing parallelized simulations.

2.2. Phase field model: Concentration

In the hybrid model, concentration evolves based on the following rate equation:

$$\frac{dC}{dt} = M_c(\nabla^2 \frac{\partial E}{\partial C} - \kappa_c \nabla^4 C) \tag{1}$$

In this equation, M_c represents the mobility of the solute, κ_c determines the magnitude of the interfacial energy (excess energy due to the gradient of C), and C is the concentration field, represented in the MMSP Grid. Lastly, E_v is the bulk free energy, which is given by the paper to be:

$$E_v = 0.3((C-0.25)^2 + (C-0.75)^2) \tag{2}$$

$$+ 0.5(C-0.05)^2(1-\phi) + 0.5(C-0.95)^2\phi \tag{3}$$

Where ϕ is the value of the phase (0 for α, 1 for β). This term is added so the free energy curve is different between phases.Using this expression, we may evaluate the partial derivative, and simplify it to the following equation:

$$\frac{\partial E_v}{\partial C} = 2.2C - 0.65 - 0.9\phi \tag{4}$$

The equilibrium is found where the derivative of the energy is zero. Therefore, when ϕ=0, the equilibrium is at $C = 0.2955$, and when $\phi = 1$, equilibrium is at $C = 0.7045$. Therefore, we expect to see values

of the concentration reasonably close to these values within the bulk of the grains.

To evaluate the rate equation across the entire grid, we first compute the laplacian of both concentration, and the partial derivative of energy w.r.t. concentration. This is done using a simple 1-2-1 stencil in both dimensions (as this is a 2-D simulation). Therefore, at some grid point (i,j), these are computed as:

$$\nabla^2 C_{i,j} = \frac{C_{i+1,j} + C_{i-1,j} + C_{i,j+1} + C_{i,j-1} - 4C_{i,j}}{\Delta x^2} \tag{5}$$

$$\tag{6}$$

A similar expression can be arrived at for the derivative of energy with respect to concentration. The value of Δx is specified within the grid class, equal to the distance between the neighboring grid points. After computing these laplacians, they are stored into a second auxiliary grid, which is then ghost-swapped. This is so that the values may be reused to compute the fourth gradient.

$$\nabla^4 C_{i,j} = \frac{1}{\Delta x^2} \left(\nabla^2 C_{i+1,j} \right. \tag{7}$$

$$\left. + \nabla^2 C_{i-1,j} + \nabla^2 C_{i,j+1} + \nabla^2 C_{i,j-1} - 4\nabla^2 C_{i,j} \right) \tag{8}$$

Finally, the value of the rate equation is computed for each grid point, then multiplied by the time-step, Δt, to determine the change in concentration after each step. So long as the time-step satisfies the Courant condition, the simulation will be stable in this explicit finite difference method.

2.3. Monte Carlo Potts Model : Grain growth and Phase Change

The Potts model is a statistical mechanical model that uses Monte Carlo method to evolve an ensemble of particles or sites defined on a lattice. These sites represent the microstructure of a material. Each lattice site assumes a spin, q_i that represents a given membership of a grain on the microstructure. The total energy of the system is given by the sum of the bulk energy, E_v, at each site and the neighboring lattice sites. The total energy is given by:

$$E_p = \sum_{i=1}^{N} \left(E_v(q_i) + \sum_{i=j}^{n} J(q_i, q_j) \right)$$

where N is the total number of lattice sites and n is the neighboring lattice site of the given ste. J is defined by:

$$J = \begin{cases} 1 & \text{for } q_i \neq q_j \\ 0 & \text{for } q_i = q_j \end{cases} \tag{9}$$

In serial, the Potts model algorithm is as below:

- Q possible grain IDs
- Randomly choose a site i.
- Randomly choose a new grain ID at site i.
- Compute the change in total system energy ΔE
- The probability of accepting the new orientation one is given by Metropolis transition function (eq 10)
- One Monte Carlo Step (MCS) is defined as N resetting grain ID attempts.

