LEIPZIG UNIVERSITY

FACULTY OF PHYSICS AND EARTH SCIENCES
INSTITUTE OF THEORETICAL PHYSICS
COMPUTATIONAL QUANTUM FIELD THEORY

MASTER THESIS

Two Perspectives on the Condensation-Evaporation Transition of the Lennard-Jones Gas in 2D

---

submitted by:
Franz Paul Spitzner

Supervisors:
Prof. Dr. Wolfhard Janke
Dr. Johannes Zierenberg

A thesis submitted in partial fulfillment of
the requirements for the degree of

Master of Science (M. Sc.)
in Physics (IPSP)

June 4, 2017

Declaration of Authorship

I, Franz Paul Spitzner, declare that this thesis titled, "Two Perspectives on the Condensation-Evaporation Transition of the Lennard-Jones Gas in 2D" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

Contents

List of Figures

1 Introduction

In both, science and industry, phase transitions play an important role; especially the condensation-evaporation transition, also studied in this thesis, has shown to be of long-lasting interest [1]. One may say it is one of the prime examples when talking about nucleation theory, which itself is a vast and well established field of Statistical Physics [2–4]. Consequently, current research involving this or similar phase transitions is as multifaceted as it is meticulously scrutinised by the scientific community.

Exemplary, objectives range from increasingly sophisticated procedures such as surface growth, where condensation is an underlying process [5–7], to theoretical reconsiderations of the universality class of the transition [8]. However, the most apparent applications evolving, amongst others, around our transition are weather models and forecasts, the long history and economic relevance of which are indisputable. Not only due to the historic evolution but also because of their complexity, many of the involved processes, such as nucleation rates, are modelled via efficient parametrisations [9], some of which are heuristically obtained. While this is a justified approach, the optimisation of parameters and microphysics still requires a bottom-up understanding.

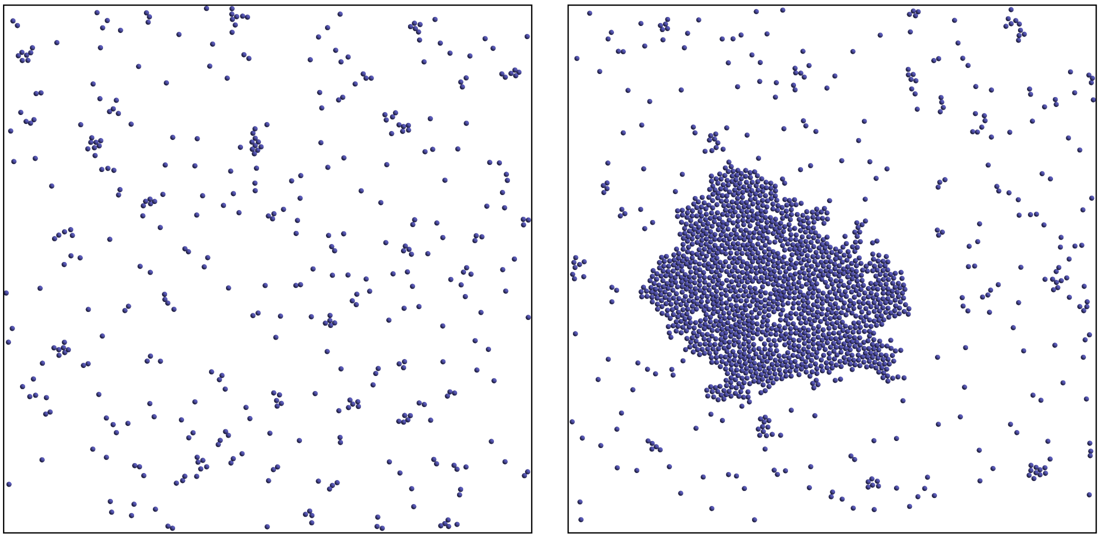

FIGURE 1.1: Snapshots of the two phases we are interested in: a homogeneous gas phase on the left and an inhomogeneous phase featuring a liquid droplet with surrounding gas on the right. Note that in the chosen, simplified model, particles have no actual volume or mass. Thus, the size of the point-particles in the figure is arbitrary.

This leads us exactly to that one thing physicists love to do: describing a complex process or phenomenon through a basic model to focus on essentials. So, in order to look at the condensation-evaporation transition, we employ a Lennard-Jones gas that consists of a short-range interaction potential acting between massless point-particles as a representation of atoms or molecules. As illustrated in Fig. 1.1, we can then distinguish between a homogeneous gas phase (left) and a mixed phase consisting of a liquid condensate with surrounding vapour (right). Despite having simplified the problem significantly, we can only simulate systems of limited complexity and *size* in computers, even with today's ever more capable hardware. For this reason we are going to apply a finite-size scaling analysis [10–12], which allows to predict the macroscopic behaviour (corresponding to an infinitely large system) from a number of finite but differently sized systems.

In our context, two formulations to obtain the finite-size scaling have been established, both of which yield the same leading order behaviour. One of them was originated by Binder et al. [13–18] and works by identifying the transition through the chemical potential, where the surface free energy is minimised during the derivation.

The other formulation, presented by Biskup et al. [19, 20] utilises similar assumptions (including free-energy contributions due to the condensate and gas) but uses the influence of the droplet size on its probability to introduce an excess fraction that describes the amount of excess particles within the single equilibrium droplet, if it exists at all. Thereby, an explicit treatment of the chemical potential is not necessary. The charm of this approach is its frugality; a generic although genuine derivation is presented, which suffices with *bare necessities*. By only considering local density fluctuations in the gas, as well as the surface contribution of the droplet, a self-consistent formalism arises without specific requirements on the model. It *feels* truly coherent.

We will proceed with this thesis by giving an overview of our used models and methods. Within the theory section, we will then follow Biskup et al. and recapitulate the formalism, which has already been verified for the discrete lattice gas [21–23]. In the given references, Monte Carlo simulations were used along with the equivalence between lattice gas and Ising model [24], but we will stay in continuous space and validate the theory for the Lennard-Jones fluid. To that end, grand canonical Monte Carlo (GCMC) simulations are performed, which allow us to evaluate density fluctuations. Various similar investigations involving the grand canonical ensemble have been conducted by Wilding et al. [25–29], although with a focus on the behaviour near the critical point.

So far, all mentioned references and considerations took place in what one could call the *native regime*. Therein, temperature is assumed to be fixed at all times and the phase transition is driven by density; we expect a known particle excess, and the remaining question is whether it contributes to local density fluctuations in the gas or a liquid condensate is formed. When keeping one's focus on density and temperature, a second, *orthogonal regime* can be studied, as it was done by Zierenberg and co-workers [30–34]. In this regime, the transition is driven by temperature while keeping the density fixed. The according leading-order scaling of the transition temperature was analytically derived based on Biskup's formalism and cross-checked with multicanonical simulations. Since data was only created for the 2D and 3D lattice gas as well as the 3D Lennard-Jones gas, this leaves a vacant spot to fill and provides a great starting point to have a closer look at the Lennard-Jones fluid in two dimensions.

All in all, we have two goals in this thesis. Firstly, we want to fill the gap just mentioned and possibly give a comparison between the results for the different models and dimensions. Since the simulations of the fixed-density scheme will take place in the (multi-)canonical ensemble, this should be a diligent but rather routine piece of work. Having said that, it opens the door for our second concern; once we manage to also create a similar set of data in the fixed-temperature scheme, we can provide a truly well-balanced look at the condensation-evaporation transition. In an ideal scenario, this includes the respective theory, Monte Carlo simulations and subsequent finite-size scaling analysis for both regimes.

2 Model and Methods

2.1 The Lennard-Jones Gas

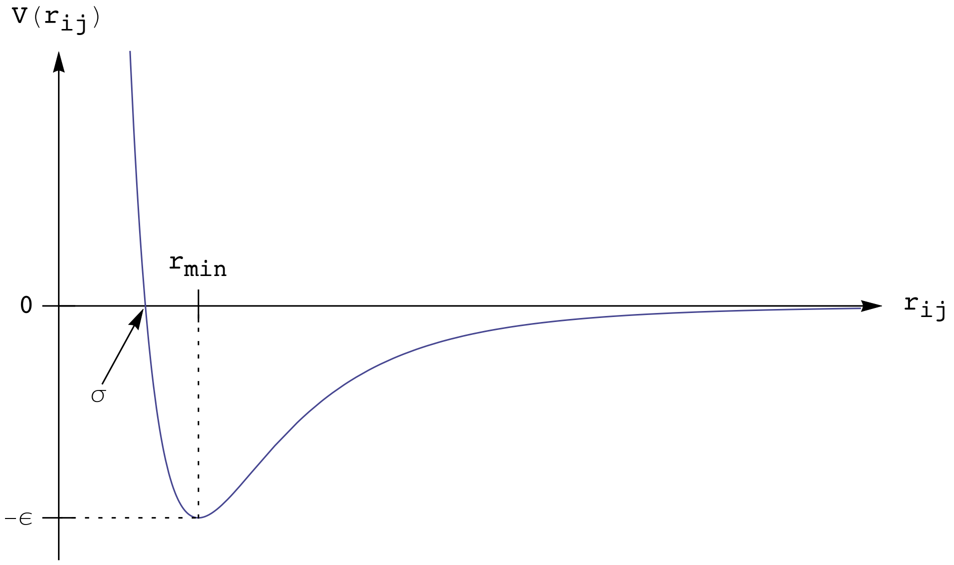

FIGURE 2.1: Schematic of the Lennard-Jones potential. The characteristic length and energy scale are set by σ and ϵ, respectively.

This thesis focuses on the Lennard-Jones gas, a continuous particle model. The according potential, in the literature also referred to as Lennard-Jones or 12-6 potential, describes the interactions between individual particles:

$$V(r_{ij}) = 4\epsilon\left(\left(\frac{\sigma}{r_{ij}}\right)^{12} - \left(\frac{\sigma}{r_{ij}}\right)^{6}\right) , \tag{2.1}$$

where r_{ij} is the distance between the i-th and j-th particle. We see that the energy scale is given via ϵ, while σ sets the characteristic length scale as illustrated in Fig. 2.1. Furthermore, the potential may be expressed in terms of $r_{\mathrm{min}} = 2^{1/6}\sigma$, the distance at which energy is minimised

$$V(r_{ij}) = \epsilon\left(\left(\frac{r_{\mathrm{min}}}{r_{ij}}\right)^{12} - 2\left(\frac{r_{\mathrm{min}}}{r_{ij}}\right)^{6}\right) . \tag{2.2}$$

In order to take all particles of the system into account, we have to calculate the pairwise interaction and include all pairs to obtain the Hamiltonian

$$\mathcal{H} = \frac{1}{2}\sum_{i\neq j} V(r_{ij}) . \tag{2.3}$$

The Lennard-Jones potential is only one part of the description of the system we are going to investigate. We also assume a number of massless and dimensionless particles in a two dimensional simulation volume. Moreover, we look at particle numbers for which the fluid is sufficiently dilute. In other words, we adjust the particle number per volume or particle density to be somewhere in the region of typical gas and liquid phases. While this may seem trivial, it has some important implications; since we do not take solid configurations into account, geometric constraints such as optimal packing ratio do not play a role. Besides, there is a typical scenario where one particle will only have very few other particles close to itself and the majority is rather far away.

2.1.1 Domain Decomposition

Without any modification, the interactions due to the potential have infinite range but are formally short range. The immediate consequence is that we have to do computations of order N^2 for a system containing N particles. In addition, we have to calculate even the tiniest of interactions for particles very far away, which on average, do not contribute much to the total energy. Long story short, in most simulations done today a cutoff radius r_c is introduced, beyond which the potential is ignored; the so called domain decomposition is then a useful tool to further speed up the computation. It is shown schematically in Fig. 2.2. In the left sketch, one has to calculate interactions between the blue particle of interest with all other particles, no matter how far away they are. On the right-hand side, the whole volume is split into boxes that are sized according to the cutoff radius. Hence, to calculate the energy due to the blue particle, it is sufficient to consider the black particles located nearby within neighbouring boxes.

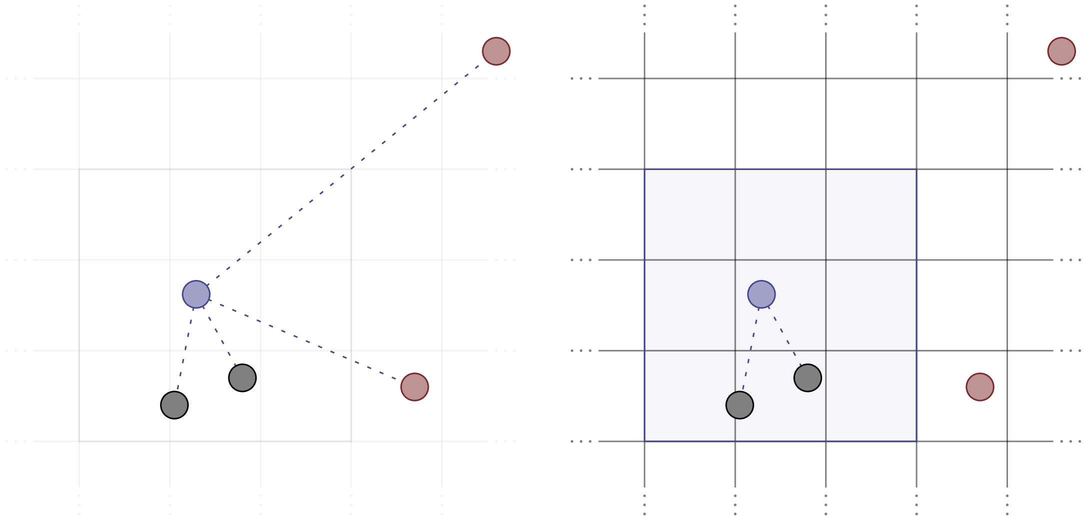

FIGURE 2.2: Comparison of required interactions for a (blue) particle of interest: without decomposition, all particles have to be considered (left sketch). When the domain decomposition is used (right sketch), only particles within the neighbouring domains need to be included into calculations.

When limiting the interaction range, the potential has to be modified. One way is to just truncate for radii larger than cutoff radius, which will cause an actual jump in the value of the potential when crossing said radius. The straight forward way to avoid this is to shift the original potential by its cutoff value. Both situations are illustrated in Fig. 2.3. Subsequently, one has to decompose the volume. Wilding [25, 26], for instance, chose to define the system size in integer multiples of the cutoff radius, thus, domain length and r_c are precisely equal. This integer division was then understood as a reduced unit for the length scale. However, we wanted to employ real numbers as the system size, therefore, our domain length was determined at runtime to be as small as possible but still larger than the cutoff radius.

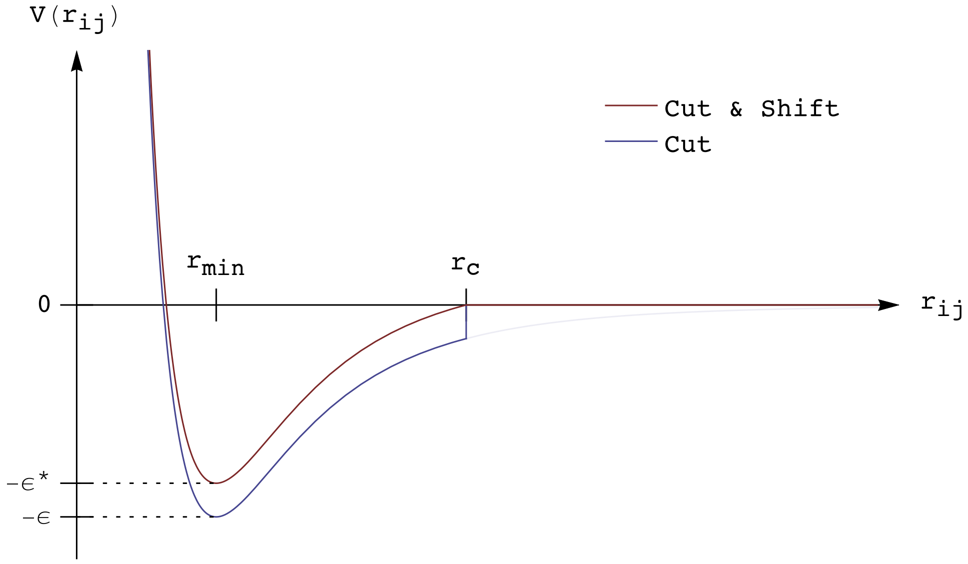

FIGURE 2.3: Illustration of two different ways to limit the interaction range of the potential. Other methods exist to smooth the discontinuity at r_c (for Molecular Dynamic simulations) but were not needed for our Monte Carlo based studies.

A further complication when comparing results with aforesaid reference is that a truncated but not shifted potential was used there. This leads to systematic deviations for most observables. Apart from test runs to verify the program with references, our simulations were performed using the following truncated and shifted potential:

$$V^*(r_{ij}) = \begin{cases} V(r_{ij}) - V(r_c) & r_{ij} < r_c = 2.5\sigma \\ 0 & \text{else} \end{cases} . \tag{2.4}$$

2.2 Monte Carlo Simulations

When talking about computer simulations utilising Monte Carlo methods, it is beneficial to recapitulate the most essential underlying concepts. As noble and well intended as this sounds, one walks a thin line between delivering necessary background and getting lost in detail. Hence, we shall, while referring to the literature, favour verbal summary over explicit mathematical formulation in this section.

Since Statistical Physics deals with many-particle systems, one needs to distinguish between microscopic and macroscopic descriptions; while the former treat individual particles or the interaction between them, think equations of motion, the latter deal with properties of the system as a whole, e.g. its specific heat. One can measure a number of such observables, or the *macroscopic state* of the system through a real experiment. The *microscopic state*, on the other hand, is an exact configuration of the system, corresponding to a single point in phase space, which in turn is spanned by the generalised momenta $\bar{p}$ and positions $\bar{q}$ of all particles. For an N-particle system in three dimensions, phase space is consequently $6N$ dimensional and using the Hamiltonian $\mathcal{H}(\bar{p}, \bar{q})$ one can calculate the energy of any such configuration. One can also track the trajectory of the system through phase space via the equations of motion in order to obtain an average over time of a given observable, say $\mathcal{O}$.

However, there is often no way to know the exact configuration at a given time. Instead, we focus on the macroscopic state and consider a large number of systems; specifically, we look at an *ensemble* of systems in all possible microstates but require every system to occupy the same macroscopic state. Furthermore, taking the average over the ensembles at set time instead of taking the time average of a single system will yield the same result. This is called the *ergodic* hypothesis; it is also commonly quoted by saying that all points in phase space are accessible to a given system. Now, what is that probability to be at a certain point? Let us denote the particular point in phase space or microstate with $\{\phi\}$ to specify all the required coordinates. The probability of the microstate or the *configuration weight* is then $W_X(\{\phi\})$, with a probability density

$$P_X(\{\phi\}) = \frac{1}{\mathcal{Z}_X} W_X(\{\phi\}) \; . \tag{2.5}$$

Here X denotes the chosen ensemble, often expressed through its natural variables, for instance, NVT. The *partition function* $\mathcal{Z}_X$ is the integral over all the weights over the whole phase space:

$$\mathcal{Z}_X = \int_{\{\phi\}} \mathrm{d}\{\phi\} \; W_X(\{\phi\}) \; . \tag{2.6}$$

We want to point out that the actual configuration weight is dependant upon the chosen ensemble, and it is not just an intrinsic property of the system. Further, the partition function, while seemingly simple deserves additional notice, as it is often times straight forward to write down but impossible to solve.

Using the introduced formalism, we can identify the ensemble average of the observable $\mathcal{O}$ with the expectation value (denoted by the triangular brackets)

$$\langle \mathcal{O} \rangle_X = \frac{1}{\mathcal{Z}_X} \int_{\{\phi\}} \mathrm{d}\{\phi\} \; \mathcal{O}(\{\phi\}) \; W_X(\{\phi\}) = \int_{\{\phi\}} \mathrm{d}\{\phi\} \; \mathcal{O}(\{\phi\}) \; P_X(\{\phi\}) \; . \tag{2.7}$$

2.2.1 Importance Sampling

As we just mentioned, in Monte Carlo simulations one deals with ensemble averages over a series of measurements taken from different configurations of the studied system. Those measurements are called a *time series*, somewhat contradictory to the abandoned time evolution. Nonetheless, they are snapshots of the systems evolution through phase space. This probabilistic evolution is called a *Markov chain*: the system undergoes a large number of Monte Carlo moves, each of which attempting to change the system from one configuration $\{\phi\}$ to another one, say $\{\phi^\star\}$. This could be done using *simple sampling*, where configurations are drawn completely randomly. What we want to do instead is to draw new configurations from the a priori unknown probability distribution $P(\{\phi\})$ as of Eq. (2.5). Thereby, using *importance sampling*, microstates shall get chosen according to their weight and a transition has the probability

$$p(\{\phi\} \to \{\phi^\star\}) \; , \tag{2.8}$$

which is only dependant upon the two involved configurations $\{\phi\}$ and $\{\phi^\star\}$. Additionally, we want the updates to be ergodic and the system to have no memory of previous configurations. When looking at systems in equilibrium, the *master equation*

$$\frac{\mathrm{d}}{\mathrm{d}t}P(\{\phi\},t) = \int_{\{\phi^\star\}} \mathrm{d}\{\phi^\star\} \left[P(\{\phi\},t)\, p(\{\phi\} \to \{\phi^\star\}) - P(\{\phi^\star\},t)\, p(\{\phi^\star\} \to \{\phi\})\right] \tag{2.9}$$

is time independent

$$\frac{\mathrm{d}}{\mathrm{d}t}P(\{\phi\},t) = 0\ , \tag{2.10}$$

thus,

$$\int_{\{\phi^\star\}} \mathrm{d}\{\phi^\star\}\ P(\{\phi\})\, p(\{\phi\} \to \{\phi^\star\}) = \int_{\{\phi^\star\}} \mathrm{d}\{\phi^\star\}\ P(\{\phi^\star\})\, p(\{\phi^\star\} \to \{\phi\})\ . \tag{2.11}$$

One possible solution is the so called *detailed balance* condition

$$\frac{P(\{\phi\})}{P(\{\phi^\star\})} = \frac{p(\{\phi^\star\} \to \{\phi\})}{p(\{\phi\} \to \{\phi^\star\})}\ . \tag{2.12}$$

Equivalently, we could say that the system has the same probability to go from $\{\phi\}$ to $\{\phi^\star\}$ and vice versa. In order to implement this criterion, we usually consider the *total* transition probability $p(\{\phi\} \to \{\phi^\star\})$ as the product of the probability $s(\{\phi\} \to \{\phi^\star\})$ to *suggest* a new configuration and the probability $a(\{\phi\} \to \{\phi^\star\})$ to *accept* it. Using this ansatz, we can solve Eq. (2.12) by

$$a(\{\phi\} \to \{\phi^\star\}) = \min\left(1,\ \frac{P(\{\phi^\star\})\ s(\{\phi^\star\} \to \{\phi\})}{P(\{\phi\})\ s(\{\phi\} \to \{\phi^\star\})}\right)\ , \tag{2.13}$$

which is often called the *acceptance test* or acceptance ratio. Hinting at the following pages, we want to mention that the suggestion probabilities can either be symmetric or antisymmetric, the latter ones leading to different acceptance tests for different Monte Carlo moves.

As a consequence of importance sampling, we may express the expectation value of an observable in Eq. (2.7) differently than through the integral over the ensemble. Actually, we already stressed out that configurations are drawn following the probability distribution $P(\{\phi\})$. Hence, we have probability $P(\{\phi\})$ to measure the observable $\mathcal{O}(\{\phi\})$ that belongs to the configuration, and it makes sense to approximate the expectation value by the mean of all n measurements taken

$$\langle \mathcal{O} \rangle_X \approx \frac{1}{n}\sum_{i=1}^{n} \mathcal{O}_i\ . \tag{2.14}$$

For later use when comparing data between ensembles, one has to remember that $\mathcal{O}_i$ was measured in ensemble X and implicitly contains the configuration weight W_X of that ensemble.

2.2.2 Canonical Ensemble and Neglecting Momenta

The canonical ensemble is used to characterise a closed system that is in contact with a much larger system, that in turn functions as a heat bath. The transfer of energy shall be permitted, while particle number, volume and temperature are constant, hence, the canonical ensemble is often referred to as NVT ensemble. One can easily see its relevance, justified by the similarity to real life experiments, for which it is far easier to control temperature, than internal energy, for instance. Using another ensemble, i.e. the very fundamental microcanonical NVE ensemble, one can show [35, 36] that the configuration weight for the NVT ensemble is the Boltzmann weight $e^{-\beta\mathcal{H}}$, where $\beta = 1/k_{\mathrm{B}}T$ is the inverse temperature. Unless explicitly stated otherwise, we set $k_{\mathrm{B}} = 1$ in this thesis.

