



universität
wien

MASTERARBEIT / MASTER'S THESIS

Titel der Masterarbeit / Title of the Master's Thesis

Tensor networks for hard-core lattice gases

verfasst von / submitted by

Bernhard Georg Mitterwallner BSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien / Vienna, 2016

Studienkennzahl lt. Studienblatt /
degree programme code as it appears on
the student record sheet:

A 066 876

Studienrichtung lt. Studienblatt /
degree programme as it appears on
the student record sheet:

Masterstudium Physik UG2002

Betreut von / Supervisor:

Univ.-Prof. Dr. Frank Verstraete

Acknowledgements

I wish to express my gratitude to all the people who have supported and encouraged me over the years.

Especially, I would like to thank my sister Birgit, my friends and flat mates Philipp and Christoph, as well as my Aikido teacher and friend Manfred, for providing me with invaluable advice and encouragement during my final year. Over the course of my studies, Aikido has become a constant in my life, and I thank all members of the Aikido Union Wien for giving me the chance to train and hone my skills. And of course, I want to thank my parents, Georg and Jutta, for their support.

I gratefully acknowledge, that Valentin Stauber certainly answered most of my stupid questions concerning tensor networks. I would like to thank Frank Verstraete for his support and for providing me with a very interesting and challenging problem. Furthermore, I want to thank Benoît Descamps, Damian Draxler, Matthew Fishman, Martin Veznik and Burak Şahinöglü, for stimulating discussions. I am also grateful to Andrej Gendiar for his encouragement and very helpful and interesting discussions.

Contents

Summary	7
Zusammenfassung	7
1 Motivation	8
2 Lattice gases	9
2.1 Statistical mechanics on the lattice	9
2.2 Lattice gases and spin models	13
2.3 Athermal systems: Hard-core lattice gases	15
2.4 An “ideal” lattice gas	19
2.5 The k NN model	20
3 Tensor networks and classical systems	27
3.1 k NN in $d = 1$ and the transfer matrix	28
3.2 k NN in $d = 2$	34
3.3 Local observables and environments	37
4 Tensor networks and quantum systems	42
4.1 Entanglement	42
4.2 Matrix product states	45
4.3 Matrix product operators	51
4.4 The relationship between classical and quantum systems . . .	52
5 Algorithms for contracting tensor networks	57
5.1 Computational cost	57
5.2 Finding dominant eigenvalues	58
5.3 Boundary MPS methods: iTEBD	60
5.4 Spontaneous breaking of sublattice symmetry	64
5.5 Real-space renormalization: TRG	67
6 Numerical results	75
6.1 1NN	76
6.2 2NN	80
6.3 3NN	84
7 Conclusion and outlook	88
A Some linear algebra	90
A.1 Singular value decomposition	90
A.2 Non-negative and convergent matrices	90

Summary

This work is concerned with the study of classical lattice gases whose constituent particles interact solely via hard-core repulsion. We investigate the k nearest neighbour exclusion model (k NN model) on the square lattice, which is the discrete version of a system of hard disks in the plane. After a review of the statistical mechanics of hard core lattice gases as well as previous results, tensor networks are introduced in the context of classical partition functions. A tensor network formulation of the k NN model is presented. Applications of tensor networks to quantum many-body problems are reviewed, with a special emphasis on matrix product states and matrix product operators and their relevance to classical problems. Two algorithms for the contraction of two dimensional tensor networks are discussed: *infinite time evolved block decimation* (iTEBD) and *tensor renormalization group* (TRG). Numerical results for the cases $k = 1, 2, 3$ are presented and found to be in agreement with previous studies. Finally, challenges in tensor network treatment of hard-core lattice gases are outlined and possible directions for further studies are proposed.

Zusammenfassung

Diese Arbeit behandelt klassische Gittergase, welche aus ausgedehnten, nicht überlappenden Teilchen bestehen und ansonsten nicht wechselwirken. Wir untersuchen das k NN Modell (engl. *k next nearest neighbours model*) auf dem quadratischen Gitter. Dieses Modell ist die diskrete Variante eines Systems harter Scheiben in der Ebene. Nach einer Zusammenfassung der klassischen statistischen Mechanik von Gittergasen und einem Überblick über bisherige Resultate, werden Tensornetzwerke im Kontext klassischer Zustandssummen eingeführt. Eine Tensornetzwerk-Formulierung des k NN Modells wird präsentiert. Anwendungen von Tensornetzwerken auf Quanten-Vielteilchenprobleme werden erläutert, mit einem besonderen Fokus auf Matrixproduktzustände und Matrixproduktoperatoren, sowie deren Relevanz für klassische Probleme. Zwei Algorithmen für die Kontraktion zweidimensionaler Tensornezwerke werden diskutiert: *infinite time evolved block decimation* (iTEBD) und *tensor renormalization group* (TRG). Numerische Ergebnisse für die Fälle $k = 1, 2, 3$ werden präsentiert und für übereinstimmend mit bisherigen Resultaten befunden. Abschließend werden Herausforderung des Zugangs über Tensornetzwerke besprochen und mögliche Richtungen für weiterführende Arbeiten vorgeschlagen.

1 Motivation

Hard-core lattice gases have received attention over many decades from theoretical physics and chemistry alike. Unfortunately, most of them are extremely difficult to solve exactly in more than one dimension, the hard hexagon model on the triangular lattice being a notable exception: It was finally solved exactly by Baxter in 1980 [1]. For other cases, approximation methods are the only tools available. Series expansions, as well as Monte Carlo[23] and transfer matrix algorithms are probably the most prominent examples [4, 2, 15, 7, 11, 30, 24].

In recent years, tensor networks have emerged as a novel concept in the study of quantum many body systems [27]. Aside from encoding ground states of quantum systems, tensor networks turn out to offer a straightforward way of expressing the partition function of many classical lattice systems. Once formulated in tensor network language, algorithms may be employed which have a crucial advantage over Monte Carlo and conventional matrix methods: They allow for studying systems which are either in the thermodynamic limit to begin with, as with boundary MPS algorithms exploiting translational invariance [26], or at least exponentially large, as with tensor network based real-space renormalization schemes [21]. Just like transfer matrix studies of finite sized systems, tensor network approaches give direct access to the system's partition function - a quantity which is cumbersome to obtain using Monte Carlo. This in turn allows to calculate the system's entropy.

2 Lattice gases

A classical lattice gas is a model of interacting particles whose positions are constrained to coincide with the vertices of a graph. They arise naturally out of an attempt to discretise a more realistic model i.e one in which particles can occupy continuous positions. On the other hand, for some systems, lattice gases are realistic models by themselves. For instance, the absorption of atoms on the surface of a crystal may be modelled by a lattice gas, e.g [44].

Aside from (often merely crudely) approximating the microscopic physics of real world systems, lattice gases are interesting in their own right for they possibly allow us to study thermodynamic phenomena such as phase transitions [3, 13].

In any case, the study of lattice models is often motivated by the hope that discrete models are somehow easier to deal with. Even if they elude exact solution in most cases, perturbative and numerical methods may give access to thermodynamic properties.

2.1 Statistical mechanics on the lattice

Before going into the details of lattice gases, the statistical mechanics of lattice models in general shall be reviewed.

The underlying geometric structure - the lattice - of any lattice model is a connected graph consisting of $N_0 \in \mathbb{N}$ vertices which are called lattice sites. We assign an index $i = 1, \dots, N_0$ to each such lattice site. Let $\nu_i(k)$ denote the number of sites that can be reached by a path originating at site i and whose length is not exceeding k . The *coordination number* of lattice site i is defined as $\nu_i(1)$. Note, that for all lattices depicted in fig. 1, the coordination number is the same for all lattice sites. In accordance with [3, chap. 1], a lattice is said to be of dimension d if and only if

$$\nu_i(k) = \mathcal{O}(k^d) \quad (k \rightarrow \infty), \quad \forall i \quad .$$

This definition implicitly assumes an infinite lattice size, $N_0 = \infty$. However, the lattices considered will always exhibit some periodic structure. We can therefore also talk about the dimension of a finite sized periodic lattice if we define its dimension to be equal to the dimension of its infinite sized counterpart.

The system may have N degrees of freedom (DOF) which sit on the vertices or edges of the lattice, for spin and vertex-models respectively. Let $\mathbf{q} = \{q_1, \dots, q_N\}$ denote a configuration of the system. Each degree of freedom q_i may take on a finite set of discrete values. The set of all possible configurations $\{\mathbf{q}\}$ forms the configuration space.

The state of the system is described by a probability distribution p on the




Figure 1: Various periodic lattices in two dimensions: Square (a), triangular (b) and honeycomb (c) lattice.

configuration space. Since we are only interested in equilibrium states, p maximises the entropy under a given set of constraints specific to the ensemble in which the system is regarded. The probability for a configuration $\mathbf{q}$ is written as

$$p(\mathbf{q}) = \frac{w(\mathbf{q})}{Z} \quad (1)$$

where

$$Z = \text{Tr}_q w(\mathbf{q}) \quad (2)$$

is the partition function of the system. The symbol Tr_q denotes the sum over all configurations:

$$\text{Tr}_q := \sum_{q_1 \in Q} \sum_{q_2 \in Q} \cdots \sum_{q_N \in Q} = \sum_{\mathbf{q} \in \mathcal{Q}}$$

Expectation values of observables $o(\mathbf{q})$ can then be written as:

$$\langle o \rangle = \frac{1}{Z} \text{Tr}_q o(\mathbf{q}) w(\mathbf{q}) \quad (3)$$

Of particular importance are local observables, i.e observables for finite regions of the lattice such as single-site observables. An example is the expectation value $\langle c_i \rangle$ of a configuration c_i of site i , i.e. its spin value or its vertex configuration, depending on the kind of model we are considering. We will also be interested in correlations between local observables. The correlation function of the system is defined as:

$$G_{ij} = \langle c_i c_j \rangle - \langle c_i \rangle \langle c_j \rangle$$

We restrict ourselves to the case where the system is homogeneous and denote the relative distance between sites i and j by r . Along a fixed direction, the asymptotic behaviour may then be described by the ansatz[3]:

$$G(r) \rightarrow a(r) e^{-r/\xi} \quad (r \rightarrow \infty)$$

Here, $a(r)$ is some algebraic function and ξ is called the *correlation length* of the system. Purely algebraically decaying correlation thus correspond to an infinite correlation length.

Let $H(\mathbf{q})$ be the Hamiltonian of the system and $\beta = 1/k_B T$ the inverse temperature, where k_B denotes Boltzmann's constant. The canonical partition function is given by the sum of all Boltzmann weights:

$$Z^c(\beta) = \text{Tr}_{\mathbf{q}} e^{-\beta H(\mathbf{q})} \quad ,$$

If the notion of a particle number is present, the grand canonical ensemble may be more convenient to work with. Let $n(\mathbf{q})$ denote the number of particles associated with configuration $\mathbf{q}$. For a given chemical potential μ , the grand canonical partition function is defined as:

$$Z^{\text{gc}}(\beta, \mu) := \text{Tr}_{\mathbf{q}} e^{-\beta H(\mathbf{q}) + \beta \mu n(\mathbf{q})} = \text{Tr}_{\mathbf{q}} z^{n(\mathbf{q})} e^{-\beta H(\mathbf{q})} \quad (4)$$

The quantity

$$z = e^{\beta \mu}$$

is called fugacity or activity in the physical and chemical literature respectively.

A necessary condition for the occurrence of a phase transition in a classical system is that it must be in the thermodynamic limit [13]. In the case of lattice models, this means that the lattice size goes to infinity i.e $N_0 \rightarrow \infty$ while number of degrees of freedom per lattice site N/N_0 remains constant. Since according to (2) Z generally diverges in the thermodynamic limit, it is custom to define a partition function per site κ , such as to still be able to talk about partition functions in a meaningful way.

$$\kappa := \lim_{N_0 \rightarrow \infty} Z^{1/N_0}$$

For a given partition function Z the (grand) free energy is given by

$$F = -\ln Z \quad (5)$$

with Z being the (grand) canonical partition function¹. From the definition of κ , the free energy per site is given by:

$$f = -\ln \kappa \quad (6)$$

The regions of parameter space in which f is analytic define the phase diagram of the system. If the first derivative of the free energy with respect

¹Note that this definition of the free energy differs from the usual presentation $\beta F = -\ln Z$, in that the inverse temperature β has been absorbed into F .

to some external parameter is discontinuous, we speak of a first order phase transition. The appearance of discontinuity in any higher order derivative marks a continuous or second order phase transition. First order phase transitions are accompanied by jump discontinuities in several thermodynamic quantities (e.g. the density), indicating phase coexistence. In the case of second order transitions however, many quantities diverge with a power law and the system exhibits scale invariance. This is also reflected in diverging correlation lengths which implies that correlations decay algebraically at the point of a continuous phase transition, which is called the *critical point*. This power-law behaviour near the critical point will be further discussed in section 2.3 with special emphasis of the class of systems relevant to this work.

From a microscopic point of view, a phase transition in a classical system is always accompanied by a breaking of ergodicity[13]: It turns out, that as the system undergoes a phase transition, configuration space breaks into disjoint regions in one of which the system becomes “trapped”. This means, that there is no longer a unique equilibrium state describing the system! In many cases, phase transitions are accompanied by the breaking of a symmetry²[13]: In the broken phase, the disjoint regions of configuration space are related by the action of a particular group G : Given a configuration $\mathbf{q}_A$ belonging to region A , there is a $g \in G$ such that $l_g(\mathbf{q}_A) = \mathbf{q}_B$, where $\mathbf{q}_B$ now belongs to region B and l is the group action. Imagine some local observable o whose expectation value is zero in the disordered phase and becomes non-zero as the system undergoes a phase transition. Furthermore, $\langle o \rangle$ is unique to the particular equilibrium state the system gets trapped in. Such an observable is called an *order parameter*. The order parameter therefore belongs to a faithful representation of the relevant symmetry group G . Sometimes however, it will only be of interest whether the system ordered at all - i.e. we are interested in a quantity that is non-zero in the ordered phase and zero in the disordered one.

It is very important to realise, that the actual state of the system is no longer described by (1) which treats all possible equilibrium states on equal footing. Therefore care must be taken when computing local observables. In particular, expectation values of order parameters would always be zero if calculated according to (3), which is deeply unphysical! It is therefore of great importance to be aware of the consequences of ergodicity breaking.

²There are also phase transitions in which no local symmetries are broken. These “topological” phase transitions are however still accompanied by a breaking of ergodicity. Topological phase transitions are beyond the scope of this work.

2.2 Lattice gases and spin models

As mentioned earlier, a lattice gas is a model of particles which are constrained to be located on the sites of a lattice. Consider a system of N particles on a lattice with N_0 sites, where

$$\mathbf{x}_N = \{x_1, \dots, x_N\}$$

denotes a configuration of the system. Each particle, which is assigned an index $i = 1, \dots, N$, has one discrete positional degree of freedom x_i , which is simply the index of the lattice site it is located on.

If the interaction between particles is mediated by the pair potential v and there are no external fields, the Hamiltonian reads:

$$H(\mathbf{x}_N) = \frac{1}{2} \sum_{i \neq j} v(x_i, x_j) \quad (7)$$

Here we encounter the following issue: We want to impose the constraint that a lattice site may be occupied by at most one particle, i.e. $x_i \neq x_j$, $\forall i \neq j$. A pair potential $V(\vec{r}_1, \vec{r}_2)$, $\vec{r}_1, \vec{r}_2 \in \mathbb{R}^3$ for a system in the continuum usually has some repulsive core, i.e. a positive divergence,

$$V(\vec{r}_1, \vec{r}_2) \rightarrow +\infty, \quad (|\vec{r}_1 - \vec{r}_2| \rightarrow 0)$$

that prevents particles with finite energies from getting arbitrarily close to each other³. One could, of course, set

$$v(x, x) = +\infty$$

For the Boltzman weight $e^{-\beta H(\mathbf{x}_N)}$ of a given configuration $\mathbf{x}_N$, we'd therefore have:

$$e^{-\beta H(\mathbf{x}_N)} = 0 \quad \Longleftrightarrow \quad \mathbf{x}_N \text{ contains at least two elements equal to each other}$$

Spin models can be mapped to lattice gases and vice versa. Inspired by the treatment for the Ising model in [20] and [13, chap. 2], this relationship shall now be discussed for the more general case.

In a classical spin model, a discrete degree of freedom σ_i (a “spin”) is assigned to each lattice site i . Therefore the number of degrees of freedom N is equal to the number of lattice sites:

$$N = N_0$$

³There are soft condensed matter systems for which the effective pair potential between constituent particles is actually bounded[22]. The lattice version of such a system would require a formulation different from the one presented here.

Let σ denote a spin-configuration, where the subscript N has been conveniently dropped. The Hamiltonian $H(\sigma)$ can generally be expanded in sums of n -site contributions $v_{i_1 \dots i_n}(\sigma_{i_1}, \dots, \sigma_{i_n})$:

$$H(\sigma) = \sum_i v_i(\sigma_i) + \sum_{i \neq j} v_{ij}(\sigma_i, \sigma_j) + \sum_{i \neq j \neq k} v_{ijk}(\sigma_i, \sigma_j, \sigma_k) + \dots$$

Defining the “free Hamiltonian”

$$H_0(\sigma) = \sum_i v_1(\sigma_i)$$

and grouping the remaining terms into the interaction part H_{int} , the Hamiltonian reads:

$$H(\sigma) = H_0(\sigma) + H_{\text{int}}(\sigma) \quad (8)$$

Let us now consider spins which can take on values $\sigma_i \in \{0, 1\}$ and a homogeneous external field $H_0(\sigma) = h \sum_i \sigma_i$. The canonical partition function for the spin system then reads:

$$Z_{\text{spin}}^c = \sum_{\sigma_1=0,1} \dots \sum_{\sigma_n=0,1} e^{-\beta h \sum_i \sigma_i - \beta H_{\text{int}}(\sigma)} = \text{Tr}_{\sigma} e^{-\beta h n(\sigma)} e^{-\beta H_{\text{int}}(\sigma)}$$

Here $n(\sigma) = \sum_i \sigma_i$ is the number of spins equal to 1, and the trace Tr_{σ} is taken over all spin degrees of freedom. If we interpret the possible spin values $\{0, 1\}$ as occupation numbers (i.e site i is occupied $\iff \sigma_i = 1$), and the external field h as a chemical potential $\mu = -h$, we can regard the canonical partition function Z_{spin}^c of the spin system as the *grand* canonical partition function of a lattice gas:

$$Z_{\text{spin}}^c(\beta, h) = Z_{\text{lg}}^{\text{gc}}(\beta, -h) \quad (9)$$

Note, that the constraint that a site may be occupied by at most one particle is automatically implied by the range of possible spin values $\{0, 1\}$. The trace is taken over *all* degrees of freedom irrespective of the number of particles, which is now controlled by an external parameter μ .

To summarise the correspondence:

spin model	$\leftrightarrow$	lattice gas
canonical	$\leftrightarrow$	grand canonical
external field	$\leftrightarrow$	chemical potential
spin values	$\leftrightarrow$	occupation numbers

Also note, that the number of degrees of freedom N is no longer equal to the particle number n . Instead, to each spin-configuration σ , a particle number $n(\sigma)$ is assigned. This allows to keep the number of DOF fixed while allowing n to vary, thus simplifying the calculation of Z^{gc} .




Figure 2: Hard core particles (here on the square lattice) may not overlap but do not interact otherwise.

2.3 Athermal systems: Hard-core lattice gases

Consider a lattice gas of identical finite sized particles, all of which now have a shape: Each occupied site is surrounded by a region of neighbouring sites forming the shape of the particle. Shapes of different particles may not overlap (hard-core repulsion). This is best illustrated by a picture, see fig. 2. Aside from this non-overlapping condition, particles don't interact with each other. Such a system will be called “hard core lattice gas” (HCLG). The entire configuration space of the system can be regarded as being split into a set X of valid configurations and its complement $\bar{X}$ of invalid ones, where a “valid configuration” is one with no overlapping particles, while for an “invalid” one at least 2 particles overlap. For a system composed of N_0 lattice sites this constraint is expressed by

$$H(\mathbf{q}) = \begin{cases} 0 & \mathbf{q} \in X \\ \infty & \text{else} \end{cases}.$$

For finite energies, the system cannot leave X , hence the internal energy $U = \langle H \rangle$ must vanish and the free energy is thus entirely entropic:

$$F = -S$$

This is confirmed by looking at the statistical mechanics. The canonical partition function reads:

$$Z^c(\beta) = \text{Tr}_{\mathbf{q}} e^{-\beta H(\mathbf{q})}$$

The Boltzmann-weight $e^{-\beta H(\mathbf{q})}$ reduces to a combinatorial factor Γ , which merely determines whether the configuration is possible or not.

$$e^{-\beta H(\mathbf{q})} = \begin{cases} 1 & \mathbf{q} \in X \\ 0 & \text{else} \end{cases} =: \Gamma(\boldsymbol{\sigma})$$

This expression is independent of $\beta \in (0, +\infty)$ and thus Z^c itself is temperature-independent, for the weight Γ solely depends on the configuration σ . Such a system is called *athermal* [11, 19] and it follows, that its canonical partition function Z^c is just the number of allowed configurations and therefore equal to the microcanonical partition function Ω .

$$Z = |X| =: \Omega \quad \Rightarrow \quad F = -S$$

Let us now restrict our attention to hard-core lattice gases again. For the grand partition function we have:

$$Z(\beta, \mu) = \text{Tr}_{\sigma} \Gamma(\sigma) e^{\mu n(\sigma)}$$

where β has been conveniently absorbed into μ . Configurations σ are distributed according to

$$p(\sigma) = \frac{1}{Z} \Gamma(\sigma) e^{\mu n(\sigma)},$$

from which the (Shannon) entropy may be calculated⁴:

$$\begin{aligned} S &= - \sum_{\sigma} p(\sigma) \ln p(\sigma) \\ &= - \sum_{\sigma} p(\sigma) \ln \Gamma(\sigma) - \mu \sum_{\sigma} p(\sigma) n(\sigma) + \ln Z \sum_{\sigma} p(\sigma) \\ &= - \frac{1}{Z} \sum_{\sigma} e^{\mu n(\sigma)} \Gamma(\sigma) \ln \Gamma(\sigma) - \mu \langle n \rangle + \ln Z \\ &= \ln Z - \mu \langle n \rangle \end{aligned}$$

Here, $\langle n \rangle$ and ρ denote the expectation value of the particle number and the particle density, respectively:

$$\langle n \rangle = N_0 \rho = \sum_{\sigma} p(\sigma) n(\sigma) = \sum_{\sigma} p(\sigma) \sum_i \sigma_i = \sum_i \langle \sigma_i \rangle \quad (10)$$

Also note the relationship between the density and the first derivative of the grand free energy:

$$\rho = - \frac{1}{N_0} \frac{\partial F}{\partial \mu} = \frac{1}{N_0} \frac{1}{Z} \frac{\partial Z}{\partial \mu} \rightarrow \frac{1}{\kappa} \frac{\partial \kappa}{\partial \mu} \quad (N_0 \rightarrow \infty)$$

⁴ Aside from the normalisation condition $\text{Tr}_{\sigma} p(\sigma) = 1$, it is used that

$$\Gamma(\sigma) \ln \Gamma(\sigma) = 0, \quad \forall \sigma$$

This is justified if the ill-defined expression $0 \ln 0$ is interpreted as a limit:

$$\lim_{x \rightarrow 0} x \ln x = 0$$

As with the free energy, the entropy may also be defined per site:

$$s := \lim_{N_0 \rightarrow \infty} \frac{S}{N_0} = \lim_{N_0 \rightarrow \infty} \frac{1}{N_0} (\ln \kappa^{N_0} - \mu \langle n \rangle) = \ln \kappa - \mu \rho$$

Unless stated otherwise, systems are always regarded in the grand canonical picture from now on. Grand partition function and free energy per site will be denoted by the symbols κ and f respectively and simply be referred to as *the* partition function/free energy.

