

Atomistic modeling of the solid-liquid interface of metals and alloys

Présentée le 24 janvier 2020

à la Faculté des sciences et techniques de l'ingénieur
Laboratoire de science computationnelle et modélisation
Programme doctoral en science et génie des matériaux

pour l'obtention du grade de Docteur ès Sciences

par

Edoardo BALDI

Acceptée sur proposition du jury

Prof. D. Damjanovic, président du jury
Prof. M. Ceriotti, directeur de thèse
Dr S. Angioletti-Uberti, rapporteur
Dr M. Salvalaglio, rapporteur
Prof. R. Logé, rapporteur

Quin etiam passim nostris in versibus ipsis
multa elementa vides multis communia verbis,
cum tamen inter se versus ac verba necessest
confiteare et re et sonitu distare sonanti.
Tantum elementa queunt permutato ordine solo;
at rerum quae sunt primordia, plura adhibere
possunt unde queant variae res quaeque creari.

(Lucrezio, *De Rerum Natura*, I, 823–829)

To both my grandmothers, Teresa and Maria. One convinced me to begin, and one
accompanied me with her thoughts till the end.

Acknowledgements

If experience always counts more than the result, I am delighted to have to thank some people who have made these four years of Ph.D. a remarkable experience.

First and foremost, my supervisor, Prof. Michele Ceriotti, who not only bet on me in 2015 but, above all, guided me, taught me, and endured me for four years. Even if my future career doesn't continue in the world of academic research, he made me understand what it means to do research in a scientific field.

Four years is not a long time, but it's already enough to take root in a place and establish interpersonal bonds. It was so for the group in which I worked on this thesis, the Laboratory of Computational Science and Modeling (COSMO). When I arrived, we were less than a dozen; in October 2019, the group now numbered about twenty people. Each of these people taught me something, and – I am convinced – allowed me to improve at least a little as a human being. Outside of the world of equations, simulations, and computers, Anne Roy also helped me a lot and not infrequently. Who knows how many administrative problems I would have got into without her help.

My doctoral project has been part of the “AM3 Materials Challenge,” supported by the Competence Center for Materials Science and Technology (CCMX), and it was made possible also thanks to the financial support of some companies that I have to mention here: Oerlikon Metco, Heraeus Materials, Rolex, Richemont–Varinor, Swatch–Asulab, Patek Philippe, Audemars Piguet.

This is the third time I mention Dr. Gabriele Sosso in a thesis, and I couldn't be happier to do so. Although much more sporadically and almost only by email, he has again advised and helped me to find the positive where I only saw the opposite.

It took me almost three years to call Lausanne “home.” This was only possible thanks to the people I was lucky enough to meet and know, many of whom had nothing to do with the world of my doctorate. Some of them accompanied me on more than one hiking to the Swiss Alps. Some of them put up with me in my ramblings after a movie at the cinema. Others even decided to share an apartment with me. It takes guts, so thank you.

Invaluable support has never been lacking from a tiny group of people who continue to live in a couple of small towns near Milan. I would be ashamed to quantify the ridiculous distance that separated us – and I am well aware of the luck of having my family so close – but I am sure that my sister, my mother, my father, and a few very close friends down there would have backed me up with the same perseverance even if my Ph.D. had been in the Chatham Islands.

It may seem even uncouth that very few names appear in these lines because the people who

Acknowledgements

have given me so much in these four years are many. But after all, these are just words written as they say “black on white” and they count for very little. What’s worth more is the bond I’ve had the good fortune to build with all these people. I’ll do anything I can to keep them in my life in the future, beyond words, in the real world.

Zürich, 25 November 2019

E. B.

Abstract

Solidification is a phase transformation of utmost importance in material science for it largely controls materials' microstructure on which a wide range of mechanical properties depends. Almost every human artifact undergoes a transformation that leads to a solid phase, be it via well-established manufacturing processes as casting or forging, or more recent technological advances such as 3D printing. This thesis aims to study some fundamental aspects of solidification, focussing in particular on metals and alloys.

Despite being a phenomenon investigated for a long time with both experiments and theoretical models, solidification still involves a very complex set of phenomena that requires a multi-scale approach in order to provide useful insights. For example, properties relating to the thermodynamics and kinetics of solid-liquid interfaces play a crucial role in micro-scale modeling of solidification, yet are particularly challenging to assess with experimental techniques. Our main instrument of investigation has been a set of well-established computer simulation techniques at the atomic scale, and we applied and extended some of these methods, with the primary goal of improving the reliability of some results related to the properties of solid-liquid interfaces and being able to study systems whose level of complexity comes closer to that of interest to some real manufacturing processes.

The approach of atomistic simulations involves several technical and theoretical problems. A first issue is related to the size of the systems that it is necessary to simulate in order to obtain results of some relevance despite an inherent statistical error that is often substantial. Moreover, some theoretical subtleties — such as the arbitrariness in the definition of the concept of interface — inevitably emerge from an approach at the atomic scale and lead to problems whose solution is anything but univocal; often, in fact, different formulations of these problems provide different results. The first contribution of this thesis focused on extending a computational method that serves to calculate a fundamental quantity known as *interface free energy*, and in particular to decrease its computational cost by reducing the number of particles the simulations contain. The second contribution addressed the study of crystal-melt interface properties of a particular metallic binary alloy. The idea behind this part of the work was to combine different techniques of atomistic simulations whose outcomes make it possible to obtain an exhaustive description of both the thermodynamics and the dynamics of the interfaces. We have based this approach on a precise definition of dividing surface, and we have derived all our results in a consistent way which allowed us to avoid some arbitrary choices that similar kinds of simulations usually entail. A part of the results obtained confirmed the reliability of our approach, showing satisfactory agreement with some other established results. However, challenges remain associated with

Abstract

the accuracy of the interatomic potential, the presence of significant finite-size effects, and the difficulty in converging to satisfactory statistical accuracy the thermodynamic and dynamical properties of solid-liquid interfaces.

Keywords Molecular dynamics, Monte Carlo, enhanced sampling, solid-liquid interfaces, thermodynamics, solidification, binary alloys, Additive Manufacturing.

Sommario

La solidificazione è una transizione di fase di primaria importanza nella scienza dei materiali, in quanto controlla in larga misura la microstruttura dei materiali da cui dipende un'ampia gamma di proprietà meccaniche. Pressoché ogni manufatto prodotto dall'uomo subisce una trasformazione che porta a una fase solida, sia attraverso processi produttivi consolidati come la fusione o la forgiatura, sia attraverso i più recenti progressi tecnologici come la stampa 3D. Questa tesi si propone di studiare alcuni aspetti fondamentali della solidificazione, concentrandosi in particolare su metalli e leghe.

Pur essendo un fenomeno studiato da tempo sia con esperimenti che con modelli teorici, la solidificazione comporta un insieme di problemi molto complesso; un approccio multi-scala rappresenta l'unica via possibile per ottenere soluzioni, spesso approssimate, ad alcuni di questi problemi. Ad esempio, le proprietà relative alla termodinamica e alla cinetica delle interfacce solido-liquido giocano un ruolo cruciale per descrivere la solidificazione su scala atomica, ma sono particolarmente difficili da misurare con tecniche sperimentali. Il nostro principale strumento di indagine è stato un insieme di tecniche di simulazione atomistica ben consolidate. Abbiamo adoperato ed esteso alcuni di questi metodi con l'obiettivo principale di migliorare l'affidabilità dei risultati di certe proprietà delle interfacce solido-liquido e studiare dei sistemi il cui livello di complessità si avvicinasse maggiormente a quello dei materiali impiegati in alcuni processi produttivi reali.

L'approccio delle simulazioni atomistiche comporta diverse problematiche sia tecniche che teoriche. Un primo problema è legato alle dimensioni dei sistemi che è necessario simulare per ottenere risultati affidabili, nonostante l'errore statistico associato a queste grandezze sia spesso considerevole. Inoltre, alcune sottigliezze teoriche — ad esempio, l'arbitrarietà nella definizione di un'interfaccia — emergono inevitabilmente da un approccio atomistico e generano problemi la cui soluzione è tutt'altro che univoca; spesso, infatti, formulazioni diverse di questi problemi forniscono risultati che possono essere molto diversi fra loro. Il primo contributo del lavoro di questa tesi è stato di estendere un metodo computazionale che serve a calcolare una quantità fondamentale nota come *energia libera di interfaccia*, e in particolare a diminuire il suo costo computazionale riducendo il numero di particelle che le simulazioni richiedono. Il secondo contributo ha riguardato lo studio delle proprietà di interfaccia solido-liquido di una particolare lega binaria metallica. L'idea alla base di questa parte del lavoro è stata quella di combinare diverse tecniche di simulazione atomistica i cui risultati permettono di ottenere una descrizione esaustiva sia della termodinamica che della dinamica delle interfacce. Abbiamo basato questo approccio su una precisa definizione di interfaccia, derivando tutti i nostri risultati all'interno

Sommario

di un framework consistente ed evitando alcune assunzioni arbitrarie che simulazioni di questo tipo solitamente comportano. Una parte dei risultati ottenuti ha confermato l'affidabilità del nostro approccio, mostrando un soddisfacente accordo con dei risultati della letteratura scientifica. Tuttavia, rimangono ancora molte problematiche associate all'accuratezza del potenziale interatomico, alla presenza di significativi effetti di scala, e alla difficoltà di convergere con soddisfacente precisione statistica i risultati delle proprietà termodinamiche e dinamiche delle interfacce solido-liquido.

Parole chiave Dinamica molecolare, metodi Monte Carlo, metodi di campionamento accelerato, interfacce solido-liquido, termodinamica, solidificazione, leghe binarie, stampa 3D.

Contents

Acknowledgements	v
Abstract	vii
Sommario	ix
List of Figures	xiii
List of Tables	xix
1 Introduction	1
1.1 Solidification phenomena	1
1.2 Additive manufacturing	2
1.3 Materials	3
1.4 Numerical modeling of solidification	6
I Theory	13
2 Statistical mechanics	15
2.1 Phase space and esembles	15
2.2 Common ensembles	17
2.3 The NVE ensemble	17
2.4 The NVT ensemble	18
2.5 The NPT and μ VT ensembles	18
2.6 The ergodic hypothesis	19
2.7 Statistical errors	20
3 Computer experiments	23
3.1 Interatomic potentials	24
3.2 Molecular dynamics	29
3.3 Monte Carlo methods	34
3.4 Enhanced sampling	35
4 Thermodynamics	39

Contents

4.1	Fundamentals	39
4.2	Binary systems	43
4.3	Thermodynamics at interfaces	58
II	Results	73
5	Anisotropy and free energy of a smooth fluctuating dividing surface	75
5.1	Theoretical background	75
5.2	System and simulation details	77
5.3	Results and discussion	79
5.4	Details about the order parameter for FCC symmetry	85
5.5	Implementation	87
6	Solid–liquid properties of binary alloys from atomistic simulations	95
6.1	Theoretical background	95
6.2	Systems and computational details	102
6.3	Equilibrium thermodynamic properties	103
6.4	Interface properties	108
7	Conclusions	119
III	Appendices	125
A	Interface properties of a Lennard-Jones model	127
A.1	Interface stiffness and free energy	127
A.2	Interface dynamics	130
B	Block averaging in practice	135

List of Figures

1.1	A pie chart illustrating the research interests of Laser Melting Additive Manufacturing applied to different classes of materials. Data are based on publications indexed by Web Of Science and ScienceDirect from 1999 to 2014 [4].	3
1.2	Portion of the NiAlTi ternary phase diagram at 1423 K where the regions of the γ and γ' phases have been labeled. The full phase diagram of this system has been calculated with the CALPHAD software using proprietary thermodynamic data [9].	5
1.3	Schematic plot showing time and length scales of several modeling techniques that should be combined to study solidification.	9
4.1	Gibbs free energies of mixture (G_1), of the ideal solution (G_2), and the configurational contribution to entropy due to mixing. The Gibbs free energy of the pure component A can be taken equal to zero, while $G_B > G_A$	45
4.2	Example of free energy curves in an ideal system. At this temperature, the solid and the liquid phase have the same free energy only for $x_B = 0$, that is, pure A. The temperature therefore indicates the equilibrium melting temperature for the pure component A.	46
4.3	Free energy curves for an ideal system as a function of the fraction of B. M is the lowest free energy for the alloy with composition x	47
4.4	Schematic plot of the molar Gibbs free energy of a regular solution as a function of the fractional amount of component A. Without loss of generality, we set $G_{A,m} = 0$ and $G_{B,m} = 10$. The dashed line is the common tangent line. The two black dots define the compositional range in which the Gibbs free energy of the solution is greater than that of a two-phase aggregate with compositions $x_{A,\alpha}$ and $x_{A,\beta}$. Within this range of compositions, the binary mixture separates into two mixtures: β , richer in B, and α , richer in A.	51
4.5	Gibbs free energy for a regular solution as a function of the fractional amount of component A plotted for several temperatures. The lowest curve, labeled with T_c , is the graph of G_m at precisely the critical temperature. The upper curve represent increasing temperatures. For each curve, the two pairs of black and red dots represent the points of common tangency and the inflections points, respectively. At $T = T_c$, the compositional range for which phase separation occurs completely disappear, and this is seen by the overlapping dots.	52

List of Figures

4.6	Schematic plot of free energy curves as a function of the composition for two terminal phases (α and γ) and a β phase that indicates a large dissimilarity between the components.	53
4.7	Schematic plot of free energy curves as a function of the composition in which the phase β represent an intermediate phase of a solid–solution type.	54
4.8	A schematic phase diagram of a binary system which exhibits a liquidus minimum and a solid solubility gap.	55
4.9	(a) Hypothetical phase diagram when the two solid phase region of fig. 4.8 intersects the upper part showing the liquidus minimum. (b) The phase diagram on the left correctly redrawn according to Gibbs’ phase rule, as explained in the text.	56
4.10	Enlarged portion of fig. 4.9b around the point e that explains what happens during eutectic solidification.	58
4.11	Schematic representation of (a) a real, heterogeneous system comprising two phases in contact, and (b) a model system in which a Gibbs dividing surface is defined by setting zero surface excess of the molar volume.	59
4.12	Schematic plot of the Gibbs free energy of a unary system versus the value of a suitably chosen collective variable (CV). The equal depth of the two minima, that correspond to single–phase configurations of the system, implies that the chemical potential of the two phases is the same, that is, the system is at equilibrium temperature and pressure. The energy cost associated to a phase transition that brings the system from bulk solid to bulk liquid is precisely $\gamma_{sl}A$, where A is the interface area. Moreover, being at equilibrium means that the value of γ_{sl} obtained with this method is independent of the choice of the CV. The same statement is not valid in out of equilibrium conditions, although it is still possible to employ metadynamics simulations to extract γ_{sl} as a function of temperature [74].	68
5.1	Fitted Fourier spectra along the k_x direction for a range of values of the Gaussian kernel coarse–graining length ξ (see eq. (5.4)).	79
5.2	Fitted values for the stiffness tensor for the (100) interface as a function of the radius of the cutoff in k –space	81
5.3	Fitted values for the stiffness tensor of the (110) interface as a function of the radius of the cutoff in k –space	82
5.4	Plot illustrating the auto–correlation time τ for $A(k_x, k_y)$ in Lennard–Jones reduced units of $\sqrt{m\sigma^2/\epsilon}$ as a function of k_x (red circles) and k_y (blue diamonds) for the (100) interface.	82
5.5	Fitted values for the stiffness tensor for the (110) interface as a function of the radius of the cutoff in k –space, obtained by performing fits using eq. (5.8) . . .	83

5.6	Graphical representation that illustrated the dependence of the free energy of the interface on its orientation. The high symmetry directions are highlighted and it is clear that the (100) has a higher interface free energy than the (111). A perfect sphere would indicate a fully isotropic interface free energy.	85
5.7	Probability density distribution function for the order parameter that is described in this section. The histograms shown in red and blue in this figure were calculated from a snapshot of the (110) simulation. The dashed line shows the switching function of eq. (5.11), that was used to convert the order parameter values into a continuous quantity that better distinguished the solid from the liquid.	87
5.8	Average value of the FCCUBIC symmetry function defined in eqs. (4.22) and (5.11) as a function of x , the direction parallel to the interface normal.	88
5.9	A snapshot from one of molecular dynamics simulations of a solid–liquid interface in a Lennard–Jones system. In this figure the atoms are colored according to the value they have for the symmetry function defined in eqs. (4.22) and (5.11). Those colored blue have a high value for Φ_i and are thus in a solid–like environment. Those colored in red have a low value for this quantity. The structure surrounding these atoms is thus more liquid–like.	89
5.10	A snapshot from one of molecular dynamics simulations of a solid–liquid interface in a Lennard–Jones system with the atoms colored as described in the caption of fig. 5.9. The Willard–Chandler surface that separates the solid and liquid phases is shown in yellow in this figure.	91
5.11	A contour plot showing the height of the interface between the solid and liquid phases of Lennard–Jones versus x and y for the snapshot of the trajectory shown in fig. 5.9.	92
6.1	The upper part of panel (a) shows a plot of the average order parameter as a function of the coordinate along z perpendicular to the interface. The average has been calculated with a Gaussian kernel density estimation. The colored areas show the bulk regions: solid (blue) and liquid (red). The dividing surfaces' locations, which are found to be around -40 \AA and 40 \AA , are determined according to the Gibbs construction (see eq. (6.8)). The lower part of panel (a) shows a snapshot of the solid-liquid simulation from which the upper plot has been calculated. Panel (b) shows a histogram of the distribution of the order parameters colored according to the degree of crystallinity.	104

6.2	Concentration profiles across the interface obtained as described in the text for (a) silver and (b) copper at three different temperatures. The zero of the z coordinate coincides with the location of the Gibbs dividing surface. In both figures, the solid phase lies in the $z < 0$ region, while the liquid phase lies in the $z > 0$ portion of the plot. The values of the concentrations where the profiles plateau off correspond to the equilibrium compositions of the bulk phases. From these profiles, it is evident that the solid solubility limits are both very similar at all temperatures and close to 0.1 and 0.9, suggesting that the corresponding solid solutions are primarily composed by copper and silver, respectively. . . .	105
6.3	Relative surface excesses a function of temperature. The blue curve shows the adsorption coefficient of copper relative to silver, thus it has been computed for total concentrations of silver $x_{Ag} > 0.5$ where it is reasonable to assume that copper is the solute. The orange curve instead shows the adsorption coefficient of silver relative to copper when $x_{Ag} < 0.5$	107
6.4	Phase diagram of the silver-copper binary alloy above the eutectic. The red circles are the equilibrium compositions of the solid and liquid phase as obtained using the simulation framework described in the text. The blue, dashed lines are the results obtained by Williams, Mishin, and Hamilton [127] with Grand-Canonical Monte Carlo (GC-MC) simulations. The black, dotted lines are the experimental phase diagram [134, 135]. The comparison shows that the EAM potential predicts satisfactorily the melting temperatures of the pure elements within about 30 K of the experimental values, while the estimated eutectic composition (0.458) and temperature (935 K) deviate more from the experimentally measured values of 0.601 and 1053 K.	107
6.5	Plot of the quantity $k_B T / S / \langle A(k_x, k_y) ^2 \rangle$ versus \mathbf{k} for the (100) interface of (a) pure silver and (b) pure copper. The simulations were run at the melting temperatures predicted by the EAM potential employed. The blue curve is along k_x ($k_y = 0$), while the orange curve is along k_y ($k_x = 0$). The isotropy of the (100) interface can be seen from the two lines of each plot having roughly the same intercept with the y axes.	109
6.6	Plot of the results of γ_{sl} (mJ/m ²) of the silver-copper alloy as a function of the total fraction of silver, x_{Ag} . Both the values calculated for pure copper and pure silver are comparable with results found in the literature [83, 138]. The two sets of results at the eutectic composition refer to independent simulations in which the concentrations of the solid phases were set to the solubility limits at the eutectic; in particular, in the lower set of points x_s was equal to 0.933 (silver-rich solid-solution), whereas in the upper set $x_s = 0.03$ (copper-rich solid-solution). As discussed in the text, Gibbs adsorption coefficients provide us with only one of the ingredients required to justify the observed dependence of γ_{sl} on composition, yet the two sets of results at the eutectic composition suggest that this dependence is not an artifact of our simulations.	111

6.7	Anisotropy magnitude as defined by eq. (6.20) as a function of the total fractional amount of silver, x_{Ag} . The results for pure copper and pure silver are also shown. The latter value is in good agreement with that predicted by Hoyt and Asta [83], in which they employed a different interatomic potential. Within their statistical uncertainty, these results suggest that for alloyed compositions γ_{sl} is almost isotropic with an anisotropy magnitude very close to zero. Analogously to fig. 6.6, the two points at the eutectic composition refer to two independent simulations where the solid phase was either a copper-rich or silver-rich solid-solution.	112
6.8	Kinetic coefficient M ($\text{\AA ps}^{-1} \text{ eV}^{-1}$) of the (100) interface estimated from eq. (6.12) as a function of the total fractional amount of silver, x_{Ag} . The values for the pure elements are of the right order of magnitudes we expect when comparing them to similar results found in the literature [83, 102].	114
6.9	Plot of the auto-correlation functions of f for the simulations both at intermediate compositions and for the pure elements.	115
6.10	Solid fraction as monitored during the simulations at the total compositions of 0.165, 0.805, and 0.458 (eutectic). The drift in the solid fraction, in particular for the simulation with $x_{\text{Ag}} = 0.805$, makes it clear why it was not possible to fit the corresponding auto-correlation function with a simple exponential model.	116
A.1	Plots of $\langle A(k_x, k_y) ^2 \rangle$ versus k_x ($k_y = 0$) for two different cells with a one-dimensional geometry: (a) $28 \times 4 \times 28$ cell and (b) $40 \times 7 \times 40$ cell. Both these plots show a k -dependence of the solid bulk contributions when they are calculated from bulk slabs instead of full cells.	128
A.2	Plot of $\langle A(k_x, k_y) ^2 \rangle$ versus k_x ($k_y = 0$) for a simulation cell with two-dimensional geometry with size $40 \times 40 \times 24$. Although the bulk contributions of the liquid phase are now independent of whether they are evaluated from a slab section or a full cell, the solid bulk contributions still show a remarkable difference especially at small k , suggesting that this “slab effect” is affecting more prominently the solid phase.	129
A.3	Plot of the liquid and solid bulk contributions to $\langle A(k_x, k_y) ^2 \rangle$ as a function of k for different simulation cells sizes, with both one and two-dimensional geometries. As the cell size increases, the difference between the contributions of liquid and solid in the same simulation and between the same bulk phase in two different simulations increases. In particular, this effect is particularly pronounced for the solid phase, where there is a difference of about two orders of magnitude in the absolute value of $\langle A(k_x, k_y) ^2 \rangle$ between the smallest cell, $28 \times 4 \times 28$, and the largest $40 \times 40 \times 24$	130
A.4	Plot of the solid fraction as a function of time monitored during two pinned-interface simulations of Lennard–Jones in which two different values of the strength of the harmonic bias applied have been used.	132

- A.5 Plot of the number of solid atoms N_s as a function of time during several non-equilibrium simulations with different values of undercooling ranging from 2.0% (rightmost blue curve) to 5.0% (leftmost brown curve). In these simulations, temperature was fixed by a local, Fast-forward Langevin thermostat [147]. The plateau regions for all but the smallest undercooling indicate that our simulations were long enough to observe a complete solidification, and the linear growth of N_s confirm that constant non-equilibrium conditions have been attained for all the simulations. The dashed vertical lines indicate the cutoffs up to which the linear fits were performed, with the rightmost vertical line referring to the smallest undercooling and so on. 133
- A.6 Crystal growth rate as a function of temperature below the melting temperature of $0.62 \epsilon/k_B$. The two curves represent two sets of simulations where different thermostats were keeping the temperature constant during the non-equilibrium run: the orange curve reports the results from the simulation in which a global, stochastic velocity-rescaling thermostat [50] was used, while the blue curve reports results from simulations with a local, Fast-forward Langevin [147] thermostat. The values of R and their error bars are calculated from five independent simulations with the same undercooling. Besides the shift between the two curves, the linear fits (black dashed lines) suggest two consistent values of the mobility, which turned out to be 2.51 ± 0.18 (stochastic thermostat) and 2.18 ± 0.15 (Langevin thermostat). 134

List of Tables

3.1	Lattice properties of Ni and Al calculated with the potential developed by Mishin [37] in comparison with experimental data [40, 41]. The perfect match between the predicted and experimental equilibrium energy and lattice parameter is due to the inclusion of these quantities in the dataset used to fit the EAM parameters.	28
3.2	Comparison table between predicted values of EAM potentials for gold discussed in this section. The table is partially reproduced from Grochola, Russo, and Snook [42]. Elastic constants are in GPa, energies in eV and lattice parameters in Å.	29
5.1	Dimensions of the simulations and the number of atoms in each of the cells. The fourth column gives the dimension along the z axis, perpendicular to the interface normal. All lengths are in units of σ	78
5.2	Stiffness values for the (100) and (110) interfaces calculated using our method at a reduced temperature of 0.6185. The values of the stiffness are given in units of ϵ/σ^2 . For comparison, the values in parenthesis are those obtained by Becker et al. [115] at the same temperature.	84
5.3	Specific free energy for the solid–liquid interface in a Lennard–Jones system at the melting temperature of 0.6185. The results obtained here are compared with those in previous work of Becker et al. [115] and Cheng, Tribello, and Ceriotti [117]. The free energy is in units of ϵ/σ^2	86
6.1	Composition (and temperature) dependence of γ_{sl} (mJ/m ²) of the (100), (110), and (111) interfaces as as function of the total fractional amount of silver, x_{Ag} . The results obtained for pure silver and pure copper are also shown. The reported errors have been estimated from the standard errors in the fitting parameters of the models $\sigma(k_x^2 + k_y^2)$, used for the (100), and $\sigma_{11}k_x^2 + \sigma_{22}k_y^2$, used for the (110).	110

1 Introduction

1.1 Solidification phenomena

The solid-state of matter is the one with which everyone is most familiar. Solidification, therefore, is a transformation of matter ubiquitous and well known to all, even if only for a natural phenomenon such as the formation of ice crystals that make up the snow. It is therefore of fundamental importance in material science since almost every human artifact undergoes a transformation that leads to a solid phase. Solidification as a manufacturing step could occur via either well-established methods like casting and forging, or more recent technological processes such as Additive Manufacturing, also known as 3D printing. It should not also be overlooked the recent proliferation of new and advanced numerical simulation methods. Thanks to the availability of computers with extremely efficient parallel architectures, it is now realistic to model portions of condensed matter containing up to billions of atoms on time scales of several nanoseconds [1–3]. The work discussed in this thesis focused on the application and, in small part, the extension of atomistic simulation methods to study the phenomenon of solidification.

The family of processes going under the name of *casting* gathers the most widely adopted methods by the metalworking industry to produce metallic objects. If the precursor materials do not have prohibitively high melting points, casting methods are often a better choice of manufacturing object with respect to forging. For example, the production rate of casting is usually higher once the mould – the object representing the negative of the desired shape – is ready. Furthermore, casting usually requires very few secondary machining processes and allows more freedom and higher level of detail in the final shape. On the other hand, detailed reports¹ on the comparison between casting and forging measured a ultimate tensile strength that is about 25% higher in forged parts than in cast parts, and showed that forging leads to objects that are going to last longer due to a higher fatigue strength.

Even though casting would be a better manufacturing choice when it could meet the desired final product's quality and mechanical properties, it poses multiple challenging problems which pertain to solidification process's conditions. These, in turn, influence the material microstructure on which mechanical properties of the final object strongly depend. For example, the tensile strength is strongly hindered by the average distance between microstructure features such as dendritic branches. Being able to control this kind of properties which are results of the solidification process is key to obtain a finished product with desired characteristics.

The oldest known object produced by casting with a technique known as *investment casting* (or *lost-wax casting*) is dated to be roughly 6000 years old. This technique was probably born as a natural consequence of clay processing used to produce mainly decorative artifacts. The possibility of manufacturing metal objects — copper, initially, followed by bronze of the first objects found in China and dated around 5000 BCE, and then iron around 500 BCE — has ushered the use of casting techniques in areas such as agriculture, architecture and warfare.

As Kurz states in his *Fundamentals of solidification*, it is interesting to note how little the technological importance of the solidification process is appreciated. There are many examples of modern technological applications where nothing would have been possible if the problems of understanding solidification had not been addressed with the same analytical approach that science has dedicated to many other fundamental problems. Prominent examples are microelectronic and solid-state technologies, where solidification of highly pure crystalline materials is of utmost importance. The well-established integrated circuits industry requires the preparation of large amounts of single crystals of silicon without impurities, except those purposely introduced to modify the material's electronic properties necessary for the operation of devices such as the transistor. For this reason, the need to understand solidification in depth has led the knowledge of this phenomenon from being a collection of rather empirical procedures and observations to the status of an independent branch of modern material science.

1.2 Additive manufacturing

Until recently, being able to make objects with varying shapes without going through an excessively long preparation phase pertained mostly to rapid prototyping. However, the latest technological developments on techniques such as AM opened a whole new realm of possibilities.

Several industrial sectors, ranging from the energy one to high precision watchmaking, have for some time shown increasing interest in the possibility of introducing AM in their manufacturing processes. The foremost reason for such interest is the possibility to design the final object employing a three-dimensional modeling software and then fabricate it directly with almost no need of any other intermediate design stage. Other manufacturing processes would require a detailed analysis of the final object to determine, for example, the order in which each part has to be built and later assembled.

AM technologies are considered as a step-change revolution in product development and manufacturing, and they could redefine what is meant today by a “manufacturing process.” The optimism of these forecasts underscores the two main advantages of AM technologies. First, the rapid character of the whole process: it is not just in terms of overall time needed to build an object, but all the development process experiences a speedup, primarily because computer's aid is used throughout.

The second advantage is the reduction in the number of process's steps despite the possible complexity of the final object. AM usually involves one single step, while standard manufacturing processes would require a particular sequence of multiple and iterative stages to be carried out. Even a relatively simple change in the project may require significant modification in the fabrication process, which, in turn, increases the time required to obtain the final product.

Other drawbacks are the often high processing costs due to the amount of waste material and that required to build components necessary to the entire process – e.g., molds in the specific case of casting.

Current research efforts at both the academic and industrial level focused on a more in-depth understanding and improvement of the three fundamental laser-based processing that developed in parallel as AM matured. These are *laser melting* (LM), *laser sintering* (LS), and *laser metal deposition* (LMD). Although the doctoral project focused on the phenomenon of complete solidification which occurs only in the case of laser melting, these three variants of the manufacturing process allow working with different classes of materials, of which the next section will give a brief overview.

1.3 Materials

All the Additive Manufacturing techniques mentioned previously can be in theory applied to any metallic material, be it a pure metal or an alloy (see fig. 1.1). After a brief discussion about the usage of both pure metals and alloys in AM processing – with a particular focus to the *selective laser melting* technique – we give a more detailed description of nickel-based superalloys that are particularly relevant for several high-performance applications.

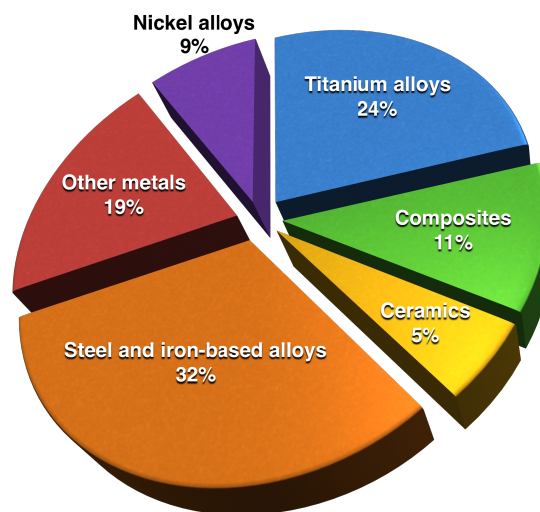


Figure 1.1 – A pie chart illustrating the research interests of Laser Melting Additive Manufacturing applied to different classes of materials. Data are based on publications indexed by Web Of Science and ScienceDirect from 1999 to 2014 [4].

1.3.1 General considerations on pure and alloyed metals

Pure metals have been applied for various AM processes since the beginning. However, they are no more the focus of future development of AM technology mainly because of two reasons: firstly, pure metals are known to have relatively poor mechanical properties and, from a chemical point of view, they are more prone to corrosion and oxidation. Secondly, early attempts to

processing pure metals with the laser sintering technique have been largely unsuccessful (i.e. poor overall quality of final products) and a real improvement has been possible only with first successful applications of LM. For example [5], LS processing of pure Ti samples showed that final products typically have an heterogeneous microstructure, in which powder grains had cores showing incomplete melting, melted surfaces and voids; this means that the density of parts so obtained can be quite far from the level required to ensure good mechanical properties. For these reasons, currently industrial applications moved from LS to LM or LMD to build nonferrous pure metals components. However, recent studies [6, 7] on samples of pure metals such as Ti and Ta that underwent only partial melting have demonstrated the possibility of producing porous structures with complex shapes and controlled porosity that could be interesting for biomedical applications.

The majority of the actual research (and subsequent applications) has been focused on powder alloys based on metals such as Ti, Ni and Fe. The next research step might focus on aluminum-base alloys since they pose well-known difficulties related to their high reflectivity to laser energy. LS and LMD are the most widely adopted AM techniques to process almost all the alloys powder where complete melting is required. Therefore, it is necessary to utilize high power laser sources in order to obtain parts with the desired density, except in cases where the porosity of the final material is a required characteristic. Examples of lasers used for this purpose are powered lasers, Nd:YAG and fiber lasers. Besides the setup of the processing parameters, residual stresses and microstructure are two important aspects to pay attention to, as they are largely affected by the substantial undercooling during rapid solidification after the laser melting step. As proposed by Abe et al. [8], a possible way to solve these issues is the use of a dual laser scanning system; essentially, in this way one of the two laser can be used during a preheating phase before the actual melting or afterwards in a reheating phase.

1.3.2 Nickel alloys

The fundamental solutes in nickel based superalloys are aluminium and titanium, with a total concentration which is typically less than 10 atomic percent. This generates a two-phase equilibrium microstructure, consisting of a γ and a γ' phase. It is the γ' which is largely responsible for the elevated temperature strength of the material and its incredible resistance to creep deformation. The amount of γ' depends on the chemical composition and temperature.

The γ phase is a solid solution with a Face Centered Cubic lattice where different chemical species are randomly distributed. By contrast, γ' is an intermetallic phase with composition $\text{Ni}_3(\text{Al}, \text{Ti})$ and has a simple cubic Bravais lattice where one nickel atom lies at the center of each face of the cube, while aluminium or titanium atoms are at the cube corners.

The following fig. 1.2 shows a portion of the ternary NiAlTi phase diagram at 1423 K. The $(\gamma+\gamma')/\gamma'$ phase boundary suggests that the phase is not strictly stoichiometric: sub-lattice vacancies or substitutional defects might lead to deviations from stoichiometry.

One of the most important properties of these alloys is their strength up to very high temperatures. More precisely, these alloys show high yield strengths and at first these values are independent on temperature. An explanation of this peculiar property lies in the structure




Figure 1.2 – Portion of the NiAlTi ternary phase diagram at 1423 K where the regions of the γ and γ' phases have been labeled. The full phase diagram of this system has been calculated with the CALPHAD software using proprietary thermodynamic data [9].

of the two phases: both have a cubic crystal structure with similar lattice parameters, and the γ phase forms the matrix in which the γ' precipitates. This transition preserves the symmetry of the cubic lattice, and thus the unit cell edges of the γ' remains parallel to the corresponding edges of the γ phase. Moreover, when the precipitate's size is small, the lattice misfit between the two structures is also small. Dislocations that form in the γ phase do not penetrate easily into the γ' , partly because the latter is an intermetallic crystalline phase. The mechanism of propagation of the dislocations is therefore partially prevented by the crystalline order of the precipitated phase and this contributes considerably to strengthening the alloy.

The lattice mismatch between the γ and the γ' phase plays quite an important role: it promotes the stability of the microstructure at high temperature. This is mainly due to the fact that the γ and the γ' belong to almost equal crystal structure and thus the γ/γ' interface energy is quite low. The mechanism of precipitate coarsening (i.e. the Oswald ripening) is thermodynamically driven by the minimization of the interface energy, and thus the more coherent is the interface – the smaller its surface free energy – the more stable is the microstructure, especially because, at high temperatures, it does not happen that a precipitate of a solid phase grows at the expense of another one.

Nickel superalloys of some industrial interest contain more than just Ni, Al or Ti. Essential resistance to oxidation is promoted also by the presence of chromium and aluminum, while small quantities of yttrium help the oxide layer to be commensurate with the substrate. Furthermore, when dealing with polycrystalline superalloys, the addition of boron and zirconium results in reducing the grain boundary energy, which in turn improve overall creep strength.

1.4 Numerical modeling of solidification

Solidification phenomena offer an excellent example of a rather complex set of problems both from a physical and a mathematical standpoint. Despite knowing precisely the fundamental equations behind them, for the results to be of some practical utility, the only hope is to rely on numerical modeling methods.

There are several reasons why this is the case, but the underlying one is that an analytical solution to the fundamentals equations of solidification is bounded by very restrictive conditions that rarely apply for real systems. To give a few examples:

- Exact solutions often assume overly simplified geometries, or they may be obtained only for one-dimensional systems.
- Material properties — e.g., density, heat conductivity, or diffusion coefficients — cannot be taken as uniform quantities even for the case of elemental systems.
- Phenomena such as heat transfer or mass transport may be decoupled only in very few and simple cases.

A numerical solution lies at the heart of the fundamental technique of molecular dynamics, that has been extensively used to produce all the results discussed in this thesis. The main interest of many numerical approaches to solve complicated problems is to understand how the parameters of the simulation affect the accuracy of the solution. Analytical solutions that exist for model systems can serve as benchmarks against the results obtained by the chosen numerical approximation, thus giving precisely a measure of how accurate the numerical approach could be when applied to real systems of interest.

An essential task of whichever approach is chosen to model solidification is to find a suitable description of the problem that allows following the phase separation front in both time and space as the material solidifies, as it provides fundamental insights on the microstructure of the resulting phase. Numerical methods of this kind belong to two categories, which differentiate by the type of grid employed to discretize both time and spatial dependence. These two classes are *fixed grid* methods and *front tracking* methods.

1.4.1 Finite difference, volume and element methods

The most direct approach belonging to fixed grid methods is the *finite difference method*. The idea behind does not restrict to model solidification problems, but it is far more general: it decomposes both the time and spatial domain upon which some variable of interest depends, and it chooses two discrete units of time and space: the evolution of the variable can then be followed through a sequence of time-steps.

As an example, consider the one-dimensional equation of heat conduction

$$\frac{\partial T}{\partial t} = \alpha \frac{\partial^2 T}{\partial x^2}, \quad (1.1)$$

where the spatial variable x is bounded in the interval $[0, L]$ and $t > 0$. One can define the discrete evolution of T as follows

$$T_i^n = T((i-1)\Delta x, (n-1)\Delta t),$$

where i runs over all the points in the spatial dimension of the grid, while n is a positive integer. One can then write explicitly the expressions for T at different positions, and expand these as Taylor series about the value of temperature at the current point. The expansions are truncated at a certain order in Δx , which tells about how the approximate solution will converge (diverge) when the grid spacing decreases (increases). With this procedure, one obtains an expression for the second-order spatial derivative of T

$$\frac{\partial^2 T}{\partial x^2} = \frac{T_{i-1} - 2T_i + T_{i+1}}{\Delta x^2} + O(\Delta x^2).$$

By following an identical procedure to discretize also the time dependence, one arrives to an equation that relates the dependence of T in the time domain to that in the spatial domain. This relation is the central result of any finite difference method, and it is usually called a *stencil* because successive applications lead to the full² – yet approximate – dependence of the variable on time and space.

The intrinsic limit of the finite difference method is that, although both Δx and Δt need not be constant, it is not evident how the accuracy changes for an irregular grid, because Δx and Δt would depend explicitly on time and space.

An alternative method to obtain time and spatial dependence of T that guarantees an accuracy equal to that of the Taylor expansions even with irregular grids is the *finite volume method*. The idea of this method is to split the spatial domain into control volumes, which may not well be constant, and retain a finite difference scheme for the time domain. The fundamental equation underlying this method is a balance equation as the following

$$\int_V \rho c_p \frac{\partial T}{\partial t} dV = - \int_A \mathbf{q} \cdot \hat{\mathbf{n}} dA, \quad (1.2)$$

where \mathbf{q} is the heat flux vector and $\hat{\mathbf{n}}$ is the normal to the surface enclosing a control volume; ρ and c_p represent the density and the specific heat at constant pressure, respectively. The Fourier's law of heat conduction for an isotropic medium reads

$$\mathbf{q} = -\kappa \nabla T,$$

so the $\partial T / \partial x$ term on the right-hand side of eq. (1.2) can be expressed, at time t_n , via the finite difference approximations explained previously, as well as the time derivative on the left-hand side. Although finite volume method does not constrain control volumes to be mapped onto a structured grid, it can be shown³ that for a two-dimensional Cartesian grid the finite volume method leads to the same time and spatial evolution equations as the finite difference method. The real advantage over the latter is that the accuracy is still $O(\Delta x^2)$ regardless of whether the grid is regular or not.

A third widely used strategy to approach this kind of modeling problems is known as *finite element method* and it is closely related to finite volume. Both these two methods start by decomposing the spatial domain in volumes that are defined by the points connecting one element to its neighbors. In the finite element method the variable of interest is defined for every point in each control volume using some interpolation scheme, usually a polynomial one. This strategy allows avoiding the use of finite difference approximations, except for the term expressing the time derivative.

The finite volume method can be considered a special case of the finite element method, in which the interpolating functions are just piecewise, constant functions, and the value of the variable on one grid point represents the entire volume to which the point belongs. A clear advantage of the finite volume method over its parent one is that it deals directly with fluxes and quantities integrated over every volume of the grid – expressed by relation similar to eq. (1.2). This approach bears a clearer physical significance, while the mathematics behind finite element methods may become a little cumbersome, especially for very complex domain geometries. On the other hand, finite element methods can be generalized to higher orders of approximation in a straightforward manner, since the polynomial interpolation provides an easier way to estimate accuracy.

1.4.2 Zooming in on the system: phase field and atomistic modeling

As mentioned previously, the number of real–world technological advances in which modern material science has an impact is uncountable. In almost all these fields of application, the importance of solidification is that it ultimately controls the underlying microstructure of materials.

Given the nature of solidification phenomena, it should not be surprising that in both pure science and computational modeling the best approach to the study of this class of problems is a hierarchy of simulations techniques that operate on multiple scales. Figure 1.3 is a typical plot of this interplay between different methods. It does not only restate that each method has its intrinsic limitations, but, more importantly, that any information obtained by, say, atomistic modeling is a valuable input for one technique whose target is to simulate the behavior of materials at the micro–scale.

A prominent example of a method that has emerged in both engineering and material science is the so–called *phase field method*. The features of this method are better discussed with an example. In the previous section, we outlined the most common numerical methods starting from the interest in solving the fundamental heat conduction equation. To study even the apparently simple problem of solidification of a pure material, that equation is coupled with two other conditions to form what is known as the *Stephan problem*. One of these conditions is of particular interest since it expresses the correction that one should introduce in describing the temperature field away from the solidification front – that is, the interphase boundary. This is the so–called *Gibbs–Thomson equation*⁴, which specifies the interface temperature. It reads

$$T_{\text{int}} = T_m - \frac{T_m}{L_f} \gamma \kappa - \frac{V_n}{\mu}, \quad (1.3)$$

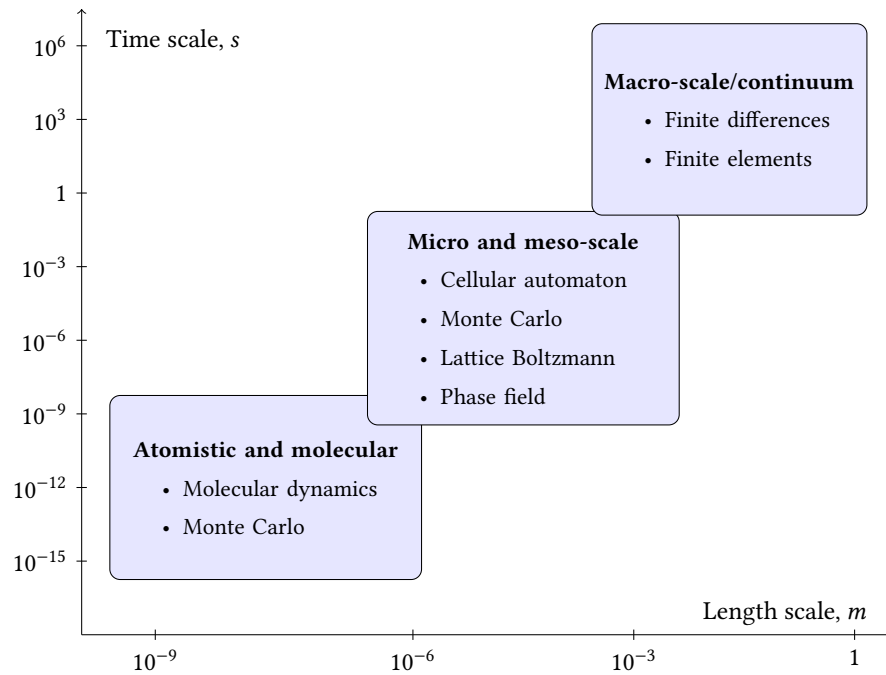


Figure 1.3 – Schematic plot showing time and length scales of several modeling techniques that should be combined to study solidification.

where T_m is the melting temperature, γ is the solid–liquid interface free energy and κ is interface local curvature and L_f the latent heat of fusion. V_n is the velocity of the interface along its normal direction and μ the atomic mobility, sometimes also referred to as kinetic coefficient. It is necessary to note that eq. (1.3) has been greatly simplified by omitting the dependency of γ and μ from the specific orientation of the interface normal. This aspect will be extensively discussed later on since it is essential to include the anisotropy of these two quantities to be able to model accurately real solidification phenomena.