Expressed through the generalised phase space coordinates $\{\phi\}$ the partition function then reads

$$\mathcal{Z}_{\mathrm{NVT}} = \int_{\{\phi\}} \mathrm{d}\{\phi\}\; e^{-\beta\mathcal{H}(\{\phi\})} \;, \tag{2.15}$$

which is of course equivalent to doing the integration for momenta $\bar{p}$ and position $\bar{q}$ explicitly

$$\mathcal{Z}_{\mathrm{NVT}} = \int \mathrm{d}\bar{p} \int \mathrm{d}\bar{q}\; e^{-\beta\mathcal{H}(\bar{p},\bar{q})} \;. \tag{2.16}$$

This rewriting allows us to highlight an important trick played in most Monte Carlo simulations; they do not treat momenta. When looking at Eq. (2.16) we see that for a Hamiltonian that allows to separate momenta and positions, such as for any classical system in d dimensions, the integrations can be done independently

$$\begin{aligned}
\mathcal{Z}_{\mathrm{NVT}} &= \int \mathrm{d}\bar{p} \int \mathrm{d}\bar{q}\; \exp\left[-\beta\left(\sum_{i=1}^{N}\frac{\bar{p}_i^2}{2m} + V(\bar{q})\right)\right] \\
&= \int \mathrm{d}\bar{p}\; \exp\left[-\beta\sum_{i=1}^{dN}\frac{p_i^2}{2m}\right] \int \mathrm{d}\bar{q}\; e^{-\beta V(\bar{q})} \\
&= (2\pi m k_{\mathrm{B}}T)^{dN/2} \int \mathrm{d}\bar{q}\; e^{-\beta V(\bar{q})} \;.
\end{aligned} \tag{2.17}$$

When the observables contributing to an expectation value are no function of momentum, the same trick can be played and the factor cancels. However, one has to be careful: for observables that depend on the distribution of energy, such as involved in free-energy barriers, an explicit treatment is necessary [33]. Nevertheless, we shall from now on only consider the configuration part of phase space, unless explicitly noted otherwise. The configurational partition function of the canonical ensemble is

$$\hat{\mathcal{Z}}_{\mathrm{NVT}} = \int \mathrm{d}\bar{q}\; e^{-\beta V(\bar{q})} \;, \tag{2.18}$$

where the circumflex in the notation (indicating the difference to the full canonical partition function) will be neglected. The according probability density is given by

$$P_{\mathrm{NVT}}(\bar{q}) = \frac{1}{\mathcal{Z}_{\mathrm{NVT}}} e^{-\beta V(\bar{q})} \;. \tag{2.19}$$

Since we have effectively reduced phase space, it seems reasonable to revisit the notation to express microstates via $\{\phi\}$. Clearly, such a notation does not allow to distinguish if momenta are included or left out. While this seems disadvantageous, it is a handy feature; as $\{\phi\}$ contains *all* coordinates that matter in the chosen context of ensemble and phase space, it enables general expressions for partition functions and states to stay well arranged and to focus on individual coordinates, whilst others of less interest are neatly grouped.

Let us take a look at the acceptance ratio. The typical Monte Carlo move in this ensemble is the displacement of one of the many particles in the system. Clearly, if a random particle is chosen and moved to a random position, the suggestion probability is symmetric, i.e.

$$s(\{\phi\} \to \{\phi^\star\}) = s(\{\phi^\star\} \to \{\phi\}) \, . \tag{2.20}$$

Since the probability density is just the Boltzmann weight $P(\{\phi\}) \propto \exp\left[-\beta V(\{\phi\})\right]$, we can now simplify Eq. (2.13) to

$$a(\{\phi\} \to \{\phi^\star\}) = \min\left(1, \, e^{-\beta \Delta E}\right) \, . \tag{2.21}$$

Here $\Delta E = V(\{\phi^\star\}) - V(\{\phi\})$ is the difference in potential after and before the move and we have just obtained the acceptance criterion of the METROPOLIS algorithm [37].

2.2.3 Density of States

Having discussed the canonical ensemble, this is a good moment to talk about another concept: the density of states. Starting again with phase space, in which a configuration is described by a set of variables $\{\phi\}$, we can calculate the energy of such a state with the Hamiltonian. We further employ the *postulate of equal weights*, which states that for isolated systems in equilibrium, such as considered here, all configurations of equal energy have the same probability. Under this assumption, one can hypothetically group and *count* all the configurations with equal energies up to some infinitesimal deviation $E \leq \mathcal{H}(\{\phi\}) \leq E + \delta E$.

This number is what is called the *density of states* $\Omega(E)$; it is the amount of phase space configurations with potential energy E. Often, one finds the expression for the canonical partition function involving the the density of states as

$$\mathcal{Z}_{\mathrm{NVT}} = \int \mathrm{d}E \; \Omega(E) \; e^{-\beta E} \; , \tag{2.22}$$

where instead of integrating over the configurational phase space $\{\phi\}$, one now integrates over the possible energy range. Of course, this can be done in any ensemble where states can be expressed through their potential energy. The general formulation of a partition function for ensemble X in terms of the density of states reads

$$\mathcal{Z}_X = \int \mathrm{d}E \; \Omega(E) \; W_X(E) \; . \tag{2.23}$$

This formulation assumes the configuration weight to be only dependent on energy. If this also holds for an observable of interest, i.e. $\mathcal{O}(\{\phi\}) = \mathcal{O}(E)$ with $E = E(\{\phi\})$ we can calculate the expectation value of that observable using the density of states

$$\langle \mathcal{O} \rangle_X = \frac{1}{\mathcal{Z}_X} \int \mathrm{d}E \; \Omega(E) \; \mathcal{O}(E) \; W_X(E) \; . \tag{2.24}$$

From here we can identify the energy probability distribution

$$P_X(E) = \frac{1}{\mathcal{Z}_X} \Omega(E) W_X(E) \ , \tag{2.25}$$

which completes the transition. We went from treating a system through its microstates to a macroscopic representation, where the state variable is energy and the link between the macroscopic and the microscopic scale is made through the density of states. As a concluding remark we want to note that the density of states itself fits the description of a partition function, specifically that of the configurational NVE ensemble.

2.2.4 Grand Canonical Ensemble

Imagine a large volume in which a simulation in a canonical ensemble takes place with fixed particle number. So, basically, particles are moving around in a large, closed box. If we now divide the box into smaller ones, we will have a varying particle number in each sub-volume. This very idea has actually been implemented to evaluate varying local densities while keeping the total density fixed [38, 39]. Having said that, it would not make sense to run a simulation of the smaller box using the canonical ensemble. Thus, let us use this as a motivation for the grand canonical ensemble: it enables us to look at particle and density fluctuations. The according (so called GCMC) simulations are especially useful for studying inhomogeneous systems [36].

Instead of a constant particle number N as in the canonical case, we fix the *chemical potential*, yielding the ensemble description μVT. The chemical potential quantitatively describes how favourable it is to add another particle into the simulation box and thereby implicitly influences the expected particle number. Since N is not a constant, another dimension in phase space has to be considered. We adjust the variables describing the microstates to $(N, \{\phi_N\})$, where $\{\phi_N\}$ includes the configuration of all the N particles currently present. Minimising entropy under the constraints given by the ensemble, one can show [35] that

$$W_{\mu\mathrm{VT}}(N, \{\phi_N\}) = e^{-\beta V(\{\phi_N\})} e^{\beta \mu N} \tag{2.26}$$

is the configuration weight in the grand canonical ensemble. The partition function is

$$\mathcal{Z}_{\mu\mathrm{VT}} = \sum_{N=0}^{\infty} \int_{\{\phi_N\}} \mathrm{d}\{\phi_N\} \ e^{-\beta V(\{\phi_N\})} e^{\beta \mu N} \ , \tag{2.27}$$

which can also be related to the canonical partition function:

$$\mathcal{Z}_{\mu\mathrm{VT}} = \sum_{N=0}^{\infty} \mathcal{Z}_{\mathrm{NVT}} \ e^{\beta \mu N} \ . \tag{2.28}$$

We have already taken into account that discrete particle numbers imply summation instead of integration. The probability distribution reads

$$P_{\mu\mathrm{VT}}(N, \{\phi_N\}) = \frac{1}{\mathcal{Z}_{\mu\mathrm{VT}}} e^{-\beta V(\{\phi_N\})} e^{\beta \mu N} \ . \tag{2.29}$$

Furthermore, we want to give the acceptance ratio for update moves. Moving a particle around within the box is the same as for a canonical simulation, but we have to be more careful when *inserting* particles into the system or *removing* them. The suggestion probability of an insertion (for which the particle number changes from $(N) \to (N+1)$ and the spacial configuration changes from $\{\phi\} \to \{\phi^\star\}$) is just

$$s_+(\{\phi\} \to \{\phi^\star\}) = \frac{1}{V} \,. \tag{2.30}$$

We randomly pick a position in the simulation volume V and place the particle. Admittedly, Eq. (2.30) is unphysical in the sense that one should treat infinitesimal volume increments. Anyhow, this is not possible due to numerical precision limits in computers and the derived acceptance criterion turns out the same. For the particle deletion from $(N+1) \to (N)$ with $\{\phi^\star\} \to \{\phi\}$, the suggestion probability is

$$s_-(\{\phi^\star\} \to \{\phi\}) = \frac{1}{N+1} \,, \tag{2.31}$$

because we have to select one particle to remove, independently of its position. If we now require detailed balance as of Eq. (2.12), we see that

$$\frac{a_+(\{\phi\} \to \{\phi^\star\})}{a_-(\{\phi^\star\} \to \{\phi\})} = \frac{s_-(\{\phi^\star\} \to \{\phi\})}{s_+(\{\phi\} \to \{\phi^\star\})} \frac{P(N+1, \{\phi^\star\})}{P(N, \{\phi\})} \,, \tag{2.32}$$

hence,

$$\frac{a_+(\{\phi\} \to \{\phi^\star\})}{a_-(\{\phi^\star\} \to \{\phi\})} = \frac{V}{N+1} \frac{e^{-\beta V(\{\phi^\star\})} e^{\beta\mu(N+1)}}{e^{-\beta V(\{\phi\})} e^{\beta\mu N}} = \frac{V}{N+1} e^{-\beta\Delta E + \beta\mu} \,. \tag{2.33}$$

As earlier, employing the expression for the energy difference before and after the update move $\Delta E = V(\{\phi^\star\}) - V(\{\phi\})$, we arrive at the acceptance criterion for the insertion

$$a_+(\{\phi\} \to \{\phi^\star\}) = \min\left(1, \frac{V}{N+1} e^{-\beta\Delta E + \beta\mu}\right) \,. \tag{2.34}$$

Following similar reasoning we get the acceptance test for a deletion, i.e. going from $(N) \to (N-1)$ and the now different target configuration $\{\phi\} \to \{\phi'\}$ with ΔE adjusted accordingly

$$a_-(\{\phi\} \to \{\phi'\}) = \min\left(1, \frac{N}{V} e^{-\beta\Delta E - \beta\mu}\right) \,. \tag{2.35}$$

2.2.5 Update Moves

When we talk about update moves in the context of Monte Carlo simulations, we refer to the complete process of changing the system from one configuration to another. That includes the following tasks: first, a possible new configuration, depending only on the old one, is proposed. Secondly, the program has to evaluate the necessary changes in variables needed to check the acceptance criterion; the check itself is then done by comparing to a (pseudo) random number, which is drawn from a uniform distribution $[0, 1)$. Depending on the outcome, the new or old configuration is kept. Having said that, such an update move can be realised in a variety of ways, from flipping a single spin in the Ising model, to decreasing the volume of an isobaric simulation, or to very complex cluster updates, that change multiple spins as a group.

Across the simulations in this thesis, a total of four different update moves were implemented. Two of them are particle displacement moves with symmetric suggestion probabilities, while the other two, specifically the insertion and its reverse, the deletion of a particle, require different acceptance criteria due to their different probabilities to be suggested.

The first move is a simple particle *shift*. One of the present particles is selected using a random number and moved by a small amount in a random direction. This *small* distance is also determined randomly but so that it is not larger than a given threshold, which was eventually chosen to be 0.8σ. While the prefactor was obtained rather heuristically, to optimise performance, it is reasonable to adjust the range according to the characteristic length scale σ. Intuitively, the particle shift seems to be the most natural update move in the sense that particles only move small steps at a time.

In contrast, the second displacement move (a *jump*) also picks a random particle but attempts to place it in a completely random position within the simulation volume, independently of the original position. For both displacement moves, the energy due to the selected particle has to be calculated twice; once for the old configuration and once for the new one, hence the effort is proportional up to at most $2N$. Due to the nature of this move, it is most efficient in sampling the gas phase rather than the liquid phase.

Third in line, but only of interest for the grand canonical simulations, is the *insertion*. A position is selected and the interaction a particle would have in the position is calculated. Note that because no particle was present before, we only have to calculate N interactions this time. Now, the particle number N also influences the likelihood to be accepted according to the chemical potential. We discovered quite late that it is *imperative* to not only choose the coordinates of the particle randomly, but also to insert it in a random position into the list of present particles. Methodically adding particles to the end of the memory causes a systematic deviation when reweighting to the canonical ensemble. The reason is that the deletion move would tend to select *old* particles more often for possible deletion than freshly inserted ones. For more details and an illustration of this particular issue, please see Sec. 4.1.

The *deletion* move simply does the opposite; it selects some of the present particles with probability $1/N$ and calculates its interaction energy, which is subtracted if the deletion move is accepted.

Both displacement moves were tested in all simulations, employing different ratios to choose which move to propose. It turned out that values of measured observables were not influenced. Nonetheless, the choice which move to propose matters, as the acceptance ratios and the time after which the weights for MUCA and MUGC converged were strongly dependant upon the suggestion ratio. The two algorithms are outlined in the following section.

2.3 Advanced Simulation Techniques

A common obstacle when dealing with first-order phase transitions, such as the condensation-evaporation transition studied in this thesis, is the very strong suppression of certain states. This can be explained by the phase coexistence at the transition, which generates a double peak in the probability distribution. Figure 2.4 illustrates this in terms of energy. Each peak belongs to one phase and states of those energies are very likely, while configurations with intermediate energies hardly occur at all. In order to obtain a good estimate of a distribution, our Monte Carlo simulation is required to sample each energy sufficiently, or to be more precise, it has to gather statistic in both phases. However, the region of suppressed intermediate states, often called *the barrier*, prevents the simulation from going from one phase to the other within an acceptable amount of computation time. This section focuses on algorithms and the involved methods to increase the occurrence of such *tunnel events*. There are methods that, in simple words, try to walk around the barrier, such as parallel tempering, but we implemented two versions of the multicanonical method, MUCA for short. It beats the barrier by sampling from a different, in fact flat probability distribution and subsequently uses modified configuration weights. The trick is that the original and desired distribution can be regained mathematically exact in a post-production step using so called reweighting techniques, as will be discussed in Sec. 2.3.4.

2.3.1 The Multicanonical Method

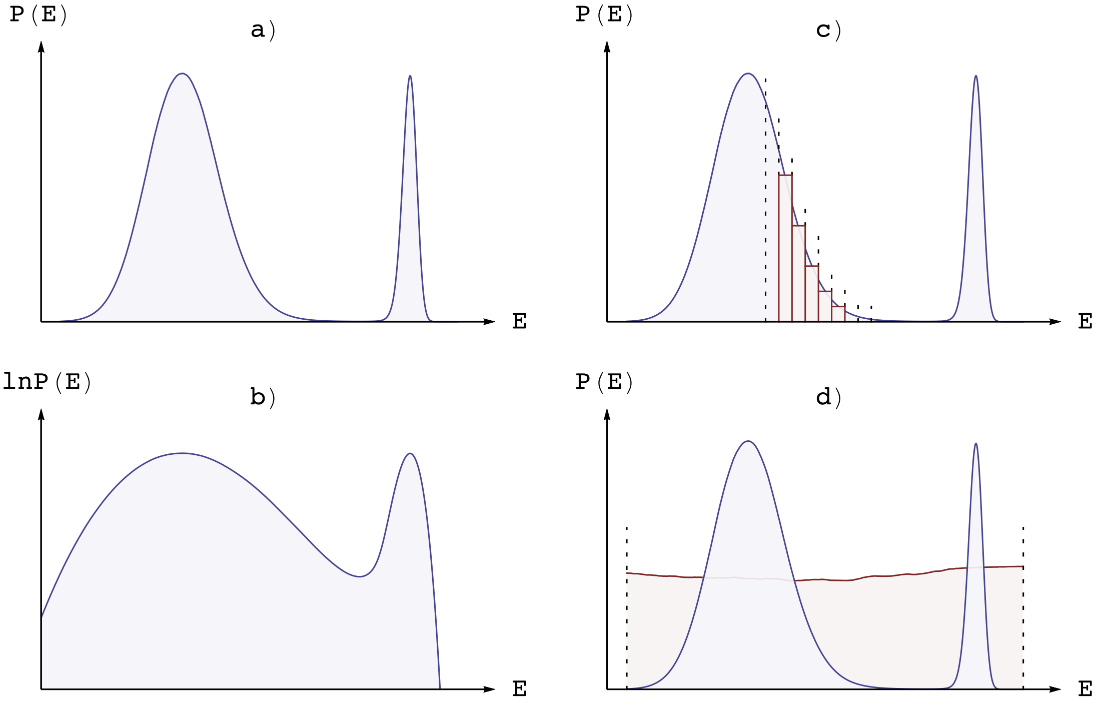

FIGURE 2.4: Probability distribution with suppressed states between a double peak: **a)** On a linear scale - **b)** On a logarithmic scale - **c)** Dividing the continuous distribution into equally sized bins of discrete values - **d)** Flat MUCA histogram with manually set boundaries.

The multicanonical method [40–43] is derived from the canonical ensemble with fixed particle number, volume and temperature. We have seen before that for this ensemble the partition function is often written in terms of the density of states. Thereby, the energy E is used as governing variable. So far, energy was treated as a continuous quantity, which it is of course, especially for off-lattice systems such as the one at hand. Nonetheless, computers have finite precision, so let us discretise energy into small intervals or *bins* of size ΔE. This allows us to express the sampled probability distribution via a histogram $H(E)$. Whenever a microstate with $E - \Delta E/2 < E < E + \Delta E/2$ is measured, an entry into the according bin of the histogram is made. Since the microstates are drawn from $P_{\mathrm{NVT}}(E)$ we see that

$$H(E) \propto P_{\mathrm{NVT}}(E) = \frac{1}{\mathcal{Z}_{\mathrm{NVT}}} \Omega(E)\, W_{\mathrm{NVT}}(E) \ . \tag{2.36}$$

In the canonical ensemble, the configuration weight $W(E)$ is the Boltzmann weight $e^{-\beta E}$ and the partition function reads

$$\mathcal{Z}_{\mathrm{NVT}} = \int \mathrm{d}E\; \Omega(E)\; e^{-\beta E} \;, \tag{2.37}$$

where for the aforementioned discrete energies, the integral is just a sum. In order to gain statistics for unfavoured states we now want the probability distribution to be flat, or in other words, we want

$$H(E) \approx \mathrm{const} \;. \tag{2.38}$$

To that end, we replace the Boltzmann weight with a multicanonical configuration weight, for which we first *assume* that it gives a flat histogram. Formally, we can write the multicanonical partition function as

$$\mathcal{Z}_{\mathrm{muca}} = \int \mathrm{d}E\; \Omega(E)\; W_{\mathrm{muca}}(E) = \int \mathrm{d}E\; \Omega(E)\; \mathcal{W}(E) \;. \tag{2.39}$$

Note the two different notations. While $W_{\mathrm{muca}}(E)$ is the formal configuration weight of this new, multicanonical ensemble, we also introduced the curly weight $\mathcal{W}(E)$. This was done to emphasise that those weights are used in the actual simulation and that they are obtained iteratively in advance, but it mainly highlights similarities with MUGC, as of Sec. 2.3.3. Again, we want to ensure detailed balance and obtain the acceptance criterion in terms of the simulation weights. Due to symmetric suggestion probabilities, a particle displacement that changes the energy from $E \to E^{\star}$ will be accepted with

$$a(E \to E^{\star}) = \min\left(1,\; \frac{\mathcal{W}(E^{\star})}{\mathcal{W}(E)}\right) \;. \tag{2.40}$$

The remaining question is how the iteration leading to the final simulation weights $\mathcal{W}(E)$ is done. Let us, therefore, look at an iteration step (t) that was performed with non-final weights $\mathcal{W}^{(t)}(E)$ and produced the histogram $H^{(t)}(E)$. The *trivial modification* is then done with

$$\mathcal{W}^{(t+1)}(E) = \frac{\mathcal{W}^{(t)}(E)}{H^{(t)}(E)} \;. \tag{2.41}$$

We see that states that were occupied very often, with an accordingly large number of entries in the histogram, get a lower weight in the next iteration. Unfavoured states, on the other hand, are assigned a larger weight and should therefore score more often in the next iteration. This update actually corresponds to the current estimate of the density of states for the energy bin of interest

$$\Omega(E) \approx \frac{H(E)}{\mathcal{W}(E)} \;. \tag{2.42}$$

We self-consistently update the weights until they match the inverse density of states well enough, so that $H(E)$ fulfils a certain flatness criterion. Employing the final weights, we will have a flat histogram and can relate

$$\mathcal{W}(E) \approx \frac{1}{\Omega(E)} = \frac{1}{\mathcal{Z}_{\mathrm{NVE}}} \;. \tag{2.43}$$

This also gives some understanding of the weights themselves; they are just an iteratively improving estimate for the inverse density of states.

While this method is rather sophisticated and allows to save effort in the production run, the iteration process to obtain the weights is accordingly challenging and computationally intensive. It is key to have sufficient statistic for every bin of every iteration step. In the implementations used in this thesis, we aimed for approximately 10^6 entries in every bin during iterations. Typically, the algorithm only occupies very few bins at first and then *eats* its way towards a flat and broader histogram in later iterations. This behaviour is actually advantageous and an easy way to tune the performance was realised by increasing the amount of statistic, depending on the expected numbers of populated bins. Also, the choice of the initial weight configuration is of some interest. One could simply initialise the weights according to the Boltzmann weights of infinite temperature, hence $\mathcal{W}^{(t=0)}(E) = \mathrm{const}$, but it turns out that choosing a finite temperature may be favourable. When dealing with phase transitions, the initial temperature should be sufficiently above or below the critical temperature. This can be justified by considering that we need to manually set the lower and upper limit of the energy range, for which the histogram shall be flat. This requires a preliminary understanding of the system, but it successfully avoids the iteration from attempting to flatten the histogram for an infinite energy range. From Eq. (2.41) we see that in order to avoid zero divisions, one has to find a way to deal with empty bins. These are commonly encountered at the edges of the energy range occupied by the current iteration. Either the weights at such unoccupied energies are not updated, or one can extrapolate from the last hit bin with the aforementioned appropriate temperatures, sufficiently below and above the critical temperature. By doing so, we can usually ensure that the iteration starts at one border and successively enlarges the sampled energy region until arriving at the other predefined border.

A major disadvantage of the *trivial modification* is a certain loss of information gathered in previous iterations, upon each update. When looking at Eq. (2.41) again, we see that only the last histogram $H^{(t)}(E)$ contributes to the change. While there is of course some implicit memory of past iterations encapsulated in the previous weight, we have no measure how many data points actually contributed in total. Hence, a few *unluckily* sampled iterations can ruin a large portion of the previous work.

The solution to this particular issue is the *error weighted* or *recursive weight modification* [44]. Using this approach, a memory of all the previous statistics is kept and the weighted average is used to determine how severely weights are changed. To this end, we introduce the quality parameter $q^{(t)}(E)$, which measures how flat the histogram is locally. This is done by comparing two neighbouring bins, namely those at E and $E + \Delta E$ through the relation

$$q^{(t)}(E) = \frac{H^{(t)}(E+\Delta E)\, H^{(t)}(E)}{H^{(t)}(E+\Delta E) + H^{(t)}(E)} \,. \tag{2.44}$$

Doing a quick back of the envelope calculation, one can be convinced that $q^{(t)}(E)$ is small if the entries in neighbouring bins are different, while a large parameter value, corresponding to good quality, occurs when neighbouring entries are approximately the same. Note that this parameter only represents the latest iteration. Therefore, we shall accumulate those to

$$Q^{(t)}(E) = \sum_{i}^{t-1} q^{(i)}(E) \,, \tag{2.45}$$

which offers the desired memory over all previous iterations, excluding the most recent one. By introducing $q^{(t)}(E)$ and $Q^{(t)}(E)$ we moved the algorithm away from looking at individual bins towards a representation where the ratio *between* bins is considered. In the same spirit we can express a ratio between the iteration weights as

$$R^{(t)}(E) = \frac{\mathcal{W}^{(t)}(E + \Delta E)}{\mathcal{W}^{(t)}(E)} \ . \tag{2.46}$$

We now desire to update this ratio somehow depending on the determined quality of the last and all previous iterations. Having done so, we then obtain the new weights from the new ratio. The updated ratio is

$$R^{(t+1)}(E) = R^{(t)}(E) \left(\frac{H^{(t)}(E)}{H^{(t)}(E + \Delta E)} \right)^{\kappa} \ , \tag{2.47}$$

with

$$\kappa = \frac{q^{(t)}(E)}{Q^{(t)}(E) + q^{(t)}(E)} \ . \tag{2.48}$$

Before relating the weight ratio to its physical meaning, let us have a closer look at three possible outcomes of such an update, in a quite heuristic way:

1. If the neighbouring bins are approximately equal, as it is the case if the used weights are satisfying, we want to keep them. We see that $H^{(t)}(E)/H^{(t)}(E + \Delta E) \approx 1$ in Eq. (2.47). Thus, only small corrections are made to the ratio.
2. If both histogram entries are rather different, the exponent κ becomes more interesting. This is the case when all previous weights were actually good but due to *unlucky* sampling in the last iteration, the most recent histogram entries differ greatly. Due to good previous results, the accumulated quality $Q^{(t)}(E)$ will be large, while $q^{(t)}(E)$ is neglectable. As a consequence, $\kappa \to 0$ and the ratio will not be changed much. The memory helped to compensate for random fluctuations.
3. On the other hand, if most previous weights led to unwanted histograms but the new iterations provided better results, this would reflect in $q^{(t)}(E)$ being the dominant contribution. Accordingly, $\kappa \to 1$ and the ratio will be tuned heavily to incorporate the new information.

We see that the direct update of weights could not check the flatness of histograms as efficiently as the error weighted recursion. The flatness criterion was implemented rather implicitly through repetitive sampling with adapted configuration weights. However, there are certain disadvantages to the more advanced method. Due to the long-lasting memory implemented in the quality parameter, it is nearly impossible to compensate for large scale fluctuations. By way of example, if the iteration accidentally jumps to the second set boundary very early, one can systematically observe oscillations in the weights and histograms as the energy range is enlarged from two sides. Those oscillations then prevent the weights from converging to the required criterion. Of course, this can be avoided by carefully setting the initial weights as described earlier, to ensure a steady and step-wise but slow exploration of the energy regime.