On physical grounds, we may expect the entropy to take on a global maximum for zero chemical potential, since in this case all allowed configurations have the same weight. We can prove this by making use of the convexity of s , and the fact that ρ is a strictly monotonously increasing function of the chemical potential. Setting $\kappa' := \partial \kappa / \partial \mu$ and using $\rho = \kappa' / \kappa$, we have:

$$\frac{\partial s}{\partial \mu} = -\mu \frac{\partial \rho}{\partial \mu}$$

The only way for the r.h.s. to vanish is that $\mu = 0$, since the derivative of ρ is strictly positive. To summarise, for athermal systems we have:

$$\begin{aligned} f(\mu) &= -\ln \kappa(\mu) = -s(\mu) - \mu \rho(\mu) \\ s_{\max} &= s(0) \end{aligned}$$

In a thermal system, an equilibrium state for $\beta \rightarrow \infty$ - i.e. the zero temperature limit - is usually called a “ground state” of the system. In analogy, a ground state of a HCLG shall be defined as an equilibrium state in the limit of infinite chemical potential, $\mu \rightarrow \infty$. Since ρ is a strictly monotonously increasing function of μ , the limit $\mu \rightarrow \infty$ corresponds to states of maximum density $\rho_{\max}$, or equivalently maximum particle number $n_{\max}$. These states live entirely on a (in most cases) tiny subset of the overall configuration space, since for $\mu \rightarrow \infty$ the probability of a configuration σ must necessarily vanish if $n(\sigma) < n_{\max}$. This subset is the set of what will be called *fully packed configurations* (FPC)- which as their name suggests, realise the highest possible particle number, i.e the densest possible packing of the extended particles. A system which has $g \geq 1$ fully packed configurations will be said to have a g -fold ground state degeneracy. Configurations belonging to different ground states will usually be related by some symmetry transformations. Therefore, if there are multiple ground states, the overall symmetry of the system will be broken for $\mu \rightarrow \infty$. Under the assumption that phase transitions - i.e symmetry breaking - only take place for $\mu < \infty$, the ground state is in the same phase as the high-density states. Furthermore, we shall assume that ground states, and thus by assumption the high-density phase, exhibit the least symmetry. Thus, any possible intermediate phases between low and high densities only partially break the relevant symmetry. Putting all these considerations together, we can guess

the overall symmetry group of the system from its ground states. Unfortunately, the ground states are generally unknown. However, since they must be an average over fully packed configurations, we can at least construct all possible candidates for ground states: Any subset of all FPCs $\{\sigma_1^{\text{fp}}, \dots, \sigma_g^{\text{fp}}\}$ generates a possible ground state. Thus generally only a subset - i.e. a sublattice - of all lattice sites will be occupied in the ground state, while all the other sites remain unoccupied. For the corresponding high-density states with $\mu < \infty$, this means that particular sublattice will become *preferentially* occupied. The density will therefore be site-dependent and sublattice symmetry will be broken. If the high density phase exhibits a density pattern corresponding to a particular FPC, i.e. the corresponding ground state consists only of a single configuration, we may speak of a *solid* phase. Order parameters may be constructed from (the differences of) these sublattice-densities [11, 30, 24]. The symmetry can also be explicitly broken by making the chemical potential site dependent according to a pattern corresponding to a particular ordered state in which some sites B are preferentially occupied, while the remaining sites A are preferentially empty:

$$\mu_A = \mu \quad \mu_B = \mu + h$$

The parameter h is then nothing but an external field which breaks the symmetry.

We shall now discuss some more material properties of athermal systems and their relationship to (local) observables.

As already mentioned, the density ρ is related to the expectation value of the particle number, see eq. (10). Since we will be regarding homogeneous systems, we have

$$\rho = \langle \sigma_i \rangle, \quad \forall i$$

as long as translation invariance is not broken. If however, the system undergoes a phase transition which is accompanied by a breaking of sublattice symmetry, ρ is given by the average over all sublattice densities ρ_j :

$$\rho = \frac{1}{N_s} \sum_{j=1}^{N_s} \rho_j, \quad \rho_j = \langle \sigma_i \rangle, \quad \forall i \in L_j$$

Here, N_s is the number of distinct sublattices and L_j , ($j = 1, \dots, N_s$) are the sets of lattice points belonging to the different sublattices.

In analogy to the specific heat of a thermal system, we shall introduce the quantity

$$C := \frac{\partial \rho}{\partial \mu} = \frac{\partial}{\partial \mu} \frac{\kappa'}{\kappa} = \frac{\kappa''}{\kappa} - \rho^2$$

which is related to the *isothermal compressibility* $K_T = C/\rho^2$ [4]. However, C shall simply be referred to as *compressibility*, for convenience. In terms

of expectation values, C is the sum over all density-density correlators

$$C = \sum_{ij} \langle \sigma_i \sigma_j \rangle - \rho^2 = \frac{1}{N_0^2} \left(\langle n^2 \rangle - \langle n \rangle^2 \right)$$

which is just the fluctuation of the density. Since

$$\sum_{ij} \langle \sigma_i \sigma_j \rangle = \left\langle \sum_{ij} \sigma_i \sigma_j \right\rangle \quad ,$$

the compressibility is given by a non-local observable.

At a critical point we expect several quantities to exhibit power law behaviour. (We won't go into the details but instead refer to the vast body of literature on critical phenomena, e.g. [13].) As in [24], we define a reduced chemical potential

$$\tilde{\mu} := \frac{\mu - \mu_C}{\mu_C}$$

which is positive (negative) above (below) the critical point μ_C . In the limit of $\tilde{\mu} \rightarrow 0$, the critical exponents α, γ, ν are defined as:

$$C \propto |\tilde{\mu}|^{-\alpha}$$

$$\xi \propto |\tilde{\mu}|^{-\nu}$$

For $\tilde{\mu} > 0$, there is also a critical exponent β for the order parameter:

$$|q| \propto \tilde{\mu}^\beta$$

As already mentioned, the correlation function at the critical point itself, i.e. $\tilde{\mu} = 0$, is expected to decay purely algebraically. To describe this behaviour, yet another exponent η is introduced. Together with the dimensionality d of the system, the asymptotic behaviour of the correlation function may then be written as:

$$G(r) \propto r^{-d+2-\eta}$$

2.4 An “ideal” lattice gas

The most simple HCLG is one with no interaction between different particles at all, the only feature being that a lattice site can be occupied by at most one particle. In the spin-picture, this the case for all possible configurations σ and therefore

$$\Gamma(\sigma) = 1, \quad \forall \sigma \quad .$$

The partition function then reads:

$$Z(\mu) = \text{Tr}_{\sigma} e^{\mu n(\sigma)} = \text{Tr}_{\sigma} e^{\mu \sum_i \sigma_i} = \text{Tr}_{\sigma} \prod_i e^{\mu \sigma_i} = \prod_i \sum_{\sigma_i=0,1} e^{\mu \sigma_i} = (1 + e^{\mu})^N$$




Figure 3: A system composed of hard disks of radius r (left) is equivalent to a system of point particles which are surrounded by a disk-shaped excluded region of radius $2r$ (right).

For the free energy per site f , one therefore obtains:

$$f = -\ln(1 + e^\mu)$$

This expression is analytic $\forall \mu$ which confirms the absence of phase transitions. For the density and the entropy per site one obtains:

$$\rho(\mu) = -\frac{\partial f}{\partial \mu} = \frac{1}{1 + e^{-\mu}}$$

$$s(\mu) = \ln(1 + e^\mu) - \frac{\mu}{1 + e^{-\mu}}$$

Therefore the maximum entropy of the ideal lattice gas is given by:

$$s_{\max} = \ln 2$$

We expect it to upper-bound the maximum entropies of all interacting lattice gases. Note, that the Hamiltonian (and thus the partition function) is completely independent of the underlying lattice, which implies that the system behaves the same in any dimension. As in the continuous case, we expect the free energy of this ideal lattice gas to lower-bound the free energy of any interacting lattice gas.

2.5 The k NN model

The k NN model is the discretised version of the famous hard spheres model in d spatial dimensions. As the name suggests, the latter is a system of d -dimensional spheres of radius r which solely interact via hard-core repulsion, i.e. two given spheres may not overlap. Equivalently, this model can also be regarded as describing point particles which must always be separated by a distance greater than $2r$ and do not interact in any other way. In the case of two spatial dimensions, the hard spheres model is also called the *hard disks* model. It was the first system to be studied by the Metropolis Monte




Figure 4: Equivalence of hard-core repulsion and excluded volume models: (a) Hard disks in the continuum (b) 3NN model.

Carlo algorithm [23]. However, determining its phase diagram remained a long standing open problem. The most recent studies [5, 9] provide strong evidence for the existence of three distinct phases, namely a solid, a fluid⁵ and an intermediate hexatic phase. Just like the fluid phase, the latter lacks positional order, i.e continuous translational symmetry is preserved, while the the system exhibits an “orientation” (see the references for details). Discretising the hard-disks model, that is putting it on a lattice, leads to what is called the k NN model: It describes a HCLG for which the first k nearest neighbours around a particle are excluded, i.e. cannot be occupied by another particle, see fig. 4. In order to come up with a definition of nearest neighbour which makes sense for discretising the hard-disks problem, the notion of distance between lattice sites needs to be defined with care. Up until now, the definition from graph theory was used, where the distance between two sites is defined as the number of connected edges forming the shortest path between those two sites. But since the lattice is regarded as a discretisation of euclidean space, the distance d_{ij} between two lattice sites i and j may also be defined as the euclidean distance between the corresponding points x_i and x_j :

$$d_{ij} = |x_i - x_j|, \quad x_i, x_j \in \mathbb{R}^d$$

Since there are at most countably many lattice sites, there are also only countably many different d_{ij} :

$$d_{ij} \in \{r_0, r_1, r_2, \dots\}, \quad r_i < r_{i+1}$$

The sites i and j are now called k th nearest neighbours if and only if

$$d_{ij} = r_k \quad .$$

⁵Since hard-core interaction is purely repulsive, there is no distinction between liquid and gaseous phases.




Figure 5: The $k = 1, \dots, 7$ nearest neighbours of the site in the centre (black dot).




Figure 6: The k NN model in $d = 1$ is equivalent to model of hard rods of length $k + 1$. Occupied (empty) sites are represented by filled (empty) circles.

All k th nearest neighbours of a given lattice site i thus lie on the surface of a sphere of radius r_k , with x_i at its centre. This is precisely what we want for discretising the hard spheres/disks problem. For the square lattice, and thus two dimensions, this is illustrated in fig. 5. In the limit $k \rightarrow \infty$, the k NN model could intuitively be expected to exhibit the same behaviour as its continuous counterpart, the hard spheres model. For the $d = 2$ case in particular, this means that for large k the model would exhibit the same two-step melting behaviour as the hard disks model. Note however, that there is a fundamental difference: The hard disks model exhibits a *continuous* symmetry while the symmetry of a k NN model is discrete.

As a first example, consider the k NN model in one spatial dimension. It is equivalent to a model of discrete hard rods which all have a length of $k + 1$ lattice sites, see fig. 6. The hard rods model undergoes no phase transitions




Figure 7: Shapes of hard core particles in the k NN model for $k = 1, \dots, 5$: Filled (empty) circles denote occupied (excluded) lattice sites.

for $\mu < \infty$ and can be solved exactly. A detailed discussion of technical details is postponed to section 3.1. The absence of a phase transition implies that there is no solid phase for finite chemical potential: All sublattice densities are the same, and translation-invariance is preserved. However, the system's correlation length diverges for $\mu \rightarrow \infty$ which means that there is a phase transition at infinite chemical potential.

The k NN model in $d = 2$ exhibits a much richer behaviour. This model has received considerable attention over the years (see [11, 30, 24] for recent contributions). The following discussion builds upon these references.

With increasing density, the systems generally undergo phase transitions where the sublattice symmetry breaks: One or more sublattices become preferentially occupied, and the sublattice densities are no longer equal to each other. These sublattice densities will be used to construct the order parameter.

Shapes associated with the particles for a given k may be obtained by looking at fully packed configurations, see 7. Since the model is discrete, the shapes associated with the particles for a given k are no longer perfect disks.

Consider the 1NN model. It is equivalent to a model of hard squares which are tilted by $\pi/4$. There are two fully packed configurations, and thus the ground state is twofold degenerate. These two configurations are related




Figure 8: The 1NN model has two different fully-packed configurations: Either sublattice A or B is fully occupied while the other one is completely empty.

by a translation of all particle positions by one lattice site in either the vertical or horizontal direction, (see fig. 8). The set of lattice sites can be split into two subsets, A and B , such that either of these sublattices is fully occupied in the ground state, while the other one remains empty. Densities ρ_A and ρ_B may be defined separately for both sublattices. Thus in a fully packed configuration either $\rho_A = 1 \wedge \rho_B = 0$ or $\rho_A = 0 \wedge \rho_B = 1$. More generally, either of the two sublattices becomes preferentially occupied in the high-density regime i.e either $\rho_A < \rho_B$ or $\rho_B < \rho_A$. For very low densities however, the sublattice densities are equal to each other. Thus a phase transition takes place at a certain density (and thus chemical potential). As an order parameter the difference in sublattice-densities may be chosen:

$$q_1 = \rho_A - \rho_B \quad (11)$$

From numerical studies it is known [15], that at a critical chemical potential of

$$\mu_c = 1.33401510027774(1)$$

the 1NN model undergoes a second order phase transition. The discrete symmetry of translations by one site is broken for $\mu > \mu_c$.

The 2NN model can also be considered as a model of 2 by 2 hard squares. Without loss of generality, a particle is always in the top left corner of such a square, see fig. 7. The 2NN model has infinitely many fully packed configurations: The squares may align themselves along either rows or columns which may slide freely along each other in either the horizontal or vertical direction, respectively. This so called sliding instability gives rise to the infinitely degenerate ground state, see fig. 9. Previous studies [11, 30] provide strong evidence that the 2NN model exhibits two distinct phases: A low-density disordered phase and columnar phase for high densities. In the latter, the squares preferentially align along either rows or columns of lattice sites. The transition between those phases has been found to be of second order, with a critical chemical potential of [11]

$$\mu_c \approx 4.574$$




Figure 9: Examples of possible fully packed configurations of the 2NN model: They are either column- (top) or row-ordered (bottom).

Due to the size of the squares, only either even or odd rows/columns become preferentially occupied for $\mu > \mu_c$. Therefore there are four different ways in which the system may order. They are related by rotations of $\pi/4$ of the entire lattice. It suffices to divide the lattice into four different sublattices A , B , C and D to characterise the way in which the system orders. In the ordered phase, one and only one of the following relations between the four sublattice-densities holds:

$$\begin{aligned} \text{row-ordered} \quad & \begin{cases} (\rho_A = \rho_B) > (\rho_D = \rho_C) \\ (\rho_A = \rho_B) < (\rho_D = \rho_C) \end{cases} \\ \text{column-ordered} \quad & \begin{cases} (\rho_A = \rho_C) > (\rho_B = \rho_D) \\ (\rho_A = \rho_C) < (\rho_B = \rho_D) \end{cases} \end{aligned}$$

The maximum mean density is $1/4$ and the maximum sublattice-density is $1/2$, due to the sliding instability. Since there are four different sublattices, we obtain six different density differences:

$$\begin{aligned} \rho_A - \rho_B, \quad \rho_A - \rho_C, \quad \rho_A - \rho_D, \\ \rho_B - \rho_C, \quad \rho_B - \rho_D, \quad \rho_C - \rho_D, \end{aligned}$$

An order parameter q_2 can be defined by picking the terms corresponding to diagonal sites:

$$q_2 = \sqrt{2}(\rho_A - \rho_C + i(\rho_B - \rho_D)) \quad (12)$$




Figure 10: Two fully packed configurations of the 3NN model which are related by mirror symmetry. For any fully packed configuration a labelling can be chosen, such that all particles sit in one and only one sublattice. This remains true if the system is translated in the horizontal or vertical direction. In the mirrored system however, the average occupation number of each sublattice site is $1/5$.

The four possible phases $\pi/4, 3\pi/4, 5\pi/4, 7\pi/4$ of q_2 in the ordered phase correspond to the four different orderings [30].

The 3NN model is equivalent to a HCLG of cross-shaped pentamers. Five fully packed configurations may be found, which are related by translations of one to five lattice sites in either horizontal or vertical direction. The mirror images of these five configuration however, are again fully packed configurations which are different from the original ones. Thus the system also breaks mirror symmetry, and its ground state is therefore ten-fold degenerate, see fig. 10. It is therefore no longer possible to choose a sublattice-labelling, such that sites with the same label become preferentially occupied for all ways in which the symmetry may break. Rather there are two such labellings - one being the mirror image of the other [11]. The particle density for a given sublattice now also depends on the chosen labelling denoted by L or R . For denoting sublattice densities, the symbols $\rho_{l,L}$ and $\rho_{l,R}$ are used, with $l = A, \dots, E$ indexing the sublattice. For a chosen labelling, say R , there are 10 distinct density differences $\{(\rho_{A,R} - \rho_{B,R}), (\rho_{A,R} - \rho_{C,R}), \dots, (\rho_{D,R} - \rho_{E,R})\}$ from which an order parameter may be constructed. In [11] it was found, that at a critical chemical potential of

$$\mu_c \approx 3.67$$

the 3NN model undergoes a first order phase transition from a fluid to a solid phase.

Since no results for $k > 3$ were obtained in this work, the reader is referred to [11, 24] (and references therein) for recent results concerning these cases. It should be noted, that already for $k = 4$ two phase transitions were found.




Figure 11: Various examples of the diagrammatic tensor notation: (i) scalar, (ii) vector, (iii) matrix, (iv) scalar product of two vectors, (v) product of two matrices, (vi) trace of a matrix.

3 Tensor networks and classical systems

Tensor networks originated in the study of the quantum many-body problem. However it turns out, that the partition function of a classical lattice system with only local interactions, may be formulated in terms of a tensor network quite straightforwardly. Since this work is mostly concerned with classical systems, tensor networks shall be introduced from this perspective. Before doing so, the diagrammatic notation used in the field of tensor networks, is reviewed briefly.

The term *tensor* shall refer to any multi-indexed object $T \in \mathbb{C}^{d_1 \times d_2 \times \dots \times d_N}$. The number of indices of a given tensor, is called its order⁶. In this terminology, scalars, vectors and matrices are tensors of order zero, one and two, respectively. In the field of tensor networks, heavy use is made of a diagrammatic notation: Tensors are represented by dots and their indices by lines. An order four tensor, for example, may be written as:

$$T_{abcd} = \begin{array}{c} b \\ | \\ a - \text{blue dot} - c \\ | \\ d \end{array}$$


Most of the time, letters labelling the indices will be omitted. If both ends of a line are attached to a dot representing a tensor, this simply means that the corresponding index is contracted, i.e. summed over. A few examples are given in fig. 11. From the diagrammatic notation the meaning of the term “tensor network” may already be guessed. To quote [27]:

“A Tensor Network (TN) is a set of tensors where some, or all, of its indices are contracted according to some pattern.”

If all indices of a TN are contracted, it will also be called *fully* contracted.

⁶The word “rank” is also used frequently. In this work, this term is reserved for the rank of matrices, to avoid confusion.

For this case the symbol tTr is reserved:



Note that this notation is merely symbolic - tTr does not specify the exact way in which way the tensors are contracted. Its only purpose is to tell us which tensors are in the network and that all indices have been contracted.

3.1 k NN in $d = 1$ and the transfer matrix

As a first example, the tensor network representation of classical partition functions shall be illustrated by reviewing the well-known transfer matrix method for one dimensional systems. The already mentioned model of hard rods will serve as an illustration. Many issues of setting up the correct tensor network representation are already encountered in this case which is however easier to deal with. Furthermore, the concept of the transfer matrix is also very important for studying matrix product states (a particular kind of TN representation for quantum states) which will be introduced later.