The above equation contains several parameters that are both crucial and usually unaccessible by modeling methods at the macroscopic scale. However, the real difficulty is that the formulation of this problem is far from universal, and there are numerous examples of solidification–related problems for which the Stefan problem is unknown [10]. Moreover, searching for a solution to eq. (1.1) or eq. (1.3) by means of approximate methods as those described in the previous section usually involves numerical simulations with not so straightforward mathematics and rather complex algorithmic implementations.

Phase field modeling methods choose as their key quantity a continuous field analogous to the temperature term that appears in every macroscopic model. The phase field can be seen as a quantity that, from the standpoint of solid state physics, describes the degree of order in a certain phase; as such, it assumes constant values in the bulk and interpolates in a continuous fashion across the interface between different phases. A fundamental advantage over other continuum models is that it can be directly related to a description of the interface as a diffuse region between two portions of matter that are uniform on the atomic scale. Yet, the most interesting

feature of phase fields quantities is that they are connected to fundamental thermodynamics quantities known as free energy functionals⁵, from which the order parameter dynamics, heat and mass transfer can be derived by solving a set of partial differential equation. The parameters governing these equations are closely related to those appearing in a macroscopic model, and sometimes even adjusted so as to match their counterparts in the equations of the aforementioned Stefan problem.

However, nothing comes for free: the advantage of having a first-hand connection with atomistic methods — on which this thesis is entirely based — brings with it several limitations and challenges. For example, phase field methods are not capable of rendering interfaces that extend to the mesoscopic scale, which are necessary to describe real solidification microstructures. Moreover, as phase field models can be interpreted as the next step from atomistic simulations, they also have the intrinsic limit of the time scales they can cover. Interface kinetics strongly influences phase field models, which therefore cannot hope to transfer their results on time-scales that describe real phenomena.

Thesis outline

The rest of the thesis is grouped in two parts. Part 1 deals with all the established theoretical background. It covers a few fundamental concepts of statistical mechanics (chapter 2). It then describes the key ingredients of computer simulations (chapter 3), trying to balance theoretical and practical aspects. The third chapter deals with both equilibrium thermodynamics of binary systems (section 4.2) and fundamental thermodynamics of interfaces (section 4.3), with an outlook to out-of-equilibrium conditions.

Part 2 contains the original contributions of the thesis: the first chapter (chapter 5) describes the adaptation of a well-known method to compute interface free energies to instantaneous fluctuating dividing surfaces, while the second chapter (chapter 6) presents a versatile framework to determine interface and solidification related properties of binary alloys.

Notes

1. See, for example, Williams and Fatemi [11].
2. A stencil represents a discretize version of the solution to a second-order partial differential equation, which require two boundary conditions to be specified.
3. For a more detailed discussion see, for example, chapter 6 (section 2.2) of Dantzig and Rappaz [12].
4. The Gibbs–Thomson equation relates a change in pressure, temperature, or chemical potential across the interface between two phases. A first approach to describe phase transitions with textbook thermodynamics neglects any additional contribution to the free energy due to the presence of an interface, but this approximation becomes insufficient if the surface free energy is comparable to the bulk energy. A justification of the Gibbs–Thomson equation can be given considering a simple system where a crystalline phase with chemical potential μ_s is a sphere of radius R surrounded by a liquid phase with chemical potential μ_l . If γ is the specific free energy of the surface, then the total surface energy of the system is simply $4\pi R^2\gamma$. Assuming that the solid is thermodynamically more stable than the liquid, if dn moles of particles (atoms or molecules) transfer from the liquid phase to the solid, then the free energy of the system will be reduced by $(\mu_l - \mu_s) dn$. As a consequence, the radius of the spherical crystalline

phase will change by an amount given by $dR = V_m dn / 4\pi R^2$, where V_m is the molar volume of the solid. An increase in the solid phase will also change the surface energy by $\gamma dA = \gamma 8\pi R^2 dR$. These two contributions are opposite in sign, therefore thermodynamic equilibrium corresponds to the condition where the two free energy terms are equal. By comparing the two equations, one gets

$$(\mu_l - \mu_s) dn = 8\pi R \gamma \frac{V_m}{4\pi R^2} dn,$$

that is equivalent to

$$\mu_l - \mu_s = \frac{2\gamma V_m}{R},$$

and this is precisely the Gibbs–Thomson equation. Even though we assumed a spherical crystalline phase, the reasoning holds also for a non-uniform distribution of the particles that undergo the phase transition, and it is always possible to express the difference in chemical potential between the two phases in terms of a shift in temperature, pressure, or concentration from their corresponding equilibrium values. The additional term on the left-hand side of eq. (1.3), V_n/μ , is a kinetic term which is needed to account for the fact that the interface does not respond instantaneously to a change in temperature or concentration. For rough interfaces, it can be shown [13] that a linear term in the interface velocity corresponds to an exponential relaxation of the interface due to a perturbation, and this is usually the approximation that leads to form of the Gibbs–Thomson equation discussed in this thesis.

5. We used the term “functional” on purpose because phase field models depend on quantities which in turn depend on functions on a entire domain and not just on individual points. For example, a quantity written as $F = \int f(\phi(x)) dx$ defines a functional of $\phi(x)$, and it is usually denoted with $F[\phi(x)]$. Functionals are well known to the physics of condensed matter and have a fundamental role in *density functional theory*, a keystone of modern computational physics. It is perhaps interesting to note that a certain type of phase field models can be shown to derive from special cases of the so-called *classical density functional theory*. This link has the appealing notion to draw parameters of phase field models directly from atomistic properties. References [14, 15] are examples of the applicability of classical density functional theory to solidification problems.

Theory **Part I**

2 Statistical mechanics¹

2.1 Phase space and ensembles

When studying a portion of condensed matter, most common questions we are often interested in answering are, for example, *what is the value of that property* (e.g., the diffusion coefficient) *at a given pressure and temperature? How the structural ordering* – which could be captured by a quantity like the structure factor – *would change if an external perturbation is turned on?* These examples are trivial, but answers to all this kind of questions generally depend on the state of the system which is defined by a certain number of parameters – such as the temperature T , the pressure p , and the number of particles N ¹ – that can be more or less controlled at will. At the microscopic (yet not subatomic) level, each particle composing the system is uniquely identified once its position and momentum are defined. However, any macroscopic quantity we might want to study, even those providing information at smaller scales, is not dictated by the astronomically large number of positions and momenta that define the mechanical state of the system. The connecting bridge between the over-detailed realm of positions and momenta – a multidimensional coordinate space that is called the *phase space* – and the macroscopic world was built by *statistical mechanics*.

With a system being in a particular state, given by a particular point in phase space Γ , the instantaneous value of any property A will depend on Γ . Since $\Gamma \equiv \Gamma(t)$, A is also going to change over time. The natural definition of a measured value of A would be obtained by collecting multiple observations of A in a given time interval and computing the average. In the limit of an infinite number of observations, an average so calculated would match the true average. With a limited number of measurements,

$$A_{\text{obs}} = \langle A \rangle_{\text{time}} = \frac{1}{N_{\text{obs}}} \sum_{i=1}^{N_{\text{obs}}} A(\Gamma(t)).$$

From this definition, one could determine the average value of A by letting the system evolve in time-steps of fixed length, thus obtaining a discrete time evolution of A .

Instead of calculating $\langle A \rangle$ with a time average of $A(\Gamma(t))$, another potential way relies on the concept of *ensemble average*, which has been originally coined by Gibbs. In this context, the

¹Additional readings for some of the topics of this chapter: Ruelle [16], Allen and Tildesley [17], and Ruelle [18].

term *ensemble* indicates a set of points, or configurations, in phase space. These configurations are not distributed randomly, but they are so according to some probability distribution function, $\mathcal{P}(\Gamma)$, which strictly depends on the set of macroscopic quantities that define the state of the system. For example, we may choose to study a system whose temperature and pressure are held constant, thereby defining the so-called NPT ensemble. A different thermodynamic state of the system changes the type of ensemble we are working with; hence, the probability distribution may be labeled accordingly as \mathcal{P}_{ens} . More importantly, since every point in phase space evolves in time, \mathcal{P} will change with time as well.

Thermodynamics (chapter 4) deals with systems attaining equilibrium. These particular states of the system are particularly relevant for the probability distribution: by definition, equilibrium means not changing with time; hence $\partial\mathcal{P}(\Gamma)/\partial t$ must vanish. Among the infinitely many ways in which a system may evolve over time, there is a kind of such processions that is peculiar: the one that allows the system to explore all the region of phase space. These evolutions are termed *ergodic*, and we shall return to the importance of this concept. For the moment, this picture of the system as a point in phase space that jumps from one configuration to another suggests a way to define the ensemble average: since the likelihood that the system visits a particular configuration is determined by the probability distribution function, which in turn depends on the ensemble, we may compute the average value of A as

$$A_{\text{obs}} = \langle A \rangle_{\text{ens}} = \sum_{\Gamma} A(\Gamma) \mathcal{P}_{\text{ens}}(\Gamma).$$

Since in general not all the configurations are equally likely, the probability distribution function is often replaced with an unnormalized function that expresses the absolute weight $w_{\text{ens}}(\Gamma)$ of that particular configuration. To make this change consistent, the weights must sum to a quantity defined in such a way to guarantee that $\sum_{\Gamma} \mathcal{P}(\Gamma) = 1$. This fundamental quantity is called *partition function* and it clearly depends on the kind of ensemble as well. A formal definition of ensemble average is then

$$\langle A \rangle_{\text{ens}} = \frac{\sum_{\Gamma} w_{\text{ens}}(\Gamma) A(\Gamma)}{Z_{\text{ens}}}, \quad (2.1)$$

where $Z_{\text{ens}} = \sum_{\Gamma} w_{\text{ens}}$ is the partition function. From the partition function, one may build the connection with classical thermodynamics by introducing a *thermodynamic potential* of particular importance for one ensemble,

$$\Psi_{\text{ens}} = -\ln Z_{\text{ens}}.$$

Ψ_{ens} is not a new quantity introduced with no purpose, but it is the function that attains its minimum value when the system is at thermodynamic equilibrium. For an ensemble in which the total energy is constant, Ψ_{ens} might be defined as the entropy taken with opposite sign, or, as the chapter on thermodynamics will show, it might coincide with a free energy function. These functions are particularly relevant for ensembles most relevant for typical laboratory conditions, where thermodynamic variables such as the pressure and the temperature are usually held

constant.

2.2 Common ensembles

The definition of the partition function given by eq. (2.1) suggests a possible approach to computing thermodynamics properties of a system: it would involve a way to determine directly Z_{ens} . However, the system contains too many particles, and therefore there are too many states (i.e., configurations in phase space) that will have a negligible weight. Ideally, we would discard those configurations if only we could know beforehand which are those states. Unfortunately, this direct way of approaching the problem is unfeasible. Therefore, a computer simulation experiment consists of a time-evolution equation that takes potentially any initial configuration of the system and, given sufficiently long sampling time, allows computing an ensemble average that converges to the average one would obtain if it were possible to pick initial configurations with the right statistical weight.

The choice of the ensemble should be dictated mostly by convenience – that is, it should be the most suitable choice for the kind of system the performed simulation should reproduce. However, a further question to address is whether the average of property A calculated in one ensemble is equal, in some circumstances, to the average of the same property in another ensemble. Assuming that the conditions for a positive answer to this question are satisfied, the rest of the section summarizes four of the most common ensembles adopted for computer simulations. One should bear in mind that, once some thermodynamic properties are fixed by definition in one ensemble, all the other quantities that depend of the configurations in phase space will deviate from their average.

2.3 The NVE ensemble

A system composed by classical particles obeys Newton's equations of motion, which conserve the total energy, the total linear and angular momentum. If the initial total energy is E , then the probability distribution function should represent all those configurations that have that particular value of E . With $\mathcal{H}(\Gamma)$ being the Hamiltonian, the partition function is given by

$$Z_{\text{NVE}} = \sum_{\Gamma} \delta[\mathcal{H}(\Gamma) - E], \quad (2.2)$$

where δ is a function that picks only those configurations of the system with the given value of the total energy E . Dealing with macroscopic systems means having an almost infinite number of possible states, and the six-dimensional coordinates in phase space vary in a continuous fashion². The partition function of any ensemble is thus better expressed with an integral

$$Z_{\text{NVE}} = \frac{C}{N!} \int \delta[\mathcal{H}(\mathbf{r}, \mathbf{p}) - E] d\mathbf{r}d\mathbf{p}, \quad (2.3)$$

where C is a constant required to make the partition function dimensionless, and $1/N!$ takes into account the indistinguishability of the N particles³.

A straightforward implementation of molecular dynamics (section 3.2) will sample exactly this ensemble – also called *microcanonical ensemble* – and the thermodynamic potential that tells whether the system is evolving towards equilibrium is the negative entropy, $-S/k_B T$, where k_B is the Boltzmann constant.

2.4 The NVT ensemble

A more common ensemble is that in which the temperature is held constant, instead of the total energy. This is the *canonical ensemble* and its probability distribution function is proportional to

$$\exp[-\mathcal{H}(\mathbf{r}, \mathbf{p})/k_B T].$$

The partition function is, then

$$Z_{\text{NVT}} = \frac{C}{N!} \int \exp[-\mathcal{H}(\mathbf{r}, \mathbf{p})/k_B T] \, d\mathbf{r} d\mathbf{p}. \quad (2.4)$$

Newton equations of motion evolve in time initial configurations in such a way that they will explore only those regions of phase space with total energies close to their initial values, and the temperature will then fluctuate. A different strategy to evolve configurations from the NVT ensemble is needed, one that takes into account the possibility for any of these configurations to exchange energy. A couple of approaches to devise such integration schemes will be discussed in the section on thermostats and barostats.

The final important detail about this ensemble is which thermodynamic potential is more naturally to consider. Being V and T fixed, the thermodynamic function is the Helmholtz free energy, A

$$\frac{A}{k_B T} = -\ln Z_{\text{NVT}}.$$

2.5 The NPT and μ VT ensembles

The ensemble in which pressure and temperature are fixed is usually called isobaric–isothermal ensemble. Avoiding for the moment any technical discussion on how this ensemble is implemented in a molecular simulation, the substantial difference with *NVT* is in the partition function. Since now the total volume may vary, an additional integration over all the possible volumes must be added, introducing a normalizing constant volume V_0 needed to have a dimensionless partition function.

$$Z_{\text{NPT}} = \frac{C}{N!} \frac{1}{V_0} \int dV \int \exp[-(\mathcal{H}(\mathbf{r}, \mathbf{p}) - pV)/k_B T] \, d\mathbf{r} d\mathbf{p}. \quad (2.5)$$

The fundamental thermodynamic function associated with this ensemble will return many times in the following since it is of great importance in chemical thermodynamics. It is the Gibbs free energy, which is the equivalent to Helmholtz free energy when pressure is constant

instead of volume:

$$\frac{G}{k_B T} = -\ln Z_{\text{NPT}}.$$

It always a good choice to fix at least one extensive variable to avoid pathological ensembles in which the system could grow without limit. The extensive variable to constrain could be the volume, while the number of particles can change. The grand canonical – or μ VT – ensemble allows the total number of particles to change and fix the total chemical potential μ . The partition function is very similar to that of the NPT ensemble with an additional summation over N :

$$Z_{\mu\text{VT}} = \sum_N \frac{C}{N!} \exp(\mu N/k_B T) \int \exp(-\mathcal{H}(\mathbf{r}, \mathbf{p})/k_B T) d\mathbf{r} d\mathbf{p}. \quad (2.6)$$

In a similar way in which in the NVT and NPT ensembles the total energy – or the total energy and the volume – can vary, the scheme adopted to generate and evolve configurations from the μ VT ensemble must allow the system to exchange particles with a reservoir, and it must keep the overall chemical potential constant during these exchanges.

Although the discussion was limited to a system composed by only one type of atoms, extending the grand canonical ensemble to a system containing C components, each comprising a set of N_i particles, is rather straightforward: the single summation over N of eq. (2.6) is replaced by a multiple summation over the atoms of all the possible kinds, and the exponential term outside the integral becomes

$$\prod_{i=1}^C \frac{\exp(\mu_i N_i/k_B T)}{N_i!}.$$

A slightly modified alternative to the grand canonical is the *semi-grand canonical* ensemble. This ensemble is particularly useful to study systems which contains both multiple chemical species and several phases, each with a given composition. Such a system could be well represented by a semi-grand ensemble in which the total number of particles $N = \sum_i N_i$ is fixed, but the chemical species of each particles can change. In this case, an arbitrary reference for chemical potentials must be defined – say μ_1 , the chemical potential of species 1 – while the differences $\mu_i - \mu_1$ with the potentials of all the $N - 1$ remaining species are held constant.

2.6 The ergodic hypothesis

As described at the beginning of this chapter, in statistical mechanics physical observables are obtained by ensemble averages, that is, averages calculated on all the states accessible to the system. However, when dealing with real experiments, quantities are usually the outcome of an average of a series of measurements collected over a given time interval. As it will be shown in the following sections, the aim of well-established computational techniques such as *molecular dynamics* (MD) is precisely that of providing a reliable numerical approximation of the time evolution of a many-particle system. The question is more about whether there is a link between ensemble averages and those averages obtained when running real experiments.

It turns out that a link exists, even though it cannot be proven right for any arbitrary system under study. The idea, known as the *ergodic hypothesis*, is due to Ludwig Boltzmann, one of the physicists who gave birth to modern statistical mechanics. Boltzmann thought that, for an isolated system that cannot exchange neither energy nor particles with the surroundings, the system would visit, over the course of time, all the possible energy configurations. This means that each particle in the system would have to assume all the values of position and velocity that are compatible with a given total energy and, waiting long enough, these configurations would all be visited. In other words, saying that a system's evolution is ergodic means that the system will return to an arbitrarily small distance from any permitted energy configuration.

Since a point in the phase space represents a particular state of the system, its time evolution is simply the motion of that point in the phase space. The physical content of the ergodic hypothesis is that, while moving in the phase space, the point representing our system spends in any region of the space a time proportional to the volume of that region.

In practice, if the time evolution of the system is ergodic, an ensemble average of any dynamical property A is equivalent to its time average taken over the system trajectory, if the latter is available for a sufficiently long time:

$$\langle A \rangle_{\text{ens}} = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t A(X(t')) dt'.$$

2.7 Statistical errors

The previous section underscored the importance of the ergodic hypothesis: if it is applicable to the time evolution of our system, then it provides a practical way to estimate averages out of trajectories obtained with MD or MC. If calculating an average quantity seems now to be straightforward, estimating the statistical errors is a more delicate topic.

If a simulation runs for a total time length T , for a ergodic time evolution

$$\langle A \rangle = A_T = \frac{1}{T} \int_0^T A(t) dt$$

where the subscript distinguishes the time-averaged A from its ensemble average. The variance is given by

$$\begin{aligned} \sigma^2(A) &= \langle A_T^2 \rangle - \langle A_T \rangle^2 \\ &= \frac{1}{T^2} \int_0^T \int_0^T \langle [A(t) - \langle A \rangle][A(t') - \langle A \rangle] \rangle dt dt' \end{aligned} \quad (2.7)$$

where the integrand is the correlation function of time fluctuations in the quantity A , $C_A(t - t')$. Assuming that T represents a much greater time scale than that characterizing the decay of

fluctuations, then the integral of eq. (2.7) can be approximated as

$$\begin{aligned}\sigma^2(A) &\approx \frac{1}{T} \int_{-\infty}^{+\infty} C_A(t) \\ &= \frac{2\tau_C}{T} C_A(0),\end{aligned}\tag{2.8}$$

in which the usual definition of the decay time τ_C as half the integral of the normalized correlation function $C_A(t)/C_A(0)$ has been used. Therefore, the variance relative to the mean is given by

$$\frac{\sigma^2(A)}{\langle A \rangle^2} \approx \frac{2\tau_C}{T} \frac{\langle A^2 \rangle - \langle A \rangle^2}{\langle A \rangle^2}.\tag{2.9}$$

The meaning of eq. (2.9) restates the known result that the error in the mean decreases as the number of measurements taken long after the typical decay time of the fluctuations of A increases. This number is given by the ratio T/τ_C , and it gives a precise meaning to the vague concept of “running a simulation long enough.” A practical problem arises when the sampled quantity involves fluctuations decaying over a long wavelength limit, which require to perform simulations of systems with a box size much larger than all the characteristic length scales of such fluctuations. This fact reformulates the problem in terms of the computational cost of the simulation: since the computational cost scales with the number of particles at best, the time required by such a calculation rapidly increases as a power law of the simulation box length. It is therefore not a problem of being able to reach prohibitive time scales, but rather that of having sufficient computational resources to study a system large enough to sample all the relevant fluctuations of the observable of interest. If a too small system size suppresses some of these fluctuations, then the observable will be affected by finite-size effects.

The relation given by eq. (2.9) suggests a way of verifying the ergodicity of a time evolution of the system: by measuring both the correlation time⁴ and the time average of the observable A as the simulation proceeds, the error in the mean would get smaller and smaller for an ergodic evolution of A . There are sampling techniques that go beyond simple MD or MC that aim at minimizing τ_C for the observables under study. However, there are a few subtle problems worth mentioning.

1. First of all, computing the auto-correlation time from simulations is not trivial. A reason for this is that τ_C is usually estimated from the integral of the auto-correlation function, and the latter is determined without knowing exactly the value of $\langle A \rangle$, which is inferred from the simulation trajectory itself. This is true for the majority of real cases, therefore the resulting auto-correlation time will be underestimated.
2. If several observables are monitored at the same time, it often happens that the two quantities have very different auto-correlation times: running a simulation for a time much greater than the shortest τ does not guarantee that the simulation has already lost memory of the fluctuations of the observable with a longer auto-correlation time.
3. An observable might be the sum of many contributions that exhibit different time scales. Therefore, care must be taken in judging the simulation to be long enough just because

the auto-correlation function shows a rapid, initial decay: integrating only this first portion of $C_A(t)$ might give a wrong estimate of τ because the long-time tail of $C_A(t)$ is discarded.

2.7.1 Block averages

Since it is not possible to run simulations for an infinitely long time, the concerns related to the estimation of errors are all but irrelevant: when reporting a measured quantity by means of a computer simulation, a quantitative estimation of its uncertainty must be given. It is easy to show that, if eq. (2.9) involves a quantity A for which time correlation plays no role, then $\langle A \rangle$ and its error are readily computed with standard formulas of statistics, and thus it is expected that the statistical error is inversely proportional to \sqrt{N} , where N is the number of measurements⁵.

When samples of A drawn from an ensemble distribution are instead correlated, a natural question arises: is there an alternative method to estimate the average and a correct uncertainty of $\langle A \rangle$, given that it is far more easier to compute from a simulation mean values rather than auto-correlation times? A smart way of obtaining reliable error estimates of time-correlated quantities has been proposed by Flyvbjerg and Petersen [19], and is known as *block average*. As its authors put it, the idea “combines maximum rigor with minimum computation,” because it involves an iterative algorithm that only requires a series of values of A drawn from a simulation’s trajectory. In fact, obtaining the ensemble average of an observable and its error through block averaging is equivalent to estimate the observable’s auto-correlation time.

Notes

1. For the present discussion, we are considering only one-component systems.
2. A macroscopic portion of matter can be considered as a molecular or atomic system in the so-called *thermodynamic limit*, that is, when the extensive variables such as the number of particles and the volume are brought to infinity while still having a well-defined, finite ratio. Thus, a system whose phase space contains an almost infinite number of states does not represent the thermodynamic limit of a system with a countable number of accessible states.
3. The integral $\int dr dp$ is a shorthand notation to indicate that the integration variables are the full set of spatial coordinates and momenta.
4. In this context, since we are studying correlations over time of an observable with itself, τ_C is also known as *auto-correlation time*.
5. This is true even when dealing with correlated samples. The only requirement is that one must suitably choose the time span that should separate consecutive samples of the observable. As showed by eq. (2.9) this spacing is given by T/τ_C .

3 Computer experiments

In most of this thesis, the simulations run to study the systems of interest belonged essentially to the field of *classical methods*. In this context, the term *classic* means that the motion of the particles making up the system obey the laws of Newtonian mechanics: once the initial conditions of the system are set and the law describing the interactions between particles is known, it is possible to determine the temporal evolution of the positions and momenta of each particle. Although the laws of quantum mechanics are necessary to explain the behavior of matter at the atomic scale, a simple argument entirely justifies the classical approximation made above: atoms are made up of nuclei which are much more massive than electrons. It is thus possible to decouple the dynamics of the former from that of the latter, and describe the dynamics of the system as a function of the positions and momenta of the nuclei only¹. The electrons with a much faster dynamics are assumed to follow instantaneously any displacement of the nuclei and to be always in their ground state.

These two assumptions lead to a description of the system in terms of its Hamiltonian \mathcal{H} , a function of $6N$ degrees of freedom that depends on the set of coordinates \mathbf{q} and momenta \mathbf{p}^2

$$\mathcal{H}(\mathbf{q}, \mathbf{p}) = \mathcal{K}(\mathbf{p}) + \mathcal{V}(\mathbf{q}). \quad (3.1)$$

The set of generalized coordinates \mathbf{q} and momenta \mathbf{p} depend on the kind of particles we are studying. If they are molecules that can be treated as rigid bodies, then the coordinates may be more suitably expressed as the positions of the centres of mass together with a number of variables describing the orientation of each molecule in space – for example, a set of Euler angles. When dealing with a simple atomic system, the kinetic term in the Hamiltonian is the simplest one, since it is just a sum over single-particle kinetic energies

$$\mathcal{K} = \sum_{i=1}^N \sum_{\alpha=x,y,z} \frac{p_{i,\alpha}^2}{2m_i}.$$

Since \mathcal{K} depends only on coordinates and mass of each particle, it does not contain any interesting information about the interactions between different particles, which is instead captured by the potential energy term, \mathcal{V} . Assuming that \mathcal{V} can be expressed in some form – which will almost always include a series of approximations more or less drastic depending on the problem – then it will be possible to write the full form of \mathcal{H} , from which we derive not

only the temporal evolution of the system, but also the equilibrium distribution of coordinates and momenta.

It is evident the the potential energy term is the key ingredient for any kind of computer experiments one wants to perform. Before diving into the details of the two standard techniques of molecular dynamics and Monte Carlo methods, the following section explores more in depth the available possibilities to construct interatomic potentials. It focusses more on the class of functions known as *empirical potentials*: these descriptions of the interactions between atoms or molecules are built as sufficiently flexible interpolations of experimental data or calculations obtained via *ab initio* methods³. To this latter class belong all those computational schemes that choose to avoid an effective interatomic potential, but they rather attempt to solve⁴ the associated many-body Schrödinger equation⁵.

It is worth mentioning another set of versatile methods that, even more recently than *ab initio* ones, have seen an incredible growth in the number of applications to the field of computational science: it is the broad class of *machine learning methods*⁶.

3.1 Interatomic potentials

Recent years have witnessed a spectacular growth in both the number and the capabilities of computational methods that allow investigating more deeply complex phenomena of condensed matter. One reason is undoubtedly the very rapid evolution of the available computational resources and power. At the same time, computer science has developed models – such as very specific programming languages – and increasingly efficient tools to make the most of all those resources.

However, in the scientific field, a motivation is undoubtedly linked to the more and more precise strategies for modeling the interactions between the constituents of a generic many-particle system. As mentioned earlier, the term indicated with \mathcal{V} contains all the interesting information about this aspect. Models like the *embedded atom method* (EAM) – described in detail the following section and most suited for systems showing metallic bonding – were developed less than thirty years ago [20–22], but have been continuously improved and extended in order to reduce the discrepancy between their predicted results and those obtained with much more consolidated experimental measurements.

Despite all this progress, many systems of interest for a variety of scientific fields are still too complex and demand reliable but approximate methods. A unified framework to explain all these possible interactions is both impossible and useless, because one would have to deal with an enormous quantity of information. Yet, any attempt to construct an interatomic potential should start by considering which terms play a relevant role in describing the properties of our system. For an atomic system of N particles, a general form for \mathcal{V} may be written as

$$\mathcal{V} = \sum_i v_1(\mathbf{r}_i) + \sum_i \sum_{j>i} v_2(\mathbf{r}_i, \mathbf{r}_j) + \sum_i \sum_{j>i} \sum_{k>j} v_3(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + \dots$$

In such expression one can identify: a first term representing an external field, which may be just the fixed walls of the constant-volume vessel where the system is contained. A second,

pairwise term that, albeit written as a function of individual atoms' positions \mathbf{r}_i and \mathbf{r}_j , depends only on the distance $r_{ij} = |\mathbf{r}_{ij}| = |\mathbf{r}_i - \mathbf{r}_j|$ between two particles. The third term captures the interactions between triplets, and it is essential when modeling molecular or even more complex systems (e.g., polymers or proteins), where bonds between atoms stretch and bend. One could also include higher-order terms, but both for a lesser physical importance and computational efficiency, most empirical schemes of interatomic potentials do not go beyond the third term.

Since already the three-body term usually needs time-consuming algorithms when implemented in some computer simulation program, the two-body term is often replaced by an effective term that collects all the many-body effects including those of v_3 . The obvious trade-off is how many approximations (and at what extent) this "effective" term introduces. The expression for \mathcal{V} then becomes

$$\mathcal{V} \approx \sum_i v_1(\mathbf{r}_i) + \sum_i \sum_{j>i} v_2^{\text{eff}}(r_{ij}).$$

The next section describes in more detail a scheme that has proved very effective in reproducing the metallic bond. Given the importance of simple metals and alloys in this thesis' scope, the section will give a brief review of the fundamentals and discuss some examples where the model obtained results with excellent agreement with experimental or first-principles calculations.

3.1.1 The embedded atom method

The question that sparked the idea on which the EAM model is based was: how could we formulate a scheme that allows avoiding both the need to solve the many-electron Schrödinger equation and the assumption of the cohesive energy of a solid as a pairwise sum?

The problem in considering the cohesive energy as a sum of pairwise terms is that it treats every bond between a pair of particles as independent from all the others. This assumption has been disproved in many ways by comparing the predicted and expected trends of some fundamental properties of solids. Daw, Foiles and Baskes [23] bring the example of the elastic constants of solids with a cubic crystal structure: if it were true that the cohesive energy could be well approximated as a sum of independent two-body terms, then the ratio between two elastic constants C_{12}/C_{44} would be approximately 1, the value it takes for a purely pairwise interaction. It is the case for rare gases, but when turning to simple metals such as palladium, silver and gold, that ratio lies more often between 2 and 4.

Once verified the reasonable guess that many-body interactions have a decisive contribution already for systems as simple as pure metals, it was necessary to find an approximate but accurate way to represent these interactions, a choice that did not require to solve explicitly the Schrödinger equation. In 1984, Daw and Baskes [20] proposed the following expression for the cohesive energy of a metal

$$E_{\text{coh}} = \sum_i \left[G_i \left(\sum_{j \neq i} \rho_j^a(r_{ij}) \right) + \frac{1}{2} \sum_{j \neq i} U_{ij}(r_{ij}) \right], \quad (3.2)$$

where G is the *embedding energy* defined as the interaction of atom i with the background

electron density originating from all the remaining atoms' electrons, ρ_j^a is an average atomic electron density, and U_{ij} is a two-body electrostatic interaction.

The derivation of eq. (3.2) will be only summarized, since many extensive works [21, 24] contain an integral discussion on its justification and verification. The starting point is the expression of the cohesive energy according to the density-functional formalism [25]

$$E_{\text{coh}} = G[\rho] + \frac{1}{2} \sum_{\substack{i,j \\ j \neq i}} \frac{Z_i Z_j}{R_{ij}} - \sum_i \int d\mathbf{r} \frac{Z_i \rho}{|\mathbf{r} - \mathbf{R}_i|} + \frac{1}{2} \int d\mathbf{r}_1 \int d\mathbf{r}_2 \frac{\rho(\mathbf{r}_1) \rho(\mathbf{r}_2)}{r_{12}} - E_{\text{atoms}}. \quad (3.3)$$

This expression contains: the repulsive interaction energy of the nuclei (2nd term), the interaction between the electron density and the positive nuclei (3rd term), and the Hartree energy. Everything else, that is, kinetic, exchange and correlations energies, are grouped in the $G[\rho]$.

The first assumption to obtain an expression easier to evaluate from eq. (3.3) is to write the electron density as a sum over atomic densities only⁷. With this assumption, the first term of eq. (3.3) splits into a difference between the energy of the solid and that of the single atom

$$G \left[\sum_i \rho_i^a \right] - \sum_i G [\rho_i^a],$$

where it is made explicit that the embedding energy G is a functional of the electron density. A way to approximate this expression is to assume that the background density of the atom i , $\rho_i^b(\mathbf{r}) \equiv \sum_{j \neq i} \rho_j^a(\mathbf{r} - \mathbf{R}_j)$, can be replaced by some constant value $\bar{\rho}_i$. By defining the energy of an atom embedded in a constant electron density field $\bar{\rho}_i$ as

$$G_i(\bar{\rho}_i) \equiv G[\rho_i^a + \bar{\rho}_i] - G[\rho_i^a] - G[\bar{\rho}_i], \quad (3.4)$$

one introduces an error E_{err} into eq. (3.3); the optimal value for $\bar{\rho}_i$ is that which sets $E_{\text{err}} = 0$. If the system contains atoms of a single chemical species and no defects, then only one value of $\bar{\rho}$ is needed. However, in the most general case the constant density term depends on the local environment of atom i . The considerable advantage of this approximation is that, instead of having to evaluate the functional $G[\rho]$, only the embedding function $G(\bar{\rho}_i)$ is necessary. The latter depends solely on how complicated the relationship between $\bar{\rho}_i$ and ρ_i^b is. Although many technical and relevant details have been skipped here, this justification of eq. (3.2) underscores at least two scenarios in which the approximations behind EAM would manifest immediately: one is when it is not possible to replace the total electron density with a linear combination of atomic densities⁸, and the second would be if there is no plausible form of the background density that is straightforward enough to evaluate.

In more recent years, there have been numerous works that extended and improved the method, and have made it possible to describe those interactions in metal systems that were not correctly captured by the original model. For example, Finnis and Sinclair [22] derived a different expression for the embedding energy term as a second moment⁹ approximation within the well-known *tight binding framework*. It can be proven that the bond energy term – i.e., the

embedding energy term – is proportional to $-\sqrt{\phi_i}$, where ϕ_i is proportional to

$$\phi_i = A_i \sum_j \beta_{ij}(r_{ij}),$$

where the suffixes ij accounts for the types of the atoms, leading to contributions to the bond energy that depend on the chemical nature of each atom, but without considering any orientational dependence of the bonds.

A second well-known example EAM extensions is the *modified embedded atom method* [26], where the term of the background density is not calculated as an isotropic average, but takes into account the possible angular dependencies of the electronic density [27, 28] – and therefore also that of the chemical bonds.

EAM applied to alloys

Interatomic potentials developed within the EAM method have been applied to study various properties of alloyed metals. Among the advantages of EAM over other semi-empirical approaches there is that the atomic interaction can capture the chemical nature of the atoms. In the Finnis–Sinclair model, the embedding energy is a function of electron density terms that explicitly depend on the atom types: changing the chemical identity of the neighbors will affect the electron densities. In principle, the EAM framework allows modeling an alloy once known the embedding functions that are able to describe correctly the metallic bond of the pure elements composing the alloy. However, even though in the original work of Daw and Baskes were given mixing rules to obtain $\rho_{\alpha\beta}$ and $\phi_{\alpha\beta}$ for any alloy once known the corresponding functions for the elemental systems, it was soon made clear that this approximation was not enough to produce reliable and transferable potentials with the desired accuracy, and the EAM was then revisited to fit explicitly properties of the complex system of interest.

One should also be aware of possible errors when applying EAM potentials to complex systems like alloys. For example, EAM is not capable of modeling system with large charge transfers, as well as any effect arising from particular features of the Fermi surface.

This section will present some useful and recent applications of the EAM to study different properties of metal alloys, with a particular focus on nickel and gold alloys.

Nickel ternary alloys As described previously, Ni-based superalloys are subject of many studies for their high performance applications as structural materials with excellent strength and creep resistance. These alloys are generally made of Ni, Ti and Al and the precipitation of the γ' ordered phase within the Ni FCC matrix is mainly responsible for the high strength. Also, plasticity response is thought to be due to formations of defects such as vacancies and impurities that affect dislocations motion [29, 30]. There is therefore an extensive literature that investigates the applicability of EAM to reproduce crucial properties of Ni-based alloys.

The study of nickel ternary alloys typically requires reliable potentials to model the interactions between Ni–Al, Ni–Ti and Al–Ti. The Ni–Al system has been already thoroughly studied, and several EAM potentials have been developed [31] for both the Ni₃Al system [32, 33] and

Table 3.1 – Lattice properties of Ni and Al calculated with the potential developed by Mishin [37] in comparison with experimental data [40, 41]. The perfect match between the predicted and experimental equilibrium energy and lattice parameter is due to the inclusion of these quantities in the dataset used to fit the EAM parameters.

Property	Ni		Al	
	Experiment	EAM	Experiment	EAM
a_0 (nm)	0.352	0.352	0.405	0.405
E_0 (eV)	-4.45	-4.45	-3.36	-3.36
<i>Elastic constants</i> (GPa):				
B	181.0	181.0	79.0	79.0
c_{11}	246.5	241.3	114.0	116.8
c_{12}	147.3	150.8	61.9	60.1
c_{44}	124.7	127.3	31.6	31.7

the NiAl system [34, 35]. Farkas [36] has developed and compared interatomic potentials for the Ti-Al system that included the non-central forces and the best potential available that did not include any angular term. In that work, Farkas demonstrated that “the strongest effect of the angular forces are seen in the values of the elastic constants,”¹⁰ and that both cohesive and stacking faults energies were much less affected by the addition of angle-dependent terms.

In the work of Mishin [37], a new EAM potential has been developed for the NiAl system by fitting to both experimental data and *ab initio* calculations. A series of recent works, aimed also at understanding the role of defects formation in plasticity response of Ni ternary alloys [38], have considered this new potential to be among the most reliable ones for modeling γ and γ' phases of Ni superalloys, showing a good accuracy in predicting γ/γ' interface energy.

Lastly, a ternary potential for the system Ni-Al-W has been constructed [39] for the Ni-based single crystal superalloys to model the doping of Ni-Al system with tungsten, also comparing results with other calculations when the doping elements were Re or Co¹¹. The potential predicts that tungsten solute atoms do not form clusters within the γ phase, which is consistent with experiments. Furthermore, it correctly predicts that increasing the amount of W reduces the lattice mismatch between γ and γ' phase.

Gold alloys Literature on EAM potential on gold and related alloys is less abundant than nickel (or titanium), especially when going beyond the simple metal system. Nevertheless, as we explained previously, the EAM framework allows to combine different single-element potentials that reproduce with the desired accuracy the interested properties, and build from these an effective potential for the alloy system.

The table below summarizes the accuracy of a new EAM potential for gold developed by Grochola, Russo, and Snook [42], in which the authors report to have used “an improved methodology to fit to lattice parameters of the solid at high temperatures as well as liquid densities.” This potential showed a better agreement with several experimental data if compared

Table 3.2 – Comparison table between predicted values of EAM potentials for gold discussed in this section. The table is partially reproduced from Grochola, Russo, and Snook [42]. Elastic constants are in GPa, energies in eV and lattice parameters in Å.

Property	Potential developed by				Exp.
	Grochola	Ercolessi	Johnson	Foiles	
Cohesive energy	-3.924	-3.78	-3.930	-3.927	-3.93
Lattice constant	4.0701	4.0704	4.0806	4.0805	4.07
Bulk modulus	1.8026	1.8037	1.6987	1.6673	1.803
$c_{11} - c_{12}$	0.3207	0.5998	0.2687	0.2454	0.319
c_{44}	0.4594	0.5998	0.4069	0.4524	0.454
Melting point (K)	1159	1338	1053	1121	1337

with previous potentials of Foiles, Baskes, and Daw [43], Johnson [44] and that obtained by Ercolessi, Parrinello, and Tosatti [45] within the so-called glue model. Grochola and coworkers emphasize how attempting a fit to all the properties for which an experimental data is available is by far unsuccessful. However, their results suggest that other metallic elements may gain an overall improvements by employing a different fitting methodology.

3.2 Molecular dynamics¹

Molecular dynamics (MD) is a computational technique that makes it possible to determine the equilibrium and dynamical properties of many-body classical systems.

The accuracy of the outcomes of MD simulations, the size of the systems and the typical time-scale on which the simulation can be performed all strongly depend of the kind of method employed to model the interaction between particles in the system. For example, an empirical potential allows the simulation of millions of particles on a time-scale of microseconds.

Evaluating the forces acting on every particle in a system from the knowledge of the interatomic potential is straightforward on paper, but it is, at the same time, the portion of a molecular simulation that represent the majority of its computational cost. Therefore, it must require a time that is orders of magnitude lower than that necessary to determine the potential of interaction by explicitly dealing with the electronic problem (i.e., solving Schrödinger equation). On the other hand, the accuracy and reliability of a classical MD simulation strictly depend on the goodness of the function chosen to represent the interatomic potential.

A simulation of MD can be thought as a computer version of a real experiment: initially a model of the system is prepared consisting of N particles; we determine the temporal evolution of the system, that is, we calculate positions and momenta for every particle as a function of time by solving Newton's equations of motion; finally we determine the properties of interest of the system.

¹More extensive discussions on the topics of this section can be found in Allen and Tildesley [17], Tuckerman [46], and Frenkel and Smit [47].

3.2.1 Equations of motion

There are many equivalent formulations of equations of motion that describe Newtonian dynamics of a system of particles whose interaction is expressed by a potential \mathcal{V} . Sticking to the formalism of generalized coordinates \mathbf{q} and \mathbf{p} , one may start from Hamilton's equations

$$\begin{aligned}\dot{\mathbf{q}} &= \partial\mathcal{H}/\partial\mathbf{p} \\ \dot{\mathbf{p}} &= -\partial\mathcal{H}/\partial\mathbf{q}\end{aligned}\tag{3.5}$$

where \mathcal{H} is the Hamiltonian. In eq. (3.1), the two contributions to the \mathcal{H} are explicit, while a more formal derivation is obtained from the Lagrangian $\mathcal{L} = \mathcal{K} - \mathcal{V}$

$$\mathcal{H}(\mathbf{q}, \mathbf{p}) = \sum_k \dot{q}_k \frac{\partial \mathcal{L}}{\partial \dot{q}_k} - \mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}).\tag{3.6}$$

From the above equation it is easy to show that, if there are no time-dependent forces acting on the system, the Hamiltonian is a constant of motion. Since \mathcal{H} represents the total energy of a closed system, it means that total energy is conserved¹².

Before showing the most common strategy to obtain an approximate solution to equations of motion, it is useful to point out that issues may arise if the potential energy term is not a sufficiently smooth function of particles' positions. This is because $\partial\mathcal{H}/\partial\mathbf{q}$ appears explicitly in Hamilton's equations, therefore at least the first derivative of $\mathcal{V}(\mathbf{q})$ needs to be continuous.

3.2.2 The Verlet algorithm

As mentioned in the introduction (section 1.4), finite differences methods are among the most employed techniques to obtain an approximate solution to problems described by several differential equations. Molecular dynamics belongs to this group.

A key quantity of any algorithm that implements finite differences is the timestep, Δt . Its choice depends on the details of the algorithm's implementation, but generally one would like to be able to use a timestep long enough to avoid excessively expensive simulations to study the evolution of the system on the time scale of interest. Moreover, in the specific case of MD, one would be able to choose an algorithm that, although the error introduced by a finite difference approximation will produce trajectories that differ considerably from the exact ones, does not violate the conservation of total energy. Lastly, the algorithm should be time-reversible and reasonably easy to implement in a computer code¹³.

