UNIVERSITÄT
DES
SAARLANDES

UNIVERSITÀ
DEGLI STUDI
DI PADOVA

TREE TENSOR NETWORKS FOR HIGH-DIMENSIONAL QUANTUM SYSTEMS AND BEYOND

By

Timo Felser

A DISSERTATION SUBMITTED FOR THE DEGREE OF
Doktor der Naturwissenschaften (Dr. rer. nat.)
TO THE NATURWISSENSCHAFTLICH-TECHNISCHEN FAKULTÄT AT THE UNIVERSITY OF
SAARLAND
AND FOR THE DEGREE OF
Dottore di Ricerca di Fisica
TO THE DIPARTIMENTO DI FISICA E ASTRONOMIA "GALILEO GALILEI" OF THE
UNIVERSITY OF PADOVA

IN MAY 2021

This doctoral thesis was prepared under the joint supervision of

Prof. Dr. Giovanna Morigi, Saarland University
and of
Prof. Dr. Simone Montangero, University of Padova

Tag des Kolloquiums:	30.11.2021
Dekan:	Prof. Dr. Jörn Erik Walter
Berichterstatter(-innen):	Prof. Dr. Giovanna Morigi
	Prof. Dr. Simone Montangero
	Prof. Dr. Karl Jansen
	Prof. Dr. Matteo Rizzi
Weitere Mitglieder:	Prof. Dr. Marco Zanetti
Vorsitz:	Prof. Dr. Christoph Becher
Akad. Beisitz:	PD Dr. Reza Sheabani

Statement of Contributions to Jointly Authored Works Contained in this Thesis

This thesis is partly of cumulative nature including publications that were the result of the work I have carried out since the commencement of my candidature for a higher research degree. In particular, it contains two publications in which I have co-authored as first author and three publications in which I have co-authored as a secondary author. In the following, I present the statements of contributions to the jointly authored works contained in this thesis.

First authorship

Two-Dimensional Quantum-Link Lattice Quantum Electrodynamics at Finite Density *Physical Review X*, **10** [37] - Incorporated as *Chapter 4*

In this publication, we developed and applied a Tensor Network to study lattice gauge theories in two spacial dimensions. This study includes two main challenges: On one hand, there is a theoretical challenge of reformulating a two-dimensional lattice gauge theory with fermionic matter into a system of local gauge-invariant dressed sites which can be used as a computational basis for the numerical analysis (i.e. does not exhibit any Jordan-Wigner strings outside the dressed sites). On the other hand, this reformulation resulted in a complex optimisation problem with numerous optimisation constraints which happens to be a numerical challenge for current Tensor Network algorithms.

Prof. Dr. Simone Montangero and Dr. Pietro Silvi initialised the project and defined the problem together with the conceptual strategy on how to deploy Tensor Networks for this particular application. Prof. Simone Montangero further supervised the execution. The above-mentioned theoretical challenge was mainly solved by Dr. Pietro Silvi while I developed the Tensor Network software addressing the above-mentioned numerical challenge. Dr. Mario Collura defined further physical cases validating the numerical correctness of the developed TN algorithm and managed the writing of the manuscript. The numerical analysis was executed by myself under the supervision of Dr. Mario Collura and with a joined interpretation of the results from all the authors.

Dr. Mario Collura prepared the first draft of the manuscript. All of the authors contributed to refining the manuscript by reviewing and editing accordingly. Further, in the comprehensive appendix, the *Appendices A,B,D,E* were mainly written by myself, *Appendices C,F,G* by Dr. Pietro Silvi and the *Appendices H-I* by Dr. Mario Collura. The Introduction and conclusion were written in mutual agreement with Dr. Pietro Silvi, Dr. Mario Collura and Prof. Simone Montangero as main contributors.

Efficient Tensor Network ansatz for high-dimensional quantum many-body problems *Physical Review Letters*, **126** [23] - Incorporated in *Chapter 5*

In this publication, we developed a novel Tensor Networks geometry, the augmented Tree Tensor Network (aTTN), for studying quantum many-body problems in higher dimensions. The clear goal here was to find a Tensor Networks geometry which on one hand encodes the so-called *area-law* for quantum states in higher dimensions and which on the other hand comes with efficient use of computational resources, i.e. is efficiently contractable and optimisable.

Prof. Dr. Simone Montangero initialised the project and defined the goal for this work. The theory and the conceptual idea behind the novel TN geometry were developed by myself together with Prof. Simone Montangero in supervision. Further, I developed the Tensor

Network software implementing the aTTN geometry, including the conception, planning, development and testing of the software. I benchmarked the software in the numerical study of the Ising model and the Heisenberg model, comparing it with alternative numerical methods. Dr. Simone Notarnicola further introduced the study of interacting Rydberg-atoms in 2D as a novel physical application where he defined the problem and applied the TN software in an extensive study. Prof. Dr. Simone Montangero supervised and merged this research.

The first draft on the introduction of the developed aTTN was written by myself including the theoretical sections *Augmented Tree Tensor Network* and *Area Law in aTTN* as well as the benchmark sections *Ising model* and *Heisenberg model*. Dr. Simone Notarnicola wrote the parts regarding the analysis on the Rydberg-atoms resulting in the section *Interacting Rydberg atoms*. All of the authors contributed to refining the manuscript by reviewing and editing accordingly. Further, Dr. Simone Notarnicola wrote the section *Rydberg atom phase diagram analysis* in the appendix, while I wrote the remaining parts including among other more technical details about the aTTN.