So far, we made a point of showing the algorithm in a straight forward way and showcasing self-consistency. We shall now give some physical interpretation from which the update rule Eq. (2.47) can actually be derived in terms of entropy. Remembering the microcanonical ensemble, we can relate entropy and the density of states

$$S(E) = k_{\mathrm{B}} \ln \Omega(E) \ , \tag{2.49}$$

which then allows an identification with the configuration weights

$$\ln \mathcal{W}(E) \approx -S(E)\ . \tag{2.50}$$

We see that the configuration weights are our temporarily best local estimates for the microcanonical entropy. Asking for the difference in entropy between two bins of energy E and $E + \Delta E$

$$\Delta S(E) = S(E + \Delta E) - S(E)\ , \tag{2.51}$$

we realise that this is nothing else than the logarithm of the inverse weight ratio, introduced in Eq. (2.46)

$$\Delta S^{(t)}(E) = \ln\left[\frac{\mathcal{W}^{(t)}(E)}{\mathcal{W}^{(t)}(E + \Delta E)}\right] = -\ln R^{(t)}(E)\ . \tag{2.52}$$

Furthermore, one can see from Eq. (2.51) that for small energy bins, $\Delta S(E)/\Delta E$ is the derivative of entropy with respect to energy, which is nothing else than the microcanonical inverse temperature $\beta(E)$. Without going into any detail, we want to point out that the quality factor given in Eq. (2.44) is in fact just a weight of the estimate for $\beta(E)$ that analytically minimises its error for each iteration, justifying the name of this update method.

When implementing any of the two methods, it is highly commended to store the weights in a logarithmic fashion and accordingly, to employ logarithmic arithmetic. The reason for this is again the precision limit of computers when, for example, adding very large numbers to very small numbers, ranging over orders of magnitude. Further, when we refer to a *flat* histogram, we usually mean that entries in all bins do not vary by more than 10% to 20%. This is for the most part sufficient, but more sophisticated methods exist, which measure and predict the flatness over multiple iterations. In simple words: just because the last iteration led to a flat histogram, one cannot be sure that the next and possible crucial one will again do so. Other variations of MUCA attempt to *predict* suitable weights [45] to reduce the iterations needed. We considered and tested such an ansatz for the thesis but failed to implement in a timely manner. In our final program, we simply check for sufficient tunnel events, which are a description for the system moving from one phase to another. A tunnel event is observed when the simulation reaches an energy value near to one of the set borders after previously visiting the other border. This also ensures that the whole desired energy range is sampled.

2.3.2 Parallel Muca

Even when utilising the error weighted recursion, the iterative process of repeatedly sampling with updated weights is often equally or more time consuming than the actual production run. The parallel implementation of the multicanonical algorithm [46] can speed up this process by gathering the statistics on multiple threads at the same time.

As illustrated in Fig. 2.5, we have $i = 1, ..., n$ threads, each generating its own histogram $H_i^{(t)}(E)$ using the *shared* weights $\mathcal{W}^{(t)}(E)$, which were obtained from the last iteration. After the sampling, the histograms are collected to a total histogram $H_\Sigma^{(t)} = \sum_i^n H_i^{(t)}(E)$ on the host, which is then processed to obtain the weights $\mathcal{W}^{(t+1)}(E)$ for the next sampling step using either update rule. The new weights are then send to all threads to create the next histograms $H_i^{(t+1)}(E)$ and so forth. Not only does the approach allow

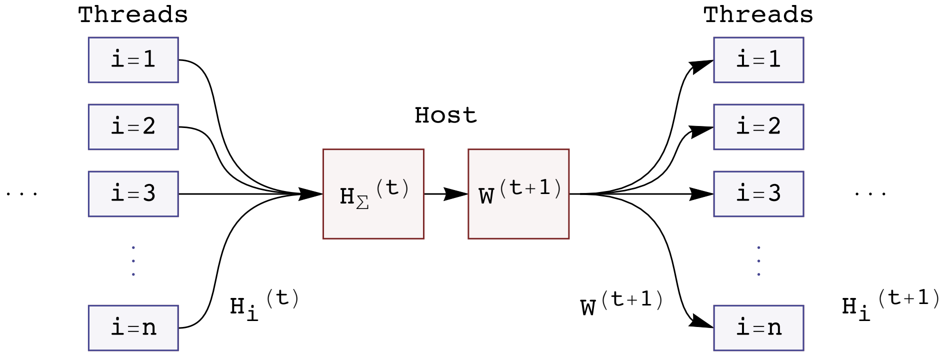

FIGURE 2.5: Illustration of the parallel MUCA. Each thread creates an individual Markov chain yielding separate histograms $H_i^{(t)}(E)$ but the threads are still correlated through the shared weights.

to gather statistics more quickly, but each thread is actually independent and (mostly) uncorrelated of the others. The only correlation is created through the shared weights, effectively pushing all the threads in the same direction. When deciding if the weights are sufficiently flat, we also check for tunnel events, the amount of which should now exceed the thread count. Thereby, one can be decently sure that every thread sampled the total range.

One could then use the final weights and let a single thread perform the simulation run, yielding a single time series. Instead, we usually just divide the desired amount of statistics onto a number of threads to gain separate time series. Those are then imported into the post processing routine, each making up one of the blocks from which the jackknife bins are created. Hence, we choose the number of threads of the production to match the desired number of jackknife bins; in this work usually 128, see Sec. 2.5 for details. It has to be mentioned that due to separating the Markov chain, a difference between the observables obtained from parallel and single-thread implementations are expected. However, it has been shown in the reference that the (very small) relative deviation stays constant for an increasing number of threads, while the total amount of measurements was unchanged.

2.3.3 Muca in the Grand Canonical Ensemble

We now want to adapt the multicanonical method to the grand canonical ensemble and conveniently call it MUGC. That is, while the original algorithm provided a flat energy probability, we now aim for a flat probability distribution of the particle number N. This idea is rather similar to the multimagnetic method, which can be implemented on lattice systems.

Using Eq. (2.22) and (2.28) we can express the grand canonical partition function in terms of the density of states

$$\mathcal{Z}_{\mu\mathrm{VT}} = \sum_{N=0}^{\infty} \int \mathrm{d}E_N \; \Omega(E_N) \; e^{-\beta E_N} \; e^{\beta\mu N} \; . \tag{2.53}$$

Beware of the footnote N for every occurrence of energy to clarify the explicit dependence. We now redo the formalism of replacing the original configuration weight $W_{\mu\mathrm{VT}}(N, E_N) = e^{-\beta E_N} e^{\beta\mu N}$ with a chosen, custom one: $W_{\mathrm{mugc}}(N, E_N)$. In contrast to the multicanonical method, the iteratively obtained weights $\mathcal{W}(N)$ only replace part of the full configuration weight:

$$\mathcal{Z}_{\text{mugc}} = \sum_{N=0}^{\infty} \int \mathrm{d}E_N \; \Omega(E_N) \; W_{\text{mugc}}(N, E_N) \tag{2.54}$$

$$= \sum_{N=0}^{\infty} \int \mathrm{d}E_N \; \Omega(E_N) \; e^{-\beta E_N} \; \mathcal{W}(N) \; , \tag{2.55}$$

or in terms of the canonical partition function

$$\mathcal{Z}_{\text{mugc}} = \sum_{N=0}^{\infty} \mathcal{Z}_{\text{NVT}} \; \mathcal{W}(N) \; . \tag{2.56}$$

We see that the weights used in the production $\mathcal{W}(N)$ only ensure equal likelihood for different values of N but not for energies, since the original Boltzmann factor is kept. Hence, this procedure only defeats the barrier in one of the two variables. Further, the weights are only valid for one specific inverse temperature β. The role of inverse temperature of the multicanonical ensemble is filled by the chemical potential.

Using the same reasoning as in Sec. 2.2.4, the acceptance criteria for insertions and deletions turn out to be different form each other due to asymmetric suggestion probabilities. An update that attempts to insert a particle and thereby goes from state $\{\phi\} \to \{\phi^\star\}$ at some $(N) \to (N+1)$ with according energy change $\Delta E = E_{N+1} - E_N$ is accepted with

$$a_+(\{\phi\} \to \{\phi^\star\}) = \min\left(1, \; \frac{V}{N+1} \frac{\mathcal{W}(N+1)}{\mathcal{W}(N)} e^{-\beta \Delta E}\right) \tag{2.57}$$

$$= \min\left(1, \; \exp\left[-\beta \Delta E + \ln\left[\frac{V}{N+1}\right] + \ln \mathcal{W}(N+1) - \ln \mathcal{W}(N)\right]\right) \; . \tag{2.58}$$

The deletion attempt changing the system from $\{\phi\} \to \{\phi'\}$ with $(N) \to (N-1)$ and $\Delta E = E_{N-1} - E_N$ is accepted with

$$a_-(\{\phi\} \to \{\phi'\}) = \min\left(1, \; \frac{N}{V} \frac{\mathcal{W}(N-1)}{\mathcal{W}(N)} e^{-\beta \Delta E}\right) \tag{2.59}$$

$$= \min\left(1, \; \exp\left[-\beta \Delta E + \ln\left[\frac{N}{V}\right] + \ln \mathcal{W}(N-1) - \ln \mathcal{W}(N)\right]\right) \; . \tag{2.60}$$

We explicitly give the criteria using the logarithmic representation as those were eventually implemented. Retrospectively, we highly recommend doing so from the beginning, as it removes one possible cause of problems when debugging. While we could not isolate troublesome issues to be solely dependant upon the lacking precision when using non logarithmic scales, there is no performance disadvantage, especially when combining the more precise way with a lookup table.

Given below are two ways to express the probability distribution, which emphasise the remaining barrier in energy. Comparing to the probability in dependence of particle number and energy for the grand canonical ensemble

$$P_{\mu\mathrm{VT}}(N, E_N) = \frac{1}{\mathcal{Z}_{\mu\mathrm{VT}}} \Omega(E_N)\, e^{-\beta E_N}\, e^{\beta\mu N} \;, \tag{2.61}$$

we give the analogon for MUGC as

$$P_{\mathrm{mugc}}(N, E_N) = \frac{1}{\mathcal{Z}_{\mathrm{mugc}}} \Omega(E_N)\, e^{-\beta E_N}\, \mathcal{W}(N) \;. \tag{2.62}$$

Having done so, we remember that we only achieve a flat histogram $H(N)$ for the particle number. This corresponds to

$$P_{\mathrm{mugc}}(N) = \frac{1}{\mathcal{Z}_{\mathrm{mugc}}} \int \mathrm{d}E_N\; \Omega(E_N)\, e^{-\beta E_N}\, \mathcal{W}(N) \tag{2.63}$$

$$= \frac{1}{\mathcal{Z}_{\mathrm{mugc}}} \mathcal{Z}_{\mathrm{NVT}}\; \mathcal{W}(N) \;, \tag{2.64}$$

where we had to integrate over all energies E_N, possible with the particle number at hand. This is helpful in the sense that it again warrants the association with the canonical partition function and we see the similarity to the original MUCA through Eq. (2.36) and (2.43)

$$\mathcal{W}(N) \approx \frac{1}{\mathcal{Z}_{\mathrm{NVT}}} \;. \tag{2.65}$$

Furthermore, we come to recognise one of the intrinsic challenges with this method. The iteratively obtained weights have to be based on a sufficiently sampled phase space. That is, for each of the many particle numbers, enough positions have to be visited to sample the density of states.

Formally, the presented technique is rather similar to that used by Wilding [28] or MacDowell et al. [15, 18]. The major difference is that we employ the sophisticated error weighted recursion to iteratively obtain the production weights $\mathcal{W}(N)$ for some chosen temperature, whereas Wilding's process systematically moves away from the critical point, further lowering the temperature in each step. Near the critical point there is no suppression of the later unfavoured states, as the two phases are indistinguishable. Consequently, the weights can be used from the previous temperature, yielding a sufficiently flat histogram from which the weights of the next iteration are obtained. One minor advantage of our approach is abandoning the factor involving explicit values of the chemical potential. We incorporate this into the weights. Hence, no a priori knowledge of the actual value of the chemical potential is required. Numerical estimates are gained from reweighting to the equal height chemical potential later, as one would do with inverse temperature in the MUCA formalism.

In order to acquire the weights within passable amounts of time, rather well tuned parameters of the simulation are necessary. Firstly, we set the inverse temperature to be $\beta = 2.5$ so that $T \approx 0.9\, T_{\mathrm{crit}}$ is sufficiently below the critical temperature. Secondly and more importantly, a maximum particle density $\rho_{\mathrm{max}} = 0.82$ should not be exceeded to avoid a jamming transition. This corresponds to the right-hand (or upper) boundary of the simulation range and was pretty much determined by trial and error. The problem with larger densities is the transition into the solid phase, where further insertions become impossible due to constraints such as the packing ratio. The system becomes so packed that there is no position in which an insertion would cause an acceptable energy change. This chain of thought also leads to the third and last imperative consideration.

In contemplation of reaching the before said maximum density, we have to enable shift moves in the simulation. Theoretically, and this is in fact done by Wilding, one could suffice without the shift move and realise particle movement through consecutive deletions and insertions. But for our weights to converge, we need many entries in every bin of the histogram and the deletion with subsequent insertion is not efficient enough in changing the configuration of a very dense liquid to accept further insertions. The shift moves allow the system to finely adjust configurations through very local changes, which correspond to subtle density fluctuations. Having said that, even with all the tweaking, the algorithm only postpones the critical slowing down [15]. Not only are we neglecting the energy contribution to the joint (density and energy) probability distribution [27], but to sample the whole phase space efficiently, further factors such as the actual geometry of the condensate need to be considered [29].

2.3.4 Reweighting between Ensembles

The MUCA and MUGC methods are our preferred choice over common METROPOLIS simulations. Nonetheless, they involve an additional step of post processing to obtain the canonical estimates for observables that we are usually interested in.

To start with a general case, say a simulation was done in ensemble Y and we want to obtain expectation values of an observable $\mathcal{O}$ in another ensemble X. Looking at Eq. (2.14) and keeping in mind that each observation is expected with the configuration weight of the chosen ensemble, we can reweight with

$$\langle \mathcal{O} \rangle_X = \frac{\langle \mathcal{O}\, W_X / W_Y \rangle_Y}{\langle W_X / W_Y \rangle_Y} \ . \tag{2.66}$$

Here, the configuration weight used in the simulation W_Y is divided out and the desired target weight W_X is multiplied. It seems trivial that the series of measurements was taken in ensemble Y, as denoted by the outer index, but it is important because it implicates the underlying phase space including the spanning variables.

For the multicanonical ensemble, the configuration weight is just the production weight

$$W_{\text{muca}}(E) = \mathcal{W}(E) \ , \tag{2.67}$$

and observables are only dependant upon energy, so that we can obtain canonical estimates with

$$\langle \mathcal{O} \rangle_{\text{NVT}} = \frac{\left\langle \mathcal{O}(E)\ e^{-\beta E} / \mathcal{W}(E) \right\rangle_{\text{muca}}}{\langle e^{-\beta E} / \mathcal{W}(E) \rangle_{\text{muca}}} \ . \tag{2.68}$$

Note that the underlying phase space for the canonical ensemble is identical to that of the multicanonical ensemble. As we can only make a finite amount of measurements $\mathcal{O}_i$, we again approximate the expectation value with a sum over all n measurements $\langle \cdot \rangle \to \frac{1}{n} \sum_i^n$ and

$$\langle \mathcal{O} \rangle_{\text{NVT}} = \frac{\sum_i^n \frac{\mathcal{O}_i}{\mathcal{W}(E_i)}\ e^{-\beta E_i}}{\sum_i^n \frac{1}{\mathcal{W}(E_i)}\ e^{-\beta E_i}} \ , \tag{2.69}$$

where the prefactor was already cancelled. Reweighting from MUGC to obtain a canonical estimate seems possible when taking care of the variables of phase space. Since the sampling took place for different values of particle number N, phase space of the grand canonical or MUGC ensemble is larger than that of the canonical ensemble and has to be projected. The target configuration weight is again just the Boltzmann weight $e^{-\beta E_N}$

for the chosen particle number

$$W_{\mathrm{NVT}}(E_N) = e^{-\beta E_N} \ , \tag{2.70}$$

and the simulated weight was

$$W_{\mathrm{mugc}}(N, E_N) = e^{-\beta E_N} \mathcal{W}(N) \ . \tag{2.71}$$

We see that the Boltzmann factor cancels, as it was already part of the simulation weight

$$\langle \mathcal{O} \rangle_{\mathrm{NVT}} = \frac{\langle \mathcal{O}(N, E_N)/\mathcal{W}(N) \rangle_{\mathrm{mugc}}}{\langle 1/\mathcal{W}(N) \rangle_{\mathrm{mugc}}} \ . \tag{2.72}$$

The projection from the grand canonical to the canonical phase space is done using the Kronecker delta, so that only measurements taken for the target particle number are included in the average

$$\langle \mathcal{O} \rangle_{\mathrm{NVT}} = \frac{\sum_i^n \frac{\mathcal{O}_i}{\mathcal{W}(N_i)} \, \delta_{N_i N}}{\sum_i^n \frac{1}{\mathcal{W}(N_i)} \, \delta_{N_i N}} = \frac{1}{m} \sum_i^n \mathcal{O}_i \, \delta_{N_i N} \ , \tag{2.73}$$

where $m = \sum_i^n \delta_{N_i N}$ is simply the number of contributing time-series entries. Through this equation, one can assume that the grand canonical simulation is theoretically just multiple canonical simulations stitched and mixed together. In fact, one should be able to recover a valid canonical time series for a chosen observable just by reduction to one fixed particle number. In order to make this a feasible approach, the amount of generated statistics of such a grand canonical simulation has to be increased accordingly.

Thus far, this section has only dealt with obtaining expectation values from time series, which is sometimes referred to as *time-series reweighting*. For observables of a continuous nature, such as energy, this is the preferred method, as it is mathematically exact and allows to recover expectation values to full precision. The second, often more efficient reweighting technique is the so called *histogram reweighting*. It relies on collecting the measurements of the time series into a histogram and requires to discretise the quantity of interest in order to sort into bins. Picking up energy as an example, this would lead to a small systematic bias. Assuming we obtained a histogram $H(E)$ from a MUCA simulation, we can obtain a probability distribution using

$$P_{\mathrm{NVT}}(E) = \frac{H(E) \, e^{-\beta E}/\mathcal{W}(E)}{\sum_E H(E) \, e^{-\beta E}/\mathcal{W}(E)} \ . \tag{2.74}$$

We can also estimate the canonical expectation value

$$\langle \mathcal{O} \rangle_{\mathrm{NVT}} = \frac{\sum_E \mathcal{O}_{\mathrm{acc}}(E) \, e^{-\beta E}/\mathcal{W}(E)}{\sum_E H(E) \, e^{-\beta E}/\mathcal{W}(E)} \ , \tag{2.75}$$

where $\mathcal{O}_{\mathrm{acc}}(E)$ is the collection of all observables belonging to the energy range assigned to the bin. Since the energy contributing to the bin may vary

$$\mathcal{O}_{\mathrm{acc}}(E) = \sum_{\{\phi\}} \mathcal{O}(\{\phi\}) \, \delta_{E(\{\phi\})E} \ , \tag{2.76}$$

which is practically done by picking out those $\mathcal{O}_i(E_i)$ of the time series for which E_i falls into the bin $H(E)$.

Adapting the histogram reweighting to particle numbers as the binning variable, we discover immediately that the bias due to discretisation is not present, making it the preferred choice over time-series reweighting.

2.4 Observables

We are interested in quite a number of physical quantities required for the understanding of the condensation-evaporation transition. Having said so, there are only three observables that we directly measure during the simulation.

1. The Energy E is an important property of the system due to its current configuration. We measure the energy after a certain number of sweeps. Before making the entry into the time series, a synchronisation of energy is performed. To be more precise, the interaction for all particles is calculated explicitly using the Lennard-Jones potential. This compensates for very small deviations that can occur due to our update moves, which only change the total energy according to the energy change of the individually moved particle. Additionally, a histogram is created and updated after every single Monte Carlo move. While the energy will be sorted into bins of a certain width, again causing unwanted bias, this provides way more statistics than could be stored in the time series.
2. The particle number N, while constant in the canonical ensemble, is the observable of main interest in grand canonical simulations. It is recorded in the same fashion as the energy; note, however, that this histogram is unbiased. Using the volume, which is constant in all our test cases, one can directly obtain the particle number density ρ.
3. The number of particles *within* the droplet N_D is of major interest as it is directly measurable and provides the means to calculate the particle excess in the droplet. As we outline in Sec. 2.4.1, the measurement of particles in the droplet is costly, which is why we limit ourselves to recording it in the time series, after each full sweep.

By employing the canonical expectation value and using the recorded data of the three direct observables, we derive, for instance, the thermal derivative

$$\frac{\partial \langle \mathcal{O} \rangle}{\partial T} = \frac{1}{k_\mathrm{B} T^2} \left(\langle E\, \mathcal{O} \rangle - \langle E \rangle \langle \mathcal{O} \rangle \right) \, , \tag{2.77}$$

such as the specific heat C_V: for $\mathcal{O} = E$, this is just the variance of the energy. Similarly, the *reduced isothermal compressibility* $\hat{\kappa}$, see Sec. 3.3, is obtained from the fluctuation of particle number. From the particle number of the grand canonical simulation, we can also gather the density in the liquid ρ_l and gas phase ρ_g as well as the fraction of particles in the largest cluster $\eta = N_\mathrm{D}/N$. The reweighted histogram of the particle number distribution provides the according probability distribution from which we can estimate the planar surface tension, which along with geometric reasoning [22] allows to obtain the surface tension of the droplet τ. We repeat the simulations and measurements for systems of different volume and/or particle number, so that we can study the finite-size scaling of the transition, especially the transition temperature T_c^V. The lower and upper indices are chosen to highlight that we are looking at the *condensation* at finite *volume*. This transition can be spotted in a number of ways, but we identify T_c^V through the peak of the specific heat, as is characteristic for first-order phase transitions.

2.4.1 Cluster-Size Identification

We want to know the size of the liquid droplet, which is of course essential to determine the current phase of the system. To put a name on it, we are interested in the number of particles within the droplet; from that number we can also estimate the droplet size using the liquid density.

In the simulation volume, the largest droplet is just the largest cluster of particles. We decide if a particle belongs to a cluster depending on its distance to another particle, namely if the distance is below a threshold of

$$r_{ij} \leq 2\sigma \ . \tag{2.78}$$

A straight forward and simple implementation was chosen in favour of more efficient but complicated cluster identification algorithms. We naively iterate through all particles of the system assigning a cluster index and then checking the distance to all other particles. If close enough, the indices are matched and possibly the size of the so far largest cluster is updated. No domain decomposition is used. As a consequence, this task is even slower than a full computation of system energy, if clusters are large.

This is an explicit function call and we do not track the largest cluster size as a system property. An improvement could be made by updating this quantity as part of the Monte Carlo moves, for which energies (and thereby some particle distances) have to be calculated anyway. This would also allow us to make use of the already present domain decomposition.

2.5 Error Estimation

"Their value [...] does, on the other hand, agree well with our result, although since no error bars were quoted it is impossible to tell to what extent the accord is meaningful." [26]

Whenever we present data generated by computer simulations, an error estimate should be provided to judge the credibility of the data. A variety of ways exist to obtain error bars. The results shown in this thesis use the probably most famous one: the jackknife method [47]. Therein, errors are obtained in a self-consistent manner, measuring fluctuations. Assuming a time series of a certain length, the series is cut into N_B equally long blocks or *bins* of size k. Using those, the so called *jackknife bins* are created by using all data of the time series except one bin. For each of those jackknife bins, an estimator $\mathcal{O}_{J,n}$ for the observable of interest is obtained. From those one can calculate the average jackknife estimate

$$\overline{\mathcal{O}}_J = \frac{1}{N_\mathrm{B}} \sum_{n}^{N_\mathrm{B}} \mathcal{O}_{J,n} \ . \tag{2.79}$$

If the estimator is simply the mean, this is identical to the average of the whole original time series. The squared error is then calculated using the bias corrected estimator

$$\epsilon^2_{\overline{\mathcal{O}}_J} = \sigma^2_{\overline{\mathcal{O}}_J} = \frac{N_\mathrm{B}-1}{N_\mathrm{B}} \sum_{n=1}^{N_\mathrm{B}} \left(\mathcal{O}_{J,n} - \overline{\mathcal{O}}_J\right)^2 \ , \tag{2.80}$$

which also includes the inflation factor to correct for the repeated occurrence of the same data in most jackknife bins. Evidently, it is straight forward to use the jackknife, which is probably why many publications do not elaborate on the details. On the contrary, as we attempt to support with Fig. 2.6, using the jackknife correctly can be tricky. The derivation of Eq. (2.80) is partly based on the binning method, which assumes that each of the N_B blocks is larger than the autocorrelation time. Therefore, all blocks are taken to be uncorrelated. As we see in the figure, if we choose the size k (of the excluded block) too small, we will underestimate the error. Conversely, using the the largest possible block sizes reduces the sum in Eq. (2.80) to very few addends, which causes heavy fluctuations exceeding the actual error value. All in all, it is desirable to compromise and to choose a block length in the a priori unknown stable region.

