Efficient techniques for bridging from atomic to mesoscopic scale in phase-field simulations

Marco Berghoff[*1], Michael Selzer[1,2], Abhik Choudhury[1] and Britta Nestler[1,2]

[1]Institute of Applied Materials, Karlsruhe Institute of Technology, Germany
[2]Institute of Materials and Processes, Karlsruhe University of Applied Sciences, Germany

August 29, 2013

Abstract

We show simulation techniques that enable to computationally bridge from the atomistic up to the mesoscopic scale. To evaluate the efficiency, we use a phase-field method to simulate symmetric thermal dendrites and start the computation at an early stage of solidification related to nucleation. The early structures are taken from the molecular dynamics method. The proposed techniques allow to simulate three dimensional dendrites from the state of nuclei ($\approx 50\,\text{Å}$) up to a size proportional to μm where ternary side-arms start to grow.

Keywords: efficient computational algorithms, phase-field, scale bridging, Hill tetrahedrons, dendrite growth

1 Introduction

In material science, a modern computational materials characterization and development requires multiscale simulation approaches in order to capture the physical processes acting on both scales, atomistic and mesoscopic length and time scales [1]. An understanding of the effects on each scale allows to transfer information, material data and physical properties across different, by now, well-established methods. Various attempts have recently been reported. Bragard

---

*Corresponding author: Karlsruhe Institute of Technology, Institute of Applied Materials - Reliability of Components and Systems, Haid-und-Neu-Str. 7, 76131 Karlsruhe, Germany, tel: +49 721 608 45314, fax: +49 721 925 2348, e-mail: marco.berghoff@kit.edu

et al. [2], for example, make use of parameters delivered by molecular dynamics (MD) simulations to compute Nickel dendrites on a micrometer length scale with the phase-field (PF) method. Another combined approach coupling cellular automaton model to a finite element approach is discussed by Gandin et al. [3]. Provatas et al. [4, 5] use a phase-field crystal model as an alternative concept on the atomistic length scale compared to MD to provide input data for PF simulations [6, 7].

To facilitate a computational description ranging from $1\,\text{Å}$ to several micrometers, improved numerical algorithms and solution schemes are necessary. The present paper introduces a number of efficiency optimized techniques as feasible solutions to be applied e. g. to dendritic morphologies. The methods chosen to establish a sequential multiscale chain combine MD and PF simulations in order to transfer information from a smaller scale to a bigger scale. As the presented methods rely on MD and PF simulations, we briefly recall recent review articles to provide a compressed overview. On the atomistic scale, MD is established with reliable results, however it has high computational effort such that big simulations are not very practical [8–12]. The PF method has emerged as a powerful computational approach to describe phase transition phenomena and microstructure formation with complex interfacial topologies typically occurring on a mesoscopic length scale of 1 to several hundred micrometers (e. g. [13–18]). Model derivatives are formulated to study solidification in pure and multicomponent systems e. g. [19–28], microstructure evolution in solid phase systems [29–31] and general motion of grain boundaries, e. g. [32].

Due to the complexity of the underlying physical model derivations, the individual methods operating on the different length scales are on its own computationally very extensive. The challenge of an integrated computational modelling for materials development and engineering with a multiscale information and data transfer essentially requires new techniques and algorithms to efficiently reduce computing time and memory. The aim of the present paper is to contribute to this challenge in developing numerical solution schemes based on symmetry arguments and rescaling procedures.

The considered simulation framework is based on an explicit finite difference scheme solved within a regular cubic grid. We choose the evolution of a typical dendritic solidification morphology to illustrate the necessity and computational advantages of particularly designed algorithms to bridge from atomistic to mesoscopic scales. We begin on the atomistic scale employing a MD simulation [33] and set a nucleus of $50\,\text{Å}$ in diameter in an undercooled melt. The nucleus grows in its surrounded melt to a large atom cluster. If solidification proceeds interfacial instabilities as a result of the limited temperature diffusion form and finally a dendritic shape evolves with a primary trunk, with secondary and ternary side-arms in the direction of the crystal anisotropy on a mesoscopic scale. The setup of a scale-bridging simulation needs a huge resolution of the numerical grid so that both, the initial nucleus can be resolved sufficiently fine as well as the morphology of the final dendrite can be contained in the simulation domain. The interface Péclet number (the ratio of the interface width and the diffusion length) has to be small for the thermodynamic consistency of the PF model

[34]. In detail, a resolution of 1 Å needs 20 000 cells in each direction to hold a dendrite of 2 μm length. For a 3D application, this rough estimation yields the need of $8 \cdot 10^{12}$ cells and, hence 58 TB for each field assuming double precision with 8 Byte per value. If we consider the simplest case of a thermal dendrite solidification in a pure substance, a PF model consists of two computed fields, the phase variable defining the phase state of the system and the temperature field. Based on the above considerations, 116 TB are to be stored and executed. Even with the employment of high performance clusters, the sketched simulation is, without the incorporation of further efficiency techniques, basically not practical as the equipment of high performance environments typically covers 2 GB per CPU. In total, the execution of the scale-bridging simulation of a thermal dendrite requires the access to 60 000 CPUs. We even remark that newer clusters commonly consist of 512 MB memory chips. Widely used techniques to reduce the computational effort employ adaptive mesh refinements which simultaneously reduce the memory demand if finite element structures are considered [1, 35, 36]. The random walker methods also allow to reduce the number of grid points in regions further away from the moving interfaces [37]. Another approach has been suggested in [32] and is based on a hybrid GPU–CPU parallelization.

In addition to these well-known techniques, we describe algorithms and iteration schemes - (i) which are particularly suited to improve the efficiency of the diffuse interface nature of PF models and - (ii) which particularly exploit the symmetry characteristics of crystal anisotropy. The phase field is only independent of its interface width within a bounded range, so that it is not arbitrary scaleable. We discuss a re-scaling technique to reduce the number of grid points in each direction from 20 000 cells to 4 000 cells. As a result, each field variable needs memory resources of 477 GB which is already in a range computable on high performance clusters. The rescaling of the grid resolution can be employed for different types of microstructure. To further improve the computational effort, we explore the cubic symmetry of the dendrite by applying a geometrical Hill tetrahedron construction to downsize the computational domain. Together with the trivial symmetry argument of considering only 1/8 of the domain, the Hill construction allows to reduce the eighth to an additional subdomain of 1/6. In total, the whole structure can well be resolved in 1/48 of the original domain which enables to halve the used number of CPUs and the computing time.

The structure of the paper is the following: In the next section, we introduce the PF model and define the material parameter obtained from MD simulations of pure Nickel [38, 39] that are transferred to the successive PF simulations. Section 3 builds the main body of the paper, in which we provide several simulation techniques that reduce the computational time and memory required to conduct the multiscale simulation. Among the presented algorithmic schemes, we exploit the symmetry of the dendrite and compute only 1/48 of the total structure. Further, we suggest an iterative up-scaling of the computational domain accommodated to the size of the crystal during the growth evolution. We apply the efficient techniques to implement a scale-bridging simulation from a nucleus of 5 Å up to a several micrometer large 3D structure that shows the typ-

ical morphological properties of dendrites. In section 4, we conclude by carving out the resulting speed up of the proposed methods.