Second authorship

On the descriptive power of Neural-Networks as constrained Tensor Networks with exponentially large bond dimension   *SciPost Physics Core*, **4** [98] - Incorporated as *Chapter 6*

In this publication, we compared Tensor Networks against Neural Networks in their ability to represent quantum wave-functions. All authors agreed on the approach to pursuing this work and contributed to the interpretation of the results and to the writing of the manuscript. Prof. Dr. Luca Dell'Anna conceived the work. Dr. Mario Collura and I performed the numerical simulations. Prof. Dr. Simone Montangero supervised and merged the research.

In particular, the implementation of the Tensor Network software (MPS) used for the comparison against neural network states in one dimension was done by Dr. Mario Collura, while I developed the Tree Tensor Network software used for the comparison in the two-dimensional Heisenberg model. Accordingly, Dr. Mario Collura performed the numerical study in 1D, while I performed the numerical simulations for the 2D scenario.

Dr. Mario Collura and Prof. Dr. Luca Dell'Anna wrote most of the theoretical *section 2* introducing the "constrained" Matrix Product State and highlighting the differences between RBM and uRBM. Dr. Mario Collura further wrote in *section 3* the numerical analysis for the one-dimensional comparison (*section 3.1*) while I strongly contributed to the writing for the two-dimensional comparison (*section 3.2*). The Introduction and conclusion were written with the mutual contribution of all authors.

Lattice Quantum Electrodynamics in (3+1)-dimensions at finite density with Tensor Networks   accepted for publication in *Nature Communications* [144] - Incorporated as *Appendix A*

Following up on our publication *Two-Dimensional Quantum-Link Lattice Quantum Electrodynamics at Finite Density* dealing with lattice gauge theories, we here applied a Tensor Network to study lattice gauge theories in three spatial dimensions.

Prof. Dr. Simone Montangero initialised the project and supervised the execution. Dr. Pietro Silvi provided the theoretical strategy of mapping the lattice gauge theory into the proper computational basis. Dr. Giuseppe Magnifico applied this strategy for the three-dimensional study while I developed the Tensor Network software used to solve the underlying numerical optimisation problem. The numerical analysis was executed by Dr. Giuseppe

Magnifico with me in a supporting role on how to best apply the Tensor Network software. The interpretation of the results was mainly done by Dr. Giuseppe Magnifico, Dr. Pietro Silvi and Prof. Dr. Simone Montangero.

Dr. Giuseppe Magnifico prepared the first draft of the manuscript. All of the authors contributed to refining the manuscript by reviewing and editing accordingly. The Introduction and conclusion were written in mutual agreement with Dr. Pietro Silvi, Dr. Giuseppe Magnifico and Prof. Simone Montangero as main contributors.

Optimizing Radiotherapy Plans for Cancer Treatment with Tensor Networks [99]- Incorporated as *Appendix B*

In this manuscript, we applied Tensor Networks in Intensity-Modulated Radiation Therapy, a particular medical treatment technique for curing cancer. Concretely, we optimised the underlying radiotherapy plans for this therapy which included two main challenges: *(i)* Mapping the radiotherapy optimisation problem to a many-body system, and *(ii)* numerically finding the ground state of this system with Tensor Networks.

Prof. Dr. Simone Montangero, Dr. Marta Paiusco and Samuele Cavinato proposed the project. Dr. Marta Paiusco and Dr. Marco Fusella defined the underlying optimisation problem while all authors worked out the conceptual strategy on how to deploy Tensor Networks for this particular application together. Samuele Cavinato developed the theoretical strategy of mapping the radiotherapy optimisation problem to a many-body system. I developed the Tensor Network software used for the numerical study. The numerical analysis was mainly executed by Samuele Cavinato with myself supporting the application of the Tensor Network software while all the authors contributed to the interpretation of the results.

The manuscript was initialised and structured by myself while Samuele Cavinato mainly wrote the content of *section II.* (Describing the radiotherapy optimisation problem), *section III.A* (mapping the problem) and *section IV.* (numerical analysis); whereas I wrote the section *III.B* (solving the mapped problem via Tensor Networks). All of the authors contributed to refining the manuscript by reviewing and editing accordingly. The Introduction and conclusion were written with the mutual contribution of all authors.

Acknowledgement of content originality

My own contributions to the above-listed publications have partly been reused in the discussions of the thesis including the introduction and the conclusion. Whenever major parts are directly taken from one of the publications, it is indicated accordingly. In particular, Chapt. 5 includes an in-depth discussion of the therein presented paper and aims to provide a deeper understanding of the developed method. Therefore, the section of this discussion (Sec. 5.2) has to be understood as an addition to the original publication and, for sake of clarity for the reader, includes selected parts and material of the original publication.

Chapt. 2 introduces the theoretical foundations of the field of Tensor Networks. Since this topic has already been widely discussed in such an introductory manner by previous publications, this chapter is inspired by the References [24, 77, 108, 172].

Declaration of original authorship

I hereby declare that this dissertation is my own original work except where otherwise indicated. All data or concepts drawn directly or indirectly from other sources have been correctly acknowledged to the best of my abilities. This dissertation has not been submitted

in its present or similar form to any other academic institution either in Germany, Italy or abroad for the award of any other degree.