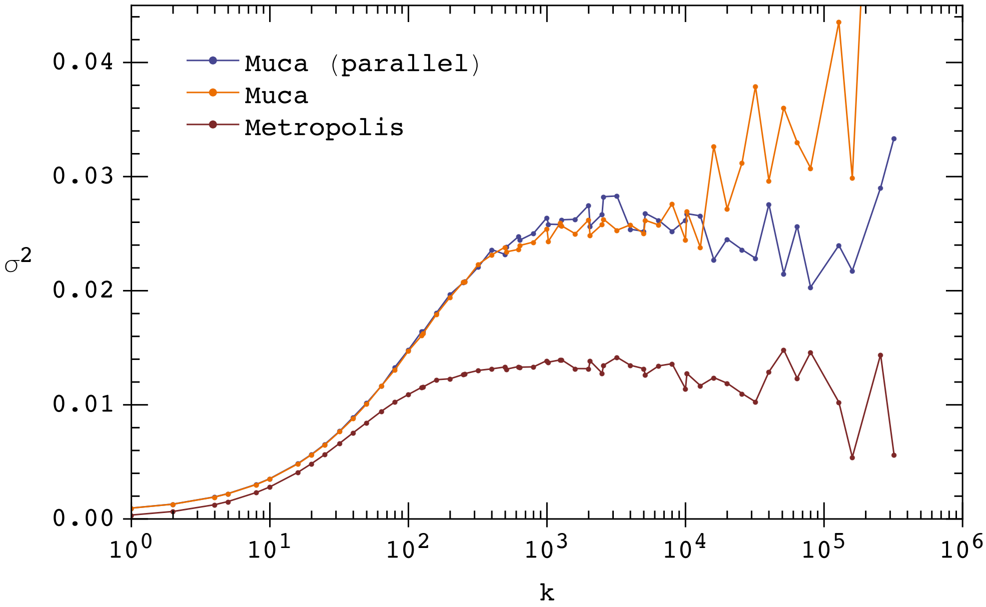

FIGURE 2.6: Jackknife variance for increasingly large bins. To avoid underestimation, a certain minimum block size k is needed. Choosing k too large leads to strong fluctuations. Note the difference in the error between MUCA and METROPOLIS for equally long time series. This can be explained by the larger energy domain sampled with MUCA; a lot of statistic does not contribute to the estimator of the target temperature, leading to a larger variance in the actually contributing data.

To illustrate this, we created three time series for the same system but using different simulation techniques. Those are a METROPOLIS (indicated red in the figure), a multicanonical production run utilising a single thread (yellow) and a parallel multicanonical run on 128 threads (blue). The length of the total time series obtained for all simulations is 1.28×10^6, where the 128 seperate time series of the parallel production were concatenated. This was only done in this test case to compare with the other two methods. In Fig. 2.6 we can clearly identify the stable region for an error estimate. Also note that the METROPOLIS variance is smaller than that obtained from MUCA, which can be explained qualitatively by the fact that MUCA samples the whole energy range, including the possible double peak in the probability distribution. METROPOLIS will only sample one of them. In addition, the figure justifies our choice of bin size $k = 10^4$ for this thesis, which we automatically obtain from the parallel MUCA simulations. The individual time series of each thread are imported as a single bin from which the jackknife bins are created.

3 Theory

3.1 Finite-Size Scaling

The main goal of finite-size scaling theory is to obtain information about infinite-size systems, such as critical exponents and transition points, from the finite systems one can actually simulate. To this end, the correlation length ξ is of major interest. It is defined to describe the length scale (determining the exponential decay) over which critical fluctuations are correlated, or plainly, it measures over what distance variables at different positions are related.

For second-order phase transitions, the correlation length diverges when approaching the critical temperature from either side, leading to power-law singularities. Using the specific heat as an example, we know $C \sim |1 - T/T_{\text{crit}}|^{-\alpha}$, but also that the correlation length scales as $\xi \sim |1 - T/T_{\text{crit}}|^{-\nu}$, where α and ν are the respective critical exponents. Evidently specific heat and correlation length relate through

$$C \sim \xi^{\frac{\alpha}{\nu}} \ . \tag{3.1}$$

While finite-size effects are negligible as long as ξ/L is small, they become apparent for $\xi \approx L$. Furthermore, as the correlation length *exceeds* the system dimension L near the critical temperature, the system size effectively limits the range of correlation and takes the role of ξ, yielding the finite-size scaling ansatz

$$C_{\max} \sim L^{\frac{\alpha}{\nu}} \ . \tag{3.2}$$

For $\xi \gg L$, the system is effectively zero-dimensional and cannot exhibit a singularity [11]. The apparent finite-size effect is the rounding of the singularity over a region of the driving variable, temperature in the given example.

For first-order phase transitions, such as the evaporation-condensation transition, the correlation length ξ stays finite at the transition temperature T_{tran} [12]. As a consequence, no universal scaling as the one outlined above occurs for first-order transitions. Instead, we observe a jump in energy, or in other words, the discontinuity occurs in the first temperature derivative of free energy. This also leads to an increasingly sharp peak in specific heat for increasing system sizes. The properties of the peak still scale in a systematic way, which we shall inquire.

3.2 The Condensation-Evaporation Transition

The condensation-evaporation transition is a first-order phase transition that takes place when crossing the coexistence curve of the liquid and gas phase, as indicated by the red line in Fig. 3.1. Since we are looking at the transition between a pure gas phase and a mixed phase consisting of a liquid droplet with surrounding gas, this is a sub-transition of the whole gas-liquid transition. It can be observed for sufficiently dilute systems at a temperature below the critical temperature. In this thesis, we are looking at two different ways to cross the coexistence curve, as indicated by the dashed arrows in the figure below.

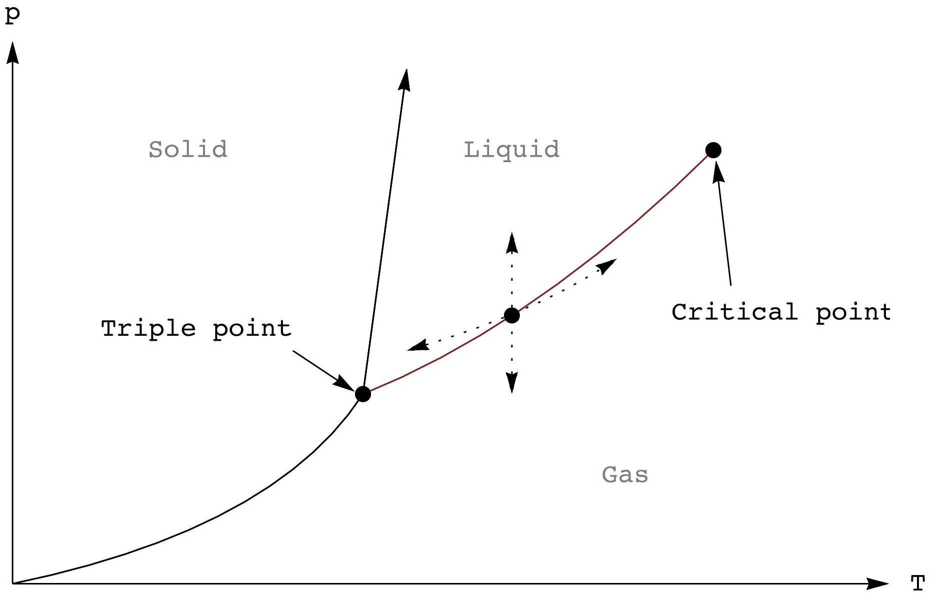

FIGURE 3.1: An idealised pressure-temperature phase diagram. The line of coexistence relevant to us is drawn in red. Note the two crossing schemes indicated by the arrows: crossing either at fixed temperature or at fixed density (which is a non-trivial function of pressure).

The native approach (or transition regime) is to cross at fixed temperature by varying the density, which we will shortly discuss in Sec. 3.2.1. This will be accomplished by increasing the particle number in the volume above the native gas density, thereby creating a particle excess. The orthogonal approach as of Sec. 3.2.2 keeps the density fixed and crosses by changing the temperature. While the two regimes seem rather different physically speaking, it was shown by Zierenberg and Janke [30, 31] that the finite-size scaling up to first order is in fact the same for both regimes.

3.2.1 Crossing at Fixed Temperature

In this section, we will discuss the theoretical background of the condensation-evaporation transition for a fixed temperature below the critical temperature. Looking at Fig. 3.3, we see that the condensation takes place when increasing the particle number or density, after starting of in the very dilute gas phase, close to a vacuum. In correspondence to the energy barrier, the probability to observe any of the consecutive mixed phases is even more suppressed. By further increasing the density, the system enters a striped phase that minimises surface free energy due to the periodic boundary conditions. An even further increase leads to a phase inverse to our droplet, specifically, a gas bubble inside a surrounding liquid that eventually vanishes for large enough densities.

Let us continue by considering a gas with variable particle number and assume a fixed chemical potential at fixed temperature below T_{crit}. In terms of a grand canonical ensemble, this provides an equilibrium particle number or *background contribution* N_{g}. In

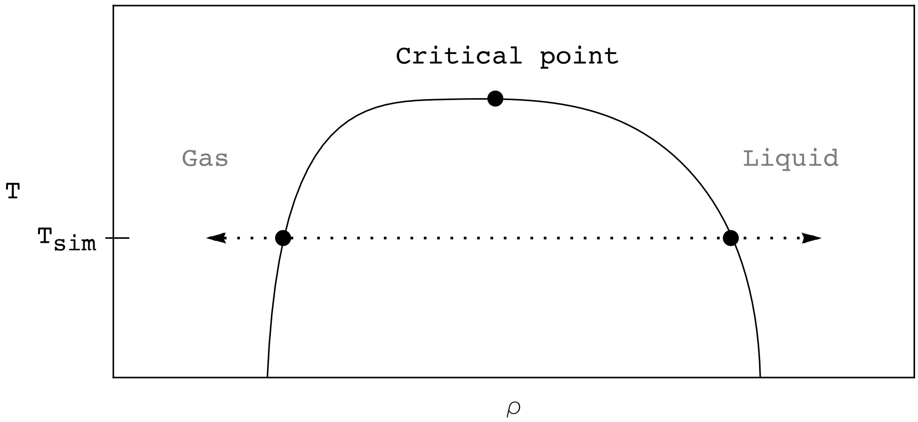

FIGURE 3.2: Sketch of the transition in density-temperature space. Adapted from Zierenberg and Janke [30]. The solid line is the infinite-size transition, which is here treated as a function $\rho_g(T)$ on the left and $\rho_l(T)$ on the right. Note the asymmetry of the curve, in contrast to that of an Ising gas. The two lower black bullets are the infinite-size transition points that coincide with the respective pure phase, while the dashed arrow indicates the finite-size scaling direction for a fixed simulation temperature $T_{\rm sim}$.

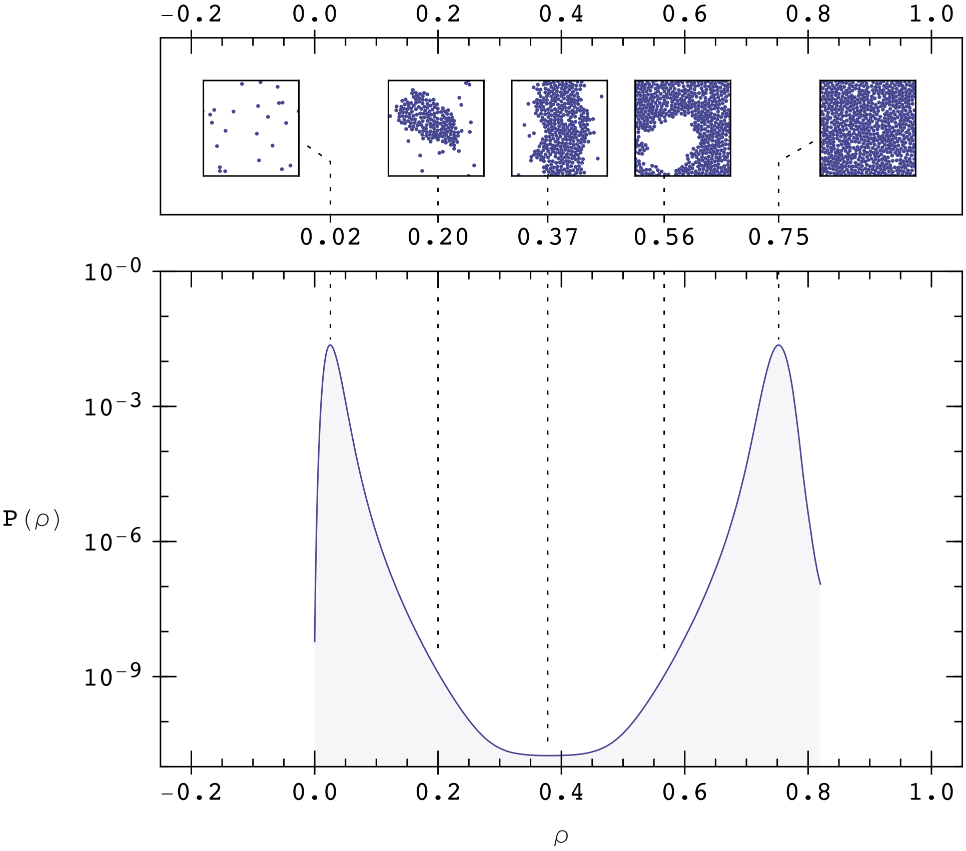

FIGURE 3.3: Snapshots of a test system of linear size $L = 30$ in dependence of density and the according probability distribution. At the set equal-height chemical potential, the full liquid-gas transition takes place from $0 \leq \rho \leq 0.82$, while for the condensation-evaporation transition, the region $\rho \leq 0.2$ is of interest. The peak positions correspond to the equilibrium gas density $\rho_g(T)$ and liquid density $\rho_l(T)$ and are almost independent of system size. Snapshots and probability look similar to the Ising gas but the distribution is not perfectly symmetric.

order to make sure of a particle excess, which corresponds to the desired supersaturation, it is now necessary to *momentarily fix* the particle number to be greater than the background contribution, i.e. $N > N_\mathrm{g}$. This translates to a canonical ensemble and one can define the particle excess $\delta N = N - N_\mathrm{g}$; if said particle excess is large enough, a droplet formation will occur. Alternatively, below the transition density ρ_c^V, corresponding to too little excess, all excess will go into the gas phase. The equilibrium droplet formation has been covered in great detail, amongst others, by Biskup [20] and Binder [15]. One can show [19] that the probability of the formation of multiple intermediately sized droplets, as opposed to a single large droplet, effectively vanishes. Thus, the problem at hand is reduced to the scenario of a homogeneous gas phase and an inhomogeneous mixed phase featuring a liquid droplet in equilibrium with the surrounding vapour. In this situation, one can identify the competition between entropy maximisation in the gas phase versus the minimisation of energy in the droplet-vapour phase. On the following pages we provide a recapitulation of the results by Biskup that are valid for the Lennard-Jones gas, while omitting the findings that concern the Ising gas, only.

Having fixed the temperature permits an inspection of the transition in terms of fixed thermal fluctuations. For the supersaturated particle gas with excess δN, the free energy of those fluctuations can be approximated with an idealised Gaussian ansatz

$$F_\mathrm{F} = \frac{(\delta N)^2}{2\hat{\kappa} V} \,, \tag{3.3}$$

where $\hat{\kappa}$ is the reduced isothermal compressibility (see Sec. 3.3), which in our context mainly describes the width of the Gaussian. As usual, $\beta = 1/k_\mathrm{B}T$ is the inverse temperature and V is the volume.

The other contributions to the free energy stem from the macroscopic droplet with volume V_D. They are approximated by

$$F_\mathrm{D} = \tau \left(V_\mathrm{D}\right)^{\frac{d-1}{d}} \,, \tag{3.4}$$

where τ is the surface free energy of a droplet of unit volume. In order to obtain τ, we follow a technique that determines the planar interface tension by comparing the height of the maximum probability in Fig. 3.3 with the minimum probability that belongs to the stripe state [48–50]. This difference in probability is identified as the free energy barrier. In the stripe state, exactly two interfaces, each of length L are present, which subsequently allows to extrapolate the surface free energy of a roughly round droplet.

According to the two different free-energy contributions, we now want to decompose the total particle excess, which we defined as the difference between the actual particle number N and the background N_g. Keeping in mind that N is set in a canonical ensemble, while the background contribution N_g corresponds to the expectation value of the particle number in a grand canonical ensemble, we separate δN into the particle excess inside the droplet δN_D and the particle excess in the fluctuating phase δN_F

$$\delta N = N - N_\mathrm{g} = \delta N_\mathrm{D} + \delta N_\mathrm{F} \,. \tag{3.5}$$

Looking at the transition in terms of the relevant *densities*, we can restrict ourselves to two variables. On the one hand, we have the (grand canonical) background density of the the gas phase ρ_g, which satisfies $\rho_\mathrm{g} = N_\mathrm{g}/V$ for an infinite-size system. On the other hand, we have the background density in the liquid phase ρ_l, as is present within the macroscopic droplet. This consideration allows linking the particle excess inside the droplet to its volume

$$\delta N_{\mathrm{D}} = (\rho_{\mathrm{l}} - \rho_{\mathrm{g}})\, V_{\mathrm{D}} \; . \tag{3.6}$$

Furthermore, we introduce what can be seen as a normalised droplet size, that is to say the scalar fraction of *particle excess inside the droplet*

$$\lambda = \frac{\delta N_{\mathrm{D}}}{\delta N} \; , \tag{3.7}$$

so that $\delta N_{\mathrm{D}} = \lambda \delta N$ and $\delta N_{\mathrm{F}} = (1-\lambda)\, \delta N$. Using this definition of λ, the total free energy $F = F_{\mathrm{D}} + F_{\mathrm{F}}$ becomes

$$F = \tau \left(\frac{\lambda \delta N}{\rho_{\mathrm{l}} - \rho_{\mathrm{g}}} \right)^{\frac{d-1}{d}} + \frac{(1-\lambda)^2\, (\delta N)^2}{2 \hat{\kappa} V} \; . \tag{3.8}$$

By utilising the dimensionless *density parameter*

$$\Delta = \frac{(\rho_{\mathrm{l}} - \rho_{\mathrm{g}})^{\frac{d-1}{d}}}{2 \hat{\kappa} \tau} \frac{(\delta N)^{\frac{d+1}{d}}}{V} \; , \tag{3.9}$$

we can rewrite the total free energy Eq. (3.8) as

$$F = \tau \left(\frac{\delta N}{\rho_{\mathrm{l}} - \rho_{\mathrm{g}}} \right)^{\frac{d-1}{d}} \left(\lambda^{\frac{d-1}{d}} + \Delta\, (1-\lambda)^2 \right) \; . \tag{3.10}$$

For the presumed fixed temperature, many of the relevant variables remain constant, in particular ρ_{l}, ρ_{g}, $\hat{\kappa}$ and τ. In the limit of large systems, we may minimise the last expression with respect to the fraction of excess λ within the one droplet. This allows us to formulate λ as a function almost exclusively dependant on the density parameter Δ. Hence, when only regarding the λ-depending part in Eq. (3.10), the consideration boils down to the following functional:

$$\Phi_{\Delta}(\lambda) = \lambda^{\frac{d-1}{d}} + \Delta\, (1-\lambda)^2 \; . \tag{3.11}$$

This can either be minimised analytically as shown by Biskup et al. [20] or numerically, yielding $\lambda(\Delta)$. Instead of going through this intensive derivation we want to support the claim using Fig. 3.4, which depicts the reduced free-energy function $\Phi_{\Delta}(\lambda)$ in $d = 2$ dimensions. There exists a constant threshold

$$\Delta_{\mathrm{c}} = \frac{1}{d} \left(\frac{d+1}{2} \right)^{\frac{d+1}{d}} \; , \tag{3.12}$$

which determines if a condensate forms. For $\Delta < \Delta_{\mathrm{c}}$ no condensation occurs ($\lambda = 0$), while for $\Delta > \Delta_{\mathrm{c}}$ the previously mentioned formation of a single macroscopic droplet containing a non-trivial excess fraction ($\lambda > \lambda_{\mathrm{c}}$) takes place. Directly at the threshold one gets

$$\lambda_{\mathrm{c}} = \frac{2}{d+1} \; . \tag{3.13}$$

For this expression, the self-consistency of Δ_{c} is easy to check. In the two dimensional case the numerical values are $\Delta_{\mathrm{c}}^{\mathrm{2D}} \approx 0.9186$ and $\lambda_{\mathrm{c}}^{\mathrm{2D}} = 2/3$.

Thanks to the above formalism, the condensation-evaporation transition can be clearly identified through the dimensionless density parameter, at $\Delta = \Delta_{\mathrm{c}}$.

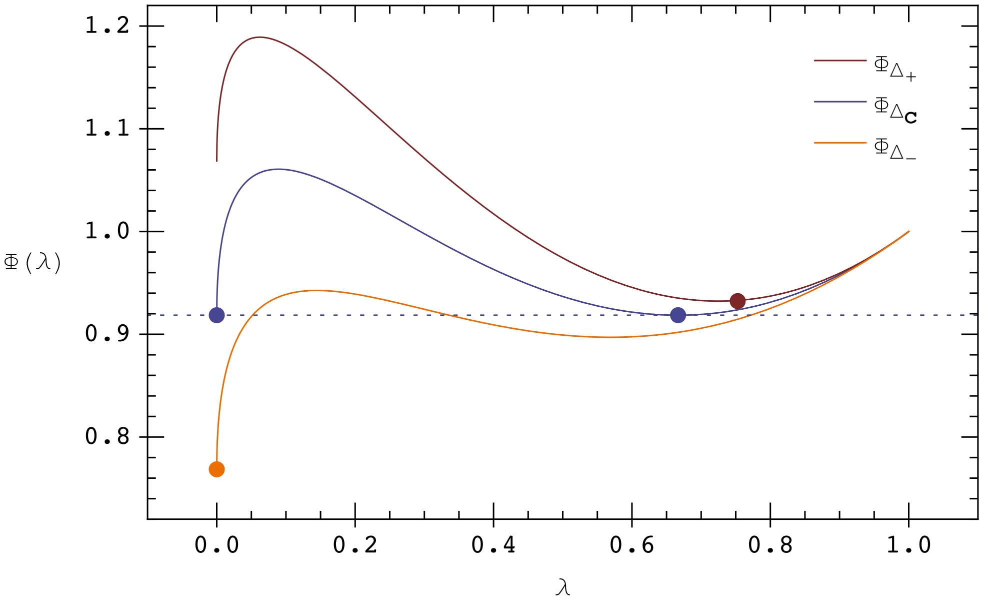

FIGURE 3.4: The reduced free-energy function $\Phi_\Delta(\lambda)$ in two dimensions for different values of the density parameter Δ. Directly at the threshold Δ_c, the functional is minimised for $\lambda = 0$ and for $\lambda = \lambda_c$. Below the threshold, i.e. $\Delta_- < \Delta_c$, no droplet formation occurs as $\lambda = 0$. Above the threshold, for $\Delta_+ > \Delta_c$, the only minimum is at some non-trivial $\lambda > \lambda_c$.

By construction, the density parameter contains the leading-order finite-size corrections to the transition density. Thus, it may be understood as a shifted density or particle number N, that accounts for corrections due to system size and dimensionality. Furthermore, we can interpret $\lambda(\Delta)$ as the expectation value of the equilibrium droplet size, which we measure in simulations through the number of particles inside the droplet N_D. Due to the assumptions stressed out before, also λ includes finite-size corrections of first order.

Aiming for an explicit look at the finite-size behaviour, Eq. (3.4) can be written in terms of the actual particle density $\rho = N/V$ instead of the particle excess:

$$\Delta = \frac{(\rho_\mathrm{l} - \rho_\mathrm{g})^{\frac{d-1}{d}}}{2\hat{\kappa}\tau} (\rho - \rho_\mathrm{g})^{\frac{d+1}{d}} V^{\frac{1}{d}} . \tag{3.14}$$

We see that the only quantity defining the transition in the previous equation is the finite-size transition density $\rho = \rho_\mathrm{c}^V$, where the (lower and upper) indices shall emphasize that the *condensation* takes place at finite *volume*. Consequently, it holds that $\Delta(\rho_\mathrm{c}^V) = \Delta_\mathrm{c}$, which yields

$$\rho_\mathrm{c}^V = \rho_\mathrm{g} + \left(\frac{2\hat{\kappa}\tau\Delta_\mathrm{c}}{(\rho_\mathrm{l} - \rho_\mathrm{g})^{\frac{d-1}{d}}} \right)^{\frac{d}{d+1}} V^{-\frac{1}{d+1}} . \tag{3.15}$$

Therein, we can clearly identify the leading-order behaviour with respect to the infinite-size transition density, which coincides with the gas density $\rho_\mathrm{g} = \rho_\mathrm{c}^\infty$. This allows us to conclude the theoretical background of the fixed temperature, yielding the finite-size behaviour in leading order as

$$\rho_\mathrm{c}^V - \rho_\mathrm{g} \propto V^{-\frac{1}{d+1}} . \tag{3.16}$$

While the description of a known, fixed particle excess over some background density was done using a canonical framework, the approach of this section actually requires grand canonical simulations. Those allow to measure densities through their expectation values, providing numerical values for the liquid and gas density, which we can confirm to stay constant for changing system sizes. Moreover, we will qualitatively justify the temperature dependence of infinite-size quantities such as $\rho_{\mathrm{l}}(T)$ and $\rho_{\mathrm{g}}(T)$ in our results.

3.2.2 Crossing at Fixed Density

The previous discussion covered the liquid vapour system at fixed temperature. Free energy was approximated through contributions from fluctuations of the density and from the surface free energy of the droplet. It was key to describe the contributions in terms of infinite-size quantities that do not change with system size, as long as the temperature is indeed kept constant.