No algorithm, however sophisticated, can provide a precise solution for an indefinite simulation time. The error made in the approximation will accumulate with each iteration, so it is necessary that the algorithm does not depend at least linearly on the timestep. Moreover, with MD we do not care at all to obtain an infinitely long trajectory, but rather *a sufficiently long one* with respect to the characteristic fluctuations of the process that we want to simulate¹⁴. If energy is not well conserved, for example because of a too large timestep, than the trajectory describing the system moving in phase space would not be sampling the NVE ensemble, and therefore any ensemble averaged property would bring no information at all.

An algorithm that satisfies all the aforementioned features and has been by far the most widely used in MD implementations is the *Verlet algorithm* [48] in the so-called velocity form.

$$\begin{aligned} \mathbf{v}\left(t + \frac{1}{2}\Delta t\right) &= \mathbf{v}(t) + \frac{1}{2}\Delta t\mathbf{a}(t) \\ \mathbf{r}(t + \Delta t) &= \mathbf{r}(t) + \Delta t\mathbf{v}\left(t + \frac{1}{2}\Delta t\right) \\ \mathbf{v}(t + \Delta t) &= \mathbf{v}\left(t + \frac{1}{2}\Delta t\right) + \frac{1}{2}\Delta t\mathbf{a}(t + \Delta t) \end{aligned} \quad (3.7)$$

This algorithm operates in three steps: after evaluating the forces to obtain accelerations at time t , velocities are calculated at time $t + \Delta t/2$. Then, positions are updated up to time $t + \Delta t$. Lastly, a second evaluation of forces allow to update also the velocities at time $t + \Delta t$. The velocity-Verlet algorithm is known to accurately conserve energy with a root mean-square error proportional to Δt^2 .

In the case of the microcanonical ensemble, it is worth repeating one last time that it is of fundamental importance to check that the algorithm chosen conserves the total energy, even if in fact small fluctuations of the energy are an inevitable consequence of the method of finite differences on which all the algorithms are based. For the velocity-Verlet algorithm, for example, it is known that the error should be proportional to Δt^2 , so a good strategy to check that everything is working as expected is to monitor the root mean square error on the total energy, $\sqrt{\langle \mathcal{H}^2 \rangle - \langle \mathcal{H} \rangle^2}$, versus Δt^2 ; as the timestep increases, the relationship should be linear. If that is not the case, a closer inspection to the trend of the total energy as a function of time could reveal a suspicious drift. One may then assess the nature of this drift by running a couple of short simulations of the same total time but with a smaller timestep: if energy drifts in the same way regardless of the timestep, then the problem probably lies in the particular configuration in which the system was prepared. On the other hand, a energy drift that changed substantially¹⁵ could indicate a problem in the implementation of the integration algorithm.

3.2.3 Isothermal molecular dynamics

We have seen that solutions to Newton's equations at the heart of standard molecular dynamics let the system explore constant energy surfaces in phase space. However, it is often more interesting to perform simulation by keeping constant other intensive variables, such as temperature or pressure. The reason for this is simple: many interesting phenomena, including solidification, actually happen in these conditions. The question is then how a MD algorithm can be modified in order to produce a system evolution in a different ensemble. We will briefly discuss some possibilities to obtain simulations at constant temperature and/or pressure in this section.

Before describing approaches that modify Hamilton equations and allow energy fluctuations, it is important to clarify what is meant by constant temperature. From the point of view of thermodynamics, the answer is contained in the Zeroth principle: one body can be brought at the same temperature of a second one simply by putting them in contact. In other words, we bring our system in contact with a much larger thermal bath whose temperature is the target

value we want our system to reach at equilibrium. In the context of atomistic simulations, on the other hand, it is straightforward to define a “kinetic” temperature from the equipartition theorem

$$2\langle\mathcal{K}\rangle = 3Nk_B T,$$

where N particles are assumed to have only 3 translational degrees of freedom. In such case, the average kinetic energy per particle is given by

$$k_B T = m\langle v_\alpha \rangle,$$

where α indicates the Cartesian component of the particle’s velocity and m its mass. One might therefore be tempted to conclude that keeping the kinetic energy per particle constant also forces the temperature to remain constant. It is not difficult to prove that this is not the case at all [47]: one can easily calculate the relative variance of kinetic energy for a system in contact with a thermal bath at a fixed temperature and note that it is nonzero. This means that directly manipulating particles’ velocities (and therefore their momenta) does not mean to sample the canonical ensemble. All these approaches are instead called *isokinetic* or *velocity-rescaling* schemes, and in general they only deviate negligibly from a true constant temperature simulation.

There are many properly designed velocity-rescaling schemes that let the system sample correctly the canonical ensemble: two examples are the scheme proposed by Evans and Morriss [49] and, more recently, a stochastic version of velocity-rescaling developed by Bussi, Donadio and Parrinello [50].

Andersen thermostat The Andersen thermostat [51] couples the system with a heat bath at a fixed, target temperature. This is achieved by alternating standard Hamiltonian dynamics for a certain number of timesteps with new assignments of momenta of a randomly chosen set of particles. The new momenta of these particles are sampled from the corresponding distribution at the target temperature, that is, Maxwell–Boltzmann distribution. The system’s dynamics differs from that of the microcanonical ensemble in that the system periodically receive a “kick” from the heat bath that makes it shift to a different constant energy surface.

Besides the target temperature, a key parameter of any implementation of Andersen thermostat is the collision frequency of the system with the heat bath. It is important to note that whichever frequency is chosen, the perturbation to the system’s dynamic introduced by these collisions are clearly unphysical, because every time the momenta of some particles are sampled anew, a sudden decorrelation affects the velocities of those particles. As a consequence, if the constant temperature MD simulation aims at determining some kind of dynamical properties – e.g., the self-diffusion coefficient – Andersen thermostat can lead to very different results for different values of its intervention frequency. An average of such dynamical properties would be close enough to the correct value only in the limit of very small collision frequency.

Nosé–Hoover thermostat The approach introduced by Nosé [52], whose most known formulation is that proposed by Hoover [53], provides a deterministic way – as opposed to the

stochastic approach by Andersen — to sample the canonical ensemble.

The idea is more convoluted than Andersen's in that there is no evident physical interpretation. Nosé–Hoover thermostat involves modifying the Hamiltonian and introducing two virtual quantities representing the position and momentum of an additional degree of freedom: the coupling of the system with the dynamics of this fictitious degree of freedom is what guarantees a canonical sampling of the phase space. As a determinist approach, the advantage of Nosé–Hoover scheme over Andersen thermostat is that a conserved quantity can be associated to the former, in the same way we showed that trajectories that follow Hamilton equations conserve the total energy.

The detailed derivation of Nosé Hamiltonian can be found in many standard textbooks [46, 47], but here it is more relevant to highlight the well-known issues of a deterministic thermostat like Nosé–Hoover's. First of all, the integration of equations of motion is not as straightforward as for original Hamilton equations, and it might even be the case that the integrator used is the most computationally expensive step in during the simulation. Moreover, it can be shown [54] that, in some circumstances, energy is not the only conserved quantity if the system evolves according to Nosé–Hoover's equations of motion. This in turn means that the partition function of the actual ensemble will not be the canonical one; additional degrees of freedom must be introduced to recover the correct ensemble, and therefore time integration schemes become more complex.

3.2.4 Isobaric molecular dynamics

If the system is supposed to simulate a process happening at constant pressure, then it is evident that the volume must fluctuate. Moreover, there are generally two kinds of volume fluctuations that may be more suitable for different simulations: isotropic, where the simulation box changes only in shape, and anisotropic. In the latter case, each dimension is usually uncoupled from the others, but sometimes one might want to hold the cross-sectional area of the simulation cell fixed, while letting the third dimension fluctuate. This could be the case for a multi-phase system, where a solid–liquid coexistence simulation is studied with a planar, slab-like geometry.

The very first proposition on how to perform constant pressure MD simulations is again due to Andersen [51]. His idea was to make volume a dynamic variable — to which to associate a mass that simulates the existence of a piston with its own inertia — and add a term PV term to the potential energy. An extended Hamiltonian can be derived in very similar way to Nosé–Hoover approach to the isothermal ensemble, and trajectories obtained from the associated equations of motion would sample the isobaric–isoenthalpic ensemble (NPH). The NPT ensemble is in practice the combination of this ensemble with a thermostat of choice.

3.3 Monte Carlo methods²

Monte Carlo methods are based on a smart intuition that Metropolis, Ulam and von Neumann had to solve complex problems of mathematical physics. The gist of their idea is that the solution to certain problems may be obtained in a more practical way by finding a probabilistic analogue, and then solving those problems with stochastic sampling.

A very simple example, that is nonetheless relevant to the calculation of quantities like the partition function, is to exploit Monte Carlo to evaluate complex integrals. The idea is to take the following prototype integral

$$F = \int_a^b f(x) dx,$$

and rewrite it in a way that highlights how the points of the integration domain are distributed

$$F = \int_a^b \left(\frac{f(x)}{\rho(x)} \right) \rho(x) dx.$$

One may then approximate the value of the integral by extracting N_{trials} random numbers from the probability distribution $\rho(x)$ and evaluating the function on this set of points. The simplest choice of a probability distribution function is a uniform distribution, that is $\rho(x) = (b - a)^{-1}$ with $a \leq x \leq b$. The integral would be then approximated by

$$F \approx \frac{b - a}{N_{\text{trials}}} \sum_{i=1}^{N_{\text{trials}}} f(x_i). \quad (3.8)$$

It is not hard to convince oneself that this is the only feasible way to estimate averages such as that of eq. (2.1). For example, the configurational part¹⁶ of Z_{NVT} can be written in a similar fashion to eq. (3.8)

$$Z_{\text{NVT}} \approx \frac{V^N}{N_{\text{trials}}} \sum_{i=1}^{N_{\text{trials}}} \exp(-\mathcal{V}_i/k_B T), \quad (3.9)$$

where i indicates one among the N_{trials} randomly chosen configurations. The more trial configurations we use, the better the approximation of the integral. Unfortunately, this idea is doomed to fail for a simple reason that we have already pointed out: there is no guarantee that many of the randomly chosen configurations have a not negligible exponential factor.

3.3.1 Importance sampling and the Metropolis scheme

The only viable way to estimate ensemble averages of the kind

$$\langle A \rangle_{\text{NVT}} = \frac{\int A \exp(-\mathcal{V}/k_B T) d\mathbf{r}}{\int \exp(-\mathcal{V}/k_B T) d\mathbf{r}}$$

²Further discussion about this topic can be found in Tuckerman [46] (chapter 7) and Frenkel and Smit [47] (chapter 3, 5, and 13).

is to sample configurations directly from the distribution of interest: this method is called *importance sampling*. The integrand of the above equation is just $\rho_{\text{NVT}}(\Gamma) A(\Gamma)$, and it will be significant only for those configurations where ρ is far from zero. The initial problem has not disappeared, it has only been reformulated in the following question: how can we generate configurations that follow a specific probability distribution? In general, we are looking for a iterative procedure that takes a initial configuration \mathbf{q} to \mathbf{q}' with the requirement that

$$\frac{\rho(\mathbf{q})}{\rho(\mathbf{q}')} = \exp[-(\mathcal{V}(\mathbf{q}) - \mathcal{V}(\mathbf{q}'))/k_B T].$$

The prescription to generate any next configuration is in essence a transition rule which tells the probability \mathcal{P} that the initial state $\mathbf{q}_i = \mathbf{q}$ transforms into the state $\mathbf{q}_{i+1} = \mathbf{q}'$. A necessary requirement of such rule is that it leaves invariant the underlying probability distribution. An alternative prescription, which is often the choice for many practical Monte Carlo schemes, is the so-called *detailed balance*. This rule simply asserts that the probability of forward and backward transitions between any pair of configurations is the same

$$\rho(\mathbf{q})\mathcal{P}(\mathbf{q} \rightarrow \mathbf{q}') = \rho(\mathbf{q}')\mathcal{P}(\mathbf{q}' \rightarrow \mathbf{q}).$$

A physical justification of the requirement of detailed balance is that, since the goal is to sample an equilibrium distribution, the transition rule must not destroy such equilibrium once reached. The strategy proposed by Metropolis [55] splits the transition rule in two steps: the first is a trial move, which could mean any attempt to change the initial configuration of the system, not only by displacing spatial coordinates; the second step is the acceptance–rejection of the attempted move. In this way, the overall transition probability is the product of the probability of each step: $\mathcal{P}(\mathbf{q} \rightarrow \mathbf{q}') = \mathcal{T}(\mathbf{q} \rightarrow \mathbf{q}') \times \mathcal{A}(\mathbf{q} \rightarrow \mathbf{q}')$, where \mathcal{T} and \mathcal{A} stand for the transition and acceptance probabilities. In practice, \mathcal{T} is taken to be symmetric, $\mathcal{T}(\mathbf{q} \rightarrow \mathbf{q}') = \mathcal{T}(\mathbf{q}' \rightarrow \mathbf{q})$, and the acceptance probability according to the Metropolis scheme is

$$\mathcal{A}(\mathbf{q} \rightarrow \mathbf{q}') = \min\left(1, \frac{\rho(\mathbf{q}')}{\rho(\mathbf{q})}\right). \quad (3.10)$$

In each case, $\mathcal{P}(\mathbf{q} \rightarrow \mathbf{q}')$ satisfies detailed balance. An important remark is that eq. (3.10) never requires to evaluate the full probability distribution, which would imply knowing the partition function, but only the ratio $\rho(\mathbf{q}')/\rho(\mathbf{q}) = \exp[-(\mathcal{V}(\mathbf{q}') - \mathcal{V}(\mathbf{q}))/k_B T]$, whose calculation only depends on the cost of determining the potential energy of a given configuration.

3.4 Enhanced sampling

Under the term *enhanced sampling* we may gather all those advanced simulation techniques that allow either going beyond the standard Metropolis prescription (section 3.3.1) – in which the states that the system explores during the simulation have Boltzmann-like probabilities – and extracting information on the free energy of the system, or sampling the so-called rare events, that is, those particular evolutions of the system usually characterized by energy barriers: the

system would naturally reside in a small number of points in phase space, while transitions from one point to an adjacent one cannot typically be observed with standard MD or MC.

Broadly speaking, the computational methods developed to tackle the problem of studying rare events can be grouped into three categories: methods that study the system in equilibrium, non-equilibrium techniques [56], and a third one which gathers those methods that choose a particularly suited set of reaction coordinates — also referred to as *collective variables* — and calculate the free energy as a function of those coordinates¹⁷.

Restricting the discussion to the first class, that is, methods that sample the system in equilibrium, to allow the system to overcome an energy barrier one can either alter its energy to reduce the barrier, or constrain the system to sample the phase space only along those degrees of freedom that do not characterize the transition over the barrier.

The latter method is the so-called *thermodynamic integration* [57–59], which we only briefly comment here. Recalling that free energy functions are the thermodynamic potentials associated with different ensembles, it is clear Helmholtz free energy, for example, cannot be determined as a simple ensemble average. Instead, one could determine the derivative of the free energy with respect to some thermodynamic variable such as the volume or temperature: these are quantities that can be obtained readily from a numerical simulation. Then, once these quantities are known, the free energy can be evaluated by integration over a given path where the thermodynamic variables of choice were allowed to change. When dealing with computer experiments, one is free to choose any parameter, not only those characterizing a physical transformation that could be studied with a real experiment. The only requirement is that, for a given value of the chosen parameter, the system should correspond to one for which the free energy is known either analytically or numerically with reasonable accuracy.

Changing the potential energy surface that characterizes the system is the main ingredient of *biased sampling techniques*. A well-established method of this kind is the so-called *umbrella sampling* developed by Torrie and Valleau [60, 61]. Their idea was to estimate free energy differences by altering the way the system explores the phase space; this is done by adding an external bias potential on top of that describing the physical interaction between particles of the system. The name of this technique comes precisely from the presence of the external bias, which connects regions of the phase space that are separated by energy barriers.

A crucial step of this method is to define the appropriate bias potential in terms of some reaction coordinate, which could be defined either on geometrical considerations — as when dealing with a chemical reaction in which products and reactants differ in conformation or configuration — or according to an order parameter that distinguishes clearly between the two states among which one wants to calculate the free energy difference. Moreover, it must be possible to subsequently remove the effect of the external bias on the probability distribution function and to calculate the free energy of the system as a function of the reaction coordinate.

If ξ is the chosen reaction coordinate, then the probability distribution function for the canonical ensemble depends on ξ according to

$$\rho(\xi) = \frac{\int \delta[\xi'(\mathbf{r}) - \xi] \exp[-\beta\mathcal{V}] d\mathbf{r}}{\int \exp[-\beta\mathcal{V}] d\mathbf{r}}, \quad (3.11)$$

where the integration is carried out over all degrees of freedom except ξ , and the Boltzmann factor $\beta = 1/k_B T$ has been introduced. The free energy – the Helmholtz free energy, but the argument that follows is valid even for a constant pressure ensemble in which the thermodynamic potential is the Gibbs free energy – will also depend on ξ . The bias potential term $W(\xi)$ is, by definition, the difference between the potential energy of the biased system and the unbiased one

$$\mathcal{V}(\mathbf{r})^B - \mathcal{V}(\mathbf{r})^U = W(\xi),$$

where the superscripts U and B indicate a quantity in the unbiased or biased system, respectively, and a term with no superscript always represents an unbiased quantity. The probability distribution function of the biased system is, analogously to eq. (3.11),

$$\rho^B(\xi) = \frac{\int \delta[\xi'(\mathbf{r}) - \xi] \exp\{-[\beta\mathcal{V}(\mathbf{r}) + W(\xi'(\mathbf{r}))]\} d\mathbf{r}}{\int \exp\{-\beta[\mathcal{V}(\mathbf{r}) + W(\xi'(\mathbf{r}))]\} d\mathbf{r}}, \quad (3.12)$$

Since the bias potential depends only on ξ , one has

$$\rho^B(\xi) = \exp[-\beta W(\xi)] \times \frac{\int \exp[-\beta\mathcal{V}(\mathbf{r})] \delta[\xi'(\mathbf{r}) - \xi] d\mathbf{r}}{\int \exp\{-\beta[\mathcal{V}(\mathbf{r}) + W(\xi'(\mathbf{r}))]\} d\mathbf{r}}. \quad (3.13)$$

Combining eqs. (3.11) and (3.13), one may write the unbiased probability distribution function as

$$\rho^U(\xi) = \rho^B(\xi) \exp[-\beta W(\xi)] \times \frac{\int \exp[-\beta\mathcal{V}(\mathbf{r})] \exp[-\beta W(\xi'(\mathbf{r}))] d\mathbf{r}}{\int \exp[-\beta\mathcal{V}(\mathbf{r})] d\mathbf{r}}, \quad (3.14)$$

in which the fractional term is the ensemble average of $\exp[-\beta W(\xi)]$. Hence, finally

$$\rho^U(\xi) = \rho^B(\xi) \exp[-\beta W(\xi)] \langle \exp[-\beta W(\xi)] \rangle, \quad (3.15)$$

from which the free energy as a function of ξ may be written as

$$A(\xi) = -1/\beta \ln \rho(\xi)^B - W(\xi) - 1/\beta \ln \{ \langle \exp[\beta W(\xi)] \rangle \}. \quad (3.16)$$

The previous relation allows determining the free energy function from the knowledge of the probability distribution function in the presence of the bias and the value of the bias itself.

Notes

1. This is the so-called *Born–Oppenheimer approximation*.
2. A comment on notation: wherever a symbol \mathbf{q} or \mathbf{p} appears, it means the set of N vector coordinates or momenta, that is, \mathbf{q} (or equivalently \mathbf{p}) stands for (q_1, q_2, \dots, q_N) .
3. More often than not, empirical potentials are built on a joint dataset, where many interesting physical properties that the potential should be able to reproduce have been assessed with either experimental measurements or first-principles simulations, or sometimes even both of them.

Chapter 3. Computer experiments

4. We used the word “attempt” to underscore that in almost all the practical cases of interest a solution to the many-body Schrödinger equation is an approximate solution.
5. Many of these methods rely on density functional theory. See for example Martin [62] for a theoretical discussion with a focus on computational aspects, or the reference textbook by Marx and Hutter on ab initio molecular dynamics [63].
6. This field has grown so vast recently that it is almost impossible to give an exhaustive selection of introductory references. Therefore, I will point the interested reader to the recent review of Ceriotti [64] on some of the possibilities of machine learning techniques applied to atomistic simulations data.
7. This assumption can be relaxed (see Daw [24]) without invalidating the expression of eq. (3.2).
8. This would be the case of all transition metals, except those at the very beginning and end of the series.
9. The second moment of the the local density of states, $D_i(\epsilon)$, which contains the contributions from all the orbitals of an atom at site I . The second moment is then: $\mu_2^I = \int_{-\infty}^{+\infty} \epsilon^2 D_i(\epsilon) d\epsilon$. See Finnis [65] for a complete discussion on second moment models.
10. See again Farkas [36].
11. W, Re, Co, Mo and Va are common additives for Ni superalloys since they behave as solid-solution strengtheners both for γ and γ' phases.
12. This can be readily seen by writing the total derivative of any function $F(\mathbf{q}, \mathbf{p})$ as

$$\frac{dF}{dt} = \sum_i \left(\frac{\partial F}{\partial \mathbf{q}_i} \dot{\mathbf{q}}_i + \frac{\partial F}{\partial \mathbf{p}_i} \dot{\mathbf{p}}_i \right) + \frac{\partial F}{\partial t}.$$

In our case, $F \equiv \mathcal{H}$, and if \mathcal{H} has no explicit time dependence, $\partial \mathcal{H} / \partial t = 0$. Using eq. (3.5) one gets $d\mathcal{H} / dt = 0$, which means that the Hamiltonian is a constant of motion.

13. It should be easy to translate in any procedural programming language. Besides, it should be fast and require the storage of as little memory as possible.
14. The important concept to stress here is that the purpose of molecular dynamics is not that of providing *exact* trajectories of the system, but those that are *statistically* meaningful, for we are only interested in quantities calculated as ensemble averages.
15. A smaller timestep would generally reduce such energy drift.
16. For the canonical ensemble, it is always possible to factorize the energy in a kinetic plus a potential energy term. Hence, the partition function of eq. (2.4) is the product of one integral over the momenta and one over the positions. The latter is usually referred to as *configurational integral*.
17. A well-known example of this family of methods is *metadynamics* [66], whose application to the calculation of interface properties will be discussed later (section 4.3.2).

4 Thermodynamics¹

4.1 Fundamentals

In 1824 the French physicist Sadi Carnot published his only one book, some hundred pages collected under the name of *Réflexions sur la puissance motrice du feu*. He was the son of the eminent military engineer Lazare Carnot, and was fascinated by the possibility of transforming a rather useless form of energy, heat, into one that could be exploited for a purpose, mechanical work. Like Nicolaus Copernicus, with his treatise on heat as a workforce Carnot laid the foundations for a new science that later took the name of *thermodynamics*.

Although Carnot drew several of his conclusions from completely wrong concepts — such as the idea that heat was some sort of a fluid that was literally flowing from one body to another — it is noteworthy how he was still able to arrive to the fundamental law that states the limitations of the transformation of heat into work; a law that, after many rephrasing, has been given the name of Second Law of thermodynamics.

In the days of Carnot and the scientists who followed, the atomic theory had not yet been fully accepted. Today we know with certainty that matter is made up of atoms, which are in turn “divisible” into elementary particles — so quantities such as temperature or pressure can be explained by knowing at any time the positions and velocities of a set of atoms in their continuous, disordered motion. But neglecting the difficulty of having to consider the mechanics of a system with an incalculable number of independent particles, well before the counterintuitive theories of the 20th century, Maxwell, Gibbs and Boltzmann formulated the laws of a branch of physics, called statistical mechanics (chapter 2), which made it possible to link the postulates of thermodynamics, assumed on the basis of experimental evidence, to the idea that any portion of matter is nothing more than a practically infinite set of atoms.

The results achieved by following the principles of thermodynamics may seem astonishing if one considers that the laws on which it is based completely ignore the origin of the phenomena for which an explanation is being sought. Yet, through the concepts of functions and variables of state, spontaneity or reversibility of a transformation, thermodynamics results are in general highly accurate in predicting in which state a system should be found, once set some external and internal constraints.

¹A more in-depth discussion of some aspects covered in this chapter can be found in the textbooks by Fermi [67], Gordon [68], and Atkins and Paula [69].

4.1.1 State quantities

Thermodynamics aims at describing the transformations that a system undergoes when subject to some interaction with its surroundings. Even without knowing the details of such transformations, they are described in terms of changes of *state variables*, quantities that can be manipulated almost directly by altering the condition (or the *state*) of the system. Temperature and pressure are two of these variables, but in other situations, also mechanical or electro-magnetic potentials may be relevant.

The important concept is that state variables are not independent of one another, but there is always a relation — an *equation of state* — constraining some of them once the others are set. Moreover, some quantities play a more prominent role since they are *only* a function of state variables. No matter which kind of transformation the systems undergoes: these quantities depend exclusively on the values of state variables before and after the transformation. Alongside these simple definitions, there are two laws which, although they can be expressed in different forms, assert the following:

1. If a system that cannot exchange particles with its surroundings in internal equilibrium undergoes a transformation, the difference between the heat absorbed by the system, Q , and the work W it performs¹ is equal to the change in a state function called *internal energy*, E :

$$dE = dQ - dW. \quad (4.1)$$

This statement, known as the *First Law of thermodynamics*, is nothing but the fundamental principle of the conservation of energy.

2. There exists another state function, the *entropy*, S , whose variation during a transformation taking place in an isolated system² is always positive: $dS \geq 0$. Although entropy is only partially defined by classical thermodynamics, this statement links the concept of reversible transformation to its spontaneity: an irreversible process is a spontaneous change. Such a fundamental concept deserved the name of *Second Law of thermodynamics*

As seen from the equation expressing the Second Law, to understand whether a process occurs as a natural change, one must calculate the total change in entropy, that of the system and the surroundings. However, it is possible to introduce another fundamental quantity that gives the same information but only in terms of state variables of the system. This quantity is called *Gibbs free energy*, G , and it is defined as

$$G \equiv H + pV,$$

where H is enthalpy. The Second Law recast in term of Gibbs free energy reads

$$dG_{p,T} \leq 0, \quad (4.2)$$

if the transformation is carried out at constant temperature and pressure, which are the most common conditions in real laboratories and computer experiments. Thus, combining the


First and the Second Laws, it is clear that a process at constant temperature and pressure is spontaneous if associated to a decrease in Gibbs free energy, and also that the equilibrium state is the one where G reaches its minimum.

4.1.2 Phases and components

The seemingly elementary conclusion summarized by eq. (4.2) happens to be the fundamental rule to explain the behavior of more complex systems, even if it was derived assuming the system to be *closed* and of *fixed composition*. To see why this is the case, it is useful to recall the definitions of *phase* and *component*.

In his seminal work that laid the foundations of chemical thermodynamics, Gibbs defined a phase as a state of matter uniform in both chemical composition and physical state³. For example, a solid in a given crystal structure, a mixture of gases or two perfectly miscible liquids form a single phase. A subtle example of a system comprising two phases is a *dispersion*. Although the term brings to mind the liquid state of matter, the concept of dispersion is also crucial in metallurgy, where the controlled precipitation of a particular crystalline phase in a matrix having another structure is of fundamental importance for some advanced applications.

A component is defined as a chemically independent species of a system. Thus, the number of components is the minimum number of chemical constituents needed to define the composition of *all* the phases present in the system. As an example, the reaction of decomposition of calcium carbonate



where the number of phases P is 3, while the number of components C is 2: although there are 3 chemical species, only the amounts of two of them are needed because the last is always fixed by the stoichiometry of the reaction.

Even if a transformation as simple as that of the example above may take place in a laboratory vessel not allowing for any exchange of mass, the system may be subject to variations in the relative compositions of its constituents. It may therefore be useful to focus for a moment on a particular phase in the system and consider it as an open system that can exchange matter with the environment, that is, the other phases. To each transfer of matter can be associated a variation of the internal energy; therefore, for a system with C components we have

$$dE = \sum_{i=1}^C \frac{\partial E}{\partial n_i} dn_i = \sum_{i=1}^C \mu_i dn_i,$$

where μ_i is called the *chemical potential* of the component i , and dn_i is the infinitesimal change in the amount – be it in kilograms, number of moles or atoms – of that component.

A useful combination of the First and Second Principle can be written as $dE = TdS - pdV^4$, to which must be added the contributions to internal energy due to changes in the components that occur in an open system such as the one considered now. It then becomes

$$dE = TdS - pdV + \sum_i \mu_i dn_i, \quad (4.3)$$

in which the chemical potentials are at constant S , V and the amounts of all the other components. Therefore, the quantity μ_i shows how the internal energy changes when the composition changes: it is the *partial molar internal energy* at constant entropy, volume and amounts of other components.

By writing the total differential of Gibbs free energy and comparing it with eq. (4.3), one arrives at

$$dG = Vdp - SdT + \sum \mu_i dn_i. \quad (4.4)$$

Moreover, since Gibbs free energy is a state function, this relationship is valid for any type of transformation, reversible or irreversible. This is why eq. (4.4) is usually referred to as the *fundamental equation of chemical thermodynamics*.

One of the simple consequences of thermodynamic equilibrium is that the temperature is constant at every point of the system. This principle extends to pressure as well. Therefore, several phases of a multicomponent system will necessarily have to have the same temperature and pressure. The fundamental equation we have just arrived at introduces a third, important principle on which depends an indispensable tool in the study of multi-component systems: the phase diagram. This principle states that, at equilibrium, the chemical potential of each component of a system will be the same in all phases.

To see why this is the case, consider two phases α and β within a system at equilibrium. If dn_i moles of a component transform from α phase into β , holding the temperature, the pressure and the amounts of any other component constant, the total change in Gibbs free energy is

$$dG = \mu_i^\alpha dn_i - \mu_i^\beta dn_i = (\mu_i^\alpha - \mu_i^\beta) dn_i.$$

Since the composition and mass of the system as a whole have not changed, at equilibrium $dG = 0$, which implies that

$$\mu_i^\alpha = \mu_i^\beta. \quad (4.5)$$

This argument directly applies to any other couple of phases in the system, thus proving the above statement that at equilibrium each component's chemical potential is the same throughout all the phases present.

One of the most elegantly simple arguments proved by Gibbs is the well-known *phase rule*. This relation ties the variance of a system, F – that is, the number of intensive variables that can be changed independently from one another – to the number of components, C , and phases, P , at equilibrium:

$$F = C - P + 2. \quad (4.6)$$

A justification of the phase rule goes as follows. For a one-component system ($C = 1$) with two phases in equilibrium, since the chemical potentials are functions of the temperature and pressure, eq. (4.5) reads $\mu^\alpha(p, T) = \mu^\beta(p, T)$. That is, only one variable among temperature and pressure is independent. If the system consists of three phases in equilibrium,

$$\mu^\alpha(p, T) = \mu^\beta(p, T) = \mu^\gamma(p, T).$$

Since the latter is a system of two equations with two unknowns, only a single value of T and p can be a solution to that system. Should there be a system with three equations and two unknowns, no solution would exist, and this is consistent with the fact that a single component system cannot have four phases in mutual equilibrium.

Consider now the general case. Pressure and temperature count as 2. Each phase's composition is fully determined once $C - 1$ variables for the amounts of the components are given. Thus, with P phases, the total number of composition variables is $P(C - 1)$. At equilibrium, the chemical potential of a component is equal in every phase, and this accounts for $P - 1$ equations of the kind of eq. (4.5). With C components, there are $C(P - 1)$ constraints fixed by such equations. Eventually, the number of independent intensive variables left is $F = 2 + P(C - 1) - C(P - 1)$, which is exactly Gibbs' phase rule.

The phase rule is of primary importance in the study of phase diagrams, of which a more in-depth discussion will be the subject of the next section. However, the phase rule already shows some interesting scenarios in the case of a two-component system. For example, unlike a single-component system, it is possible that two phases coexist in a range of temperatures and compositions, or that at a fixed pressure three phases coexist in equilibrium⁵.

4.2 Binary systems

As briefly discussed in the previous section, binary systems have an additional state variable, the total composition, on which the Gibbs free energy function depends. A phase diagram of a two-components system will therefore depend on three variables, that is, it is represented by a three-dimensional surface. Since many interesting phase transformations of the solid matter occur at a fixed pressure, it is customary to describe how the Gibbs free energy changes upon a change in temperature and/or compositions.

The simplest binary system is the one in which the two constituents are present in the solid phase either as pure elements or as a solution of the two in each other. Moreover, if each solid shows only one crystalline structure, a solid solution will form at any composition.

If x_A and x_B represent the mole fractions of the two components, the total Gibbs free energy is simply $G = x_A G_A + x_B G_B$. The formation of a solid solution brings about a contribution to the total Gibbs free energy due to mixing. Since G can be always written as an enthalpy minus an entropy term ($G = H - TS$), before the mixing

$$G_1 = x_A H_A + x_B H_B - T(x_A S_A + x_B S_B),$$

while after the formation of the solid solution

$$G_2 = x_A \bar{H}_A + x_B \bar{H}_B - T(x_A \bar{S}_A + x_B \bar{S}_B) - T \Delta_{\text{conf}} S,$$

where Y indicates a molar property of a pure component and \bar{Y} the same property when it is in solution. The term labelled as "conf" expresses the configurational contribution to entropy and distinguishes between entropy due to atomic thermal vibrations. $\Delta_{\text{conf}} S$ arises from a purely statistical argument: when there are two, distinguishable species, the number of ways of

arranging them is far greater than one – the trivial case of the pure substances before mixing.

4.2.1 Ideal solutions

The concept of ideal solution is closely related to that of ideal gas: it is a rather extreme approximation which never holds for real systems, though ideality is approached in certain conditions. An ideal solution results by definition from the assumption that there is no change between both the enthalpy and entropy of a pure component and its solution with the other one. Hence

$$\bar{H}_A - H_A = \bar{H}_B - H_B = 0,$$

and

$$\bar{S}_A - S_A = \bar{S}_B - S_B = 0.$$

This means that the species of the two components behave as two identical entities both chemically and physically, yet they can be treated as distinguishable particles. The first assumption requires that $\Delta_{\text{mix}}H = 0$, that is, there is no change in enthalpy upon mixing. However, even if the second assumption requires a zero net change in entropy between the pure component and its solution, $\Delta_{\text{conf}}S \neq 0$. Since an ideal solution is analogous to a mixture of perfect gases – although the interactions between species cannot be simply neglected as in a mixture of ideal gases – the entropy of mixing is exactly equal to the configurational term. To obtain the entropy of mixing of an ideal mixture, one may start from its definition in terms of chemical potential: an ideal mixture is one for which Raoult's law – that is $p_A = x_A p_A^*$, where the asterisk denotes the pure component and p_A the partial pressure – is valid throughout the composition range from pure A to pure B. The chemical potential is then

$$\mu_A = \mu_A^* + RT \ln x_A.$$

Consider now the Gibbs free energy of two ideal components before mixing. This is

$$G_i = n_A \mu_A^* + n_B \mu_B^*,$$

while after they have mixed

$$G_f = n_A(\mu_A^* + RT \ln x_A) + n_B(\mu_B^* + RT \ln x_B).$$

It follows that the Gibbs free energy of mixing is

$$\Delta_{\text{mix}}G = nRT(x_A \ln x_A + x_B \ln x_B),$$

and that the ideal entropy of mixing, since $\partial G/\partial T = -S$,

$$\Delta_{\text{mix}}S = -nR(x_A \ln x_A + x_B \ln x_B) \equiv \Delta_{\text{conf}}S,$$

where n here stands for the total number of moles, and R is the universal gas constant.

It is easy to see why an ideal solution is more stable than any mixture of solutions with different compositions. If an ideal solid solution with overall composition x^1 exists as a mixture of two solutions with compositions x^α and x^β , recalling eq. (4.4)

$$dG^\alpha = V^\alpha dp - S^\alpha dT + \mu_A^\alpha dx_A^\alpha + \mu_B^\alpha dx_B^\alpha.$$

Since $dx_A^\alpha = -dx_B^\alpha$, at constant temperature and pressure

$$\left(\frac{dG}{dx_B}\right)^\alpha = \mu_B^\alpha - \mu_A^\alpha,$$

and similarly for the second phase β

$$\left(\frac{dG}{dx_B}\right)^\beta = \mu_B^\beta - \mu_A^\beta.$$

At equilibrium the chemical potentials of each component must be equal in every phase, hence

$$\left(\frac{dG}{dx_B}\right)^\alpha = \left(\frac{dG}{dx_B}\right)^\beta, \quad (4.7)$$

that is, the slopes of the curve of the Gibbs free energy as a function of x_B must be the same at the two compositions x^α and x^β . Since the free energy of an ideal solution will resemble the curve labeled G_2 fig. 4.1, we are led to conclude that eq. (4.7), which is known as *common tangent condition*, is true only if $x^\alpha = x^\beta = x^1$. This relation implies that a single ideal solution, no matter of which total composition, will be the equilibrium state of any mixture of two solutions with different compositions.




Figure 4.1 – Gibbs free energies of mixture (G_1), of the ideal solution (G_2), and the configurational contribution to entropy due to mixing. The Gibbs free energy of the pure component A can be taken equal to zero, while $G_B > G_A$.

If it were not already evident, the concept of the ideal solution has to be loosened if one wants to be able to describe real binary systems, where multiple phases with different compositions

may well coexist.

4.2.2 Equilibrium phase diagrams of ideal systems

Before describing a more sophisticated model than that of ideal solutions — one that is at least able to approximate the complex behavior of real systems — it is useful to review how one can obtain an equilibrium phase diagram for an ideal solution. The procedure to construct such a phase diagram is entirely based on the determination of those compositions which, on the free energy curve as a function of the total composition, correspond to points with the same tangent line — the already mentioned method of the common tangent.

Consider a system in which both the liquid and the solid state may exist as ideal solutions. For each phase, the molar Gibbs free energy is given by

$$G_m = x_A G_{A,m} + x_B G_{B,m} + \Delta_{\text{mix}} G_m,$$

and it has been shown that for an ideal solution $\Delta_{\text{mix}} G_m = -RT \Delta_{\text{mix}} S_m$. Plotting such free energy curves as a function of either x_A or x_B results in a straight line at zero temperature, while at $T > 0$ they will be more and more sharply concave due to the increasing entropic term. Moreover, since at fixed composition and pressure

$$\left(\frac{\partial G}{\partial T} \right)_p = -S,$$

the rate of increase of G_m will be greater for the liquid phase than for the solid, and eventually the two curves will intersect. The intersection may happen as in fig. 4.2: in this case, the free energy of the pure solid A and pure liquid A are equal, that is, the temperature is the melting temperature of A as a pure substance.




Figure 4.2 – Example of free energy curves in an ideal system. At this temperature, the solid and the liquid phase have the same free energy only for $x_B = 0$, that is, pure A. The temperature therefore indicates the equilibrium melting temperature for the pure component A.

If the intersection happens at some intermediate composition, then it can be shown that

within a certain range of compositions the equilibrium is a two-phase system constituted by a solid and a liquid solutions. This is consistent with the phase rule, for the variance F of such a region of the phase diagram is 1: one intensive variable can be change arbitrarily without altering the phases at equilibrium.




Figure 4.3 – Free energy curves for an ideal system as a function of the fraction of B. M is the lowest free energy for the alloy with composition x .

To see why the two-phase system is more stable than either a solid or a liquid solution alone, consider fig. 4.3: if the total composition is x , a solid solution at this composition would have the free energy corresponding to point a , while point b would give the free energy of a liquid solution. However, the lowest free energy attainable would be that of point M , which corresponds to a mixture of two solutions, a solid one with composition given at point d and a liquid one with composition given at point e .

Suppose that the alloy exists as a mixture of a solid solution and a liquid solutions with composition x^S and x^L , respectively. By definition, the chemical potential is the molar free energy of a component in solution,

$$G_m = \mu_A x_A + \mu_B x_B.$$

Given that $dx_A = -dx_B$,

$$\mu_A = G_m - x_B \frac{\partial G_m}{\partial x_B}, \quad (4.8)$$

and an analogous relation for the chemical potential of B:

$$\mu_B = G_m + (1 - x_B) \frac{\partial G_m}{\partial x_B}. \quad (4.9)$$

These two relations are valid for both the solid and the liquid phase.

Comparing the relation for μ_A expressed for the liquid phase with the diagram,

$$\mu_A^L = \overline{HF} + x_B^L \frac{\overline{CH}}{x_B^L} = \overline{HF} + \overline{CH} = \overline{CF}.$$

An analogous reasoning for the solid phase gives $\mu_A^S = \overline{JF}$. Since at equilibrium the two chemical potential must be the same, $\overline{CF} = \overline{JF}$, which means that *C* and *J* are the same point. This condition, together with that of eq. (4.7), tells us that, at equilibrium, the free energy curves must have the same slope at the points given by the two compositions x^S and x^L , and that the two tangent lines must intersect the vertical axis at the same point: these conditions are fulfilled only if the two tangent lines are the same straight line.

This equilibrium condition is valid for every value of the overall composition x lying in the range between the points of common tangency — indicated by *d* and *e* in fig. 4.3: at this temperature and for any composition in the range, the system will always tend to form a mixture of two solutions with composition given by points *d* and *e*, regardless of the total composition. If the total composition lies on the left of point *d*, then the system returns to a single-phase region, in particular a single solid solution; the same goes for total compositions on the right of point *e*, but in this case a single liquid solution is formed.

4.2.3 The lever rule

An equilibrium phase diagram tells us not only within which range of compositions, at a given temperature, two or more phases may coexist. One can also extract quantitatively the amount of each phase present. This is done by means of a relation known as *lever rule*. Consider a system containing n moles and with two phases, α and β , at equilibrium; if x_A is the total composition, the total amount of *A* is given by

$$nx_A = n^\alpha x_A^\alpha + n^\beta x_A^\beta.$$

Since the total number of moles $n = n^\alpha + n^\beta$ is fixed,

$$nx_A = n^\alpha x_A + n^\beta x_A$$

from which it follows that

$$\frac{n^\alpha}{n^\beta} = \frac{x_A - x_A^\beta}{x_A^\alpha - x_A}. \quad (4.10)$$

The practical importance of the lever rule is that both the numerator and the denominator of eq. (4.10) can be readily obtained by inspection of a phase diagram.

4.2.4 Regular solutions

Real solutions are obviously not ideal. A first, intuitive reason for this is that the interactions between particles of species A and B are all different and they can vary substantially across different phases. In contrast to ideal liquid or solid solutions where no enthalpy or volume changes occur, mixing of real solutions may be accompanied by an additional contribution to entropy, besides that arising from configurational freedom. For example, if the entropy change is negative – because particles organize in an orderly mixture – and the enthalpy change is large and positive, then the free energy of mixing may be positive: the two liquids or solid may then be only *partially miscible*, that is, they will form a single phase only over a certain range of compositions.

A more complex – albeit approximated – model which is suitable to explain departures from ideality is that of *regular solutions*. This model rests on the two following assumptions:

1. The internal energy depends only on two-body interactions between nearest neighbor particles. Hence, these interactions – whose energies are labelled ϵ_{AA} , ϵ_{BB} and ϵ_{AB} – are independent of the composition of the mixture.
2. For a given crystal structure of the solid phase, the probability of occupation of a lattice site by a particle A or B, as well as the probability that two adjacent sites have occupation AA, BB or AB, depend only on the overall composition and not on the local environment of a particle. In other words, this corresponds to a completely random distribution of the two species on the available sites, which is equivalent to having an ideal solution: therefore, the entropy of mixing is identical to that of ideal solutions.

With N_A particles of species A and N_B of species B, x_A and x_B being the usual atomic fractions, the energies of a system made up of pure A is

$$E_A = \frac{1}{2}N_A Z \epsilon_{AA},$$

where Z stands for the coordination number of a site. An equivalent expression is valid for the species B.

In the mixture, a particle on a site (either A or B) has on average Zx_A first neighbors of type A and Zx_B first neighbors of type B. Therefore, the energy of the mixture will be equal to

$$\begin{aligned} E_{AB} &= \frac{1}{2}N_A (x_A Z \epsilon_{AA} + x_B Z \epsilon_{AB}) + \frac{1}{2}N_B (x_B Z \epsilon_{BB} + x_A Z \epsilon_{AB}) \\ &= \frac{1}{2}NZ (x_A^2 \epsilon_{AA} + 2x_A x_B \epsilon_{AB} + x_B^2 \epsilon_{BB}). \end{aligned}$$

The energy of mixing is obtained by subtracting from the above expression the energies of the pure substances, therefore

$$\begin{aligned} \Delta_{\text{mix}}E &= E_{AB} - E_A - E_B = \frac{1}{2}NZ(2\epsilon_{AB} - \epsilon_{AA} - \epsilon_{BB})x_A x_B \\ &= NZ\epsilon x_A x_B, \end{aligned}$$

where

$$\varepsilon = \varepsilon_{AB} - \frac{\varepsilon_{AA} + \varepsilon_{BB}}{2}.$$

For condensed phases, the term pV that has to be added to E to obtain enthalpy is usually negligible at ordinary pressures, so assuming that $\Delta_{\text{mix}}H = \Delta_{\text{mix}}E$ is a legitimate approximation. The molar enthalpy of mixing is then

$$\Delta_{\text{mix}}H_m = \mathcal{N}Z\varepsilon x_A x_B = Bx_A x_B, \quad (4.11)$$

where the Avogadro's constant is denoted with \mathcal{N} to avoid confusion.