_____
Place, Date   (Signature)

Abstract

This thesis presents the development of a numerical simulation technique, the *Tree Tensor Network*, aiming to overcome current limitations in the simulation of two- and higher-dimensional quantum many-body systems. The development and application of methods based on *Tensor Networks* (TNs) for such systems are one of the most relevant challenges of the current decade with the potential to promote research and technologies in a broad range of fields ranging from condensed matter physics, high-energy physics, and quantum chemistry to quantum computation and quantum simulation. The particular challenge for TNs is the combination of accuracy and scalability which to date are only met for one-dimensional systems by other established TN techniques. This thesis first describes the interdisciplinary field of TN by combining mathematical modelling, computational science, and quantum information before it illustrates the limitations of standard TN techniques in higher-dimensional cases. Following a description of the newly developed Tree Tensor Network (TTN), the thesis then presents its application to study a lattice gauge theory approximating the low-energy behaviour of quantum electrodynamics, demonstrating the successful applicability of TTNs for high-dimensional gauge theories. Subsequently, a novel TN is introduced augmenting the TTN for efficient simulations of high-dimensional systems. Along the way, the TTN is applied to problems from various fields ranging from low-energy to high-energy up to medical physics.

Kurzbeschreibung

In dieser Arbeit wird die Entwicklung einer numerischen Simulationstechnik, dem *Tree Tensor Network* (TTN), vorgestellt, die darauf abzielt, die derzeitigen Limitationen bei der Simulation von zwei- und höherdimensionalen Quanten-Vielteilchensystemen zu überwinden. Die Weiterentwicklung von auf *Tensor-Netzwerken* (TN) basierenden Methoden für solche Systeme ist eine der aktuellsten und relevantesten Herausforderungen. Sie birgt das Potential, Forschung und Technologien in einem breiten Spektrum zu fördern, welches sich von der Physik der kondensierten Materie, der Hochenergiephysik und der Quantenchemie bis hin zur Quantenberechnung und Quantensimulation erstreckt. Die besondere Herausforderung für TN ist die Kombination von Genauigkeit und Skalierbarkeit, die bisher nur für eindimensionale Systeme erfüllt wird. Diese Arbeit beschreibt zunächst das interdisziplinäre Gebiet der TN als eine Kombination von mathematischer Modellierung, Computational Science und Quanteninformation, um dann die Grenzen der Standard-TN-Techniken in höherdimensionalen Fällen aufzuzeigen. Nach einer Beschreibung des neu entwickelten TTN stellt die Arbeit dessen Anwendung zur Untersuchung einer Gittereichtheorie vor, die das Niederenergieverhalten der Quantenelektrodynamik approximiert und somit die erfolgreiche Anwendbarkeit von TTNs für hochdimensionale Eichtheorien demonstriert. Anschließend wird ein neuartiges TN eingeführt, welches das TTN für effiziente Simulationen hochdimensionaler Systeme erweitert. Zusätzlich wird das TTN auf diverse Probleme angewandt, die von Niederenergie- über Hochenergie- bis hin zur medizinischen Physik reichen.
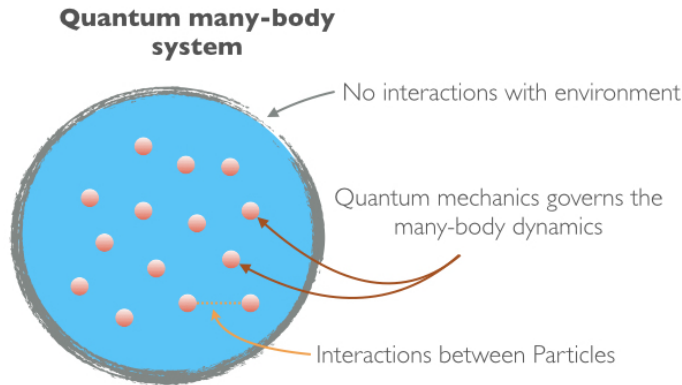
Contents

Quantum many-body system

No interactions with environment

Quantum mechanics governs the many-body dynamics

Interactions between Particles

1

Introduction

Quantum physics describes the core theory behind our understanding of nature and recently gave rise to various technologies [1, 2]. The understanding of quantum many-body systems plays a key role in unveiling novel insights in a broad range of modern research, since many fundamental phenomena only occur for a large number of interacting particles [3]. The underlying quantum physics encompasses phenomena in Condensed Matter Physics, such as high-temperature superconductivity [4, 5], governs applications in Quantum Information [6, 7] and forms the foundation in High Energy Physics to promote our understanding of the fundamental elements of our universe [8, 9]. Furthermore, quantum many-body problems occur in applied science, among others when investigating electronic structures in molecular systems [10, 11], in the development of quantum technologies, such as quantum computers, quantum sensors, or quantum simulators [12–15], and in material science for studying and engineering material properties [16].

Consequently, it is fundamental to have theoretical and numerical methods for describing such systems accurately, not only to advance numerous research areas in science but further to enable the development of novel cutting-edge quantum technologies impacting our daily life [17]. However, the investigation progress in the field of quantum many-body physics is often hindered by the lack of general analytical solutions. Especially strong interacting systems, where perturbation theory fails, seek a stable and scalable numerical method to unveil the rationale behind these complex quantum systems.

In the last decades, Tensor Network methods have proven to successfully solve quantum many-body problems on a classical computer and rapidly became the state-of-the-art numerical tool for one-dimensional (1D) systems. They offer an efficient representation of a quantum many-body state by taking advantage of the underlying entanglement properties of the system. Crucially, in this Tensor Network representation, the amount of information can be controllably compressed regarding the available computation resources towards an accurate approximation [18]. At the current state of research, this class of numerical methods is especially relevant for problems where competing methods fail, such as *Quantum Monte Carlo* methods which are limited by the infamous *sign-problem* [19].