The idea of the transfer matrix method is to rewrite the global weight $w(\boldsymbol{\sigma})$ of a given configuration $\boldsymbol{\sigma}$ as a product of local weights. This follows quite naturally for most models, since their Hamiltonian is a sum of local terms to begin with. Consider for instance, a spin model which has a Hamiltonian which is the sum of nearest-neighbour terms:

$$\begin{aligned} \mu n(\boldsymbol{\sigma}) - H(\boldsymbol{\sigma}) &= \sum_i \tilde{H}(\sigma_i, \sigma_{i+1}) \\ \Rightarrow w(\boldsymbol{\sigma}) &= e^{\mu n(\boldsymbol{\sigma}) - H(\boldsymbol{\sigma})} = \prod_i e^{\tilde{H}(\sigma_i, \sigma_{i+1})} =: \prod_i T(\sigma_i, \sigma_{i+1}) \end{aligned}$$

The weight $T(\sigma_i, \sigma_{i+1})$ has a quite straightforward interpretation: Without loss of generality, spins may be labelled from left to right, and thus $T(\sigma_i, \sigma_{i+1})$ is the statistical weight that a spin on site i with value σ_i , has a spin with value σ_{i+1} sitting to its right.

For an athermal system in one dimension, such as the hard rods model, we may look directly at the weight and make the ansatz, that for a certain configuration the overall weight can be broken down into a product of local ones:

$$w(\boldsymbol{\sigma}) = e^{\mu n(\boldsymbol{\sigma})} \Gamma(\boldsymbol{\sigma}) = \prod_{i=1}^N T(\sigma_i, \sigma_{i+1}) \quad (13)$$




Figure 12: “Spin” configurations for the hard rods model.

In order to obtain the desired properties for the 1NN model, T must fulfill:

$$T(0,0) = 1 \quad T(1,1) = 0 \quad T(0,1)T(1,0) = e^\mu$$

To see this, consider a configuration σ with n occupied sites i.e. $n = \sum_i \sigma_i$. If there are at least two adjacent occupied sites (representing overlapping rods), the term $T(1,1)$ appears and the entire expression collapses to zero. Otherwise, it is equal to the product $(e^\mu)^n 1^{N-n}$. Thus:

$$\prod_{i=1}^N T(\sigma_i, \sigma_{i+1}) = \begin{cases} 0 & \text{overlapping rods} \\ e^{\mu n} & \text{else} \end{cases}$$

One specific choice of weights is:

$$T(0,0) = 1 = T(1,0) \quad T(0,1) = e^\mu \quad T(1,1) = 0$$

Note that each weight may be thought of as corresponding to a particular bond between two spins - or alternatively to the vertices of the *dual* lattice⁷. The crucial insight of the transfer matrix method, is to interpret these weights as the entries of a matrix T , with indices σ and σ' :

$$T_{\sigma\sigma'} := T(\sigma, \sigma')$$

The partition function for a system with periodic boundary conditions

$$\sigma_{N+1} = \sigma_1$$

⁷It may be interesting to note, that in the latter interpretation, the 1NN model on the line is actually nothing but a system of dimers.




Figure 13: Partition function of a one dimensional system for (a) open and (b) periodic boundary conditions. The red crosses mark the physical degrees of freedom.

then turns out to be the trace of T to the power of the number of lattice sites:

$$Z(\mu) = \text{Tr}_{\boldsymbol{\sigma}} \prod_{i=1}^N T(\sigma_i, \sigma_{i+1}) = \sum_{\sigma_1} \cdots \sum_{\sigma_N} T_{\sigma_1 \sigma_2} \cdots T_{\sigma_{N-1} \sigma_N} T_{\sigma_N \sigma_1} = \text{tr} (T^N)$$

Since the trace of a product of matrices is nothing but a fully contracted network of order two tensors, we have achieved our initial goal of rewriting the partition function as a tensor network (with $\text{tTr} = \text{tr}$), at least for the case of periodic boundary conditions (PBC). For open or fixed boundary conditions, the introduction of left and right boundary vectors $x_L^\dagger$ and x_R is necessary:

$$Z(\mu) = x_L^\dagger T^{N-1} x_R \quad (14)$$

Open boundary conditions (OBC) mean that the sum is taken over all possible configurations at the boundary, while fixed boundary conditions (FBC) enforce a particular configuration at the boundary. Thus the respective boundary vectors are given by:

$$\begin{aligned} (x_L^\dagger)_j &= \sum_i T_{ij} & (x_R)_i &= \sum_j T_{ij} & (\text{PBC}) \\ (x_L^\dagger)_j &= \sum_i \delta_{aj} T_{ij} & (x_R)_i &= \sum_j \delta_{ib} T_{ij} & (\text{FBC}) \end{aligned}$$

In fig. 13 the transfer matrix formulation is expressed in TN notation.

In the case of $k > 1$, the transfer matrix needs to somehow enforce the condition that if a site i is occupied, i.e $\sigma_i = 1$, the following k sites are empty. But since the only possible occupation numbers are 0 and 1, there is no way of telling just by looking at a single site whether, if say a particular site is occupied, the previous k sites have all been unoccupied. This problem can be solved by extending the range of values for the occupation numbers to $\sigma_i = 0, \dots, k$ (see fig. 12). For arbitrary k , the transfer matrix can then be chosen to have the following non-zero entries:

$$T_{0,0} = T_{k+1,0} = 1, \quad T_{i,i+1} = e^{\mu/k} \quad (0 < i < k)$$

Written down explicitly, T reads:

$$T = \begin{pmatrix} 1 & e^{\mu/k} & 0 & 0 & \dots & 0 \\ 0 & 0 & e^{\mu/k} & 0 & \dots & 0 \\ 0 & 0 & 0 & e^{\mu/k} & \dots & 0 \\ \vdots & & & & \ddots & \vdots \\ & & & & & e^{\mu/k} \\ 1 & 0 & 0 & 0 & \dots & 0 \end{pmatrix} \quad (15)$$

The only constraint on the elements in the super diagonal is that their product must be equal to the fugacity. Note how the characteristic polynomial

$$\lambda^{k+1} - \lambda^k - e^\mu = 0 \quad (16)$$

is left invariant by the specific choice of “distributing” the fugacity among the entries in the superdiagonal: One could also choose

$$T_{0,1} = e^\mu, \quad 1 = T_{0,0} = T_{k+1,0} = T_{i,i+1} \quad (0 < i < k)$$

and obtain the same characteristic polynomial.

The power of the transfer matrix method lies in the fact, that it turns the calculation of the partition function into an eigenvalue problem: It turns out, that in the thermodynamic limit of $N \rightarrow \infty$, the partition function per site κ is in fact equal to the *dominant eigenvalue* λ_0 of the transfer matrix:

$$\kappa = \lim_{N \rightarrow \infty} Z^{1/N} = \lambda_0 \quad (17)$$

Assuming the existence of a dominant eigenvalue (more on that shortly), we can see this by considering the Jordan canonical form of the transfer matrix

$$T = SJS^{-1}$$

and defining

$$\tilde{T} := \frac{1}{\lambda_0} T = S\tilde{J}S^{-1}, \quad \tilde{J} := \frac{1}{\lambda_0} J \quad .$$

We then have

$$\lim_{N \rightarrow \infty} \left(\tilde{T} \right)^N = v_R v_L^\dagger =: E$$

with the dominant left and right eigenvectors $v_L^\dagger$ and v_R (see A.2). This result can be readily applied to prove (17) for the case of periodic boundary conditions:

$$Z = \text{tr}(T^N) = \lambda_0^N \text{tr} \tilde{T}^N \rightarrow \lambda_0^N \quad (N \rightarrow \infty)$$

where it was used that $\text{tr} E = 1$. For the case of open/fixed boundary conditions, we have:

$$Z = x_L^\dagger T^N x_R \rightarrow \lambda_0^N (x_L^\dagger E x_R), \quad (N \rightarrow \infty)$$

The partition function for large N is therefore equal to the case of PBC, up to a constant pre-factor. For the partition function per site we thus again obtain:

$$\kappa = \lambda_0$$

Furthermore, the correlation length of the system may readily be obtained from the two largest magnitude eigenvalues of the transfer matrix.

$$\xi = \frac{1}{\ln |\lambda_0/\lambda_1|} \quad (18)$$

To see this, we consider the two-point correlator of two one-site observables o and o' which are separated by a distance of r lattice sites:

$$G(o_i, o'_{i+r}) = \langle o_i o'_{i+r} \rangle - \langle o_i \rangle \langle o'_{i+r} \rangle = \frac{v_L^\dagger A T^r B v_R}{\lambda_0^{r+2}} - \frac{v_L^\dagger A v_R}{\lambda_0} \frac{v_L^\dagger B v_R}{\lambda_0} \quad (19)$$

T is then decomposed into “dominant” and “orthogonal” part

$$T = S J S^{-1} = \lambda_0 E + \lambda_1 S \tilde{J}_\perp S^{-1} = \lambda_0 E + \lambda_1 \tilde{T}_\perp,$$

where

$$\tilde{J}_\perp := \frac{1}{\lambda_1} J_\perp := \begin{pmatrix} 0 & & & 0 \\ & \frac{1}{\lambda_1} J_{k_1}(\lambda_1) & & \\ & & \ddots & \\ 0 & & & \frac{1}{\lambda_1} J_{k_n}(\lambda_n) \end{pmatrix}, \quad \tilde{T}_\perp := \frac{1}{\lambda_1} T_\perp = S \tilde{J}_\perp S^{-1} \quad .$$

By noting that

$$T^N = \lambda_0^N E + \lambda_1^N S \tilde{J}_\perp^N S^{-1} = \lambda_0^N E + \lambda_1^N \tilde{T}_\perp^N$$

we have

$$\begin{aligned}
\frac{v_L^\dagger A T^r B v_R}{\lambda_0^{r+2}} &= \frac{\lambda_0^r v_L^\dagger A E^r B v_R + \lambda_1^r v_L^\dagger A \tilde{T}_\perp^r B v_R}{\lambda_0^{r+2}} \\
&= \frac{v_L^\dagger A v_R v_L^\dagger B v_R}{\lambda_0^2} + \left(\frac{\lambda_1}{\lambda_0}\right)^r \frac{v_L^\dagger A \tilde{T}_\perp^r B v_R}{\lambda_0^2} \\
&= \langle o_i \rangle \langle o'_{i+r} \rangle + \left(\frac{\lambda_1}{\lambda_0}\right)^r \frac{v_L^\dagger A \tilde{T}_\perp^r B v_R}{\lambda_0^2} \\
\Rightarrow \quad G(o_i, o'_{i+r}) &= e^{i\varphi r} \left| \frac{\lambda_1}{\lambda_0} \right|^r \frac{v_L^\dagger A \tilde{T}_\perp^r B v_R}{\lambda_0^2}, \quad e^{i\varphi} = \lambda_1/|\lambda_1|
\end{aligned}$$

In the limit of large r , we have

$$\frac{v_L^\dagger A \tilde{T}_\perp^r B v_R}{\lambda_0^2} \rightarrow C$$

for some constant $C \in \mathbb{C}$, since

$$\tilde{T}_\perp^r = S \tilde{J}_\perp^r S^{-1} \rightarrow \begin{pmatrix} 0 & & \\ & \mathbb{1}_{k_1} & \\ & 0 & \\ & & \ddots \end{pmatrix}, \quad (r \rightarrow \infty)$$

Thus we finally obtain

$$G(o_i, o'_{i+r}) \propto \left| \frac{\lambda_1}{\lambda_0} \right|^r = e^{-r/\xi}$$

up to some phase $e^{i\varphi r}$ and in the limit of large r .

For the $1NN$ model in $d = 1$ for example, we obtain for the partition function and the correlation length:

$$\kappa = 1 + \sqrt{1 + 4e^\mu} \tag{20}$$

$$\xi = [\ln(1 + \sqrt{1 + 4e^\mu}) - \ln(1 - \sqrt{1 + 4e^\mu})]^{-1} \tag{21}$$

For this case, the correlation length thus diverges in the limit $\mu \rightarrow \infty$.

Since the transfer matrix for the hard rods example is not positive but merely non-negative, it is not immediately clear, whether it is irreducible and thus has a dominant eigenvalue for all k . If T is irreducible, then

$$(\mathbb{1}_{k+1} + T)^k$$

is strictly positive. To see that this is in fact the case, we take a look at the sparsity pattern:

$$(\mathbb{1}_{k+1} + T)^N = \begin{pmatrix} * & * & & & & \\ & * & * & & & \\ & & * & * & & \\ & & & \ddots & \ddots & \\ & & & & * & * \\ * & & & & & * \end{pmatrix}^N$$

It follows that $(\mathbb{1}_{k+1} + T)^N$ is positive for $N \geq k$. Thus for the hard rods model in one dimension, there always exists a non-degenerate largest eigenvalue for the transfer matrix, and therefore no phase transitions take place. Provided that the number of tensors in the network is an integer multiple of k , the TN may also be coarse grained:

$$Z = \text{tr } T^N = \text{tr } T'^{N/k}, \quad T' = T^k$$

As was just shown, the new transfer matrix T' is now strictly positive.

3.2 k NN in $d = 2$

In the $d = 1$ case, the partition function was written as the trace of a matrix product. In principle, this transfer matrix formalism can immediately be extended to higher dimensions by approximating a two dimensional system by a stripe (or a cylinder, for PBC) which is finite in one dimension and infinite in the other. Such a system is effectively one dimensional and the transfer matrix formalism applies. A drawback of this approach is, that the transfer matrix grows with system size (i.e the “width” of the stripe or the circumference of the cylinder). In the past, the k NN model has been studied via this approach (e.g. [4]). Our aim however, is to treat higher dimensional systems by writing the partition function as a fully contracted network of local tensors T whose order may now be greater than two.

$$Z(\mu) = \text{Tr}_{\boldsymbol{\sigma}} e^{\mu n(\boldsymbol{\sigma})} \Gamma(\boldsymbol{\sigma}) = \text{tTr } T(\mu)^N \quad . \quad (22)$$

As for the $d = 1$ transfer matrix, the trace over all tensor indices is just the sum over all physical degrees of freedom:

$$\text{tTr} = \text{Tr}_{\boldsymbol{\sigma}}$$

Tensor networks and classical vertex models are closely related. To see this, consider the graphical notation of a TN and regard tensors as lattice sites and their legs/indices as edges. A particular entry of a tensor is now nothing




Figure 14: Network used for encoding the 1NN model: On every other plaquette there is a four-leg tensor whose indices correspond to occupation numbers $\{0, 1\}$. There are thus twice as many lattice sites as tensors. Red crosses mark the physical degrees of freedom.

but a weight for a given vertex, whose configuration is fully determined by the indices of the tensor. Summing over all indices is therefore equivalent to summing over all possible vertex configurations. Thus vertex models are most naturally expressed as tensor networks, and finding a TN representation of a classical system basically boils down to finding an appropriate mapping to a vertex model. However, this mapping and therefore the tensor network representation, are generally not unique, and for the sake of computational speed, it is generally desirable to have tensors with low bond dimensions.

Consider again the k NN model, but this time on the square lattice. We shall at first treat the simplest case, the 1NN model, which can be encoded into a TN similar to the one used in [10], see fig. 14. Occupation numbers directly correspond to the indices of an order four local tensor T , which may be chosen to have the following non-zero entries:

$$\begin{aligned}
 T_{ijkl} : \quad & \begin{array}{c} 0 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 0 \end{array} = 1 \\
 & \begin{array}{c} 1 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 0 \end{array} = \begin{array}{c} 0 \quad 1 \\ \diagdown \quad \diagup \\ 0 \quad 0 \end{array} = \begin{array}{c} 0 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 1 \end{array} = \begin{array}{c} 0 \quad 0 \\ \diagdown \quad \diagup \\ 1 \quad 0 \end{array} = e^{\mu/2} \\
 & \begin{array}{c} 1 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 1 \end{array} = \begin{array}{c} 0 \quad 1 \\ \diagdown \quad \diagup \\ 1 \quad 0 \end{array} = e^{\mu}
 \end{aligned} \tag{23}$$

As in the one dimensional case, the problem arises of how to “communicate” the shape of a particle across the tensor network for $k > 1$. For the hard rods model this was solved by assigning an unique spin value to each lattice site being part of a hard rod. In principle, something similar could be done for hard-core particles on the square lattice. However, the number of different spin values and thus the bond dimension of the local tensors, now





Figure 15: Tensor network for a vertex model on the 15a dual and the 15b original lattice.

grows with the area of the particles as compared to growing with the length of the rods as in $d = 1$. For the k NN model this would mean $D = \mathcal{O}(k^2)$ in the plane as compared to $D = \mathcal{O}(k)$ on the line.

A more efficient way of encoding a HCLG is by finding a vertex model which only encodes the trails making up the borders of the particles. This allows for a TN representation with a bond dimension $D = \mathcal{O}(l)$, where l is the length of the trail i.e the number of edges it consists of.

As a first example, the k NN model for $k = 2, 3$ is considered. First, to each edge connecting two adjacent lattice sites, a bond index $\sigma \in \{0, 1, 2\}$ is assigned according to the following rules:



Filled (empty) dots represent occupied (empty) lattice sites, as usual. Note, that the case of two adjacent occupied sites is not included, since the overall weight for such a case must be zero anyway according to the model. Including this forbidden case would however require to extend the bond dimension to four, without changing model. For the 2NN model, the tensors sit on the sites of the dual lattice and their non-zero entries are given by

$$\begin{array}{c} 0 \\ | \\ 0 \text{---} \text{---} 0 \\ | \\ 0 \end{array} = 1, \quad \begin{array}{c} 1 \\ | \\ 0 \text{---} \text{---} 1 \\ | \\ 0 \end{array} = 0 \begin{array}{c} 0 \\ | \\ 0 \text{---} \text{---} 2 \\ | \\ 1 \end{array} = 2 \begin{array}{c} 0 \\ | \\ 2 \text{---} \text{---} 0 \\ | \\ 2 \end{array} = 1 \begin{array}{c} 2 \\ | \\ 1 \text{---} \text{---} 0 \\ | \\ 0 \end{array} = e^{\mu/4} \quad (24)$$

while the 3NN model is encoded into a TN with tensors on the vertices of the lattice (see fig. 15b), with non-zero entries:

$$\begin{array}{c} 0 \\ | \\ 0 \text{---} \text{---} 0 \\ | \\ 0 \end{array} = 1, \quad \begin{array}{c} 1 \\ | \\ 0 \text{---} \text{---} 0 \\ | \\ 0 \end{array} = 0 \begin{array}{c} 0 \\ | \\ 0 \text{---} \text{---} 0 \\ | \\ 2 \end{array} = 2 \begin{array}{c} 0 \\ | \\ 2 \text{---} \text{---} 0 \\ | \\ 0 \end{array} = 0 \begin{array}{c} 0 \\ | \\ 0 \text{---} \text{---} 1 \\ | \\ 0 \end{array} = 1 \begin{array}{c} 2 \\ | \\ 1 \text{---} \text{---} 2 \\ | \\ 1 \end{array} = e^{\mu/5}$$




Figure 16: Representation of the 16a 2NN and 16a 3NN model as a vertex model. Black edges correspond to a bond index = 0 while red edges to bond indices $\neq 0$.

This idea can be readily generalised to higher k or other HCLG. For a given model, the set of possible vertex configurations can be regarded as building blocks for trails. They are chosen in a way so that there is one and only one trail that may be constructed and its shape is such, that it reproduces the shape of the relevant hard-core particle (see 16). To enforce this uniqueness of trails, each bond index except the one representing “vacuum” (usually 0) may only be used once along each direction. For the square lattice, bonds are either vertically or horizontally aligned, and therefore a trail consisting of l_h horizontal and l_v vertical edges can be represented by a tensor with bond dimension

$$D = 1 + \max\{l_h, l_v\} \quad .$$

The 5NN model is encoded by the following non-zero weights:

$$\begin{aligned}
 0 \begin{array}{c} 0 \\ | \\ 0 \end{array} 0 &= 1, & 0 \begin{array}{c} 1 \\ | \\ 0 \end{array} 1 &= 0 \begin{array}{c} 0 \\ | \\ 2 \end{array} 2 = 3 \begin{array}{c} 0 \\ | \\ 3 \end{array} 0 = 4 \begin{array}{c} 4 \\ | \\ 0 \end{array} 0 = e^{\mu/8} \\
 0 \begin{array}{c} 2 \\ | \\ 1 \end{array} 0 &= 2 \begin{array}{c} 0 \\ | \\ 0 \end{array} 3 = 0 \begin{array}{c} 3 \\ | \\ 4 \end{array} 0 = 1 \begin{array}{c} 0 \\ | \\ 4 \end{array} 4 = e^{\mu/8}
 \end{aligned}$$

By also adding

$$3 \begin{array}{c} 1 \\ | \\ 3 \end{array} 1 = 4 \begin{array}{c} 4 \\ | \\ 2 \end{array} 2 = e^{\mu/4}$$

the 4NN model is obtained, see fig. 17.

3.3 Local observables and environments

Once the partition function (per site) has been found, the density and other macroscopic quantities may be obtained via differentiation of the free energy/partition function. However, taking derivatives numerically, is very




Figure 17: Examples of fully packed configurations of the 17a 4NN and 17b 5NN model.

ill-conditioned and should be avoided if possible. Furthermore, it is not possible to evaluate sub-lattice densities this way which will be necessary for computing order parameters. Thus it is also desirable to have a means of evaluating the density as a local observable. Assuming for simplicity that each tensor corresponds to one and only one lattice site and thus $N = N_0$, the expectation value of an one-site observable o is given by

$$\langle o \rangle = \frac{\text{tTr}(T^{N-1}A)}{\text{tTr} T^N} . \quad (25)$$

Here, the “impurity” tensor A encodes the observable o . A very important example is the density operator R which may be obtained [43] by taking the derivative of the fully contracted tensor network:

$$\frac{\partial Z}{\partial \mu} = N \text{tTr}(T^{N-1}R), \quad R := \frac{\partial T}{\partial \mu}$$

The (local) density is thus given by

$$\rho = \frac{N}{N_0} \frac{\text{tTr}(T^{N-1}R)}{\text{tTr} T^N}$$

Note however, that this is the mean density operator in the sense that its expectation value is the average over all sublattice densities. As a first illustration, consider the k NN model in $d = 1$. For the encoding (15) the non-zero entries of the density operator are given by

$$R_{i,i+1} = \frac{1}{k} e^{\mu/k} \quad (0 < i < k)$$

However, the k NN model in $d = 1$ does not exhibit sublattice symmetry breaking, and there is thus no need for evaluating sublattice densities.