The probability of accepting the reorientation is given by:

$$p(\Delta x) = \begin{cases} 1 & \text{it } \Delta E \leq 0 \\ exp(-\Delta E/kT) & \text{if } \Delta E > 0 \end{cases} \tag{10}$$

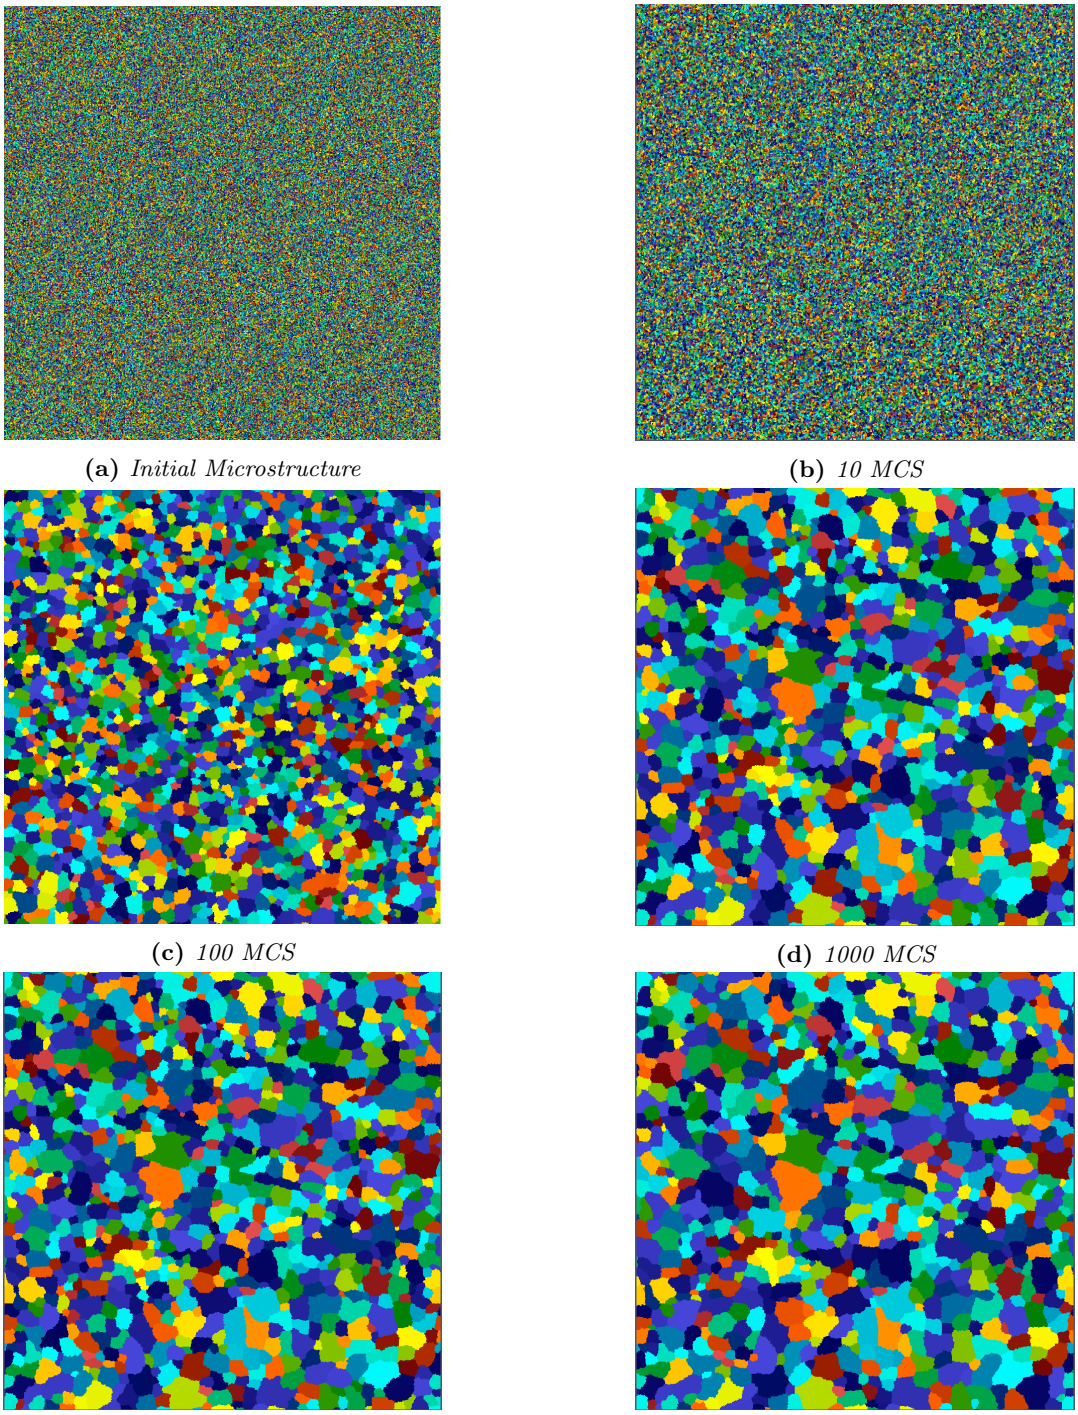

(a) *Initial Microstructure*

(b) *10 MCS*

(c) *100 MCS*

(d) *1000 MCS*

(e) *10000 MCS*

(f) *20000 MCS*

Figure 1: *Evolution of grain structure over 20000 Monte Carlo steps (MCS)*

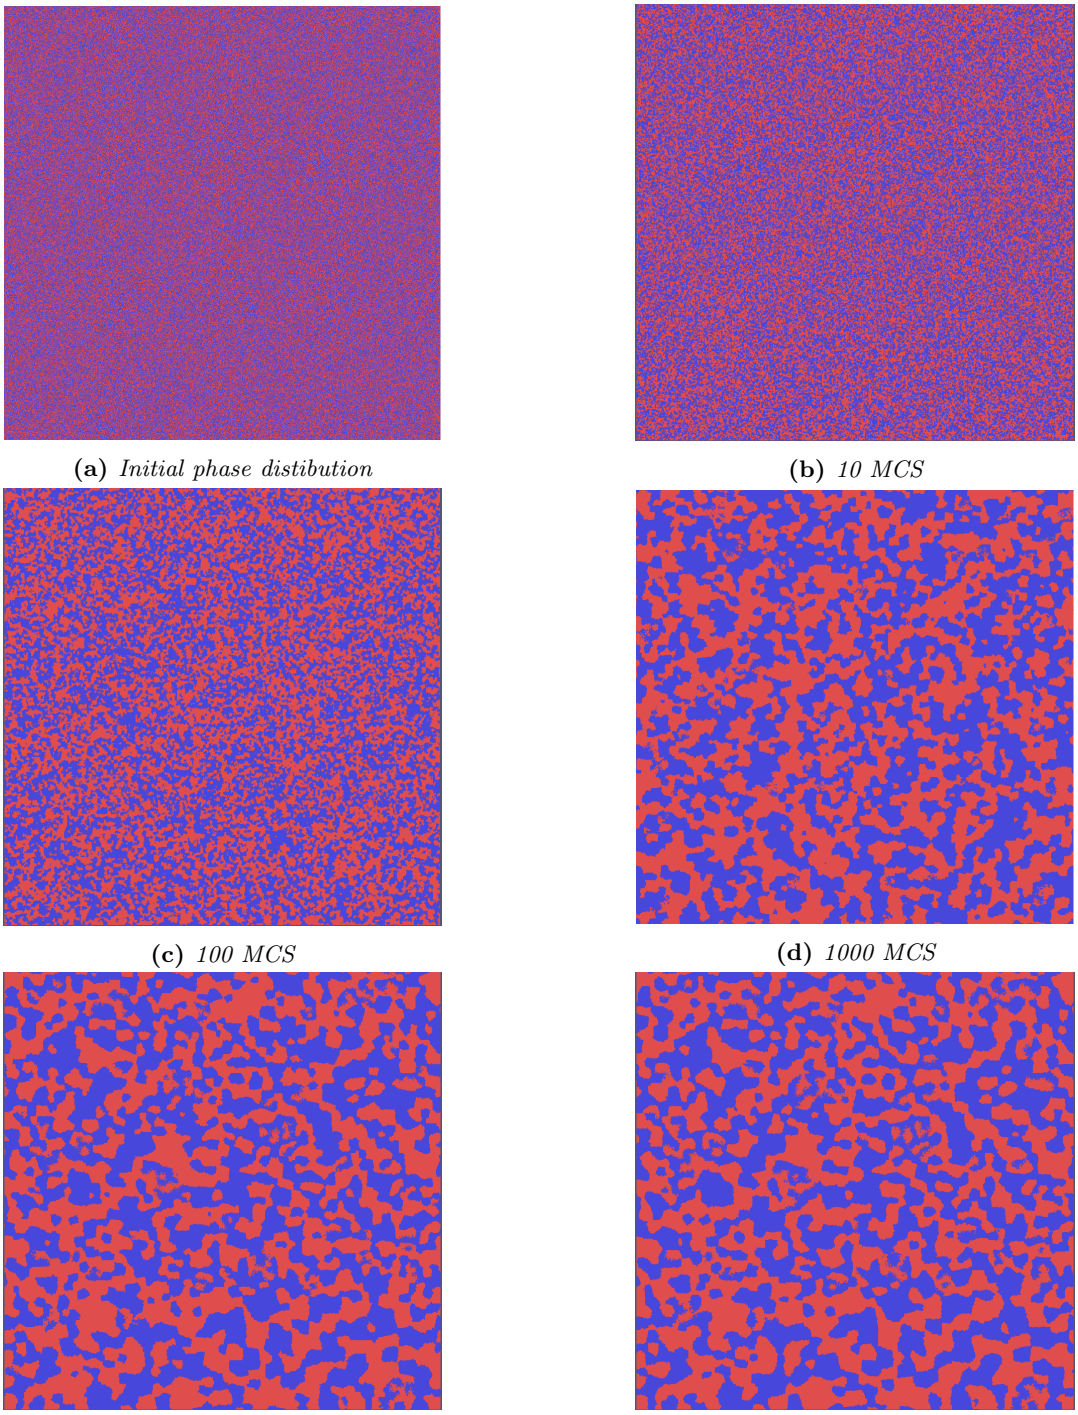

(a) *Initial phase distibution*

(b) *10 MCS*

(c) *100 MCS*

(d) *1000 MCS*

(e) *10000 MCS*

(f) *20000 MCS*

Figure 2: *Phase Change over 20000 MCS, α is denoted by blue and β is denoted by red*

2.4. Hybrid Model

As the team consists of two individuals with differing simulation experiences (Phase field for Scott, Monte Carlo for Sagar), the team will tackle a hybrid model from the literature that combines these two methodologies. Because of this, each member may work on the component of the model that is more closely related to their own research, allowing for a more efficient division of labor, where the experiences gained through this project may be more adequately applied to each member's research in the future.