However, most two-dimensional problems still pose a non-trivial challenge for current Tensor Network simulations, especially when operating on standard computational resources [20–24]. The realisation of efficient methods for the simulation of two- or even higher-dimensional quantum many-body systems is not only a highly challenging but as well crucially important task for numerous scientific cases: Among others, such methods can be used to benchmark future quantum computer algorithms, or quantum simulators [25–33]. Moreover, their realisation enables the observation of topologically ordered phases [34], the study of the quantum Hall effect [35], and the classification of two-dimensional phases of matter based on the entanglement entropy [36], all of which being topics at the heart of the research in statistical and condensed matter physics. Consequently, the study of high-dimensional quantum systems could open up new fields of research, for instance in the context of the numerical implementation of Lattice Gauge Theories at finite density [37–39] or open quantum systems [40, 41], giving access to a huge amount of very rich and promising physics almost unexplored today.

Thus, by continuing the development of Tree Tensor Networks for high-dimensional systems, this thesis aims to bridge the existing gap in the investigation of quantum many-body physics towards systems that are currently out of reach. The first concrete applications of the developed algorithm are in the field of Lattice Gauge Theories [37], the discretised version of quantum field theory, which represents a fundamental theoretical tool for our understanding of the microscopic processes governing the dynamics of elementary particles (see Chapt. 4). In fact, many of the collective phenomena arising in these theories, including the phase diagram, have not been fully characterised yet [37, 42]. This is especially the case for Lattice Gauge Theories in higher spatial dimensions and at finite charge density where the sign problem prevents the use of Monte Carlo methods [19, 39].

Furthermore, in the course of this thesis, we introduced a novel Tensor Network geometry based on a Tree Tensor Network which combines accuracy and scalability for simulations of high-dimensional systems [23]: This *augmented Tree Tensor Network* (aTTN) enables to reach larger system sizes with more degrees of freedom while maintaining a precise description of the systems. Consequently, the aTTN overcomes the current limitation of Tensor Network methods and therefore provides a powerful pathway to boost the research in all aforementioned fields of physics. Finaly, this direction of development can potentially have a cutting-edge impact on the theoretical understanding of the quantum nature of complex many-body systems and, at the same time, support the experimental realization of physical devices in different technological fields.

On the shoulders of Quantum many-body systems

In its most original form, the quantum many-body problem describes the dynamics of electrons in matter as a continuum problem in three dimensions with strongly interacting particles. From a mathematical point-of-view, we know the theory of everything behind this problem since we simply have to solve the corresponding Schroedinger Equation $\mathcal{H}|\psi\rangle = E|\psi\rangle$[1]. In other words, we *just* need to solve an underlying eigenvalue problem while all properties of the quantum system are enclosed in the eigenstates $|\psi\rangle$ of the systems Hamiltonian $\mathcal{H}$. However, the Hilbert space $\mathscr{H}$ of the system and thereby the underlying eigenvalue problem grows exponentially with system size: As an example, for N interacting spin-$\frac{1}{2}$ particles, the number of possible basis states is 2^N. Thus, for a system with $N > 130$ sites, the number of amplitudes required to describe the complete wave function is greater than the number of atoms in the known universe (estimated in the order $\mathcal{O}(10^{80})$). So we could use all atoms

---

[1] or $i\partial_t|\psi\rangle = \mathcal{H}|\psi\rangle$ accordingly for dynamical problems

of the complete universe to store the information of $N = 130$ sites, but we would already need two universes for $N = 131$ sites. Looking at a *small* physcial system with $N \sim 10^{23}$ particles (which is the order of the Avogadro constant), the total state space would be of the order $\dim\{\mathcal{H}\} = \mathcal{O}(2^{10^{23}})$ which is even exponentially larger than the number of atoms in the universe. Thus, to rephrase the problem in the words of Paul Dirac: "The underlying physical laws necessary for the mathematical theory [are] completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble" [43].

Basically, we know exactly what to do analytically, but we do not have a simple way to access the information directly due to the, in theory, exponentially growing Hilbert space. But can nature itself really work this way and store such exponentially large information in a single quantum many-body system? Or in other words, are the amplitudes of a general, realistic many-body wave-function all different or is there, in reality, some kind of simplifying structure that makes the wave-function a feasible concept to work with? And in fact, in the last 30 years, a physical phenomenon has increasingly been linked as the key to understanding an underlying structure of quantum many-body wave-functions: quantum entanglement [44–48].

Now, knowing that there is an underlying structure, could we just try to guess the wave-function? Even though this might seem far-fetched, an educated guess can indeed work: In this way, Laughlin famously found ground-state and excited-state wave functions describing the condensation of a two-dimensional electron gas [49]. More sophisticated methods, however, include techniques which rely on perturbation theory, such as *series expansion techniques*. Still, this analytical approach famously works well only for weakly-coupled systems and breakes down for stronger interactions which are typical for complex quantum many-body systems. Another well-studied approach for such systems is the *Mean field theory* which can be successfully deployed for certain strongly interacting quantum systems but fails to faithfully represent the quantum correlations in the systems. *Quantum Monte Carlo*, on the other hand, has established itself early on as one of the most successful methods without the draw-backs of all the above-mentioned techniques: By stochastically reconstructing the partition function of the system, it can accurately describe the quantum properties of a system making it today an extremely successful tool for a broad scale of applications [50–53]. Nevertheless, this technique faces severe difficulties when analysing systems for which the stochastic distribution to be sampled over becomes negative, or even complex, known as the *sign-problem* [19]. At this point, Tensor Networks offer a fundamental tool to faithfully represent the information and by capturing the underlying structure of a quantum many-body state.