For $d = 2$, we first consider the 1NN model. Its encoding (23) has the advantage that the order parameter is a one-site observable. To see this,

note that all diagonal leg-indices correspond to spins on one of the two sublattices, say w.l.o.g. A , while the anti-diagonal leg-indices correspond to spins on sublattice B . The density operator $R := \partial T / \partial \mu$ may be written as

$$R = R^A + R^B$$

where R^A and R^B are the density operators for sublattices A and B . Their non-zero entries are:

$$\begin{aligned} R_{ijkl}^A : \quad & \begin{array}{c} 1 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 0 \end{array} = \begin{array}{c} 0 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 1 \end{array} = \frac{1}{2} e^{\mu/2}, & \begin{array}{c} 1 \quad 0 \\ \diagdown \quad \diagup \\ 0 \quad 1 \end{array} = e^{\mu} \\ R_{ijkl}^B : \quad & \begin{array}{c} 0 \quad 1 \\ \diagdown \quad \diagup \\ 0 \quad 0 \end{array} = \begin{array}{c} 0 \quad 0 \\ \diagdown \quad \diagup \\ 1 \quad 0 \end{array} = \frac{1}{2} e^{\mu/2}, & \begin{array}{c} 0 \quad 1 \\ \diagdown \quad \diagup \\ 1 \quad 0 \end{array} = e^{\mu} \end{aligned}$$

The order parameter q_1 is now given by the expectation value of the operator

$$Q_1 = R^A - R^B$$


The situation is more complicated for the 2NN model, since in the encoding (24) the order parameter cannot be written as a single-site observable. Nevertheless, the operator for the mean density is still given by $R := \partial T / \partial \mu$, which can be written as

$$R = R^{NW} + R^{NE} + R^{SE} + R^{SW}$$

where the four tensors on the r.h.s. only have a single non-zero entry each:

$$\begin{aligned} R^{NW} : \quad & \begin{array}{c} 2 \\ \text{red corner} \\ 1 \quad 0 \\ | \\ 0 \end{array} = \frac{1}{4} e^{\mu/4} & R^{NE} : \quad \begin{array}{c} 1 \\ \text{red corner} \\ 0 \quad 1 \\ | \\ 0 \end{array} = \frac{1}{4} e^{\mu/4} \\ R^{SW} : \quad & \begin{array}{c} 0 \\ \text{red corner} \\ 2 \quad 0 \\ | \\ 2 \end{array} = \frac{1}{4} e^{\mu/4} & R^{SE} : \quad \begin{array}{c} 0 \\ \text{red corner} \\ 0 \quad 2 \\ | \\ 1 \end{array} = \frac{1}{4} e^{\mu/4} \end{aligned}$$

Note that these tensors are *not* operators for the four sublattice densities⁸! In order to come up with the correct operators, the network can be coarse-grained by blocking four tensors together. A choice is made in regard to the labelling of the surrounding sites on the original (not the dual) lattice. For example:



⁸ For example, if R^{NW} is treated as a single site observable, its expectation value is merely the mean occupation number of the “north west” or top-left corner of *any* plaquette. But since any site on the original lattice is the top left corner of some plaquette, the expectation value of R^{NW} is nothing but one quarter of the overall density.




Figure 18: One-site environment of a tensor network making up a square lattice.

Now that the coarse-grained tensor T' has been introduced, the coarse grained density operator is given by $R' = \partial T' / \partial \mu$. (Note, that the expectation value of R' is no longer the density of a single site!) The approach just described may readily be generalised to the cases where $k > 2$. We are now in a position to come up with an operator for the four sublattice densities. This is achieved by formally introducing a distinct chemical potential for each sublattice, and taking the derivative with respect to the chemical potential for the sublattice of interest. This leads to an expression for the density operator of, e.g., sublattice A :

$$R^A = \frac{1}{4} \frac{\partial}{\partial \mu_A} \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \bullet \quad \bullet \\ | \quad | \\ \text{---} \text{---} \end{array}$$

$$\frac{\partial}{\partial \mu_A} \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \bullet \quad \bullet \\ | \quad | \\ \text{---} \text{---} \end{array} = \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \bullet \quad \bullet \\ | \quad | \\ \text{---} \text{---} \end{array} + \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \text{---} \text{---} \end{array} + \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \text{---} \text{---} \end{array} + \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \text{---} \text{---} \end{array} + \begin{array}{c} \text{---} \text{---} \\ | \quad | \\ \text{---} \text{---} \end{array}$$

The impurities in the corners are just the tensors R^{NW} , R^{NE} , R^{SE} and R^{SW} . Having obtained density operators for sublattices A to D , we are in the position to construct an operator Q_2 for the order parameter:

$$Q_2 = \sqrt{2} [R^A - R^C + i (R^B - R^D)]$$

A few general remarks about expectation values of observables for tensor networks are in order. Computation of local observables is closely related to the computation of *environments*. An environment of a particular tensor in a fully contracted network is obtained by removing this tensor and leaving the legs of the tensors it was previously connected to, open. By contracting all the other indices, an environment tensor E is obtained (see fig. 18).

$$Z = \text{tTr } T^N = \text{tTr } ET$$

The entries of E generally grow with system size. For actual calculations, E must therefore be normalized. It is convenient, to choose the normalization

4 Tensor networks and quantum systems

Tensor networks originated in the study of quantum many body physics out of attempt to handle states living in exponentially large Hilbert spaces. A tensor network *state* is an encoding of a quantum state into a tensor network. Consider a prototypical N -body state

$$|\Psi\rangle = \sum_{\{s_i\}} c_{s_1, \dots, s_N} |s_1\rangle \dots |s_N\rangle \in \mathcal{H}_N, \quad (27)$$

where $\mathcal{H}_N$ denotes the N -body Hilbert space. The individual degrees of freedom $|s_i\rangle$ live in copies of a local, finite dimensional Hilbert space $\mathcal{H} \simeq \mathbb{C}^d$:

$$|s_i\rangle \in \mathcal{H}, \quad \mathcal{H}_N = \mathcal{H}^{\otimes N}$$

The wave function of a general $|\Psi\rangle$ is made up of d^N coefficients $c_{s_1, \dots, s_N}$, i.e. exponentially many. One way of looking at tensor networks for quantum many-body states is as a means of compressing a state, so that it can be expressed by merely polynomially many coefficients. This is achieved by rewriting the coefficients as a product of local tensors. The concept of quantum-entanglement plays a central role. Before elaborating on tensor networks in more detail, this important concept from quantum information theory is therefore reviewed briefly.

4.1 Entanglement

Consider a bipartite quantum system composed of subsystems A and B . Let $\{|\varphi_i\rangle\}$ and $\{|\chi_j\rangle\}$ span the Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively. The state $|\Psi\rangle$ of the overall system can then be written as:

$$|\Psi\rangle = \sum_{ij} c_{ij} |\varphi_i\rangle |\chi_j\rangle, \quad c_{ij} \in \mathbb{C}$$

The subsystems A and B are said to be *entangled* if there exist no local bases $\{|\tilde{\varphi}_i\rangle\}$ and $\{|\tilde{\chi}_j\rangle\}$ such that $|\Psi\rangle$ can be written as a product state, i.e

$$|\Psi\rangle = |\tilde{\varphi}_i\rangle |\tilde{\chi}_j\rangle$$

for some specific choice of i, j . If a state is not entangled, it is called *separable*. As an example of an entangled state, consider one of the four so-called Bell states

$$|\Psi^-\rangle = \frac{|0\rangle |1\rangle - |1\rangle |0\rangle}{\sqrt{2}},$$

where the vectors $|1\rangle$ and $|0\rangle$ span a two-dimensional Hilbert space. Entanglement in a composite system is not simply present or absent. Two states may be somewhere between separable and what will be called maximally

entangled. To make this precise, it is necessary to come up with a way to quantify entanglement. Consider the density matrix of the system:

$$\rho = |\Psi\rangle \langle \Psi| = \sum_{ijkl} c_{ij} c_{kl}^* |\varphi_i\rangle \langle \varphi_k| \otimes |\chi_j\rangle \langle \chi_l|$$

The reduced density matrices ρ_A and ρ_B are given by tracing out the B and A part of the system respectively:

$$\begin{aligned} \rho_A &:= \text{Tr}_B \rho = \sum_m \langle \chi_m | \rho | \chi_m \rangle = \sum_{ijk} c_{ji} c_{ki}^* |\varphi_j\rangle \langle \varphi_k| \\ \rho_B &:= \text{Tr}_A \rho = \sum_m \langle \varphi_m | \rho | \varphi_m \rangle = \sum_{ijk} c_{ij} c_{ik}^* |\chi_j\rangle \langle \chi_k| \end{aligned}$$

Note, how they generally no longer describe pure states! The reduced density matrices therefore have a non-vanishing von Neumann entropy in general, which for an arbitrary density matrix is defined as:

$$S(\rho) := -\text{Tr} \rho \ln \rho = -\sum_{i=1}^D \rho_i \ln \rho_i \quad (28)$$

Here, ρ_i are the eigenvalues of the density matrix. For pure states of the overall system - and only those are considered in this work - the entanglement can be quantified by the entanglement entropy, which is the von Neumann entropy of a reduced density matrix. The von Neumann entropy has the following upper bound:

$$S(\rho) \leq \ln D$$

where D is the dimension of the Hilbert space. This can be seen by maximizing the r.h.s of (28) under the constraint $\sum_{i=1}^D \rho_i = 1$.

As an illustration, consider the already mentioned Bell-state:

$$\rho = \frac{1}{2} (|0\rangle \langle 0| \otimes |1\rangle \langle 1| - |0\rangle \langle 1| \otimes |1\rangle \langle 0| - |1\rangle \langle 0| \otimes |0\rangle \langle 1| + |1\rangle \langle 1| \otimes |0\rangle \langle 0|)$$

$$\rho_A = \frac{1}{2} (|0\rangle \langle 0| + |1\rangle \langle 1|) = \rho_B$$

From this follows, that the entanglement entropy of the Bell-state is $S = \log(2)$. This state is therefore maximally entangled.

When working with reduced density matrices, a particular choice of bases is convenient to work with. Although it is not possible to write a general state of a bipartite system as a product state, it is possible to choose orthonormal bases $\{|\tilde{\varphi}_i\rangle\}$ and $\{|\tilde{\chi}_i\rangle\}$, such that

$$|\Psi\rangle = \sum_{i=1}^D \lambda_i |\tilde{\varphi}_i\rangle |\tilde{\chi}_i\rangle, \quad \lambda_i \neq 0 \quad ,$$

where

$$D \leq \min \{ \dim \mathcal{H}_A, \dim \mathcal{H}_B \} \quad .$$

If additionally the state is normalised, the coefficients $\{\lambda_i\}$ fulfill:

$$\sum_i |\lambda_i|^2 = 1$$

This is the so-called Schmidt decomposition [8] and the λ 's are called Schmidt values. Its existence can be seen by singular value decomposition of the coefficient matrix[33]: $c_{ij} = \sum_k s_k U_{ik} V_{jk}^*$. Without loss of generality, we can assume that $\{|\varphi_i\rangle\}$ and $\{|\chi_i\rangle\}$ are ONBs:

$$|\Psi\rangle = \sum_k s_k \sum_i U_{ik} |\varphi_i\rangle \sum_j V_{jk}^* |\chi_j\rangle = \sum_k s_k |\tilde{\varphi}_k\rangle |\tilde{\chi}_k\rangle$$

Then $\{|\tilde{\varphi}_i\rangle\}$ and $\{|\tilde{\chi}_i\rangle\}$ are again ONBs since $U^\dagger U = \mathbb{1}_A$ and $V^\dagger V = \mathbb{1}_B$. Thus the Schmidt-values are just the non-zero singular values of c_{ij} ⁹:

$$\begin{aligned} \lambda_i &= s_i, \quad s_i \neq 0 \\ D &= \text{rank}(c_{ij}) \end{aligned}$$

In the Schmidt-basis, the reduced density matrices have the convenient property that they are diagonal:

$$\begin{aligned} \rho_A &= \sum_{i=1}^n |\lambda_i|^2 |\varphi_i\rangle \langle \varphi_i| \\ \rho_B &= \sum_{i=1}^n |\lambda_i|^2 |\chi_i\rangle \langle \chi_i| \end{aligned}$$

The entanglement entropy is therefore fully determined by the Schmidt-spectrum

$$S = S(\rho_A) = S(\rho_B) = - \sum_{i=1}^n |\lambda_i|^2 \ln |\lambda_i|^2$$

and is the same for both reduced density matrices

Note, that the rank of the reduced density matrices is equal to the rank D of the coefficient matrix. It is known, that the optimal way of approximating¹⁰ a given rank- D matrix A by a matrix B of rank $D' < D$ is by

⁹In the case of a “degenerate” Schmidt value (i.e. some of the λ_i are equal), the Schmidt-decomposition is not unique since the corresponding Schmidt vectors may be chosen as arbitrary linear combinations of each other, as long as the overall normalisation remains unchanged.

¹⁰In the sense that $\|A - B\|_F = \min$.

keeping only the D' largest singular values and setting the others to zero (see appendix). This is of great importance, since the size (and therefore the rank) of a matrix is limited for actual numerical calculations. In this sense, the optimal way of compressing the aforementioned state $|\Psi\rangle$ is by keeping only D' of its largest Schmidt values. Equivalently, only the D' largest eigenvalues of the reduced density matrix are retained. This way, the error introduced into the entanglement entropy becomes minimal. In this particular sense, an optimal way of compressing the state of a bipartite system has been found.

4.2 Matrix product states

One of the most important examples of tensor network states are the so-called matrix product states (MPS). A matrix product state (with periodic boundary conditions) is of the following form:

$$|\Psi\rangle = \sum_{\{s_i\}} \text{tr} \left(A_1^{[s_1]} A_2^{[s_2]} \cdots A_N^{[s_N]} \right) |s_1\rangle |s_2\rangle \cdots |s_N\rangle$$

The objects $A_k^{[s_k]}$ are matrices labelled by s_k which is called the “physical” index, since it labels the basis vector $|s_k\rangle$ in the local Hilbert space of site k . By fixing s_k , the object A_k becomes an order two tensor, that is, a matrix. The coefficients $c_{i_1, \dots, i_N}$ of the original many-body state (27) have therefore been expressed as the trace of a matrix product (hence the name):

$$c_{i_1, \dots, i_N} = \text{tr} \left(A_1^{[s_1]} A_2^{[s_2]} \cdots A_N^{[s_N]} \right)$$

Each entry of a given A_k is identified by the physical index s_k and the two indices of the matrix. Thus A_k is a tensor of order three, and using the TN notation we may write:

$$\left(A_k^{[s_k]} \right)_{ij} = \begin{array}{c} A_k \\ \text{---} i \text{---} \bullet \text{---} j \\ | \\ s_k \end{array}$$

$$c_{i_1, \dots, i_N} = \text{tr} \left(A_1^{[s_1]} A_2^{[s_2]} \cdots A_N^{[s_N]} \right) = \begin{array}{c} \boxed{A_1 \quad A_2 \quad \cdots \quad A_N} \\ \text{---} \bullet \text{---} \bullet \text{---} \cdots \text{---} \bullet \\ | \quad | \quad \quad \quad | \\ s_1 \quad s_2 \quad \quad \quad s_N \end{array}$$

The matrix indices i and j are called “bond” or “virtual” indices, and the range of possible values for the indices is called the bond dimension. In principle, any quantum state living in $\mathcal{H}_N$ can be written as an MPS, although the bond dimensions may become quite large. MPS are however most naturally suited for the description of one non-critical one dimensional quantum systems. It is very important to understand, that the set of all MPS of a

fixed bond dimension D forms a manifold which generally is *not* a subspace of the overall Hilbert space $\mathcal{H}_N$ [39, 16]. Just consider the sum of two MPS, $|\Psi_A\rangle$ and $|\Psi_B\rangle$ which both have bond dimension D :

$$\begin{aligned} |\Psi_A\rangle + |\Psi_B\rangle &= \\ &= \sum_{\{s_i\}} \left(\text{tr} [A_1^{[s_1]} \dots A_N^{[s_N]}] + \text{tr} [B_1^{[s_1]} \dots B_N^{[s_N]}] \right) |s_1\rangle \dots |s_N\rangle = \\ &= \sum_{\{s_i\}} \text{tr} [(A_1^{[s_1]} \oplus B_1^{[s_1]}) \dots (A_N^{[s_N]} \oplus B_N^{[s_N]})] |s_1\rangle \dots |s_N\rangle \end{aligned}$$

The resulting MPS $|\Psi\rangle$ is generally of bond dimension $D' = 2D$ and therefore no longer lives on the same manifold as $|\Psi_A\rangle$ and $|\Psi_B\rangle$.

Open boundary conditions may also be chosen:

$$c_{i_1, \dots, i_N} = A_1^{[s_1]} A_2^{[s_2]} \dots A_N^{[s_N]} = \begin{array}{c} A_1 \quad A_2 \quad \dots \quad A_N \\ | \quad | \quad \dots \quad | \end{array}$$

This will in fact be the most common case, since MPS with open boundary conditions can be gauged in a certain desirable way, soon to be discussed. Note that the tensors on the boundary, A_1 and A_N , are now of order two. Overlaps and expectation values can readily be computed from MPS:

$$\begin{aligned} \langle \Psi | \Phi \rangle &= \sum_{\{r_i\}} \sum_{\{s_i\}} c_{r_1, \dots, r_N}^* b_{s_1, \dots, s_N} \langle r_1 | s_1 \rangle \dots \langle r_N | s_N \rangle = \\ &= \sum_{\{s_i\}} c_{s_1, \dots, s_N}^* b_{s_1, \dots, s_N} = \begin{array}{c} |\Phi\rangle \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \\ | \\ \dots \\ \langle \Psi| \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \end{array} \end{aligned}$$

For a normalised state $|\Psi\rangle$ and a local observable O_k at site k , we have

$$\langle \Psi | O_k | \Psi \rangle = \begin{array}{c} \dots \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \\ | \\ \dots \text{---} \text{blue circle} \text{---} \text{brown circle} \text{---} \text{blue circle} \text{---} \\ | \\ \dots \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \end{array} O_k \dots$$

In general, N different MPS-tensors are required for the description of a system composed of N sites. If however the state the system exhibits a periodic structure, i.e. it is invariant under translations of p sites, only p different tensors are required for its description. Consider for instance an MPS with a $p = 3$ - site unit cell:

$$\dots \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \text{blue circle} \text{---} \dots$$

$C \quad A \quad B \quad C \quad A \quad B \quad C$

By exploiting periodicity, systems of arbitrary size can be described by a fixed number of tensors, thus providing a way to treat “infinite” systems¹¹. Such an infinite MPS will be called iMPS [40]. Consider the overlap of two one-site unit cell iMPS:

Diagram illustrating the inner product of two states, $\langle \Phi | \Psi \rangle$, using a quantum circuit. The circuit consists of two horizontal lines representing qubits. The top line starts with state $|\Psi\rangle$ and the bottom line with state $\langle \Phi|$. The circuit includes a series of controlled operations (gates) acting on the two qubits. The first operation is a controlled- U gate, followed by a controlled- S gate, which are enclosed in a red dashed box. The circuit continues with more controlled- U and controlled- S gates. The final output is the inner product $\langle \Phi | \Psi \rangle$.

It is in fact nothing but an infinite product of what will be called (mixed) MPS transfer matrices:

$$T = \text{blue circle with two horizontal lines} := \text{two blue circles connected by a vertical line, each with a horizontal line} \quad (29)$$

Thus the transfer matrix formalism from the preceding chapter immediately applies to overlaps of iMPS! We may define the overlap *per site* as the dominant eigenvalue of the MPS transfer matrix. If the bra and ket iMPSs are equal to each other, this overlap is just the squared norm per site. Local observables, correlation functions and the correlation length can in principle be obtained via the transfer matrix formalism. Note however, that grouping a bra and a ket tensor as in (29) should not be done in actual computations, since the MPS structure is concealed in this way and computations may get more expensive. It merely served to illustrate the connection to the 1D transfer matrix formalism.

The MPS representation of a certain state is not unique: It is always possible to insert resolutions of identity $UU^{-1} = \mathbb{1}$ between two MPS tensors:

[illegible]

This way, new MPS tensors can be defined without changing the state they describe:

The diagram shows an equality between two circuit components. On the left is a blue circle labeled A' with a horizontal line passing through its center and a vertical line extending downwards from its bottom. On the right is a sequence of three components connected horizontally: a small brown circle labeled U^{-1} , a blue circle labeled A , and another small brown circle labeled U . The blue circle A has a horizontal line passing through its center and a vertical line extending downwards from its bottom. The brown circles U^{-1} and U have horizontal lines passing through their centers. An equals sign is placed between the two sides of the diagram.