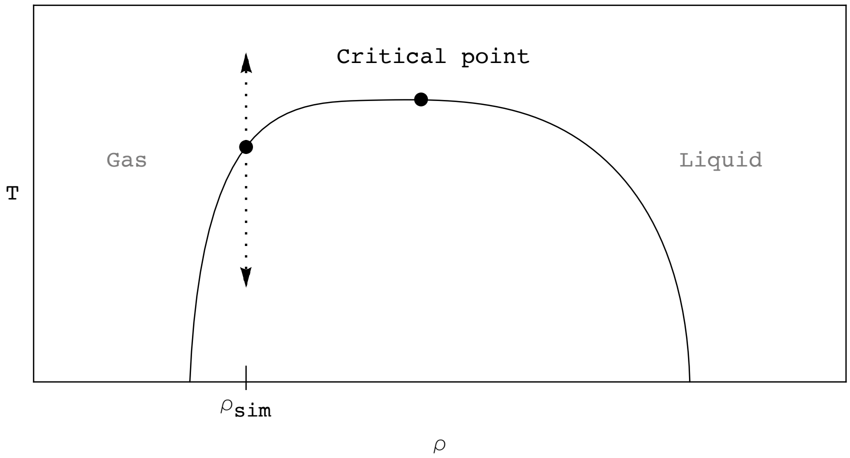

FIGURE 3.5: Illustration of the infinite-size transition (solid line). This time, the curve is treated as a function $T_{\mathrm{g}}(\rho)$ with the according (dashed) finite-size scaling direction at fixed simulation density ρ_{sim}. The scaling is orthogonal to the previous, fixed-temperature approach. Note that the representation through a transition line only holds for the infinite-size system, as the transition gets smeared out for finite systems. This is discussed in the next section, which treats the transition rounding.

Looking at the other alternative, we arrive at the orthogonal scheme, which is depicted in Fig. 3.5. This time, one keeps the density fixed. Namely, let us start again by contemplating the liquid-vapour system with N particles and finite periodic volume V. For a now chosen density, the transition temperature T_{c}^{V} separates the phase of the droplet surrounded with bulk gas from the homogeneous supersaturated gas phase. For large enough systems, the transition temperature approaches the infinite-size transition temperature $T_{\mathrm{g}} = T_{\mathrm{c}}^{\infty}$, as indicated in Fig. 3.5 by the lower black dot.

We realise that the infinite-size transition line can either be regarded as a function of temperature $\rho_{\mathrm{g}}(T)$, which we had done in Fig. 3.2, or as a function of density $T_{\mathrm{g}}(\rho)$. Considering the chosen, fixed density, therefore, enforces the latter scheme, in which a change in system size causes a shift in the transition temperature. The finite-size scaling direction is again marked by the dashed arrow.

Moreover, both transition schemes work for every point (ρ, T) of the phase diagram. This becomes clear when picking any point, e.g. a finite-size transition point $(\rho_{\mathrm{c}}^{V}, T_{\mathrm{c}}^{V})$

and treating it in a canonical ensemble, for which it belongs to one fixed-temperature and one fixed-density scheme, simultaneously. The same holds for any canonical function $f(\rho, T)$ and the orthogonal crossing schemes are equivalent. This translates to a functional dependence for which the ansatz $f(\rho, T)V^{\alpha} = 1$ allows to change from one scheme to another by a Taylor series expansion. Around some $T^{\star}$ this yields in first order

$$V^{-\alpha} = f(\rho, T^{\star}) + f'(\rho, T^{\star})\,(T - T^{\star}) + \dots \,. \tag{3.17}$$

This can be solved for T and the remaining task is to find suitable functional dependencies [30]. To this end, let us reconsider the dimensionless density parameter Δ as of Eq. (3.9). In the present scheme, the infinite-size quantities (that were assumed to be constant in the previous section) become functions of temperature. This includes the background densities $\rho_{\mathrm{g}}(T)$ and $\rho_{\mathrm{l}}(T)$, the reduced isothermal compressibility $\hat{\kappa}(T)$, the surface free energy $\tau(T)$ and also the particle excess $\delta N(T) = (\rho - \rho_{\mathrm{g}}(T))\,V$. Considering those dependencies, Eq. (3.9) reads

$$\Delta = \frac{(\rho_{\mathrm{l}}(T) - \rho_{\mathrm{g}}(T))^{\frac{d-1}{d}}}{2\hat{\kappa}(T)\tau(T)} \frac{(\delta N(T))^{\frac{d+1}{d}}}{V} \,. \tag{3.18}$$

Exact computation of $f(\rho, T)$ is generally impossible, but it turns out that rewriting Eq. (3.18) to the form

$$\Delta^{\frac{d}{d+1}} V^{-\frac{1}{d+1}} = f(\rho, T) = \frac{\rho - \rho_{\mathrm{g}}(T)}{\rho_{\mathrm{l}}(T) - \rho_{\mathrm{g}}(T)} \left(\frac{(\rho_{\mathrm{l}}(T) - \rho_{\mathrm{g}}(T))^2}{2\hat{\kappa}(T)\tau(T)} \right)^{\frac{d}{d+1}} \tag{3.19}$$

allows sufficient simplification, when expanded. At the condensation transition $\Delta = \Delta_{\mathrm{c}}$ is constant and the left-hand side of the equation is only depending on the system size V. Therefore, by fixing the system size, one can find a suitable combination of T and ρ that solves Eq. (3.19) and provides the transition point $(\rho, T) = (\rho_{\mathrm{c}}^{V}, T_{\mathrm{c}}^{V})$. This can be done numerically, but we will instead carry out the Taylor expansion of Eq. (3.19) around the infinite-size transition temperature $T^{\star} = T_{\mathrm{g}}$, keeping $\rho = N/V$ constant:

$$f(\rho, T) = f(\rho, T_{\mathrm{g}}) + f'(\rho, T_{\mathrm{g}})\,(T - T_{\mathrm{g}}) + \dots \,. \tag{3.20}$$

Revisiting Fig. 3.5, one readily sees that $\rho_{\mathrm{g}}(T_{\mathrm{g}}) = \rho$, which causes the first term $f(\rho, T_{\mathrm{g}})$ to vanish. Solving for the finite-size transition temperature $T = T_{\mathrm{c}}^{V}$ yields

$$T_{\mathrm{c}}^{V} = T_{\mathrm{g}} + \frac{\Delta_{\mathrm{c}}^{\frac{d}{d+1}}}{f'(\rho, T_{\mathrm{g}})} V^{-\frac{1}{d+1}} + \dots \,, \tag{3.21}$$

where $f'(\rho, T_{\mathrm{g}})$ remains unsolved but we obtained the leading-order behaviour of T_{c}^{V}. A comparison with Eq. (3.16) yields the same finite-size scaling for both regimes, in terms of system volume

$$T_{\mathrm{c}}^{V} - T_{\mathrm{g}} \propto V^{-\frac{1}{d+1}} \,. \tag{3.22}$$

For fixed density, this also holds in terms of the particle number

$$T_{\mathrm{c}}^{N} - T_{\mathrm{g}} \propto N^{-\frac{1}{d+1}} \,. \tag{3.23}$$

3.2.3 Transition Rounding

As a motivation on why to think about rounding, we know from thermodynamic considerations how, for instance, the fluctuations of energy around the expectation value scale with the particle number in a canonical ensemble. When thinking of the according energy probability distribution, this corresponds to the width or *the rounding* of the curve. Particularly, one can show that $\Delta E/\langle E\rangle \propto 1/\sqrt{N}$, which shows that in the thermodynamic limit ($N \to \infty$) the fluctuations vanish. The energy distribution is peaked so sharply around the average, that the canonical ensemble becomes equivalent to the microcanonical ensemble, which has no energy fluctuations by definition.

So far, we have treated the leading-order finite-size scaling of the finite-size transition density ρ_{c}^{V} and temperature T_{c}^{V} with regard to the respective infinite-size transition quantities ρ_{g} and T_{g}. In this section, the leading-order behaviour of the transition rounding will be discussed, i.e. $\Delta\rho = (\rho - \rho_{\mathrm{c}}^{V})$ and $\Delta T = (T - T_{\mathrm{c}}^{V})$, as was done in the work by Binder [13] as well as Zierenberg and Janke [30, 31, 34]. It is clear, that in order to provide a meaningful discussion of the transition region, a choice has to be made about which system states *actually belong* to that region and which do not.

Thus, one starts with a two-state approximation of the system, where the first state corresponds to the system in the condensed phase with a droplet, which has probability $P_{\text{cond}} \propto e^{-\beta F_{\text{cond}}}$. The second state, describing the completely evaporated phase, has probability $P_{\text{evap}} \propto e^{-\beta F_{\text{evap}}}$. It is worth mentioning that F_{cond} and F_{evap} are not to be confused with the free-energy contributions of fluctuations and the droplet as given in Section 3.2.1. Instead, they evaluate the total free energy Eq. (3.8) for the system in the two different states. The crossing of the phase boundary from the homogeneous (evaporated) to the inhomogeneous (condensed) state is then linked to the difference in free energy

$$\Delta F = F_{\text{cond}} - F_{\text{evap}} \, . \tag{3.24}$$

With this in mind, we can write the expectation value of any observable in terms of the probabilities of the two states

$$\langle \mathcal{O} \rangle = \frac{\mathcal{O}_{\text{evap}}\, e^{-\beta F_{\text{evap}}} + \mathcal{O}_{\text{cond}}\, e^{-\beta F_{\text{cond}}}}{e^{-\beta F_{\text{evap}}} + e^{-\beta F_{\text{cond}}}} \tag{3.25}$$

$$= \frac{\mathcal{O}_{\text{evap}} + \mathcal{O}_{\text{cond}}\, e^{-\beta \Delta F}}{1 + e^{-\beta \Delta F}} \, . \tag{3.26}$$

Choosing $|\beta\Delta F| \approx 1$ in the above expression implies that both phases have equally significant contributions to the expectation value, which is a decent criterion for the transition region. Of course, this is the condition applied in the mentioned references. Interchangeably, one requires that free-energy fluctuations are of the order $k_{\mathrm{B}}T$,

$$\Delta F \approx k_{\mathrm{B}} T \, . \tag{3.27}$$

The description up to this point is valid for the fixed-temperature as well as the fixed-density scheme, and the derived expressions for the rounding will in fact be similar, but the considerations to justify them are different.

Rounding for Fixed Temperature

In order to estimate the transition rounding at fixed temperature, we start by evaluating the total free energy Eq. (3.8), for the two different states. Firstly, we know that for the condensed phase $\lambda \neq 0$, which leads to

$$\beta F_{\text{cond}} = a\,\lambda^{\frac{d-1}{d}}\,(\delta N)^{\frac{d-1}{d}} + b\,(1-\lambda)^2\,\frac{(\delta N)^2}{V}\;, \tag{3.28}$$

where $a = \beta\tau/(\rho_{\text{l}} - \rho_{\text{g}})^{\frac{d-1}{d}}$ and $b = \beta/2\hat{\kappa}$ contain the (in this scheme) constant infinite-size quantities. On the other hand, we have $\lambda = 0$ in the gas phase, thus,

$$\beta F_{\text{evap}} = b\frac{(\delta N)^2}{V}\;, \tag{3.29}$$

which consistently recovers Eq. (3.3) for the homogeneous state. Using those two expressions and the particle excess $\delta N = (\rho - \rho_{\text{g}})\,V$, the free-energy difference Eq. (3.24) can be rewritten in terms of the density

$$\begin{aligned} \beta\Delta F(\rho) &= a\,\lambda^{\frac{d-1}{d}}\,(\rho-\rho_{\text{g}})^{\frac{d-1}{d}}\,V^{\frac{d-1}{d}} + b\left((1-\lambda)^2 - 1\right)(\rho-\rho_{\text{g}})^2\,V \qquad (3.30)\\ &= a\,(\rho-\rho_{\text{g}})^{\frac{d-1}{d}}\,V^{\frac{d-1}{d}}\left[\lambda^{\frac{d-1}{d}} + \frac{b}{a}\,(\rho-\rho_{\text{g}})^{\frac{d+1}{d}}\,V^{\frac{1}{d}}\left((1-\lambda)^2-1\right)\right]\;. \qquad (3.31) \end{aligned}$$

Directly at the transition, the free-energy difference $\beta\Delta F(\rho_{\text{c}}^V) = 0$, but the transition density $\rho_{\text{c}}^V \neq \rho_{\text{g}}$ since the system is of finite size. This implies that the factor proceeding the expression in square brackets in the previous equation cannot be zero. It follows that

$$\underbrace{\beta\Delta F(\rho_{\text{c}}^V)}_{=0} = \underbrace{a\left(\rho_{\text{c}}^V-\rho_{\text{g}}\right)^{\frac{d-1}{d}}\,V^{\frac{d-1}{d}}}_{\neq 0}\underbrace{\left[\lambda^{\frac{d-1}{d}} + \frac{b}{a}\left(\rho_{\text{c}}^V-\rho_{\text{g}}\right)^{\frac{d+1}{d}}\,V^{\frac{1}{d}}\left((1-\lambda)^2-1\right)\right]}_{=0}\;. \tag{3.32}$$

Aiming for the desired expression containing $\Delta\rho = (\rho - \rho_{\text{c}}^V)$, we expand (3.31) around the transition density $f(\rho)|_{\rho_{\text{c}}^V} \approx f(\rho_{\text{c}}^V) + f'(\rho_{\text{c}}^V)(\rho - \rho_{\text{c}}^V) + \ldots$, where $f(\rho_{\text{c}}^V) = 0$ and the first term of the differentiation vanishes due to Eq. (3.32). The second term yields

$$\beta\Delta F(\rho) \approx b\left((1-\lambda)^2-1\right)\frac{d+1}{d}\left(\rho_{\text{c}}^V-\rho_{\text{g}}\right)V\left(\rho-\rho_{\text{c}}^V\right) + \ldots \tag{3.33}$$

and with the leading-order behaviour for ρ_{c}^V from Eq. (3.16) we arrive at

$$\beta\Delta F(\rho) \propto V^{\frac{d}{d+1}}\Delta\rho\;. \tag{3.34}$$

At last, one concludes that for a free-energy difference of order unity the finite-size rounding scales as

$$\Delta\rho \propto V^{-\frac{d}{d+1}}\;. \tag{3.35}$$

Rounding for Fixed Density

We have stated before that the infinite-size quantities become functions of temperature in the fixed-density approach. This impedes an evaluation of Eq. (3.8) and we instead relate to the scaling of the droplet volume V_D, which is known [51]. We start with the most general expression for $\beta\Delta F(T)$, this time treated as a function of temperature, and expand around the condensation temperature (in regard to earlier notation, here we chose to write T_c^N over T_c^V in order to stress the fixed density). We identify $F = E - TS$, which entails $\frac{\partial}{\partial T}(\beta F) = -\frac{1}{k_\mathrm{B}T^2}E$, so that the expansion reads

$$\beta\Delta F(T)\big|_{T_\mathrm{c}^N} \approx \beta\Delta F(T_\mathrm{c}^N) - \left(\frac{1}{k_\mathrm{B}T^2}\Delta E\right)\Bigg|_{T_\mathrm{c}^N}(T - T_\mathrm{c}^N) + \ldots\ . \tag{3.36}$$

Next, we recognise $\Delta T = (T - T_\mathrm{c}^N)$ as the desired measure for the rounding and consider the scaling behaviour of $\Delta E = E_\mathrm{cond} - E_\mathrm{evap}$. Evidently, we can neglect the contribution due to the evaporated particles by thinking of the dilute gas; particles are separated further than the relevant interaction length-scale, in our case, the cutoff radius of the Lennard-Jones potential. The energy contribution of the condensate E_cond, on the other hand, scales with the droplet volume V_D, which itself is dependant on the droplet radius and was shown [51] to grow in leading order as

$$E_\mathrm{cond} \propto V_\mathrm{D} \propto N^{\frac{d}{d+1}}\ . \tag{3.37}$$

Plugging the scaling for T_c^N into Eq. (3.36) and approaching the infinite-size transition temperature T_g, we see that the corrections to T_c^N in the denominator only contribute to the energy scaling in higher orders and that

$$\beta\Delta F \propto \left(\frac{1}{k_\mathrm{B}T_\mathrm{g}^2}N^{\frac{d}{d+1}}\right)\Delta T\ . \tag{3.38}$$

By again applying our condition $|\beta\Delta F| \approx 1$, we arrive at the leading-order finite-size scaling of the transition rounding for fixed density

$$\Delta T \propto N^{-\frac{d}{d+1}}\ , \tag{3.39}$$

or equivalently, in terms of volume

$$\Delta T \propto V^{-\frac{d}{d+1}}\ . \tag{3.40}$$

3.3 Isothermal Compressibility

We superficially used the reduced isothermal compressibility $\hat{\kappa}$ without any explanation, when we followed the formalism of Biskup. As it turns out though, a closer investigation is needed to reproduce numerically matching results for the dimensionless density parameter Δ, if a relationship to a real physical quantity is desired, too.

In its widest physical sense, compressibility is a description for the amount of volume change that occurs due to a given pressure change. This process can either be adiabatic or isothermal, which in terms of thermodynamic variables defines if either entropy, or in our case, temperature, is kept constant for the partial derivative. Formally the isothermal compressibility κ can be written as:

$$\kappa = -\frac{1}{V}\left(\frac{\partial V}{\partial p}\right)_T \quad . \tag{3.41}$$

While this expression is rather intuitive, physically speaking, it is not immediately clear how to relate it to particle fluctuations. Starting with the grand canonical partition function

$$\mathcal{Z}_{\mu \mathrm{VT}} = \sum_{N=0}^{\infty} \int_{\{\phi\}} \mathrm{d}\{\phi\}\ e^{-\beta V(\{\phi\})} e^{\beta \mu N} \ , \tag{3.42}$$

where $\beta = 1/k_\mathrm{B}T$ as usual denotes inverse temperature, one can derive the expression for an expectation value of the particle number $\langle N \rangle$ through differentiation with respect to the chemical potential μ. Remembering the general expression for the ensemble average Eq. (2.7), we see that

$$\langle N \rangle = \frac{\sum_N \int_{\{\phi\}} \mathrm{d}\{\phi\}\ N\ e^{-\beta V(\{\phi\})} e^{\beta \mu N}}{\sum_N \int_{\{\phi\}} \mathrm{d}\{\phi\}\ e^{-\beta V(\{\phi\})} e^{\beta \mu N}} \ . \tag{3.43}$$

For constant volume and temperature it follows that

$$\langle N \rangle = \frac{1}{\beta \mathcal{Z}_{\mu \mathrm{VT}}} \frac{\partial \mathcal{Z}_{\mu \mathrm{VT}}}{\partial \mu} \ , \tag{3.44}$$

$$\left\langle N^2 \right\rangle = \frac{1}{\beta^2 \mathcal{Z}_{\mu \mathrm{VT}}} \frac{\partial^2 \mathcal{Z}_{\mu \mathrm{VT}}}{\partial \mu^2} \ . \tag{3.45}$$

This can be combined in the usual way to obtain an expression motivated by the variance:

$$\left\langle N^2 \right\rangle - \langle N \rangle^2 = \left(\frac{1}{\beta^2 \mathcal{Z}_{\mu \mathrm{VT}}} \frac{\partial^2 \mathcal{Z}_{\mu \mathrm{VT}}}{\partial \mu^2} - \left(\frac{1}{\beta \mathcal{Z}_{\mu \mathrm{VT}}} \frac{\partial \mathcal{Z}_{\mu \mathrm{VT}}}{\partial \mu} \right)^2 \right)_{T,V} \tag{3.46}$$

$$= \left(\frac{1}{\beta^2} \frac{\partial^2 \ln \mathcal{Z}_{\mu \mathrm{VT}}}{\partial \mu^2} \right)_{T,V} \ . \tag{3.47}$$

Next, we need to employ the potential according to the grand canonical partition function, namely the Landau free energy (or grand potential)

$$J(\mu, V, T) = F - G = -pV \tag{3.48}$$

with the total differential

$$\mathrm{d}J = -S\,\mathrm{d}T - p\,\mathrm{d}V - N\,\mathrm{d}\mu\ , \tag{3.49}$$

where F and G are the Helmholtz and Gibbs free energy, respectively. Partition function and potential are related via

$$J(\mu, V, T) = -\frac{1}{\beta}\ln \mathcal{Z}_{\mu\mathrm{VT}} = -pV\ . \tag{3.50}$$

For fluids of constant volume and temperature we can then write

$$\frac{\partial^2 \ln \mathcal{Z}_{\mu\mathrm{VT}}}{\partial\mu^2} = \beta V \frac{\partial^2 p}{\partial\mu^2}\ , \tag{3.51}$$

which allows to relate the particle fluctuations as in Eq. (3.47) with the partial derivative $\partial p/\partial\mu$ so that

$$\frac{\partial^2 p}{\partial\mu^2} = \frac{\beta}{V}\left(\langle N^2\rangle - \langle N\rangle^2\right)\ . \tag{3.52}$$

What remains to do is to establish the link to the compressibility, or in other words to rewrite the partial derivative. Dividing Eq. (3.48) by N and substituting $G = N\mu$ we can write

$$\frac{F}{N} - \mu = -\frac{pV}{N}\ . \tag{3.53}$$

If we now introduce a *volume per particle* $v = V/N$ and assume that the Helmholtz free energy can also be expressed per particle

$$f(v, T) = \frac{1}{N}\,F(N, V, T)\ , \tag{3.54}$$

one can write the chemical potential as

$$\mu(T, p) = f(v, T) + pv\ . \tag{3.55}$$

Since $f(v, T)$ is no function of pressure, it follows that $\partial\mu/\partial p = v$ and

$$\frac{\partial\mu}{\partial v} = \frac{\partial\mu}{\partial p}\frac{\partial p}{\partial v} = v\frac{\partial p}{\partial v} = -\frac{1}{\kappa}\ , \tag{3.56}$$

as well as

$$\frac{\partial^2 p}{\partial\mu^2} = \frac{\partial}{\partial\mu}\left(\frac{1}{v}\right) = -\frac{1}{v^2}\frac{\partial v}{\partial\mu} = \frac{\kappa}{v^2}\ , \tag{3.57}$$

where the commonly practised sketchy formalism of inverse partials was used. Nonetheless, this allows to combine all the pieces using Eq. (3.52) and (3.57) to obtain the well known expression for isothermal compressibility, see e.g. [52]:

$$\kappa = \frac{v\beta}{\langle N\rangle}\left(\langle N^2\rangle - \langle N\rangle^2\right)\ . \tag{3.58}$$

Having omitted some finer points so far, this equation deserves further discussion. The attentive reader has noted the appearance of $\langle N\rangle$ in the denominator, that has not been accounted for. We introduced the volume per particle, assuming a strictly known particle number N, and implicitly used this idea to treat the Gibbs free energy and

Helmholtz free energy, which are of course the potentials belonging to the canonical ensemble, where this is perfectly consistent. However, when using the ansatz to establish a relation with the grand canonical ensemble, where particle insertion and deletion are possible, the volume per particle has to be adjusted to $v = V/\langle N\rangle$ and it is helpful to rather treat it as an inverse density $\rho = \partial p/\partial\mu = 1/v$ instead. This clarifies that v does not describe the size of a particle but the simulation volume available to that particle. Additionally, the direct correspondence of particle number and density is emphasised.

Further, we strictly assumed the system to be homogeneous, which is not the case throughout all our simulations; the coexistence of the pure gas with the inhomogeneous phase of a liquid droplet with surrounding gas are studied. It has been shown quantitatively (e.g. using NpT simulations [53]), that the compressibility is larger for systems in the gas phase than in the liquid phase. This is just what one would expect from everyday language; gases are easier to compress than liquids. Using the expression as in Eq. (3.58), we were able to qualitatively confirm this behaviour, where the values for the Lennard-Jones system differ by one to two orders of magnitude between the liquid and the gas phase.

Reduced Isothermal Compressibility

After discussing this general approach, we now relate to the response function that is used by Biskup et al. [20], which follows a different reasoning. More precisely, let us recall the probability of particle excess δN was given as

$$P \sim \exp\left[-\frac{(\delta N)^2}{2\kappa V}\right] , \tag{3.59}$$

which *"is just the leading-order asymptotic of a full-fledged Gaussian (central limit) distribution"* [19]. By direct comparison, we realise that the denominator should equate to two times the variance, and thereby, the width of the considered peak in the probability distribution. So, let us introduce a *reduced isothermal compressibility*

$$\hat{\kappa} = \frac{\beta}{V}\left(\langle N^2\rangle - \langle N\rangle^2\right) . \tag{3.60}$$

When substituted for κ, this gives the desired expression for the variance up to the factor of β. In order to have a closer look at this quantity, it is useful to include the Ising gas into the discussion due to its analogy to our Lennard-Jones gas. On the one hand, this identification helps to appreciate the similarity between compressibility and its Ising equivalent, the *magnetic susceptibility*. On the other hand, there is a tight correspondence between the two respective models themselves. Magnetic susceptibility is commonly denoted as

$$\chi = \frac{\beta}{V}\left(\langle M^2\rangle - \langle M\rangle^2\right) , \tag{3.61}$$

where the volume V is just the total number of spins. The magnetisation M is the sum of all spin values, each representing either an occupied or unoccupied lattice site. Hence, M corresponds to the particle number in the grand canonical Lennard-Jones system and we see that both, $\hat{\kappa}$ and χ describe the normalised variance of the according observable.

This also means that we can determine each of the two quantities directly through the width of one of the two peaks of the respective probability distribution, as illustrated in Fig. 3.6 for the reduced compressibility. If we were to look at this distribution for an Ising gas, namely the probability of magnetisation, we would expect perfectly symmetric and equally broad peaks for the gas and liquid phase. This is due to *particle-hole symmetry*, which can be summarised by saying the Ising system is invariant under a transformation that flips all spins at once. In the gas language, this manifests by vacant sites becoming occupied and vice versa, while only the sign of magnetisation changes. One of the consequences is that both probability peaks (or both phases) yield the same value for magnetic susceptibility.