A brief history of Tensor Networks

The history of Tensors in theoretical physics dates back to 1898 in the mathematical context of researching the physical properties of crystals [54]. However, their application as variational methods in numerical analysis has started way after computers have established themselves as a workhorse for complex computations. In fact, in the early stages of complex numerical analysis in classical physics, Monte Carlo methods have established themselves as the main simulation tool. These methods have been well developed early on and still to date present one of the most efficient numerical methods available. However, as mentioned above, they suffer from the sign problem under certain conditions which makes the method diverge, and thus are infeasible, for simulating some of the highly relevant quantum systems [19].

A particular key development for Tensor Networks was the *Density Matrix Renormalisation Group* (DMRG) developed by S. White in the early 90s [55, 56]. In its introduction,

this technique was not yet brought into the context of Tensors but rather reformulates the established *numerical renormalisation group* method [57, 58] towards an efficient application on quantum many-body systems: Instead of performing the renormalisation procedure in real space, White performed it in the density matrix eigenspace. This change of space enabled to efficiently identify the *physically* relevant degrees of freedom in the renormalisation procedure for computing the low-energy eigenstates of one-dimensional Hamiltonians. Thereby, White laid the foundation of Tensor Networks, as he demonstrated that the relevant information for representing quantum many-body states is to be found in the degrees of freedom corresponding to the *entanglement* of the quantum wave-function. Eventually, it was unveiled that all states obtained by the DMRG algorithm are so-called *finitely-correlated states*, also referred to as Matrix-Product-States (MPS), which can be represented by a set of connected Tensors [44, 59, 60]. Thus, the fundamental idea of Tensor Networks was built: They parameterise the wavefunction of interest by recasting the system's degrees of freedom as a Network of several connected Tensors (see Fig. 1.1). In fact, in this framework, it became evident that the DMRG can be understood as a variational optimisation over an MPS [59, 61]. The success of the DMRG was a breakthrough in the analysis of quantum systems and rapidly became the state-of-the-art tool for low-energy states of one-dimensional systems.

The second boom of development came in the mid-2000s: Further investigations in quantum theory helped to increasingly understand that there was a certain structure of the entanglement. It was proven that gapped one-dimensional systems obey the so-called *area law* and that this exact entanglement structure is naturally encoded in MPS [48]. On the other hand, the MPS representation allowed to expand the available techniques from finding low-energy states in equilibrium to concepts for out-of-equilibrium evolution [62, 63] and finite temperatures [64, 65]. Further, novel Tensor Network structures, such as the *Multi-scale Entanglement Renormalisation Ansatz* (MERA) [66] and the *Tree Tensor Network* (TTN) [67, 68], have been developed addressing different problems ranging from the analysis of critical systems [66, 68] to systems with periodic boundary conditions [69, 70]. In this wave of development, the natural desire for an equivalently efficient representation of quantum many-body states in two-dimensions arose which gave rise to the *Projected Entangled Pair States* (PEPS) [71]. Further, the established network geometries, i.e. MPS, MERA and TTN, have been applied to two dimensional systems with various degrees of success [22, 72, 73]. However, as we will see in more detail within this thesis, the developments of Tensor Network methods is still ongoing for two- or even higher-dimensional quantum systems, since combining of accuracy and scalability in simulating such high-dimensional systems represents an open challenge for Tensor Networks [21, 24].

Tensor Networks in a Nutshell

The concept behind Tensor Network is a very general yet simple idea: As illustrated in Fig. 1.1, they decompose a large tensor into a set of smaller tensors that are connected over some auxiliary indices being summed over, called *bond-links*. The number of bond-links that are left unsummed denotes the indices of the original tensor. In this representation, the bond-links of the network can be upper-bounded by a controllable *bond-dimension* m. Increasing m allows to control the amount of information captured in the Tensor Network. Thus, formally both, the complete tensor and the Tensor Network, live in the same space while the latter just occupies a corner of that space. Representing the information in such a Tensor Network gives raise to major benefits, such as an exponential reduction in memory, an exponential speedup of
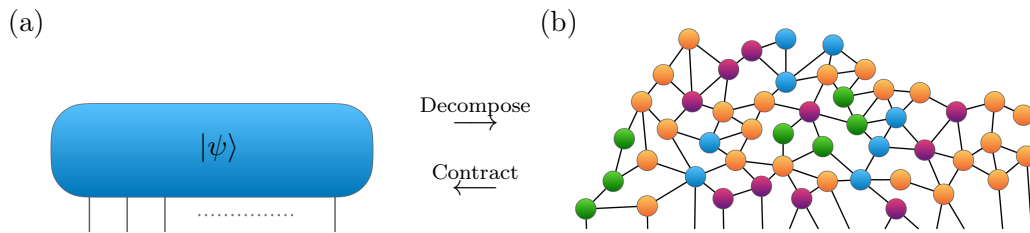
FIGURE 1.1: (a) High-order tensor $|\psi\rangle$ describing a quantum many-body wavefunction exactly with all exponentially many coefficients. (b) A possible Tensor Network decomposition of the same wavefunction where the tensors are illustrated in different colors and connected with each other over *internal bond-links*. The dimension of the bond-links can be control to compress the information represented within the network.

computations, a theoretical insight and interpretation, an estimation of missing or corrupted entries and many optimisation algorithms and strategies. Consequently, Tensor Networks are, in general, an efficient linear algebra tool for problems in exponentially high-dimensional spaces.