Note however, that all the U 's must be equal to each other so as to preserve translation invariance. This gauge freedom can be exploited to cast an MPS into a form that may be convenient for certain purposes. An example is the

¹¹Actually, there is some subtlety involved here, since the construction of the Hilbert space for infinitely large systems is problematic. The mathematically correct way of treating infinite systems is by employing C*-algebras. We won't elaborate on this any further and merely note, that the iMPS ansatz works in practice.

so-called canonical-form of an MPS [26]. Roughly speaking, it is a Schmidt-decomposition of the MPS at each bond and thus gives immediate access to the entanglement properties of the state. Both finite and infinite MPS can be brought into canonical form, however only the latter shall be considered here. Consider an iMPS which has a single-site unit cell holding the MPS tensor A . Its canonical form is determined by a different order 3 tensor Γ and a diagonal matrix λ holding the Schmidt values:

$$\dots \text{---} \overset{A}{\bullet} \text{---} \overset{A}{\bullet} \text{---} \overset{A}{\bullet} \text{---} \dots = \dots \text{---} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \text{---} \dots$$

$$\lambda_i \delta_{ij} = \text{---} \overset{\lambda}{\bullet} \text{---} i \text{---} j$$

At each bond, the iMPS can be decomposed into a left and right part:

$$|\Psi\rangle = \sum_i \lambda_i |\Psi_i^L\rangle |\Psi_i^R\rangle = \dots \text{---} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \text{---} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \text{---} \dots$$

The left and right Schmidt-vectors are labelled by their bond indices:

$$\dots \text{---} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \text{---} i = |\Psi_i^L\rangle, \quad |\Psi_i^R\rangle = i \text{---} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \text{---} \dots$$

Since they are orthonormal, we have:

$$\dots \begin{array}{c} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \text{---} i \\ \text{---} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \text{---} j \\ \text{---} \overset{\lambda}{\bullet} \overset{\Gamma^*}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma^*}{\bullet} \text{---} j \end{array} = \left(\begin{array}{c} i \\ j \end{array} = \delta_{ij} = \begin{array}{c} i \\ j \end{array} \right) = \begin{array}{c} i \text{---} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \text{---} \\ j \text{---} \overset{\Gamma^*}{\bullet} \overset{\lambda}{\bullet} \overset{\Gamma^*}{\bullet} \overset{\lambda}{\bullet} \text{---} \end{array} \dots$$

The identity matrix is therefore a left (right) eigenvector of the transfer matrix corresponding to the MPS in the left (right) gauge:

$$\overset{A_L}{\bullet} := \text{---} \overset{\lambda}{\bullet} \overset{\Gamma}{\bullet} \text{---} \quad \overset{A_R}{\bullet} := \text{---} \overset{\Gamma}{\bullet} \overset{\lambda}{\bullet} \text{---}$$

Note, that it was implicitly assumed, that the boundary conditions are not periodic. This is necessary, since cutting a bond of an (infinite) MPS with periodic boundary conditions does *not* produce a bipartition of the system (see fig. 19). Hence, doing the Schmidt decomposition is not possible in this case.




Figure 19: Cutting a bond of an MPS with OBC produces a bipartition (i), while this is no longer the true in the case of PBC (ii).

The canonical form allows for a very efficient way of computing local observables. Consider for instance a single-site observable O :

$$\langle \hat{O} \rangle = \frac{\langle \Psi_D | \text{[MPS with } O \text{]} | \Psi_U \rangle}{\langle \Psi_D | \text{[MPS]} | \Psi_U \rangle} = \frac{\text{[Contracted MPS with } O \text{]}(u_L^\dagger, u_R)}{\text{[Contracted MPS]}(u_L^\dagger, u_R)}$$

The diagram shows the calculation of the expectation value of a single-site observable O (represented by a red circle) in a matrix product state $|\Psi\rangle$. The state is represented as a grid of blue circles (sites) with horizontal and vertical bonds. The numerator shows the state with the observable O inserted at a central site. The denominator shows the state without the observable. The right side of the equation shows the contraction of the indices $u_L^\dagger$ and u_R on the left and right boundaries, respectively, resulting in a single-site contraction of the observable O with the MPS tensors.

The last equality holds because due to orthonormality, everything to the right and left collapses to identity.

Another nice feature of the canonical form is that it provides a straightforward way of approximating an MPS of bond dimension D by another MPS with lower bond dimension $D' < D$ [26]: As outlined in section 4.1, the state is *optimally* approximated by only keeping the D' largest Schmidt values at a particular bond, thus effectively reducing the bond dimension. For an iMPS with single site unit cell, this strategy may be employed *at each bond simultaneously*. While the rigorous statements about optimally approximating a state made in 4.1 may no longer apply, this works very well in practice, as was also noted in [26].

Up until now, it hasn't been explained how to actually bring an iMPS into canonical form. One way of doing this is described in [26], and a slightly modified version of this procedure shall be outlined in the following¹². Consider an arbitrary iMPS $|\Psi\rangle$ with a single-site unit cell. At first, we find the left and right dominant eigenvectors (actually they are matrices) of the MPS transfer matrix:

$$v_L^\dagger \text{[MPS tensor]} = \eta \text{[MPS tensor]} v_L^\dagger \quad \text{[MPS tensor]} v_R = \eta \text{[MPS tensor]} v_R$$

The diagram shows the equations for the left and right dominant eigenvectors of the MPS transfer matrix. The left equation shows $v_L^\dagger$ (a pink circle) contracted with the left index of a blue MPS tensor, resulting in η times the tensor contracted with $v_L^\dagger$ on the right index. The right equation shows v_R (a green circle) contracted with the right index of a blue MPS tensor, resulting in η times the tensor contracted with v_R on the left index.

¹²There are also methods for gauging an MPS based on the QR decomposition of matrices [17]. These algorithms may have the advantage of being more accurate in the case of small Schmidt values.

(The explanation of how to do this in an efficient way is postponed.) As already mentioned, η is the squared norm per site, i.e. $\eta = 1$ if $|\Psi\rangle$ is normalized. Next, we exploit that $v_L^\dagger$ and v_R must necessarily be of the form

$$\left(\begin{array}{c} \bullet \\ \bullet \end{array} \right) = \begin{array}{c} Y \\ \bullet \\ Y^\dagger \end{array} \quad \left(\begin{array}{c} \bullet \\ \bullet \end{array} \right) v_R = \begin{array}{c} X \\ \bullet \\ X^\dagger \end{array}$$

due to the form of the scalar product¹³. Thus an eigendecomposition of $v_L^\dagger$ and v_R exists, which we can in turn use to find X and Y :

$$\begin{aligned} v_R &= W_x D_x W_x^\dagger, & X &= W_x D_x^{1/2}, & X^{-1} &= D_x^{-1/2} W_x^\dagger \\ v_L^\dagger &= W_y D_y W_y^\dagger, & Y &= D_y^{1/2} W_y, & Y^{-1} &= W_y^\dagger D_y^{-1/2} \end{aligned}$$

We then exploit the gauge freedom of the iMPS and insert

$$\mathbb{1} = Y^{-1} Y X X^{-1}$$

between the bonds:

$$\dots \text{---} \overset{A}{\bullet} \text{---} \overset{A}{\bullet} \text{---} \dots = \dots \text{---} \overset{Y^{-1}}{\bullet} \overset{A}{\bullet} \overset{X^{-1}}{\bullet} \overset{X}{\bullet} \overset{Y}{\bullet} \overset{Y^{-1}}{\bullet} \overset{A}{\bullet} \overset{X^{-1}}{\bullet} \text{---} \dots$$

After doing a singular value decomposition of YX

$$\text{---} \overset{X}{\bullet} \overset{Y}{\bullet} \text{---} = \text{---} \overset{U}{\bullet} \overset{\lambda}{\bullet} \overset{V^\dagger}{\bullet} \text{---}$$

the iMPS reads:

$$\dots \text{---} \overset{A}{\bullet} \text{---} \overset{A}{\bullet} \text{---} \dots = \dots \text{---} \overset{V^\dagger}{\bullet} \overset{Y^{-1}}{\bullet} \overset{A}{\bullet} \overset{X^{-1}}{\bullet} \overset{U}{\bullet} \overset{\lambda}{\bullet} \overset{V^\dagger}{\bullet} \overset{Y^{-1}}{\bullet} \overset{A}{\bullet} \overset{X^{-1}}{\bullet} \text{---} \dots$$

The singular values turn out to be nothing but the Schmidt values and Γ is given by:

$$\text{---} \overset{\Gamma}{\bullet} \text{---} = \text{---} \overset{V^\dagger}{\bullet} \overset{Y^{-1}}{\bullet} \overset{A}{\bullet} \overset{X^{-1}}{\bullet} \text{---}$$

To see this, we prove that indeed a left canonical form can be constructed

¹³Due to finite numerical precision this may be slightly violated in practice. One can however stabilise the algorithm via $v_R \mapsto 1/2(v_R + v_R^\dagger)$.

from λ and Γ as in (4.2)

where $\eta = 1$ if the MPS is normalised. The proof for the right orthonormal case is analogous. Note, that the matrix multiplications are always understood to be done in clockwise order, e.g.:

4.3 Matrix product operators


The idea of decomposing a large tensor into a product of smaller ones can also be applied to operators. Provided a basis of the local Hilbert spaces has been chosen, a general operator acting on states living in the many-body Hilbert space may be written as:

$$T = \sum a_{r_1, \dots, r_N}^{s_1, \dots, s_N} |s_1\rangle \langle r_1| \otimes |s_2\rangle \langle r_2| \otimes \dots \otimes |s_N\rangle \langle r_N|$$


As in the case of states, such an operator is again generally represented by an exponentially large number of coefficients $a_{r_1, \dots, r_N}^{s_1, \dots, s_N}$. Consider for instance an operator being applied to a MPS:

The resulting state is now no longer a MPS. If we want to work with MPS, we must come up with operators which preserve the MPS structure. In order to find such operators, we note that it is always possible to decompose an operator (which is nothing but a very large tensor once a local basis has been chosen) into network of smaller tensors:

In general, the bond dimension χ of such a matrix product *operator* (MPO) exponentially grows with system size. However, if the operator is a sum of local terms, the bond dimension may be a constant [39]. Note that the physical dimensions of the MPO must be the same as the physical dimension of the MPS. Furthermore, the bond dimension D of an MPS increases to $D\chi$ after an MPO with bond dimension χ acts on it:



However, the MPS structure is preserved regardless of the bond dimension of the MPO. As for MPS, MPOs may be defined for infinite systems by exploiting translation invariance. Similar to the MPS case, one can for instance write down a three-site “unit cell” infinite MPO:



Such “iMPOs” are of relevance to this work, since they can represent infinitely large one-dimensional transfer matrices of homogeneous classical systems [39], as will be elaborated on in detail in the remainder of this work. If the underlying classical system only has finite-range interactions, the resulting transfer matrices can be represented exactly by an MPO which has a finite bond dimension.

4.4 The relationship between classical and quantum systems

There is a deep connection between quantum and classical systems. In 1976 Suzuki showed [35], that a general quantum system in d dimensions can be mapped to a certain classical one in $d + 1$ dimensions. Even more important, it turns out that the ground state (i.e zero temperature) energy of a quantum system equals the free energy of a higher dimensional classical system in the thermodynamic limit. Thus quantum phase transitions which are non-analyticities in the ground state energy, can be studied by looking at phase transitions of a higher dimensional classical system [32], [41]. For the purpose of this work, the existence of such a correspondence is interesting by itself.

Consider a quantum system with Hamiltonian $\hat{H}$ which has a discrete energy spectrum $\{E_n\}$:

$$\hat{H} |\Psi_n\rangle = E_n |\Psi_n\rangle, \quad n = 0, 1, 2, \dots$$

Its ground state $|\Psi_0\rangle$ (which is assumed to be non-degenerate, for the sake of simplicity) can be obtained by evolving an arbitrary state $|\Psi\rangle$ with $\langle\Psi|\Psi_0\rangle \neq$

0 in imaginary time $t \mapsto it$:

$$\lim_{t \rightarrow \infty} \frac{e^{-t\hat{H}} |\Psi\rangle}{\|e^{-t\hat{H}} |\Psi\rangle\|} = |\Psi_0\rangle \quad (30)$$

This is implied by the asymptotic behaviour of the imaginary time evolution operator: Since

$$e^{-t\hat{H}} = \sum_n e^{-tE_n} |\Psi_n\rangle \langle \Psi_n| \quad \Rightarrow \quad \text{tr } e^{-t\hat{H}} = \sum_n e^{-tE_n}$$

and by assuming that $E_0 < E_n \leq E_{n+1}$, it follows that

$$e^{-t\hat{H}} \rightarrow e^{-tE_0} |\Psi_0\rangle \langle \Psi_0| \quad (t \rightarrow \infty)$$

which in turn implies (30).

For the trace of $e^{-t\hat{H}}$ we have:

$$\text{tr } e^{-t\hat{H}} \rightarrow e^{-tE_0} \quad (t \rightarrow \infty) \quad (31)$$

However, the l.h.s. of the previous expression is nothing but the partition function of the system if t is interpreted as inverse temperature.

$$Z_{\text{q.}}(\beta) = \text{tr } e^{-\beta\hat{H}}$$

Since (31) is equivalent to

$$\lim_{t \rightarrow \infty} \frac{\text{tr } e^{-t\hat{H}}}{e^{-tE_0}} = 1 \quad ,$$

the relationship between ground state energy and zero temperature limit of the free energy is obtained:

$$E_0 = - \lim_{\beta \rightarrow \infty} \frac{1}{\beta} \ln \text{tr } e^{-\beta\hat{H}} = - \lim_{\beta \rightarrow \infty} \frac{1}{\beta} \ln Z_{\text{q.}}(\beta) \quad (32)$$

In the following it shall be demonstrated, that E_0 is equal to the free energy of a classical system at a *finite* $\beta_{\text{cl.}} < \infty$, where the “classical” inverse temperature $\beta_{\text{cl.}}$ is yet to be specified. The discussion is based on the original work of Suzuki [35].

It is assumed, that $\hat{H}$ is a sum of two non-commuting parts:

$$\hat{H} = \hat{H}_1 + \hat{H}_2$$

By using Trotter’s theorem, $e^{-\beta\hat{H}}$ can be rewritten as:

$$e^{-\beta\hat{H}} = \exp\left(-\frac{\beta}{n}\hat{H}_1 - \frac{\beta}{n}\hat{H}_2\right)^n = \lim_{n \rightarrow \infty} \left(\exp\left(-\frac{\beta}{n}\hat{H}_1\right) \exp\left(-\frac{\beta}{n}\hat{H}_2\right)\right)^n$$

The free energy of the quantum system then reads:

$$-\frac{1}{\beta} \ln \text{tr} e^{-\beta \hat{H}} = - \lim_{n \rightarrow \infty} \frac{1}{\beta} \ln \left[\sum_{\varphi_1} \langle \varphi_1 | \left[\exp \left(-\frac{\beta}{n} \hat{H}_1 \right) \exp \left(-\frac{\beta}{n} \hat{H}_2 \right) \right]^n | \varphi_1 \rangle \right]$$

Eventually we want to take the limit $\beta \rightarrow \infty$. This can be done simultaneously with taking $n \rightarrow \infty$, by fixing the ratio β/n . Thereby a new parameter $g = \beta/n$ with dimension $[\text{energy}]^{-1}$ is introduced. As shall soon become clear, it is convenient to set g equal to a characteristic inverse coupling strength of $\hat{H}$. The zero temperature limit for a certain value of g is then defined as follows:

$$\begin{aligned} - \lim_{\beta \rightarrow \infty} \left| \frac{1}{\beta} \ln \text{tr} e^{-\beta \hat{H}} \right| &:= - \frac{1}{g} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \text{tr} e^{-ng \hat{H}} = \\ &= - \frac{1}{g} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \left[\sum_{\varphi_1} \langle \varphi_1 | \left(e^{-g \hat{H}_1} e^{-g \hat{H}_2} \right)^n | \varphi_1 \rangle \right] \end{aligned}$$

The expression in brackets is equal to the partition function of the quantum system at inverse temperature $\beta_n = ng$ up to a Trotter error $\mathcal{O}(\frac{1}{n^2})$:

$$\sum_{\varphi_1} \langle \varphi_1 | \left(e^{-g \hat{H}_1} e^{-g \hat{H}_2} \right)^n | \varphi_1 \rangle = Z_q(g, \beta_n) + \mathcal{O}(\frac{1}{n^2})$$

Inserting resolutions of the identity,

$$\mathbb{1} = \sum_{\varphi} |\varphi\rangle \langle \varphi|$$

for which $\hat{H}_1$ is diagonal, we obtain:

$$\begin{aligned} Z_q(g, \beta_n) + \mathcal{O}(\frac{1}{n^2}) &= \sum_{\varphi_1} \langle \varphi_1 | \left(e^{-g \hat{H}_1} e^{-g \hat{H}_2} \right)^n | \varphi_1 \rangle = \\ &= \sum_{\{\varphi_k\}} \langle \varphi_1 | e^{-g \hat{H}_1} e^{-g \hat{H}_2} | \varphi_2 \rangle \langle \varphi_2 | e^{-g \hat{H}_1} e^{-g \hat{H}_2} | \varphi_3 \rangle \dots \langle \varphi_n | e^{-g \hat{H}_1} e^{-g \hat{H}_2} | \varphi_1 \rangle = \\ &= \sum_{\{\varphi_k\}} \prod_{k=1}^n e^{-g H_1(\varphi_k)} \langle \varphi_k | e^{-g \hat{H}_2} | \varphi_{k+1} \rangle = \\ &= \sum_{\{\varphi_k\}} e^{H_{\text{cl.}}^{(n)}(\varphi, g)} =: Z_{\text{cl.}}(n, g) \end{aligned}$$

Where we set $\varphi_{n+1} = \varphi_1$ and introduced the *effective* classical Hamiltonian¹⁴:

$$H_{\text{cl.}}^{(n)}(\varphi, g) := \sum_{j=1}^n \left[g H_1(\varphi_j) - \ln \langle \varphi_j | e^{-g \hat{H}_2} | \varphi_{j+1} \rangle \right]$$

¹⁴It should however be noted, that there is no guarantee that this effective Hamiltonian is real in general.

This is the main result of the quantum to classical mapping: The quantum system with Hamiltonian $\hat{H}$ at zero temperature, can effectively be described by a classical one with Hamiltonian $H_{\text{cl.}}^{(n)}$. The mapping becomes exact (i.e. the Trotter error vanishes) as $n \rightarrow \infty$. The parameter $g < \infty$ can be interpreted as the temperature of the classical system:

$$\beta_{\text{cl.}} := g$$

Note, that a new spatial dimension has been introduced: If the original quantum system lives on a d -dimensional lattice, then the underlying lattice of the classical system is composed of n copies of the original lattice stacked on top of each other along the “ $d + 1$ -direction”. Furthermore, the classical system obeys periodic boundary conditions in the newly introduced dimension. For instance, a one dimensional quantum chain with open (periodic) boundary conditions, corresponds to a classical system living on the surface of a cylinder (torus). The state of layer k is given by φ_k and two neighbouring layers are coupled by the term

$$\ln \langle \varphi_j | e^{-g\hat{H}_2} | \varphi_{j+1} \rangle \quad . \quad (33)$$

Note, that if $\hat{H}_2$ was also diagonal with respect to $\{|\varphi\rangle\}$ (i.e. $[H_1, H_2] = 0$), there wouldn’t be any coupling between neighbouring layers since (33) would collapse to $-gH_2(\varphi_j)\delta(\varphi_j, \varphi_{j+1})$. In this case, the classical system would effectively only be d -dimensional, since its partition function would be a product of n identical factors.

Together with (32), the ground state energy $E_0(g) = E_0(\beta_{\text{cl.}})$ turns out to be equal to the free energy per layer of the classical system:

$$E_0(g) = \frac{1}{g} \lim_{n \rightarrow \infty} \ln (Z_{\text{cl.}}(n, g))^{1/n}$$

This in turn allows for an interpretation of the zero temperature limit in terms of the thermodynamic limit of a higher dimensional system at finite temperature.

Since it is always possible to find an ONB which diagonalises the *entire* Hamiltonian, the previous considerations may seem somewhat artificial. However, the particular choice of the ONB which only diagonalises one part of $\hat{H}$ introduces couplings between the hyperplanes in the $d + 1$ -direction thus making the correspondence to a higher dimensional classical system visible.

The procedure of mapping a quantum to higher dimensional classical system can also be turned upside down, as shall be illustrated on a two-dimensional classical system on a square lattice: Having found its formulation in terms of a tensor network, we may regard a single layer of local tensors (which is nothing but a one-dimensional transfer matrix) as an MPO,




Figure 20: The partition function $Z_{\text{cl.}}$ of a two-dimensional classical system with local interactions, may be expressed as a product of MPOs, M .

see fig. 20. Interpreting this MPO as an infinitesimal step in imaginary time evolution, the dominant eigenstate of this MPO may then be regarded as the ground state of a quantum system. Describing this ground state as an MPS is the main idea behind boundary-MPS methods.

At the beginning of this discussion, entanglement was mentioned. Now that the quantum-classical correspondence has been established, one could ask how entanglement in the quantum system translates into properties of its higher dimensional classical counterpart. For a given classical system, boundary-MPS methods give access to the entanglement properties of the corresponding (in the sense just described) quantum system and may thus provide us with some insights.

5 Algorithms for contracting tensor networks


Having found a TN representation of a classical partition function, the question of how to actually compute it, i.e. contract the tensor network, remains. Furthermore, we would like to compute environments in order to obtain expectation values of local observables. This chapter outlines ways in which to achieve these goals in a computationally efficient manner.

5.1 Computational cost

Computational cost is measured in the number of floating point operations which are required to carry out a given calculation. Consider the product of matrices $A \in \mathbb{C}^{l \times m}$ and $B \in \mathbb{C}^{m \times n}$. The result is an l by n matrix, each of whose entries takes m multiplications and $m - 1$ additions to compute¹⁵. Since the number of multiplications and additions is $\mathcal{O}(m)$ in both cases, we may treat them on equal footing. Thus the cost of the operation

$$AB = C$$

is $\mathcal{O}(lmn)$. This result also applies to the case of *pairwise contractions* of tensors (i.e. summing over all indices connecting a given pair of tensors), which are nothing but products of matrices when the indices are grouped together. As an example, consider two tensors all of whose indices have the same bond dimension D :




By grouping the indices together, the pairwise contraction can be regarded as the product of a D^3 by D^2 and a D^2 by D^4 matrix, which has a cost of $\mathcal{O}(D^9)$.


Given a finite sized TN, one can of course contract it by the brute force approach of summing over all indices. This procedure is generally referred to as contracting a TN *exactly*. The computational cost of doing so heavily depends on the exact order in which the summation is performed. First, it turns out that pairwise contractions (i.e. summing over all indices connecting a given set of two tensors before doing any other summations) of tensors are generally cheaper [29]. Second, the order in which these pairwise contractions are performed has a great influence on computational cost. As an example, consider the following network which appears in the eigenvalue

¹⁵Actually $2m$ and $2(m - 1)$ operations respectively, since we are dealing with complex numbers and therefore need to treat imaginary and real parts separately


equation of an MPS transfer matrix:



d and D is the physical and virtual bond dimension, respectively. There are three choices for the order in which pairwise contractions may be performed, only two of which differ in computational cost, however. The contraction sequence



has a cost of $\mathcal{O}(D^4 d)$, while



merely scales as $\mathcal{O}(D^3 d)$. Note that the first contraction scheme corresponds to merging the bra and ket tensors of the MPS into a matrix first, while the second approach exploits the tensor network structure to gain a computational advantage.