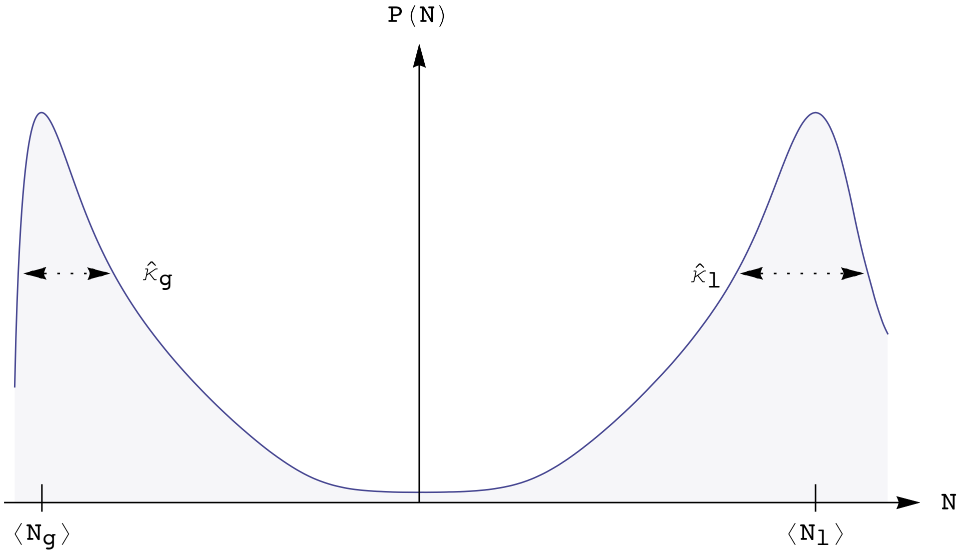

FIGURE 3.6: Logarithmic probability distribution of density, expressed through the measured particle number N at constant volume for a grand canonical simulation. The peak-width allows us to measure the respective reduced isothermal compressibility $\hat{\kappa}_g$ and $\hat{\kappa}_l$. Note the broken symmetry as well as the larger peak-width in the liquid phase (right-hand side, large N). Due to particle-hole symmetry, both peaks would be symmetric and equally broad for an Ising gas; the Lennard-Jones simulation, on the other hand, samples a partly solid phase indistinguishably from the liquid phase.

Regarding the Lennard-Jones gas now, one can see in Fig. 3.6 that the symmetry is broken. The right-hand peak, belonging to the liquid phase, is broader than the left one, which represents the gas phase. Apart from the absent particle-hole symmetry, this can be explained through the solid phase, which occurs for even larger densities and cannot trivially be extracted from the measurement of the fluid phase. In fact, another phase transition takes place and effects such as jamming and close-packing constraints become relevant. Those effects can become a problem from the simulation point of view, resulting in the necessity to finely tune the MUGC algorithm with an upper limit of density. It becomes more and more unlikely to insert particles into increasingly dense configurations.

We see in the figure that the reduced compressibility is clearly larger in the right-hand peak. Even though we expected and explained this behaviour, it certainly does not differ from the left-hand one by several orders of magnitude as seen of the *real* isothermal compressibility. Ergo, let us relate Eq. (3.58) and (3.60).

To that end, we already identified the volume per particle with inverse density $v = 1/\rho$. We can also relate the total system volume with the expectation value of the particle number via $V = \langle N \rangle / \rho$ and write Eq. (3.60) as

$$\hat{\kappa} = \frac{\rho\beta}{\langle N \rangle} \left(\langle N^2 \rangle - \langle N \rangle^2 \right) = \rho^2 \kappa \; . \tag{3.62}$$

Here, it is key to treat the density as an infinite-size quantity, which is constant for systems of any size in the respective phase. Hence, we have to treat the two peaks, phases and corresponding (real) compressibilities separately. To be more precise, we get

$$\kappa_{\mathrm{g}} = \frac{1}{\rho_{\mathrm{g}}^2} \hat{\kappa}_{\mathrm{g}} \tag{3.63}$$

for the gas phase and

$$\kappa_{\mathrm{l}} = \frac{1}{\rho_{\mathrm{l}}^2} \hat{\kappa}_{\mathrm{l}} \tag{3.64}$$

for the liquid phase. As we have reasoned, $\hat{\kappa}_{\mathrm{g}}$ and $\hat{\kappa}_{\mathrm{l}}$ are of the same order, which implies that the large difference between the real compressibilities of the two phases is caused by the density factor.

4 Results

4.1 Simulation Details

At the beginning of this chapter, we want to shortly summarise the details and parameters used in the simulations that lead to the results of this thesis. For all the presented data: $k_{\mathrm{B}} = 1$, $\epsilon = 1$, $\sigma = 1$ and the Lennard-Jones potential was limited to $r_{\mathrm{c}} = 2.5\sigma$ along with the shift by $-V(r_{\mathrm{c}})$ to compensate the offset.

First of all, we performed MUGC simulations at fixed inverse temperature $\beta = 2.5$, which corresponds to roughly $T \approx 0.9\,T_{\mathrm{crit}}$. System sizes went from $L = 5$ up to $L = 70$, while covering the full density range $0 \leq \rho \leq 0.82$ to obtain the infinite-size quantities, where $\rho = N/L^2$. A choice of the value of the chemical potential μ is not required. However, in order to obtain MUGC weights that actually converge during the iteration, it was necessary to perform shift moves (instead of doing only particle insertions and deletions). The used suggestion rates were $s_{\mathrm{shift}} = 0.5$, $s_{\mathrm{ins}} = 0.25$ and $s_{\mathrm{del}} = 0.25$. The shift solves the difficulty related to the strong decrease in the acceptance chance of insertions and deletions for high densities.

The parallel MUGC production run provided 128 time series, each of length 10^4 after thermalisation for 100 sweeps, and featured a flat probability distribution of the particle number N. Sweeps were set to scale as $m \times L^2$, usually with a prefactor of $m \approx 20$. Time-series entries were recorded after 5 sweeps, histogram entries were made after every update move. In a post-production program, the probability distribution was then reweighted to fulfil the equal-height criterion: starting with large negative values of $\mu = -5$, the flat MUGC histogram was reweighted to the final μ_{eqh}, by iteratively increasing or decreasing the chemical potential until both peaks had the same height. The resulting grand canonical distribution was then analysed, supplying estimators of ρ_{g}, ρ_{l}, $\hat{\kappa}$ and σ, the planar surface tension. This was repeated for every system size.

In the orthogonal scheme, parallel MUCA simulations were performed at a fixed density of $\rho = 0.01$. Thus, the system size could be described through the particle number, ranging from $N = 64$ to $N = 2435$. This allows a straight forward comparison to the lattice gas language. The chosen energy range was $-2N \leq E \leq 0$ with bin sizes between $\Delta E = 0.5$ and $\Delta E = 2.0$. Sweeps were defined as N^2, time-series entries were made after 5 sweeps and the energy histogram was updated after every move. Subsequent to the production run, again yielding 128 time series each of length 10^4, canonical estimates for temperatures of interest $0.25 \leq T \leq 0.4$ were gathered through reweighting.

The estimate of the constant infinite-size quantities from the full-range MUGC enabled us to guess suitable values of N in dependence of system size, which were needed to perform canonical METROPOLIS simulations. Due to the double peak nature, one can be sure that a thermalised system would stay in its current phase. For the smaller systems of $L \leq 180$, a sweep was defined again as N^2; one out of the 1.28×10^6 time-series entries was recorded after every sweep and thermalisation consisted of 1000 sweeps.

This was repeated 12 times in each phase to obtain independent estimators and errors. For larger systems, it is infeasible to produce such an amount of statistics. The time-series length was shortened to 1.28×10^4 for $L = 360$, thermalisation was cut to 100 sweeps and estimators per phase were reduced from 12 to 3. For even larger systems (work in progress), more ways to optimise need to be considered; sweeps of length N^2 become impossible to realise and it grows incredibly difficult to obtain a properly thermalised state to begin production.

Along with the METROPOLIS simulations leading to the rescaled representation, another set of parallel MUGC simulations over a now smaller density range $0 \leq \rho \lesssim 0.1$ (which roughly corresponds to $\Delta < 3$) was performed solely to reweight to the canonical ensemble. Revisiting Eq. (2.73), we see that knowledge of neither the density probability peak-position nor the chemical potential is required for this. To reduce fluctuations, the length of the 128 time series had to be increased to 10^5 and beyond, depending on the system size. Apart from gaining a smooth curve of λ as a function of the density parameter Δ to connect the individual entries from the METROPOLIS simulations, this served as a consistency check for the MUGC.

The MUGC results need to match the METROPOLIS data, which in turn was also exemplary confirmed by MUCA simulations. A systematic deviation, where MUGC underestimated λ or equivalently, the number of particles inside the droplet N_{D}, lead us to reconsider how the insertion and deletion moves work. We want the deletion to uniformly select a random particle at index i in memory out of the N currently present ones with probability $1/N$. Figure 4.1 illustrates this through histograms $H(i)$, representing a probability distribution of suggested indices i without normalisation. The blue histograms were recorded for a certain or *single* particle number each, for instance, whenever the system contained 80 particles and a deletion was proposed, we made an entry at deletion index i, yielding a distribution that represents the index probability. As expected, distributions are flat for every single but set particle number. However, looking at the *total* distribution, drawn red in the figure, we see that lower indices are more likely to be selected for deletion. The total histogram was recorded independently of the current particle number, or in other words, it is the sum of all histograms for single particle numbers.

While this behaviour is perfectly fine and consistent, it requires the particles to be spread out randomly in memory, so that the index does not have physical implications. Consequently, when inserting particles, one must not do so in a systematic manner at one or the other side of the memory container, but one rather needs to select an index for insertion arbitrarily. Our initial method: systematically inserting particles at the end, caused a relation between the index and the time (or number of sweeps), for which the particles have been in the system. Thus, *old* particles were deleted more often, causing the deviation from canonical simulations. This was a rather difficult issue to track as the grand canonical results, especially the infinite-size quantities, seem unaffected by the offset.

4.2 Fixed Temperature

The finally chosen inverse temperature for this regime as well as the grand canonical and METROPOLIS simulations was $\beta = 2.5$, which corresponds to about $T \approx 0.9\,T_{\mathrm{crit}}$. The implications of the chosen temperature are sketched in Fig. 4.2, showing probability distributions of density for different inverse temperatures. Starting sufficiently below

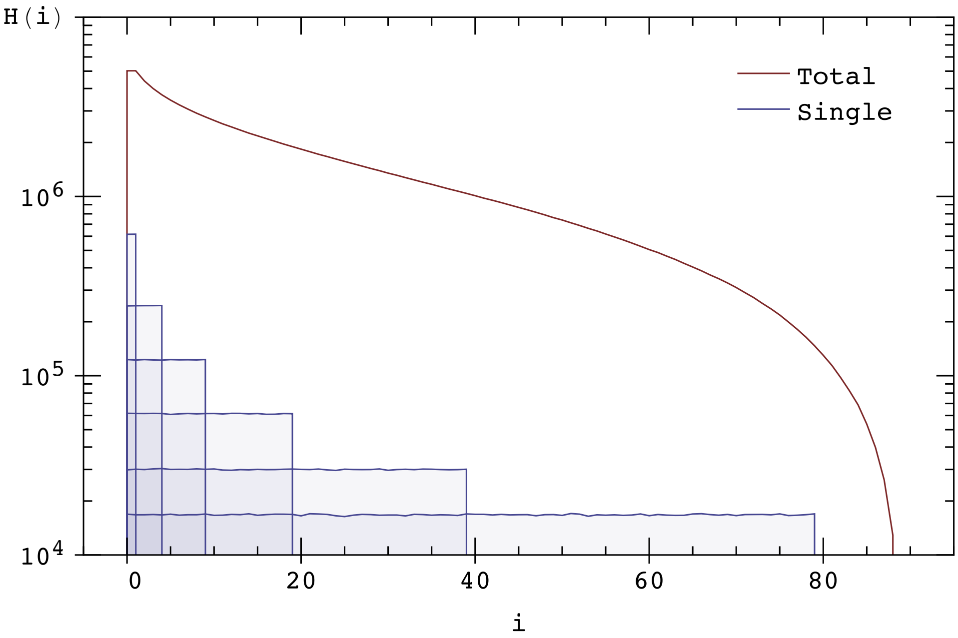

FIGURE 4.1: Histograms $H(i)$ of the randomly selected index i of a particle deletion. **Blue:** For a certain or *single* number of currently present particles, the distribution is flat. **Red:** The *total* distribution, recorded for all possible particle numbers, shows that lower indices are selected more often. Hence, the insertion move must not insert particles systematically at the end of the memory container but rather choose the index for insertion randomly. Otherwise, *old* particles have lower indices and are deleted more often than *recently* inserted ones, causing a deviation from canonical results.

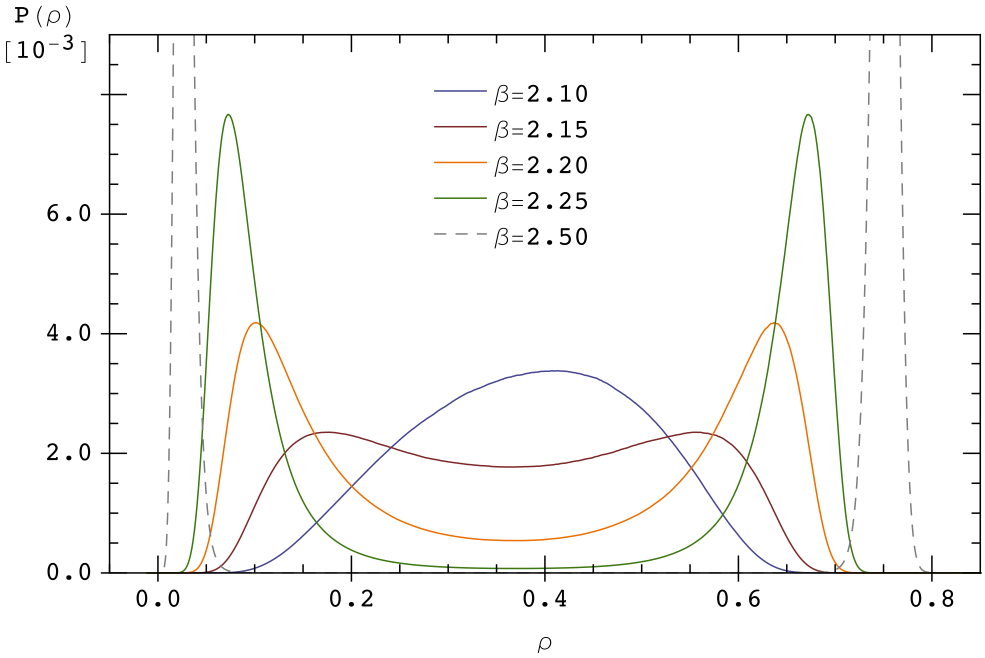

FIGURE 4.2: Probability distribution of density for different temperatures, reweighted to the equal-height chemical potential. For decreasing inverse temperatures β, the peak positions shift inwards and the probability difference between favoured and suppressed states reduces. Past the critical point, the equal-height criterion cannot be fulfilled and any chosen chemical potential produces a single peak. Our final choice for the inverse simulation temperature of $\beta = 2.5$ is sufficiently far away from the critical point (grey, dashed).

the critical point, or in other words, with $\beta = 2.5 > \beta_{\text{crit}}$, we see the well established double-peak nature, characteristic for phase coexistence.

As the inverse temperature is decreased and the critical point is approached, the two peaks become less pronounced and intermediate states in mixed phases become more and more likely to be observed. Additionally, the peak positions wander inwards. This is consistent with Fig. 3.2, which back in the theory section, illustrated the transition in the density-temperature plane: as the (real) temperature is fixed at increasingly large values, the liquid and gas density of the infinite-size system grow closer and closer until the critical temperature is reached and the phases are not distinguishable anymore.

Evidently, we observe only one individual peak above the critical temperature, the position of which is not defined in our context. Note that the peak of the blue curve ($\beta < \beta_{\text{crit}}$) in Fig. 4.2 is not located where it could be expected from symmetry. Our criterion to find the chemical potential, which identifies the phase transition through the equal height of two separate peaks, cannot be fulfilled and the single-peak position remains arbitrary. However, it should be mentioned that it is in fact possible to obtain a suitable value for the single-peak chemical potential by extending gathered points along the measured scaling directions in μ-T space [25, 26].

Chemical Potential

Since we are not primarily interested in the critical point, we return to the chosen temperature $\beta = 2.5$ and proceed to showcasing some possible situations that can be encountered when looking for the equal-height chemical potential μ_{eqh}.

The reweighting program starts at some negative and large chemical potential $\mu_{-} < \mu_{\text{eqh}}$, reweights to that value and measures the height and the position of the (at first) single peak, positioned far on the left. In figure 4.3 this corresponds to the feint yellow peak near $\rho \approx 0$. Depending on which of the (possibly two) peaks below and above a (manually chosen) threshold density is larger, the iteration either increases or decreases μ. If only one peak is present, increasing the chemical potential will move it to the right. Similarly, when two peaks are already present, raising μ enlarges the right-hand one and lowers the left, while the respective positions stay roughly the same. When the chemical potential is build up to exceed μ_{eqh}, increasingly dense states are favoured. This includes mixed phases consisting of solid and liquid contributions. In the figure we can identify this behaviour by the occurrence of another double peak-like structure for larger densities near $\rho \approx 0.8$. Due to the manually set upper boundary, the solid states are not sampled by MUGC and the rightmost peak appears to be cut off. While it would certainly be interesting to also investigate the sold-liquid transition, the grand canonical method involving particle insertions is not suited for the task.

System Size

In order to support the subsequent plots of the infinite-size quantities, another set of probability distributions is shown in Fig. 4.4, this time for differently sized systems with otherwise equal parameters.

As we assumed before, the peak positions indeed match for all system sizes, leading to similar, almost size independent estimates for the background gas and liquid density. The peak width on the other hand, lessens as L is increased. Nonetheless, due

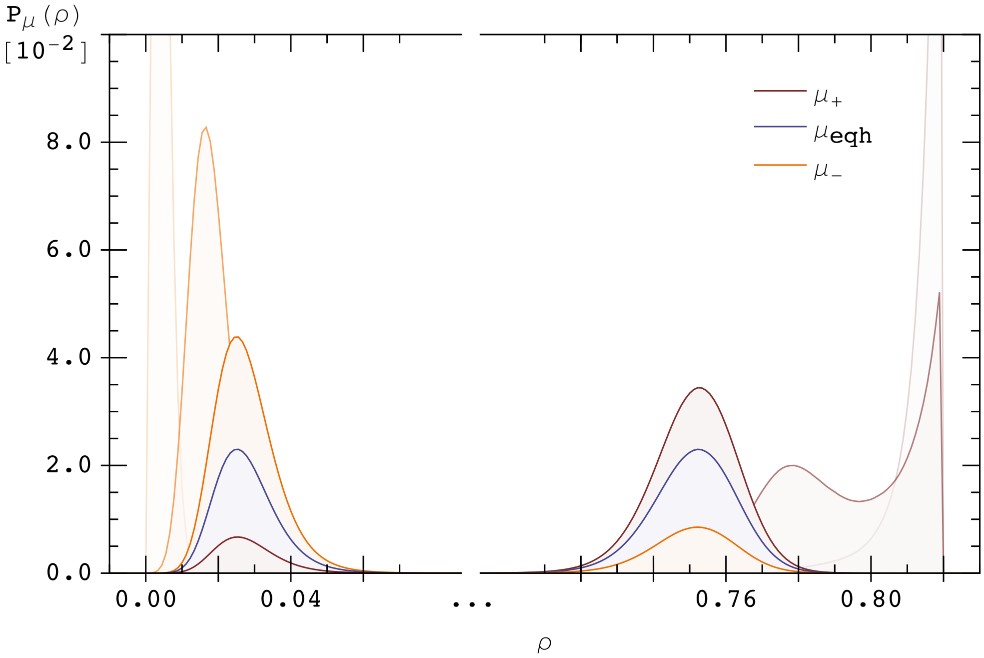

FIGURE 4.3: Influence of the chemical potential on the probability distribution of density. Values further away from equal height in either direction are drawn increasingly transparent. **Yellow:** For large negative values $\mu_- < \mu_{\text{eqh}}$, particle insertions become unlikely, thereby increasing the likelihood of the gas phase, which is observed through the larger left-hand peak. **Blue:** At $\mu = \mu_{\text{eqh}}$, the two peaks have the same height and both phases are equally likely. **Red:** For least negative chemical potentials $\mu_+ > \mu_{\text{eqh}}$, dense states are favoured. Note the suppressed region near $\rho \approx 0.8$ where a liquid-solid transition takes place.

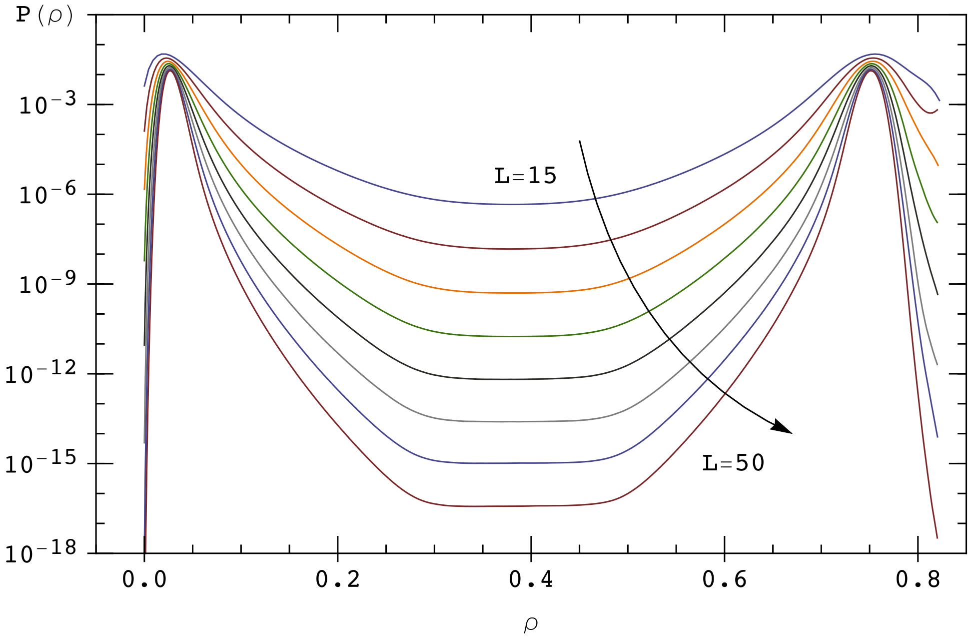

FIGURE 4.4: Probability distribution of density for different linear system sizes L. The peak positions, representing the expectation values of the densities of the homogeneous phases, stay constant. In the region of suppressed states, note the stepwise change in slope for the largest systems as opposed to the smooth transition for small systems. The systematically located slope changes indicate the sub-transitions between different phases.

to the volume scaling factor in the definition of the reduced isothermal compressibility, the respective estimates will also turn out to be size independent. Unmistakably, the probability suppression of mixed phases heightens in proportion to system size. While the suppressed interval seems like one smooth region for small systems, we can clearly identify individual sections for the largest systems. Having a close look at the curve for $L = 50$, one can see two rather sudden changes in slope to either side of the flat minimum, before reaching the peak positions. Those regions indicate the subtransitions: starting in the left peak with a pure gas phase, the systems transitions to the mixed phase featuring a liquid droplet with surrounding gas, superseded by the striped liquid-gas phase at the probability minimum. As the density is increased further, a so-to-say inverse droplet phase occurs, consisting of a gas bubble in a liquid surrounding, before the pure liquid phase is attained. For snapshots of those phases, please refer to Fig. 3.3 in Sec. 3.2.1.

4.2.1 Infinite-Size Quantities

As we have stressed before, certain quantities or observables of the system stay constant at fixed temperature, independently of the system size. Consequently, they match the value that the infinite-size system would yield, accounting for this section's title. Apart from a general insight into the system, they provide the means to formulate the density and the particle excess through the dimensionless parameters Δ and $\lambda(\Delta)$, respectively.