In quantum many-body physics, Tensor Networks are used to efficiently represent the exponentially large wavefunction by decomposing the complete quantum state vector. In this representation, the bond-dimension m can be seen as a control parameter which allows to interpolate between a product state ($m = 1$) and the exact, but inefficient, representation. Based on the exact decomposition of the complete, exponentially large state, different network geometries can be defined. In a nutshell, the Matrix Product States (MPS) are the established Tensor Network geometry for equilibrium and, in many cases, even for out-of-equilibrium problems of one-dimensional systems while MERA provides an efficient representation for critical systems [66, 74–76]. The most prominent Tensor Network representations for higher dimensions are at the current stage the Projected Entangled Pair States (PEPS) [45, 71, 77–79] for two-dimensional systems, and the Tree Tensor Networks (TTN) [22, 24, 35, 37, 67, 68, 80] which, in principle, can be defined for any system dimensionality (see Sec. 2.2.2 for more details).

Even though Tensor Networks can efficiently represent a quantum many-body state, the underlying simulations still rapidly reach the computational limits of current classical machines, especially for highly entangled systems. One powerful way to drastically reduce the required numerical resources can be achieved by exploiting the symmetries of the underlying quantum system. Thus, the incorporation of symmetries has become a standard add-on for many Tensor Network algorithms [59, 81] in order to significantly push forward the capabilities of numerical methods for studying quantum many-body systems (see Sec. 2.3). Consequently, the implementation and development of Tensor Networks has become an interdisciplinary field combining elements from mathematics, such as linear algebra and group theory, with High-Performance computing and Quantum Information.

Black Holes and Baby Universes and Tensor Networks beyond quantum systems

By now, you might have noticed that all section titles in this introduction are variations of books by Stephen Hawking and reading this title you might wonder what Tensor Networks have to do with Black Holes and Baby Universes. While Tensor Networks have originally mostly been applied to problems between Quantum Information and Condensed Matter
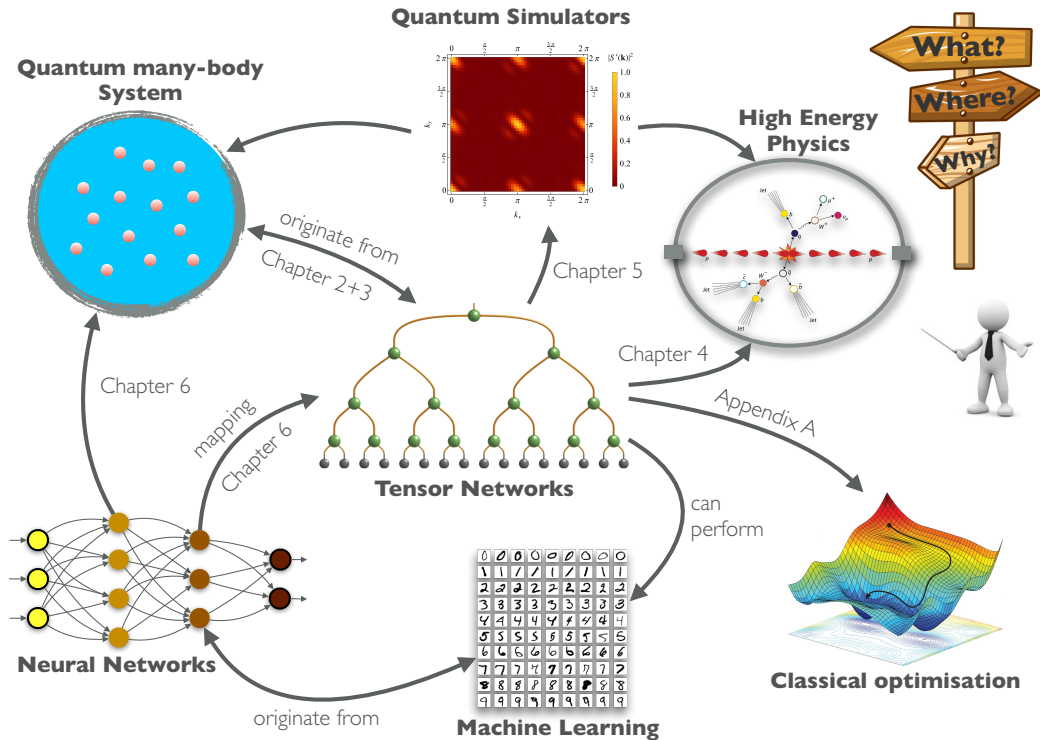
FIGURE 1.2: Overview of some of the interdisciplinary applications of Tensor Networks in Quantum Physics and beyond which are touched on in this thesis. Subpicture *High Energy Physics* adapted from Ref. [94], *Classical Optimisation* taken from Ref. [95] and *Machine Learning* adapted from Ref. [96].