5.2 Finding dominant eigenvalues

The problem of finding dominant eigenvectors frequently occurs in MPS algorithms. The standard approach are iterative eigenvalue solvers which are based on the *power method*. Its basic idea is to apply the matrix to a random start vector which is then normalized and again chosen as a new start vector and so forth. The procedure is iterated until certain convergence criterion is satisfied, e.g. the corresponding dominant eigenvalue has converged to within a predefined precision. The power method suffers from several drawbacks and is not used for finding eigenvalues of finite sized matrices. Instead, the Arnoldi iteration (which also has the nice property that it allows for finding the first n largest eigenvalues) is used for non-hermitian eigenvalue problems. We will not go into the details, but instead refer to [12]. However, the most basic form of the power method shall be treated in full detail, since the idea behind will be generalised to finding the dominant “eigen-MPS” of a given MPO.

Let $A \in \mathbb{C}^{n \times n}$ have a real and positive dominant eigenvalue $\lambda_0 = |\lambda_0|$, with $v \in \mathbb{C}^n$ being a corresponding normalised eigenvector. Furthermore,

let $v_0 \in \mathbb{C}^n$ be an arbitrary “start vector”, with $\langle v | v_0 \rangle = \alpha \neq 0$. Then

$$\lim_{N \rightarrow \infty} \frac{A^N v_0}{\|A^N v_0\|_2} = \frac{\alpha}{|\alpha|} v$$

Proof. The start vector v_0 , can be written as $v_0 = \alpha v + v_\perp$ $\alpha \in \mathbb{C}$. Therefore

$$\left(\frac{1}{\lambda_0} A\right)^N v_0 = \alpha v + \tilde{A}^N v_\perp, \quad \tilde{A} := \frac{1}{\lambda_0} A$$

Note that, since $\langle v | \tilde{A}^N v_\perp \rangle = 0$:

$$\Rightarrow \quad \left\| \tilde{A}^N v_0 \right\|_2 \geq |\alpha|$$

Since $\lambda_0 = |\lambda_0|$:

$$\frac{A^N v_0}{\|A^N v_0\|_2} = \frac{\tilde{A}^N v_0}{\|\tilde{A}^N v_0\|_2}$$

Putting it all together we obtain:

$$\begin{aligned} & \left\| \frac{A^N v_0}{\|A^N v_0\|_2} - \frac{\alpha}{|\alpha|} v \right\| = \frac{1}{\|\tilde{A}^N v_0\|_2} \left\| \tilde{A}^N v_0 - \frac{\alpha}{|\alpha|} \|\tilde{A}^N v_0\|_2 v \right\| \\ & \leq \frac{1}{|\alpha|} \left\| \tilde{A}^N v_0 - \frac{\alpha}{|\alpha|} \|\tilde{A}^N v_0\|_2 v \right\| = \frac{1}{|\alpha|} \left\| \left(\alpha - \frac{\alpha}{|\alpha|} \|\alpha v + \tilde{A}^N v_\perp\|_2 \right) v + \tilde{A}^N v_\perp \right\| \\ & \leq \left\| \left(1 - \frac{1}{|\alpha|} \|\alpha v + \tilde{A}^N v_\perp\|_2 \right) v \right\| + \frac{1}{|\alpha|} \left\| \tilde{A}^N v_\perp \right\| \\ & = \left| 1 - \frac{1}{|\alpha|} \|\alpha v + \tilde{A}^N v_\perp\|_2 \right| + \frac{1}{|\alpha|} \left\| \tilde{A}^N v_\perp \right\| \\ & \rightarrow 0 \quad (N \rightarrow \infty) \end{aligned}$$

where $\|\cdot\|$ is some arbitrary norm on $\mathbb{C}^n$. □

The approximate dominant eigenvalue $\lambda_0^{[N]}$ after N iterations can be calculated by

$$\lambda_0^{[N]} = \frac{\langle u | A^{N+1} | v_0 \rangle}{\langle u | A^N | v_0 \rangle}.$$

Note, that the speed of convergence depends on $|\lambda_1/\lambda_0|$ ¹⁶ : the closer the ratio is to unity, the slower the convergence. This result has an immediate consequence for the simulation of many-body systems, since the correlation length is determined by the ratio of the two largest eigenvalues of the transfer matrix. Thus, as the correlation length diverges, so does the time needed for the power method to converge.

¹⁶More generally, it depends on the spacing of all eigenvalues.




Figure 21: Two dimensional TN composed of layers of equal MPOs. One such layer forms what will also be called the one dimensional transfer matrix of the TN.

5.3 Boundary MPS methods: iTEBD

Infinite, one dimensional translation invariant tensor networks can be contracted exactly by finding the dominant eigenvalue of their transfer matrix, see section 3.1. Boundary MPS methods aim to generalise this idea to the case of two dimensional tensor networks. The remainder of this section closely follows [26]. Without loss of generality, we may assume that two dimensional TN we are considering forms a square lattice¹⁷. The TN can then be regarded as nothing but infinitely many iMPOs stacked on top of each other, see fig. 21. Assuming that all tensors are equal (or at least repeat themselves according to some periodic pattern), the iMPO can be regarded as infinitely large transfer matrix M . Since only one MPO tensor is needed to fully describe this transfer matrix, we can treat an infinitely large system with only finite computational resources. In complete analogy to the $d = 1$ case, the fully contracted TN is then given by

$$\lim_{N \rightarrow \infty} \text{tTr } T^N = \lim_{l \rightarrow \infty} \langle \Phi_D | M^l | \Phi_U \rangle$$

where $\langle \Phi_D |$ and $|\Phi_U \rangle$ are some infinitely large boundary vectors. The idea is now to express boundary vectors as iMPS and find a “dominant MPS” $|\Psi_U \rangle$ and a corresponding dominant eigenvalue θ of M .

$$M |\Psi_U \rangle = \theta |\Psi_U \rangle \quad (34)$$

Note, that θ may diverge as the number of sites m of the MPS goes to infinity. Instead, we may consider the dominant eigenvalue per site, κ :

$$\theta = \kappa^m$$

Thus we can express the fully contracted network in terms of κ :

$$\lim_{N \rightarrow \infty} \text{tTr } T^N = \lim_{l, m \rightarrow \infty} \kappa^{lm}$$

¹⁷If it doesn't, we may in principle obtain a square lattice description by appropriately grouping tensors together, etc.

For a TN describing the partition function of a classical system, κ is just the partition function per site. Since θ is given by

$$\theta = \|M|\Psi_0\rangle\|, \quad \| |\Psi_0\rangle \| = 1$$

the dominant eigenvalue ω is equal to the norm per site of $M|\Psi_U\rangle$, i.e. the square root of the dominant eigenvalue of the MPS TM given by $M|\Psi_U\rangle$ (see section 4.2). There is however one big problem: The application of an MPO to an MPS $|\Psi\rangle$ yields an MPS $|\Psi'\rangle$ with a higher bond dimension. Strictly speaking, the eigenvalue equation (34) is therefore problematic, since the MPSs $|\Psi_U\rangle$ and $T_{\text{MPO}}|\Psi_U\rangle$ live on different MPS-manifolds. What we are actually looking for is some map

$$\varphi(|\Psi\rangle) = |\Psi'\rangle$$

which lets the MPO act onto the MPS and then projects it back onto the manifold of MPS of bond dimension D_{max} . The dominant MPS is then given by the fixed point (assuming that it exists) of this map:

$$\varphi(|\Psi_U\rangle) = \theta|\Psi_U\rangle$$

The iTEBD algorithm by Orùs and Vidal [26] does exactly that: It is a generalisation of the power method to the case of iMPS. The MPO is successively applied to a random start MPS until a certain convergence criterion is met. To prevent the bond dimension of the boundary MPS from growing beyond D_{max} , the MPS is *truncated* after each step, by bringing it into canonical form and discarding the smallest Schmidt values. As already described in section 4.2, this is a good way of approximating an MPS by another MPS with a lower bond dimension. The higher the maximum bond dimension D_{max} is chosen, the more Schmidt values can be retained, which increases the accuracy of the algorithm.

The iTEBD method can be summarised by the following steps:

1. Apply the MPO to get a new MPS $|\Psi'\rangle$ with bond dimension D_{χ} . For an MPS that already is in canonical form, this can be done in two steps:

This results in an MPS described by Γ' and λ' which is however no longer in canonical form.

2. Bring the MPS into canonical form. As a by-product, the dominant eigenvalue η of the MPS transfer matrix and thus $\kappa = \sqrt{\eta}$ is obtained.
3. Truncate to D_{max} .




Figure 22: Number of iterations needed for convergence vs. correlation length ξ in the disordered phase of the 1NN model ($-5 \leq \mu \leq 1.2$) at a maximum bond dimension of $D_{\max} = 40$.

4. Repeat until convergence.

In practice, Schmidt values below a certain threshold, e.g. 10^{-12} are truncated regardless of whether $D' > D_{\max}$ or not, both for stability and speed. The algorithm terminates when both all Schmidt values $\{\lambda_i\}$ and κ have converged to within a certain precision. iTEBD is essentially a power method, and we therefore expect a similar convergence behaviour. Numerically it is observed, that the number of iterations needed for convergence scales linearly with the system's correlation length, see fig. 22.

If one is only interested in κ , the “up” fixed point MPS $|\Psi_U\rangle$ is actually of no further use. For the calculation of environments however, we also need the “down” fixed point $\langle\Psi_D|$:

$$\begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \end{array} \approx \theta \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \end{array} \quad \begin{array}{c} A_U A_U A_U A_U \\ \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \end{array} \approx \theta \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \\ \text{---} \text{---} \text{---} \text{---} \end{array}$$

where the Schmidt values have been absorbed into the Γ s for convenience, e.g. according to what is called the *central gauge*:

$$\text{---} \text{---} \text{---} \text{---} := \text{---} \text{---} \text{---} \text{---}$$

For some single-site observable $\hat{O}$ we then have:

$$\langle \hat{O} \rangle = \frac{\langle \Psi_D | \begin{array}{c} \text{Diagram: A 3x4 grid of sites. The top row has blue circles with right-pointing semi-circles. The middle row has cyan circles with left-pointing semi-circles. The bottom row has blue circles with right-pointing semi-circles. A red circle is at the second column, second row. } \end{array} | \Psi_U \rangle}{\langle \Psi_D | \begin{array}{c} \text{Diagram: A 3x4 grid of sites. The top row has blue circles with right-pointing semi-circles. The middle row has cyan circles with left-pointing semi-circles. The bottom row has blue circles with right-pointing semi-circles. } \end{array} | \Psi_U \rangle} = \frac{\begin{array}{c} \text{Diagram: A single site with a red circle in the center, connected to four green circles (top, bottom, left, right). The left green circle is labeled } u_L^\dagger \text{ and the right green circle is labeled } u_R. \end{array}}{\begin{array}{c} \text{Diagram: A single site with a cyan circle in the center, connected to four blue circles (top, bottom, left, right). The left blue circle is labeled } u_L^\dagger \text{ and the right blue circle is labeled } u_R. \end{array}}$$

where $u_L^\dagger$ and u_R are the left and right dominant eigenvectors of the zero dimensional transfer matrix given by $\langle \Psi_D | M | \Psi_U \rangle$:

$$\begin{array}{c} \text{Diagram: A single site with a red circle in the center, connected to four green circles (top, bottom, left, right). The left green circle is labeled } u_L^\dagger. \end{array} = \omega_0 \begin{array}{c} \text{Diagram: A single green circle with a left-pointing semi-circle. The left semi-circle is labeled } u_L^\dagger. \end{array} \quad \begin{array}{c} \text{Diagram: A single site with a red circle in the center, connected to four green circles (top, bottom, left, right). The right green circle is labeled } u_R. \end{array} = \omega_0 \begin{array}{c} \text{Diagram: A single green circle with a right-pointing semi-circle. The right semi-circle is labeled } u_R. \end{array}$$

The problem has thus been reduced to the one dimensional case with transfer matrix $T = \langle \Psi_D | M | \Psi_U \rangle$. It so happens, that the dominant eigenvalue ω_0 of T is actually the partition function per site, i.e. $\omega_0 = \kappa$, and the correlation length of the system can immediately be obtained by also computing the second largest eigenvalue ω_1 . Note however, that T should not be confused with a transfer matrix for a one dimensional classical system. Therefore, the statement about the absence of phase transitions made in section 3.1 does not apply. However, MPS are also ill-suited for describing critical systems: MPS with finite bond dimension cannot handle polynomially decaying correlations [27]. It can however be shown, that MPS with a fixed bond dimension can approximate the ground state of a gapped local¹⁸ Hamiltonian efficiently [39].

So far, nothing has been said about uniqueness of fixed points of one dimensional transfer matrices. It turns out, that this question is intimately related with spontaneous symmetry breaking. Consider a Hamiltonian for a one dimensional quantum system with a degenerate ground state space spanned by the ONB $\{|\Psi_A\rangle, |\Psi_B\rangle\}$. Any (normalized) superposition of these two states is thus again a ground state. This is of immediate consequence for finding ground states with imaginary time evolution: The resulting ground state depends on the starting vector! We expect that the appearance of a degenerate largest eigenvalue of a transfer matrix (i.e. a degenerate “ground state” if we interpret the TM as a step in imaginary time evolution) marks the occurrence of symmetry breaking and thus a phase transition. Each of the dominant eigenvectors should encode a different equilibrium state corresponding to a particular way in which the symmetry may break. Furthermore, these dominant eigenvectors must be orthogonal to each other, since

¹⁸A *local* Hamiltonian is given by a sum of local terms.

they correspond to different ways in which the system may order. These arguments are far from rigorous of course, but the intuition is that there must exist a preferred ONB of the degenerate ground state space. This allows for the computation of order parameters. Any superposition of these basis vectors however, while still being a ground state, would certainly give an unphysical expectation value for the order parameter! We therefore need a way to target these preferred basis vectors. What happens now if the ground state is approximated by an MPS? In this case it turns out, that boundary MPS algorithms target a preferred basis of the degenerate ground state space: The algorithm converges only to MPSs which not only minimize the energy (or equivalently, maximize the partition function per site) but also the entanglement entropy. It turns out, that these MPSs are precisely the approximations to the preferred basis vectors we were looking for.

5.4 Spontaneous breaking of sublattice symmetry

A priori, hard core lattice gases are translation-invariant which is reflected in the fact, that the two dimensional TN is made up of equal tensors. However, for the k NN model, this translation invariance may spontaneously break in the high density regime - at least single-site translation invariance: As outlined in section 2.5, the fully packed configurations still exhibit a periodic structure, albeit with a larger “unit cell”.

This poses a challenge to boundary MPS methods. The issue could in principle be overcome by using multi-site boundary MPSs [26] which is however further complicated by the fact, that translation invariance generally breaks both in the horizontal (i.e. parallel to the boundary MPS) and the vertical direction. It would therefore be nice to always have a translation-invariant system.

Single-site translation invariance across the entire phase diagram can be recovered by blocking tensors which make up a unit cell together - thereby coarse-graining the TN, see section 3.3. As already mentioned, this comes at the cost of a very large bond dimension of the coarse grained tensors. However, in the case of hard core lattice gases, we may be in the fortunate situation that the coarse grained tensor T' has very large “null spaces”. By this we mean the null spaces of the matrices obtained by picking a particular index of T' and grouping the remaining indices together: We can therefore introduce tensors into the network which project out the null-spaces of the coarse grained tensors, thereby reducing their bond dimension while not changing the the partition function, see fig. 23. In other words, we can come up with a lossless compression scheme for the tensors making up the coarse grained network.

We shall employ a coarse-graining procedure where tensors are contracted pairwise. After that, the tensors V_R and V_L are found which reduce the




Figure 23: Coarse graining of a TN by grouping four tensors together to form a new one. The “compressors” (represented by red triangles) are chosen such that they (partially) remove the null space of the tensors (in the sense explained in the text) thereby reducing the bond dimension of the coarse grained TN without losing any information.

bond dimension:

(35)

This procedure may then be iterated until the desired number of original tensors has been included into the final coarse grained tensor, see fig. 23 and 24.

One way of obtaining V_R and V_L is via singular value decomposition. In a first step, the problem is re-expressed in terms of matrices:

(36)

Consider the product AB of two matrices $A \in \mathbb{C}^{l \times m}$ and $B \in \mathbb{C}^{m \times n}$ which are assumed to be rank deficient, i.e. $\text{rank}(A) = r_A < m$ and $\text{rank}(B) = r_B < m$. Our goal is now to find an isometry $U \in \mathbb{C}^{m \times r}$, where $r = \min\{r_A, r_B\}$, such that

$$AB = AU^\dagger UB = \tilde{A}\tilde{B}, \quad \tilde{A} \in \mathbb{C}^{l \times r}, \tilde{B} \in \mathbb{C}^{r \times n} \quad .$$

We consider the SVD of A and B

$$AB = U_A \Lambda_A V_A^\dagger U_B \Lambda_B V_B^\dagger$$

and note that the rank of a matrix is equal to the number of its non-zero singular values. Without changing AB , we can therefore replace $V_A^\dagger$ by $\tilde{V}_A^\dagger \in \mathbb{C}^{r_A \times m}$ which is obtained from $V_A^\dagger$ by only retaining the first $l - r_A$ rows. This is equivalent to inserting $\tilde{V}_A \tilde{V}_A^\dagger$ between A and B . For the case





Figure 24: Coarse graining by blocking a 3 by 3 square of tensors together. This can be used to study the 5NN model which has a 3 by 3 unit cell in the ordered phase.

$r_A < r_B$ we have therefore found our isometry $U = \tilde{V}_A^\dagger$. For $r_B < r_A$ we choose $U = \tilde{U}_B^\dagger$, where $\tilde{U}_B$ is obtained from U_B by retaining the first $l - r_B$ columns. To summarise:

$$U = \begin{cases} \tilde{V}_A^\dagger & r_A < r_B \\ \tilde{U}_B^\dagger & \text{else} \end{cases}$$

We can now obtain the tensors V_R and V_L by reshaping the matrices $U^\dagger$ and U into order three tensors. This procedure of compressing a TN is actually very similar to a single step of the Higher Order TRG algorithm [42]. The main difference is that the coarse graining is carried out exactly.

What happens if a single tensor in the original network is replaced by some impurity which encodes a local observable? Then the following problem arises: Blocking a uniform tensor and an impurity together may result in a tensor which has different null spaces than a tensor that is the result of blocking two uniform tensors. Therefore we can no longer use the isometries we found for the uniform TN using the procedure just described. We may however employ a different strategy by exploiting a property of the particular models we are dealing with. Consider the following network made up of tensors encoding some k NN model:



In general, there will be particular pairs of indices (i, i') and (j, j') for which the entire expression becomes zero, because all possible choices for the remaining indices correspond to entries equal to zero. By simply deleting these unnecessary index pairs, the bond dimension can be reduced. (Put differently, the matrices A and B in (36) may contain columns (A) or rows (B) which are only populated by zeros.) The tensors V_R and V_L in (35)

can now be chosen such as to merely eliminate these unnecessary indices. It should be pointed out however, that this procedure does not guarantee that the entire null-space has been removed. However it has another big advantage: Since the only property exploited for compressing the TN was the sparsity pattern of its constituent tensors, operators for local observables may now conveniently be obtained by replacing uniform tensors in the coarse graining network by appropriate impurities. This is true, since all the impurities we will be considering have a sparsity pattern very similar to the uniform tensors: Indices for non-zero entries of the impurity tensor always correspond to non-zero entries in the uniform tensors. Thus, the network of “compressing” tensors used to produce a single coarse grained tensor may serve as a scaffold for creating coarse grained observables.

It is important to realize, that the computational cost of the iTEBD algorithm increases quite fast with the bond dimension of the MPO: The bottleneck is the orthonormalization procedure which has a leading cost of $\mathcal{O}(D_{\max}^3 d^4)$ for an MPO with bond dimension d and a boundary MPS with a retained bond dimension of $D_{\max}$, see section 5.1. This renders iTEBD intractable for large unit cells, despite them being compressed.