2 Phase-field model for dendritic solidification in pure materials

As the central objective of this paper is to establish a general algorithm framework for scale-bridging computations from a nucleation state to a micrometer microstructure, we exemplarily elaborate a simple PF model for thermal dendritic solidification of pure materials.

The crystal growth in a pure material is modeled by the variables of the internal energy density e and of a phase-field $\phi(\vec{x}, t)$ with $0 \leq \phi \leq 1$ in the interfacial region and with $\phi = 1$ for the solid phase and $\phi = 0$ for the liquid phase. The variable $\phi(\vec{x}, t)$ denotes the local fraction of the solid phase at position $\vec{x}$ and at time t.

The PF model is based on an entropy functional to ensure consistency with classical irreversible thermodynamics

$$\mathcal{S}(e, \phi) = \int_{\Omega} s(e, \phi) - \left(\varepsilon a(\phi, \nabla \phi) + \frac{1}{\varepsilon} w(\phi) \right) \mathrm{d}x. \tag{1}$$

The bulk entropy density s depends on the internal energy density e and the phase-field variable $\phi(\vec{x}, t)$. The functions $a(\phi, \nabla \phi)$ and $w(\phi)$ reflect the thermodynamics of the interfaces and ε is a small length scale parameter related to the thickness of the diffuse interface.

The gradient entropy $a(\phi, \nabla \phi)$ incorporates a formulation of crystal anisotropy by the factor $(A(\hat{n}))^2$ depending on the orientation of the interface. We employ the expression

$$a(\phi, \nabla \phi) = a(\nabla \phi) = \frac{\gamma_0}{T} \left(A(\hat{n}) \right)^2 |\nabla \phi|^2 = \gamma(\hat{n}) |\nabla \phi|^2 \tag{2}$$

where $\hat{n} = (n_1, n_2, n_3)^T = -\nabla \phi / |\nabla \phi|$ is the normalized gradient vector of the solid–liquid interface.

The function $\omega(\phi) = \frac{16}{\pi^2} \frac{\gamma_0}{T} \phi(1 - \phi)$ is an obstacle type potential. We set $\omega(\phi) = \infty$ for $\phi \notin [0, 1]$. Note that $\frac{\gamma_0}{T}$ is the interface entropy density. The obstacle potential is chosen for numerical reasons allowing the phase-field equation to be solved only in a finite diffuse interface region.

From Eq. (1) one derives the equations for the non-conserved phase-field variable ϕ and for the energy conservations by taking the functional derivatives in the form

$$\tau \varepsilon \frac{\partial \phi}{\partial t} = \frac{\delta \mathcal{S}}{\delta \phi},$$
$$\frac{\partial e}{\partial t} = -\nabla \cdot L_{00}(T, \phi) \nabla \frac{\delta \mathcal{S}}{\delta e}, \tag{3}$$

where $\tau = \tau(\phi, \nabla\phi)$ is an anisotropic relaxation parameter depending on the temperature according to $\tau = \frac{L(T)}{TT_m k(\hat{n})}$. By Eq. (4), τ is a function of the kinetic coefficient k_0.

The mobility coefficient $L_{00}(T, \phi)$ is related to the thermal conductivity $K(\phi)$. We assume the simple form $L_{00} = KT^2$ for a phase independent constant $K(\phi) = K$.

The evolution of the phase-field variable is described by

$$\tau\varepsilon\frac{\partial\phi}{\partial t} = \varepsilon\left(\nabla\cdot\frac{\partial a(\nabla\phi)}{\partial\nabla\phi}\right) - \frac{1}{\varepsilon}\frac{\partial w(\phi)}{\partial\phi} + \frac{\partial s(e,\phi)}{\partial\phi}.$$

The thermodynamic relation $e = f + Ts$ yields $\frac{\partial s(e,\phi)}{\partial\phi} = -\frac{1}{T}\frac{\partial f(T,\phi)}{\partial\phi}$ (for details see [40]), so that the phase-field equation can be expressed in terms of the free energy density $f(T, \phi)$. We define

$$f(T, \phi) = L(T)\frac{T - T_m}{T_m}h(\phi) + \sum_\alpha\left(\int_{T_m}^T C_v^\alpha(\tilde{T})\mathrm{d}\tilde{T} - T\int_{T_m}^T C_v^\alpha(\tilde{T})\frac{\mathrm{d}\tilde{T}}{\tilde{T}}\right)h(\phi_\alpha)$$

where $h(\phi)$ is a monotonous function on $[0, 1]$ with $h(0) = 0$, $h(1) = 1$ and $h'(0) = h'(1) = 0$. We choose $h(\phi) = \phi^3(6\phi^2 - 15\phi + 10)$. Note $\alpha \in \{\text{solid}, \text{liquid}\}$ and $\phi_{\text{solid}} = \phi = 1 - \phi_{\text{liquid}}$. L is the latent heat and $C_v^\alpha(T)$ is the volumetric heat capacity of phase α, both depending on the temperature T. T_m is the melting temperature.

The evolution of the temperature is derived from the entropy density $s = -\frac{\partial f}{\partial T}$ and the thermodynamic relation

$$e(T, \phi) = f(T, \phi) - T\frac{\partial f(T, \phi)}{\partial T}$$

with T and ϕ both functions of time t. According to Eq. (3), the time derivative of e reads

$$\frac{\partial e}{\partial t} = \left(-T\frac{\partial^2 f}{\partial T^2}\right)\frac{\partial T}{\partial t} + \left(\frac{\partial f}{\partial\phi} - T\frac{\partial^2 f}{\partial T\partial\phi}\right)\frac{\partial\phi}{\partial t} = -\nabla\cdot L_{00}(T, \phi)\nabla\frac{1}{T}$$

with $\frac{\delta S}{\delta e} = \frac{1}{T}$ for pure substances. Rearrangement results in the evolution equation for the temperature field

$$\frac{\partial T}{\partial t} = \left(-\nabla\cdot L_{00}(T, \phi)\nabla\frac{1}{T} - \left(\frac{\partial f}{\partial\phi} - T\frac{\partial^2 f}{\partial T\partial\phi}\right)\frac{\partial\phi}{\partial t}\right)\cdot\left(-T\frac{\partial^2 f}{\partial T^2}\right)^{-1}.$$

We transfer the properties of Ni from molecular dynamics computations based on an embedded atom method (EAM) potential by Foiles (F85) [38] and thermophysical properties from [39]. According to the atomistic simulations, the density satisfies the function $\rho(T) = a + b\cdot T + c\cdot T^2$, with $a = 8901\,\mathrm{kg/m^3}$, $b = -0.2038\,\mathrm{kg/m^3 K}$ and $c = -6.142\cdot 10^{-5}\,\mathrm{kg/m^3 K^2}$ for the fcc solid phase and

$a = 8992 \,{}^{\text{kg}}/\text{m}^3$, $b = -0.667 \,{}^{\text{kg}}/\text{m}^3\text{K}$ and $c = 3.316 \cdot 10^{-5} \,{}^{\text{kg}}/\text{m}^3\text{K}^2$ for liquid. The latent heat is $L_0(T) = \tilde{L}_0(T) \cdot \rho(T)$ with $\tilde{L}_0(T) = a + b \cdot T + c \cdot T^2$ where $a = -1.598 \cdot 10^4 \,{}^{\text{J}}/\text{kg}$, $b = 324.7 \,{}^{\text{J}}/\text{kgK}$ and $c = -8.11 \cdot 10^{-2} \,{}^{\text{J}}/\text{kgK}^2$.