This model is derived from the research paper "Hybrid Potts-phase field model for coupled microstructural-compositional evolution" by Eric Homer *et al.* [1] Within the model, the primary parameters of interest are concentration (which atoms are where), phase (whether the material is the "α" solid, or the "β" solid, and Grain ID (which original grain does a given grid point belong to). Therefore, the MMSP grid we use for this model will contain the value for these three parameters at every grid location.

Our methodology for implementing this model is to create an arbitrary initial grid, using Voronoi tessellation to split the grid into many grains, then simulating the evolution of the three parameters over time. The evolution of the phase and Grain ID will be discrete, using Monte Carlo method, while the concentration will continuously evolve between 0 and 1, using the phase field model.

The concentration derived at each lattice site using the phase-field model is used to compute the overall free energy which is use in Monte Carlo model to define the probability of accepting the change in phase and grain membership.

3. Results and Discussion

For the project, the goal was to simulate a 2-D grid, containing 1000x1000 points. This grid will be simulated at a varying number of time steps, to observe the changes in the various fields over time. Also, for the strong scaling test of the Blue Gene, the simulation will be for a constant 1000 time steps, with varying numbers of nodes/ranks. In total, the simulation was run for 1 to 128 nodes in powers of two (due to resource restrictions on 'small' partition of BG/Q, we were unable to evaluate the performance metrics on 64 nodes), with 64 ranks per node.

As can be seen in fig. 2, the size of the grains increase as the simulation progresses. Furthermore, the nature of this increase roughly follows the typical grain growth behavior: when the average grain size is small, grains grow quickly, but when the grains are large, this growth is greatly slowed. There is quite a large change in the average grain size between 0 and 10000 time steps, but ultimately very little change between 10000 and 20000.