Since the entropy of mixing is the same as an ideal solution, the Gibbs free energy of mixing is given by

$$\Delta_{\text{mix}}G_m = \Delta_{\text{mix}}H_m - T\Delta_{\text{mix}}S_m = Bx_A x_B + RT(x_A \ln x_A + x_B \ln x_B), \quad (4.12)$$

and the total Gibbs free energy of the regular mixture is then

$$\begin{aligned} G_m &= x_A G_{A,m} + x_B G_{B,m} + \Delta_{\text{mix}}G_m = \\ &= G_{B,m} + (G_{A,m} - G_{B,m})x_A + Bx_A(1 - x_A) + RT [x_A \ln x_A + (1 - x_A) \ln (1 - x_A)]. \end{aligned}$$

By applying eqs. (4.8) and (4.9), it is possible to obtain the chemical potentials of the two components of the regular mixture

$$\begin{aligned} \mu_A &= G_{A,m} + B(1 - x_A)^2 + RT \ln x_A \\ \mu_B &= G_{B,m} + B(1 - x_B)^2 + RT \ln x_B \end{aligned} \quad (4.13)$$

As explained previously, there exists no combination of two solid (liquid) ideal solutions that is more stable than a single solid (liquid) solution with some intermediate composition. This was a direct consequence of neglecting the enthalpy of mixing. From eq. (4.11), the enthalpy of mixing of a regular solution can be either negative or positive, depending on the sign of $B = \mathcal{N}Z\varepsilon$. In the first case ($B < 0$), the interaction between particles of different species is more favored than those between identical particles; hence, the enthalpy term lowers $\Delta_{\text{mix}}G_m$ – the entropic term is always negative – further stabilizing the solution. In the second case $B > 0$, which means that AB interactions are energetically less favored than AA or BB ones: this can give rise to the transformation from a homogenous solid phase into two phases with different compositions, if the temperature is lower than a certain critical value, T_c . When $B > 0$ and $T > T_c$, the Gibbs free energy curve always shows upward concavity like the case of $B < 0$. When, instead, $B > 0$ but $T < T_c$, the repulsive energy term generates an interval of intermediate compositions in which the curve for $\Delta_{\text{mix}}G_m$ has downwards concavity, delimited by two points of inflection, and a wider interval delimited by the two points of minimum.

In these conditions, the curve $G_m(x_A)$ looks like fig. 4.4. There exists only one common tangent line, whose abscissas are indicated by $x_{A,\beta}$ and $x_{A,\alpha}$. For all compositions between $x_{A,\beta}$ and $x_{A,\alpha}$, the value of G_m of the mixture is greater than that of the heterogeneous aggregate of two phases with compositions $x_{A,\beta}$ and $x_{A,\alpha}$, represented by the ordinate of the corresponding




Figure 4.4 – Schematic plot of the molar Gibbs free energy of a regular solution as a function of the fractional amount of component A. Without loss of generality, we set $G_{A,m} = 0$ and $G_{B,m} = 10$. The dashed line is the common tangent line. The two black dots define the compositional range in which the Gibbs free energy of the solution is greater than that of a two-phase aggregate with compositions $x_{A,\alpha}$ and $x_{A,\beta}$. Within this range of compositions, the binary mixture separates into two mixtures: β , richer in B, and α , richer in A.

tangent segment. In this composition range, the mixture is separated to form a biphasic aggregate of mixtures β , richer in B, and α , richer in A. Recalling that the intercepts of the tangent line with the vertical axis at $x_A = 0$ and $x_A = 1$ give the chemical potentials μ_B and μ_A , respectively, it follows that at the points of tangency the relations of biphasic thermodynamic equilibrium are valid

$$\mu_A^\alpha = \mu_A^\beta \quad (4.14)$$

$$\mu_B^\alpha = \mu_B^\beta. \quad (4.15)$$

By imposing the condition of existence of an inflection point, namely

$$\partial^2 G_m / \partial x_A^2 = 0, \quad (4.16)$$

one finds the composition range within which the two-phase decomposition appears, provided that $T < T_c$. The solutions to eq. (4.16) are given by

$$x_A = \frac{1 \pm \sqrt{1 - 2\frac{RT}{B}}}{2}, \quad (4.17)$$

that correspond to real solutions in the interval $0 < x_A < 1$ only if $B > 2RT$, thus setting the value of the critical temperature:

$$T_c = \frac{B}{2R}.$$

The two-phase states of the system are represented by all the points lying on the segment

of the tangent line bounded by the pair of black dots of fig. 4.5 in which the G_m is plotted for several values of $T < T_c$. The amounts of the two phases α and β are given by the lever rule.




Figure 4.5 – Gibbs free energy for a regular solution as a function of the fractional amount of component A plotted for several temperatures. The lowest curve, labeled with T_c , is the graph of G_m at precisely the critical temperature. The upper curve represent increasing temperatures. For each curve, the two pairs of black and red dots represent the points of common tangency and the inflections points, respectively. At $T = T_c$, the compositional range for which phase separation occurs completely disappear, and this is seen by the overlapping dots.

The *phase separation curve* is the one connecting all the points of common tangency at different temperatures (see again fig. 4.5); the equation of this curve, composed by two branches $T(x_A, \alpha)$ and $T(x_A, \beta)$, is obtained by solving equations eqs. (4.14) and (4.15) simultaneously after substituting eqs. (4.8) and (4.9) as the expression for the chemical potentials of the regular solutions. The solution of such a system, however, can only be obtained numerically.

The interval of compositions delimited by the red dots of fig. 4.5 – which correspond to the inflection points of G_m – gives the range of *spinodal decomposition*, and the spinodal curve is obtained by connecting all those points at different temperatures; its equation can be derived from eq. (4.17) by writing T as a function of x_A . This curve, which lies within the phase separation curve, is of particular interest from a kinetic point of view. For its internal compositions, the curve of G_m has downward concavity: the solid solution is therefore not only unstable, but cannot be observed under metastable conditions either. Instead, for compositions between the spinodal curve and the phase separation curve, G_m presents upward concavity, and therefore the solid solution, though less stable than the mixture of the two phases, can be metastable.

As a concluding remark, it is worth noting that a single curve is able to represent the Gibbs free energy of the system in both single-phase and two-phase conditions. This is possible because the crystalline structure of the two solid solutions α and β is the same. If we wanted to compare two solid solutions with different structures, it would be necessary to represent them with two distinct G_m^α and G_m^β functions.

4.2.5 Invariant reactions

So far, the discussion on binary systems only considered the coexistence of solid solution phases with similar crystal structure. In practice, this circumstance is verified only if the two components are both chemically and physically very much alike. In turn, as several examples of phase diagrams showed, very similar components will tend to form only *terminal phases*, that is, a phase that can exist as a pure component. When the components start to show an appreciable difference in chemical properties, it is possible that, at fixed pressure, the free energy curves look like those in fig. 4.6: here it is assumed that for a given temperature range, three solid phases with different crystal structure coexist. The phase β is called *intermediate phase*.




Figure 4.6 – Schematic plot of free energy curves as a function of the composition for two terminal phases (α and γ) and a β phase that indicates a large dissimilarity between the components.

In contrast to the case of isomorphous systems, that is, systems where no structural difference between the solid phases is present, free energy curves of the kind depicted in fig. 4.6 will never join smoothly because there must be a point, corresponding to a given temperature and composition, where a phase change accompanied by a crystal structure change occurs.

An intermediate phase such as the β phase of fig. 4.6 has several interesting characteristics: since it involves rather dissimilar atoms, interatomic forces are typically of the kind of ionic or covalent bonds, thus leading to highly ordered arrangements of atoms. This implies that both the enthalpy and entropy of mixing are large and negative, which results in a solid phase with a very low free energy, i.e. highly stable. Moreover, ordered crystalline structures usually have atomic fractions proportional to small integer numbers, which means that intermediate phases will tend to have strictly fixed stoichiometric ratios of the species, such as A_2B , AB_2 or A_2B_3 , which are usually referred to as *compound-like* phases. Since the enthalpy depends strongly on the composition, any deviation from the equilibrium stoichiometry corresponds to a steep rise of the free energy.

An intermediate phase may also form when the atomic interactions between similar atoms

are stronger than those between dissimilar ones. This case is even more likely when the crystal structure of the pure components are very different. Such intermediate phases may have lower free energies than both the terminal phases and a mixture of those: this case is schematically shown in fig. 4.7. Phases of this kind tend to be random solid solutions, therefore their negative entropy of mixing stabilizes the mixture across a broad range of compositions; as a consequence, the formation of solid-solution-like intermediate phases limits the solubility of the components, although the effect is much more evident with a compound-like phase.




Figure 4.7 – Schematic plot of free energy curves as a function of the composition in which the phase β represent an intermediate phase of a solid-solution type.

In this section, the phase diagrams discussed differ from those of the previous sections because they contain one or more *invariant reactions*, that is, transformations during which the system goes through regions of the phase diagram where the number of coexisting phases exhausts the number of degrees of freedom, as defined by Gibbs' phase rule. This class of transformations is particularly important as almost any real phase diagram of some interest presents at least one of these phase transformations. It has been already noted that, for a two-component system at constant pressure, the maximum number of phases that may coexist at equilibrium is 3; in fact

$$F' = C - P + 1 \Rightarrow 0 = 2 - P + 1 \Rightarrow P = 3, \quad (4.18)$$

where F' emphasizes that the pressure has been already fixed.

There are not many kinds of invariant reactions that can occur in a two-component systems, and those transformations that do take place fall in two general categories:

1. The invariant reaction proceeds through the cooling of a single liquid phase which separates to form two solid phases. This is called *eutectic reaction* and it is by far the most common.
2. The reaction consists on the transformation upon cooling between a solid and a liquid phase to form a second solid phase. The invariant point the system goes through is called *peritectic*.

Given the importance of the first kind of invariant reactions, the remaining of this section will detail the thermodynamics how the solidification process occurs in binary systems that exhibit eutectic transformations.

4.2.6 Eutectic systems

Eutectic systems are binary systems that contain generally two solid solutions with terminal compositions and only one liquid phase. The regular solution model (section 4.2.4) is already able to account for eutectic invariant reactions. Recalling that a regular solution is that in which the $\Delta_{\text{ex}}G \neq 0$ but $\Delta_{\text{ex}}S = 0$ – i.e., entropy is that of an ideal solution – eutectic invariant reactions happen when ΔG_{ex} is positive, with the contribution from the solid much larger than that of the liquid. All phase diagrams that are characterized by an eutectic transformation will show three one-phase regions, three two-phase regions and one three-phase region. Since all these phase diagrams represent the equilibrium of the system at constant pressure, the three-phase region is one-dimensional: it is the tie line – the common tangent line – joining the compositions of the three phases in mutual equilibrium.




Figure 4.8 – A schematic phase diagram of a binary system which exhibits a liquidus minimum and a solid solubility gap.

Qualitatively, a eutectic phase diagram occurs when a two-solid region like the one depicted in fig. 4.8 intersects the upper two-phase region which displays a point where the liquid phase is in equilibrium with the solid at the same composition. The eutectic phase diagram of fig. 4.9b is the only possible choice for a system in which there is such intersection, and this can be elegantly seen by applying Gibb's phase rule. Consider fig. 4.9a: at temperature T phase α at composition a is in equilibrium with phase β at composition b' , according to the equilibria described by the solvus lines ca and $b'd$. The line fh tells that phase α is also in equilibrium with the liquid with composition e : that is, phase α , β and the liquid are all in mutual equilibrium at temperature T . An identical argument can be applied for at temperature T' , where the three phases will be in equilibrium with compositions a' , e' and b , respectively. However, eq. (4.18)

tells us that in a binary system at constant pressure a three-phase equilibrium must be an invariant point, which means that no intensive variable is allowed to change without disrupting the phases at equilibrium. The only possible conclusion is that $a = a'$, $b = b'$ and $e = e'$, and of course that $T = T' = T_e$. Figure 4.9b results from appropriately redrawing fig. 4.9a according to the just proven argument.




Figure 4.9 – (a) Hypothetical phase diagram when the two solid phase region of fig. 4.8 intersects the upper part showing the liquidus minimum. (b) The phase diagram on the left correctly redrawn according to Gibbs' phase rule, as explained in the text.

4.2.7 Solidification in eutectic systems

Eutectic invariant reactions are essentially a solidification process. Yet the interesting aspects of an eutectic system arise from the possibility to undergo different types of solidifications.

The first kind is known as *primary solidification* and occurs when, during the cooling of the liquid phase, the latter crosses only one two-phase region. Once the solid has finished precipitating, the system undergoes no further transformation and the system at that composition exhibits complete solid solubility.

The second possibility involves *solid-solid precipitation*. This happens when the systems with no more liquid content enters the second two-phase region, where two solid phases are

only partially miscible. Further cooling both the A and B solid solutions become richer in their respective phases, α and β respectively.

If we consider a portion of phase β embedded in an environment of α , there are two possible mechanisms by which expulsion of particles of type A may occur. The first is a diffusive process: particles of type A will first have to “migrate” to the interface between α and β . The second mechanism involves the nucleation of a small seed of α phase within the portion of β . Which of the two processes is most likely depends on many factors, including the temperature that controls the diffusive processes, the size of the β phase portion and, not least, the presence of defects (e.g., grain boundaries and cracks). The latter are fundamental for a thermodynamic reason: the nucleation of the seed of type A involves the formation of an interface to which is associated a term of positive free energy, often very large, therefore preexisting interfaces due to defects usually behave as preferential sites where nucleation is favored.

The third kind of solidification process is *eutectic solidification* and, as its name suggests, is what happens when a liquid of the right composition — i.e. the composition labeled by e in fig. 4.9b — is cooled until the eutectic temperature T_e . Once reached this temperature, the alloy will start its solidification by forming some phase α or β . However, in many real alloys that exhibit eutectic solidification this is not the case, and a solidified sample of the alloy will not show randomly alternating regions of primary α and β .

To understand the reason behind this experimental evidence, one should think once more to Gibbs’ phase rule: since an eutectic is an invariant point, if solidification occurs under equilibrium condition, both temperature and phase compositions cannot change as far as the three coexisting phases are present. Consider again fig. 4.9b and imagine extending the equilibrium lines between the liquid and the two phase α and β down to a temperature $T_1 = T_e - dT$. Then, referring to fig. 4.10 which zooms in on the region around the eutectic of fig. 4.9b, if some amount of α forms, the liquid would have to shift its composition towards the point m in the figure. On the other hand, if some β starts forming somewhere else, the liquid would have to match the composition given by point h . Since equilibrium solidification happens so slowly that diffusion of particles A and B may be considered instantaneous, these two opposite shifts will keep the composition of the liquid around e and the temperature at T_e ⁷. Real solidification only occurs when the degree of undercooling is finite and measurable, and we could rightly define the conditions in which the real solidification of an eutectic takes place as *quasi-equilibrium solidification*.

Quasi and non-equilibrium solidifications are key to enable real manufacturing and engineering processes obtaining so-called *metastable phases*, such as martensite in steel. In fact, quasi-equilibrium solidification is responsible of phenomena that simply cannot happen with an equilibrium phase transformation. An example is *constitutional undercooling*, that is, an increase in the degree of undercooling caused by a change in composition rather than temperature. At $T_1 < T_e$ (see again fig. 4.10) the alloy is represented by the point k , and when an amount of α forms, it will have a composition very close to that of point a in fig. 4.9b. The undercooling associated to the α phase, ΔT_α is given by the distance between the extension of liquidus line el and the horizontal line indicating the temperature T_1 . In a similar way, the undercooling of β (ΔT_β) is defined by the distance to the liquidus line relative to B. The solidification of particles




Figure 4.10 – Enlarged portion of fig. 4.9b around the point e that explains what happens during eutectic solidification.

of α phase enriches the liquid of B and, as a consequence, ΔT_α decreases while ΔT_β increases, giving rise to constitutional undercooling. This effect makes less and less favorable the growth of α and, at the same time, the growth of β more conducive. At some point the α phase stops growing completely because its driving force is zero. An analogous process takes place when β particles start nucleating, and eventually the alternation of these two processes is responsible for the typical growth patterns showing up in eutectic solidification. Interestingly, this kind of solidification happens also with an undercooling so small that it can be mistakenly confused with an equilibrium process. This would not be possible, however, because at equilibrium the diffusion of atoms of either A or B is fast enough to prevent any amount of A or B to accumulate in the liquid⁸. The result of equilibrium solidification would be therefore very much like two solid phases undergoing primary solidification, and we would end up with a solid consisting predominantly of a uniform mixture of primary phases.

4.3 Thermodynamics at interfaces

As seen when discussing eutectic solidification, departure from equilibrium, albeit complicating the picture, is crucial for many real-world applications. The case of AM is even more extreme, one may say, because it is characterized by several phenomena that are truly out-of-equilibrium.

A legitimate question is then: what would be the use of all the fundamentals of thermodynamics if applicable conditions are never really met? The answer is not as pessimistic as it might seem because, under some circumstances, it is still possible to apply results obtained under equilibrium conditions to those regions of the system that can be considered in *local equilibrium*, that is, the *interfaces*. These regions are the most interesting and, at the same time, elusive to study and describe, yet they can be still considered in equilibrium even if the bulk phases they separate are characterized by gradients in both temperature and composition.

Gradients in the various phases present are evidently only an example of what “non-

equilibrium” really involves. There are many circumstances in which even the interface deviates remarkably away from local equilibrium. These situations may concern both static and moving interfaces when a phase transition is underway, i.e., when the interface is not only subject to fluctuations.

This section outlines some of the most relevant concepts to study interface thermodynamics and their dynamical properties, with a particular outlook on atomistic modeling techniques that can provide useful insights on these properties.

4.3.1 How to define an interface?

Excess quantities and Gibbs dividing surface

A rigorous treatment of thermodynamics of solids can be very complex, especially for heterogeneous systems with many constituents. Should one be interested only in bulk properties, then a first approximation may neglect the effects of the interfaces separating the phases present.

However, a more thorough approach which includes the effect of interfaces is to assume that they behave like geometrical surfaces dividing two entirely homogeneous phases. Although this remains an approximation, Gibbs already proved⁹ how this idea of replacing the physical interface with a model system, in which the interface is only a geometric boundary, provides a way to study several interface properties (i.e. surface tension, adsorption coefficients), and that the same equilibrium conditions of equal temperatures and chemical potentials in each species must hold even when considering heterogeneous interfaces.

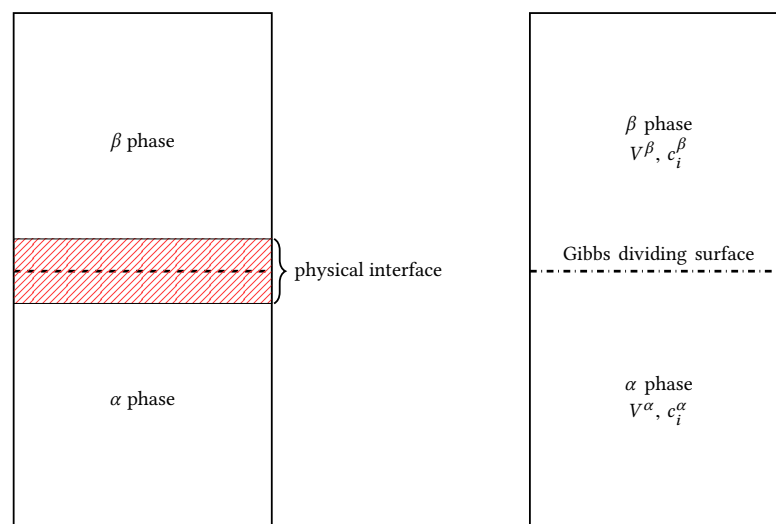


Figure 4.11 – Schematic representation of (a) a real, heterogeneous system comprising two phases in contact, and (b) a model system in which a Gibbs dividing surface is defined by setting zero surface excess of the molar volume.

If we assume that throughout the two phases α and β (see fig. 4.11) the concentration of the

component c_i is constant, then in the model system we have¹⁰

$$\begin{aligned}n_i^\alpha &= c_i^\alpha V^\alpha, \\n_i^\beta &= c_i^\beta V^\beta.\end{aligned}$$

If the real and the model system are to be equivalent chemically and physically, we have to introduce a term that expresses the number of moles contained in the dividing surface, n_i^S . Hence, the following requirement must always hold

$$n_i = n_i^\alpha + n_i^\beta + n_i^S,$$

and the number of moles associated with the interface is then

$$n_i^S = n_i - c_i^\alpha V^\alpha - c_i^\beta V^\beta,$$

where n_i^S is usually written as $\Gamma_i A$ to emphasize its dependence on the interfacial area. The prescription of Gibbs to define a dividing surface is to *explicitly* set $\Gamma = 0$. In other words, once an extensive reference quantity is chosen, a Gibbs dividing surface is assumed to have no associated surface excess of that quantity.

Evidently, there is a great degree of arbitrariness in choosing which extensive quantity is more suitable. With a similar argument one could define a ‘‘Gibbs free energy excess,’’ as the Gibbs free energies of the bulk phases α and β are not the same. Since at the interface there is mismatch between structural orderings of the two phases in contact, the interface itself has an intrinsic excess free energy. Thus, integrating the Gibbs free energy excess over the thickness of the interface naturally defines a quantity known as *solid–liquid interfacial free energy*

$$\gamma_{sl} \propto \int \Delta G^{\text{ex}}(z) dz, \quad (4.19)$$

where z indicates an arbitrary direction along the normal to the interface, and the proportionality constant has the dimension of a molar volume, which must take into account that the two phases in contact may have different values of V_m . The relevant point to remark is that all these excess quantities will depend on the exact location of the dividing surface, thus making its definition inevitably arbitrary.

Order parameters

Although there is no absolute choice of an extensive quantity to which a surface can be associated, it should still be interesting to look for criteria that could guide that choice. A firsthand possibility is to pick some thermodynamic variable of bulk phases, such as the volumes of each phase in the above example. However, these variables usually capture a global state of the system, while during certain phase transformations condensed phases often show consistent changes in rotational and translational degrees of freedom. The obvious example is a solid phase transition involving a change in the crystalline structure.

When studying solid–liquid phase transitions, changes in positional and/or rotational order usually entails that a solid phase is invariant only to a subset of all the possible symmetry operations, while a liquid phase is by definition isotropic, that is, it will be identical to itself regardless which symmetry operation is carried out. It therefore seems quite natural to include in the description of a phase transition a quantity that is capable of detecting such symmetry differences. Quantities of this kind are known as *order parameters*: Landau theory of phase transitions, for example, treat the Gibbs free energy as explicitly dependent on one or more order parameters, which become state variables as pressure, temperature and volume are.

There are several phase transformations in which the geometrical ordering does not change and those can be still faithfully described by, say, the difference in the density between the two phases. However, especially in the context of atomistic simulations, there are many situations where it is far more useful to describe the phase transition by means of order parameters that are able not only to distinguish unlike crystalline structures, but also the same crystalline lattice with different spatial orientation — that is, a generic rotation of the lattice represents a symmetry breaking.

In the field of molecular simulations, an order parameter is a more or less elaborate function of all the atomic coordinates, and it usually gives a quantitative information of the structure of each atom’s neighborhood. Many of these functions have been developed in order to face different problems; probably, the most known order parameters are the *Steinhardt–Nelson* functions of different orders: Q_3 , Q_4 and Q_6 [70]. Their choice was to exploit spherical harmonics to capture the symmetry of an atomic environment. The order parameter is then a discrete function that assigns a value to every particle in the system and is given by the following expression

$$q_{lm}^i = \frac{1}{N_n^i} \sum_{j=1}^{N_n^i} Y_{lm}(r_{ij}), \quad (4.20)$$

where N_n^i represents the number of neighbors of the particle i . To be able to distinguish different crystalline structure, Steinhardt proposed to average the complex quantities of eq. (4.20) by summing over different the index m

$$q_m^i = \sqrt{\frac{4\pi}{2l+1} \sum_{m=-l}^l |q_{lm}^i|^2}, \quad (4.21)$$

in which choosing different values of l enables to distinguish different lattices. There have been many improvements of this kind of order parameters — for example, by introducing a *locally averaged* order parameter [71] — but the common feature of these functions is that, since they are built as appropriate combinations (eq. (4.21)) of spherical harmonics, they are all rotationally invariant.

A particularly versatile choice of order parameter is that proposed by Angioletti-Uberti et al. [72] and has been used in some subsequent works [73, 74] featuring the study of interfaces between a liquid and a solid phase with a face–centered cubic symmetry. This order parameter

(henceforth called FCCUBIC, Φ) is defined as

$$\Phi_i = \frac{\sum_{j \neq i} C_r(r_{ij}) C_\alpha(\mathbf{r}_j - \mathbf{r}_i)}{\sum_{j \neq i} C_r(r_{ij})}. \quad (4.22)$$

In practice, Φ is a weighted average running over all pairs of atoms of an angular term (C_α) and a weight (C_r) that guarantees the continuity of Φ . The angular part depends on an Euler rotation matrix \mathcal{R} defined by the tuple of angles (φ, ψ, θ) and can be written as

$$C_\alpha(\mathbf{r}_{ij}; \varphi, \psi, \theta) = C_\alpha(\mathcal{R}(\varphi, \psi, \theta) \cdot \mathbf{r}_{ij}). \quad (4.23)$$

The angular part of Φ recognizes by construction the features of the FCC environment, and its peaks point in the directions where the first nearest-neighbors lie. Moreover, it is *not* rotationally invariant and thus it is ideally suited for identifying different orientations of the same crystal structure. To deal with arbitrary orientation it suffices specifying a suitable rotation matrix that transforms standard Cartesian axes into the desired set of orthonormal vectors.

4.3.2 Interface thermodynamics

A surface separating two phases can be seen as an extended, two-dimensional defect. Atoms in the vicinity of an interface must accommodate the structural order of both phases – if they are not both liquids or amorphous.

As eq. (4.19) shows, the tendency to minimize the surface in contact follows from the free energy cost associated with the formation of a surface with unit area. When a small crystalline seed that in an undercooled liquid forms and grows, if the interface free energy is isotropic, it is not surprising that the seed will increase its size as a growing sphere. In this case, γ_{sl} is usually identified with the surface tension, and the shape is given by the minimization of the surface area with the constraint of a fixed volume.

However, when crystalline solids are involved, the interface free energy varies substantially for different crystal orientations, that is, γ_{sl} is an anisotropic quantity. Therefore, returning to the example of the crystalline seed, one would soon observe perturbations in the growing shape that express the preferred orientations of the underlying crystal structure and its symmetries. The anisotropy of γ_{sl} is of fundamental importance when studying solidification since, together with the anisotropy of interfacial growth kinetics, it is at the origin of the so-called *dendritic growth mechanism* that strongly influences the microstructure of the final solid.

Anisotropy

While the shape of a solid seed for a material with isotropic γ would be a sphere, real materials show deviations of the seed's shape from a perfect sphere that are due to anisotropic behavior of γ , which is best written in terms of spherical harmonics, $Y_l^m(\theta, \varphi)$. A common expression

is [75]

$$\gamma_{\text{sl}}(\theta, \varphi) = \gamma_{\text{sl}}^0 \left[1 + \frac{1}{2} \sum_{l,m} \delta_l^m Y_l^m(\theta, \varphi) \right], \quad (4.24)$$

where $\theta \in [0, \pi]$ and $\varphi \in [0, 2\pi]$ are the polar and azimuthal coordinates that identify a point on the surface of a sphere with fixed radius, and δ_l^m represents the magnitude of the deviations from constant γ .

Physical intuition could often simplify generic expressions as the one above. Since not every symmetry operation of a spherical shape is compatible with a given crystal structure, a symmetry-adapted basis set that is a linear combination of spherical harmonics may yield a much more meaningful and usable expression. For the case of the face-centered cubic lattice, Fehlner and Vosko [76] proposed a basis that would satisfy the symmetry of this crystal structure, in particular its cubic symmetry operations. This basis is built on two terms that are function of θ, φ and the Cartesian components of the unit normal vector¹¹, $\mathbf{n} = (n_x, n_y, n_z)$. These functions are

$$\begin{aligned} Q &= n_x^4 + n_y^4 + n_z^4 = \sin^4 \theta (\cos^4 \varphi + \sin^4 \varphi) + \cos^4 \theta, \\ S &= n_x^2 n_y^2 n_z^2 = \sin^4 \theta \cos^2 \theta \sin^2 \varphi \cos^2 \varphi. \end{aligned} \quad (4.25)$$

These functions can be shown to be equivalent to linear combinations of some spherical harmonics. In particular, if one considers only the first three terms of the sum in eq. (4.24) – those with $l = 4, 6, 8$ – these terms are linear and quadratic functions of Q and S . Therefore, one may write a new expression for γ_{sl} as follows

$$\gamma_{\text{sl}}(\vec{n}) = \gamma_{\text{sl}}^0 \left[1 + \varepsilon_1 \left(Q - \frac{3}{5} \right) + \varepsilon_2 \left(3Q + 66S - \frac{17}{7} \right) + \varepsilon_3 (65Q^2 - 94Q - 208S + 33) \right], \quad (4.26)$$

where the $\varepsilon_1, \varepsilon_2$ and ε_3 are usually referred to as *anisotropy parameters*. The latter is a form that describes the anisotropy of γ_{sl} which will prove to be particularly appropriate for our atomistic simulations. It should be noted that eq. (4.26) represents an expansion on the chosen symmetry-adapted basis up to the third order, but could easily include higher order terms. However, in most of the cases, using a second-order expansion is already sufficient to fully describe the anisotropy accurately¹².

Curvature and stiffness

The anisotropy parameters in eq. (4.26) are just mathematical coefficients of a power expansion, yet the anisotropy itself is a quantity that plays a central role during early stages of solidification. From a geometric standpoint, one may describe any deviation of an initially spherical shape in term of the surface's curvature radii¹³. The information provided by the mean curvature, for example, directly enters the Gibbs–Thomson equation (see eq. (1.3)), therefore it is an essential parameter to describe interface evolution.

There is a second key physical quantity that is usually considered when trying to justify the preferred growth directions of a crystalline solid. This quantity is known as *stiffness* and it is strictly related to curvature. If one wanted to give the statement “preferred growth direction”

a precise meaning, it would be that the most favored directions are those along which the interface attains its maximum mean curvature; equivalently, one could say that those directions correspond to the minima of the interfacial stiffness. This statement is the only robust rule one can adopt to determine the truly preferential growth directions: if one considers the expression given by eq. (4.26), without knowing whether the anisotropy parameters will be positive or negative, it is impossible to predict whether the $\langle 100 \rangle$ directions are preferred over the $\langle 110 \rangle$. It has been actually shown [77] that many FCC metals display a positive ε_1 , while ε_2 is negative; therefore, besides particular cases with very strong anisotropies, one is not able to say *a priori* which direction tends to grow spontaneously because of the mixing of opposite weights in the expansion of γ_{sl} .

The stiffness may be defined as the free energy cost to bend away the interface from its local equilibrium. A smaller stiffness means that the surface can be more easily deformed along a particular direction, which then turn out to be one of the preferential growth directions. In mathematical terms, Herring [78] gave the first formal expression of the stiffness – usually indicated with σ_{ij} , where the two indexes stress the tensorial nature of the stiffness for a three-dimensional crystal – and related the stiffness to the chemical potential difference between the solid and the liquid phase. If one defines two angles, θ_i and θ_j , that measure the deviations of the interface normal vector \hat{n} with respect to two orthogonal directions in the tangent plane, the stiffness can be written as

$$\sigma_{ij} = \gamma \delta_{ij} + \frac{\partial^2 \gamma}{\partial \theta_i \partial \theta_j}, \quad (4.27)$$

where δ_{ij} is the Kronecker delta. Once a suitable coordinate system is fixed, it is straightforward to calculate the analytical expression of σ_{ij} for different crystallographic orientations starting from an expansion such as eq. (4.26). Symmetry arguments can simplify the derivation of these expressions by imposing constraints. For example, for a crystal with the face-centered cubic symmetry, the presence of at least a threefold rotational symmetry for interfaces such as (100) and (111) reduces the stiffness tensor to be diagonal ($\sigma_{ij} = 0$ if $i \neq j$) and requires that $\sigma_{ii} = \sigma_{jj}$. This means that the stiffness is simply a scalar, that is, an isotropic quantity.

The capillary fluctuation method

The very first attempt to measure γ_{sl} for a solid–liquid interface by means of atomistic simulations is due to Broughton and Gilmer [79]. They developed the so-called *interface cleaving method* (CM) to estimate the surface excess free energy by determining the work needed to cleave both a bulk liquid and bulk solid – the latter done in such a way to expose the desired interface – and then joining these two halves back again. Calculating the total reversible work per unit area along this pathway gives the value of γ_{sl} .

In a recent work, Davidchack and Laird [80] suggested several improvements to the way in which it is possible to build a cleaving potential, but it remains the difficulty that, to ensure a reversible transformation, such potentials are neither functions determinable *a priori*, nor particularly simple. Moreover, even having a method to generate these potentials that do not significantly perturb the system, it is necessary to generate them for each crystalline orientation

of the interface whose free energy is to be calculated.

An alternative route described for the first time by Hoyt, Asta, and Karma [81] allows to overcome many of the limitations of the cleaving method and has quickly become one of the reference methods for extracting information on free surface energy and anisotropy from atomistic simulations.

If the solid has cubic symmetry, a commonly used expression for the interfacial free energy is

$$\gamma_{sl} = \gamma_{sl}^0(1 + \varepsilon_4 \cos 4\theta).$$

By simply applying eq. (4.27), since there is only one angle, θ , that defines the tilt of the equilibrium interface's normal, one can see that the stiffness is

$$\sigma = \gamma_{sl} + \frac{\partial^2 \gamma_{sl}}{\partial \theta^2} = \gamma_{sl}^0(1 - 15\varepsilon_4 \cos 4\theta).$$

The anisotropy of γ_{sl} is controlled by ε_4 , but it is evident that the stiffness is 15 times more anisotropic than the interfacial free energy. This argument is not restricted to the cases where one could write such a simple expression of γ_{sl} , and it lies at the heart of the method introduced by Hoyt and coworkers known as the *capillary fluctuation method* (CFM). It exploits atomistic simulations to directly determine σ and, once the anisotropy parameters are known, obtain the value of γ_{sl} for any particular interface. Capillary wave theory asserts [13, 82] that fluctuations of an interface in thermodynamic equilibrium are directly related to the stiffness via

$$\langle |A(k)|^2 \rangle = \frac{k_B T_m}{S \sigma k^2}, \quad (4.28)$$

where $\langle |A(k)|^2 \rangle$ are the mean square Fourier coefficients of the expansion

$$h(x) = \sum_k A(k) \exp(ikx), \quad (4.29)$$

where $h(x)$ is the profile of the fluctuating interface. For eq. (4.28) to be applicable, the wavelength $\lambda = 2\pi/k$ associated to each Fourier mode with frequency k must be much larger than the typical atomic spacing, in order to sample only fluctuations of the atomically rough interface and not phonon vibrations.

The quantity on the left hand side of eq. (4.28) can be extracted from a standard MD simulation provided that one has an appropriate way of locating the dividing surface during the simulation trajectory. One possible strategy is that explained previously (section 5.1) and has been applied to a standard Lennard–Jones system to extract its interface anisotropy. An alternative approach that is based on the Gibbs dividing surface prescription will be showed in the next section. Once the information of the location of the interface is available, applying a discrete Fourier transform algorithm directly provides the Fourier amplitudes $|A(k, t)|^2$. The static fluctuation spectrum is computed by performing a standard ensemble (time) average. As it has been noted previously, a simulation's length should only be long enough to provide the temporal evolution of the system for a time greater than the typical auto–correlation time of the process under study. A subtle

detail of performing such averaging on an MD trajectory to extract $\langle |A(k)|^2 \rangle$ is that different Fourier modes typically have very different auto-correlation times, therefore a simulation whose goal is to obtain such fluctuation spectrum needs to last from about hundreds of picoseconds to several nanoseconds. This requirement may pose a practical problem that cannot always be avoided if, for example, a considerable portion of the simulation's computational cost is due to the particular choice of the interatomic potential used.

Remembering eq. (4.26), it is evident that to be able to fully determine the anisotropy parameters one needs at least three independent simulations from which to extract the stiffness of the corresponding interfaces. A viable alternative is to perform simulations with two-dimensional cells where both the dimensions of the interface are comparable. In this way, the interface can sustain fluctuations along two independent directions in the tangent plane¹⁴ and the stiffness tensor along these two directions can be determined from a single simulation. Although such simulations require boxes with many more atoms than those where the interface has a ribbon-like geometry — that is, where fluctuations are sampled along only one tangent direction — they are particularly useful for determining the stiffness tensor of anisotropic interfaces, that is, those where $\sigma_{ii} \neq \sigma_{jj}$. An expression equivalent to eq. (4.28) can be written for these two-dimensional interfaces

$$\langle |A(\mathbf{k})|^2 \rangle = \frac{k_B T}{S (\sigma_{11} k_x^2 + \sigma_{22} k_y^2 + 2\sigma_{12} k_x k_y)}, \quad (4.30)$$

where now \mathbf{k} is a two-dimensional vector in Fourier space that represents a single mode in the expansion

$$h(x, y) = \sum_{k_x, k_y} A(k_x, k_y) \exp [i(k_x x + k_y y)]. \quad (4.31)$$

The CFM has been applied successfully to several pure metals modeled by means of an EAM potential, such as gold and silver [83], aluminum [84, 85], and nickel [81, 86]; yet predictions of the interface stiffness via the CFM are not limited to elemental systems [37] or those modeled within the EAM framework [87, 88].

It is worth noting that results for the interface free energy and anisotropy of a Lennard-Jones system [89] have shown excellent agreement with those obtained with the interface cleaving method [79, 90], thus proving that the CFM is a reliable method for extracting solid-liquid interfacial properties from atomistic simulations. It is interesting therefore to highlight few advantages and drawbacks of both methods.

The main advantage of CM is the requirement of smaller simulations in terms of the number of particles. The work by Davidchack and Laird [90] reports a system with 10^4 particles, while applications of the CFM to the same system required roughly as many as 5 times the number of atoms. Moreover, Davidchack and Laird [91] reports a higher statistical accuracy on the results of the interface free energy.

On the other hand, since the CFM directly measures the stiffness, it can be considered the method of choice to assess the anisotropy with much higher precision. The same information obtained with CM is inevitably associated with a larger error since it is the result of the difference between two numbers — e.g., $\gamma_{sl}^{(100)} - \gamma_{sl}^{(110)} / \gamma_{sl}^{(100)}$ — that can be very close to each other. Lastly, the CFM does not need the construction of any cleaving potentials, but relies only on the

accuracy of the interatomic potentials employed. Therefore, the CFM is the preferable choice for calculating interface free energies of alloys [92], although the large number of particles and the more complex interatomic potential require more computationally demanding simulations.

Interface free energy from metadynamics simulations

It is worth briefly discussing another computational method that exploits an enhanced sampling technique known as metadynamics. Without entering the details of this technique¹⁵, it is interesting to see how the method by Angioletti-Uberti compares with the previously mentioned CFM and CM methods. The central idea of the method aims at reconstructing the free energy surface of the system under study and then to explore transitions from single-phase to two-phase configurations: evaluating the free energy difference between such configurations is it possible to obtain an estimate of γ_{sl} .

For a unary system at its equilibrium melting temperature, we may write the following relation

$$G_{ex} = G_{sl}(p, T_m) - G_s(p, T_m),$$

where the term denoted as “ex” represents the (free) energy cost of introducing a dividing surface. The specific interface free energy is simply the excess term divided by the surface area, that is, $G_{ex} = \gamma_{sl}A$.

A metadynamics simulation allows the system to explore its free energy landscape as a function of one or more collective variables chosen beforehand that well discern between different configurations. Therefore, for the above relation to yield a correct result, one must be able to unambiguously define the free energy of a solid-liquid configuration of the system. Specific details and related issues of the method have been thoroughly discussed in the work by Angioletti-Uberti et al. [72]; here we want to reiterate the gist: for a rather simple interatomic potential such as a Lennard-Jones one, a qualitative representation of the free energy profile along a suitable collective variable would be like fig. 4.12. The two wells correspond to the single-phase configurations of the system, and their equal depths tells us that the system is indeed at its equilibrium melting temperature. Thus G_{sl} can be taken as the value of the free energy in an arbitrary point of the plateau region.

Neglecting for the time being some inherent computational subtleties of metadynamics simulations, when compared to both CFM and CM, the method described has several advantages worth mentioning. The first one, at least from a computational point of view, is the requirement of simulations of a few thousands of atoms. Simulation cells of this size showed [93] to give reliable estimates of γ_{sl} and are about two orders of magnitude smaller than those required by CFM. With respect to CM, which still requires simulation cells containing roughly 1×10^4 atoms, the advantage is that, during the simulation, the system undergoes the phase transition many times and, therefore, it mitigates any hysteresis effect that is instead a fundamental issue of CM [80, 91, 94].

An additional advantage has to do with the accuracy of the results. All three methods discussed so far can be an equally valid choice for obtaining a γ_{sl} estimate that somehow mitigates the huge error associated with measuring this quantity experimentally. However,




Figure 4.12 – Schematic plot of the Gibbs free energy of a unary system versus the value of a suitably chosen collective variable (CV). The equal depth of the two minima, that correspond to single-phase configurations of the system, implies that the chemical potential of the two phases is the same, that is, the system is at equilibrium temperature and pressure. The energy cost associated to a phase transition that brings the system from bulk solid to bulk liquid is precisely $\gamma_{sl} A$, where A is the interface area. Moreover, being at equilibrium means that the value of γ_{sl} obtained with this method is independent of the choice of the CV. The same statement is not valid in out of equilibrium conditions, although it is still possible to employ metadynamics simulations to extract γ_{sl} as a function of temperature [74].

this method based on metadynamics might even aim to provide results with an accuracy at the level of *ab initio* methods. Molecular dynamics simulations of this type are feasible for systems containing hundreds of atoms and for a duration of several hundred picoseconds. These simulations would provide estimates of γ_{sl} with a statistical accuracy far better than that achievable with CFM, mostly because of the more reliable description of atomic interactions. On the other hand, such simulations are more affected by larger statistical errors, due to much shorter simulation times, and finite-size effects. An alternative approach could be to apply the method starting from a simple interatomic potential and to obtain only a rough assessment of the free energy. The next step would employ a more sophisticated potential that more accurately reproduces the real system and use the first free energy estimate as a starting point. One may repeat this procedure several times, including more and more complex potentials, but without having to recalculate each time the free energy from scratch, which will only be refined at each iteration.

In conclusion, it is interesting to note that this method allows to obtain γ_{sl} in out-of-equilibrium conditions [74], even if in this case the value obtained will depend of the collective variable (or order parameter) chosen, and it is necessary to introduce a corrective term that depends on the difference in chemical potential between the two phases.

4.3.3 Interface dynamics

The solidification process is not only controlled by thermodynamics. In fact, because it is a phenomenon that can only begin in conditions of quasi-equilibrium – in the presence of a small

undercooling – kinetics plays a predominant role. If we consider again the Gibbs-Thomson equation, now that it is clear how relevant it is to take into account the anisotropy of the interface properties, we can rewrite it as follows

$$T_{\text{int}} = T_m - \frac{T_m}{L_f} \sum_{i=1,2} \left[\gamma(\hat{n}) + \frac{\partial^2 \gamma(\hat{n})}{\partial \theta_i^2} \right] \frac{1}{R_i} - \frac{V_n}{\mu(\hat{n})}, \quad (4.32)$$

where the dependence of both γ and μ appears explicitly. The term in square brackets represents the stiffness defined with respect to the angles θ_i between the normal \hat{n} and the two principal directions of the interface – R_i are the associated radii of curvature. The last term expresses the undercooling measured at the interface due to the kinetics of atoms attachment. It is evident that interface mobility μ is the quantity that links the average velocity of a planar interface with the undercooling, that is, $V = \mu \Delta T$. Even if not written explicitly, it is reasonable to assume that both the velocity and interface mobility will depend on temperature, therefore assuming linear relationship is an approximation valid only in certain circumstances.