The specific heat capacity is given by $c(T) = \tilde{c}(T) \cdot \rho(T)$ with $\tilde{c}(T) = a + b \cdot T$ where $a = 392.1 \,{}^{\text{J}}/\text{kgK}$ and $b = 8.525 \cdot 10^{-2} \,{}^{\text{J}}/\text{kgK}^2$ for the solid and $a = 673.4 \,{}^{\text{J}}/\text{kgK}$ and $b = -3.126 \cdot 10^{-2} \,{}^{\text{J}}/\text{kgK}^2$ for the liquid.

The thermal diffusivity proposed by the MD model is calibrated to value $2.1 \cdot 10^{-7} \,{}^{\text{m}^2}/\text{s}$ and differs from the experimental value $170 \cdot 10^{-7} \,{}^{\text{m}^2}/\text{s}$ by two orders of magnitude [39]. The melting temperature is $T_m = 1748 \,\text{K}$.

The interfacial energy γ depending on the orientation $\hat{n}$ is represented by a harmonic cubic expansion

$$\frac{\gamma(\hat{n})}{\gamma_0} = 1 + \varepsilon_1 \left(Q - \frac{3}{5} \right) + \varepsilon_2 \left(3Q + 66S - \frac{17}{7} \right) + \varepsilon_3 \left(5Q^2 - 16S - \frac{94}{13}Q + \frac{85}{13} \right)$$

with $Q = n_1^4 + n_2^4 + n_3^4$ and $S = n_1^2 n_2^2 n_3^2$. The mean interfacial free energy averaged over the orientations is $\gamma_0 = 0.302 \,{}^{\text{J}}/\text{m}^2$ and the anisotropy strengths are $\varepsilon_1 = 0.102$, $\varepsilon_2 = -1.34 \cdot 10^{-3}$ and $\varepsilon_3 = 8.76 \cdot 10^{-3}$.

The orientation-dependent kinetic coefficient is represented by

$$\frac{k(\hat{n})}{k_0} = 1 - 3\varepsilon_k + 4Q + \delta_k \left(R + 30S \right) \tag{4}$$

with $R = n_1^6 + n_2^6 + n_3^6$. The mean kinetic coefficient averaged over the orientations is $k_0 = 0.3192 \,{}^{\text{m}}/\text{sK}$ and the strengths of the anisotropy are $\varepsilon_k = -0.1965$ and $\delta_k = 0.2303$.

3   Techniques to optimize computational efficiency

We solve the phase-field and temperature equation using an explicit finite differences scheme on a regular cubic grid. To initialize the simulation domain, a crystalline cluster of about $5 \,\text{nm}$ in diameter is appropriately converted from molecular dynamics. Therefore we calculate the $q_6 q_6$ order parameter [41] and map the atom positions to a mesh of $1 \,\text{Å}$ width. The converted cluster is used as the starting condition in the PF simulations. Fig. 1 shows the discrete atom positions of MD and the converted continuous field of PF. However the phase-field evolution smoothens the cluster approximating a sphere. We benefit from this by assuming a symmetric cluster and simulating only one eighth of the cluster. To exploit this property, the cluster is set in one corner of a simulation domain called *reduced cube* (see Fig. 2 upper left) and the borders act as reflection planes. For each border we use Neumann boundary conditions ensuring the symmetry on the reflection planes. The temperature field evolves very quickly. To avoid boundary effects of the temperature, we set the simulation box twice the size needed for the resulting crystal.

To establish a multiscale approach, we develop different techniques to reduce the computational time. The algorithms are described in the following sections.

3.1 Iterative up-scaling method

As a first method, we introduce an iterative up-scaling method. The simulation is conducted in a small simulation box with 145^3 cells of size $1\,\text{Å}$. We use the thermal diffusivity of $2.1 \cdot 10^{-7\,\text{m}^2/\text{s}}$ from MD to benefit from a slower temperature field. The temperature field is initialized with $1450\,\text{K}$. The dendrite starts growing and before boundary effects occur, we interrupt the simulation. The grid size is multiplied by a factor 2. The dendrite is scaled to this size so that it fits in a 72^3 box, the remaining part of the simulating box is filled with the initial state. We use a trilinear interpolation to scale the domain. The physical cell width is doubled, the physical size of the evolved nucleus remains unchanged. Following, that we continue the simulation with a time step size chosen 4 times larger. This algorithmic scheme is repeated 4 times in total, resulting in a physical cell width of $16\,\text{Å}$ in the final simulation sequence. The time step size of the last simulation period is 256 times larger than in the first simulation step.

Fig. 2 shows the simulation starting from an MD nucleus and evolving to a full mesoscopic dendrite. In the different up-scaling steps, we observe that the surface of the nucleus changes from a convex curvature to a concave shape. The dendrite approaches $1\,\mu\text{m}$ in size, which is small for a pure nickel dendrite, but the simulation contains all parameters from MD including a thermal diffusivity of $2.1 \cdot 10^{-7\,\text{m}^2/\text{s}}$.

In this procedure, we also change the value of the diffuse interface thickness ε, so that we have an interface of about 8 cells for each simulation part. This implies a change of the physical interface width, which has little effect on the growth rate. As tested in 1D simulations for varying numbers of ε, the influence is small i. e. approximately $1.5\,\%$ for the chosen values of ε (ε, 2ε and 4ε) for early state solidification. For an ε which is 8 times larger than the physical interface, the growth is less than $3\,\%$ faster, for ε 16 times than the size of the interface, the effect is still less than $7\,\%$. This deviation results in a non-symmetric interface so that large values $\Delta x > 4\,\text{Å}$ lead to a non-expected behavior. The range of a symmetric interface depends on the driving force and hence on the velocity of growing, so that for another undercooling this method may produce good results.

The up-scaling is very fast because of the small simulation box, but, as discussed, limited by the largest Δx providing a fine interface profile.