The behavior of the grains can be roughly seen in the images for the change of the phase (fig. 2) over the simulation window. As each grain has a unique phase associated with it, as the grains grow, the regions of one continuous phase will also increase in a similar manner.

As MMSP get run with more and more ranks, the performance increases roughly linearly, as can be expected due to it being coded for high parallelization. However, as the number of ranks increases, the number of ghost rows similarly increases. Therefore, despite having to compute fewer grid points every step, the computer must shuffle more data between ranks to swap the ghost rows after each iteration. Therefore, the "normalized" performance decrease can be explained as seen in the performance graphs.

The performance metrics we have evaluated here include the time taken for simulation on the same initial grid size, the speedup and strong Scaling efficiency , given by:

$$\eta = \frac{p_1 \times T_1}{p_n \times T_n}$$

Where, p_i and T_i denotes number of processes and time taken respectively. The result of the strong scaling experiments are given in table 1.

Table 1: *Performance metrics for strong scaling experiments of the Hybrid model*

Number of MPI Ranks	Time Taken(s)	Speed Up	Strong Scaling Efficiency (%)
64	1055.104	1.0000	100
128	541.910	1.9470	97.35
256	276.632	3.8141	95.35
512	135.412	7.7918	97.39
1024	68.190	15.4729	96.71
2048	39.639	26.6175	83.18
8192	11.423	92.3612	72.16

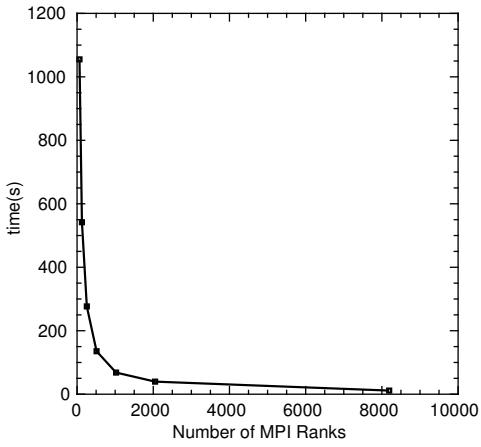

Figure 3: *Time taken for simulation over a 1000 × 1000 grid*

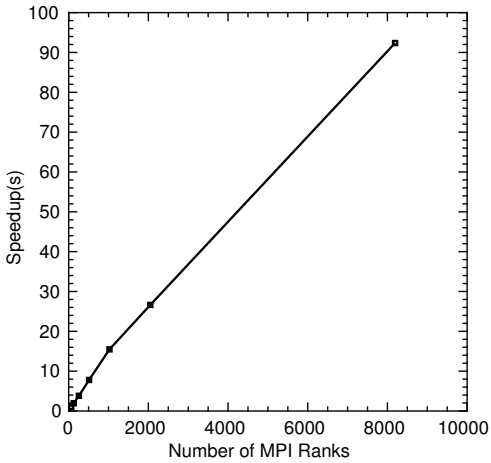

Figure 4: *Speed up*

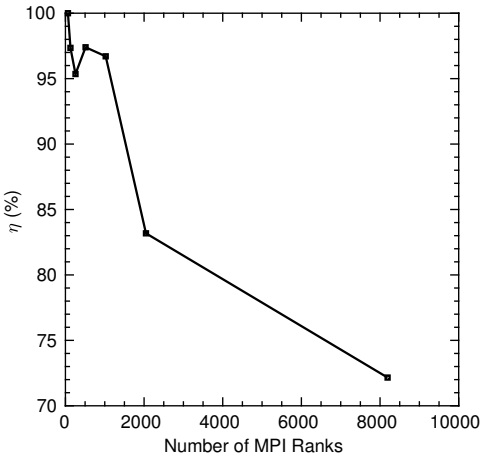

Figure 5: *Strong Scaling efficiency*

NOTE: The final code along with this report and README can be found on Sagar's kratos home folder: PPCbhatts8/ForExaminer_FinalCode/

References

[1] Eric R Homer, Veena Tikare, and Elizabeth A Holm. Hybrid potts-phase field model for coupled microstructural–compositional evolution. *Computational Materials Science*, 69:414–423, 2013.