As already mentioned, the difficulty with quantities like γ and μ is twofold: they cannot be neglected even if their anisotropies are small – which can still greatly affect solidification in real materials – and, at the same time, they are extremely elusive to be measured with enough accuracy by experiments. This is an interesting premise to call into play atomistic modeling, and the first attempt to apply molecular dynamics simulation to extract information on μ and its dependence on orientation is due to Broughton, Gilmer, and Jackson [95]. In their work, they proposed to modify an even earlier model of crystallization kinetics by Frenkel [96] and Wilson [97] and wrote the dependence of the interface velocity on temperature as

$$V_n \propto \sqrt{\frac{3k_B T}{m}} e^{-\Delta S/k_B} \left[1 - e^{-\Delta \mu/k_B T} \right]. \quad (4.33)$$

In their model, Broughton, Gilmer and Jackson added a temperature-dependent prefactor – the $3k_B T/m$ term, with m being the mass of an atom – while in Frenkel–Wilson’s model that term depended only on the diffusion coefficient in the liquid phase. The physical justification of this change can be explained by considering which kind of chemical bonding better describes the material under consideration [98]. This model has been verified for a number of systems modeled with the EAM potential [99], therefore it is interesting to briefly discuss how several methods that rely on molecular dynamics could provide information on interface dynamics.

Forced velocity simulations A first attempt to employ molecular dynamics simulations to extract information about the solidification rate has been done by Broughton, Gilmer, and Jackson [95] and Burke, Broughton, and Gilmer [100]. The method they devised is known as *forced velocity simulations* since an external constant velocity along the normal of the solidification front is added to the system at each timestep during its time evolution.

The simulation cell consists of three regions: two static ones whose temperature is controlled by a thermostat – the top region is solid at $T < T_m$ while the bottom one is liquid held at some temperature above melting – and a dynamic region at the center containing the solid–liquid

interface whose temperature was left uncontrolled. During the simulation, the dimensions of the cell that are parallel to the interface are fixed to their values that guarantee the zero-pressure crystalline lattice parameter, while the third dimension is free to fluctuate in order to keep the pressure along the z direction equal to zero.

By monitoring the temperature close to the interface for different values of the external velocity, one is able to extract the relation between the undercooling and the crystal growth rate. The difficulty is that there is no *a priori* rule to find the correct value of the imposed velocity, which must be changed, by trial and error, until one finds that value for which the solid-liquid interface remains approximately in the center of the simulation box.

The principal result that Broughton, Gilmer, and Jackson [95] obtained by applying this method to simulations with very different undercooling contradicted with a previous model of Frenkel [96] and Wilson [97], who suggested that crystal growth was a thermally activated process. On the contrary, in their work Broughton, Gilmer and Jackson found that crystallization rates were still significant even for very low undercooling.

In a more recent work, Celestini and Debierre [101] proposed a slightly more versatile alternative of the force velocity simulation method. The main advantage was the possibility to specify an arbitrary value for the external velocity. The simulation cell was set up in such a way to have four regions whose temperatures were controlled by four different thermal baths, and periodic boundary conditions guaranteed that solid atoms leaving the cell from its top would appear immediately in the solid region at the bottom of the cell. In such configuration of the simulation cell, the steady state is reached once a constant temperature profile establishes across the whole cell. At the steady state, one can obtain the relation between undercooling and velocity by comparing the constant temperature profile and the final positions of the two solid-liquid interfaces.

Free solidification method The free solidification method has the advantage of a much simpler simulation's setup, if compared with the previous forced velocity scheme. A typical simulation involves preparing a box the majority of which consists of a liquid phase, while only a small portion of the system is solid. After equilibrating such starting configuration, the temperature is lowered below the melting temperature and thus solidification spontaneously occurs. Since solidification is an exothermic process, the latent heat released upon the phase transformation must be constantly removed by coupling the system to some kind of thermostat.

The most direct way of estimating interface kinetic coefficient is to track the position of the interface versus the applied undercooling during the solidification. This can be done by monitoring either the total potential energy of the system as described by Hoyt et al. [102], or the number of crystalline atoms N_s according to some kind of order parameter that clearly distinguishes the solid from the liquid phase. In the latter case, an estimate of the crystal growth rate is

$$R = \frac{v_s}{2S} \frac{dN_s}{dt},$$

where v_s is the equilibrium molar volume of the solid and S the interface area. For the free solidification method to work, it is essential to employ the NPT ensemble, where there pressure

is maintained at zero during throughout the simulation. This is required because of the usually significant difference between the densities of the solid and the liquid: a constant volume simulation would lead to a steady increase of pressure which would eventually halt solidification.

Two final remarks are useful to make regarding free solidification simulations. The first is that this method completely ignores the effect of temperature gradients on interface motion, that is, the mobility estimate extracted from free solidification simulations captures only the kinetics of atoms in the liquid phase attaching to the growing solid interface. In real solidification, however, temperature varies across the interface because the latent heat released during solidification does not diffuse instantaneously, and indeed may take longer than the time it takes for the interface to thicken by some crystalline plane. Hoyt, Asta, and Karma [99] have verified, by considering the results of a phase field model on pure melt solidification [103], that the contribution to the kinetic coefficient of temperature variation across the interface in free solidification simulations is about 2% at least for metallic systems, and it is therefore justified to ignore this small correction.

The second remark underscores one drawback of the free solidification method which pertains to statistical accuracy. At low undercooling, the solidification front might move by only very few lattice planes even in long MD simulations. This means that estimating the crystal growth rate from the slope of N_s versus time leads to large uncertainties on R and, in turn, on the kinetic coefficient. Therefore, a method of obtaining information on interface kinetic from a simulation at the melting temperature could be a viable alternative in the low undercooling regime, which may be relevant in some experimental studies of solidification.

Random walk technique A method that allows extracting the kinetic coefficient from monitoring the interface fluctuations has been proposed by Briels and Tepper [104]. The authors studied the fluctuations of the number of crystalline particle for the (100) interface of a solid–liquid Lennard–Jones system at coexistence temperature. The key ingredient of this method is that simulations are performed within the NVT ensemble which is required to keep the average interface position fixed. In particular, this approach proposes that the instantaneous variation of the number of solid atoms from its equilibrium value, ΔN_s , is described by the following relation

$$\frac{\Delta N_s}{t} = \frac{1}{\tau} \Delta N_s,$$

where the time constant τ represents the relaxation time that the interface takes to recover its equilibrium position and it is explicitly linked to the kinetic coefficient. The relaxation time can be obtained from the auto-correlation function of $\Delta N_s(t)$, that is,

$$\frac{\langle \Delta N_s(t) \Delta N_s(0) \rangle}{\langle \Delta N_s(0)^2 \rangle} = e^{-t/\tau}.$$

The assumption made by Briels and Tepper is that the relationship between the interface velocity and the chemical potential difference is a linear one, therefore they wrote

$$V = \frac{k}{k_B T} (\mu_l - \mu_s), \quad (4.34)$$

where k is their definition of the kinetic coefficient. An alternative approach which avoids monitoring explicitly the interface position and exploits the central result of CFM will be discussed in the Results.

Notes

1. The sign adopted for quantities such as work or heat transferred from and to the system is, of course, arbitrary. We are thus choosing the convention by which any amount of energy that *leaves* the system is negative, while any quantity *entering* the system is positive. Hence, the reason for the minus sign in front of the term dW of eq. (4.1).
2. That is, a system that cannot exchange both mass and energy (work or heat) with the surroundings. The isolated system by definition is the Universe, which always comprises any particular system of interest and its environment.
3. «In considering the different homogeneous bodies which can be formed out of any set of component substances, it is convenient to have a term which shall refer solely to the composition and thermodynamics state of any such body without regard to its size or form. The word *phase* has been chosen to this purpose.» [105]
4. This is readily seen by recalling that, for a reversible change in a closed system at constant composition, the First Law reads $dE = dQ + dW$. If there is only expansion work, then $dW_{\text{rev}} = -pdV$ and from the definition of entropy $dQ_{\text{rev}} = TdS$. Combining these two, one obtains $dE = TdS - pdV$, which must be valid regardless what kind of transformation is considered, since the internal energy E is a function of state.
5. These are the so-called *invariant points*. In a single component system, invariant points are outside any external control, for they occur at a definite temperature and pressure that depend only on the substance.
6. The superscript “ex” indicate an excess quantity. In this context it labels a difference between a quantity calculated for a real system and an ideal one.
7. This conclusion perfectly makes sense with the definition of equilibrium temperature from a thermodynamic standpoint: it is that temperature at which the driving force for solidification vanishes.
8. This is because the degree of undercooling affects the velocity of the solid–liquid interface but not atomic diffusion at a given temperature.
9. Gibbs formulated a thorough study of the thermodynamics of surfaces in multicomponent systems in his famous paper entitled “On the Equilibrium of Heterogeneous Substances” [105, pag. 55].
10. For an extensive discussion on surface excess properties, see Lupis [107].
11. The unit normal vector associated to the mathematical surface that coincides with the solid–liquid interface.
12. There have been studies reporting the best agreement with experiments only when including the third–order expansion. See Niederberger, Michler, and Jacot [108].
13. For every point on the surface, there is always a choice of a particular set of orthogonal vectors defining the principal radii of curvature.
14. If \hat{n} is the interface normal, the tangent plane is the space of all the tangent vectors, each of which describes a direction in which the interface is curved.
15. The interested reader may find all the core concepts and details comprehensively discussed in Refs. [109, 110].

Results **Part II**

5 Anisotropy and free energy of a smooth fluctuating dividing surface

This chapter presents and discusses the core results of the paper entitled “Extracting the interfacial free energy and anisotropy from a smooth fluctuating dividing surface” [73]. Both the relevance of interface properties and the related difficulties in determining them have been previously discussed, together with several atomistic simulation approaches. In the research article subject of this chapter, we present a generalization of the construction proposed initially by Willard and Chandler [111]. This method allows one to obtain a smooth dividing surface that follows the ever-changing shape of interfaces that are present in a system with multiple phases, such as solid-liquid systems. The work shows how this construction can be used to study the surface that separates a solid material from its melt in terms of interface’s capillary fluctuations. By analyzing the Fourier transform of this instantaneous dividing surface, one can extract the specific free energy excess of the interface and its dependence on the orientation relative to the bulk phases. The efficiency of this approach is discussed in terms of system size and statistical sampling.

The first section explains the theoretical background behind the idea of a coarse-grained order parameter field and how it naturally defines an instantaneous solid-liquid dividing surface.

The second section presents the system chosen as a benchmark against which to test the framework and how in practice the simulations of solid-liquid interfaces were prepared and run with all the relevant computational details.

The fundamental differences with the original method of Willard and Chandler are discussed in the third section, investigating the sensitivity of the procedure described in the first section in terms of the two main ingredients of the Gaussian kernel density estimation. A detailed presentation of the results for the interface free energy obtained from fitting the capillary fluctuation spectrum follows, together with a comparison with literature works that already investigated the same properties for an identical system.

The last section describes the implementation of the framework we developed into the widely-used simulation package PLUMED [112], that constituted a substantial part of this work.

5.1 Theoretical background

Picking the right order parameter may well be a difficult task, for one choice may end up correctly distinguishing two phases that one would want to ignore. A typical situation could

be the spontaneous formation of defects, or the appearance of an unexpected phase that is identified as different by the order parameter. For example, since the FCCUBIC order parameter is not invariant with respect to rotations or reflexions, it will immediately detect defects such as stacking faults or twin boundaries that can emerge rather easily when studying planar interfaces.

On the other hand, their great flexibility makes them the most convenient choice in atomistic simulations, precisely because they are atom-centered quantities that are, in most cases, very easy and cheap to compute. A Gibbs dividing surface approach is then straightforward: once known the mean values of the order parameter in the bulk phases – e.g., with $\bar{\phi}_l$ and $\bar{\phi}_s$ indicating a liquid and solid phase – then, for a system of N particles one could define

$$\Phi = \sum_{i=1}^N \phi_i = N_s \bar{\phi}_s + (N - N_s) \bar{\phi}_l, \quad (5.1)$$

where N_s stands for the number of solid atoms. This definition implicitly determines the location of the dividing surface between the solid and liquid by requiring a vanishing surface excess for this particular choice of ϕ . It also implies that any other macroscopic observable X will have a non-zero surface excess quantity, since in general $X \neq N_s \bar{X}_s + (N - N_s) \bar{X}_l$.

Gibbs' approach is undoubtedly elegant and simple, but many other approaches¹ have been proposed to define a surface in a way that is sometimes more convenient – yet always arbitrary. Here we want to discuss the idea that Willard and Chandler [111] introduced and that has some analogy with phase field methods mentioned in the introduction. Their approach was the starting point for the method used to extract interfacial free energies from fluctuating surfaces [73], whose results will be discussed in the next part of the thesis.

The basic ingredient of Willard and Chandler's idea is a continuous density field defined as follows

$$\rho(\mathbf{r}, t) = \sum_i \delta(\mathbf{r} - \mathbf{r}_i(t)), \quad (5.2)$$

where the summation runs over the number of particles in the system, and δ is simply a Dirac delta. A coarse-grained version of this field is obtained by a convolution with normalized kernel functions, which are taken to be Gaussians of the form

$$K(\mathbf{r}; \xi) = (2\pi\xi^2)^{-3/2} \exp(-r^2/2\xi^2),$$

where ξ is the coarse-graining parameter that should be chosen appropriately based on the physical system under study. The convolution yields

$$\bar{\rho}(\mathbf{r}, t) = \sum_i K(|\mathbf{r} - \mathbf{r}_i(t)|; \xi). \quad (5.3)$$

One may employ a slightly different field that we may define as an *order parameter density field*

$$\rho_\phi(\mathbf{r}, t) = \sum_i \phi_i \delta(\mathbf{r} - \mathbf{r}_i(t)), \quad (5.4)$$

where Φ can be the FCCUBIC or any other choice of a suitable order parameter for the description of the system under consideration. Once the coarse-grained field has been normalized,

$$\bar{\rho}_\phi(\mathbf{r}, t) = \frac{\sum_i \phi_i K(|\mathbf{r} - \mathbf{r}_i|; \xi)}{\sum_i K(|\mathbf{r} - \mathbf{r}_i|; \xi)}, \quad (5.5)$$

then the interface between the phases that Φ discriminates is defined by the two-dimensional surface satisfying

$$I \equiv \{(x, y) | \bar{\rho}_\phi(\mathbf{r}, t) = \phi_c\}, \quad (5.6)$$

where ϕ_c is a constant value to be fixed. In other words, the instantaneous dividing surface is the locus of points in space where the order parameter density field assumes the value Φ_c . It is clear that the choice of this constant is arbitrary, but it should rely on some meaningful physical consideration. Since the coarse-grained field depends explicitly on time through atoms' coordinates, the interface will also be a function of time: $I \equiv I(x, y, t)$.

This approach is particularly well suited for molecular dynamics which naturally provides the time evolution of a system by means of positions and velocities of all its particles. In practice, eq. (5.6) can be solved on the fly during a molecular dynamics simulation by employing one among many root finding algorithms.

5.2 System and simulation details

The simulations presented in this paper were of an interface between the solid and liquid phases of Lennard-Jones atoms. A truncated form of the Lennard-Jones potential [113] was used throughout, and thus the pair potential was given by the following expression:

$$V(r) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right] + C_1 & r \leq 2.3\sigma \\ C_2 \left(\frac{\sigma}{r}\right)^{12} + C_3 \left(\frac{\sigma}{r}\right)^6 + C_4 \left(\frac{r}{\sigma}\right)^2 + C_5 & 2.3\sigma < r < 2.5\sigma, \\ 0 & r \geq 2.5\sigma \end{cases}, \quad (5.7)$$

where $C_1 = 0.016132\epsilon$, $C_2 = 3136.6\epsilon$, $C_3 = -68.069\epsilon$, $C_4 = -0.083312\epsilon$ and $C_5 = 0.74689\epsilon$. As it is customary when employing an interatomic potential of this kind, all the relevant units are expressed in term of quantities derived from ϵ , the unit of energy, and σ , the unit of distance.

All molecular dynamics simulations were performed using LAMMPS [114] at the equilibrium melting temperature of 0.6185 Lennard-Jones units. In the simulations the temperature and pressure were kept constant using a Nosé-Hoover thermostat and barostat. To monitor which atoms were in the solid parts of the simulation box and which atoms were liquid we used a cubic harmonic symmetry function already described previously (section 4.3.1) and that has been used in a number of other recent works [72, 74, 93].

To generate an initial configuration containing a solid-liquid interface we took a perfect crystalline FCC unit cell with a lattice parameter that was consistent with the density at the chosen temperature (i.e. the melting temperature $T_m = 0.6185$). Since we wanted to simulate solid-liquid interfaces with multiple different orientations, we aligned the z axis of

Table 5.1 – Dimensions of the simulations and the number of atoms in each of the cells. The fourth column gives the dimension along the z axis, perpendicular to the interface normal. All lengths are in units of σ .

Orientation	L_x	L_y	L_z	N
(100)	32.354	32.354	80.886	80 000
(110) large	45.756	32.354	68.634	96 000
(110) small	22.878	19.413	54.907	23 040

the coordinate system for our cell with the normal to the desired surface. We then replicated the unit cell in the xy plane and along the z direction. When doing this we used a large number of replicas in the z direction to ensure that we would have reasonably thick solid and liquid regions. This reduces the severity of finite-size effects and prevents the system from melting or freezing in its entirety.

The initial stages of our molecular dynamics simulations were used to generate the solid-liquid interface. In these early simulations the atoms with z positions in a particular range were held fixed. The size of this fixed region depended on the length of the supercell in the z direction but in general it was set between $L_z/3$ and $2L_z/3$, where L_z is the total length of the box in the z -direction. These constraints were kept in place during an MD simulation of approximately 100 LJ time units, with a timestep of 0.004 time units, that was run in NVT ensemble with the temperature fixed well above T_m . During this simulation the unconstrained portion of the supercell was observed to melt completely so that two solid-liquid interfaces were formed. With these interfaces in place, the system was then equilibrated at the melting temperature, T_m , for a further 100 time units in the NPT ensemble with the constraints on the atoms in the solid parts of the system still in place. Lastly, the constraints on the solid atoms were removed entirely and a final 100 time units of NVT equilibration of the system were performed. Once this final equilibration had completed a 10 000 time-units run, NPT production run was performed.

Interfaces between the (100) and the (110) surfaces and the melt were generated using the procedure outlined in the previous paragraph. For the simulations of the (100) interface a supercell of of $20 \times 20 \times 50$ unit cells along the standard x , y and z crystallographic axes was generated.

For the simulations of the (110) interface supercells with the x , y and z axes aligned with the $[\bar{1}10]$, $[00\bar{2}]$ and $[110]$ crystallographic directions respectively were generated. There were thus 80 000 atoms in our simulations of the (100) interface and 96 000 and atoms in the first of our simulations of the (110) interface. In order to assess the validity of our proposed model for fitting the Fourier spectrum we also prepared a cell with a (110) interface that contained roughly a quarter of the atoms that were in this large cell. Table 5.1 provides a summary of the various simulation cells that were used in this work.

5.3 Results and discussion

We began our analysis by investigating the sensitivity of the procedure outlined in the previous section to the two main parameters in the recipe; namely, the bandwidth of the Gaussian kernel functions in eq. (5.5), ξ , and the spacing between the grid points at which this function is evaluated. These initial tests were performed on trajectories in which the (100) interface was simulated. We found that the ensemble average for the Fourier transform of the phase field was almost identical for grid spacings of 0.25 and 0.5 Lennard–Jones units and thus opted to use the larger spacing in order to lower the total number of grid points and to thus reduce the computational cost.




Figure 5.1 – Fitted Fourier spectra along the k_x direction for a range of values of the Gaussian kernel coarse-graining length ξ (see eq. (5.4)).

Figure 5.1 shows how the x -component of the ensemble average for the Fourier transform changes as the bandwidth, ξ , is adjusted. According to the capillary fluctuation relation eq. (4.30), $A(k)$ should decay following $A(k) \propto 1/k^2$. We find that this rule is followed for small k , but that there are clear deviations from this behavior at large k . These deviations at large k are unsurprising, however, as the capillary fluctuation dispersion is known to be affected by artifacts that depend on how the interface profile is constructed. These effects are dramatic at wave-vectors that describe fluctuations over length-scales of individual atoms, but deviations from linearity can also be apparent at much smaller values of k and can thus affect any estimate of the interface stiffness coefficients. When it comes to these deviations there is a clear advantage in defining the position of the interface from the isocontour points of a smooth kernel density as we have some indication as to the form these artifacts should take. The definition of coarse-grained density field given by eq. (5.3) describes a convolution of the atomic density with the chosen kernel function; so the Fourier transform of the smoothed density $\bar{\rho}(\mathbf{r}, t)$ is the point-wise product of the Fourier transform of the density field and the Fourier transform of the kernel function. The onset of the deviations from linearity induced by the smoothing should therefore depend on the value of the bandwidth ξ . In particular, as ξ increases, the onset of the deviation should move to smaller k -vectors.

To account and correct for the impact of the density smoothing on the capillary fluctuation

dispersion, we introduced a different form of eq. (4.30)

$$\langle |A(\mathbf{k})|^2 \rangle = \frac{k_B T \exp(-k_x^2/2\lambda_x^2 - k_y^2/2\lambda_y^2)}{S (\sigma_{11}k_x^2 + \sigma_{22}k_y^2 + 2\sigma_{12}k_xk_y)}. \quad (5.8)$$

This is the form that would result from convoluting the “ideal” interface profile with a Gaussian function. However, this argument is evidently an oversimplification of the matter as the expression for the phase field equation (eq. (5.3)) is not a simple convolution of the atomic order parameters. In fact, we are finding the location of the dividing surface from an isocontour of a smoothed order parameter density field and are not smoothening the isocontour. For this reason, we considered the two parameters λ_x and λ_y as fitting parameters. We also found that the best fit values typically differ by less than 25% from the value of $1/\xi$ that would be expected if the Gaussian convolution argument could be considered strictly valid. Better still, we observed that when eq. (5.8) rather than eq. (4.30) is used as the model to fit the data, and λ_x and λ_y are just two additional fitting parameters, the artifacts due to the smoothening are taken into account, and the final values obtained for the stiffness tensor components, σ_{11} , σ_{22} and σ_{12} , do not depend significantly on the value chosen as the Gaussian kernel bandwidth.

The fact that results of the fit based on the model given by eq. (5.8) ended up being relatively insensitive to the manner in which the dividing surface is constructed lead to significant computational advantages. For example, as discussed in the previous paragraph, eq. (4.30) is only valid in long wavelength limit of the capillary wave spectrum since the amplitude of large \mathbf{k} depends more strongly on the way in which the height profile is obtained. In previous simulation works on CFM [115], very large simulation cell have been used and $\langle |A(k_x, k_y)|^2 \rangle$ has been calculated for as many small \mathbf{k} values as possible. The fact that we can correct for the artifacts induced by the smoothening of the phase field ensures that, by contrast, we are able to extract meaningful information from larger values of k , and thus allows us to reduce the size of the simulation. Figures 5.2 and 5.3 show that we can get values for the stiffness tensor that are consistent with those in previous works even though we used a much smaller simulation cell, and that we can extend the fitting to wave-vectors with larger moduli, as long as we use eq. (5.8) when performing the fitting. In fact, figs. 5.2 and 5.3 show the values of the stiffness that we extracted as a function of the radius of the region in k -space, k_{\max} , from which we took the data. It is evident that when this radius is made larger, we include more data points in the fitting, thereby reducing the statistical uncertainty in the stiffness.

A further benefit of using the information on $\langle |A(k_x, k_y)|^2 \rangle$ for large values of k_x and k_y is that it allows reducing the amount of time for which simulations have to be run to achieve statistical convergence. It has been already shown (section 4.3.3) that the dynamics of $A(k_x, k_y)$ can be seen as that of a harmonic oscillator of frequency $\omega \propto k$ undergoing over-damped Langevin dynamics. The auto-correlation time for such processes is proportional to $1/\omega^2$ [116], so longer simulation times are required to converge the averages associated with the modes with long wavelengths. Furthermore, difficulties associated with converging the Fourier amplitude fluctuations for the longer wavelength modes might also be the cause of the non-Gaussian behavior that was observed for these modes in simulations of the interface between solid and




Figure 5.2 – Fitted values for the stiffness tensor for the (100) interface as a function of the radius of the cutoff in k -space. The points on the lines labeled *orig* were obtained by performing fits using eq. (4.30), which is the model that has been used in many other works on CFM. The points on the lines labeled *aniso* were calculated by fitting using our model (eq. (5.8)) and the points on the lines labeled *iso* were fitted using a modified version of the eq. (5.8) in which $\lambda_x = \lambda_y$. The horizontal, grey dashed line indicates the value of the stiffness obtained by Becker et al. [115]. It is clear from this figure that, when the fit is performed using the Gaussian-corrected equation, stiffness values that are consistent with the known values of this quantity can be obtained for a range of k_{\max} values. By contrast, when the fit is performed using the original model of eq. (4.30) the value of the stiffness differs substantially from the known values for all but the smallest values of k_{\max} . This is unfortunate as the results in the region close to $k = 0$ are affected by substantial errors as only a small number of points are included in the fit. For this interface, the anisotropic and isotropic version of eq. (5.8) give very similar values for the stiffness tensor as it should be, given that the (100) interface has only one independent value for the stiffness tensor.

liquid sodium chloride [88]. Figure 5.4 shows that these statements agree with the observations we made for our system. We found that the auto-correlation time for the modes increases with $1/k^2$ when k is small. In conclusion, the fact that we were able to use the information on $\langle |A(k_x, k_y)|^2 \rangle$ at relatively large values of k_x and k_y when fitting using eq. (5.8) ensures that we can extract information from simulations with both shorter timescales and fewer particles as compared to those required when using the “original” model given by eq. (4.30).

Figure 5.3 illustrated why we chose an anisotropic kernel, that is, two independent values of λ treated as fitting parameters. Our experience here is that for an isotropic surface such as the (100) the inclusion of the additional parameter makes very little difference. We get very similar values for σ_{11} and σ_{22} if we use an isotropic kernel and enforce λ_x and λ_y to be equal in the fit. However, for an anisotropic surface such as the (110) we see a weaker dependence on k_{\max} and much closer agreement between our results and the results in the literature when we include separate λ_x and λ_y parameters in the fit. Furthermore, the fitted values for σ_{11} and σ_{22} are more consistent across the range of k_{\max} values when the λ_x and λ_y parameters are allowed to vary independently during the fit.




Figure 5.3 – Fitted values for the stiffness tensor of the (110) interface as a function of the radius of the cutoff in k -space. The labels used in legend that describe each of the lines are explained in the caption to fig. 5.2. This figure once again demonstrates that fitting using eq. (4.30) only gives values for the stiffness that are consistent with the literature values when k_{\max} is set equal to a small value. If, however, the data is fit using eq. (5.8) values for the components of the stiffness tensor are consistent with the known literature values even for a wider range of values of k_{\max} . It is interesting to note that the (110) has two independent values in its stiffness. It would seem that fitting using the anisotropic version of eq. (5.8) gives values that agree more closely with those from Becker et al. [115]. When the isotropic version of this model is used, the values obtained for the stiffness tensor are pulled closer together than they should be.




Figure 5.4 – Plot illustrating the auto-correlation time τ for $A(k_x, k_y)$ in Lennard–Jones reduced units of $\sqrt{m\sigma^2/\epsilon}$ as a function of k_x (red circles) and k_y (blue diamonds) for the (100) interface.

The values of the stiffness presented in fig. 5.3 were obtained by analyzing the smaller of the two simulations that were run on this particular interface. Figure 5.5 shows the values for the stiffness that were obtained from this simulation together with the values obtained by analyzing the larger of the two simulation cells. The size of the cell in this larger simulation is comparable with the size used in the work by Becker et al. [115] from which we took the values we have




Figure 5.5 – Fitted values for the stiffness tensor for the (110) interface as a function of the radius of the cutoff in k -space. All the values on this figure were obtained by performing fits using eq. (5.8). The dashed lines give the values that were obtained from a simulation cell containing 96 000 atoms, while the solid lined are the values obtained from simulations cells with only 24 000 atoms.

used as a reference. From fig. 5.5 it is indeed evident that we can use much smaller simulation cells when we fit using eq. (5.8). The values for the components of the stiffness tensor that we obtained from these smaller simulations are within $0.02 (\epsilon/\sigma^2)$ of the reference values and are in some cases closer to these reference values than the values we obtain from the larger simulations.

The final stiffness values that we obtain are reported in table 5.2. Error bars were calculated using a block averaging technique (section 2.7.1), with a block length that was determined based on the auto-correlation times for the $(2\pi j_x/L_x, 0)$ and $(0, 2\pi j_y/L_y)$ components of the Fourier transform $A(k_x, k_y)$. The values that we obtain for the stiffness values are consistent with those from the the previous work of Becker and co-workers [115]. However, as discussed in the previous paragraphs, we were able to use much smaller simulation cells because we use a fluctuating dividing surface construction and because we perform the fit using an expression like eq. (5.8) that models the short wavelength modes of the dividing surface using a Gaussian convolution.

The values obtained for the stiffness can be used to get information on the interface free energy exploiting an expansion such as that of eq. (4.26), where one should first substitute the components of the interface normal vector. The second step is to derive analytic expressions for the interfacial free energy for each interface. For example, for the (110) interface, eq. (4.26) up to the second order yields

$$\gamma_{(110)} = \gamma_0 \left(1 + \frac{7}{5} \epsilon_1 + \frac{25}{7} \epsilon_2 \right), \quad (5.9)$$

Table 5.2 – Stiffness values for the (100) and (110) interfaces calculated using our method at a reduced temperature of 0.6185. The values of the stiffness are given in units of ϵ/σ^2 . For comparison, the values in parenthesis are those obtained by Becker et al. [115] at the same temperature.

Orientation	σ_{11}		σ_{22}	
	Our work	Becker et al.	Our work	Becker et al.
(100)	0.2897 ± 0.0008	0.2866	0.2871 ± 0.0007	0.2866
(110)	0.429 ± 0.004	0.431	0.271 ± 0.003	0.305

while by exploiting eq. (4.27) the stiffness tensor is given by

$$\sigma_{(110)} = \gamma_0 \begin{pmatrix} 1 - \frac{21}{10}\epsilon_1 + \frac{365}{14}\epsilon_2 & 0 \\ 0 & 1 + \frac{39}{10}\epsilon_1 + \frac{155}{14}\epsilon_2 \end{pmatrix}. \quad (5.10)$$

We have extracted three values for the component of the stiffness tensors of the (100) and (110) surfaces by analyzing the Fourier spectra of the capillary fluctuations in our MD simulations. We can thus extract the three anisotropy parameters for a system modeled with the chosen Lennard–Jones potential by solving simultaneously the corresponding equations that relate the stiffness to γ_0 , ϵ_1 and ϵ_2 . The solutions of this system of equations are $\gamma_0 = 0.35558 \epsilon/\sigma^2$, $\epsilon_1 = 0.064$ and $\epsilon_2 = -0.0039$. Obviously, in this context, ϵ and σ refer to the fundamental Lennard–Jones units of energy and distance, and not to anisotropy parameters or stiffness. The coincident name is unfortunate, yet it should not give rise to any ambiguity. Plugging these values into eq. (4.26) leads to the values for the interfacial free energy for high–symmetry directions. These values are shown in a spherical contour plot of fig. 5.6 and reported in table 5.3 where they are also compared with the values obtained using CFM [115] and metadynamics [66] in the work of Cheng, Tribello, and Ceriotti [117]. Despite the noticeable discrepancies in the values of the stiffness, the fitted values of γ are in near-perfect agreement with previous results based on CFM calculations. There is a more significant discrepancy between the values we obtained and the values that were not calculated by performing a fit based on stiffness, but which were instead computed directly using metadynamics simulations (section 4.3.2).

In conclusion, it is interesting to note that all these methods can capture the anisotropy between different directions, given the difficulties associated with measuring and computing solid–liquid interfacial properties and also given that the Lennard–Jones potential is known to exhibit a very isotropic interfacial free energy. Being able to compute this quantity is critical as small anisotropies underlay important phenomena such as the formation of dendrites during the solidification of metals and alloys.




Figure 5.6 – Graphical representation that illustrated the dependence of the free energy of the interface on its orientation. The high symmetry directions are highlighted and it is clear that the (100) has a higher interface free energy than the (111). A perfect sphere would indicate a fully isotropic interface free energy.

5.4 Details about the order parameter for FCC symmetry

The symmetry function that has been used in this work to distinguish atoms in the solid phase has been already introduced previously (section 4.3.1). Here we discuss few technical details of the FCCUBIC order parameter.

The angular part denoted as $C_\alpha(\mathbf{r}_{ij}; \varphi, \psi, \theta)$ in eq. (4.22), when expressed in terms of the Cartesian components of the distance vector \mathbf{r}_{ij} reads

$$C_\alpha(\mathbf{r}_{ij}; \varphi, \psi, \theta) = \frac{x^4 y^4 + x^4 z^4 + y^4 z^4}{r_{ij}^8} - \frac{\zeta x^4 y^4 z^4}{r_{ij}^4},$$

where ζ is a parameter that is set equal to 27, which is the value that allows the function to best resolve between solid and liquid configurations. The radial term $C_r(r_{ij})$ is a switching function that acts on the length of the distance vector \mathbf{r}_{ij} and has the following form

$$C_r(x) = \begin{cases} 1 & x \leq d_0 \\ 0 & x \geq d_1 \\ (y-1)^2(1+2y) & d_0 < x < d_1 \quad \text{where } y = (x-d_0)/(d_1-d_0) \end{cases}$$

where the parameters d_0 and d_1 should be suitably tuned for the system under consideration.

Table 5.3 – Specific free energy for the solid–liquid interface in a Lennard–Jones system at the melting temperature of 0.6185. The results obtained here are compared with those in previous work of Becker *et al.* [115] and Cheng, Tribello, and Ceriotti [117]. The free energy is in units of ϵ/σ^2 .

γ_{hkl}	$\gamma_{(100)}$	$\gamma_{(110)}$	$\gamma_{(111)}$
Our work	0.364	0.355	0.348
Becker <i>et al.</i>	0.363	0.354	0.350
Cheng <i>et al.</i>	0.373	0.360	0.352

For the case of the Lennard–Jones simulations whose results have been discussed in this chapter, we set $d_0 = 1.2$ and $d_1 = 1.5$.

It is possible to show that the value of Φ_i is equal to 1/16 when the atoms in the first coordination sphere around atom i sit in their ideal FCC lattice sites. Since the FCCUBIC order parameter is not rotationally invariant, $\Phi_i = 1/16$ for an atom in an ideal FCC lattice site only if the simulation cell vector are aligned with the vectors of the FCC structure. For the simulations set up to analyze the (110) interface, one should employ the correct values for the three Euler angles (φ, ψ, θ) so as to ensure that, for example, the z direction runs parallel to the [110] direction and the (110) interface coincides with the xy plane.

If the atomic environment around atom i is uniform – i.e., perfectly isotropic as it would be for a liquid – then $\Phi = \frac{143-\zeta}{5005}$. Therefore, to ensure that $\Phi = 1$ corresponds to a perfect FCC lattice, while its average value for bulk liquid is zero, a simple linear scaling

$$\Phi'_i = \frac{\Phi_i - \langle \Phi \rangle_l}{\langle \Phi \rangle_s - \langle \Phi \rangle_l}$$

is applied, where the symbol $\langle \cdot \rangle_p$ denotes an average in the p phase. Figure 5.7 shows the distributions for the values of the resulting quantities for the atoms in a bulk solid and a bulk liquid. One can see that there is very little overlap between the distribution of the values in bulk solid and bulk liquid. To make these distributions even more strongly peaked and the distinction between atoms in the liquid and solid parts even clearer, Φ is transformed by applying a non–linear switching function of the kind

$$s(\Phi') = \left[1 + (2^{a/b} - 1) \left(\frac{\Phi' - d_0}{r_0} \right)^a \right]^{-b/a}, \quad (5.11)$$

where r_0 and d_0 are tunable parameters. For $r \leq d_0$ then $s(\Phi')$ smoothly converges to 1.0, while for $r > d_0$ the function decays to 0.0. a and b are two parameters that control the shape of the switching function. The choice of these parameters is arbitrary, but one could use as a visual guide the plot of the distribution function for the chosen order parameter – that of the FCCUBIC is shown in fig. 5.7. One can then tune the values of a and b in such a way that the sigmoid function of eq. (5.11) better filters out the overlapping regions and thus enhances the separation of the solid phase from the liquid one.

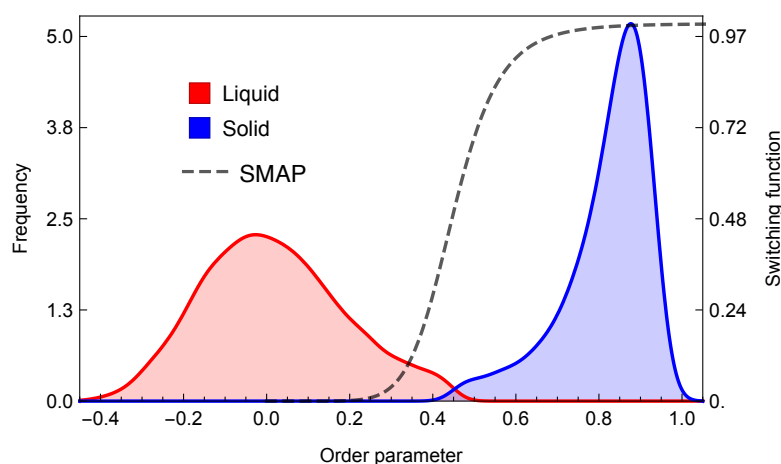


Figure 5.7 – Probability density distribution function for the order parameter that is described in this section. The histograms shown in red and blue in this figure were calculated from a snapshot of the (110) simulation. The dashed line shows the switching function of eq. (5.11), that was used to convert the order parameter values into a continuous quantity that better distinguished the solid from the liquid.

5.5 Implementation

This section illustrates in detail the implementation of the fluctuating dividing surface method in the molecular dynamics package PLUMED [112] and presents commented input files used to analyze a CFM simulation.

Once the order parameter Φ has been calculated for every atom in the simulation, the first step is to build the density field, a scalar field that allows evaluating the value of Φ in every point in space. By averaging such density field along two directions – say, x and y – it is possible to obtain a density profile along the direction perpendicular to interface. In this way, one could verify that the chosen order parameter is capable to distinguish the solid from the liquid phase. An input like the following would allow calculating one such profile with PLUMED:

```

UNITS NATURAL 1
2
FCCUBIC ... 3
LABEL=fcc 4
SPECIES=1-20736 SWITCH={CUBIC D_0=1.2 D_MAX=1.5} 5
ALPHA=27 PHI=0.0 THETA=-1.5708 PSI=-2.35619 6
... FCCUBIC 7
8
center: CENTER_OF_MULTICOLVAR DATA=fcc 9
10
MULTICOLVARDENS ... 11
LABEL=dens 12
DATA=fcc ORIGIN=center DIR=z 13

```

```
NBINS=100 BANDWIDTH=1.0 STRIDE=1 CLEAR=25
... MULTICOLVARDENS
```

14
15
16
17

```
DUMPGRID GRID=dens FILE=profile.dat STRIDE=25
```

Lines 3 to 7 of this input file tells PLUMED that we would like to calculate the Φ order parameter of eq. (4.22) for all the 20 736 atoms in the system. These Φ_i values are then inserted in the expression below

$$\langle \varphi(z) \rangle = \left\langle \frac{\sum_{i=1}^N s_i K\left(\frac{z-z_i}{\lambda}\right)}{\sum_{i=1}^N K\left(\frac{z-z_i}{\lambda}\right)} \right\rangle, \quad (5.12)$$

which is done in lines 11 to 15 of the input file. In this expression K is a normalized kernel function (in this work a Gaussian with bandwidth set equal to 1.0σ , where σ stands for the unit of length in Lennard–Jones reduced scheme) and z_i is the z coordinate of the i th atom relative to a reference position. In this work this reference position is located in the centre of the solid region of the system and is found by taking the following weighted sum

$$z_{\text{ref}} = \frac{L_z}{2\pi} \tan^{-1} \left(\frac{\sum_{i=1}^N \Phi_i \sin\left(\frac{2\pi z_i}{L_z}\right)}{\sum_{i=1}^N \Phi_i \cos\left(\frac{2\pi z_i}{L_z}\right)} \right).$$

In this expression L_z is the length of the simulation cell in the z direction, Φ_i is the value of eq. (4.22) evaluate from the positions of the atoms around atom i , whose z coordinate is z_i . This quantity is calculated by the CENTER_OF_MULTICOLVAR, which also calculates similar weighted-averages from the x and y positions of the input atoms.

The ensemble average of the order parameter is calculated by computing the function inside the angle brackets in eq. (5.12) every step and by averaging over 25 trajectory snapshots. The final ensemble average that is calculated in this way is shown in fig. 5.8. This result is unsurprising: there is a region in the center of the box where the average value of Φ_i is large and where the structure is a solid. In the remainder of the box, however, where the structure is liquid, the average value of Φ_i is small.




Figure 5.8 – Average value of the FCCUBIC symmetry function defined in eqs. (4.22) and (5.11) as a function of x , the direction parallel to the interface normal.

It is easy to adjust the PLUMED input that was used to generate fig. 5.8 to calculate two or three dimensional profiles. Furthermore, by taking advantage of the STRIDE parameter in the DUMPGRID command and the CLEAR parameter in the MULTICOLVARDENS command we can easily adjust the number of frames over which we calculate the ensemble average. Finally, we can obviously change the quantity we are averaging by using some action besides FCCUBIC to calculate the symmetry functions that are used to calculate the average field in MULTICOLVARDENS.




Figure 5.9 – A snapshot from one of molecular dynamics simulations of a solid–liquid interface in a Lennard–Jones system. In this figure the atoms are colored according to the value they have for the symmetry function defined in eqs. (4.22) and (5.11). Those colored blue have a high value for Φ_i and are thus in a solid–like environment. Those colored in red have a low value for this quantity. The structure surrounding these atoms is thus more liquid–like.

Within the CFM framework, our aim is not to simply calculate the average value of some symmetry function. Instead we would like to calculate the surface that separates the liquid phase from the solid phase and to investigate how this surface fluctuates as the simulation progresses. Figure 5.9 shows that this objective is not unreasonable. The atoms in a single snapshot of the trajectory are shown colored according to their scaled Φ_i parameter. The atoms with a high Φ_i value that therefore have an environment that resembles that found in the solid are colored in blue, while those with a low, liquid-like value are colored in red. This figure shows clearly that the box can be partitioned into liquid and solid parts and that there is clear, albeit non–flat, dividing surface between these regions. The following input to PLUMED uses our implementation of the Willard–Chandler surface to find the location of this dividing surface.

```

UNITS NATURAL 1
FCCUBIC ... 2
LABEL=fcc 3
SPECIES=1-20736 SWITCH={CUBIC D_0=1.2 D_MAX=1.5} 4
ALPHA=27 5
... FCCUBIC 6

```

```

smapfcc: MTRANSFORM_MORE DATA=fcc SWITCH={SMAP R_0=0.5 A=8 B=8}
center: CENTER_OF_MULTICOLVAR DATA=smapfcc

MULTICOLVARDENS ...
LABEL=dens
DATA=smapfcc ORIGIN=center DIR=xyz
NBINS=50,80,80 BANDWIDTH=1.0,1.0,1.0
STRIDE=1 CLEAR=1
... MULTICOLVARDENS

FIND_CONTOUR GRID=dens CONTOUR=0.5 FILE=mycontour.dat STRIDE=1

```

This input instructs PLUMED to calculate the scaled value of Φ_i using eq. (4.22) for each of the atoms in the system. These quantities are then transformed using equation eq. (5.11) and a spatial average for the resulting quantity is evaluated on a three dimensional grid that covers the whole simulation cell using equation eq. (5.4). The two-dimensional manifold of points for which the resulting spatial average of the transformed and scaled Φ_i values equals 0.5 is then found using a variant on the marching cubes algorithm. This algorithm searches over all the grid points at which the spatial average is evaluated. At each grid point the difference between the value of the spatial average at that point and the target value of 0.5 is evaluated. The sign of this quantity is then compared with the sign of the same quantity evaluated at the adjacent grid points in the x , y and z directions. If the sign of this quantity on adjacent grid points in any of these directions do not match, then the manifold separating the solid from the liquid must pass through the vector connecting these two grid points. We thus find points on the manifold separating the solid from the liquid by searching between pairs of grid points with opposing signs for $\varphi(x, y, z) - \varphi_0$ using Brent's root-finding algorithm [118] and splines that interpolate the value of the spatial average between grid points from the values that are evaluated at the grid points. The positions of the resulting points are then output to a file. A surface passing through the set of points on the manifold that were found when we analyzed the configuration shown in fig. 5.9 is shown in fig. 5.10.

It is important to note that in the PLUMED input that we used for this calculation the CLEAR and STRIDE parameters in the MULTICOLVARDENS action are set equal. Setting these two parameters to be equal ensures that the grid accumulators are reset after each calculation of the contour. The contour found thus provides information on the instantaneous location of the interface between the solid and liquid parts of the cell. If the CLEAR parameter is set larger than the STRIDE parameter, then an ensemble average for the spatial average of Φ_i is computed over a number of trajectory frames. One can then use FIND_CONTOUR to find the location of the dividing surface for this averaged profile.