3.2 Symmetry arguments

In the previous section we have shown that the thermal diffusivity taken from MD provides too small dendrites. If we use a higher thermal diffusivity, e. g. from experiments, the temperature evolves faster so that we need bigger domains to avoid boundary effects. To further increase the efficiency of computation and to enable the consideration of large mesoscopic domains, we derive symmetry arguments related to the symmetry of the crystal structure.

We explore the crystal symmetry of the nickel dendrites in order to reduce

the computational domain to a representative volume element in accordance with the structure requirements. Random noise acts only to the considered part of the dendrite so that the resulting reconstructed dendrite is perfectly symmetric. If non-symmetric dendrite arm formations are to be discovered the following method is not appropriate.

3.2.1 Hill tetrahedron construction

A cube $Q := \big\{(x, y, z) \in \mathbb{R}^3 \mathrel{\big|} 0 \leq x, y, z \leq 1\big\}$ can be split in 6 symmetric parts by cutting Q along the three bisection planes: $(\bar{1}10)$, $(\bar{1}01)$ and $(0\bar{1}1)$. This results in six Hill tetrahedrons described by the sets

$$\mathcal{H}_1 := \{(x, y, z) \in \mathbb{R}^3 \mid 0 \leq x \leq y \leq z \leq 1\},$$
$$\mathcal{H}_2 := \{(x, y, z) \in \mathbb{R}^3 \mid 0 \leq x \leq z \leq y \leq 1\},$$
$$\mathcal{H}_3 := \{(x, y, z) \in \mathbb{R}^3 \mid 0 \leq z \leq x \leq y \leq 1\},$$
$$\mathcal{H}_4 := \{(x, y, z) \in \mathbb{R}^3 \mid 0 \leq y \leq x \leq z \leq 1\},$$
$$\mathcal{H}_5 := \{(x, y, z) \in \mathbb{R}^3 \mid 0 \leq z \leq y \leq x \leq 1\},$$
$$\mathcal{H}_6 := \{(x, y, z) \in \mathbb{R}^3 \mid 0 \leq y \leq z \leq x \leq 1\}.$$

The union of these sets is the cube $Q = \bigcup_{n=1}^{6} \mathcal{H}_n$. Fig. 3 shows the cutting planes of the set Q. Every cutting plane divides the original cube in two equal parts, where each part holds three Hill tetrahedrons. A Hill tetrahedron consists of four triangles. The two non-visible triangles in Fig. 3(d) are isosceles.

We now identify the reduced cube with Q and place a nucleus into the origin of Q, so that the cube Q contains $^1\!/_8$ of the nucleus. It is obvious that the three cutting planes are projection planes for the symmetry of the nucleus. Each Hill tetrahedron contains $^1\!/_6$ of the nucleus in the reduced cube Q, hence, $\mathcal{H}_1$ contains $^1\!/_{48}$ of the nucleus. The nucleus grows with a dendrite tip in the z-direction. $^1\!/_8$ of the arm is contained in $\mathcal{H}_1$. The projection plane $(\bar{1}10)$, together with the xz- and yz-planes project the section back to the whole arm in z-direction. The arms in x- and y-direction are projected by the $(\bar{1}01)$ and $(0\bar{1}1)$ planes, respectively.

On account of the symmetry arguments, only one of the Hill tetrahedrons is necessary for the calculation. We use the discrete analogue of the Hill tetrahedron $\mathcal{H}_1$ to define the simulation domain by

$$H_1 := \big\{(i, j, k) \in \mathbb{N}^3 \mathrel{\big|} 0 \leq i \leq j \leq k \leq N\big\},$$

with N being the number of cells in one direction. Analogously $H_2, \ldots, H_6$ are the discrete set of the Hill tetrahedrons $\mathcal{H}_2, \ldots, \mathcal{H}_6$. To reconstruct the whole

cubic domain, we need to map H_1 to $H_2, \ldots, H_6$ by the mapping functions

$$A_2 : H_1 \to H_2, (i, j, k) \mapsto (i, k, j),$$
$$A_3 : H_1 \to H_3, (i, j, k) \mapsto (k, i, j),$$
$$A_4 : H_1 \to H_4, (i, j, k) \mapsto (j, i, k),$$
$$A_5 : H_1 \to H_5, (i, j, k) \mapsto (k, j, i),$$
$$A_6 : H_1 \to H_6, (i, j, k) \mapsto (j, k, i).$$

To numerically solve the evolution equations for the phase-field and the temperature, we calculate finite differences in an equidistant mesh. The discretisation contains a stencil incorporating the direct and diagonal neighbours. In the case of border cells, where neighbouring cells do not exist, a one-sided discretisation is replaced. To avoid discretisation errors due to the use of the one sided discretisation schemes we append one cell in each direction of H_1. With this construction, the inner of H_1 is calculated with the full discretisation stencil ensuring that the calculation in H_1 is the same as in a whole cubic domain. The inner domain border is given by the set

$$B := \left\{ (i, j, k) \in \mathbb{N}^3 \mid k \leq N : (i \leq j \wedge j = k + 1) \vee (i = j + 1 \wedge j \leq k) \right.$$
$$\left. \vee (i \leq j \wedge j = k + 2) \vee (i = j + 2 \wedge j \leq k) \vee (i = j + 1 \wedge j = k + 1) \right\}.$$

All cells not covered in $H_1 \cap B$ are marked as not to be calculated. For the cells in B, a suitable boundary condition needs to be formulated. The cells are updated by the mapping functions $A_2, \ldots, A_6$ after each simulation step. Consequently, the cells inside H_1 have the same neighbouring cells as in the full cubic domain. Geometrical details of the inner border B and of the data exchanging mechanism for boundary elements in a parallel simulation environment are displayed in Fig. 4.

A test simulation of a 73^3 cubic domain on a single core shows a speedup of around 6.5 more than the theoretically expected maximum of 6. The speedup of parallel simulations is listed in Table 1.

3.2.2 Reduced domain size

To initialize the simulation domain as a rectangular cuboid, we allocate the memory of the whole reduced cube, although only 1/6 of the domain is calculated. As seen in Fig. 5(a), the dendrite tip fills a small part of the Hill tetrahedron, which allows to cut the Hill tetrahedron by a plane parallel to the xz-plane at $(0, \frac{1}{2}, 0)$. This plane divides the Hill tetrahedron into two identical halfs defined by

$$\mathcal{H}_1^R := \left\{ (x, y, z) \in \mathcal{H}_1 \,\middle|\, y \leq \frac{1}{2} \right\}$$
$$= \left\{ (x, y, z) \in \mathbb{R}^3 \,\middle|\, 0 \leq z \leq 1 \wedge 0 \leq x \leq y \leq \min \left\{ \frac{1}{2}, z \right\} \right\}.$$

$\mathcal{H}_1^R$ fits in a rectangular cuboid that is one quarter of the cube Q. The cutting plane requires a boundary condition. We assume the same as for the other outer boundaries. As soon as one of the fields reaches the boundary, the result is influenced by the boundary condition. The diffusion of the fastest field, in this physical setting the temperature, has approximately the same expansion as the dendrite surface, if we use infinite relaxation kinetics as described in appendix A. Due to this property, the cutting plane is not reached before the tip approaches the boundary.