The first infinite-size quantities of interest to us are the background density of the pure liquid phase ρ_{l} and the pure gas phase ρ_{g}, individually depicted in Fig. 4.5 a) and b), respectively. We obtain the values from the corresponding peak positions of the probability distributions by calculating the expectation value for each phase (separated by the probability minimum) and assume that each peak approximately resembles a normal distribution. Errors stem from jackknifing the 128 threads, as outlined in Sec. 2.5. Note the different magnitude of the two densities. While we encounter strong finite-size effects for small systems of $L < 30$, the fit to $\rho = \mathrm{const}$ works rather well beyond that threshold. Using the GNUPLOT fit routine with corrected errors for $30 \leq L \leq 70$ we obtain

$$\rho_{\mathrm{l}} = 0.75085(4) \tag{4.1}$$

$$\rho_{\mathrm{g}} = 0.027871(2)\ . \tag{4.2}$$

In the same spirit, we measure the width of each peak as the variance around the peak position. This yields the two estimates for the reduced isothermal compressibility, $\hat{\kappa}_{\mathrm{g}}$ and $\hat{\kappa}_{\mathrm{l}}$, where only the former is used for the rescaling. Note the stronger fluctuations of the peak width of the liquid phase (responsible for $\hat{\kappa}_{\mathrm{l}}$) in Fig. 4.5 d). At a first guess, this can be explained by the partial solid-phase contributions to that peak, skewing the distribution and rendering the assumed Gaussian approximation less accurate. The fit for $L \geq 30$ yields

$$\hat{\kappa}_{\mathrm{l}} = 0.3123(9) \tag{4.3}$$

$$\hat{\kappa}_{\mathrm{g}} = 0.17116(6)\ . \tag{4.4}$$

In order to estimate the normalised surface free energy of the droplet τ, we actually measure the planar surface tension σ and map it onto a (2D) sphere. This works quite well for the Lennard-Jones system, as the surface tension of the round droplet can be

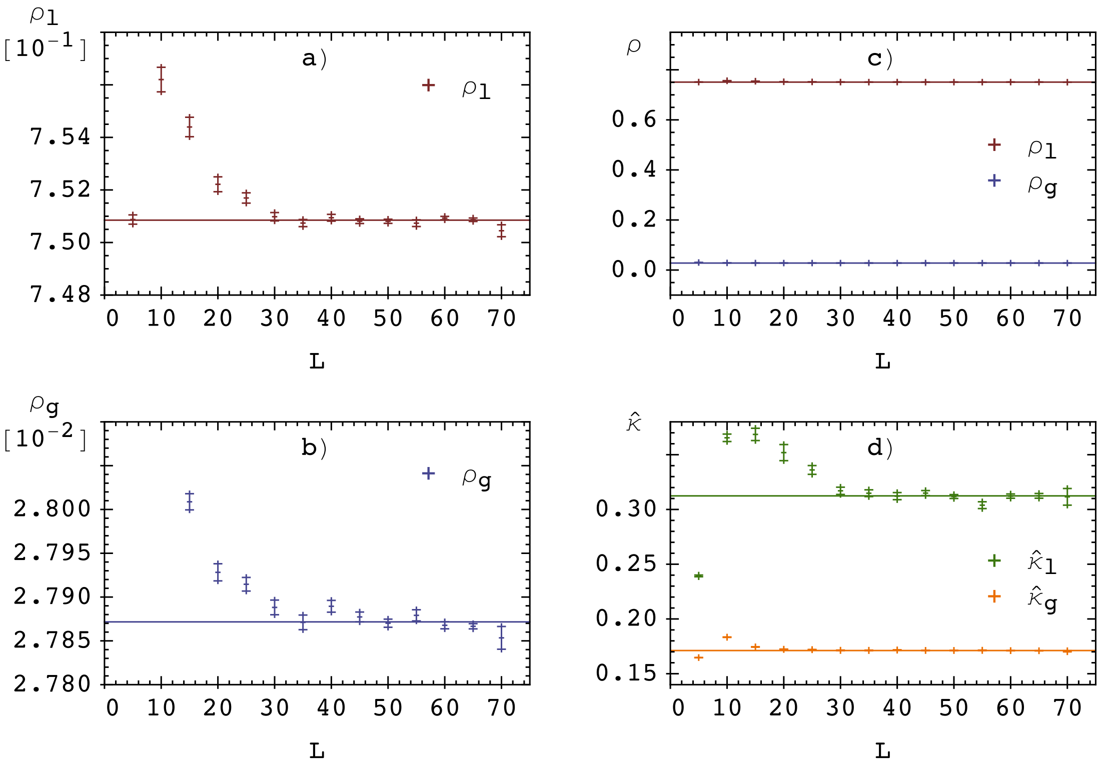

FIGURE 4.5: The infinite-size quantities almost stay constant as the system size scales: **a)** Liquid density - **b)** Gas density; note the different magnitudes and quickly vanishing finite-size effects of the two respective densities - **c)** Combined plot of both densities - **d)** Reduced isothermal compressibility, as obtained from the gas peak of the probability distribution ($\hat{\kappa}_{\mathrm{g}}$) and for completeness also from the liquid peak ($\hat{\kappa}_{\mathrm{l}}$).

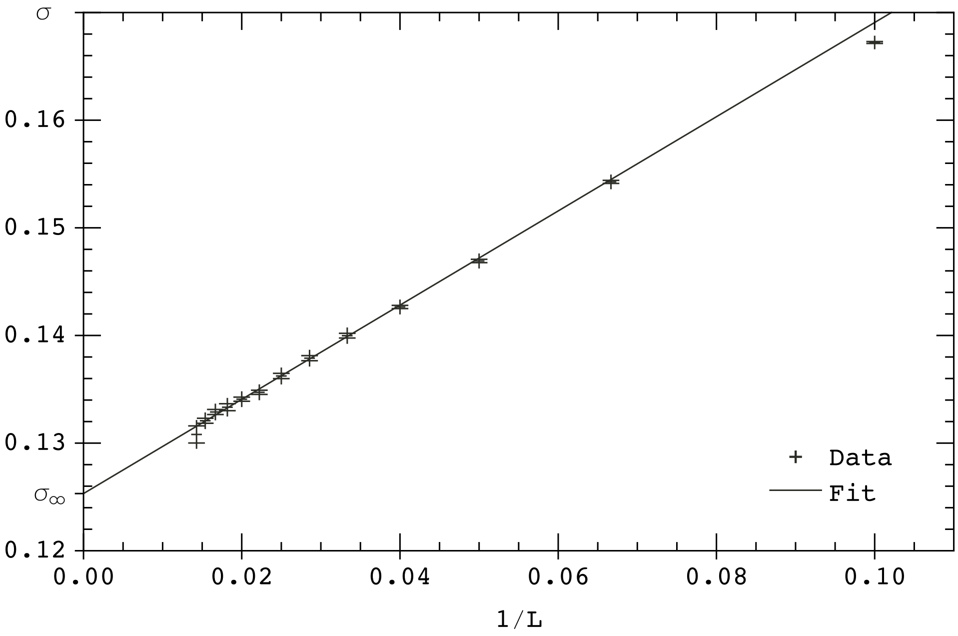

FIGURE 4.6: Scaling of the planar interface tension σ for systems of size $10 \leq L \leq 70$. In the plotted units, a linear fit was employed, ignoring possible higher-order or logarithmic corrections [54].

assumed to be isotropic for all angles, in contrast to lattice systems [49, 50]. For those, the contributing planar surface tension varies with angle; think of a diagonal interface compared to one parallel to the system boundary of a cubic lattice.

We then follow the method first outlined by Binder [48] to measure the probability of the pure-phase peaks P_{max} and the probability minimum P_{min}, in between. At the minimum, one finds a mixed, strip-like phase consisting of a liquid and a gas contribution that are separated by an interface. Due to periodic boundary conditions and the second law, the length of the two interfaces of the stripe matches the system size L. This can be exploited to measure the planar interface tension as

$$\sigma = \frac{1}{2\beta L} \ln \left[\frac{P_{\mathrm{max}}}{P_{\mathrm{min}}} \right] . \tag{4.5}$$

Figure 4.6 shows this behaviour for differently sized systems. The intersection with zero indicates the infinite-size value

$$\sigma_{\infty} = 0.1253(4) . \tag{4.6}$$

From the planar interface tension, we can deduce the surface tension of the droplet [22] via $\tau = 2\sqrt{\pi}\sigma_{\infty}$ and estimate

$$\tau = 0.444(2) . \tag{4.7}$$

While σ clearly shows a size dependence, we only use the infinite-size estimate for the rescaling. This is reasonable, since we are only interested in the leading-order behaviour and the theoretic derivation of the dimensionless parameters is tailored around this aim. Including the corrections to the interface tension would require the consistent inclusion of corrections for all other quantities.

4.2.2 Transition in Reduced Parameters

Using our estimators for the constants in this regime, we can now rescale the directly measured quantities describing the condensation. Namely, the particle number N is the measure of density, and the number of particles within the droplet N_{D} describes the droplet excess. From Eq. (3.9) we see that in $d = 2$ dimensions and with $\hat{\kappa} = \hat{\kappa}_{\mathrm{g}}$, we can express the dimensionless density parameter as

$$\Delta = \frac{\sqrt{\rho_{\mathrm{l}} - \rho_{\mathrm{g}}}}{2\hat{\kappa}_{\mathrm{g}}\tau} \frac{(N - \rho_{\mathrm{g}} V)^{\frac{3}{2}}}{V} . \tag{4.8}$$

Similarly, the fraction of particle excess within the droplet $\lambda = \delta N_{\mathrm{D}} / \delta N$ can be written as

$$\lambda = \frac{(\rho_{\mathrm{l}} - \rho_{\mathrm{g}})\, N_{\mathrm{D}}}{(N - V\rho_{\mathrm{g}})\, \rho_{\mathrm{l}}} , \tag{4.9}$$

where we assumed that the surrounding vapour is of density ρ_{g}, while the droplet of volume $V_{\mathrm{D}} = N_{\mathrm{D}} / \rho_{\mathrm{l}}$ has the liquid density ρ_{l}. Since we are using the native units and the lengthily discussed compressibility (normalised with volume), a correction factor to tweak Δ into the units intended by Biskup et al. is not required. A direct comparison with the literature is possible. The results are shown in Fig. 4.7. Individual data points were obtained from METROPOLIS simulations; the continuous lines of manageable system sizes stem from MUGC. The analytic solution [22]

$$\lambda = \begin{cases} \frac{4}{3}\cos^2\left(\frac{\pi - \arccos\left(\frac{3\sqrt{3}}{8\Delta}\right)}{3}\right) & \Delta > \Delta_c \\ 0 & \text{else} \end{cases}, \tag{4.10}$$

is also shown.

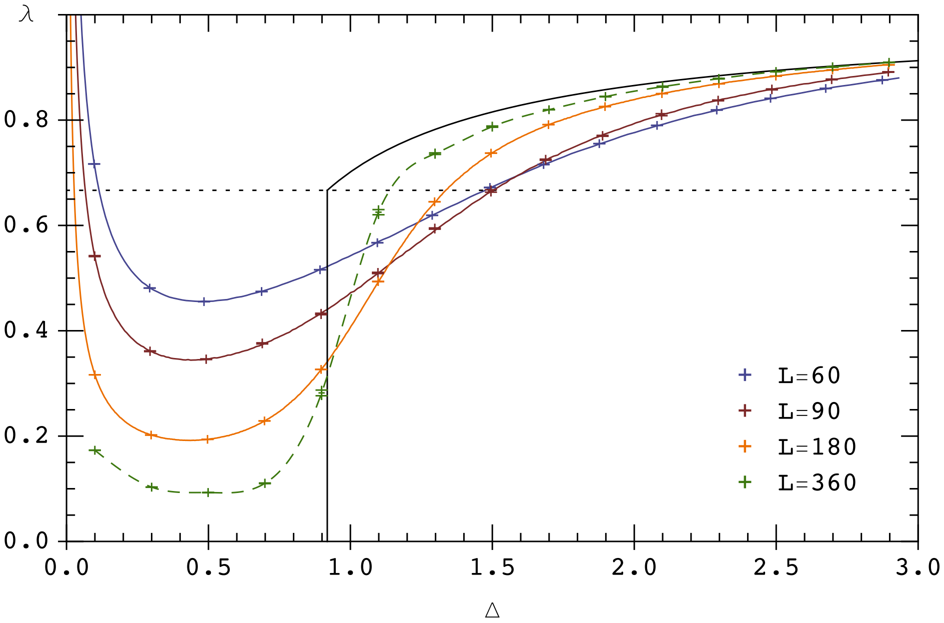

FIGURE 4.7: Transition in reduced parameters: MUGC simulations (yielding the solid lines) were only realisable up to $L = 180$ but match the METROPOLIS results (data points with error bars). The dashed line for $L = 360$ is a guide to the eye, only. The dimensionless density parameter Δ includes leading-order corrections, and λ is the excess fraction within the droplet. As anticipated, the smearing-out of the transition decreases for large systems. The measured data approaches the analytic solution (black line), confirming the rescaling. Note that the estimates for the infinite-size quantities influence the collapse point and could cause skewing in Δ and λ. We want to point out the strong overestimation of excess within the droplet as a finite-size effect, especially in the gas phase. For very small systems (e.g. $L = 30$, not shown here) the measured λ even exceeds the transition value for all densities. The reader is kindly asked to excuse the clumsy choice of system sizes.

The figure visualises that the behaviour predicted by theory is supported by the data from our Lennard-Jones system. Furthermore, we can observe a good qualitative agreement with comparable plots obtained for the lattice models [22, 32]. Especially for the largest systems, we have no excess λ within the droplet, as long as the rescaled density stays below the threshold $\Delta_c \approx 0.91$. Beyond the threshold, droplet formation occurs and a large fraction of excess is located inside the droplet, almost matching the analytic prediction $\lambda \geq \lambda_c = 2/3$. All in all, we are inclined to confirm both, the rescaled transition density Δ_c and the transition excess fraction λ_c. However, in order to validate this claim, a detailed finite-size scaling analysis should be undertaken. Continuing our quantitative description, we see the anticipated smearing-out or rounding of the transition for small systems, which is accompanied by a systematic overestimation of the excess within the droplet. For instance at $L = 30$ (not shown), the excess actually exceeds the threshold and $\lambda > \lambda_c$, over the whole density region. Lastly, the positions in Δ of the METROPOLIS simulations seem to fluctuate or to be ill chosen. This is due to the mapping from the discrete scale (N) to the continuous scale (Δ) and vanishes for sufficiently large system sizes. The continuous lines from MUGC, on the other hand, provide a better resolution; every possible N is covered, but immense computational effort is required to provide a *smooth* curve.

4.2.3 Finite-Size Behaviour at Fixed Temperature

For completeness, we want to outline the procedure to measure the finite-size scaling of density [32], which has not been completed at the time of writing due to limited system sizes. As we have seen in Fig. 4.7, the leading-order behaviour is satisfyingly covered by the rescaling. Still, we can presume that higher-order corrections exist, since the data sets do not cross or collapse at a size-independent density. Thus, measuring the actual finite-size transition density $\Delta_c(L)$ may prove helpful to identify those corrections. There are two apparent ways to do so: the peak position of the first derivative could be used to pinpoint the transition, which requires a rather fine resolution in density. Hence, this method may work for MUGC simulations but not so well for the few data points provided by METROPOLIS. Either way, a discrete derivative has to be employed. Alternatively, one would define the transition through the intersection of $\lambda(\Delta)$ with the analytic excess threshold $\lambda_c = 2/3$. From there, one can *read of* a reduced density value, corresponding to a size-dependent particle number $N_c(L)$, in the vicinity of which the transition should take place. Additional simulations with either technique should be performed, to enclose the final value from both sides. Errors are then obtained through the closest points above and below, where precision is still limited by discrete N.

4.3 Fixed Density

A considerable amount of work evolving around the fixed-density scheme has been done in the recent years, presenting sophisticated methods for treating the condensation-evaporation transition on various models [30–32]. This regime allows to natively harness the full potential of MUCA with subsequent reweighting to the canonical ensemble. In accordance with the given references, we chose the density to be $\rho = 0.01$.

Furthermore, we adapt the convention of expressing system size in terms of the particle number N, in order to ease the comparison with lattice systems. A typical probability distribution of energy is shown for a few sizes in Fig. 4.8. Having mostly treated distributions for varying particle numbers so far, where the two peaks belong to the homogeneous liquid and gas phase, we want to point out that the shown peaks now represent the homogeneous gas phase (right peak) and the mixed phase of a macroscopic liquid droplet with surrounding gas (left peak).

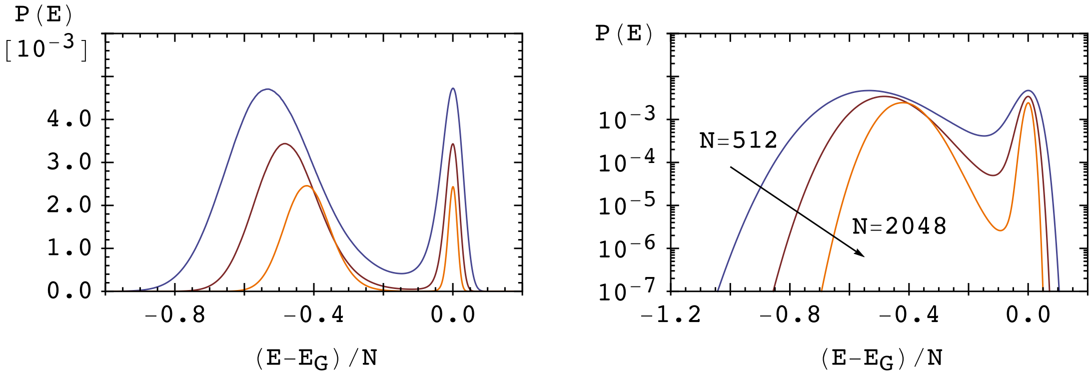

FIGURE 4.8: Probability distribution of Energy at β_{eqh} on a linear and a logarithmic scale. In contrast to previous distributions, the peaks correspond directly to the evaporation-condensation transition, where the high-energy peak at $E = E_{\mathrm{g}}$ stands for the pure gas and the low-energy peak represents the droplet with surrounding vapour. Again, the suppression of states increases with system size.

To determine the transition temperature T_{c}^{N}, the same approach as described for locating the equal-height chemical potential could be used. However, it was shown that the obtained infinite-size transition point (thermodynamic limit) is the same for different definitions of the finite-size transition temperature [55]. Since temperature is a continuously defined variable in this regime, opposed to density in its respective scheme, we preferred the procedure where the transition is located through the peak position in the first thermal derivative.

We want to support with Fig. 4.9 and 4.10 that the peak positions as well as the shown rounding are in fact the same for both observables E and η. Furthermore, when inspecting the droplet fraction $\eta = N_{\mathrm{D}}/N$, we can reconfirm the effects of the transition that were observed in the fixed-temperature scheme. Evidently, η is not a rescaled observable of temperature, hence, the transition temperature shows finite-size behaviour. Nonetheless, we can observe a similar transition rounding for small systems and the previously highlighted overestimation of droplet size across the whole range of the driving variable, which is now temperature. Apart from the estimate for droplet size being too large, which is mostly apparent in the high-temperature (gas) phase, we still have the same behaviour: above the transition point, almost no particles are in the

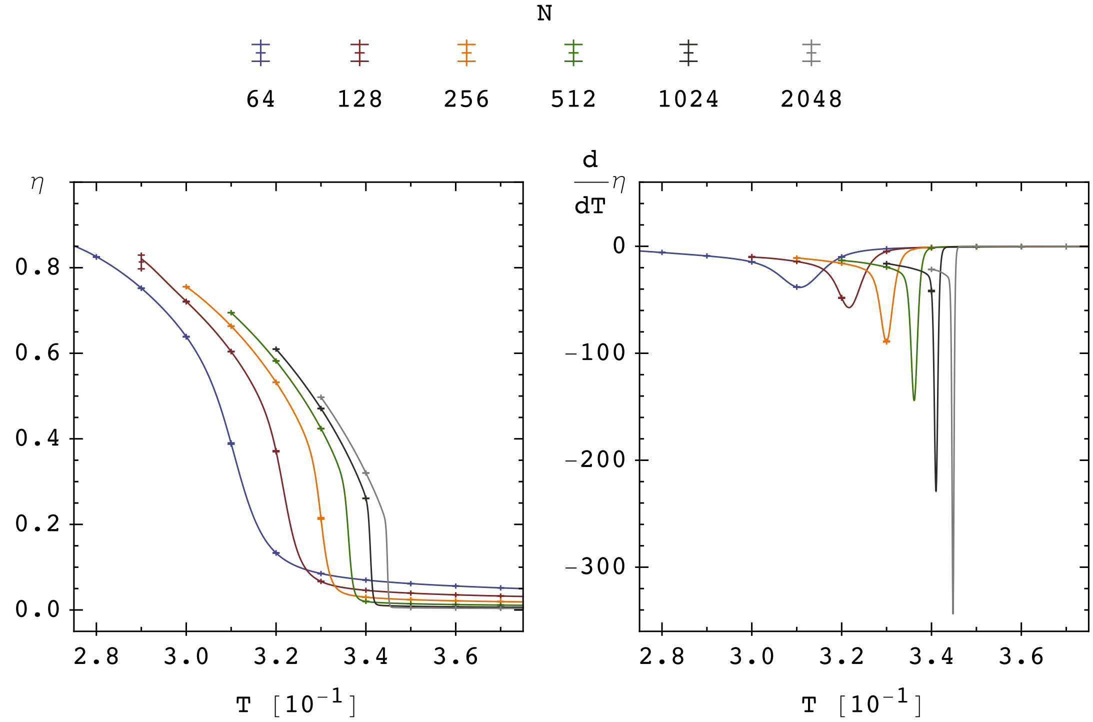

FIGURE 4.9: Droplet Fraction $\eta = N_\mathrm{D}/N$ and its first thermal derivative for increasing system sizes at fixed density. Note the *smearing-out* of the transition as a clear finite-size effect. Qualitatively, η resembles the behaviour observed in Fig. 4.7: for small systems, the droplet size is overestimated across the whole regime, including the gas phase, where no droplet is expected.

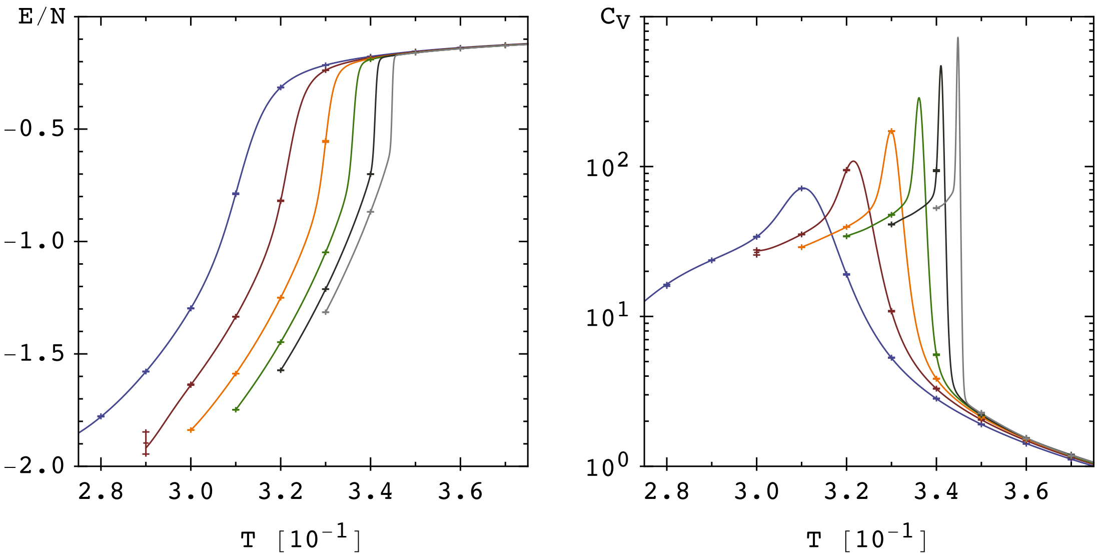

FIGURE 4.10: Energy and its first thermal derivative, the specific heat, here normalised with $1/N$. The expected first-order behaviour is apparent, including the pronounced step-like discontinuity in energy and the increasingly sharp peaks in the specific heat. Note that peak positions, from which T_c^N is obtained, nicely coincide with those of $d\eta/dT$.

droplet, but as soon as the condensation point is crossed, a suddenly very large droplet is present containing in fact *most* of the particles.

4.3.1 Finite-Size Behaviour at Fixed Density

We look at the finite-size behaviour for the transition temperature T_c^N and the transition rounding $\Delta T = \Delta T(N)$. Especially, we want to obtain an estimate for the infinite-size transition point $T_\mathrm{g} = T_\mathrm{c}^\infty$, which is commonly denoted by T_0 in the Ising language.

Using $d = 2$ dimensions, the fit to *leading order* was done as

$$\Delta T = a' N^{-\frac{d}{d+1}} \tag{4.11}$$

$$T_\mathrm{c}^N = T_\mathrm{g} + a N^{-\frac{1}{d+1}} \ . \tag{4.12}$$

Including the next *higher order*, we fitted

$$\Delta T = a' N^{-\frac{d}{d+1}} + b' N^{-\frac{2d}{d+1}} \tag{4.13}$$

$$T_\mathrm{c}^N = T_\mathrm{g} + a N^{-\frac{1}{d+1}} + b N^{-\frac{2}{d+1}} \ . \tag{4.14}$$

Additionally, the *intermediate* regime was fitted to

$$\Delta T = a' N^{-1} \tag{4.15}$$

$$T_\mathrm{c}^N = T_\mathrm{g} + a N^{-\frac{1}{d}} \ . \tag{4.16}$$

We start the discussion with the rounding ΔT, illustrated in Fig. 4.11. While we lack an analytic solution for T_g, we know that the transition rounding approaches zero in the thermodynamic limit. Moreover, in the chosen log-log plot, the leading-order scaling ($N^{-d/(d+1)}$, blue) and the intermediate fit (to N^{-1}, yellow) are both straight lines, enabling intuitive comparison. It is apparent that the leading-order behaviour is not yet exhibited, even for our largest systems up to $N = 2435$. Since no direct fit was possible, the blue line only represents a guide to the eye, the slope of which was taken from the fit including the next higher order (red line). The consideration of the next higher order is done empirically, but seems justified, as intermediate and large system sizes beyond $N = 215$ are matched within error bars. The intermediate fit, on the other hand, motivated by the general scaling of first-order transitions [56], starts to deviate for systems larger than $N = 724$. The slope is slightly too steep.

Let us have a look at the scaling of T_c^N now. In the intermediate regime, systems are still rather small, which causes the previously stressed overestimation of the droplet fraction. The reason is that almost all particles of the system contribute to the condensate, which is an effect typically observed for homopolymere aggregation [57]. In our case, the largest droplet contains at least one particle, after all. Subsequently, this causes the linear droplet size to scale as $R \propto N^{1/d}$. Although we have not shown the explicit dependence on the radius, this serves as a good motivation for the intermediate scaling of T_c^N with $N^{-1/d}$. As we see in Fig. 4.12, we are able to leave the intermediate regime, at least with our Lennard-Jones system. This is usually not the case for other systems, where interactions are increasingly expensive to compute.