Physics, the applications of Tensor Networks have expanded into numerous scientific fields of research in recent years. In Quantum Gravity, for instance, the MERA has indeed been linked to the geometry of space via the Anti-de Sitter (AdS) / Conformal Field Theory (CFT) correspondence [75, 82] which in return has been brought in context to black holes by, to some extend, resolving the black hole information paradox [83, 84]. Since Tensor Networks are a general tool for information processing based on correlations and entanglement, they can essentially flourish for any problem where a certain structure of correlations is present. Thus, it is no wonder that since the late 2000s, Tensor Networks have become increasingly of interest for the applied math community as well [85–87]. In this community, the idea of Matrix Product States was rediscovered as *Tensor Trains* [88] and Tree Tensor Networks have been proposed as *Hirarchical Tucker Trees* [89], thus highlighting the potential of Tensor Networks as a general tool for information processing. This potential makes Tensor Networks further able to perform Machine Learning tasks as a novel competitor to state-of-the-art neural networks (NN) [90, 91]. First approaches of Machine Learning with Tensor Networks already yield comparable results when performing supervised learning on standardised datasets [91–93]. Consequently, in the last decades, the development and application of Tensor Networks has rapidly evolved into an interdisciplinary field of research.

In fact, beyond developing a novel Tensor Network for high-dimensional quantum many-body systems, this thesis showcases some of these unconventional, interdisciplinary applications of Tree Tensor Networks. In particular, in the course of this thesis, the developed TTN

was used to perform simulations to solve classical optimisation problems in medical research, to analyse a system of Rydberg-atoms in the pursuit of benchmarking novel quantum technologies and to investigate Lattice Gauge systems in High-Energy physics aiming to provide a new step in our understanding of the universe. Fig. 1.2 provides a brief overview of all applications discussed in the thesis.

Brief answers to big questions - Thesis outline

The outline of this thesis is presented as a set of questions, one for each chapter and one question for the complete thesis:

Are Tensor Networks able to efficiently simulate quantum many-body systems in high-dimensions on a classical computer? The main goal of this thesis was the development of a Tensor Network algorithm capable to analyse quantum many body-systems on a classical computer in any spacial dimension. Therefore, a TTN has been implemented for simulating such quantum systems and successfully applied to high-dimensional problems in quantum physics and further to selected problems beyond quantum physics. Moreover, this Tensor Network geometry has a clear bottleneck as it fails to capture the area law for high-dimensional systems with increasing system sizes. Thus, as a result of this research, an *augmentation* of the underlying TTN, the aTTN, has been developed to address this lack of scalability of the network in higher dimensions. This novel Tensor Network opens the pathway to investigating large-scale quantum systems, enabling to promote numerous fields of physics, including quantum technologies, quantum field theories, and condensed matter.

The following provides a guidance on the structure of the thesis going from the theoretical background over the obtained results and even up to the applications on problems beyond quantum physics.

Chapter 2: What makes a Tensor Network a *good* Tensor Network state? The first chapter provides an overview of Tensor Networks, starting from the theoretical introduction of tensors and their operations as the fundamental building blocks of every Tensor Network in Sec. 2.1. By formally introducing Tensor Network states and algorithms in Sec. 2.2, we will see that a *good* Tensor Network shall satisfy the same entanglement bounds under real-space bi-partitions as the physical states they represent. Further, in this section, it is illustrated how the numerical complexity $\zeta = \mathcal{O}(poly(N)poly(m))$ naturally scales polynomial in the system size N and the network's bond-dimension m. However, by introducing the different Tensor Network geometries, we will see that finding a proper balance between sufficiently satisfying the entanglement bounds and keeping the numerical complexity feasibly low poses in practice a highly non-trivial challenge for systems in higher dimensions.

The closing section (Sec. 2.3) of this chapter gives additional insights into how to further decrease the computational costs by incorporating symmetries in the Tensor Network state. After a description of typical symmetries in quantum many-body states, and an explanation of how a Tensor Network is affected by the incorporation of such symmetries, we will see that exploiting symmetries can not only drastically improve the Tensor Network computations but further enables to target specific global symmetry sectors in an analysis, thus enabling to easily fix for instance the total charge of $U(1)$ symmetric model.

Chapter 3: What are Tree Tensor Networks and how do they work? Having gained an overview on Tensor Networks, this subsequential chapter formally and technically introduces Tree Tensor Networks and their implementation for high-dimensional quantum systems.

Starting from the reason why this thesis addresses the quantum many-body problem with a TTN, instead of alternative methods, Sec. 3.1 presents the definition and fundamental idea of TTNs. It illustrates that the natural loop-less, hierarchical structure of a TTN on the one hand greatly benefits the underlying algorithms, such as finding the ground state of a quantum many-body system, and on the other hand is the key for the unique flexibility of TTNs to be applied to various problems in any spacial dimension.

In Sec. 3.2, we will see that the concrete structure of a TTN is crucial for the analysis of systems in two spatial dimensions and beyond, which is defined by a technically non-trivial mapping problem. As a hands-on summary regarding the ground-state search with a TTN, the chapter provides a step-by-step example for a (4×4)-system. Finally, with Sec. 3.3, the chapter provides an in-depth description of the technical part of implementing TTNs for high-performance computations giving some hands-on illustration.

Chapter 4: Why are Tree Tensor Networks interesting in higher dimensions? It is reasonable to wonder why one should be interested in TTNs compared to other numerical methods, such as QMC, or even compared to competing Tensor Networks, like the PEPS. To answer this question, Chapt. 4 consists of the publication [37] which illustrates the application of the TTN to study Lattice Gauge Theories at finite density in two dimensions. In fact, this study belongs to a set of fundamental problems in which QMC simulations suffer from the sign problem, alongside quantum chromodynamics, quantum electrodynamics at finite density, and fermionic systems. Furthermore, the study of these systems with alternative Tensor Network methods, such as PEPS, poses a non-trivial challenge due to an underlying, reasonably large local Hilbert space which can streightforwardly be addressed by the natural TTN structure. Thus, this chapter shows the potential of TTNs to promote the understanding of fundamental elements of the universe in future research of High Energy Physics. Finally, it illustrates that this potential, together with the possibility to analyse systems unfeasible for QMC, makes a TTN a highly promising candidate for the analysis of general high-dimensional quantum many-body systems.