5.5 Real-space renormalization: TRG

The Tensor Renormalisation Group (TRG) introduced by Levin and Nave [21] is a real-space renormalization scheme much in spirit to Kadanoff’s block-spin transformation [18][6]. The idea is to successively coarse-grain the tensor network description of the system, which means that s tensors in the original network get replaced by just one tensor in the coarse-grained network, where it is assumed, that all tensors are equal to each other. Therefore, the number of tensors decreases by a factor of s with each iteration step:

$$Z = \text{tTr } T^N = \text{tTr } T'^{N/s}$$

T' is the renormalized tensor which is given by coarse-graining transformation:

$$T' = R(T)$$

The map R is called the renormalization group (RG) transformation and - in principle - leaves Z invariant. It can be applied successively, i.e.

$$T^{(n)} = R\left(T^{(n-1)}\right)$$

and thus produces a “flow” of tensors:

$$T =: T^{(0)} \rightarrow T^{(1)} \rightarrow T^{(2)} \rightarrow \dots \rightarrow T^{(n)}$$

After n iterations the number of tensors has decreased to

$$N_n = \frac{N}{s^n} \quad .$$

Provided that the number of iterations was sufficiently large, N_n is small enough so that contracting the tensor network

$$Z = \text{tTr } T^{(n) N_n}$$

exactly becomes feasible. The way in which the final network is contracted depends on the boundary conditions chosen. This will be explained in more detail shortly. Note, that since Z generally grows with system size, the entries of the renormalized tensors $T^{(n)}$ do so as well during the iteration process. For practical implementations, it is therefore necessary to rescale $T^{(n)}$ after each RG step:

$$\tilde{T}^{(n)} = \frac{1}{c_n} T^{(n)}$$

The factors c_n must be chosen such as to ensure

$$\left\| \tilde{T}^{(n)} \right\|_{\infty} < \infty, \quad \forall n \quad ,$$

where $\|\cdot\|_{\infty}$ gives the entry with largest absolute value. The choice of c_n is not unique. One possibility is:

$$c_n = \left\| T^{(n)} \right\|_{\infty}$$

Regardless of the specific choice of c_n , the partition function can be rewritten as:

$$\begin{aligned} Z &= \text{tTr } T^N = c_1^{N/s} \text{tTr } \tilde{T}'^{N/s} \\ &= c_1^{N/s} c_2^{N/s^2} \dots c_n^{N/s^n} \text{tTr } \tilde{T}^{(n) N/s^n} \end{aligned}$$

The partition function per site is given by $\kappa = \lim_{N \rightarrow \infty} Z^{1/N}$. After n iterations we have:

$$\kappa_n := c_1^{1/s} c_2^{1/s^2} \dots c_n^{1/s^n} \left(\text{tTr } \tilde{T}^{(n) N/s^n} \right)^{1/N} \quad (37)$$

For actual implementations, it is convenient fix the system size by the maximum number of iterations:

$$N = s^{n_{\max}}$$

After $n_{\max}$ iterations, the network then only consists of a single tensor $\tilde{T}^{(n_{\max})}$:

$$\kappa_n := c_1^{1/s} c_2^{1/s^2} \dots c_n^{1/s^n} \left(\text{tTr } \tilde{T}^{(n_{\max})} \right)^{1/N}$$





Figure 25: Exact coarse-graining of a 1D tensor network with periodic boundary conditions according to (38)

Unfortunately, finding a RG-transform R that is exact, in the sense that it leaves Z invariant, is generally not possible in higher dimensions or would rather complicate matters further. To illustrate the problem, consider the TN representation of a partition function in 2D. Naively generalising the RG transform that was used in the 1D example, would lead to something like this:




In contrast to the 1D example, the coarse-grained tensor now has the squared bond dimension of the original one which renders the procedure infeasible after a few iterations. It therefore becomes necessary to somehow truncate the coarse grained tensor after each step, which in turn introduces an error into the partition function. Therefore finding a renormalized tensor T' for which


$$\left| \text{tTr } T^N - \text{tTr } T'^{N/s} \right|$$

is small, becomes the crucial step in the RG procedure. In this work, the original TRG (as first described by Levin and Nave [21] for the hexagonal and by Gu, Levin and Wen [14] for the square lattice) is used, for it is both easy to implement and only requires comparably modest computational resources. The core idea is to regard the local tensors as D^2 by D^2 matrices and replace them by matrices of lower rank D' . The optimal way to do this

is by computing the SVD and only keeping the D' largest singular values:




There are two choices where to “cut” a tensor: along the top-left to bottom right or along the bottom-left to top-right diagonal. The cuts are made according to the following pattern:



The coarse-grained tensors T' are obtained by rewiring:

$$T' = \text{blue circle} = \text{tilted square with colored triangles}$$

Thus again a square lattice is obtained, although it is tilted by $\pi/4$ with respect to the original one:



The procedure is repeated the desired number of times until the partition function can be calculated according to (37). If the size of the original network has been chosen such, that exactly one tensor is left after $n_{\max}$ iterations, then the final contraction is given by


$$\text{tTr } \tilde{T}^{(n_{\max})} = \text{blue circle with loops} \quad \text{or} \quad \text{tTr } \tilde{T}^{(n_{\max})} = \sum_{abcd} a \text{---} \text{blue circle} \text{---} c \text{---} d \text{---} b$$

for periodic or open boundary conditions, respectively. In the actual implementation, κ is calculated after each iteration step, and the algorithm halts as soon as κ has converged to within a certain relative precision (10^{-14} in the implementation used for this work).


The calculation of local expectation values requires a more careful analysis. Consider some local expectation value $\langle \hat{O} \rangle$:

$$\langle \hat{O} \rangle = \frac{1}{2} \text{tTr } AT^{N-1}$$

In the trace in the numerator, one of the local tensors has been replaced by an “impurity” A . In the 1D case, this is again a straightforward contraction, since there is exactly one impurity tensor in the network after each step. Consider again the previous example of a 1D network composed of 8 tensors, one of which now is an impurity:



In the TRG procedure described for the 2D case, the impurity needs to be cut during coarse-graining, thus introducing two impurities in the renormalized TN. One could naively assume, that the number of impurities increases with each step which would render the final contraction and thus the computation of expectation values intractable. Luckily it turns out, that there are certain arrangements of impurity tensors that remain invariant under coarse graining steps. For instance, if the network contains four impurity tensors sitting on the corners of the same plaquette, the coarse-grained TN looks exactly the same up to being tilted:



This is also described in [14]. Thus for some arrangement of m impurities

that is invariant under coarse-graining, we have:

$$Z\langle\hat{O}\rangle = \text{tTr } A_1 \dots A_m T^{N-m} = \text{tTr } A_1^{(n)} \dots A_m^{(n)} T^{(n)N/s^n-m}$$

Note, how the number m of impurity tensors is the same after each RG step. As for the uniform tensors, the impurity tensors are rescaled after each step in the actual implementation:

$$\tilde{A}_1 \dots \tilde{A}_m = \frac{1}{a_n} A_1 \dots A_m$$

$$a_n = \prod_{j=1}^n \|A_j\|_\infty$$


Now it makes sense to choose the original network size such, that after $n_{\max}$ steps the network is composed of exactly m tensors:

$$\frac{N}{s^{n_{\max}}} = m$$


Then one obtains:

$$\begin{aligned} \langle\hat{O}\rangle &= \frac{\text{tTr } A_1^{(n_{\max})} \dots A_m^{(n_{\max})}}{\text{tTr } T^{(n_{\max})m}} \\ &= \left(\prod_{i=1}^{n_{\max}} a_n c_n^m \right) \frac{\text{tTr } \tilde{A}_1^{(n_{\max})} \dots \tilde{A}_m^{(n_{\max})}}{\text{tTr } \tilde{T}^{(n_{\max})m}} \end{aligned} \quad (39)$$

For the case with four impurity tensors sitting in the corners of the same plaquette ($m = 4$), the network can be coarse-grained until it only consists of 4 impurities. For PBC, eq. (39) can then be written as

$$\langle\hat{O}\rangle = \left(\prod_{i=1}^{n_{\max}} a_n c_n^m \right) \frac{\text{tTr } \tilde{A}_1^{(n_{\max})} \dots \tilde{A}_m^{(n_{\max})}}{\text{tTr } \tilde{T}^{(n_{\max})m}}$$


while in the case of OBC

$$\langle\hat{O}\rangle = \left(\prod_{i=1}^{n_{\max}} a_n c_n^m \right) \frac{\text{tTr } \tilde{A}_1^{(n_{\max})} \dots \tilde{A}_m^{(n_{\max})}}{\text{tTr } \tilde{T}^{(n_{\max})m}}$$


where the open indices of the renormalized tensors have simply been traced out. As for the free energy, the expectation value is computed after each iteration step and checked for convergence. In the actual implementation, the expectation value is considered converged if the relative error is less than




Figure 26: Order parameter $|q_1|$ for the 1NN model vs. number of TRG steps n , at a chemical potential of $\mu = 1.4 > \mu_C$, for a maximum bond dimension $D_{\max} = 40$. The convergence behaviour is quite different for open and periodic boundary conditions: While for the former, $|q_1|$ wildly fluctuates and suddenly comes to a halt, for periodic boundary conditions, the order parameter fluctuates only mildly around zero (indiscernible in the plot) before smoothly increasing to its final value.

10^{-14} .

What happens if the the impurity tensor encodes a local order parameter? In principle, the expectation value of the order parameter should always be zero, as long as there are no external fields breaking the symmetry. However, the symmetry may break spontaneously in the appropriate parameter regime due to the build-up of numerical rounding errors, see fig. 26.

This is also true for sublattice symmetry breaking. In order to calculate expectation values of site-dependent observables, the original region containing the impurity tensor needs to be chosen large enough to contain a unit cell, i.e. all different sublattices. This will be further elaborated on before discussing the results for the 2NN model.

6 Numerical results

All algorithms were implemented in MATLAB®[37], and heavy use was made of the `ncon()` subroutine [28] for numerically contracting tensor networks. Computations were carried out in part on the Vienna Scientific Cluster 3.

The first question that may naturally arise when studying a system is about the nature of its phase diagram: We would like to find the parameter values for which the free energy becomes non-analytic, i.e. phase transitions take place. In the case of a HCLG of identical particles in the absence of external fields, we are in the fortunate situation that the only parameter is the chemical potential. The parameter space is thus one dimensional. Since the algorithms used in this work give direct access to the free energy, a naive approach of detecting phase transitions may consist of sampling the free energy across some region in parameter space and then checking whether its numerical derivatives are continuous. Due to finite numerical precision and the fact that taking numerical derivatives is very ill-conditioned, this approach is not particularly precise. However, we are also able to compute local observables, such as the density, which are related to derivatives of the free energy. For a first order phase transition, we therefore expect to see a jump in the density. Furthermore, boundary MPS algorithms allow us to compute the correlation length which is known to diverge at a critical point. We thus have an effective means of detecting the critical chemical potential for a second order phase transition.

Furthermore, we may look at the order parameter which we expect to have a jump discontinuity for first order transitions. For continuous transitions, we expect the order parameter to continuously increase for $\mu > \mu_C$. Yet another way offered by boundary MPS methods is to take a look at the entanglement entropy S_E of the corresponding quantum system at the boundary: It turns out, that S_E as computed from the boundary MPS, exhibits a local maximum at the point of both first order and continuous phase transitions.

For continuous transitions, we may also be interested in the critical exponents of various quantities (see section 2.3). We are however facing the problem, that the methods for approximately contracting a two dimensional TN described in this work cannot capture fluctuations on all length scales. For boundary MPS methods, this is already implied by the fact that MPSs cannot handle algebraically decaying correlations - the correlation length will always be finite for a finite bond dimension. We may still be able to extract information about the critical behaviour by looking at parameter values very close to the actual critical point. It is of utmost importance to realise, that boundary MPS methods must be expected to become increasingly less accurate as the critical point is approached. In an attempt

to systematically treat errors introduced by finite bond dimensions, a *finite entanglement scaling theory* has been put forward [36, 34]: Near criticality, various quantities may exhibit scaling behaviour as the bond dimension is varied, allowing for the extraction of critical exponents. The application of this procedure to the TN treatment of the k NN model is left for future studies. We shall content ourselves with a comparison of the critical exponents obtained for the highest considered bond dimension to what is already known [11].

6.1 1NN

The most accurate results available from previous works are for the 1NN model, and it thus serves as an excellent benchmark of the methods. Using corner transfer matrix methods, the partition function and the density for zero chemical potential have been calculated up to 43 and 41 decimal digits accuracy, respectively [2]:

$$\begin{aligned}\kappa(0) &= 1.5030480824753322643220663294755536893857810 \\ \rho(0) &= 0.22657081546271468894199226347129902640080\end{aligned}\tag{40}$$

Using iTEBD, $\kappa(0)$ is obtained, where a bond dimension of $D_{\max} = 12$ is sufficient to retain all Schmidt values $> 10^{-12}$ during the entire simulation. For the specific encoding of the 1NN model used, the dominant eigenvalue of the one dimensional transfer matrix corresponds to κ^2 , and is converged to within a relative error of 10^{-15} , leading to a total error for $\kappa(0)$ of $< 10^{-15}$. We give the results for $\kappa(0)$ and $\rho(0)$ to 14 and 12 significant digits respectively.

The same quantities are computed using TRG with bond dimensions $D_{\max} = 20, 40, 60$ for the partition function and $D_{\max} = 20, 40$ for the density. The results are converged to within a relative error of $< 10^{-14}$. From the $D_{\max} = 60$ TRG results for κ , the density at $\mu = 0$ is computed from the symmetric difference of $\kappa(\pm 10^{-4})$.

	$\kappa(0)$	$\rho(0)$
iTEBD	1.50304808247533	0.2265708154627
TRG ($D_{\max} = 60$)	1.50304808247533	0.226577 (symmetric difference)
TRG ($D_{\max} = 40$)	1.503048082468	0.226570809

$\kappa(0)$ and $\rho(0)$ as obtained with iTEBD are in perfect agreement with (40), to the given precision. TRG for $D_{\max} = 60$ performs equally well for $\kappa(0)$, while $\rho(0)$ computed as a local observable for $D_{\max} = 40$ using TRG is accurate up to including the 7th digit. From κ and ρ , the entropy may also be calculated, see 27.




Figure 27: Density (left) and entropy (right) vs. chemical potential for the 1NN model. ρ monotonously increases from zero until it saturates at the fully packed density of $1/2$, while the entropy shows a clear maximum at $\mu = 0$.

Next, the compressibility C is investigated. It shows a very pronounced peak near $\mu_C \approx 1.334$, where the system is expected to undergo a continuous phase transition. Since $C = \partial^2 f / \partial \mu^2$, we expect the occurrence of a discontinuity at - and power-law behaviour $C \propto |\tilde{\mu}|^{-\alpha}$ around this point. In order to determine α , the nature of the discontinuity needs to be ascertained. The data obtained from the simulations shows a peak of C at μ_C , see fig.28. If this peak is a genuine one, the compressibility attains a true maximum which is not an artefact due to insufficient resolution near the critical point. This would imply $\alpha = 0$ and the discontinuity could be due to a jump. If however C diverges, α would necessarily be > 0 . Unfortunately, reliable data close to μ_C is hard to come by, due to the inherent limitations of our methods and limited computational resources. The best we can do is to calculate C for various bond dimensions and different methods and see whether the nature of the peak changes. Using iTEBD, the compressibility is calculated by taking the numerical derivative of the density, which has been calculated as a local observable. TRG produces unphysical oscillating behaviour for the expectation value of the density operator around the critical point. Using TRG, C is therefore computed from the free energy by taking the second derivative numerically.

The data for C obtained in this manner hints to the case that C has indeed a peak instead of a divergence, implying $\alpha = 0$. It should however be kept in mind, that data close to criticality has a large systematic error especially since the bond dimension near μ_C must be chosen rather small due to computational cost.

At the critical point, the correlation length ξ is expected to diverge according to the power law $\xi \propto |\tilde{\mu}|^{-\nu}$. As already mentioned, a boundary MPS gives immediate access to a system's correlation length. Indeed, the




Figure 28: Compressibility C of the 1NN model: The main plot shows the behaviour of C near criticality as obtained from iTEBD and TRG, with no hint of a divergence. The inset shows the data obtained using iTEBD with $D_{\max} = 20, 40, 60, 80, 100$ and a step size of $\Delta\mu = 0.01$ which is visually indistinguishable.




Figure 29: Correlation length of the 1NN model: The main plot shows the data obtained using iTEBD with $D_{\max} 20, 40, 60, 80, 100$ and a step size of $\Delta\mu = 0.01$ which is visually indistinguishable. The inset shows the behaviour near criticality showing a clear increase of ξ with growing bond dimension. Note how the peak for $D_{\max} = 40$ already exceeds the one for $D_{\max} = 20$ by one order of magnitude.

data for ξ obtained from iTEBD shows a peak near μ_C , see fig. 29. MPS can handle only finite correlation lengths. Performing the simulation for various bond dimensions, strongly hints to an increase of the peak height with growing $D_{\max}$ very close to the critical point. We perform a non-linear least squares fit of $a|\mu_C - \mu|^{-\nu}$ to the data obtained from the $D_{\max} = 80$ run, where a, μ_C and ν are fit parameters. This is done both above and below the critical point, leading to

	$\mu_C(D_{\max} = 80)$	$\nu(D_{\max} = 80)$
above	1.33389(1)	0.937(2)
below	1.33416(1)	0.988(4)

where the uncertainties originate from the fitting procedure. Assuming ν is the same both above and below the critical point, we therefore estimate

$$\nu(D_{\max} = 80) = 0.96(3)$$

which is slightly below $\nu = 1$, which is to be expected for the 1NN model is known to belong to the Ising universality class.




Figure 30: Order parameter q_1 for the 1NN model.

To obtain a more accurate estimate for the critical chemical potential as well as an estimate for the critical exponent β , we investigate the behaviour of the order parameter above μ_C , see fig. 30. The function $a|\mu_C - \mu|^{-\nu}$ was fitted to the data from the ($D_{\max} = 80$) iTEBD run and the ($D_{\max} = 40$) TRG run. The obtained estimates for the critical exponent β and the location of μ_C are as follows:

	$\mu_C(D_{\max})$	$\beta(D_{\max})$
iTEBD ($D_{\max} = 80$)	1.334016(1)	0.12477(1)
TRG ($D_{\max} = 40$)	1.334074(1)	0.1208(4)

Again, the stated uncertainties originate only from the fit. Comparing with the result of $\mu_c = 1.3340151 \dots$ from [15], the value for μ_C is accurate up to including the fifth and fourth digit, for iTEBD and TRG respectively. For Ising universality we'd have $\beta = 0.125$ which is significantly closer to the value obtained with iTEBD.

The entanglement entropy S_E also shows a peak at the transition point whose height increases with bond dimension, see fig. 31. This is to be expected, since S_E diverges at the point of a continuous phase transition.

6.2 2NN

As discussed in section 2.5, the 2NN model describes a system of hard 2 by 2 squares and has a 2 by 2 unit cell in the broken phase. In order to render the problem amenable to a single-site iTEBD treatment, the corresponding TN is first coarse-grained exactly by grouping four tensors together according




Figure 31: Entanglement entropy of the boundary MPS for the 1NN model.

to 5.4. For the coarse-grained description of the 2NN model, the dominant eigenvalue of the MPO is therefore equal to κ^4 . Using iTEBD, the TN is contracted using boundary MPSs of maximum bond dimensions $D_{\max} = 40, 60$. Going up higher turns out to be too expensive, since the bond dimension of the coarse grained (and compressed) MPO is $d = 5$ as compared to merely $d = 2$ as for the 1NN model. Numerically it is observed, that a given “top” fixed-point MPS is also a “bottom” fixed point. This convenient for the calculation of site-dependent observables, since once we have found, let’s say a “top” fixed point, we immediately have a corresponding “bottom” one, in the sense that they are in the same phase.

In a first step, we again study the system at $\mu = 0$ and observe, that a maximum bond dimension of 24 suffices to retain all Schmidt values $> 10^{-12}$. The dominant eigenvalue, i.e. κ^4 , is again converged to within a relative error of $< 10^{-14}$ which happens to coincide with the relative error for $\kappa(0)$. $\kappa(0)$ and $\rho(0)$ are also computed using TRG.

	$\kappa(0)$	$\rho(0)$
iTEBD	1.342643951124	0.1367658677
TRG ($D_{\max} = 60$)	1.3426439550	0.1367658656 (symmetric difference)
TRG ($D_{\max} = 40$)	1.342643939	0.1367661

From our data we estimate:

$$\kappa(0) = 1.34264395, \quad \rho(0) = 0.13676586$$




Figure 32: Density (left) and entropy (right) vs. chemical potential for the 2NN model.




Figure 33: Compressibility (left) for the 2NN model, showing a clear peak near the transition point. The entanglement entropy S_E shows two converging branches in the ordered phase.

The plots for the entropy and the density in fig 32 are again in perfect agreement with what is to be expected for the 2NN model, as discussed in section 2.

Unfortunately, the behaviour of the compressibility C around the critical point cannot be investigated as thoroughly as for the $k = 1$ case, due to a lack of sufficient data. Obtaining the compressibility via TRG also fails around the critical point, since both the free energy and the density are too inaccurate. We therefore cannot determine whether the peak at μ_C (see fig. 33) is actually a divergence or not. Interestingly, the compressibility shows a very clear maximum aside from the peak at μ_C at around $\mu \approx -1.3$.

The system becomes anisotropic in the ordered phase: It exhibits either row- or column-order. These rows or columns may either be parallel or or-




Figure 34: Correlation length of the 2NN model as obtained using iTEBD with maximum bond dimensions of $D_{\max} = 40, 60$. As can be seen in the left plot, ξ exhibits two different branches in the ordered phase (see main text). Near the critical point, the data hints to an increase of the correlation length with growing bond dimension. However the situation is much less clear than in the 1NN case.

thogonal to the boundary MPS, depending on the randomly chosen start MPS. Interestingly, the correlation length as obtained from the boundary MPS, shows two different branches in the ordered phase, see fig. 34: One branch seems to diverge, while the other one decays. We offer the explanation, that parallel to the direction of columnar order, the system effectively becomes one dimensional and is thus described by the 1NN model in one dimension. The latter exhibits a diverging correlation length for $\mu \rightarrow \infty$, see equation (21). In the direction orthogonal to the rows/columns, the correlation length decays, since the rows/columns may slide freely with respect to each other. Notably, the entanglement entropy also shows a branching behaviour, see fig. 33. Both branches converge to constant values with increasing μ , where one of the branches approaches zero. This particular branch likely corresponds to the case where the boundary MPS is orthogonal to the columns: The columns can move freely with respect to each other, i.e there is little correlation (or none at all for $\mu \rightarrow \infty$) between them which is reflected in vanishing entanglement entropy.

Unfortunately, the number of samples near the critical point is insufficient to make any reliable estimate for the critical exponent ν . We leave the determination of ν to further studies.

The order parameter q_2 , as defined in eq. (12), was calculated both using TRG and iTEBD, see fig. 35. From the behaviour of its absolute value $|q_2|$ near criticality, both μ_C and the exponent β may be estimated:




Figure 35: Absolute value of the order parameter q_2 for the 2NN model.

	$\mu_C(D_{\max})$	$\beta(D_{\max})$
iTEBD ($D_{\max} = 60$)	4.5827(2)	0.1017(6)
TRG ($D_{\max} = 40$)	4.4472(2)	0.103(2)

The given uncertainties are again only due to the fitting procedure. The estimate for μ_C from the iTEBD data is in much better agreement with the value of $\mu_C \approx 4.574$ obtained in [11]. In [11] it was concluded, that the 2NN model lies in the Ising universality class, and thus has an exponent $\beta = 0.125$. Later on however it was noted, that the 2NN model is expected to be *not* in the Ising universality class[30]. Although Ising universality seems to be ruled out within the uncertainties of the fit, the estimate has most likely a systematic error which is too large to settle the question.