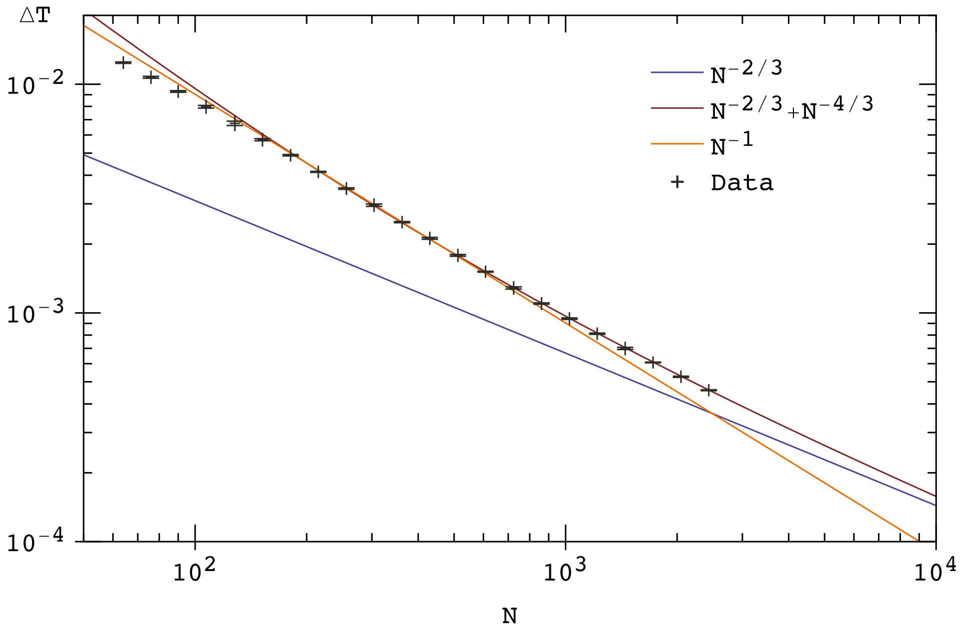

FIGURE 4.11: Finite-size scaling of the transition rounding, obtained from the half-width of the specific-heat peaks, on a log-log scale. **Blue:** Realisable systems are too small to show the leading-order behaviour. We provide a guide to the eye as a mere outlook. **Red:** Including the next higher order, intermediate and large systems are covered. **Yellow:** Within the intermediate regime, a fit to the scaling ansatz N^{-1} for first-order transitions [56], is also possible. Furthermore, this describes smaller systems slightly better than our empirical inclusion of the second order. If data for the largest systems was not available, this ansatz could have falsely been assumed to work on all length scales.

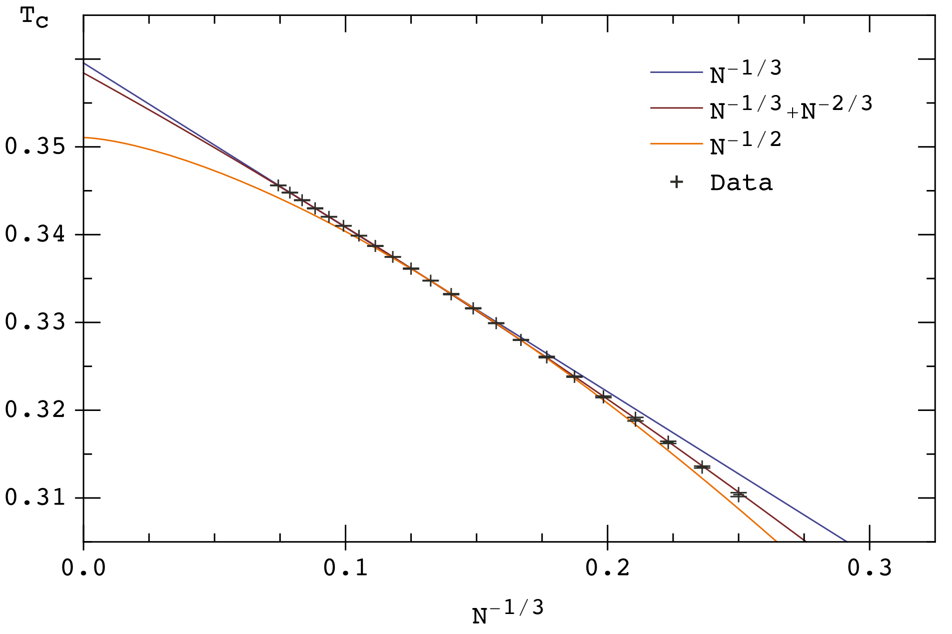

FIGURE 4.12: Finite-size scaling of the transition temperature, as obtained from the peak positions of the specific heat. **Blue:** The fit to leading order strongly deviates for small systems. **Red:** Including the next higher order into the fit, all data points are covered within error bars. **Yellow:** Up to medium system size, a fit to $N^{-1/d}$ describes the scaling; due to rather limited particle numbers, almost all particles contribute to the droplet. However, having access to larger systems, we see that this ansatz is misleading and should only be considered for the intermediate regime.

A fit of T_c^N to leading order ($N^{-1/(d+1)}$, again indicated in blue) is possible and already falls within error bars for systems with $N \geq 256$. This clearly highlights that the behaviour shown in the intermediate regime is not globally valid; especially the outlook of the thermodynamic limit and T_g is misleading (yellow line). Nonetheless, even this statement has to be considered with care; as no analytic solution exists for our system, any predictions remain questionable. Having this in mind, we only provide numerical results for the leading-order fit, even though the fit including higher-order corrections (red line) looks superior.

Taking into account all data points for systems $N \geq 256$, the least-square fit to leading order yields

$$T_\mathrm{g} = 0.35962(4) \;, \tag{4.17}$$

with $\chi^2 \approx 1.5$ and $Q \approx 0.005$. Excluding, for instance, one more system size, thereby using systems $N \geq 304$, decreases the $\chi^2 \approx 1.1$ and puts $Q \approx 0.27$. This smaller fit range yields $T_\mathrm{g} = 0.35955(4)$ as the transition temperature.

A Different Representation of the Transition Rounding

The previously discussed finite-size plots become cluttered very quickly. Identifying different scaling regimes can be especially tricky. To that end, we came up with another representation, see Fig. 4.13 and 4.14. Here we create a multitude of fits to

$$\Delta T = cN^\alpha \;, \tag{4.18}$$

each of which belongs to a different range $[N_\mathrm{min} : N_\mathrm{max}]$ of included system sizes. We then map the exponent α to a colour scale, which allows to visually distinguish the scaling regimes; dark blue belongs to a scaling with ($\alpha = -1$), as expected for intermediate systems, and red indicates the respective leading-order regime ($\alpha = -d/(d+1)$). In those charts one would expect the infinite-size limit in the upper-right corner, i.e. including only the largest available systems into the fit. A fit over the full range, containing all available data, is located in the very bottom-right, while the lower-left treats only the small systems. The additional data sets for the 3D Lennard-Jones, as well as 2D and 3D lattice models, were provided by Johannes Zierenberg [30] and are greatly appreciated.

Starting with the 2D Lennard-Jones system in Fig. 4.13, we can clearly identify the intermediate scaling behaviour and the beginning of the transition towards the leading-order behaviour for large systems. The 2D lattice model shows the same trend, but neither of the two actually reaches the leading-order regime. Additionally, we experience heavy finite-size effects on the lattice (bottom-left), where the fit exponents do not match our defined domain. Interestingly, the exponent actually exceeds the expectation ($\alpha = -1$) in the intermediate regime of the (2D and 3D) lattice systems.

When looking at the 3D systems in Fig. 4.14, we notice the leading-order behaviour is actually exhibited. Of course, since the leading-order exponent in three dimensions is numerically closer to the intermediate exponent, the colour scale covers a smaller domain and thereby *suggests* the leading-order behaviour more dominantly. Nonetheless, the additional degree of freedom helps to reach the infinite-size behaviour quicker than in two dimensions. A similar train of thought allows to reason that the 3D Lennard-Jones system transitions from the intermediate to the leading-order behaviour at smaller system sizes than the 3D lattice system; the continuous models are less sensitive to the

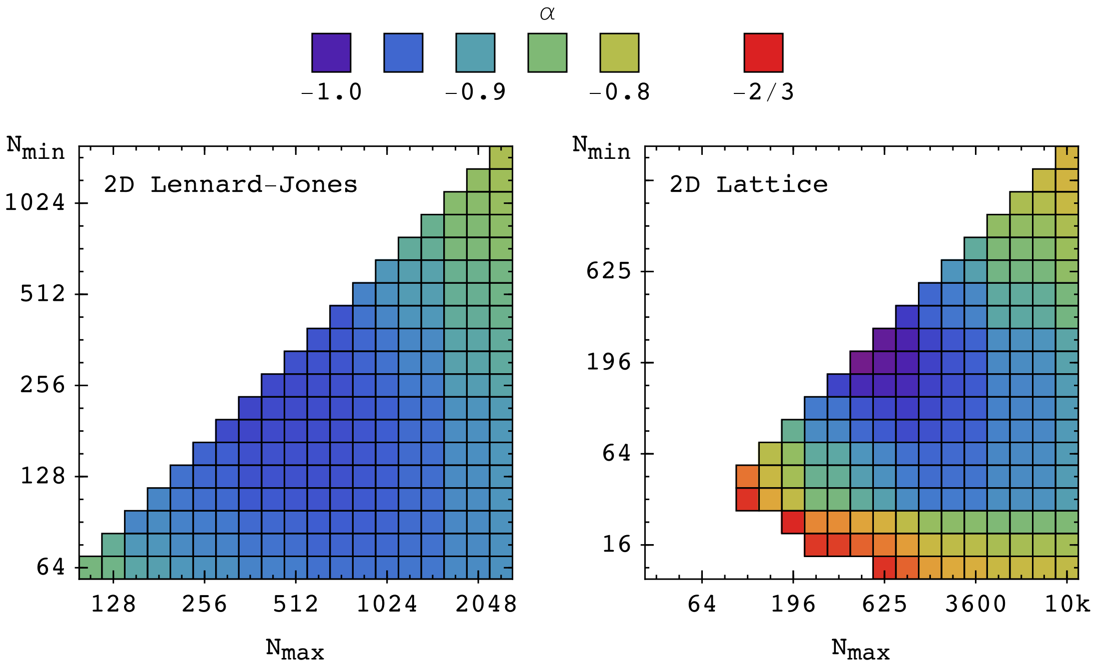

FIGURE 4.13: Matrix representation of the transition rounding ΔT for a fit to N^{α} in two dimensions. This visualises the exponent α obtained from different fit ranges and thereby allows to identify scaling regimes. While the lattice system can be simulated for way larger particle numbers, strong discretisation and finite-size effects occur for small systems. In both models, α clearly matches the intermediate regime (dark blue) for medium system sizes. The leading order behaviour (red) is not reached.

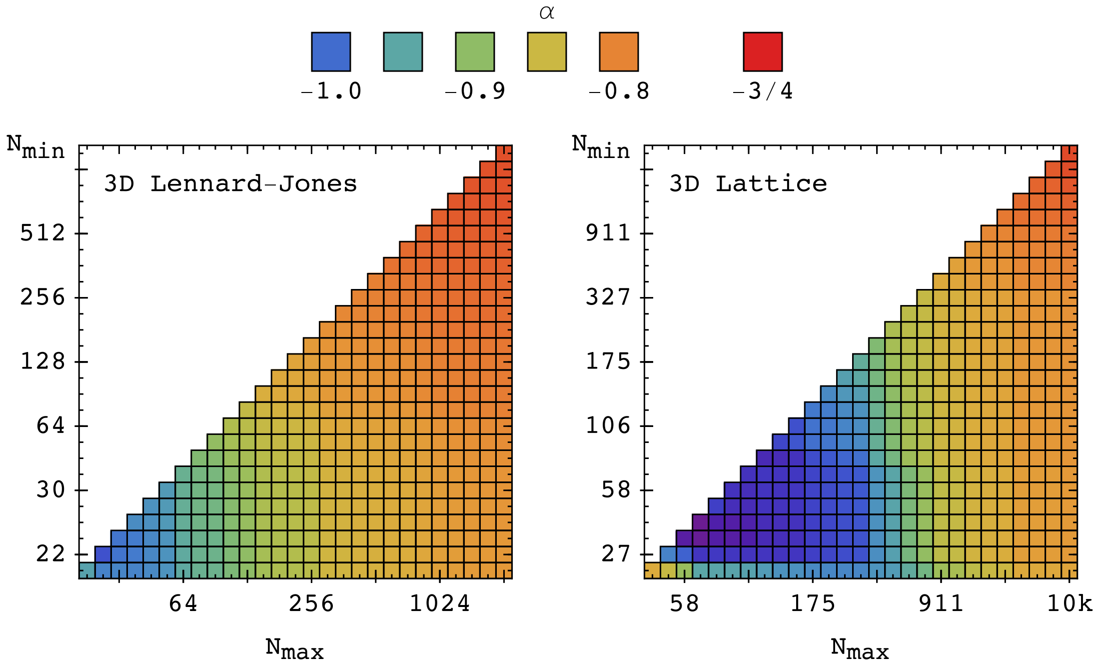

FIGURE 4.14: Matrix representation of the transition rounding in 3D. Note the different colour scale. Numerically, the leading-order exponent (red) is closer to the intermediate regime than in two dimensions. In fact, both models show leading-order behaviour for fits including only very large systems (top-right). Due to the additional degree of freedom, the discretisation effects on the lattice system are less severe. We thank Johannes Zierenberg for providing the data on the remaining three models as was published in [30].

discretisation and to configurational constraints, leading to a rather *smooth* change of colour.

Considering all four charts at once, we see that both 3D systems exhibit the intermediate scaling for much smaller particle numbers than the 2D equivalents. Hence, we have to assume that still larger sizes are necessary to actually measure the leading-order behaviour in two dimensions, which is computationally not feasible at present. Especially our 2D Lennard-Jones system is plagued by restraining factors; not only that we are confined, well, to two dimensions, but the interactions are costly too and, thereby, limit the particle number to a fraction of what is possible on the lattice.

5 Conclusion

After having given an overview of the required computational methods, we recapitulated the theory around the condensation-evaporation transition in the two different schemes. In order to acquire data at fixed temperature, allowing density fluctuations, we implemented custom Monte Carlo simulations in the grand canonical ensemble. The method is rather similar to other realisations [15, 28], but the program itself is set apart by the tight integration with the iterative MUCA algorithm. This leads to a strong similarity between the new MUGC and the original MUCA code and output; allowing easy comparison, maintenance and analysis. Well established tools such as time-series and histogram reweighting, jackknife binning and finite-size scaling were then employed to obtain the infinite-size quantities. Those estimates were required to create the rescaled and dimensionless representation in the spirit of the discussed theory. Not only could we thereby confirm the theoretic predictions and results from lattice systems [21–23, 32], but this thesis actually complements those findings with an off-lattice system.

Having said so, the presented framework helps closing the gap between the fixed-temperature and the fixed-density approach; the employed grand canonical simulations gave us the opportunity to follow the formalism of Biskup et al. [19, 20], which, to the authors knowledge, has so far only been done on lattice systems [21–23, 32].

In addition, we were able to successfully perform a cross-check between data from grand canonical MUGC and the canonical methods METROPOLIS and MUCA. This check is a detail that is only mentioned as a side note in the literature, if it is mentioned at all. Contrariwise, we found it to be insightful and useful; it helped to fix a subtle but quite influential issue of the grand canonical implementation, which we could not have uncovered otherwise.

Concerning the scheme of fixed density, we used the native MUCA algorithm to gather data of the Lennard-Jones system, with particle numbers even slightly exceeding two thousand. The subsequent finite-size scaling treatment enabled us to partially confirm the expectations [30, 31]; although we managed to simulate system sizes that clearly extend beyond the intermediate regime, the leading order behaviour was not exhibited yet. Nonetheless, as we adhered to parameters identical to the reference, the presented results aim to conclude the discussion of the condensation-evaporation transition at fixed density, at least for the set of the four systems: lattice gas and Lennard-Jones gas in two and three dimensions, respectively. Our final page of figures was intended to capture this motivation. While the matrix-style representations do not show anything new or revolutionary, they allow to identify scaling regimes and are a nice *gimmick* for visualisation.

All in all, we think that we were able to present a well rounded examination of the transition by splitting our focus in two major directions. Clearly, as each scheme deserves individual attention, we could not treat both with the initially desired amount of detail. In spite of that, we ended up with a resourceful framework and a good setup, which can serve as a starting point for future investigations.

5.1 Outlook

As a first outlook, we want to mention the finite-size analysis of the transition density at fixed temperature, which is currently work in progress. The difficulty lies in the scaling to large systems: linear system sizes greater than a thousand is what we are aiming for. In this region, thermalisation becomes a major problem, which we circumvent, at least for the METROPOLIS simulations, by preparing a state that is far enough on either side of the transition. While trivial for the gas phase, we have to *grow* the droplet by *inheriting* suitable configurations from smaller systems.

When we started working on this thesis, it seemed artificially limiting to deal with a system in two rather than in three dimensions; especially because the computational effort is virtually identical. Then again, the transition to the more complicated case is straight forward from the perspective of programming. With minor modifications, the whole code is applicable in 3D. Thus, looking at mere two dimensions was a justified choice to begin with, even more so, because the system is much easier to visualise (think of snapshots and states). Thereby, direct control is possible: the influence of changing parameters, for instance, can be observed right away, without looking for cut-planes.

Following this train of thought, we could identify mechanisms and processes at work, possibly through the microcanonical analysis [58] or by having a closer look at the free-energy barrier [23, 33, 59]. This would then allow to actually pinpoint where the nucleations takes place, not just for one particularly chosen density or temperature but as a point in phase space containing, amongst others, both parameters. Closely related to free-energy barriers, we came across two more noteworthy techniques involved with the transition.

Firstly, the string method [60] focuses on rare events and finding the most probable path to change from one metastable region to another. This is especially fascinating as we implicitly encountered this problem; MUGC only deals with the barrier in one direction of phase space (density) while struggling with another (energy).

Secondly, we may have a closer look at the formalism introduced by Binder after all; as we have teased before, the condensation-evaporation transition can be identified through the discontinuity of the chemical potential as a function of density. This seems similar to the equal-height criterion we employed during reweighting, but it is actually quite different, because we were considering the two pure phases. To list only a few of the many references with similar scope, we hope to find a conclusion on how exactly both formalisms interlink through a comparison with e.g. [16–18].

Apart from pointing out a personal interest in porting the code to run on GPUs and heterogeneous computing environments through OPENCL, we also want to give some motivation for a consideration of Molecular Dynamic simulations. Admittedly, both, MC as well as MD simulations struggle with critical slowing down and evaluating actual time scales is difficult. However, in recent years new methods have surfaced that may allow to investigate kinetic effects and to gain an intuition about the involved time scales, after all [61, 62]. Along with simulations to determine hysteresis, this may help relating to experiments. Since the 2D simulations represent a 3D system with confinement, nanotubes and surface deposition come to mind. Speaking of the comparison to experiments, a particularly straight forward thing to do would be to cross-check the gas density we measure with literature values of Argon. Having said so, even this seemingly simple comparison requires simulations in 3D, which seems like a nice aspiration to conclude the thesis.

Bibliography

[1] L. P. Kadanoff, W. Götze, D. Hamblen, R. Hecht, E. A. S. Lewis, V. V. Palciauskas, M. Rayl, J. Swift, D. Aspnes, and J. Kane, Rev. Mod. Phys. **39**, 395 (1967).

[2] J. Feder, K. C. Russell, J. Lothe, and G. M. Pound, Adv. Phys. **15**, 111 (1966).

[3] D. Kashchiev, *Nucleation: Basic Theory with Applications* (Butterworth-Heinemann, Oxford, 2000).

[4] L. D. Landau and E. M. Lifshitz, *Statistical Physics* (Butterworth-Heinemann, Oxford, 1980).

[5] A. Kraus, H. Bremers, U. Rossow, and A. Hangleiter, J. Electron. Mater. **42**, 849 (2013).

[6] A. Kraus, C. Hein, H. Bremers, U. Rossow, and A. Hangleiter, J. Appl. Phys. **119**, 235308 (2016).

[7] E. A. Ustinov and D. D. Do, Adsorption **20**, 439 (2014).

[8] M. Yarmolinsky and A. Kuklov, arXiv:`1703.00040` (2017).

[9] A. Seifert and K. D. Beheng, Meteorol. Atmos. Phys. **92**, 45 (2006).

[10] K. Binder and D. P. Landau, Phys. Rev. B **30**, 1477 (1984).

[11] J. L. Cardy, ed., *Finite-Size Scaling* (Elsevier, Amsterdam, 1988).

[12] W. Janke, "First-Order Phase Transitions", in *Computer Simulations of Surfaces and Interfaces* (Springer Netherlands, Dordrecht, 2003), pp. 111–135.

[13] K. Binder, Physica A **319**, 99 (2003).

[14] T. Neuhaus and J. S. Hager, J. Stat. Phys. **113**, 47 (2003).

[15] L. G. MacDowell, P. Virnau, M. Müller, and K. Binder, J. Chem. Phys. **120**, 5293 (2004).

[16] M. Schrader, P. Virnau, and K. Binder, Phys. Rev. E **79**, 061104 (2009).

[17] F. Schmitz, P. Virnau, and K. Binder, Phys. Rev. E **87**, 053302 (2013).

[18] L. G. MacDowell, V. K. Shen, and J. R. Errington, J. Chem. Phys. **125**, 034705 (2006).

[19] M. Biskup, L. Chayes, and R. Kotecký, Europhys. Lett. **60**, 21 (2002).

[20] M. Biskup, L. Chayes, and R. Kotecký, J. Stat. Phys. **116**, 175 (2004).

[21] A. Nußbaumer, E. Bittner, T. Neuhaus, and W. Janke, Europhys. Lett. **75**, 716 (2006).

[22] A. Nußbaumer, E. Bittner, and W. Janke, Phys. Rev. E **77**, 041109 (2008).

[23] A. Nußbaumer, E. Bittner, and W. Janke, Prog. Theor. Phys. **184**, 400 (2010).

[24] T. D. Lee and C. N. Yang, Phys. Rev. **87**, 410 (1952).

[25] N. B. Wilding and A. D. Bruce, J. Phys. Condens. Matter **4**, 3087 (1992).

[26] N. B. Wilding, Phys. Rev. E **52**, 602 (1995).

[27] N. B. Wilding and M. Müller, J. Chem. Phys. **102**, 2562 (1995).

[28] N. B. Wilding, Am. J. Phys. **69**, 1147 (2001).

[29] N. B. Wilding, J. Phys. Condens. Matter **28**, 414016 (2016).

[30] J. Zierenberg and W. Janke, Phys. Rev. E **92**, 012134 (2015).

[31] J. Zierenberg and W. Janke, J. Phys. Conf. Ser. **750**, 012017 (2016).

[32] A. Nußbaumer, J. Zierenberg, E. Bittner, and W. Janke, J. Phys. Conf. Ser. **759**, 012009 (2016).

[33] J. Zierenberg, P. Schierz, and W. Janke, Nat. Commun. **8**, 14546 (2017).

[34] J. Zierenberg, "From Particle Condensation to Polymer Aggregation: Phase Transitions and Structural Phases in Mesoscopic Systems", PhD thesis (Leipzig University, 2016).

[35] W. Nolting, *Grundkurs Theoretische Physik 6* (Springer, Berlin, Heidelberg, 2014).

[36] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, New York, 1989).

[37] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).

[38] M. Rovere, D. W. Heermann, and K. Binder, J. Phys. Condens. Matter **2**, 7009 (1990).

[39] M. Rovere, P. Nielaba, and K. Binder, Z. Phys. B: Condens. Matter **90**, 215 (1993).

[40] B. A. Berg and T. Neuhaus, Phys. Lett. B **267**, 249 (1991).

[41] B. A. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 9 (1992).

[42] W. Janke, Int. J. Mod. Phys. C **3**, 1137 (1992).

[43] W. Janke, Physica A **254**, 164 (1998).

[44] B. A. Berg, Comput. Phys. Commun. **153**, 397 (2003).

[45] J. Viana Lopes, M. D. Costa, J. M. B. Lopes dos Santos, and R. Toral, Phys. Rev. E **74**, 046702 (2006).

[46] J. Zierenberg, M. Marenz, and W. Janke, Comput. Phys. Commun. **184**, 1155 (2013).

[47] B. Efron, *The Jackknife, the Bootstrap and Other Resampling Plans* (Society for Industrial and Applied Mathematics, Philadelphia, 1982).

[48] K. Binder, Phys. Rev. A **25**, 1699 (1982).

[49] A. Nußbaumer, E. Bittner, and W. Janke, Europhys. Lett. **78**, 16004 (2007).

[50] E. Bittner, A. Nußbaumer, and W. Janke, Nucl. Phys. B **820**, 694 (2009).

[51] K. Binder and M. H. Kalos, J. Stat. Phys. **22**, 363 (1980).

[52] D. J. Adams, Mol. Phys. **29**, 307 (1975).

[53] A. Lotfi, J. Vrabec, and J. Fischer, Mol. Phys. **76**, 1319 (1992).

[54] J. J. Potoff and A. Z. Panagiotopoulos, J. Chem. Phys. **112**, 6411 (2000).

[55] C. Borgs and W. Janke, Phys. Rev. Lett. **68**, 1738 (1992).

[56] Y. Imry, Phys. Rev. B **21**, 2042 (1980).

[57] J. Zierenberg, M. Mueller, P. Schierz, M. Marenz, and W. Janke, J. Chem. Phys. **141**, 114908 (2014).

[58] P. Schierz, J. Zierenberg, and W. Janke, J. Chem. Phys. **143**, 134114 (2015).

[59] P. Schierz, J. Zierenberg, and W. Janke, Phys. Rev. E **94**, 021301 (2016).

[60] W. E, W. Ren, and E. Vanden-Eijnden, Phys. Rev. B **66**, 052301 (2002).

[61] A. Laio and M. Parrinello, Proc. Nat. Acad. Sci. U.S.A. **99**, 12562 (2002).

[62] A. Barducci, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **100**, 020603 (2008).

Acknowledgements

I thank Prof. Dr. Wolfhard Janke for giving me the opportunity to work on this topic, for allowing me to profit from his expertise in uncountable ways, for providing me with required resources and most of all, for the warm welcome into the CQT group. I further want to show my sincere gratitude to my supervisor Johannes Zierenberg. His inexhaustible support, resourcefulness and infectious motivation made this thesis become an exciting and delightful experience of mine. In addition, I would like to thank all the members of the CQT group for giving me honest feedback after talks and for sharing an incredible collective and individual knowledge with me. In particular, I thank Jonathan Gross and Stefan Schnabel for taking the time to help me tackle the insertion issue, as well as Marius Bause and Philipp Schierz for useful discussions. To conclude, I want to express my candid appreciation to my loving family and dear friends, especially Stefan Horn, who contributed greatly to my ability to code and Diana Rodriguez Bonet for her moral support at any sacrifice.