The marching cubes algorithm that was used to find points on the manifold separating the solid from the liquid in the previous section is useful as information on the expected geometry of the contour is not required by this algorithm. It can thus find the manifold separating the phases regardless of whether the phase has a slab geometry, a spherical geometry or even a




Figure 5.10 – A snapshot from one of molecular dynamics simulations of a solid–liquid interface in a Lennard–Jones system with the atoms colored as described in the caption of fig. 5.9. The Willard–Chandler surface that separates the solid and liquid phases is shown in yellow in this figure.

toroidal geometry. Although this flexibility is useful in general, it is often unnecessary and can in fact make subsequent analysis more complicated. For this reason we provide a second method for finding the location of the dividing surface. In this second method we only search for points on the dividing surface in the direction perpendicular to the interface between the two phases. Furthermore, in this second method we do not calculate the value of the spatial average on a grid that encompasses the whole simulation cell. Instead, we only calculate the value of the spatial average in the parts of the box that we suppose are close to one of the dividing surfaces. The advantage of changing the manner in which the location of the surface is found in these two ways is illustrated in figure fig. 5.11. As this figure shows this new method finds one value for the height, z , of the dividing surface for each pair of x and y values in our three dimensional grid. We can thus use a contour plot to display $z(x, y)$ and we can interpolate and integrate this interfacial-height-profile function.

The PLUMED input that was used to generate the contour plot in fig. 5.11 is given below:

```

UNITS NATURAL 1
FCCUBIC ... 2
LABEL=fcc 3
SPECIES=1-20736 SWITCH={CUBIC D_0=1.2 D_MAX=1.5} 4
ALPHA=27 5
... FCCUBIC 6
7
smapfcc: MTRANSFORM_MORE DATA=fcc SWITCH={SMAP R_0=0.5 A=8 B=8} 8
center: CENTER_OF_MULTICOLVAR DATA=smapfcc 9
10

```

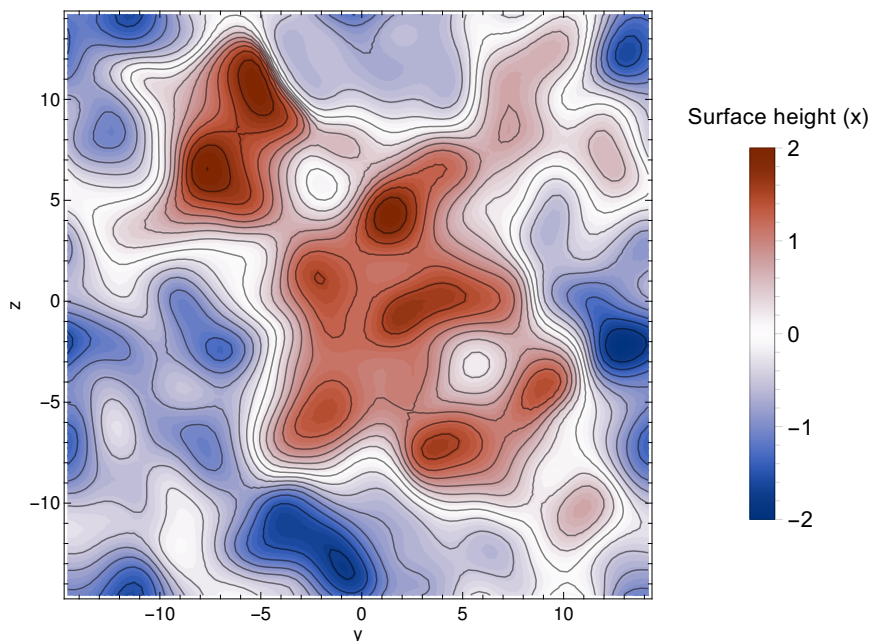


Figure 5.11 – A contour plot showing the height of the interface between the solid and liquid phases of Lennard–Jones versus x and y for the snapshot of the trajectory shown in fig. 5.9.

```

MULTICOLVARDENS ... 11
LABEL=dens 12
DATA=smapfcc ORIGIN=center DIR=xyz 13
NBINS=50,80,80 BANDWIDTH=1.0,1.0,1.0 14
XREDUCED XLOWER=0.0 XUPPER=20.0 15
STRIDE=1 CLEAR=1 16
... MULTICOLVARDENS 17

FIND_CONTOUR_SURFACE ... 18
LABEL=contour 19
GRID=dens CONTOUR=0.5 SEARCHDIR=x 20
STRIDE=1 CLEAR=1 21
... FIND_CONTOUR_SURFACE 22
23
24
DUMPGRID GRID=contour FILE=contour.dat 25

```

The keyword CONTOUR is the same as for the action FIND_CONTOUR, while the keyword SEARCHDIR tells PLUMED in which directions to search for the dividing surface. It is worth noting that the output from the contour surface found by the FIND_CONTOUR_SURFACE action is stored as a function on a grid within PLUMED. In other words, the object output by this action has the same type as the object output by MULTICOLVARDENS. Everything we can do to the output from MULTICOLVARDENS – be that printing it to a file, interpolating the function,

integrating it or even finding contours within it — can thus also be done to the output of FIND_CONTOUR_SURFACE.

In this work we took the output from FIND_CONTOUR_SURFACE and performed a discrete Fourier transform of $z(x, y)$ using the FFTW library [119]. To do this within PLUMED the following lines need to be added to the input above:

```
FOURIER_TRANSFORM ... 1
LABEL=ft 2
GRID=contour 3
STRIDE=10 CLEAR=25 4
FT_TYPE=norm FOURIER_PARAMETERS=-1,1 5
... FOURIER_TRANSFORM 6
7
DUMPGRID GRID=ft FILE=fourier.dat STRIDE=10 8
```

The keyword FT_TYPE here tells PLUMED the kind of data that is required in the Fourier transformed function that is output. The possible values control what operations are performed after the discrete Fourier transform is computed: "abs" tells PLUMED that only the complex moduli of the Fourier coefficients are required, while "norm" tells PLUMED that the norm of the complex modulus is required. The full, complex Fourier coefficients are only returned from FFTW to PLUMED when the keyword FT_TYPE is set to "complex" or left out. Furthermore, when this is done the user must think of these complex numbers on a grid as if they are vectors.

An additional keyword is implemented in the FOURIER_TRANSFORM action to control how the output should be normalized. The keyword FOURIER_PARAMETERS takes two numerical values (a, b) that define the normalization and the type of the Fourier transform, forward or backward, that is to be applied to the data. The role these parameters play in the operation can be understood by considering the following one-dimensional example, in which an array X of size N is transformed:

$$Y_k = \frac{1}{N^{(1-a)/2}} \sum_{j=0}^{N-1} X_j \exp(2\pi i bjk/N).$$

Notice once again that the output from the FOURIER_TRANSFORM command is a function evaluated at points on a grid and that within PLUMED the output from this action can thus be treated as such. Further note that in the example above the CLEAR parameter in the line defining the Fourier transform is set much larger than the STRIDE parameter. This command is thus calculating an ensemble average for the Fourier transform. In other words, the location of the dividing surface is updated based on the positions of the atoms in each of the trajectory frames and Fourier transformed. The above command thus calculates the ensemble average of this Fourier transform over multiple trajectory frames.

Notes

1. A well-known example is that proposed by Cahn and Hilliard [120].

6 Solid–liquid properties of binary alloys from atomistic simulations

The aim of this chapter is to present and discuss the results obtained from the study of solid–liquid properties of a binary alloy, specifically the silver–copper binary system. The goal of this work was to establish a workflow employing well–known atomistic simulation techniques with a twofold purpose. Firstly, to investigate relevant properties of solid–liquid systems in equilibrium, and secondly to provide a way to get reliable estimates of those properties which could be particularly useful in those contexts where modeling of solidification takes place at time and length scales other than the atomic scale.

The first section covers an additional theoretical background: it first illustrates how one can exploit the umbrella sampling technique (section 3.4) to determine the phase diagram of a unary system. It follows a description of what changes when studying a binary system, and how the method can be used in conjunction to the CFM (section 4.3.2) to extract interfacial anisotropies and free energies of such systems. Lastly, it presents a way of extracting the information on interface dynamics from an over-damped Langevin model of interface fluctuations and how this model is related to the central result of CFM.

The second section explains how MD and MC techniques have been used to study solid–liquid interfaces of the silver–copper system, and illustrates the details of the simulations carried out both to determine the equilibrium thermodynamic properties – whose results are then presented in the third section (section 6.3) of this chapter – and to study the fluctuation spectra from which to obtain information on interfacial stiffness and free energy (section 6.4) and, subsequently, mobility.

6.1 Theoretical background

6.1.1 The interface pinning method

Studying phase equilibria inevitably requires the information of the free energy as a function of temperature and, for multi–component systems, composition. As discussed in detail in the chapter on thermodynamics (section 4.2), this particular information is provided by the phase diagram, which gives an overview of the thermodynamic equilibrium conditions of a multi–component system. In other words, determining the bulk thermodynamics of the phases implies knowing their coexistence temperatures and compositions.

In particular, if the interest is the study of interface properties, methods such as the capillary fluctuation method (section 4.3.2) require simulations of planar interfaces in equilibrium conditions for a time long enough to estimate accurate ensemble averages of the interface’s fluctuation spectrum. The umbrella sampling technique explained previously (section 3.4) offers a possible solution to both the determination of bulk thermodynamics and to stabilize simulations that are both close and far to coexistence conditions.

The application of umbrella sampling to the specific problem of a flat interface separating two coexisting phases is known as *interface pinning* (IP). Pedersen et al. [121] devised a procedure to estimate free energy differences between two phases by determining the external force to apply to the system in order to keep the interface in place. This external force is linked to a harmonic bias potential built upon some order parameter apt to distinguish clearly the two phases. If the system is brought in out-of-equilibrium conditions, the bias will prevent any interface’s displacement due to, say, the solid phase growing at the expense of the liquid one, and remaps the original problem into a (quasi) equilibrium scenario.

For example, one may choose an order parameter that allows a clear-cut distinction between solid and liquid, and define consistently the fraction of solid phase f in the system. In the presence of the external potential, the total Gibbs free energy of the system will be

$$G(f)/N = \mu_s f + (1 - f)\mu_l + \frac{\kappa}{2}(f - f_0)^2 + S\gamma_{sl}/N, \quad (6.1)$$

where μ_s and μ_l are the chemical potentials of the solid and the liquid phase, and γ_{sl} is the interfacial free energy. N and S are the number of atoms and the interface area, respectively.

Pedersen showed [122] how one can determine the thermodynamic drive to solidification (or melting) – that is, the difference $(\mu_s - \mu_l)$ – by simply running different simulations at different temperatures equilibrated at zero pressure along the direction perpendicular to the interface. At equilibrium, the following relation always holds for a biased system

$$\Delta\mu = -\kappa\langle f - f_0 \rangle. \quad (6.2)$$

Moreover, one could employ the IP method to sample the full phase diagram of a one–component system: given that the coexistence points are those for which, at fixed pressure and temperature, $\Delta\mu(p, T) = 0$, the only information needed are the derivative of the chemical potential with respect to pressure and temperature, that is

$$\begin{aligned} \left(\frac{\partial \Delta\mu}{\partial p} \right)_T &= \Delta V_m \\ \left(\frac{\partial \Delta\mu}{\partial T} \right)_p &= -\Delta S_m, \end{aligned}$$

where ΔS_m and ΔV_m are the differences in molar entropy and volume, respectively. Since $\Delta\mu$ is obtained through eq. (6.2), the molar entropy difference is simply $(\Delta U_m + p\Delta V_m - \Delta\mu)/T$, where ΔU_m is the molar internal energy.

It has been shown [123] that since the bias of eq. (6.1) only acts on the average value of

the collective variable chosen to represent the bulk phases, it only affects the $\mathbf{k} = 0$ mode of interface fluctuations, and the IP method could also be used to stabilize planar interfaces in weakly out-of-equilibrium conditions from which, employing the already described CFM, one could extract the temperature dependence of the interfacial free energy.

The discussion this far has assumed that the solid and liquid phases in coexistence had a single component. Most materials of technological significance include multiple atomic or molecular components, and the chemical potential of the different phases varies depending on the composition.

Let us consider the case of a binary mixture of two constituents A and B . While many simulations of multi-component systems have been performed at constant chemical potential, here we focus on simulations performed at fixed total composition, say $x = x_A$, since they can be run without knowing the equilibrium composition of the solid and liquid phase beforehand. They can also be implemented very easily by combining molecular dynamics with a Monte Carlo algorithm in which the position of two atoms of different species is exchanged, and the move is accepted or rejected based on a Metropolis criterion.

The constant total composition constraint affects the free-energetic description of the model, resulting in a trivial, yet somewhat cumbersome, extension to the textbook treatment of the coexistence region between the solid and the liquid phases of a binary alloy. Assuming one knows the dependence on the composition of the chemical potentials of the solid and the liquid phase, $\mu_s(x_s)$ and $\mu_l(x_l)$, the Gibbs free energy function to minimize is

$$G(x_s, x_l, f) = f\mu_s(x_s) + (1 - f)\mu_l(x_l) + \frac{\kappa}{2}(f - f_0)^2,$$

subject to the composition constraint given by the well-known “lever rule” (section 4.2.3)

$$x - fx_s - (1 - f)x_l = 0.$$

Here x_s and x_l are the compositions of the solid and the liquid phase, respectively.

The derivation involves solving the constraint for x_s or, equivalently, x_l . The requirement that the composition of the system is at equilibrium, that is

$$\frac{\partial G(x_s(x_l, f), x_l, f)}{\partial x_l} = 0,$$

implies that $\partial\mu_s(x_s)/\partial x_s = \partial\mu_l(x_l)/\partial x_l$, and can be used together with the implicit function theorem to determine $\partial x_l/\partial f$ and $\partial^2 x_l/\partial f^2$, subject to this requirement. After some algebra, one finds that

$$\left(\frac{\partial G}{\partial f}\right)_{p,T,x} = (f - f_0)\kappa - (\mu_l(x_l) - \mu_s(x_s)) + (x_l - x_s)\mu'_s(x_s), \quad (6.3)$$

and

$$\left(\frac{\partial^2 G}{\partial f^2}\right)_{p,T,x} = \kappa + \frac{(x_s - x_l)^3 \mu''_l(x_l) \mu''_s(x_s)}{(x_s - x)\mu''_s(x_s) + (x - x_l)\mu''_l(x_l)}, \quad (6.4)$$

where x_s and x_l depend implicitly on f .

One can then see that if $f = f_0$, eq. (6.3) corresponds to the common tangent construction, and the compositions of the solid and the liquid phase make it possible to determine the liquidus and the solidus at the simulation temperature. By varying f_0 one can also probe the derivatives of the chemical potentials of the two phases away from the equilibrium composition.

If the composition of the liquid and the solid phase differ, then eq. (6.4) predicts that the curvature of the free energy is not solely given by the force constant associated to the umbrella potential, but contains a correction that depends on the compositions of the two coexisting phases and on the curvature of μ_s and μ_l . In practice, one can just run a pinned-interface calculation and estimate the effective umbrella potential from the distribution of interface positions.

6.1.2 Capillary fluctuations of a Gibbs dividing surface

The key ingredient needed by the CFM (section 4.3.2) to extract interface energies and their anisotropy is in essence the information of the interface position as a function of time, $h(x, y, t)$. As we have seen, defining the position of the interface is a problem that does not have a single solution. The approach in terms of excess quantities, albeit elegantly providing a simple physical interpretation, will result in as many $h(x, y, t)$ functions as the number of possible extensive variables taken as reference. In other words, there is no right choice for the extensive variable, although some will be more suitable than others depending on the problem. However, it should be kept in mind that the final results for γ_{sl} will only be consistent if the same reference quantity has been used throughout.

A method to determine directly $\langle |A(\mathbf{k})|^2 \rangle$ which is fully consistent with a Gibbs dividing surface framework has been proposed by Cheng and Ceriotti [123]. It begins with defining a continuous field based on the order parameter as a density estimation with a suitable (normalized) kernel K , e.g. a three-dimensional Gaussian. The density field has the form

$$\varphi(x, y, z) = \sum_i^N \phi_i K(x - x_i, y - y_i, z - z_i),$$

where (x_i, y_i, z_i) and ϕ_i indicate the position and the local order parameter for atom i . Imposing the zero-excess condition for $\Phi = \sum_i \phi_i$ that defines the Gibbs dividing surface means writing Φ as

$$\Phi = \int_0^{L_x} dx \int_0^{L_y} dy \int_0^{L_z} \varphi(x, y, z) dz = \int_0^{h(x,y)} \varphi_s(x, y, z) dz + \int_{h(x,y)}^{L_z} \varphi_l(x, y, z) dz.$$

Here L_z is the extent of the simulation cell in the z -direction and $h(x, y)$ is the position of the dividing surface. φ_s and φ_l are the density fields of reference calculations that only contain the solid and liquid phase, respectively.

It is then possible to write a simple expression for the average fluctuation amplitude that does not require an explicit determination of the position of the dividing surface between the

solid and liquid phase¹

$$\begin{aligned}
 & (\bar{\varphi}_s - \bar{\varphi}_l)^2 \langle |A(\mathbf{k})|^2 \rangle + \\
 & + \langle |\Phi_s(\bar{h}; \mathbf{k})|^2 \rangle + \langle |\Phi_l(L_z - \bar{h}; \mathbf{k})|^2 \rangle = \\
 & \left\langle \left[\frac{1}{S} \sum_i^N \phi_i \exp[-i(k_x x_i + k_y y_i)] \right]^2 \right\rangle,
 \end{aligned} \tag{6.5}$$

where $\langle \cdot \rangle$ denotes the usual ensemble average. It is important to stress here that the validity eq. (6.5) is assured only if one can assume that:

1. bulk fluctuations are independent from each other;
2. bulk fluctuations correlate only on a length scale λ much smaller than the cell's extent in the z -direction;
3. one takes the δ -distribution limit of the kernel function on the atoms.

The factors $\bar{\varphi}_s = \bar{\phi}_s/v_s$ and $\bar{\varphi}_l = \bar{\phi}_l/v_l$ stand for the average value of the bulk fields (v_s and v_l are the molar volumes of the solid and the liquid), and \bar{h} indicates the average thickness of the solid slab, that can be obtained in a consistent way from the Gibbs estimate of the solid fraction. The thickness of the solid slab can be obtained consistently within the Gibbs dividing surface framework as follows: from eq. (5.1) one defines the fraction of solid as

$$f = \frac{\Phi/N - \bar{\phi}_s}{\bar{\phi}_s - \bar{\phi}_l}. \tag{6.6}$$

By the same argument, one may define

$$\Phi = \frac{S}{v_s} h \bar{\phi}_s + \frac{S}{v_l} (L_z - h) \bar{\phi}_l, \tag{6.7}$$

where S is the interface area and v_s (v_l) is the molar volume of the solid (liquid). From such relation one can easily determine the thickness of the solid slab, h . Provided that Φ has been rescaled and normalized such that $\bar{\phi}_s = 1$ and $\bar{\phi}_l = 0$, then by comparing eqs. (6.6) and (6.7) one obtains

$$h = \frac{Nv_s}{S} \left[f (\bar{\phi}_s - \bar{\phi}_l) + \bar{\phi}_l \right] = \frac{Nv_s}{S} f. \tag{6.8}$$

The terms $\Phi_s(h; \mathbf{k})$ and $\Phi_l(h; \mathbf{k})$ that appear in eq. (6.5) indicate the Fourier coefficients of a bulk section of thickness h , and can be obtained – under the assumption made above of short-ranged correlations along z – from a separate simulation box of thickness λ , using the scaling relations

$$\Phi_\star(x; \mathbf{k}) = \sqrt{\frac{x}{\lambda}} \Phi_\star(\lambda; \mathbf{k}),$$

where term $\Phi_\star(\lambda; \mathbf{k})$ is evaluated identically to that on the right-hand side of eq. (6.5), with the correct normalization factors already taken into account.

Given a coexistence simulation with sufficiently large L_z in which the solid fraction does not vary much, it is always possible to define clear-cut solid and liquid regions, from which all of the bulk quantities can be estimated.

Note that eq. (6.5) refers to a single dividing surface: in an MD simulation where the cell has periodic boundary conditions, the system always contains two interface boundaries, and accounting for this is as simple as multiplying by two the first term of eq. (6.5).

The left-hand side of eq. (6.5) clearly expresses a quantity evaluated for a coexistence simulation as a sum of bulk terms and an interface excess, that describes the behavior of the interface in a way that is conceptually and operationally consistent with a Gibbs dividing surface construction.

6.1.3 A Langevin model for interface kinetic coefficient

Instead of monitoring the position of a fluctuating interface, one may exploit the central result of CFM (section 4.3.2) and extract the interface kinetic coefficient by analyzing the full fluctuation spectrum [99, 124]. This approach relies on fluctuation–dissipation relations and can be understood as resulting from the assumption that a degree of freedom $x(t)$ of the interface obeys an over-damped Langevin equation of the kind

$$\dot{x} = -\frac{D_x}{2k_B T} \frac{\partial G}{\partial x} + \sqrt{D_x} \xi \quad (6.9)$$

where ξ is the derivative of a Wiener process and D_x is the diffusion coefficient. When the system reaches a stationary state, the probability of observing a given value of x is proportional to $\exp[-G(x)/k_B T]$. A few additional results for Brownian motion are useful in this context:

1. if $F_x = -\frac{\partial G}{\partial x} = 0$, then x performs a random walk with $\partial \langle x^2 \rangle / \partial t = D_x$;
2. if $F_x = -\frac{\partial G}{\partial x}$ is constant, then $\partial \langle x(t) \rangle / \partial t = D_x F_x / 2k_B T$;
3. if $F_x = -\kappa_x x$, an harmonic restraint, the auto-correlation function of x has an exponential decay $\langle x(t)x(0) \rangle = e^{-t/\tau_x}$ with the auto-correlation time $\tau_x = 2k_B T / \kappa_x D_x$.

If x represents the interface position, in the absence of a pinning potential and in out-of-equilibrium conditions,

$$\frac{\partial G}{\partial x} = -F_x = \frac{S \Delta \mu}{v_s}.$$

Hence, the steady-state interface velocity is given by

$$\langle v \rangle = -\frac{S D_h \Delta \mu}{2k_B T v_s},$$

and the kinetic coefficient is

$$M = \frac{D_h S}{2k_B T v_s}. \quad (6.10)$$

This relation reflects the link between the interface velocity in the presence of a thermodynamic drive and the diffusion coefficient of the interface at coexistence conditions [125].

In the presence of a pinning potential, the interface fluctuates around the equilibrium position, and the dynamics is that of an over-damped harmonic oscillator, whose force constant is determined by eq. (6.1). The auto-correlation function of h is then described by a damped exponential with time constant

$$\tau_h = \frac{2k_B T}{D_h \kappa_h},$$

leading to

$$M = \frac{S}{\tau_h \kappa_h v_s}. \quad (6.11)$$

Fitting of the exponential decay of the correlation function (or, equivalently, fitting its Fourier transform to a Lorentzian) has been proposed as another way of estimating the interface mobility [126].

This approach based on the description of interface fluctuations as a Langevin dynamics applies as well within the Gibbs dividing surface framework described previously (section 4.3.1). It suffices considering the link between the solid fraction as defined in eq. (6.6) and the thickness of the solid slab, h , which is given by

$$h = \frac{N v_s}{S} f,$$

where v_s is the usual molar volume of the solid, N the total number of atoms and S the area of the interface. One may then apply eq. (6.11) which leads to the estimate of M in terms of the fluctuations of the solid fraction²

$$\begin{aligned} M &= \frac{N^2 v_s}{2S \tau_f \kappa_f} \\ &= \frac{v_s}{2S k_B T \tau_f} N^2 \langle (f - \bar{f})^2 \rangle. \end{aligned} \quad (6.12)$$

From eqs. (4.29) and (4.31) it is easy to see that the $\mathbf{k} = 0$ mode represents the average position of the dividing surface and provides the most direct route to access the interface mobility. However, it is also known [13, 124] that all the other capillary modes carry the footprint of interface dynamics. Capillary fluctuations are constrained by the energy cost associated to any deformation of the interface, that is, its stiffness. Therefore, the relaxation time of a mode with wave-vector $\mathbf{k} \neq 0$ is related to the amplitude of the fluctuation by a quantity that is proportional to the mobility according to

$$M = \frac{S \langle A(\mathbf{k})^2 \rangle}{v_s k_B T \tau(\mathbf{k})} = \frac{1}{v_s \tau(\mathbf{k}) \mathbf{k}^T \boldsymbol{\sigma} \mathbf{k}}, \quad (6.13)$$

where $\tau(\mathbf{k})$ is the time constant of the auto-correlation function of this particular mode. The second equality exploits the main result of CFM that links the average Fourier amplitude of a mode with the interface's stiffness, and it correctly takes into account the possibility that some interfaces have an anisotropic stiffness tensor.

Given that fluctuations of the order parameter in the bulk phases also contribute to the

auto-correlation functions of the interface fluctuations, one should use an expression analogous to eq. (6.5) to obtain the correlation function of the interface without the terms representing the bulk phases, that is

$$\begin{aligned}
 & (\bar{\varphi}_s - \bar{\varphi}_l)^2 \langle A(t, \mathbf{k}) A(0, \mathbf{k})^* \rangle + \\
 & + \langle \Phi_s(\bar{h}; t, \mathbf{k}) \Phi_s(\bar{h}; 0, \mathbf{k}) \rangle + \langle \Phi_l(L_z - \bar{h}; t, \mathbf{k}) \Phi_l(L_z - \bar{h}; 0, \mathbf{k}) \rangle = \\
 & \left\langle \frac{1}{S^2} \left[\sum_i^N \phi_i(t) \exp[-i(k_x x_i(t) + k_y y_i(t))] \right] \left[\sum_i^N \phi_i(0) \exp[-i(k_x x_i(0) + k_y y_i(0))] \right] \right\rangle.
 \end{aligned} \tag{6.14}$$

The use of a Gibbs construction to eliminate the spurious contributions from the bulk makes it possible to use a simpler model to determine the interface diffusivity, compared to those used in recent works [121, 123] that used a model with two coupled Langevin equations to describe explicitly the fluctuations of the bulk.

6.2 Systems and computational details

In order to study both static and dynamical properties of solid-liquid interfaces, we chose two systems: the silver-copper binary alloy and the same Lennard-Jones model employed in the previous chapter. We can give three reasons to justify the choice of the AgCu system. Firstly, the experimental phase diagram shows a significant difference in the compositions of the liquid and the solid phase at every temperature above the eutectic point. As explained in the previous section (see eq. (6.4)), a large difference between the compositions of the two phases gives rise to an appreciable contribution to the Gibbs free energy curvature due to the constant composition constraint. Secondly, from a computational point of view, the available interatomic potential developed by Williams, Mishin, and Hamilton [127] within the EAM framework showed a satisfactory agreement with experimental data. This agreement suggests reliability of the results on properties such as the interface free energy since the interatomic potential proved to be able to reproduce quite accurately some thermodynamic properties, although the latter are known to depend strongly on the kind of the interatomic potential employed. Finally, in a recent work Frolov and Mishin [128] studied the properties of the (110) solid–liquid interface of this binary system employing the same EAM potential, and that provided us a direct way to compare their results with ours.

The choice of Lennard-Jones potential as the second system was guided by the need to test both the CFM within the Gibbs dividing surface framework — in which, compared to the original implementation [123], we introduced the idea that the contributions of the bulk phases to the fluctuation spectrum can be obtained directly from a solid-liquid simulation — and the overdamped Langevin model we proposed that explains the interface dynamics, from which one can extract the kinetic coefficient of the interface. Lastly, in addition to the lower computational cost of a Lennard-Jones potential when compared to EAM, this system has been already extensively studied, and thus our results could be more easily compared with those found in scientific literature.

The simulation protocol to generate the solid-liquid supercells for both the Lennard-Jones and

silver-copper simulations is much alike the one described in a section of the previous chapter. Here we thus only highlight those additional steps required to prepare the simulations for the binary system.

We performed all MD and MC simulations with the LAMMPS software together with the PLUMED package. Plumed served both to apply the harmonic restraint to the coexistence simulations and to calculate the FCCUBIC order parameter, which is particularly suited for a crystalline solid with FCC symmetry.

The protocol to prepare the initial configuration of the supercells for the AgCu system was the following: the lattice parameter's dependence on the composition follows, with good approximation, Vegard's law, therefore we interpolated experimental data available [129] to define the lattice constant of a perfect FCC unit cell. The unit cell was then replicated along the x , y and z directions, paying attention to have enough number of replicas along z to avoid a solid slab to be sensitive to statistical fluctuations that, during an unbiased simulation, could melt or solidify the system completely. An initial, equilibration run in the NPT ensemble of about 10 thousand steps was performed at the desired temperature, using a Nosé–Hoover thermostat and barostat. Throughout the simulations, the timestep was set to 0.002 ps.

To generate a solid-liquid interface we proceeded as follows: all the atoms within the central third of the supercell were held fixed, while the remaining part of the cell was heated way above the melting temperature of both silver and copper. This part of the simulation run for about 50 thousand MD steps.

After the two outermost sections of the supercell melted completely, the constraint on the atoms in the solid phase was released, and two independent NVT runs of about 2 thousand steps equilibrated the solid and the liquid portions of the supercell. We then equilibrated the whole system at the target temperature using a stochastic thermostat instead of Nose–Hoover's, while a fully anisotropic barostat kept the pressure close to zero. During these setup and equilibration runs, a Monte Carlo algorithm was attempting ten chemical species exchanges every MD step. As explained in the following sections, the bulk properties of both the solid and liquid phase were calculated from these equilibration runs.

A similar protocol has also been used to prepare and run all the simulations needed for the capillary fluctuations analysis. The main differences were the number of unit cells along the x or y – and consequently the total number of atoms – and that the ensemble of the production runs was NP_zT instead of a fully anisotropic NPT. The simulation box's dimensions along x and y have been fixed to their average values calculated from a previous equilibration run of about 10 thousand MD steps in which the ensemble was still NPT where all the sides of the simulation cell were allowed to fluctuate with no constraints. As for the simulations to sample the phase diagram, the production runs employed a stochastic velocity rescaling thermostat.

6.3 Equilibrium thermodynamic properties

To illustrate how we calculated the equilibrium properties of the binary system, fig. 6.1 shows a snapshot during the production run of one of the AgCu simulations. In this particular simulation, the total fraction of silver was $x = 0.86$ and the temperature was 1150 K.




Figure 6.1 – The upper part of panel (a) shows a plot of the average order parameter as a function of the coordinate along z perpendicular to the interface. The average has been calculated with a Gaussian kernel density estimation. The colored areas show the bulk regions: solid (blue) and liquid (red). The dividing surfaces’ locations, which are found to be around -40 \AA and 40 \AA , are determined according to the Gibbs construction (see eq. (6.8)). The lower part of panel (a) shows a snapshot of the solid-liquid simulation from which the upper plot has been calculated. Panel (b) shows a histogram of the distribution of the order parameters colored according to the degree of crystallinity.

From a coexistence simulation such as that of fig. 6.1, we obtained an estimate of bulk properties (molar volumes, averaged order parameter, $\bar{\phi}_{s,l}$, concentration of the two species in each phase) as follows: for each frame of an equilibrated simulation in the NPT ensemble, we had to define the solid and the two liquid slabs. To this end, the very first step of the analysis has been calculating the coordinates of a center of mass based on the order parameter and re-centering the whole simulation box. Estimating the averages $\bar{\phi}_l$ and $\bar{\phi}_l$ has been done by counting the number of atoms within a thin slab of width dz lying in the solid and liquid regions and computing the average order parameters in these regions. The composition of the solid (liquid) phase can be estimated as the number density of atoms of a given species, computed in the same region of $\bar{\phi}_s$ ($\bar{\phi}_l$), divided by the density of the solid (liquid).

We then obtained an average of the order parameter which depended only on z by performing a Gaussian kernel density estimation along this direction. The result is the continuous profile shown in panel (a) of fig. 6.1. With the same procedure but considering only atoms of a given species, it is possible to estimate the concentration profile of that species across the interface. Figure 6.2 show the concentrations profiles averaged over the production simulations for both copper and silver at different temperatures. The origin of the z axis corresponds to the location of the dividing surface. Although periodic boundary conditions imply the presence of two interfaces, the profiles shown here refer to a single interface for clarity. The region of negative z corresponds to the solid phase, while the liquid phase is located in the region of positive z .

The diffuse nature of the dividing surface is clearly evident from these plots, and away from the interface, the plateaus of the concentrations profile indicate the bulk values of the compositions of the corresponding phases.




Figure 6.2 – Concentration profiles across the interface obtained as described in the text for (a) silver and (b) copper at three different temperatures. The zero of the z coordinate coincides with the location of the Gibbs dividing surface. In both figures, the solid phase lies in the $z < 0$ region, while the liquid phase lies in the $z > 0$ portion of the plot. The values of the concentrations where the profiles plateau off correspond to the equilibrium compositions of the bulk phases. From these profiles, it is evident that the solid solubility limits are both very similar at all temperatures and close to 0.1 and 0.9, suggesting that the corresponding solid solutions are primarily composed by copper and silver, respectively.

A quantity that can be estimated with numerical integration from concentration profiles such as those of fig. 6.2 is the surface excess – also known as adsorption coefficient – which expresses the absolute depletion or enrichment of one species at the interface. In particular, the

surface excess of silver can be defined as

$$2S \Gamma_{\text{Ag}} = N_{\text{Ag}} - N_{\text{Ag}}^l - N_{\text{Ag}}^s, \quad (6.15)$$

where S is the dividing surface area; an equivalent expression can be written for the absolute surface excess of copper. In our framework based on a Gibbs dividing surface construction, an estimate of the surface excess can be obtained from the constraint of constant total composition, that is

$$2S \Gamma_{\text{Ag}}/N = x - f x_s + (1 - f)x_l. \quad (6.16)$$

The above definition of surface excess depends on the location of the dividing surface since moving the dividing surface would change the compositions of the two phases. However, it is possible to define a surface excess of the solute relative to the solvent according to [115, 130]

$$\Gamma_1^{(2)} = \Gamma_1 - \Gamma_2 \frac{\rho_1^l - \rho_1^s}{\rho_2^l - \rho_2^s}, \quad (6.17)$$

where ρ_i^α stands for the atom density of the species i in the phase α . Relative adsorption coefficients are key quantities in Gibbs' thermodynamic description of planar interfaces, for they account for variations of γ_{sl} due to changes in chemical potential according to Gibbs' adsorption equation³

$$d\gamma = -S_{\text{ex}}^{(2)} dT - \Gamma_1^{(2)} d\mu_1, \quad (6.18)$$

where the labels 1 and 2 refer to the solute and the solvent⁴, respectively. $S_{\text{ex}}^{(2)}$ is the relative excess entropy which, at least for metals, is responsible of the largest contribution to changes in γ due to temperature [131, 132]. It is important to remind that, in a binary system at constant pressure, there is only one independent intensive variable in accordance to Gibbs' phase rule, and thus the two contributions of eq. (6.18) are mutually dependent. In order to apply eq. (6.18), the knowledge of the temperature dependence of Γ is a necessary ingredient, and it has been already investigated through Monte Carlo simulations for several model systems [106, 130, 133].

Once the average bulk properties were known, we ran production simulations for a total length of about 2 ns. During these simulations, a harmonic bias was applied to maintain the fraction of solid particles around its average value obtained from the unbiased simulations at the previous stage. As explained, these simulations were needed to compute the correct values of $\bar{\phi}_s$ and $\bar{\phi}_l$ which in turn yielded a consistent value of the solid fraction.

Determining the phase diagram of the system is only a matter of running multiple simulations within a given range of total compositions and at different temperatures. Since the silver-copper alloy is a simple eutectic system and the miscibility gap extends across almost the whole range of compositions, we used our simulation framework to reproduce the phase diagram above the eutectic temperature. Figure 6.4 reports a comparison between our simulations, the Grand–Canonical Monte Carlo simulations of Williams, Mishin, and Hamilton [127], and the experimental phase diagram [134, 135]. Since this binary system could be modeled with a regular solution model (section 4.2.4), a possible explanation of the deviation of the calculated phase diagram from the experimental one is suggested by Williams and coworkers: they reported




Figure 6.3 – Relative surface excesses a function of temperature. The blue curve shows the adsorption coefficient of copper relative to silver, thus it has been computed for total concentrations of silver $x_{\text{Ag}} > 0.5$ where it is reasonable to assume that copper is the solute. The orange curve instead shows the adsorption coefficient of silver relative to copper when $x_{\text{Ag}} < 0.5$.

a large discrepancy between the predicted and experimental enthalpies of mixing of liquid solutions at 1423 K. The largest value of this discrepancy was about 0.02 eV/atom, whose order of magnitude is comparable to that of the error between energies predicted with EAM and *ab initio* methods, suggesting that it could be due to an intrinsic limitation of the EAM framework.




Figure 6.4 – Phase diagram of the silver-copper binary alloy above the eutectic. The red circles are the equilibrium compositions of the solid and liquid phase as obtained using the simulation framework described in the text. The blue, dashed lines are the results obtained by Williams, Mishin, and Hamilton [127] with Gran-Canonical Monte Carlo (GC-MC) simulations. The black, dotted lines are the experimental phase diagram [134, 135]. The comparison shows that the EAM potential predicts satisfactorily the melting temperatures of the pure elements within about 30 K of the experimental values, while the estimated eutectic composition (0.458) and temperature (935 K) deviate more from the experimentally measured values of 0.601 and 1053 K.

6.4 Interface properties

In this section, we are going to discuss the results related to interface properties — that is, interface stiffness, free energy, and kinetic coefficient — of the binary system. To determine stiffness for the (100) and (110) interfaces we employed the CFM based on the Gibbs dividing surface construction described in the first section of this chapter, which avoids the need to determine explicitly the location of the dividing surface.

6.4.1 Stiffness and free energy

Figure 6.5 shows the plots of $k_B T/S / \langle |A(k_x, k_y)|^2 \rangle$ versus k along both k_x and k_y for pure copper and silver. The simulations were run at the corresponding melting temperatures predicted by the EAM potential. A fit of these plots with the model $\sigma(k_x^2 + k_y^2) = \sigma k^2$ provides the value of the interface stiffness. As it is evident from the plots, the (100) is an isotropic interface, and thus the stiffness is a scalar. An alternative way of estimating σ is to compute the average of

$$\frac{k_B T/S}{\langle |A(k_x, k_y)|^2 \rangle k^2}, \quad (6.19)$$

which corresponds again to the stiffness. If statistical errors on $|A(k_x, k_y)|^2$ have been estimated, a weighted average will correctly account for more significant errors in the Fourier amplitudes of k points close to zero. The error bars in the plots of fig. 6.5 have been computed from the corresponding simulation using the block averaging procedure explained previously (section 2.7.1). Since auto-correlation times of those k points closer to zero are larger than those of k points further away, the error bars associated with their Fourier amplitudes $|A(k_x, k_y)|^2$ are larger, as it is expected and has been discussed for similar results in the previous chapter (section 5.3).

Although it is known that different interatomic potentials can significantly influence stiffness values, we can compare our results with those of some previous works. According to the work by Hoyt and Asta [83], where the authors employed the potential developed by Voter and Chen [136], the stiffness of silver is about $(112 \pm 5) \text{ mJ/m}^2$, while the result we obtained is $(88.4 \pm 0.4) \text{ mJ/m}^2$. For copper, we obtained the results of $(110.7 \pm 0.5) \text{ mJ/m}^2$, which we may compare with $(105 \pm 14) \text{ mJ/m}^2$ [137] and 203.5 mJ/m^2 [138]. In the latter case, the authors utilized an empirical potential developed within the MEAM [26] framework and for which they modified the existing parameters for elemental copper. Even though they report a melting temperature of 1320 K, which is very close to that predicted by the EAM potential we adopted, our result is almost a factor two smaller than theirs, and a reason to explain such a discrepancy could be the comparison between two very different descriptions of the atomic interactions.

The same analysis has been applied to simulations at different temperatures and with a total fraction of silver between 0 and 1, in particular for values of $x_{\text{Ag}}(T)$ equal to 0.165 (1100 K), 0.805 (1100 K), and 0.458 (935 K). These temperature–composition points have been chosen by looking at the phase diagram of fig. 6.4 in such a way to have roughly the same amount of liquid and solid phase.

The procedure to analyze the simulations of the (110) interface differed from that for the




Figure 6.5 – Plot of the quantity $k_B T / S / \langle |A(k_x, k_y)|^2 \rangle$ versus \mathbf{k} for the (100) interface of (a) pure silver and (b) pure copper. The simulations were run at the melting temperatures predicted by the EAM potential employed. The blue curve is along k_x ($k_y = 0$), while the orange curve is along k_y ($k_x = 0$). The isotropy of the (100) interface can be seen from the two lines of each plot having roughly the same intercept with the y axes.

simulations of the (100) only by the fitting model employed. As explained previously (section 4.3.2), to capture the anisotropy of the stiffness tensor of interfaces such as the (110), there are two possible choices: the first is to run two independent one-dimensional simulations in which the longest direction in the interface plane is, in each simulation, parallel to one of two orthogonal directions to the [110]. To make a clarifying example, one would have to run the two independent simulations with the x , y , and z axes oriented along the $[1\bar{1}0]$, $[001]$, and $[110]$, with $L_x \gg L_y$ in one simulation and $L_y \gg L_x$ in the other one. Two fitting models similar to eq. (4.28) are thus employed, each of which will provide the value of the stiffness tensor of the (110) along the longest direction parallel to the interface – that is, either the $[1\bar{1}0]$ or the $[001]$.

Table 6.1 – Composition (and temperature) dependence of γ_{sl} (mJ/m²) of the (100), (110), and (111) interfaces as a function of the total fractional amount of silver, x_{Ag} . The results obtained for pure silver and pure copper are also shown. The reported errors have been estimated from the standard errors in the fitting parameters of the models $\sigma(k_x^2 + k_y^2)$, used for the (100), and $\sigma_{11}k_x^2 + \sigma_{22}k_y^2$, used for the (110).

x_{Ag}	$T(K)$	$\gamma_{(100)}$	$\gamma_{(110)}$	$\gamma_{(111)}$
0.0	1327 K	129.4 ± 0.7	127.0 ± 0.6	126.5 ± 0.6
0.165	1100 K	147.1 ± 3.3	146.5 ± 3.2	145.5 ± 3.2
0.458 ^a	935 K	88.6 ± 1.2	89.8 ± 1.2	89.3 ± 1.2
0.805	1100 K	99.5 ± 1.5	98.1 ± 1.5	97.8 ± 1.5
1.0	1267 K	104.5 ± 0.6	102.5 ± 0.6	101.9 ± 0.6

^a Simulation at the eutectic with a silver-rich solid phase ($x_s = 0.933$).

The second choice is a two-dimensional simulation, where both the axes lying in the interface plane are about the same length. In such configuration, we are able to extract the information on the interface fluctuation spectrum along both directions simultaneously and, provided that the directions parallel to the interface are orthogonal to each other, obtain both the non-zero components of the stiffness tensor by employing a fitting model similar to eq. (4.30). A manifest drawback of this choice is the larger number of particles the system will contain. Since evaluating the ensemble average $\langle |A(k_x, k_y)|^2 \rangle$ requires collecting a considerable number of “snapshots” of the fluctuation spectrum during the simulation, one has to consider the computational cost of a two-dimensional simulation and compare it to that of two smaller, one-dimensional simulations⁵.

Once the stiffness tensor for the (110) interface is known, solving the system of three equations given by eqs. (5.9) and (5.10) yields the set of anisotropy parameters, $(\gamma_0, \epsilon_1, \epsilon_2)$. The value of γ_{sl} for a given crystallographic orientation are then obtained by employing an expansion such as eq. (4.26). The above procedure is applied to the same temperature–composition points for which we already computed the stiffness of the (100); in this way it is possible to obtain an estimate of the dependence of free energy on the total fraction of silver and, therefore, on the temperature. The results for three of the most common interfaces when studying solids with an FCC symmetry are reported in table 6.1.

When studying solidification in alloys, and in particular dendritic solidification which is essentially driven by the anisotropy of γ_{sl} , an important parameter to consider is

$$\frac{\gamma_{(100)} - \gamma_{(110)}}{2\gamma_0}, \quad (6.20)$$

which provides a quantitative estimate of the magnitude of the anisotropy. For simple metals such as Cu, Ni [81], Ag [83], and Au [83], it is known that anisotropy is positive and approxi-




Figure 6.6 – Plot of the results of γ_{sl} (mJ/m²) of the silver-copper alloy as a function of the total fraction of silver, x_{Ag} . Both the values calculated for pure copper and pure silver are comparable with results found in the literature [83, 138]. The two sets of results at the eutectic composition refer to independent simulations in which the concentrations of the solid phases were set to the solubility limits at the eutectic; in particular, in the lower set of points x_s was equal to 0.933 (silver-rich solid-solution), whereas in the upper set $x_s = 0.03$ (copper-rich solid-solution). As discussed in the text, Gibbs adsorption coefficients provide us with only one of the ingredients required to justify the observed dependence of γ_{sl} on composition, yet the two sets of results at the eutectic composition suggest that this dependence is not an artifact of our simulations.

mately in the same range – between 1% and 2% – meaning that $\gamma_{(100)} > \gamma_{(110)} > \gamma_{(111)}$, which is also consistent with experimental evidences of the preferred growth directions in dendritic patterns during solidification. Figure 6.7 shows the estimated anisotropy as a function of composition. The values of pure copper (0.9 ± 0.4 %) and pure silver (1.0 ± 0.4 %) are also shown, with the latter being in excellent agreement with a previous work by Hoyt and Asta [83], where they reported a value of anisotropy magnitude of silver of 1.0 ± 0.3 %.