To simulate a dendrite with an extension of 4000 cells from tip to tip, the efficiency optimization schemes can be applied to not resolve the structure in the complete domain of 4000^3 cells. As discussed, a rectangular cuboid of $2000 \times 1000 \times 1000$ cells is fully sufficient to resolve the symmetric part of the dendrite. For volumetric data with float precision (4 Byte), the simulated box in the reduced cuboid occupies 7.45 GB for one saved timestep. In contrast, a box containing the full dendrite needs 32 times as much space, summing up to 238 GB.

3.3 Large-scale dendrite simulation

In continuation, we compute a large-scale 3D dendrite of $2\,\mu$m tip to tip size using an experimentally measured value of $170 \cdot 10^{-7}\,\mathrm{m^2/s}$ [39] for the thermal diffusivity. The evolution is shown in Fig. 6. As initial structure, we continue the growth of the rescaled dendrite in Fig. 2 as the third up-scaling step yielding a cell size $\Delta x = 5\,\text{Å}$. The simulation conducted using the symmetry optimization strategiesis described in Section 3.2, so that only $1/48$ of the dendrite is calculated.

We also apply the half Hill tetrahedron scheme so that a box of $1000 \times 1000 \times 2000$ is simulated. We employ the largest value $\Delta x = 5\,\text{Å}$ establishing a symmetric interface and apply the infinite relaxation kinetics from appendix A. The resulting dendrite is set in a 4000^3 cubic box and shows the formation of secondary and early state ternary arms.

4   Discussion and conclusion

Applying the efficiency optimization techniques of up-scaling and symmetry arguments allows to simulate huge dendrites. The up-scaling iteration scheme keeps the size of the domain and the number of grid points constant by increasing, at the same time, the physical length of the considered structure. The computational effort remains the same and does not increase with the growing structure. Scaling the length by factor 2 results in a timestep of factor 4. Similarly, a length factor of 16, as used in the last step of Fig. 2, results in a timestep of factor 256. The theoretical speedup of this method is $> 10^6$, in comparison to a domain with 16 times as many cells in each direction, and with much smaller timestep. We remark that the need of more memory and thus more CPUs has

not been considered in the current discussion. The PF method bounds the grid size so that the up-scaling method is only useful for the first steps.

The symmetry method reduces the computational effort without changing the grid size. However, the cubic domain is the same as the reduced cube but only around $1/6$ of the cells are calculated. As long as a parallelization by domain decomposition in 1 selected dimension is used, the speedup is limited by the domain size and by the number of CPUs. The maximum number of CPUs for a 145^3 cube is 145. Taking advantage of the Hill tetrahedron, a simulation exceeding over the same time only needs the elaboration of 28 CPUs. The usage of 63 CPUs allows to run simulations in half time, under the assumption that the computational load is distributed homogeneously. This leads to a speedup of > 4 by the additional symmetry without considering the trivial symmetry of the reduced cube.

The Hill tetrahedrons reduce the box size by a factor of 48 in total or by an additional factor of 6 if the trivial symmetry of the reduced cube is taken into account.

In addition, the memory of the symmetric part in the half Hill tetrahedron of the dendrite is only $1/32$. This can be reduced to $1/96$, if the non-calculated cells marked in Fig. 4 are not saved. The symmetry algorithms require a data structure that can no longer be linear. If memory problems arise, this reduction is possible.

Using the up-scaling method at the early stage of growth and continuing the simulation with the half Hill tetrahedron, allows to compute thermal dendrites of a mesoscopic scale where ternary side-arms are formed starting from a small atomic cluster. The presented techniques to improve the computational efficiency provide new insights into the physics of microstructure formation by establishing a simulation framework to conduct high resolution computations in 3D.

5   Acknowledgements

A   Infinite relaxation kinetics

In order to have infinite relaxation kinetics for the phase-field equation, we adapt the thin-interface analysis to pure material solidification, similar to the theory for alloys and under the assumption of an obstacle potential as presented in [28]. The results reflect previous work [24, 42] based on double well potentials. The

corresponding relation for the relaxation coefficient can be written as,

$$\tau = \varepsilon \frac{L^2}{TT_mK}\left(M + F\right),$$

with solvability integrals M and F, while the preceding relation ensures that the kinetic coefficient $\beta = 1/k_0 = 0$ for an isotropic interface. In case of solidification with anisotropic interfacial energy, a procedure similar to the derivations proposed by Karma et al. [43] can be formulated as,

$$\tau\left(\hat{n}\right) = \varepsilon A^2\left(\hat{n}\right)\frac{L^2}{TT_mK}\left(M + F\right),$$

where $A\left(\hat{n}\right)$ is the same as in Eq. (2). The quantitative aspects of the model have been verified by comparison of the tip velocity of a 2D dendrite for different interfacial widths. Dimensionless simulations (using length $\ell = 1\,\text{Å}$, time $\tau = 1\,\text{ps}$, temperature $\vartheta = 1748\,\text{K}$ and energy density $e = 6.6485\,\text{J/m}^3$) for Δx of $40, 50, 60, 70$ result in velocities of the tip $3.7370 \cdot 10^{-7}$, $3.7291 \cdot 10^{-7}$, $3.7276 \cdot 10^{-7}$ and $3.7332 \cdot 10^{-7}$ respectively. These values deviate by less than $0.3\,\%$ from the heoretical velocity.

References

[1] N. Provatas, M. Greenwood, B. Athreya, N. Goldenfeld, and J. Dantzig. Multiscale modeling of solidification: Phase-field methods to adaptive mesh refinement. *International Journal of Modern Physics B*, 19(31):4525–4566, 2005.

[2] Jean Bragard, Alain Karma, YoungyihH. Lee, and Mathis Plapp. Linking phase-field and atomistic simulations to model dendritic solidification in highly undercooled melts. *Interface Science*, 10:121–136, 2002.

[3] Ch-A Gandin and Michel Rappaz. A coupled finite element-cellular automaton model for the prediction of dendritic grain structures in solidification processes. *Acta metallurgica et materialia*, 42(7):2233–2246, 1994.

[4] N. Provatas, JA Dantzig, B. Athreya, P. Chan, P. Stefanovic, N. Goldenfeld, and KR Elder. Using the phase-field crystal method in the multi-scale modeling of microstructure evolution. *JOM Journal of the Minerals, Metals and Materials Society*, 59(7):83–90, 2007.

[5] PF Tupper and M. Grant. Phase field crystals as a coarse-graining in time of molecular dynamics. *EPL (Europhysics Letters)*, 81:40007, 2008.