6.3 3NN

For the 3NN model we refrain from calculating the order parameter. In principle, a coarse-graining procedure could be employed to achieve this goal, as was done in the $k = 2$ case. This is left for further study. We may however still be able to calculate local observables which are not site dependent, such as the *average* density. While this works for iTEBD, computing the density as a local observable with TRG is hampered by convergence problems, and we therefore obtain the density from the TRG data by taking the numerical derivative.

At $\mu = 0$ a boundary MPS with $D_{\max} = 58$ is necessary to retain Schmidt values $> 10^{-12}$ - bot at the “top” and the “bottom” of the non-symmetric




Figure 36: Density and entropy for the 3NN model. Both show a pronounced jump discontinuity at around $\mu_C \approx 3.7$.

3NN-MPO. We obtain:

	$\kappa(0)$	$\rho(0)$
iTEBD	1.2748320327	0.10469720191
TRG ($D_{\max} = 60$)	1.2748321675	0.10469748 (symmetric difference)
TRG ($D_{\max} = 40$)	1.2748343653	0.10469726 (symmetric difference)

From the data, $\kappa(0)$ and $\rho(0)$ may be estimated and are believed to be accurate to the given precision:

$$\kappa(0) = 1.27483, \quad \rho(0) = 0.10469$$

The clearly discernible jump in the density, strongly hints to a first order transition, see fig. 36. From our data we estimate the location of the transition point μ_C , as well as the “height” $\Delta\rho$ of the density jump:

$$\mu_C = 3.69(2) \quad \Delta\rho = 2.99(7) \cdot 10^{-2}$$

This is in agreement with the values obtained in previous works [7, 11]. The transition point suggested by the TRG data is much closer to the values $\mu_C = 3.6746$ and $\mu_C = 3.6762$ which were found in [11] and [7] respectively. The compressibility shows a divergent peak at the transition point. Furthermore, it has a local maximum of ≈ 0.0262 at around $\mu \approx -1.6$, see fig 38. The correlation length and the entanglement entropy also show jumps at the transition point, see fig. 39 and fig. 40. Note, how the peak of the correlation length has the same order of magnitude for all considered bond dimensions. It is furthermore rather small as compared to the $k = 1, 2$ cases, which are believed to be continuous transitions and the corresponding peaks in the correlation lengths are expected to actually be divergences.

Although the peak of the entanglement entropy still increases with the bond dimensions considered, the mentioned jump of S_E sets it apart from the previous two cases.




Figure 37: Density of the 3NN model near the phase transition. For iTEBD, the transition point μ_C seems to decrease with growing bond dimension.




Figure 38: Compressibility of the 3NN model as obtained with iTEBD for $D_{\max} = 40, 60$ (indiscernible in the plot).




Figure 39: Correlation length of the 3NN model near the transition point for various bond dimensions. The inset shows the overall behaviour as obtained using iTEBD with $D_{\max} = 60$.




Figure 40: Entanglement entropy of the boundary MPS for the 3NN model

7 Conclusion and outlook

In this thesis, athermal lattice systems and lattice gases in particular, have been reviewed. Special attention has been paid to the k nearest neighbour exclusion model on the square lattice, mainly building on the discussions in [11, 30, 24].

The main point however, has been the formulation of the k NN model in terms of tensor networks: While the latter have in fact been around for many decades under different names (e.g. the transfer matrix formalism), using tensor network states as variational ansätze for ground states of quantum many-body systems has sparked the development of efficient approximate contraction schemes for two dimensional TNs. Furthermore, these ideas have been extended to treat infinite translation-invariant systems. Two such examples, iTEBD and TRG, have been reviewed in this thesis. While this work was already well under-way, the author became aware of algorithms which are probably better suited (for this particular problem): The Corner Transfer Matrix Renormalization group [25] should be mentioned, which was specifically designed for classical problems. Furthermore, the iTEBD algorithm suffers from various drawbacks which are absent in different boundary MPS methods: A corner-based boundary MPS method has been developed, which is much better suited than iTEBD for MPOs with large bond dimensions and exhibits faster convergence [38].

In this light, the present study should be regarded as a proof of concept: It has been demonstrated, that infinite system size TN algorithms are indeed capable of simulating hard core lattice gases in two dimensions, with the results being in agreement with previous studies. However, there is still much room for improvement.

Historically, simulations of hard-core lattice gases have usually been done for finite system sizes: Whether the employed methods were based on Monte Carlo or transfer matrix algorithms, the parameter controlling the approximation has been the finite number of lattice sites. If it was not for a systematic treatment of finite size induced errors, the study of phase transitions using these methods would be severely hampered, and reliable estimates for critical exponents would be hard to come by.

The parameter controlling the approximation and offering a trade-off between accuracy and computational cost in the case of TN algorithms, is the retained bond dimension. Similar to finite size scaling, the introduced errors may be treated by what is called finite entanglement scaling [34].

The main challenges and possible directions in the TN treatment of hard-core lattice gases are now summarised in the following points:

- Better contraction methods: Computing environments is certainly the main computational bottleneck for most tensor network studies. Aside

from the fact that computational power is increasing as technology advances, the development of better algorithms is far from stagnant. More efficient methods will pave the way to tackle ever more complicated systems.

- Large unit cells: One of the main challenges in the tensor network treatment of hard-core lattice gases is the breaking of sublattice symmetry. This poses a problem to algorithms relying on translational invariance for the treatment of infinite system sizes. Another requirement for robust methods is therefore also the ability to treat large unit cells.
- Finite entanglement scaling: If applied to hard core lattice gases, this would certainly lead to much better estimates for critical exponents.

Finally, the author would like to express his confidence, that TN based methods have the potential to become competitive with Monte Carlo and conventional transfer matrix methods for the treatment of hard-core lattice gases. Ultimately, it would be interesting to go up to ever larger excluded volumes which would allow for extrapolating the results to the continuum, i.e. the hard disks, case. Similar to hard-square constant (40), a “hard disks” constant - i.e. the number of possible hard disks configurations in the plane - could be estimated. As far as the author is aware at the time of writing, this has not been done so far.

A Some linear algebra

In the following sections important results from linear algebra which are used throughout this work are reviewed. Proofs are given if they are short and/or deliver insights of relevance to the remainder of the text. Otherwise the reader is referred to the relevant literature [31, 12]

A.1 Singular value decomposition

Let $A \in \mathbb{C}^{m \times n}$ be an arbitrary matrix of rank r . There exists a unique decomposition of the form

$$A = U \Sigma V^\dagger$$

where $U \in \mathbb{C}^{m \times r}$ and $V^\dagger \in \mathbb{C}^{r \times n}$ are isometries, i.e

$$U^\dagger U = \mathbb{1}_m, \quad V^\dagger V = \mathbb{1}_n$$

and $\Sigma = \text{diag}(s_1, \dots, s_r)$ (without proof). This is the so-called singular value decomposition (SVD) and the $s_1 \geq s_2 \geq \dots \geq s_r > 0$ are called singular values [31]. If A is square, both U and V are actually unitaries, while in the rectangular case just one of them is.

The SVD can be used to optimally approximate (in the least-squares sense) an arbitrary matrix by a lower rank one. For a given matrix A , the Frobenius norm is defined as:

$$\|A\|_F := \sqrt{\text{Tr } A^\dagger A} = \sqrt{\sum_{ij} |A_{ij}|^2}$$

We want to find a matrix B which has the same size as A but is of lower rank, i.e $\text{rank}(B) = r' < r = \text{rank}(A)$, for which

$$\|A - B\|_F = \min$$

Let $A = U \Sigma V^\dagger$ denote the SVD of A . Then the solution turns out to be given by

$$B = U \Sigma' V^\dagger$$

where $\Sigma' = \text{diag}(s_1, \dots, s'_r, 0, \dots, 0)$. The optimal lower rank approximation is thus given by simply setting the $r - r'$ smallest singular values to zero. This result is known as the Eckart-Young theorem [12].

A.2 Non-negative and convergent matrices

An eigenvalue λ_0 of a matrix A is called *dominant*, if and only if it is equal to the spectral radius of A and non-degenerate. A corresponding eigenvector is called dominant eigenvector. It would of course be nice to have a means of telling whether a given matrix exhibits a dominant eigenvalue or not. The

Perron-Frobenius theorem does exactly that. Before stating it, some more terminology is needed.

A matrix A is called non-negative/positive, if it is non-negative/positive entry-wise. This definition is used throughout this work and should not be confused with positive *definite*. Let A be a non-negative $n \times n$ matrix, with $n > 1$. A is called *irreducible* if and only if there exists no permutation matrix P such that

$$P^T A P = \begin{pmatrix} B & C \\ 0 & D \end{pmatrix}$$

i.e A cannot be brought into block-triagonal form by permuting rows and columns.

To check whether a given non-negative matrix is irreducible, the following property may be useful: It turns out, that a non-negative $n \times n$ matrix A is irreducible $\iff (\mathbb{1}_n + A)^{n-1}$ is strictly positive.

Theorem 1. (*Perron-Frobenius*) *Let A be a non-negative irreducible matrix with spectral radius ϱ . Then:*

- $\varrho > 0$
- ϱ is the dominant eigenvalue of A .
- There exists a corresponding eigenvector x that is real and positive.

For a proof of the Perron-Frobenius theorem see eg. [31]. Note that any positive square matrix is irreducible and thus the Perron-Frobenius theorem holds for all matrices of this kind. Furthermore, the eigenvalues of a matrix are analytic functions of its entries, the only non-analyticities arising due to degeneracy. Thus if all entries are analytic functions of some parameter t , so is the largest magnitude eigenvalue - as long as the matrix in question fulfils the requirements of the Perron-Frobenius theorem for all t .

Throughout this work, powers of matrices are frequently encountered. The following discussion of convergent matrices closely follows the one presented in [31].

A matrix $A \in \mathbb{C}^{n \times n}$ is called *convergent* if and only if the entries of A^N converge to zero.

$$\lim_{N \rightarrow \infty} |(A^N)_{ij}| = 0 \quad \forall i, j = 1, \dots, n$$

Whether a given matrix is convergent or not depends on its spectral radius: Let $A \in \mathbb{C}^{n \times n}$ with spectral radius ϱ . It is convergent if and only if $\varrho < 1$. The entry-wise rate of convergence is given by

$$|(A^N)_{ij}| \leq C N^l \varrho^N$$

where l and C are some positive constants.

Proof. Let v be an eigenvector corresponding to an eigenvalue λ_0 of magnitude ϱ . If A is convergent, then

$$0 = \lim_{N \rightarrow \infty} A^N v = \lim_{N \rightarrow \infty} \lambda_0^N v \Rightarrow |\lambda_0| = \varrho < 1$$

To see the converse, consider the Jordan canonical form of A :

$$A = SJS^{-1}$$

Recall that the Jordan matrix J is the direct sum of the Jordan blocks:

$$J = \begin{pmatrix} J_{k_0}(\lambda_0) & & 0 \\ & J_{k_1}(\lambda_1) & \\ & & \ddots \\ 0 & & & J_{k_n}(\lambda_n) \end{pmatrix}$$

$\lambda_0, \dots, \lambda_n$ are the eigenvalues of A , including *geometric* multiplicities. Each Jordan block $J_k(\lambda)$ is a $k \times k$ matrix of the form:

$$J_k(\lambda) := \begin{pmatrix} \lambda & 1 & & 0 \\ & \lambda & 1 & \\ & & \ddots & \ddots \\ 0 & & & \lambda & 1 \\ & & & & \lambda \end{pmatrix}, \quad J_1(\lambda) := \lambda$$

(In the case $k_i > 1$ the algebraic multiplicity of the corresponding eigenvalue λ_i is greater than the geometric one.) Also note, that the spectral radius of a Jordan block $J_k(\lambda)$ is equal to $|\lambda|$. For A^N we therefore have:

$$A^N = SJ^N S^{-1} = S \begin{pmatrix} J_{k_0}(\lambda_0)^N & & 0 \\ & J_{k_1}(\lambda_1)^N & \\ & & \ddots \\ 0 & & & J_{k_n}(\lambda_n)^N \end{pmatrix} S^{-1}$$

It remains to show that the Jordan blocks $J_k(\lambda)$ are convergent.

$$J_k(\lambda)^N = (J_k(0) + \lambda \mathbb{1}_k)^N = \sum_{j=0}^N \binom{N}{j} J_k(0)^j \lambda^{N-j} = \sum_{j=0}^{k-1} \binom{N}{j} J_k(0)^j \lambda^{N-j}$$

In the last step it was used that $J_k(0)^j = 0$ for $j \geq k$. Note that for $0 < j < k-1$ the entries of $J_k(0)^j$ are zero everywhere except along the j th superdiagonal, where they are equal to one. Therefore each term in the sum contributes to a different set of matrix entries. It thus suffices to show that

$$\lim_{N \rightarrow \infty} \binom{N}{j} \lambda^{N-j} = 0$$

which can be seen by considering

$$\ln \left| \binom{N}{j} \lambda^{N-j} \right| \leq \ln N^j \varrho^N = j \ln N + N \ln \varrho \rightarrow -\infty \quad (N \rightarrow \infty)$$

since $\varrho < 1$ and thus $\ln \varrho < 0$.

To verify the claim for the rate of convergence, we again observe the behaviour of the non-zero entries for each Jordan block:

$$\left| \binom{N}{j} \lambda^{N-j} \right| \leq N^j |\lambda|^N \leq N^l \varrho^N$$

where l is the size of the largest Jordan block, i.e. $l = \max \{k_0, \dots, k_n\}$. Therefore

$$\left| (A^N)_{ij} \right| \leq C N^l \varrho^N =: b_N$$

where C is determined by S .

□

Note that the upper bound b_N converges linearly, with rate ϱ :

$$\frac{b_{N+1}}{b_N} = \varrho \left(1 + \frac{1}{N}\right)^l \rightarrow \varrho$$

Now consider an irreducible non-negative square matrix A with spectral radius\dominant eigenvalue ϱ . As an immediate consequence of the preceding theorem, it follows that

$$\lim_{N \rightarrow \infty} \left(\frac{1}{\varrho} A \right)^N = S \begin{pmatrix} 1 & 0 & & \\ 0 & 0 & & \\ & & \ddots & \\ & & & 0 \end{pmatrix} S^{-1}$$

Note that the first column of S and the first row of S^{-1} are just the left and right dominant eigenvectors $v_L^\dagger$ and v_R , respectively.

$$v_L^\dagger A = \lambda_0 v_L^\dagger, \quad A v_R = \lambda_0 v_R, \quad v_L^\dagger \cdot v_R = 1$$

We thus obtain

$$\lim_{N \rightarrow \infty} \left(\frac{1}{\varrho} A \right)^N = v_R v_L^\dagger =: E$$

Note that E is a projector and its trace is equal to one.

$$E^2 = E, \quad \text{tr } E = v_L^\dagger v_R = 1$$

Furthermore this implies, that for a vector $v_\perp$ orthogonal to the dominant eigenvector we have:

$$\lim_{N \rightarrow \infty} \left(\frac{1}{\varrho} A \right)^N v_\perp = E v_\perp = 0$$

References

- [1] R J Baxter. Hard hexagons: exact solution. *Journal of Physics A: Mathematical and General*, 13(3):L61, 1980.
- [2] R. J. Baxter. Planar lattice gases with nearest-neighbor exclusion. *Annals of Combinatorics*, 3(2):191–203, 1999.
- [3] Rodney J. Baxter. *Exactly Solved Models in Statistical Mechanics*. Dover Publications, 3 edition, 2007.
- [4] A. Bellemans and R. K. Nigam. Phase transitions in two-dimensional lattice gases of hard-square molecules. *The Journal of Chemical Physics*, 46(8):2922–2935, 1967.
- [5] Etienne P. Bernard and Werner Krauth. Two-step melting in two dimensions: First-order liquid-hexatic transition. *Phys. Rev. Lett.*, 107:155704, Oct 2011.
- [6] Efi Efrati, Zhe Wang, Amy Kolan, and Leo P. Kadanoff. Real-space renormalization in statistical mechanics. *Rev. Mod. Phys.*, 86:647–667, May 2014.
- [7] E. Eisenberg and A. Baram. A first-order phase transition and a super-cooled fluid in a two-dimensional lattice gas model. *EPL (Europhysics Letters)*, 71(6):900, 2005.
- [8] Artur Ekert and Peter L. Knight. Entangled quantum systems and the schmidt decomposition. *American Journal of Physics*, 63(5):415–423, 1995.
- [9] Michael Engel, Joshua A. Anderson, Sharon C. Glotzer, Masaharu Isobe, Etienne P. Bernard, and Werner Krauth. Hard-disk equation of state: First-order liquid-hexatic transition in two dimensions with three simulation methods. *Phys. Rev. E*, 87:042134, Apr 2013.
- [10] G. Evenbly and G. Vidal. Tensor network renormalization. *Phys. Rev. Lett.*, 115:180405, Oct 2015.
- [11] Heitor C. Marques Fernandes, Jeferson J. Arenzon, and Yan Levin. Monte carlo simulations of two-dimensional hard core lattice gases. *The Journal of Chemical Physics*, 126(11), 2007.
- [12] Charles F. Van Loan Gene H. Golub. *Matrix Computations*. Johns Hopkins University Press, 4 edition, 2012.
- [13] Nigel Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group*. Westview Press, 1 edition, 1992.

- [14] Zheng-Cheng Gu, Michael Levin, and Xiao-Gang Wen. Tensor-entanglement renormalization group approach as a unified method for symmetry breaking and topological phase transitions. *Phys. Rev. B*, 78:205116, Nov 2008.
- [15] Wenan Guo and Henk W. J. Blöte. Finite-size analysis of the hard-square lattice gas. *Phys. Rev. E*, 66:046140, Oct 2002.
- [16] Jutho Haegeman. *Variational Renormalization Group Methods for Extended Quantum Systems*. PhD thesis, Universiteit Gent, 2011.
- [17] Jutho Haegeman and Frank Verstraete. Diagonalizing transfer matrices and matrix product operators: a medley of exact and computational methods. 2016. to be published.
- [18] L. P. Kadanoff. Scaling laws for Ising models near $T(c)$. *Physics*, 2:263–272, 1966.
- [19] Werner Krauth. *Statistical mechanics : algorithms and computations*. Oxford University Press, 1 edition, 2006.
- [20] T. D. Lee and C. N. Yang. Statistical theory of equations of state and phase transitions. ii. lattice gas and ising model. *Phys. Rev.*, 87:410–419, Aug 1952.
- [21] Michael Levin and Cody P. Nave. Tensor renormalization group approach to two-dimensional classical lattice models. *Phys. Rev. Lett.*, 99:120601, Sep 2007.
- [22] Christos N. Likos. Effective interactions in soft condensed matter physics. *Physics Reports*, 348(4–5):267 – 439, 2001.
- [23] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [24] Trisha Nath and R. Rajesh. Multiple phase transitions in extended hard-core lattice gas models in two dimensions. *Phys. Rev. E*, 90:012120, Jul 2014.
- [25] Tomotoshi Nishino and Kouichi Okunishi. Corner transfer matrix renormalization group method. *Journal of the Physical Society of Japan*, 65(4):891–894, 1996.
- [26] R. Orús and G. Vidal. Infinite time-evolving block decimation algorithm beyond unitary evolution. *Phys. Rev. B*, 78:155117, Oct 2008.

- [27] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117 – 158, 2014.
- [28] Robert N. C. Pfeifer, Glen Evenbly, Sukhwinder Singh, and Guifre Vidal. Ncon: A tensor network contractor for matlab. *arXiv:physics.comp-ph/1402.0939*, Feb 2014.
- [29] Robert N. C. Pfeifer, Jutho Haegeman, and Frank Verstraete. Faster identification of optimal contraction sequences for tensor networks. *Phys. Rev. E*, 90:033315, Sep 2014.
- [30] Kabir Ramola and Deepak Dhar. High-activity perturbation expansion for the hard square lattice gas. *Phys. Rev. E*, 86:031135, Sep 2012.
- [31] Charles R. Johnson Roger A. Horn. *Matrix Analysis*. Cambridge University Press, 2 edition, 2013.
- [32] Subir Sachdev. *Quantum Phase Transitions*. Cambridge University Press, 3 edition, 1999.
- [33] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96 – 192, 2011. January 2011 Special Issue.
- [34] Vid Stojevic, Jutho Haegeman, I. P. McCulloch, Luca Tagliacozzo, and Frank Verstraete. Conformal data from finite entanglement scaling. *Phys. Rev. B*, 91:035120, Jan 2015.
- [35] Masuo Suzuki. Relationship between d-dimensional quantal spin systems and (d+1)-dimensional ising systems. *Progress of Theoretical Physics*, 56(5):1454, 1976.
- [36] L. Tagliacozzo, Thiago. R. de Oliveira, S. Iblisdir, and J. I. Latorre. Scaling of entanglement support for matrix product states. *Phys. Rev. B*, 78:024410, Jul 2008.
- [37] Inc. The MathWorks. *MATLAB®*, v8.5 R2015a. Natick, Massachusetts, USA, 2015.
- [38] Laurens Vanderstraeten, Jutho Haegeman, Philippe Corboz, and Frank Verstraete. Gradient methods for variational optimisation of projected entangled-pair states. *arXiv:1606.09170*, Jun 2016.
- [39] F. Verstraete, V. Murg, and J.I. Cirac. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Advances in Physics*, 57(2):143–224, 2008.

- [40] G. Vidal. Classical simulation of infinite-size quantum lattice systems in one spatial dimension. *Phys. Rev. Lett.*, 98:070201, Feb 2007.
- [41] Matthias Vojta. Quantum phase transitions. *Reports on Progress in Physics*, 66(12):2069, 2003.
- [42] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, and T. Xiang. Coarse-graining renormalization by higher-order singular value decomposition. *Phys. Rev. B*, 86:045139, Jul 2012.
- [43] Li-Ping Yang, Yuzhi Liu, Haiyuan Zou, Z. Y. Xie, and Y. Meurice. Fine structure of the entanglement entropy in the $o(2)$ model. *Phys. Rev. E*, 93:012138, Jan 2016.
- [44] Yongsheng Zhang, Volker Blum, and Karsten Reuter. Accuracy of first-principles lateral interactions: Oxygen at $pd(100)$. *Phys. Rev. B*, 75:235406, Jun 2007.

We are gods with anuses.
-Ernest Becker