The statistical accuracy of our results presented in fig. 6.6 deserves a comment that relates to both the size of our simulations and the CFM consistent with Gibbs dividing surface construction. Equation (6.5), which states that one could determine the fluctuation spectrum of the interface $\langle |A(k_x, k_y)|^2 \rangle$ by first subtracting the contributions of the bulk phases, only holds if there are no correlations between the fluctuations in the bulk phases, and that the correlations in each bulk phase have a characteristic length scale that is much smaller than the simulation cell's size in the dimension perpendicular to the interface. Since our starting simulations were of one-dimensional cells, another assumption that we had to verify was the absence of finite-size effects that influenced the bulk contributions. While, according to capillary wave theory, $\langle |A(k_x, k_y)|^2 \rangle$ of the interface follows a $1/k^2$ trend at least for small wavevectors, it can be assumed that the Fourier spectra of the bulk phases should not show a substantial dependence on k , precisely because they do not represent the fluctuations of any interface. This hypothesis had already been verified in the work of Cheng and Ceriotti [123], whose results showed that the bulk contributions were almost constant over the whole range of k -points used for the

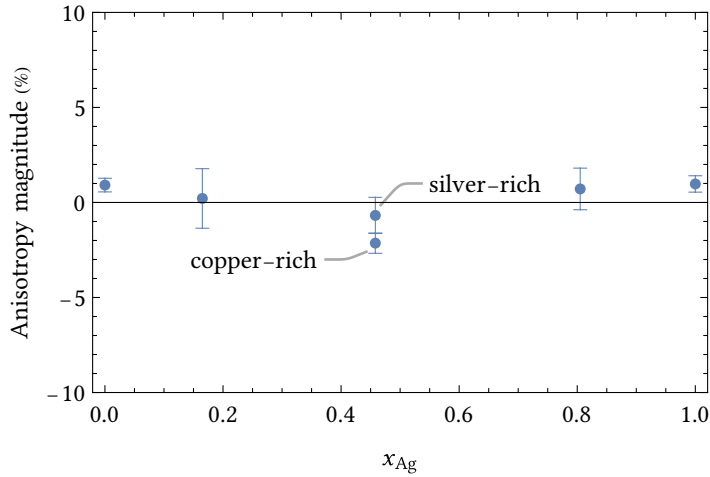


Figure 6.7 – Anisotropy magnitude as defined by eq. (6.20) as a function of the total fractional amount of silver, x_{Ag} . The results for pure copper and pure silver are also shown. The latter value is in good agreement with that predicted by Hoyt and Asta [83], in which they employed a different interatomic potential. Within their statistical uncertainty, these results suggest that for alloyed compositions γ_{sl} is almost isotropic with an anisotropy magnitude very close to zero. Analogously to fig. 6.6, the two points at the eutectic composition refer to two independent simulations where the solid phase was either a copper-rich or silver-rich solid-solution.

fitting of $\langle |A(k_x, k_y)|^2 \rangle$ from which the stiffness was extracted. However, our results suggested not only that there were apparent finite-size effects in the bulk phase contributions estimated from one-dimensional simulations, but also that their dependence on k could not be ignored entirely. When adding these two effects together, the estimate of stiffness we obtained did not agree with the results of other previous works. This led us to conclude that, in order to use this method that undoubtedly elegantly combines the Gibbs construction with the CFM, it is not possible to obtain reliable results by determining the contributions to the fluctuation spectrum of the bulk phases from a single simulation, and that it is thus necessary to perform independent simulations of sufficient size that do not exhibit any finite-size effect. A property like the interface stiffness should not in any way depend on the size of a system that should be macroscopic – that is, precisely, where all the spatial dimensions are periodic.

A last comment has to be made regarding the comparison of our results with those obtained by Frolov and Mishin [128]. Since the interatomic potential employed is identical, it is not surprising that we are able to predict the phase diagram and the concentration profiles in both the solid and the liquid with an almost perfect agreement. However, if we compare our results on γ_{sl} for the three interfaces considered (fig. 6.6), the difference is more striking and not so straightforward to explain. In particular, Frolov and Mishin found that γ_{sl} increases with decreasing temperature, while the opposite trend that our results show has recently been found only for the CuZr binary alloy [139] and it is considered an unusual behavior for this property. Even though Frolov and Mishin aimed at determining interface and related excess properties (e.g., adsorption coefficients) within Gibbs dividing surface framework and they try to devise a computational method that tries to solve the difficulties of Gibbs' approach to

multi-component alloys⁶, they adopted a thermodynamic integration scheme to determine γ_{sl} along a solid-liquid coexistence path. Besides the inherent issues of defining a Gibbs dividing surface when studying binary alloys, we calculated the temperature dependence of γ_{sl} with the CFM, therefore, without further investigations, we can only attribute to the difference in the computational methods used to determine γ_{sl} the origin of the significant discrepancy between the results.

The trend of the relative adsorption coefficient $\Gamma_{Cu}^{(Ag)}$ cannot fully explain the dependency of γ_{sl} on the composition. There are two reasons for this: firstly, according to eq. (6.18), γ_{sl} depends also on temperature, and we did not estimate the excess entropy term that captures this dependency. Secondly, if one wanted to apply eq. (6.18), for example to estimate or verify the temperature dependence of γ_{sl} , one would need to evaluate two integrals between a temperature T and a reference temperature that could be taken to be the melting (or coexistence) temperature, T_{ref} :

$$\gamma_{sl}(T) - \gamma_{sl}(T_{ref}) = - \int_{T_{ref}}^T S_{ex}^{(2)}(T') dT' - \int_{T_{ref}}^T \Gamma_1^{(2)} \frac{d\mu_1(T')}{dT'} dT'. \quad (6.21)$$

The previous relation makes clear that, even with the information on the (relative) excess entropy, one would still need (1) a dense sampling of $\Gamma_1^{(2)}(T)$ to be able to estimate the second term on the right-hand side of eq. (6.21), and (2) a valid approximation of $\mu_1(x_1(T))$ as the chemical potential depends on temperature only through the compositional variable x_1 – for this particular case, “1” refers to copper, so $x_1(T) = 1 - x_{Ag}(T)$. For example, if the solution thermodynamics appears to be close to ideal, one could model $\mu_1(T)$ as

$$\mu_1(T) = \mu_1^* + k_B T \ln [x_1(T)],$$

where μ_1^* is the molar Gibbs free energy of the pure component 1, either in the solid or in the liquid phase since at equilibrium the two must be equal. In this case, since the silver-copper phase diagram is far from ideal – yet it can be represented within the regular solution model – we cannot assume such a simple form for the dependence of the chemical potential on the composition.

6.4.2 Kinetic coefficient

To be able to model solidification microstructure it is necessary to have information on the solid–liquid interface dynamics, that is, its kinetic coefficient, or mobility, M . Mobility is usually defined as the proportionality constant between the interfacial average velocity and the undercooling, but an equally valid definition involves the ratio between interface velocity and the difference in chemical potential. The latter is particularly suited for multi–component systems where a change of one phase’s composition implies a change in chemical potential.

In the theory chapter (section 4.3.2) we outlined some of the already established computational methods to determine M , while in a previous section of this chapter (section 6.1.3) we described how from a Langevin model for interface fluctuations one can extract the information on interface mobility. In particular, eq. (6.12) predicts that the kinetic coefficient, defined as $\langle V_n \rangle = M \Delta \mu$, where V_n is the velocity normal to the interface boundary, can be obtained from

the fluctuations and the relaxation time of the auto–correlation function of the solid fraction, f .




Figure 6.8 – Kinetic coefficient M ($\text{\AA ps}^{-1} \text{ eV}^{-1}$) of the (100) interface estimated from eq. (6.12) as a function of the total fractional amount of silver, x_{Ag} . The values for the pure elements are of the right order of magnitudes we expect when comparing them to similar results found in the literature [83, 102].

Figure 6.8 shows the obtained kinetic coefficients for the (100) interface as a function of the total composition. As with the stiffness, we first verified how our results compared with those of other works in the literature. In particular, we compared the value of the kinetic coefficient of pure silver with the results of a work by Hoyt and Asta [83]. In that work, the authors calculated the kinetic coefficient for the (100), (110), and (111) interfaces by monitoring the velocity of the solid-liquid boundary as a function of the undercooling; a linear fit of such a curve yields the kinetic coefficient. Our definition of M is in terms of the the difference in chemical potential, thus to be able to compare directly the results we had to estimate the entropy of fusion of silver. Since at melting the Gibbs free energies of the solid and the liquid must be equal, we may estimate the molar entropy of fusion as

$$\Delta_{\text{fus}}S_m = \frac{\Delta_{\text{fus}}H_m}{T_{\text{fus}}},$$

where $\Delta_{\text{fus}}H_m$ is the molar latent heat of melting. This number has been calculated by Williams et al. [127], and its value of 12.45 kJ/mol is very close to the experimental value of 11.30 kJ/mol.

Our results for the kinetic coefficient of the (100) interface are shown in fig. 6.8 and were obtained by using eq. (6.12) that links the kinetic coefficient to the fluctuations of the solid fraction. We could verify that our results for silver ($48 \text{ cm s}^{-1} \text{ K}^{-1}$) and copper ($61 \text{ cm s}^{-1} \text{ K}^{-1}$) agree reasonably well with those obtained by Hoyt and Asta [83] and Hoyt et al. [102]. However, there are still some doubts on the values calculated at intermediate compositions which are found to be between 2 and 3 times smaller. To investigate this discrepancy further and understand if those results indicate either an error in the method or a problem in our simulations, we verified that we could actually describe the decay of the fluctuations of the solid fraction with a

simple exponential model. Figure 6.9 shows the plots of the auto-correlation functions of the fluctuations of the solid fraction for simulations with different compositions, including those of pure copper and silver. Particularly for the simulations at $x_{\text{Ag}} = 0.165$ and $x_{\text{Ag}} = 0.805$, it is rather evident that the single-exponential model is not able to correctly describe the behavior of the fluctuations of f .




Figure 6.9 – Plot of the auto-correlation functions of f for the simulations both at intermediate compositions and for the pure elements.

A look at the fig. 6.10, which shows the value of f monitored after an initial transient during the three simulations at intermediate compositions, suggests that the problem is related to the solid-liquid interface that does not fluctuate around its equilibrium position, and this is more prominent for those two simulations in which the auto-correlation functions of f did not follow a single exponential behavior. Therefore, determining the relaxation time τ_f assuming that these fluctuations follow a simple exponential is not correct, and the need to use a fitting model that is the sum of two decaying exponentials with two different relaxation times indicates the presence of another phenomenon that significantly affects the dynamics of the interface. A factor that undoubtedly differentiates the simulations of the binary alloy from those of the pure elements is that, within the framework of pinned interface simulations, the constraint of total composition introduces an additional term that adds up to the external harmonic bias and explicitly depends on the difference between the compositions of the two coexisting phases (see eq. (6.4)). The sum of these two terms corresponds to the total pinning potential felt by the system. A possible cause that may explain why the interface is not fluctuating around its equilibrium position could be due to a different strength in the total pinning potential. The phase diagram (fig. 6.4) shows that the difference in composition between the solid and liquid phase at the eutectic concentration is about 0.4, while for a total concentration of silver equal to 0.805 the difference is about 0.2. Since a term $(x_s - x_l)^3$ appears in eq. (6.4), the effect of the interface pinning potential is much more pronounced for the eutectic, thus preventing the interface from drifting too much off its equilibrium position during the simulation.

The second problem we had to face was to estimate the kinetic coefficient from equilibrium




Figure 6.10 – Solid fraction as monitored during the simulations at the total compositions of 0.165, 0.805, and 0.458 (eutectic). The drift in the solid fraction, in particular for the simulation with $x_{Ag} = 0.805$, makes it clear why it was not possible to fit the corresponding auto-correlation function with a simple exponential model.

fluctuations of different Fourier modes and their relaxation times. Always assuming that the dynamic of the interface follows an over-damped Langevin model, eq. (6.13) predicts that the relationship between the relaxation times of different Fourier modes and the wavevector k is of the kind $1/Mk^2$. Bearing in mind that the analysis of our simulations through the eq. (6.12) had given us results in agreement with others in the literature, we expected to be able to directly compare the values of M obtained through the eqs. (6.12) and (6.13). Unfortunately, we found an inconsistency between the values predicted by the two relations, and we are not yet in a position to justify that discrepancy objectively. However, it is appropriate to briefly discuss two sensitive issues related to the use of eq. (6.13) to estimate M . A first difficulty is the need to have extremely long simulations during which the trajectory must collect snapshots of the system with a very high frequency. In the high k limit, the relaxation times get smaller and smaller, eventually approaching the time-scale of bulk fluctuations and moving further and further away from the $1/k^2$ trend.

The second issue is related to the reliability of the results in the regime of small k . Although Monk et al. [124] have shown that the auto-correlation times of these fluctuation modes follow the expected trend of $\tau_k \propto k^{-2}$, the values of the kinetic coefficient systematically underestimate those obtained through other techniques, for example via non-equilibrium free solidification simulations. This particular problem is related to thermal conductivity which, in the case of an EAM potential, is considerably underestimated, when instead it is of fundamental importance. An incorrect thermal conductivity implies that the latent heat released during solidification is not dissipated quickly enough and ends up influencing the decay behavior of the interface fluctuations. This phenomenon had already been studied some time ago, and the proposed solution [13] modifies the proportionality relationship between M and k into $\tau_k \propto k^{-3}$, in which the proportionality constant explicitly contains the thermal diffusivity. The two different power-law behaviors of M versus k explain why it should be still justified to use the $1/k^2$

dependence of τ_k when performing MD simulations of appropriate size, even though the small wavevectors will always be harder to converge and have thus larger uncertainties, preventing one from excluding even a small contribution due to the aforementioned heat flow issue.

Notes

1. The detailed derivation of eq. (6.5) can be found in the Supporting Information of the work by Cheng and Ceriotti [123].
2. Assuming $\tau_h = \tau_f$ – a reasonable assumption since fluctuations of both h and f describe essentially the same phenomenon – we know that

$$\kappa_f = \frac{k_B T}{\langle (f - \bar{f})^2 \rangle}.$$

Using the relation between f and h one may write

$$\frac{1}{2} \kappa_f (f - f_0)^2 = \frac{1}{2} \kappa_f \left(\frac{S}{Nv_s} \right)^2 (h - h_0)^2,$$

where $\kappa_h = \kappa_f (S/Nv_s)^2$. Substituting the expressions for κ_h and κ_f into eq. (6.11) yields the result of eq. (6.12). The additional factor 1/2 comes from the presence of two interfaces fluctuating independently.

3. See, for example, chapter 3 of Chatteraj and Birdi [140].
4. There is no univocal definition of “solute” and “solvent.” For example, in a binary system, a species that has a total fractional composition less than a certain threshold could be considered the solute, yet some ambiguity remains. It is then more appropriate to study the temperature dependence of both the relative adsorption coefficients.
5. For an elemental system such as a pure metal, the choice depends almost entirely on the computational cost of the interatomic potential used. With a binary or even a multi-component system, one has not to forget that the additional Monte Carlo steps are another factor that increases the per-step simulation cost. Although we are studying a binary system, our choice was to run two-dimensional simulations, given that an EAM potential can still be considered a cheap potential if compared to more sophisticated interatomic potentials such as those based of machine learning, which aim to be almost as accurate but much faster than an *ab initio* approach.
6. A thorough analysis of this problem and a possible solution to it have been discussed by Cammarata [141].

7 Conclusions

Solidification is a phase transition of primary importance in materials science. In several industrial processes, solidification is directly responsible for the microstructure of the final material, on which a wide range of properties of most modern engineered materials depends. An example mentioned previously is that of dendritic solidification, which is particularly predominant in the determination of the microstructure of metallic materials – which are in the vast majority of actual cases with some technological significance multi-component alloys.

Solidification, however, is a very complex problem because it necessarily includes a series of phenomena that occur on very different time scales and lengths; therefore, the experimental investigations are accompanied by a series of modeling techniques that must be able to describe this phenomenon in its entirety accurately. A fruitful strategy is the one that manages to combine the results of different approaches; in this way, the information obtained through one of these techniques can guide another technique in the right direction, with the final objective of providing more and more precise indications to the actual industrial production phase in order to obtain a final product with characteristics suitable for its application.

This thesis has had as the main instrument of investigation various simulation techniques on the atomic scale. The objective was to study some fundamental aspects of solidification, with a particular interest in binary metal alloys. Although real applications rarely make use of binary alloys, these remain useful prototypes to be studied, in particular when part of the research interest concerns the development of new computational techniques or the refinement of others already well established that need to investigate in more detail an intricate phenomenon such as solidification.

Solidification is a phenomenon whose study has already been addressed in-depth and for a long time, both from an experimental and theoretical point of view. For this reason, a significant part of this thesis was dedicated to reviewing the “toolbox” of a computational material scientist: the bases of statistical mechanics that allow atomistic simulations to provide relevant results and the salient aspects – those most connected to the results presented later on – of the computational methods used. A substantial portion of this overview illustrated those essential concepts of thermodynamics that explain, even with some approximations, some fundamental aspects of solidification, particularly in binary systems and how to deal with the study of interfaces and their properties at the atomic scale.

Chapter 5, “Anisotropy and free energy of a smooth fluctuating dividing surface,” discusses

the development — and the implementation in the open-source package PLUMED [112] — of a framework for finding the location of the interface between two phases that differ either because they have a different molar volume or a different local symmetry. This framework extends the Willard–Chandler construction [111] for the liquid-vapor fluctuating interface and takes the positions of all the atoms in the system as input. It returns a two-dimensional profile that gives the height of the interface as a function of the two coordinates in the plane perpendicular to the surface normal. We showed how we could examine local fluctuations in the height of this interface using the tools of capillary wave theory and thus extract information on surface properties such as interfacial stiffness and the interfacial free energy.

We have shown that defining the height profile of the interface using a smooth density field offers significant advantages in terms of computational efficiency when it comes to determining interface free-energies and anisotropies using the capillary-fluctuation method. In particular, the form of the field suggests that a Gaussian convolution ansatz should be used to describe the deviation of $\langle |A(k_x, k_y)|^2 \rangle$ from the asymptotic $1/k^2$ behavior that occurs at short wavelengths. This ansatz makes it possible to extract meaningful information from larger values of the wave-vector. Consequently, a reasonable number of usable Fourier components can be extracted from a relatively small simulation box, thus lowering the computational cost of the CFM method. What is more, since shorter wavelength fluctuations have faster auto-correlation times, it is also possible to use shorter overall simulation times.

The fact that we can calculate stiffness values from shorter simulations containing fewer atoms combined with the availability of a robust implementation of this technique in a widely-used simulation package ensures that we are now in a position to apply these techniques when investigating more complicated and more realistic interatomic potentials. Furthermore, the fact that we can investigate the fluctuating dividing surface between stable — and even metastable — phases, using an arbitrary order parameter that is capable of distinguishing between their local atomic environments, opens up new opportunities for studying the structure and thermodynamics of interfaces.

There is considerable interest in improving the understanding of the phenomenon of solidification; in the specific context this thesis' work, the interest has focused mainly on metal alloys for applications in the field of additive manufacturing, in particular the silver-copper binary alloy. The use of atomistic simulations to investigate the problems related to the solidification of alloys like this, as well as a wide range of pure metals, has already been the subject of numerous research works, often providing results in accordance with experimental data or more accurate predictions of quantities such as interface free energy or its anisotropy, which are quantities notoriously very complicated to estimate experimentally with a reliable error.

The idea that initially motivated our work was to combine some of these standard modeling techniques and thus develop a framework that would allow to study properties related to both the thermodynamics of equilibrium (phase diagram and surface adsorption coefficients) and the interfaces between solid and liquid (interface free energy and kinetic coefficient). In order to minimize the impact of the rather arbitrary choices that one has to make in analyzing atomistic simulations (order parameters, thresholds, etc...) we decided to use a Gibbs construction of the interface as the foundation for our analyses. To study thermodynamic properties at equilibrium,

once we defined a suitable order parameter to distinguish a solid from a liquid phase, the Gibbs framework allowed us to calculate the concentrations of the two chemical species in the coexisting phases in a range of temperatures and total compositions, expressed as the amount of silver. From these results, one can obtain two extremely important pieces of information: the phase diagram (fig. 6.4) and the surface excesses (fig. 6.3), the latter being a necessary ingredient to describe how the interfacial free energy varies with the composition of the phases present, whose equilibrium intervals are given precisely by the phase diagram. Determining these equilibrium properties would not be possible using only a technique like molecular dynamics, as the time scale of diffusions far exceeds that which can be simulated. Therefore, our simulations have combined molecular dynamics and Monte Carlo and allowed studying these properties on systems of modest size (only several thousands of atoms) and for simulation's times easily accessible with modern high-performance computing facilities. In principle, the combination of these techniques is not limited to the study of binary alloys, but it could also be extended to ternary systems that are much more interesting for practical applications.

As far as interface properties are concerned, a revision of the method known as capillary fluctuation method (CFM) has allowed us to determine the interface free energy in an entirely consistent way with the Gibbs construction, avoiding the problem – which was discussed extensively in the first chapter of the results (chapter 5) – of finding a reliable operational definition of the dividing surface. In the simulations we conducted to determine the interface free energy and its anisotropy of the silver-copper binary system, we combined a method known as interface pinning [121, 122] that belongs to enhanced sampling techniques. Since we wanted to study planar interfaces, this technique was essential to be able to capture the dynamics of an interface that remained stable for a sufficiently long time to determine its interface free energy through the CFM. Moreover, since we have seen that the interface statics is not sufficient to describe the phenomenon of solidification exhaustively, we proposed a model based on Langevin's equations of damped dynamical phenomena to describe the fluctuations of the interface between solid and liquid and from which to extract information on the dynamics of the interface. This last approach is not new, but we have reformulated it in the context of Gibbs' dividing surface construction.

Stiffness, free energy and anisotropy Although both the static and dynamic properties of the interface can be considered in the context of quasi-equilibrium conditions, our results have revealed several problems that have forced us to revise much of the original idea of this unified approach.

We have already pointed out how, in the framework of the CFM within Gibbs' dividing surface construction, it is not necessary to explicitly define a method to locate the position of the interface because, before determining the fluctuation spectrum of the interface, the contributions of the two bulk phases present are appropriately subtracted from the full fluctuation spectrum (see eq. (6.5)). We have also discussed how the use of CFM does not require two-dimensional simulations: one-dimensional simulations can provide the same information about the fluctuation spectrum, and the issue is thus only related to the computational cost of the simulations to run. Although our simulations used a semi-empirical interatomic potential developed within the

EAM — therefore a relatively inexpensive potential — the initial setups of our simulations were of one-dimensional cells, one for each direction along which we wanted to extract the stiffness tensor. Furthermore, our initial idea suggested the possibility of estimating the contributions of the liquid and solid bulk phases from the simulation of the interface itself, thus reducing the number of independent simulations for each interface and greatly simplifying the workflow to obtain the results from these analyses. However, we found out that one-dimensional simulations were introducing unexpected finite-size effects that eventually led to an underestimation of the stiffness. Therefore, in order to have simulations whose results could be compared with those in the literature, we had to perform simulations with two-dimensional cells and thus hundreds of thousands of atoms. The first consequence of this choice is the real-time cost of such simulations. In order to obtain results with a statistical relevance from which one can estimate a reliable error, these simulations must be at least of the order of several nanoseconds. Moreover, to avoid underestimating the statistical error of a single simulation, it is necessary that the ensemble averages calculated from the simulations cover a time span that is at least comparable to the auto-correlation time of the observed quantity, that is, the square amplitudes of the fluctuation spectrum. Given that Fourier modes have auto-correlation times that grow as the associated wave-vector approaches zero, one faces the additional difficulty that the modes that most closely approach the thermodynamic limit are the hardest to converge in terms of statistical accuracy. The results we obtained on the interface free energy and its anisotropy shown in figs. 6.6 and 6.7 clearly indicate the limit of having simulations that are too short, although they are of the correct order of magnitude and, at least those of pure metals, in good agreement with the results of other works. We must stress again that the interatomic potential may have a sizable effect: when using different potentials, a property such as the interface free energy may vary by a factor of two or three¹, even because, for example, the melting temperature is underestimated by a few tens of kelvins.

By looking at fig. 6.6 and table 6.1, we should comment on two aspects. Firstly, we are confident that longer simulations — or performing multiple independent simulations for the same temperature-composition point — may better estimate the errors. Secondly, we can still say that the magnitude of the anisotropy of γ_{sl} predicted by our calculations falls in a meaningful interval over the whole compositional range since when a solid phase with a crystalline FCC structure is involved, the energy trend should be $\gamma_{(100)} > \gamma_{(110)} > \gamma_{(111)}$, and this has already been verified by simulations and experiments for several pure metals. Although the results obtained at the eutectic composition predict a slightly negative anisotropy of γ_{sl} — between -0.2% and 2% — they lead us to believe that the dependence on the temperature of the range is not an unphysical artifact of our simulations, but seems to grasp a fundamental aspect of an invariant equilibrium such as the eutectic, that is the coexistence of three phases with three very different compositions.

Kinetic coefficient Gibbs-Thomson's equation (eq. (4.32)) highlights how the thermal gradient that develops across the solid-liquid interface between depends as much on the interface free energy as on the propensity to solidify in a particular direction. This second contribution is thus fundamental for the accurate description of the phenomenon. We have shown that different

computational methods allow estimating the kinetic coefficient of the interface, each with its advantages and disadvantages. As part of our approach based on Gibbs' construction, the model based on an over-damped dynamics that follow Langevin's equations states that the kinetic coefficient can be obtained from two ingredients that one can get directly from a simulation of a pinned solid-liquid interface. Consider again the eq. (6.12): the first term at the numerator measures the fluctuations of the solid fraction — in practice, it is the variance of f . The second term, τ_f , is the relaxation time of these fluctuations, and it can easily be obtained provided that one has a simulation long enough during which a sufficiently fine-grained sampling of the quantity f has been made.

We have shown how the above equation can correctly predict the kinetic coefficient of the elemental metals (silver and copper) whose values are in good agreement with those presented in some previous works of the scientific literature. However, we faced the first problem when it has been evident that we could not use a model with a single exponential to describe the fluctuations of the interface in terms of the fraction of solid estimated and monitored during the simulation. In particular, we noticed that in some simulations with a specific value of the total concentration of silver were evident a drift of the interface that should instead fluctuate around its position of equilibrium — the latter controlled continuously by an external potential applied for this purpose. The origin of this drift is not yet entirely clear, but we suspect it may be related either to the need to add a Monte Carlo algorithm that serves to equilibrate the concentrations in the two phases or to the formation near the interface of another ordered solid phase.

The capillary fluctuation method (CFM) is undoubtedly a consolidated method well-suited to study both simple systems and more complex alloys as it only implies local re-equilibration of solute atoms, yet it requires simulations of considerable size and long sampling time.

We first tried to reduce the impact of finite-size effects by smoothing the interface. This approach worked but, at the same time, entailed a few arbitrary choices, such as the value of a threshold to locate the dividing surface and the extent of the smoothing procedure, which are only partly amendable. We then attempted to use a more coherent approach based on a definition of dividing surface proposed originally by Gibbs and to streamline the complex sequence of calculations that are needed to determine both equilibrium and dynamical properties of solid-liquid interfaces in alloyed systems.

While this approach showed to work well for bulk and thermodynamic interfacial properties, it suffers from the same size and time convergence issues that trouble conventional CFM studies, which are exacerbated by the need for sampling compositional fluctuations. When it comes to dynamical properties of the interface, we observed significant artifacts that arguably suggest that the cause is a combination of statistical errors, latent heat, and solute repartition among the solid and the liquid phase; all these factors eventually render the estimation of interface mobility a particularly endeavoring computational task.

When one also considers that current embedded atom method (EAM) potentials, which represent the most effective choice to model metallic systems, are often not sufficiently reliable to estimate interfacial properties — with published results that may well differ by a factor of 2 or 3 between different semi-empirical fits — one sees that more work should be devoted to the study of this problem. The surge of machine learning potentials applied to numerous

Chapter 7. Conclusions

investigations of materials properties offers an alternative approach to the definition of a more accurate interatomic potential, but the much higher cost compared to EAM makes the sampling and system size challenges even more severe. Combining the insights that computational methods as the CFM have so far provided with more recent approaches based on accelerated sampling techniques might be the most promising alley to solve this puzzle.

Appendices **Part III**

A Interface properties of a Lennard-Jones model

The purpose of this appendix is to present and discuss some additional results related to interface properties of a Lennard–Jones system. Besides choosing a Lennard–Jones model because of its cheap computational cost, the main reason was to benchmark our framework and be able to compare results on interface stiffness, free energy and kinetic coefficient with previous works of the literature that already investigated these properties in detail.

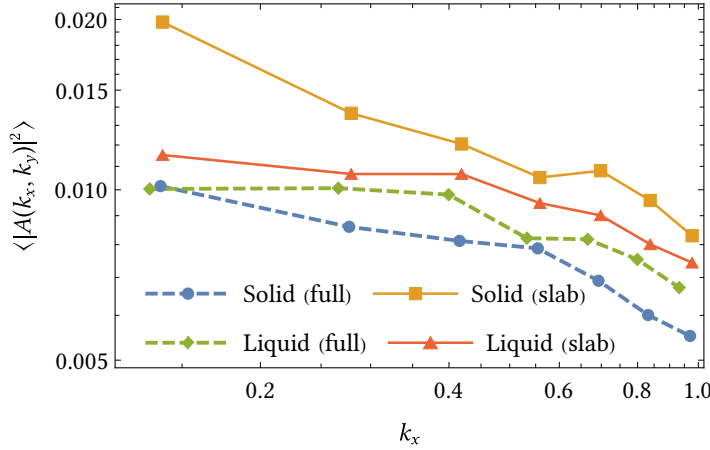
A.1 Interface stiffness and free energy

As we discussed both in the theory and result part, the main advantage of a Gibbs approach in analyzing interface fluctuations to extract the stiffness and free energy is that it sidesteps the inevitably arbitrary choice of a way to locate the dividing surface. On the other hand, it relies on two approximations that we restate here to justify the additional results discussed in the following.


The first assumption is that bulk fluctuations in the liquid are independent of those in the solid; this suggests that we could estimate the bulk contributions to the fluctuation spectrum of the interface from one simulation of a solid-liquid system. The second assumption is that, with a planar interface perpendicular to the z dimension, the characteristic length of spatial correlations of bulk fluctuations is much smaller than the extent of the simulation cell along z . This second assumption implies that we can evaluate the contributions from bulk slabs of suitably chosen widths, and there should not be any significant dependence of such contributions on the values of the slab widths, provided that the thickness of bulk regions in the supercell is large enough to consider a slab as a true bulk phase.

We have therefore verified both these two approximations. We prepared two one-dimensional simulation cells with sizes $28 \times 4 \times 28$ (12 544 atoms) and $40 \times 7 \times 40$ (44 800 atoms) and we calculated the contributions to $\langle |A(k_x, k_y)|^2 \rangle$ according to the eq. (6.5) in two ways: (1) considering two slabs, one in the solid phase and one in the liquid phase, with a thickness of about $0.2 L_z$, where L_z was the total size of the periodic cell along z ; (2) running two independent simulations for each bulk phase, each with the same dimensions of the simulation cell containing the interface. The results are shown in fig. A.1. We can immediately notice two things: in the smaller cell (fig. A.1a), the contributions of both bulk phases show a marked dependence on k , and the fact that the solid bulk is comparable with the liquid bulk suggests that this system's size is not at

Appendix A. Interface properties of a Lennard-Jones model



(a)



(b)

Figure A.1 – Plots of $\langle |A(k_x, k_y)|^2 \rangle$ versus k_x ($k_y = 0$) for two different cells with a one-dimensional geometry: (a) $28 \times 4 \times 28$ cell and (b) $40 \times 7 \times 40$ cell. Both these plots show a k -dependence of the solid bulk contributions when they are calculated from bulk slabs instead of full cells.

all sufficient to correctly sample capillary fluctuations. Although we cannot exclude *a priori* that the contributions of the bulk phases depend on k , we would expect an almost constant trend precisely because the terms of the eq. (6.5) that correspond to the bulk phases do not represent the fluctuations of any interface, but only those of the order parameter field, which in turn reflect density fluctuations. In fig. A.1b, it is evident that a larger cell does not show these artifacts, although there remains a k -dependency when bulk phase contributions are calculated from slabs and not from full bulk cells. The presence of such “slab effect” led us to conclude that it is not possible to estimate the contributions of the bulk phases correctly from the same simulation cell that contains the solid-liquid interface. Since a $1/k^2$ model to fit the fluctuation spectrum and extract the stiffness is suitable to reproduce the behavior of the spectrum only for small k , the k -dependency, more prominent for the solid bulk contribution, leads to a wrong

estimation of the interface fluctuation spectrum when subtracting the bulk terms evaluated from slabs from the “full spectrum,” that is, the right-hand side of eq. (6.5).

Finally, to rule out that this effect due to evaluating bulk contributions from slabs did not depend on the one-dimensional geometry of the simulation cells, we performed a third set of simulations with a two-dimensional cell with size $40 \times 40 \times 24$ (153 600 atoms) and verified (see fig. A.2) that the k -dependence appears whenever one considers bulk slabs instead of full cells. This conclusion is rather undesirable because it does not only force the use of simulations cells of considerable size, whose computational cost increases rapidly even for a semi-empirical and relatively cheap potential such as the EAM, but it also requires to perform three independent simulations for each interface whose stiffness is to be calculated: a simulation of the solid-liquid interface and two from which to determine correctly the contributions of the bulk phases.




Figure A.2 – Plot of $\langle |A(k_x, k_y)|^2 \rangle$ versus k_x ($k_y = 0$) for a simulation cell with two-dimensional geometry with size $40 \times 40 \times 24$. Although the bulk contributions of the liquid phase are now independent of whether they are evaluated from a slab section or a full cell, the solid bulk contributions still show a remarkable difference especially at small k , suggesting that this “slab effect” is affecting more prominently the solid phase.

One last piece of evidence that convinced us that the only way to obtain a reliable estimate of the stiffness – and therefore of the interfacial free energy – was to perform two-dimensional simulations containing at least one hundred thousand atoms was the presence of finite-size effects. From fig. A.1 it is clear that there is a substantial difference in the magnitude of $\langle |A(k_x, k_y)|^2 \rangle$. Figure A.3 illustrates this effect much more clearly by plotting together bulk phases contributions to the fluctuation spectrum as calculated from full bulk cells only. As the cell size increases, not only do the contributions of liquid and solid separate more clearly, but also the absolute value of $\langle |A(k_x, k_y)|^2 \rangle$ is reduced by about two orders of magnitude in the case of bulk solid due to a scaling effect.

Although we did these tests on a model system such as the Lennard–Jones that is much simpler than that investigated in the second part of the results of this thesis – the silver-copper binary alloy – we can reasonably conclude that this revisited CFM that avoids, thanks to Gibbs dividing surface construction, the need to find a method to define the crystal-melt boundary, forces us to simulate extensive systems with two-dimensional geometries, the latter a limit that




Figure A.3 – Plot of the liquid and solid bulk contributions to $\langle |A(k_x, k_y)|^2 \rangle$ as a function of k for different simulation cells sizes, with both one and two-dimensional geometries. As the cell size increases, the difference between the contributions of liquid and solid in the same simulation and between the same bulk phase in two different simulations increases. In particular, this effect is particularly pronounced for the solid phase, where there is a difference of about two orders of magnitude in the absolute value of $\langle |A(k_x, k_y)|^2 \rangle$ between the smallest cell, $28 \times 4 \times 28$, and the largest $40 \times 40 \times 24$.

the original CFM had not shown². Furthermore, there is the challenge of converging this type of calculations that need a long simulation time to be able to estimate an error on the measured quantities of interface stiffness and free energy that is statistically acceptable. In conclusion, it is worth noting that the need to run simulations that have to be both overly long in time and large are well recognized inherent limitations of the CFM itself that some computational trick cannot always circumvent. Nevertheless, the method continues to be a solid choice for investigating interface properties, and although it involves studying systems containing many particles that might render the method prohibitively costly, it can still provide relevant insights even on complex systems of considerable practical interest¹.

A.2 Interface dynamics

In section “Interface dynamics (section 4.3.3),” we outlined some well-established computational methods particularly suited to determine the kinetic coefficient of a crystal-melt interface from molecular dynamics simulations. These methods include both equilibrium and non-equilibrium techniques; for example, the analysis of interface fluctuations through its Fourier spectrum belongs to the former class, while free solidification simulations are an example of the latter.

¹A good example of such application of CFM is the work by Mishin [37], a comprehensive study of interface properties of the γ - γ' phase boundary in the Ni_3Al system of fundamental technological importance for high-performance alloys.

An issue that affects both equilibrium and non-equilibrium methods differently is related to the latent heat absorbed or released during melting or solidification. The direct consequence of latent heat generation is a temperature gradient across the interface, and thus the choice of the thermostat used to keep the simulation's temperature constant is a very crucial one. In practice, the less efficient the thermostat is in removing the heat generated during crystallization, the more interface dynamics will show the footprint of how quickly the heat flows away from the interface. In more precise words, this would mean that crystal growth could not be considered anymore a diffusion-limited process dominated only by the kinetics of atoms of the liquid phase attaching to the growing crystalline boundary.

The problem of the separation of atoms' diffusion and heat flow time-scales has been already investigated extensively in the literature [83, 142–145]. When simulating real systems, especially metals and alloys, this problem can have an even more pronounced effect if the interatomic potential used does not correctly describe the thermal diffusion. For metals, this is particularly true in the case of EAM potentials, which can underestimate up to 100 times the thermal diffusivity [146]. Although Briels and Tepper [104] have shown that the thermostatting strategy can have an effect on crystallization in a Lennard-Jones system, the latter is not affected by the problem of the underestimation of thermal conductivity, and it is thus possible to compare results of the kinetic coefficient obtained with different methods.

To this end, we prepared and ran two sets of simulations of a Lennard–Jones system modeled with the same potential employed for the results showed previously (section 5.3). The first set was of equilibrium simulations of a solid-liquid interface from which we extracted the kinetic coefficient by analyzing the fluctuations of the solid fraction as we did for the silver-copper system (section 6.4.2). We also investigated the effect of two different values of the strength of the harmonic bias applied to pin the solid-liquid interface. These simulations were prepared with a protocol similar to that we followed for the simulations on the silver-copper system. The supercells for these Lennard–Jones contained 134 400 atoms and where run for about 2000 Lennard–Jones time-units, and solid fraction fluctuations were analyzed on a time span of roughly 1200 time-units. Both simulations were run in the NP_zT ensemble with a stochastic velocity-rescaling thermostat [50] controlling the temperature.

The plot of fig. A.4 shows the monitored value of the solid fraction in two pinned-interface simulations with different values of κ , the harmonic spring constant. By looking at that plot and remembering eq. (6.12),

$$M = \frac{v_s}{2Sk_B T \tau_f} N^2 \langle (f - \bar{f})^2 \rangle,$$

we expect the auto-correlation time τ_f associated with the fluctuations in the simulation with a stiffer harmonic bias to be smaller than that of the simulation with a weaker bias. However, since the variance of f is also larger, the two simulations yield two values of the kinetic coefficient – 1.240 ($\kappa = 1 \times 10^{-4}$) and 1.741 ($\kappa = 1 \times 10^{-5}$) in Lennard–Jones units of $(m\epsilon)^{-1/2}$ – that are both comparable with the results obtained by Briels and Tepper [104], even though they studied the same model system with a different interatomic potential³.

In the second set of simulations, we employed the already described free solidification method (section 4.3.3) to estimate the kinetic coefficient from the analysis of the interface growth as

Appendix A. Interface properties of a Lennard-Jones model




Figure A.4 – Plot of the solid fraction as a function of time monitored during two pinned-interface simulations of Lennard–Jones in which two different values of the strength of the harmonic bias applied have been used.

a function of the degree of undercooling. In particular, we started from a simulation box containing two solid-liquid boundaries equilibrated at the melting temperature of $T_m = 0.62$ (in Lennard–Jones reduced units) and, during equilibration, we monitored that the number of solid atoms was fluctuating around one third of the total atoms present in the simulation cell. Subsequently, we rescaled the system’s temperature to a value below T_m according to a given undercooling between 2% and 5%. We then let the simulation run in the NPT ensemble for as many time-steps as needed after which, except for the smallest undercooling, the system was crystalline in its entirety. The following relation then yields an estimate of the crystallization rate

$$R = \frac{v_s}{2N_s a_0^2} \frac{dN_s}{dt},$$

where v_s is the atomic volume of the solid, N_s is the number of atoms per interface plane, and a_0 is the lattice parameter. The above relation is straightforward to apply since the only information needed is the number of solid atoms as a function of time, a quantity that the simulation’s trajectory directly provides. Figure A.5 shows the the number of solid atoms N_s as a function of time for different values of undercooling. Besides the plateaus which indicate that several simulations were already completely solidified after about 100 Lennard–Jones time-units, the linear growth of N_s confirms that we managed to attain steady non-equilibrium conditions.

Since it is known [104] that the choice of thermostat has a non-negligible effect on interface dynamics, we ran five independent sets of simulations within the same undercooling range with two different thermostat: the global, stochastic velocity–rescaling thermostat and a local one, a slightly modified Langevin thermostat whose implementation [147] correctly preserves the momentum distribution and showed to be more efficient than the standard implementation in the over-damped regime. The results obtained for the crystal growth rate R are shown in fig. A.6, where the two straight lines represent the linear fits whose slope is the kinetic coefficient. We




Figure A.5 – Plot of the number of solid atoms N_s as a function of time during several non-equilibrium simulations with different values of undercooling ranging from 2.0% (rightmost blue curve) to 5.0% (leftmost brown curve). In these simulations, temperature was fixed by a local, Fast-forward Langevin thermostat [147]. The plateau regions for all but the smallest undercooling indicate that our simulations were long enough to observe a complete solidification, and the linear growth of N_s confirm that constant non-equilibrium conditions have been attained for all the simulations. The dashed vertical lines indicate the cutoffs up to which the linear fits were performed, with the rightmost vertical line referring to the smallest undercooling and so on.

obtained the two values of 2.51 ± 0.18 and 2.18 ± 0.15 – in Lennard–Jones units of $k_B/(m\epsilon)^{1/2}$ – from simulations using the global and the local thermostats respectively. Although the measured crystal growth rates differ appreciably, the values of the kinetic coefficient are consistent with each other within the estimated errors.

In order to compare the kinetic coefficient obtained with free solidification simulations – where M links the interface normal velocity to the undercooling – and that obtained by the analysis of equilibrium interface fluctuations through the solid fraction, one needs the information of the dependence of the chemical potential on temperature. For this particular Lennard–Jones potential, $d\mu/dT$ has been calculated by Cheng, Tribello, and Ceriotti [117] in a study of homogeneous nucleation. The values of M in terms of the chemical potential difference then become 1.27 ± 0.09 and 1.46 ± 0.10 for the simulations with the local and the global thermostat, respectively. These two values are close enough to the results we obtained from equilibrium fluctuations analysis, suggesting that, at least for a potential that is suited to model a unary system and does not suffer from a wrong description of thermal phenomena, our approach provides consistent values of the kinetic coefficient.

Appendix A. Interface properties of a Lennard-Jones model




Figure A.6 – Crystal growth rate as a function of temperature below the melting temperature of $0.62 \epsilon/k_B$. The two curves represent two sets of simulations where different thermostats were keeping the temperature constant during the non-equilibrium run: the orange curve reports the results from the simulation in which a global, stochastic velocity-rescaling thermostat [50] was used, while the blue curve reports results from simulations with a local, Fast-forward Langevin [147] thermostat. The values of R and their error bars are calculated from five independent simulations with the same undercooling. Besides the shift between the two curves, the linear fits (black dashed lines) suggest two consistent values of the mobility, which turned out to be 2.51 ± 0.18 (stochastic thermostat) and 2.18 ± 0.15 (Langevin thermostat).