[6] JJ Hoyt, B Sadigh, M Asta, and SM Foiles. Kinetic phase field parameters for the cu–ni system derived from atomistic computations. *Acta materialia*, 47(11):3181–3187, 1999.

[7] J.J. Hoyt, M. Asta, and A. Karma. Atomistic and continuum modeling of dendritic solidification. *Materials Science and Engineering R*, 41:121–163, 2003.

[8] K. Kadau, T.C. Germann, and P.S. Lomdahl. Molecular dynamics comes of age:. 320 billion atom simulation on bluegene/l. *International Journal of Modern Physics C*, 17:1755–1761, 2006.

[9] Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications.* Academic Press, 2002.

[10] Michael P Allen and Dominic J Tildesley. *Computer simulation of liquids.* Oxford university press, 1989.

[11] Andrew R. Leach. *Molecular modelling: principles and applications.* Pearson Education, 2001.

[12] Yu Mishin, M Asta, and Ju Li. Atomistic modeling of interfaces and their impact on microstructure and properties. *Acta Materialia*, 58(4):1117–1151, 2010.

[13] L.Q. Chen. Phase-field models for microstructure evolution. *Annual Review of Materials Research*, 32(1):113–140, 2002.

[14] WJ Boettinger, JA Warren, C. Beckermann, and A. Karma. Phase-field simulation of solidification. *Annual Review of Materials Research*, 32(1):163–194, 2002.

[15] Ingo Steinbach. Phase-field models in materials science. *Modelling and Simulation in Materials Science and Engineering*, 17(7):073001, 2009.

[16] Ryo Kobayashi. Modeling and numerical simulations of dendritic crystal growth. *Physica D: Nonlinear Phenomena*, 63(3):410–423, 1993.

[17] AA Wheeler, BT Murray, and RJ Schaefer. Computation of dendrites using a phase field model. *Physica D: Nonlinear Phenomena*, 66(1):243–262, 1993.

[18] James A Warren and William J Boettinger. Prediction of dendritic growth and microsegregation patterns in a binary alloy using the phase-field method. *Acta Metallurgica et Materialia*, 43(2):689–703, 1995.

[19] A. Karma and W.J. Rappel. Quantitative phase-field modeling of dendritic growth in two and three dimensions. *Physical Review E*, 57(4):4323, 1998.

[20] B. Nestler, H. Garcke, and B. Stinner. Multicomponent alloy solidification: Phase-field modeling and simulations. *Physical Review E*, 71(4):041609, 2005.

[21] A. Karma. Phase-field model of eutectic growth. *Physical Review E*, 49:2245–2250, 1994.

[22] B. Nestler and AA Wheeler. A multi-phase-field model of eutectic and peritectic alloys: numerical simulation of growth structures. *Physica D: Nonlinear Phenomena*, 138(1-2):114–133, 2000.

[23] B. Nestler, AA Wheeler, L. Ratke, and C. Stöcker. Phase-field model for solidification of a monotectic alloy with convection. *Physica D: Nonlinear Phenomena*, 141(1-2):133–154, 2000.

[24] B. Echebarria, R. Folch, A. Karma, and M. Plapp. Quantitative phase-field model of alloy solidification. *Physical Review E*, 70(6):061604, 2004.

[25] R Folch and M Plapp. Quantitative phase-field modeling of two-phase growth. *Physical Review E*, 72(1):011602, 2005.

[26] J Eiken, B Böttger, and I Steinbach. Multiphase-field approach for multicomponent alloys with extrapolation scheme for numerical application. *Physical Review E*, 73(6):066122, 2006.

[27] Mathis Plapp. Unified derivation of phase-field models for alloy solidification from a grand-potential functional. *Physical Review E*, 84(3):031601, 2011.

[28] A. Choudhury and B. Nestler. Grand-potential formulation for multicomponent phase transformations combined with thin-interface asymptotics of the double-obstacle potential. *Physical Review E*, 85(2):021602, 2012.

[29] Nele Moelans, Bart Blanpain, and Patrick Wollants. An introduction to phase-field modeling of microstructure evolution. *Calphad*, 32(2):268–294, 2008.

[30] Tomohiro Takaki, Y Hisakuni, T Hirouchi, Akinori Yamanaka, and Yoshihiro Tomita. Multi-phase-field simulations for dynamic recrystallization. *Computational Materials Science*, 45(4):881–888, 2009.

[31] T Takaki and Y Tomita. Static recrystalization simulations starting from predicted deformation microstructure by coupling multi-phase-field method and finite element method based on crystal plasticity. *International Journal of Mechanical Sciences*, 52(2):320–328, 2010.

[32] L. Gránásy, T. Pusztai, and J.A. Warren. Modelling polycrystalline solidification using phase field theory. *Journal of Physics: Condensed matter*, 16:R1205, 2004.

[33] The MD simulation data are provided by J. Horbach and R. Rozas.

[34] Yung-Tae Kim, Nikolas Provatas, Nigel Goldenfeld, and Jonathan Dantzig. Universal dynamics of phase-field models for dendritic growth. *Physical Review E*, 59(3), 1999.

[35] N. Provatas, N. Goldenfeld, and J. Dantzig. Adaptive mesh refinement computation of solidification microstructures using dynamic data structures. *Journal of Computational Physics*, 148(1):265–290, 1999.

[36] Tomohiro Takaki, Toshimichi Fukuoka, and Yoshihiro Tomita. Phase-field simulation during directional solidification of a binary alloy using adaptive finite element method. *Journal of crystal growth*, 283(1):263–278, 2005.

[37] Mathis Plapp and Alain Karma. Multiscale random-walk algorithm for simulating interfacial pattern formation. *Physical Review Letters*, 84(8), 2000.

[38] SM Foiles. Application of the embedded-atom method to liquid transition metals. *Physical Review B*, 32(6):3409, 1985.

[39] K.C. Mills. *Recommended values of thermophysical properties for selected commercial alloys*. Woodhead Publishing, 2002.

[40] H. Garcke, B. Nestler, and B. Stinner. A diffuse interface model for alloys with multiple components and phases. *SIAM Journal on Applied Mathematics*, pages 775–799, 2004.

[41] T. Zykova-Timan, R.E. Rozas, J. Horbach, and K. Binder. Computer simulation studies of finite-size broadening of solid–liquid interfaces: from hard spheres to nickel. *Journal of Physics: Condensed Matter*, 21:464102, 2009.

[42] A. Karma. Phase-field formulation for quantitative modeling of alloy solidification. *Physical review letters*, 87(11):115701, 2001.

[43] A. Karma and W.J. Rappel. Phase-field method for computationally efficient modeling of solidification with arbitrary interface kinetics. *Physical Review E*, 53(4):3017–3020, 1996.

	Hc3/1 node	Hc3/2 nodes	Imp/1 node	Imp/2 nodes	Workstation
Hill tetrahedron	274 s	153 s	507 s	333 s	544 s
full cube	1568 s	498 s	1835 s	1022 s	2840 s
speedup	5.7	3.2	3.6	3.1	5.3

Table 1: Runtime of a MPI parallel test simulation with a 72^3 cubic domain. The different clusters in the header are HP XC3000 at Steinbuch Centre for Computing of Karlsruher Institute of Technologies (Hc3) with 2 Quad-Core Intel Xeon E5540 (2.53 GHz) and 24 GB RAM per node and the Institute of Materials and Processes Cluster at University of Applied Science Karlsruhe (imp) with 2 Quad-Core AMD Opteron 2350 (2 GHz) and 16 GB per node. And a Workstation with Dual-Core Intel E8400 (3 GHz) and 4 GB running two workers.

List of Figures
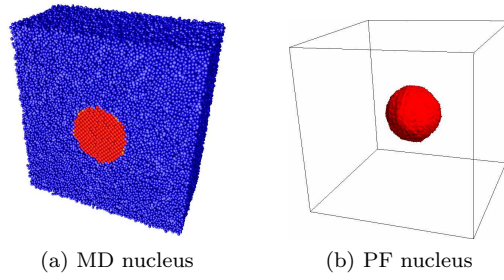
(a) MD nucleus   (b) PF nucleus

Figure 1: (a) Atomic positions of the MD nucleus, the color indicating the function of the order parameter, (b) converted PF nucleus displayed as an isosurface of the interface.
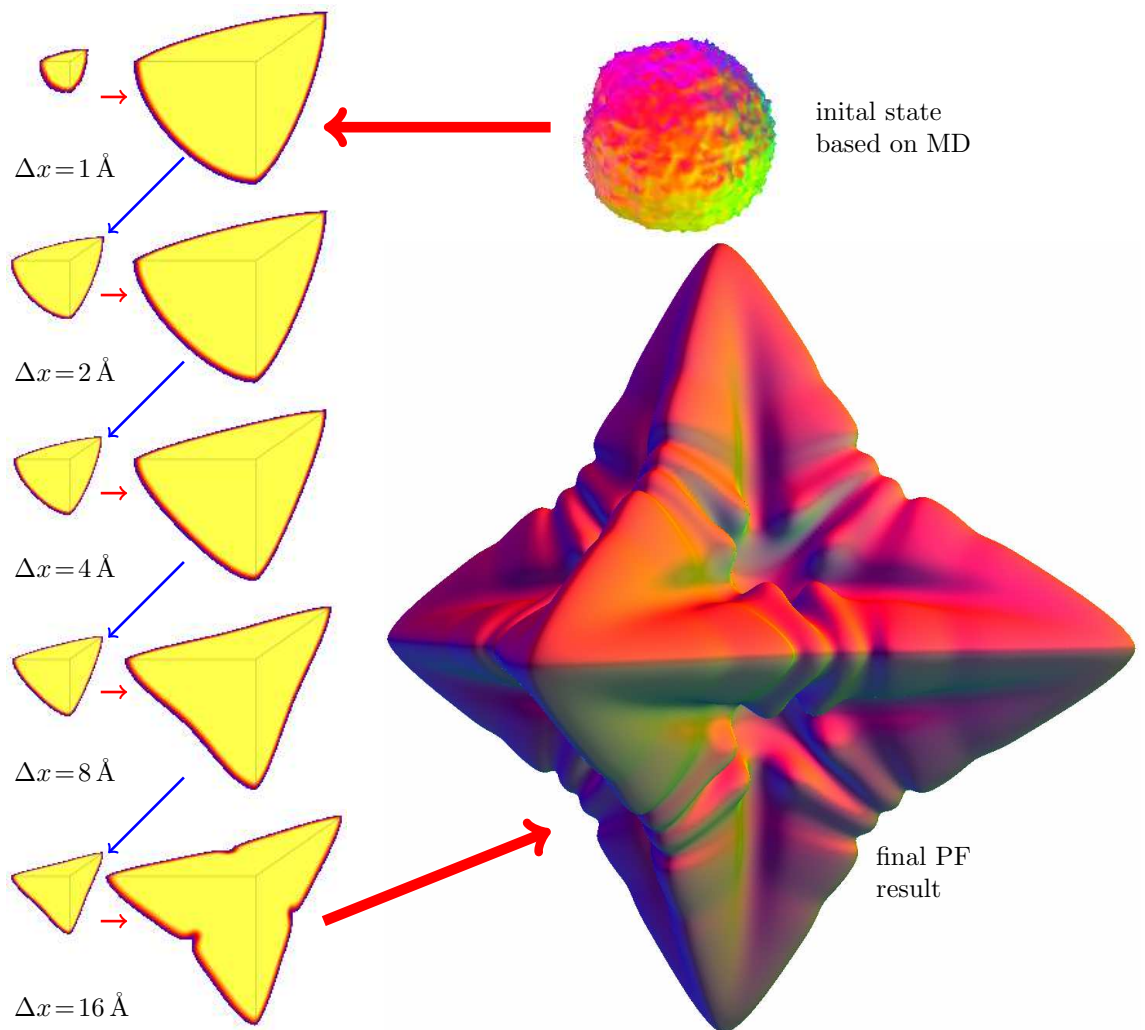
$\Delta x = 1\,\text{Å}$

$\Delta x = 2\,\text{Å}$

$\Delta x = 4\,\text{Å}$

$\Delta x = 8\,\text{Å}$

$\Delta x = 16\,\text{Å}$

inital state
based on MD

final PF
result

Figure 2: Phase-field simulation of a large dendrite (lower right) starting from a small MD nucleus (upper right) with four iterative up-scaling steps on the left side (red growing, blue modifying). The resulting dendrite has a tip to tip distance of $1\,\mu$m.

19

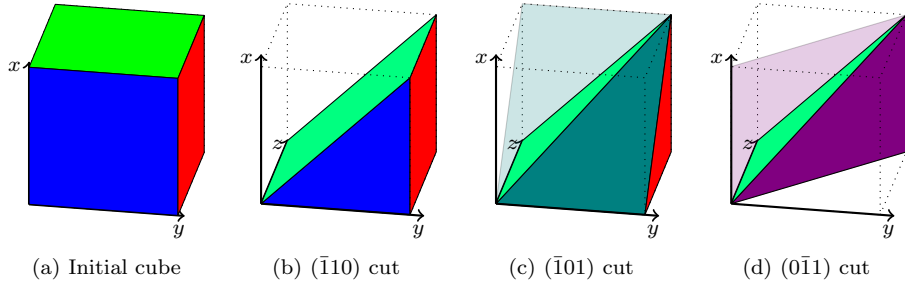(a) Initial cube   (b) ($\bar{1}$10) cut   (c) ($\bar{1}$01) cut   (d) (0$\bar{1}$1) cut