B Block averaging in practice

This brief appendix illustrates in practice how one can apply the block averaging technique discussed previously (section 2.7.1). Suppose to have a series of N measurements $A_1, A_2 \dots A_N$. Taking the average of all consecutive pairs of data, one obtains a dataset half as large: $A'_i = (A_{2i-1} + A_{2i})/2$. It can be shown that both $\langle A \rangle$ and $\sigma^2(A)$ are invariant under this blocking procedure. In particular, it is evident that $\langle A \rangle = \langle A' \rangle$ and

$$\sigma^2(A') = \frac{1}{N'} \sum_i A_i'^2 - \langle A' \rangle^2. \quad (\text{B.1})$$

Provided that the simulation has been performed for long enough, $\frac{\sigma^2(A')}{N'-1}$ converges to a constant value as the number of blocks M increases: Flyvbjerg and Petersen [19] showed that this limiting, constant value is a correct estimate of $\sigma^2(A)$ – one that takes into account the correlation between successive measurements. Hence, we can take as an estimate of the ensemble average of the observable A the value

$$\langle A \rangle \pm \frac{\sigma(A')}{\sqrt{N'-1}}. \quad (\text{B.2})$$

In practice, plotting σ^2 versus the number of blocks M gives a visual estimate of the value of M for which the averages A'_i become completely uncorrelated. Increasing further the number of blocks should not lead to any significant change in σ^2 , and one can then take eq. (B.2) as a correct estimate of $\langle A \rangle$.

Bibliography

- [1] K. Kadau, T. C. Germann, and P. S. Lomdahl. “Molecular Dynamics Comes of Age: 320 Billion Atom Simulation on BlueGene/L”. In: *Int. J. Mod. Phys. C* 17.12 (Dec. 1, 2006), pp. 1755–1761. DOI: 10.1142/S0129183106010182 (cit. on p. 1).
- [2] J. Yang, Y. Wang, and Y. Chen. “GPU Accelerated Molecular Dynamics Simulation of Thermal Conductivities”. In: *Journal of Computational Physics* 221.2 (Feb. 10, 2007), pp. 799–804. DOI: 10.1016/j.jcp.2006.06.039 (cit. on p. 1).
- [3] Y. Shibuta et al. “Heterogeneity in Homogeneous Nucleation from Billion-Atom Molecular Dynamics Simulation of Solidification of Pure Metal”. In: *Nat Commun* 8.1 (Apr. 5, 2017), pp. 1–9. DOI: 10.1038/s41467-017-00017-5 (cit. on p. 1).
- [4] C. Y. Yap et al. “Review of Selective Laser Melting: Materials and Applications”. In: *Applied Physics Reviews* 2.4 (Dec. 1, 2015), p. 041101. DOI: 10.1063/1.4935926 (cit. on p. 3).
- [5] P. Fischer et al. “Microstructure of Near-Infrared Pulsed Laser Sintered Titanium Samples”. In: *Appl Phys Mater Sci Process* 78.8 (2004), pp. 1219–1227 (cit. on p. 4).
- [6] B. V. Krishna, S. Bose, and A. Bandyopadhyay. “Low Stiffness Porous Ti Structures for Load-Bearing Implants”. In: *Acta Biomaterialia* 3.6 (Nov. 1, 2007), pp. 997–1006. DOI: 10.1016/j.actbio.2007.03.008 (cit. on p. 4).
- [7] W. Xue et al. “Processing and Biocompatibility Evaluation of Laser Processed Porous Titanium”. In: *Acta Biomaterialia* 3.6 (Nov. 1, 2007), pp. 1007–1018. DOI: 10.1016/j.actbio.2007.05.009 (cit. on p. 4).
- [8] F. Abe et al. “The Manufacturing of Hard Tools from Metallic Powders by Selective Laser Melting”. In: *Journal of Materials Processing Technology*. International Symposium on Advanced Forming and Die Manufacturing Technology 111.1 (Apr. 25, 2001), pp. 210–213. DOI: 10.1016/S0924-0136(01)00522-2 (cit. on p. 4).
- [9] R. W. Cahn. “Ni-Ti-Al Alloys”. Pat. (U.S.). Rolls-Royce plc (London, GB2). Sept. 8, 1994 (cit. on p. 5).
- [10] M. Haataja et al. “Scaling of Domain Size during Spinodal Decomposition: Dislocation Discreteness and Mobility Effects”. In: *Appl. Phys. Lett.* 87.25 (2005), p. 251901. DOI: 10.1063/1.2147732 (cit. on p. 9).
- [11] J. Williams and A. Fatemi. “Fatigue Performance of Forged Steel and Ductile Cast Iron Crankshafts”. In: *SAE Trans.* 116 (2007), pp. 339–349 (cit. on p. 10).
- [12] J. Dantzig and M. Rappaz. *Solidification*. 1st. CRC Press, 2009 (cit. on p. 10).
- [13] A. Karma. “Fluctuations in Solidification”. In: *Phys. Rev. E* 48.5 (Nov. 1, 1993), pp. 3441–3458. DOI: 10.1103/PhysRevE.48.3441 (cit. on pp. 11, 65, 101, 116).

Bibliography

- [14] J. F. Lutsko and J. Lam. “Classical Density Functional Theory, Unconstrained Crystallization, and Polymorphic Behavior”. In: *Phys. Rev. E* 98.1 (July 12, 2018), p. 012604. doi: 10.1103/PhysRevE.98.012604 (cit. on p. 11).
- [15] J. F. Lutsko and M. Baus. “Can the Thermodynamic Properties of a Solid Be Mapped onto Those of a Liquid?” In: *Phys. Rev. Lett.* 64.7 (Feb. 12, 1990), pp. 761–763. doi: 10.1103/PhysRevLett.64.761 (cit. on p. 11).
- [16] D. Ruelle. *Statistical Mechanics: Rigorous Results*. Benjamin, 1977 (cit. on p. 15).
- [17] M. Allen and D. Tildesley. *Computer Simulation of Liquids*. Second edition. Oxford University Press, 2017 (cit. on pp. 15, 29).
- [18] D. Ruelle. *Hasard et Chaos*. Paris: Éditions Odile Jacob, 1991 (cit. on p. 15).
- [19] H. Flyvbjerg and H. G. Petersen. “Error Estimates on Averages of Correlated Data”. In: 91 (1989), p. 7 (cit. on pp. 22, 135).
- [20] M. S. Daw and M. I. Baskes. “Semiempirical, Quantum Mechanical Calculation of Hydrogen Embrittlement in Metals”. In: *Phys. Rev. Lett.* 50.17 (Apr. 25, 1983), pp. 1285–1288. doi: 10.1103/PhysRevLett.50.1285 (cit. on pp. 24, 25).
- [21] M. S. Daw and M. I. Baskes. “Embedded-Atom Method: Derivation and Application to Impurities, Surfaces, and Other Defects in Metals”. In: *Phys. Rev. B* 29.12 (June 15, 1984), pp. 6443–6453. doi: 10.1103/PhysRevB.29.6443 (cit. on pp. 24, 26).
- [22] M. W. Finnis and J. E. Sinclair. “A Simple Empirical N-Body Potential for Transition Metals”. In: *Philos. Mag. A* 50.1 (1984), pp. 45–55 (cit. on pp. 24, 26).
- [23] M. S. Daw, S. M. Foiles, and M. I. Baskes. “The Embedded-Atom Method: A Review of Theory and Applications”. In: *Mater. Sci. Rep.* 9.7-8 (Mar. 1993), pp. 251–310. doi: 10.1016/0920-2307(93)90001-U (cit. on p. 25).
- [24] M. S. Daw. “Model of Metallic Cohesion: The Embedded-Atom Method”. In: *Phys. Rev. B* 39.11 (Apr. 15, 1989), pp. 7441–7452. doi: 10.1103/PhysRevB.39.7441 (cit. on pp. 26, 38).
- [25] P. Hohenberg and W. Kohn. “Inhomogeneous Electron Gas”. In: *Phys. Rev.* 136 (3B Nov. 9, 1964), B864–B871. doi: 10.1103/PhysRev.136.B864 (cit. on p. 26).
- [26] M. I. Baskes. “Modified Embedded-Atom Potentials for Cubic Materials and Impurities”. In: *Phys. Rev. B* 46.5 (1992), pp. 2727–2742. doi: 10.1103/PhysRevB.46.2727 (cit. on pp. 27, 108).
- [27] M. I. Baskes, J. S. Nelson, and A. F. Wright. “Semiempirical Modified Embedded-Atom Potentials for Silicon and Germanium”. In: *Phys. Rev. B* 40.9 (Sept. 15, 1989), pp. 6085–6100. doi: 10.1103/physrevb.40.6085 (cit. on p. 27).
- [28] M. I. Baskes. “Application of the Embedded-Atom Method to Covalent Materials: A Semiempirical Potential for Silicon”. In: *Phys. Rev. Lett.* 59.23 (Dec. 7, 1987), pp. 2666–2669. doi: 10.1103/PhysRevLett.59.2666 (cit. on p. 27).
- [29] H. Yang, M. Huang, and Z. Li. “The Influence of Vacancies Diffusion-Induced Dislocation Climb on the Creep and Plasticity Behaviors of Nickel-Based Single Crystal Superalloy”. In: *Computational Materials Science* 99 (Mar. 1, 2015), pp. 348–360. doi: 10.1016/j.commatsci.2014.12.035 (cit. on p. 27).

- [30] C. M. F. Rae and R. C. Reed. "Primary Creep in Single Crystal Superalloys: Origins, Mechanisms and Effects". In: *Acta Materialia* 55.3 (Feb. 1, 2007), pp. 1067–1081. doi: 10.1016/j.actamat.2006.09.026 (cit. on p. 27).
- [31] R. Pasianot et al. "Computer Simulation of (100) Dislocation Core Structure in NiAl". In: *Modelling Simul. Mater. Sci. Eng.* 2.3 (May 1994), pp. 383–394. doi: 10.1088/0965-0393/2/3/008 (cit. on p. 27).
- [32] A. Voter and S. Chen. "Accurate Interatomic Potentials for Ni, Al and Ni₃Al". In: *MRS Symposia Proceedings*. Vol. 82. Materials Research Society, 1987, pp. 175–180 (cit. on p. 27).
- [33] S. M. Foiles and M. S. Daw. "Application of the Embedded Atom Method to Ni₃Al". In: *J. Mater. Res.* 2.1 (Feb. 1987), pp. 5–15. doi: 10.1557/JMR.1987.0005 (cit. on p. 27).
- [34] A. Rao, C. Woodward, and T. Parthasarathy. "Empirical Interatomic Potentials for L10 TiAl and B2 NiAl". In: *Materials Research Society*. Vol. 213. 1991, pp. 125–130 (cit. on p. 28).
- [35] D. Farkas et al. "Interatomic Potentials for B2 NiAl and Martensitic Phases". In: *Modelling Simul. Mater. Sci. Eng.* 3 (Mar. 1995), pp. 201–214. doi: 10.1088/0965-0393/3/2/005 (cit. on p. 28).
- [36] D. Farkas. "Interatomic Potentials for Ti-Al with and without Angular Forces". In: *Model. Simul. Mater. Sci. Eng.* 2 (1994), pp. 975–984 (cit. on pp. 28, 38).
- [37] Y. Mishin. "Atomistic modeling of the γ and γ' for the Ni–Al system". In: *Acta Mater* 52.6 (2004), pp. 1451–1467 (cit. on pp. 28, 66, 130).
- [38] F. Bianchini, J. R. Kermode, and A. D. Vita. "Modelling Defects in Ni–Al with EAM and DFT Calculations". In: *Modelling Simul. Mater. Sci. Eng.* 24.4 (Apr. 2016), p. 045012. doi: 10/gf4wz8 (cit. on p. 28).
- [39] Q. N. Fan et al. "A Ternary Ni-Al-W EAM Potential for Ni-Based Single Crystal Superalloys". In: *Phys B Condens Matter* 456 (2015), pp. 283–292. doi: 10.1016/j.physb.2014.09.012 (cit. on p. 28).
- [40] G. Simmons and H. Wang. *Single Crystal Elastic Constants and Calculated Aggregate Properties*. MIT Press, 1977 (cit. on p. 28).
- [41] R. Weast, ed. *Handbook of Chemistry and Physics*. CRC Press, 1984 (cit. on p. 28).
- [42] G. Grochola, S. P. Russo, and I. K. Snook. "On Fitting a Gold Embedded Atom Method Potential Using the Force Matching Method." In: *J Chem Phys* 123.20 (2005), p. 204719 (cit. on pp. 28, 29).
- [43] S. M. Foiles, M. I. Baskes, and M. S. Daw. "Embedded-Atom-Method Functions for the Fcc Metals Cu, Ag, Au, Ni, Pd, Pt, and Their Alloys". In: *Phys Rev B* 33.12 (June 1986), pp. 7983–7991 (cit. on p. 29).
- [44] R. A. Johnson. "Analytic Nearest-Neighbor Model for Fcc Metals". In: *Phys Rev B* 37 (1988), pp. 3924–3931 (cit. on p. 29).
- [45] F. Ercolessi, M. Parrinello, and E. Tosatti. "Simulation of Gold in the Glue Model". In: *Philos. Mag. A* 58.1 (1988), pp. 213–226 (cit. on p. 29).
- [46] M. Tuckerman. *Statistical Mechanics: Theory and Molecular Simulation*. Oxford University Press, 2010 (cit. on pp. 29, 33, 34).

Bibliography

- [47] D. Frenkel and B. Smit. *Understanding Molecular Simulation*. Academic Press, 2002 (cit. on pp. 29, 32–34).
- [48] W. C. Swope et al. “A Computer Simulation Method for the Calculation of Equilibrium Constants for the Formation of Physical Clusters of Molecules: Application to Small Water Clusters”. In: *J Chem Phys* 76 (1982), p. 637 (cit. on p. 31).
- [49] D. J. Evans and O. P. Morriss. “Non-Newtonian Molecular Dynamics”. In: *Computer Physics Reports* 1.6 (July 1, 1984), pp. 297–343. DOI: 10.1016/0167-7977(84)90001-7 (cit. on p. 32).
- [50] G. Bussi, D. Donadio, and M. Parrinello. “Canonical Sampling through Velocity Rescaling”. In: *J Chem Phys* 126 (2007), p. 014101 (cit. on pp. 32, 131, 134).
- [51] H. C. Andersen. “Molecular Dynamics Simulations at Constant Pressure and/or Temperature”. In: *J. Chem. Phys.* 72.4 (Feb. 15, 1980), pp. 2384–2393. DOI: 10.1063/1.439486 (cit. on pp. 32, 33).
- [52] S. Nosé. “A Unified Formulation of the Constant Temperature Molecular Dynamics Methods”. In: *J. Chem. Phys.* 81.1 (July 1, 1984), pp. 511–519. DOI: 10.1063/1.447334 (cit. on p. 32).
- [53] W. G. Hoover. “Canonical Dynamics: Equilibrium Phase-Space Distributions”. In: *Phys. Rev. A* 31.3 (Mar. 1, 1985), pp. 1695–1697. DOI: 10.1103/PhysRevA.31.1695 (cit. on p. 32).
- [54] M. E. Tuckerman et al. “Non-Hamiltonian Molecular Dynamics: Generalizing Hamiltonian Phase Space Principles to Non-Hamiltonian Systems”. In: *J. Chem. Phys.* 115.4 (July 12, 2001), pp. 1678–1702. DOI: 10.1063/1.1378321 (cit. on p. 33).
- [55] N. Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. In: *J. Chem. Phys.* 21.6 (June 1, 1953), pp. 1087–1092. DOI: 10.1063/1.1699114 (cit. on p. 35).
- [56] C. Jarzynski. “Nonequilibrium Equality of Free Energy Differences”. In: *Phys. Rev. Lett.* 78 (1997), pp. 2690–2693. DOI: 10.1103/PhysRevLett.78.2690 (cit. on p. 36).
- [57] T. P. Straatsma and J. A. McCammon. “Multiconfiguration Thermodynamic Integration”. In: *J. Chem. Phys.* 95.2 (July 15, 1991), pp. 1175–1188. DOI: 10.1063/1.461148 (cit. on p. 36).
- [58] M. Sprik and G. Ciccotti. “Free Energy from Constrained Molecular Dynamics”. In: *J. Chem. Phys.* 109.18 (Nov. 8, 1998), pp. 7737–7744. DOI: 10.1063/1.477419 (cit. on p. 36).
- [59] G. Ciccotti and M. Ferrario. “Blue Moon Approach to Rare Events”. In: *Mol. Simul.* 30.11-12 (Sept. 1, 2004), pp. 787–793. DOI: 10.1080/0892702042000270214 (cit. on p. 36).
- [60] G. M. Torrie and J. P. Valleau. “Monte Carlo Free Energy Estimates Using Non-Boltzmann Sampling: Application to the Sub-Critical Lennard-Jones Fluid”. In: *Chem Phys Lett* 28 (1974), pp. 578–581. DOI: 10.1016/0009-2614(74)80109-0 (cit. on p. 36).
- [61] G. M. Torrie and J. P. Valleau. “Nonphysical Sampling Distributions in Monte Carlo Free-Energy Estimation: Umbrella Sampling”. In: *Journal of Computational Physics* 23.2 (Feb. 1, 1977), pp. 187–199. DOI: 10.1016/0021-9991(77)90121-8 (cit. on p. 36).
- [62] R. Martin. *Electronic Structure*. New York: Cambridge University Press, 2008 (cit. on p. 38).
- [63] D. Marx and J. Hutter. *Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods*. 1 edition. Cambridge University Press, May 25, 2009. 578 pp. (cit. on p. 38).

- [64] M. Ceriotti. “Unsupervised Machine Learning in Atomistic Simulations, between Predictions and Understanding”. In: *J. Chem. Phys.* 150.15 (Apr. 19, 2019), p. 150901. DOI: 10.1063/1.5091842 (cit. on p. 38).
- [65] M. W. Finnis. *Interatomic Forces in Condensed Matter*. Oxford University Press, 2003 (cit. on p. 38).
- [66] A. Laio and M. Parrinello. “Escaping Free-Energy Minima”. In: *PNAS* 99.20 (Oct. 1, 2002), pp. 12562–12566. DOI: 10.1073/pnas.202427399 (cit. on pp. 38, 84).
- [67] E. Fermi. *Thermodynamics*. Dover, 1956 (cit. on p. 39).
- [68] P. Gordon. *Principles of Phase Diagrams in Materials Systems*. McGraw-Hill, 1968 (cit. on p. 39).
- [69] P. Atkins and J. de Paula. *Physical Chemistry*. 10th ed. Oxford University Press, 2014 (cit. on p. 39).
- [70] P. J. Steinhardt, D. R. Nelson, and M. Ronchetti. “Bond-Orientational Order in Liquids and Glasses”. In: *Phys. Rev. B* 28.2 (1983), p. 784 (cit. on p. 61).
- [71] W. Lechner and C. Dellago. “Accurate Determination of Crystal Structures Based on Averaged Local Bond Order Parameters”. In: *J. Chem. Phys.* 129.11 (Sept. 17, 2008), p. 114707. DOI: 10.1063/1.2977970 (cit. on p. 61).
- [72] S. Angioletti-Uberti et al. “Solid-Liquid Interface Free Energy through Metadynamics Simulations”. In: *Phys Rev B* 81.12 (2010), p. 125416. DOI: 10.1103/PhysRevB.81.125416 (cit. on pp. 61, 67, 77).
- [73] E. Baldi, M. Ceriotti, and G. A. Tribello. “Extracting the Interfacial Free Energy and Anisotropy from a Smooth Fluctuating Dividing Surface”. In: *J. Phys. Condens. Matter* 29.44 (Nov. 8, 2017), p. 445001. DOI: 10.1088/1361-648X/aa893d (cit. on pp. 61, 75, 76).
- [74] B. Cheng, G. A. Tribello, and M. Ceriotti. “Solid-Liquid Interfacial Free Energy out of Equilibrium”. In: *Phys. Rev. B* 92.18 (Nov. 9, 2015), p. 180102. DOI: 10.1103/PhysRevB.92.180102 (cit. on pp. 61, 68, 77).
- [75] W. W. Mullins and R. F. Sekerka. “Morphological Stability of a Particle Growing by Diffusion or Heat Flow”. In: *Journal of Applied Physics* 34.2 (Feb. 1, 1963), pp. 323–329. DOI: 10.1063/1.1702607 (cit. on p. 63).
- [76] W. Fehlner and S. Vosko. “Product Representation for Cubic Harmonics and Special Directions for the Determination of the Fermi Surface and Related Properties”. In: *Can J Phys* 54 (1976), p. 2159 (cit. on p. 63).
- [77] J. J. Hoyt, M. Asta, and A. Karma. “Atomistic and Continuum Modeling of Dendritic Solidification”. In: *Materials Science and Engineering: R: Reports* 41.6 (Sept. 30, 2003), pp. 121–163. DOI: 10.1016/S0927-796X(03)00036-6 (cit. on p. 64).
- [78] C. Herring. *The Physics of Powder Metallurgy*. McGraw-Hill, 1951 (cit. on p. 64).
- [79] J. Q. Broughton and G. H. Gilmer. “Molecular Dynamics Investigation of the Crystal–Fluid Interface. VI. Excess Surface Free Energies of Crystal–Liquid Systems”. In: *J. Chem. Phys.* 84.10 (May 15, 1986), pp. 5759–5768. DOI: 10.1063/1.449884 (cit. on pp. 64, 66).

Bibliography

- [80] R. L. Davidchack and B. B. Laird. “Direct Calculation of the Hard-Sphere Crystal/Melt Interfacial Free Energy”. In: *Phys. Rev. Lett.* 85.22 (Nov. 27, 2000), pp. 4751–4754. doi: 10.1103/PhysRevLett.85.4751 (cit. on pp. 64, 67).
- [81] J. J. Hoyt, M. Asta, and A. Karma. “Method for Computing the Anisotropy of the Solid-Liquid Interfacial Free Energy”. In: *Phys. Rev. Lett.* 86.24 (June 11, 2001), pp. 5530–5533. doi: 10.1103/PhysRevLett.86.5530 (cit. on pp. 65, 66, 110).
- [82] M. P. A. Fisher, D. S. Fisher, and J. D. Weeks. “Agreement of Capillary-Wave Theory with Exact Results for the Interface Profile of the Two-Dimensional Ising Model”. In: *Phys. Rev. Lett.* 48.5 (Feb. 1, 1982), pp. 368–368. doi: 10.1103/PhysRevLett.48.368 (cit. on p. 65).
- [83] J. J. Hoyt and M. Asta. “Atomistic Computation of Liquid Diffusivity, Solid-Liquid Interfacial Free Energy, and Kinetic Coefficient in Au and Ag”. In: *Phys. Rev. B* 65.21 (June 3, 2002), p. 214106. doi: 10.1103/PhysRevB.65.214106 (cit. on pp. 66, 108, 110–112, 114, 131).
- [84] J. Morris et al. “The Anisotropic Free Energy of the Solid-Liquid Phase Boundary in Al”. In: *Interface Science* 10.2 (July 1, 2002), pp. 143–148. doi: 10.1023/A:1015800629100 (cit. on p. 66).
- [85] L. Wu et al. “Anisotropic Crystal–Melt Interfacial Energy and Stiffness of Aluminum”. In: *J. Mater. Res.* 30.11 (June 2015), pp. 1827–1835. doi: 10.1557/jmr.2015.131 (cit. on p. 66).
- [86] R. E. Rozas and J. Horbach. “Capillary Wave Analysis of Rough Solid-Liquid Interfaces in Nickel”. In: *EPL* 93.2 (Jan. 2011), p. 26006. doi: 10.1209/0295-5075/93/26006 (cit. on p. 66).
- [87] M. Ambler et al. “Solid-Liquid Interfacial Free Energy of Ice Ih, Ice Ic, and Ice 0 within a Mono-Atomic Model of Water via the Capillary Wave Method”. In: *J. Chem. Phys.* 146.7 (Feb. 21, 2017), p. 074701. doi: 10.1063/1.4975776 (cit. on p. 66).
- [88] J. Benet, L. G. Macdowell, and E. Sanz. “Interfacial Free Energy of the NaCl Crystal-Melt Interface from Capillary Wave Fluctuations”. In: *J. Chem. Phys.* 142.142 (2015), pp. 134706–134705. doi: 10.1063/1.4939641 (cit. on pp. 66, 81).
- [89] J. R. Morris and X. Song. “The Anisotropic Free Energy of the Lennard-Jones Crystal-Melt Interface”. In: *J. Chem. Phys.* 119.7 (2003), pp. 3920–3925. doi: 10.1063/1.1591725 (cit. on p. 66).
- [90] R. L. Davidchack and B. B. Laird. “Simulation of the Hard-Sphere Crystal–Melt Interface”. In: *J. Chem. Phys.* 108.22 (June 1, 1998), pp. 9452–9462. doi: 10.1063/1.476396 (cit. on p. 66).
- [91] R. L. Davidchack and B. B. Laird. “Direct Calculation of the Crystal–Melt Interfacial Free Energies for Continuous Potentials: Application to the Lennard-Jones System”. In: *J. Chem. Phys.* 118.16 (Apr. 11, 2003), pp. 7651–7657. doi: 10.1063/1.1563248 (cit. on pp. 66, 67).
- [92] M. Asta, J. J. Hoyt, and A. Karma. “Calculation of Alloy Solid-Liquid Interfacial Free Energies from Atomic-Scale Simulations”. In: *Phys Rev B* 66.10 (2001), pp. 1001011–1001014. doi: 10.1103/PhysRevB.66.100101 (cit. on p. 67).

- [93] S. Angioletti-Uberti. “The Solid–Liquid Interface Free-Energy of Pb: Comparison of Theory and Experiments”. In: *J. Phys. Condens. Matter* 23.43 (Oct. 2011), p. 435008. doi: 10.1088/0953-8984/23/43/435008 (cit. on pp. 67, 77).
- [94] R. L. Davidchack and B. B. Laird. “Crystal Structure and Interaction Dependence of the Crystal-Melt Interfacial Free Energy”. In: *Phys. Rev. Lett.* 94.8 (Mar. 1, 2005), p. 086102. doi: 10.1103/PhysRevLett.94.086102 (cit. on p. 67).
- [95] J. Q. Broughton, G. H. Gilmer, and K. A. Jackson. “Crystallization Rates of a Lennard-Jones Liquid”. In: *Phys. Rev. Lett.* 49.20 (Nov. 15, 1982), pp. 1496–1500. doi: 10.1103/PhysRevLett.49.1496 (cit. on pp. 69, 70).
- [96] J. Frenkel. In: *Phys Z Sowjetunion* 1 (1932), p. 498 (cit. on pp. 69, 70).
- [97] H. Wilson. “On the Velocity of Solidification and Viscosity of Super-Cooled Liquids”. In: *Lond. Edinb. Dublin Philos. Mag. J. Sci.* 50.303 (Aug. 1, 1900), pp. 238–250. doi: 10.1080/14786440009463908 (cit. on pp. 69, 70).
- [98] K. A. Jackson. “The Interface Kinetics of Crystal Growth Processes”. In: *Interface Science* 10.2 (July 1, 2002), pp. 159–169. doi: 10.1023/A:1015824230008 (cit. on p. 69).
- [99] J. J. Hoyt, M. Asta, and A. Karma. “Atomistic Simulation Methods for Computing the Kinetic Coefficient in Solid-Liquid Systems”. In: *Interface Sci.* 10 (2002), pp. 181–189 (cit. on pp. 69, 71, 100).
- [100] E. Burke, J. Q. Broughton, and G. H. Gilmer. “Crystallization of Fcc (111) and (100) Crystal-melt Interfaces: A Comparison by Molecular Dynamics for the Lennard-Jones System”. In: *J. Chem. Phys.* 89.2 (July 15, 1988), pp. 1030–1041. doi: 10.1063/1.455254 (cit. on p. 69).
- [101] F. Celestini and J.-M. Debierre. “Nonequilibrium Molecular Dynamics Simulation of Rapid Directional Solidification”. In: *Phys. Rev. B* 62.21 (Dec. 1, 2000), pp. 14006–14011. doi: 10.1103/PhysRevB.62.14006 (cit. on p. 70).
- [102] J. J. Hoyt et al. “Kinetic Phase Field Parameters for the Cu–Ni System Derived from Atomistic Computations”. In: *Acta Materialia* 47.11 (Sept. 8, 1999), pp. 3181–3187. doi: 10.1016/S1359-6454(99)00189-5 (cit. on pp. 70, 114).
- [103] A. Karma and W.-J. Rappel. “Quantitative Phase-Field Modeling of Dendritic Growth in Two and Three Dimensions”. In: *Phys. Rev. E* 57.4 (Apr. 1, 1998), pp. 4323–4349. doi: 10.1103/PhysRevE.57.4323 (cit. on p. 71).
- [104] W. J. Briels and H. L. Tepper. “Crystal Growth of the Lennard-Jones (100) Surface by Means of Equilibrium and Nonequilibrium Molecular Dynamics”. In: *Phys. Rev. Lett.* 79.25 (Dec. 22, 1997), pp. 5074–5077. doi: 10.1103/PhysRevLett.79.5074 (cit. on pp. 71, 131, 132).
- [105] J. Gibbs. *The Scientific Papers of J.W. Gibbs*. Vol. 1. New York: Dover, 1961 (cit. on p. 72).
- [106] C. A. Becker et al. “Equilibrium Adsorption at Crystal-Melt Interfaces in Lennard-Jones Alloys”. In: *J. Chem. Phys.* 124.16 (Apr. 26, 2006), p. 164708. doi: 10.1063/1.2185628 (cit. on pp. 72, 106).
- [107] C. Lupis. *Chemical Thermodynamics of Materials*. Elsevier, New York, 1983 (cit. on p. 72).

Bibliography

- [108] C. Niederberger, J. Michler, and A. Jacot. “Inverse Method for the Determination of a Mathematical Expression for the Anisotropy of the Solid-Liquid Interfacial Energy in Al-Zn-Si Alloys”. In: *Phys. Rev. E* 74.2 (Aug. 23, 2006), p. 021604. doi: 10.1103/PhysRevE.74.021604 (cit. on p. 72).
- [109] A. Laio and F. L. Gervasio. “Metadynamics: A Method to Simulate Rare Events and Reconstruct the Free Energy in Biophysics, Chemistry and Material Science”. In: *Rep. Prog. Phys.* 71.12 (Nov. 2008), p. 126601. doi: 10.1088/0034-4885/71/12/126601 (cit. on p. 72).
- [110] A. Barducci, M. Bonomi, and M. Parrinello. “Metadynamics”. In: *Wiley Interdiscip. Rev. Comput. Mol. Sci.* 1.5 (2011), pp. 826–843. doi: 10.1002/wcms.31 (cit. on p. 72).
- [111] A. P. Willard and D. Chandler. “Instantaneous Liquid Interfaces”. In: *J Phys Chem B* 114.5 (2010), pp. 1954–1958. doi: 10.1021/jp909219k (cit. on pp. 75, 76, 120).
- [112] G. A. Tribello et al. “PLUMED 2: New Feathers for an Old Bird”. In: *Comput. Phys. Commun.* 185.2 (2014), pp. 604–613. doi: 10.1016/j.cpc.2013.09.018 (cit. on pp. 75, 87, 120).
- [113] J. Q. Broughton and G. H. Gilmer. “Molecular Dynamics Investigation of the Crystal-Fluid Interface. I. Bulk Properties”. In: *J. Chem. Phys.* 79.10 (Nov. 15, 1983), pp. 5095–5104. doi: 10.1063/1.445633 (cit. on p. 77).
- [114] S. Plimpton. “Fast Parallel Algorithms for Short-Range Molecular Dynamics”. In: *J Comp Phys* 117 (1995), pp. 1–19 (cit. on p. 77).
- [115] C. A. Becker et al. “Atomistic Simulations of Crystal-Melt Interfaces in a Model Binary Alloy: Interfacial Free Energies, Adsorption Coefficients, and Excess Entropy”. In: *Phys Rev B* 79.5 (Feb. 2009), p. 054109. doi: 10.1103/PhysRevB.79.054109 (cit. on pp. 80–84, 86, 106).
- [116] M. Ceriotti, G. Bussi, and M. Parrinello. “Colored-Noise Thermostats à La Carte”. In: *J. Chem. Theory Comput.* 6.4 (Apr. 13, 2010), pp. 1170–1180. doi: 10.1021/ct900563s (cit. on p. 80).
- [117] B. Cheng, G. A. Tribello, and M. Ceriotti. “The Gibbs Free Energy of Homogeneous Nucleation: From Atomistic Nuclei to the Planar Limit”. In: *J. Chem. Phys.* 147.10 (Sept. 14, 2017), p. 104707. doi: 10.1063/1.4997180 (cit. on pp. 84, 86, 133).
- [118] R. P. Brent. *Algorithms for Minimization Without Derivatives*. Courier Corporation, June 10, 2013. 210 pp. (cit. on p. 90).
- [119] M. Frigo and S. G. Johnson. “The Design and Implementation of FFTW3”. In: *Proc. IEEE* 93.2 (Feb. 2005), pp. 216–231. doi: 10.1109/JPROC.2004.840301 (cit. on p. 93).
- [120] J. W. Cahn and J. E. Hilliard. “Free Energy of a Nonuniform System. I. Interfacial Free Energy”. In: *J. Chem. Phys.* 28.2 (Feb. 1, 1958), pp. 258–267. doi: 10.1063/1.1744102 (cit. on p. 94).
- [121] U. R. Pedersen et al. “Computing Gibbs Free Energy Differences by Interface Pinning”. In: *Phys. Rev. B* 88.9 (Sept. 3, 2013). doi: 10.1103/PhysRevB.88.094101 (cit. on pp. 96, 102, 121).

- [122] U. R. Pedersen. "Direct Calculation of the Solid-Liquid Gibbs Free Energy Difference in a Single Equilibrium Simulation". In: *The Journal of Chemical Physics* 139.10 (Sept. 9, 2013), p. 104102. DOI: 10.1063/1.4818747 (cit. on pp. 96, 121).
- [123] B. Cheng and M. Ceriotti. "Communication: Computing the Tolman Length for Solid-Liquid Interfaces". In: *J. Chem. Phys.* 148.23 (June 21, 2018), p. 231102. DOI: 10.1063/1.5038396 (cit. on pp. 96, 98, 102, 111, 117).
- [124] J. Monk et al. "Determination of the Crystal-Melt Interface Kinetic Coefficient from Molecular Dynamics Simulations". In: *Model. Simul. Mater. Sci. Eng.* 18.1 (2009), p. 015004. DOI: 10.1088/0965-0393/18/1/015004 (cit. on pp. 100, 101, 116).
- [125] Z. T. Trautt, M. Upmanyu, and A. Karma. "Interface Mobility from Interface Random Walk". In: *Science* 314.5799 (2006), pp. 632–635 (cit. on p. 100).
- [126] U. R. Pedersen, F. Hummel, and C. Dellago. "Computing the Crystal Growth Rate by the Interface Pinning Method". In: *J Chem Phys* 142.4 (Jan. 2015), p. 044104. DOI: 10.1063/1.4905955 (cit. on p. 101).
- [127] P. L. Williams, Y. Mishin, and J. C. Hamilton. "An Embedded-Atom Potential for the Cu-Ag System". In: *Model. Simul. Mater. Sci. Eng.* 14.5 (July 2006), pp. 817–833. DOI: 10.1088/0965-0393/14/5/002 (cit. on pp. 102, 106, 107, 114).
- [128] T. Frolov and Y. Mishin. "Solid-Liquid Interface Free Energy in Binary Systems: Theory and Atomistic Calculations for the (110) Cu–Ag Interface". In: *J. Chem. Phys.* 131.5 (Aug. 6, 2009), p. 054702. DOI: 10.1063/1.3197005 (cit. on pp. 102, 112).
- [129] W. Pearson. *A Handbook of Lattice Spacings and Structures of Metals and Alloys*. Vol. 2. Oxford Pergamon Press, 1958 (cit. on p. 103).
- [130] H. Ramalingam et al. "Atomic-Scale Simulation Study of Equilibrium Solute Adsorption at Alloy Solid-Liquid Interfaces". In: *Interface Science* 10.2 (July 1, 2002), pp. 149–158. DOI: 10.1023/A:1015889313170 (cit. on p. 106).
- [131] F. Spaepen. "A Structural Model for the Solid-Liquid Interface in Monatomic Systems". In: *Acta Metallurgica* 23.6 (June 1, 1975), pp. 729–743. DOI: 10.1016/0001-6160(75)90056-5 (cit. on p. 106).
- [132] B. B. Laird. "The Solid–Liquid Interfacial Free Energy of Close-Packed Metals: Hard-Spheres and the Turnbull Coefficient". In: *J. Chem. Phys.* 115.7 (Aug. 6, 2001), pp. 2887–2888. DOI: 10.1063/1.1391481 (cit. on p. 106).
- [133] R. Sibug-Aga and B. B. Laird. "Structure and Dynamics of the Interface between a Binary Hard-Sphere Crystal of NaCl Type and Its Coexisting Binary Fluid". In: *Phys. Rev. B* 66.14 (Oct. 15, 2002), p. 144106. DOI: 10.1103/PhysRevB.66.144106 (cit. on p. 106).
- [134] M. Hansen and K. Anderko. *Constitution of Binary Alloys*. New York: McGraw-Hill, 1958 (cit. on pp. 106, 107).
- [135] R. Hultgren and P. Desai, eds. *Selected Thermodynamics Values and Phase Diagrams for Copper and Some of Its Binary Alloys*. New York: The International Copper Research Association, 1971 (cit. on pp. 106, 107).
- [136] A. Voter and S. Chen. "Characterization of Defects in Materials". In: *MRS Symposia Proceedings*. Ed. by R. Siegel and R. Sinclair. Vol. 82. Materials Research Society, 1978 (cit. on p. 108).

Bibliography

- [137] H. Zhou et al. “Calculation of Solid–Liquid Interfacial Free Energy of Cu by Two Different Methods”. In: *Journal of Crystal Growth* 377 (Aug. 15, 2013), pp. 107–111. DOI: 10.1016/j.jcrysgro.2013.05.007 (cit. on p. 108).
- [138] E. Asadi et al. “Two-Phase Solid–Liquid Coexistence of Ni, Cu, and Al by Molecular Dynamics Simulations Using the Modified Embedded-Atom Method”. In: *Acta Materialia* 86 (Mar. 1, 2015), pp. 169–181. DOI: 10.1016/j.actamat.2014.12.010 (cit. on pp. 108, 111).
- [139] J. J. Hoyt et al. “Unusual Temperature Dependence of the Solid-Liquid Interfacial Free Energy in the Cu-Zr System”. In: *Computational Materials Science* 154 (Nov. 1, 2018), pp. 303–308. DOI: 10.1016/j.commatsci.2018.07.050 (cit. on p. 112).
- [140] D. Chattoraj and K. Birdi. *Adsorption and the Gibbs Surface Excess*. Plenum Press (New York), 1984 (cit. on p. 117).
- [141] R. C. Cammarata. “Generalized Surface Thermodynamics with Application to Nucleation”. In: *Philos. Mag.* 88.6 (Feb. 21, 2008), pp. 927–948. DOI: 10.1080/14786430802014654 (cit. on p. 117).
- [142] H. L. Tepper and W. J. Briels. “Crystal Growth and Interface Relaxation Rates from Fluctuations in an Equilibrium Simulation of the Lennard-Jones (100) Crystal-Melt System”. In: *J. Chem. Phys.* 116.12 (Mar. 12, 2002), pp. 5186–5195. DOI: 10.1063/1.1452110 (cit. on p. 131).
- [143] H. L. Tepper and W. J. Briels. “Crystallization and Melting in the Lennard-Jones System: Equilibration, Relaxation, and Long-Time Dynamics of the Moving Interface”. In: *J. Chem. Phys.* 115.20 (Nov. 13, 2001), pp. 9434–9443. DOI: 10.1063/1.1413972 (cit. on p. 131).
- [144] C. J. Tymczak and J. R. Ray. “Asymmetric Crystallization and Melting Kinetics in Sodium: A Molecular-Dynamics Study”. In: *Phys. Rev. Lett.* 64.11 (Mar. 12, 1990), pp. 1278–1281. DOI: 10.1103/PhysRevLett.64.1278 (cit. on p. 131).
- [145] M. Amini and B. B. Laird. “Kinetic Coefficient for Hard-Sphere Crystal Growth from the Melt”. In: *Phys. Rev. Lett.* 97.21 (Nov. 20, 2006), p. 216102. DOI: 10.1103/PhysRevLett.97.216102 (cit. on p. 131).
- [146] C. F. Richardson and P. Clancy. “Contribution of Thermal Conductivity to the Crystal-Regrowth Velocity of Embedded-Atom-Method-Modeled Metals and Metal Alloys”. In: *Phys. Rev. B* 45.21 (June 1, 1992), pp. 12260–12268. DOI: 10.1103/PhysRevB.45.12260 (cit. on p. 131).
- [147] M. Hijazi, D. M. Wilkins, and M. Ceriotti. “Fast-Forward Langevin Dynamics with Momentum Flips”. In: *J. Chem. Phys.* 148.18 (May 10, 2018), p. 184109. DOI: 10.1063/1.5029833 (cit. on pp. 132–134).

Edoardo Baldi

Seidenstrasse 24
8600 Dübendorf
Switzerland
☎ +41 78 662 68 00
✉ edrd.baldi@gmail.com
🌐 edoardob90

Education

- 2015–2019 **Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland.**
Doctor of Science in Computational Materials Science and Engineering
- Atomistic simulations of solid–liquid interfaces of metals and alloys
- 2012–2015 **University of Milano–Bicocca, Milano, Italy.**
Master of Science in Physics
- Atomistic simulations of crystallization process in nanowires for phase change memories
- 2009–2012 **University of Milano–Bicocca, Milano, Italy.**
Bachelor of Science in Materials Science
- Vibrational and thermal properties of a phase change material from first principles simulations

Core experience

- 2019– **Swiss Federal Laboratories for Materials Science and Technology (Empa), Dübendorf, Switzerland.**
HPC expert in the Scientific IT unit

Additional experience

- 2015–Present **IT manager of the Laboratory of Computational Science and Modelling (EPFL), Lausanne, Switzerland.**
Duties:
- basic support for efficient use of EPFL’s HPC resources (including debugging and optimization of software developed within the lab)
 - installation, customization and maintenance of the lab’s workstations
- 2015–2018 **Teaching assistant, EPFL, Lausanne, Switzerland.**
Assistant for the “Surfaces and interfaces” course of third year of the degree in Materials Science and Engineering.
- July 2017 **Summer school on “Effective high-performance computing”, CSCS–USI, Lugano, Switzerland.**
Topics:
- message passing programming model (MPI standard)
 - shared memory model (OpenMP)
 - GPU programming with CUDA and OpenACC
 - effective data I/O
 - scientific libraries
 - code optimization

January 2017 **Winter school on “Additive Manufacturing in metals”**, CCMX, Kandersteg, Switzerland.

The school was organized and supported by the Competence Centre for Materials Science and Technology (CCMX, EPFL).

June 2015 **24th summer school on Parallel Computing**, CINECA, Bologna, Italy.

Topics:

- introduction to high-performance computing and parallel architectures
- parallel programming models and methods (MPI, OpenMP, GPU programming basics)
- code profiling and optimization
- debugging

May 2013 **Workshop on “Introduction to parallel computing with MPI and OpenMP”**, CINECA, Milano, Italy.

Introductory course on parallel computing in scientific programming.

April 2013 **Workshop on “Introduction to Fortran 90”**, CINECA, Milano, Italy.

Introduction to modern paradigms of the Fortran programming language with applications to scientific calculus.

Extra-curricular activities

2016–2018 **Treasurer and committee member of the AGEPoly Ciné-club**, EPFL.

Active member of the EPFL’s students association dedicated to promoting the film culture on the campus.

Technical skills

Programming languages C/C++, Python, Fortran, CUDA

Experience with OpenMP, MPI, OpenACC, profiling and debugging tools (perf, gdb, Intel software), CVS with Git, Unix shell scripts, GNU/Linux OS

Other software Wolfram Mathematica, \LaTeX

Languages

English Fluent spoken (C1) and written (C2)

Cambridge FCE obtained in July 2008

French Intermediate spoken and written (B1)

Two intensive language courses attended at the Language Centre (EPFL)

Italian Native language

Personal informations

Date of birth 10 march 1990

Driving license B

Nationality Italian

Military obligations No

Marital status Single

Publications

E. Baldi, M. Ceriotti, and G. Tribello. Extracting the interfacial free energy and anisotropy from a smooth fluctuating dividing surface. *Journal of Physics: Condensed Matter*, 29(44):445001, 2017.

D. Campi, E. Baldi, G. Graceffa, G. Sosso, and M. Bernasconi. Electron-phonon interaction and thermal boundary resistance at the interfaces of $\text{Ge}_2\text{Sb}_2\text{Te}_5$ with metals and dielectrics. *Journal of Physics: Condensed Matter*, 27(17), 2015.

S. Gabardi, E. Baldi, E. Bosoni, D. Campi, S. Caravati, G. C. Sosso, J. Behler, and M. Bernasconi. Atomistic Simulations of the Crystallization and Aging of GeTe Nanowires. *Journal of Physical Chemistry C*, 121(42):23827–23838, 2017.