Figure 3: The initial unit cube Q in (a) is cutted by the ($\bar{1}$10), ($\bar{1}$01) and (0$\bar{1}$1) plane. The result of the first cut (b) is equivalent to the set $A_1 := \{(x, y, z) \in Q \mid x \leq y\} = \mathcal{H}_1 \cup \mathcal{H}_2 \cup \mathcal{H}_3$. The second cut (c) is equal to the set $A_2 := \{(x, y, z) \in A_1 \mid x \leq z\} = \mathcal{H}_1 \cup \mathcal{H}_2$. The third cut (d) results in $\{(x, y, z) \in A_2 \mid y \leq z\}$ and describes the Hill tetrahedron $\mathcal{H}_1$.
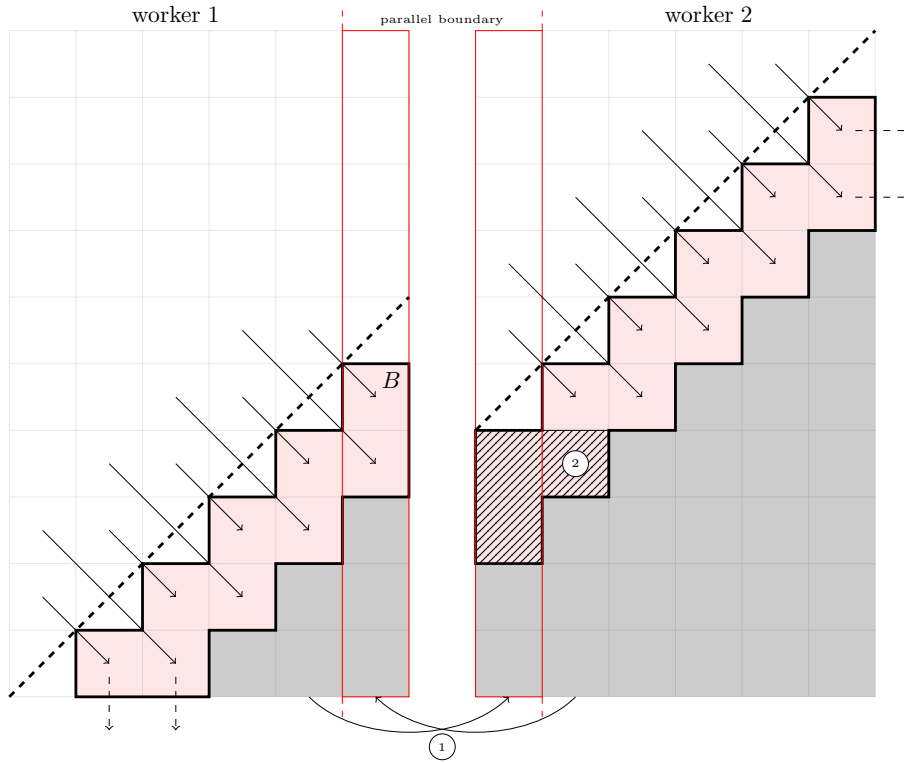
Figure 4: The scheme shows the simulation domain and copy instruction for the boundary conditions in a parallel simulation. The diagonal dashed line denotes the symmetry axis. The upper white cells are calculated, the lower gray cells are marked as non-calculated cells. After each step of the simulation, the boundary for parallel communication (red box) are exchanged by operation ①. In additional the domain boundary conditions are performed (not shown for clearness). Then the symmetry condition is archived for each inner domain border cell in B (red). The values from inside the domain are transferred by using the mapping functions $A_2, \ldots, A_6$ ($\longrightarrow$), except for the cells in ②. Here the mapping function needs data from another worker. The cells are not used for further calculation, and are hence marked as non-calculated cells (gray). If a cell in B is next to a domain boundary cell, the latter is also updated (implied by $--\rightarrow$).

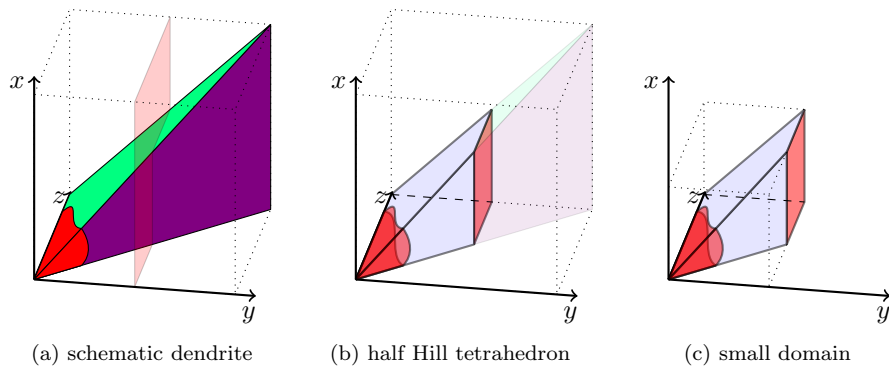(a) schematic dendrite   (b) half Hill tetrahedron   (c) small domain

Figure 5: (a) Representation of the Hill tetrahedron in Fig. 3(d) with a schematic dendrite and an illustration of the cutting plane, (b) bisected Hill tetrahedron and (c) reduced simulation domain of a rectangular cuboid of one quarter within the volume of Q.
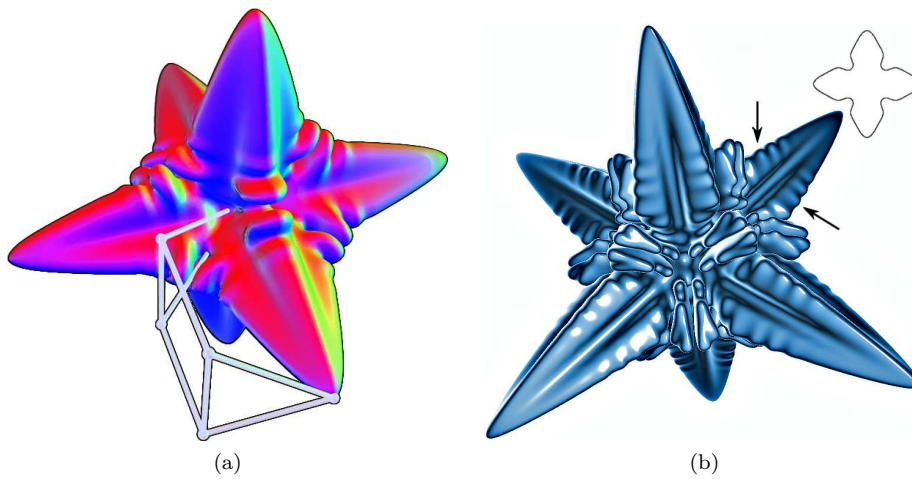
Figure 6: (a) Computed dendrite with a tip to tip distance of 2000 cells (1 μm) in its associated half Hill box. (b) Dendrite with the formation of ternary side-arms with a tip to tip distance of 4000 cells (2 μm) using experimental parameter and the kinetics $\beta = 0$. The contour of the cutting plane is marked by the arrows.