Chapter 5: Is it possible to capture area law with Tree Tensor Networks? Having gained insights into TTNs in Chapt. 3 and having presented first applications in Chapt. 4, it becomes evident that their main bottleneck is the fact that they indeed fail to incorporate the area law by design. Thus, with increasing system size and dimensionality, they eventually fail to faithfully represent a quantum many-body state. Therefore, the question naturally arises whether it is possible to capture area law in higher dimensions with TTNs by its structure?

To address this drawback, Chapt. 5 introduces the aforementioned *augmentation* of the TTN structure, the aTTN, which encodes the area law by construction. Thus, this aTTN offers better scalability and can help to overcome fundamental limitations of TTNs in high dimensional systems. In particular, this chapter includes the publication [23] in which we show that this novel technique outperforms competing methods and reaches unprecedented system sizes in two-dimensions making it the ideal candidate to, among other applications, benchmark and validate state-of-the-art quantum technologies, such as quantum simulators or quantum computers. In addition to the publication, this chapter features a more in-depth technical presentation of the novel aTTN.

Chapter 6: Can we compare Neural Networks and Tensor Networks in quantum physics? In this introduction, Tensor Networks have already been compared to Neural Networks as alternative method for performing Machine Learning tasks and, in return, Neural Networks can also be used as a valid alternative for representing a quantum many-body state [97]. Thus, a scientific comparison can bring light to the fundamental connection between the two variational methods.

Chapt. 6 consists of the publication [98] which aims to find a fair comparison between Tensor Networks and Neural Network states as a representation of quantum many-body wave functions. Along the way, it introduces a novel mapping that illustrates a mathematical connection between the networks, before analysing one-dimensional and two-dimensional critical systems with different Tensor Network approaches and Neural Network approaches. Finally, we point out that comparing quantum state representations in higher dimensions, in general, is a non-trivial task seeking profound measures for a fair comparison.

Appendix A: What about Tree Tensor Networks for three dimensional systems? In fact, the aforementioned chapters present numerical results for two-dimensional systems and include only theoretical discussion on the extension towards three-dimensional systems. However, App. A presents a manuscript that illustrates the numerical application of the developed TTN on three-dimensional systems in the field of Lattice Gauge Theory. Thus, this manuscript extends the publication presented in Chapt. 4 and, in fact, provides the first successful attempt to simulate a Lattice Gauge Theory in $(3+1)$ dimensions via Tensor Networks. Consequently, it shows that Tensor Networks, and in particular Tree Tensor Networks, can be a fruitful direction to solve fundamental long-standing problems not only in two- but even in higher dimensions.

Appendix B: So, Tree Tensor Networks seem great, but can we cure cancer with them? Well, not yet; However, App. B presents the publication [99] showing an approach on how Tensor Networks, in particular the TTN, can be deployed to optimise radiotherapy plans for cancer treatment. It provides a successful proof-of-principle for this approach including a realistic anatomical scenario simulating prostate cancer. Along the road, it defines a clear strategy to map the classical problem to simulated quantum-like hardware. Since Tensor Networks are particular examples of quantum circuits [100], this study might open the way to the application of quantum computation to cancer treatment, for example through the application of hybrid quantum-classical optimization algorithms [101].

$|\psi\rangle \overset{?}{=}$
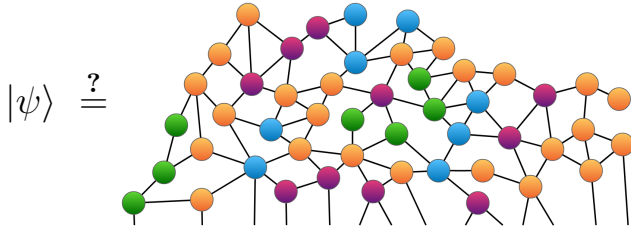
2

Tensor Networks: An overview

Tensor Networks have become increasingly popular as a numerical tool to address quantum many-body problems on classical computers [21, 24, 37, 38, 61, 102–105]. The fundamental idea in these methods is to efficiently parametrise the wavefunction of the physical system in the form of a network consisting of several interconnected tensors. Thus, as illustrated in Fig. 2.1, the wavefunction is decomposed into a set of local tensors: Like the DNA carries the genetic instructions - or information if you wish - of living organisms and like neurons constitute the main components of nervous systems such as our brain, tensors are the fundamental building blocks for representing a quantum state via Tensor Networks. Accordingly, a Tensor Network is composed out of local tensors which are connected to each other and hold the information encoded in the complete network. This construction crucially reduces the total number of parameters from an exponential $\mathcal{O}(e^N)$ to a polynomial $\mathcal{O}(poly(N))$ dependency on the system size N. In practice, each bond-link connecting two tensors defines the structure of the underlying entanglement in the quantum state where its bond-dimension m can be a quantitative measure of the amount of the present entanglement. Thus, the entanglement within the quantum system *glues* the connected tensors within the network together.

For the sake of compactness, a graphical representation for the underlying mathematical tensor notation has been established during the introduction of Tensor Networks in physics. In a nutshell, Tensor Networks are represented by circles (or similarly closed objects) for the tensors and lines for the links (see Fig. 2.1). Each link connecting two different tensors indicates a *contraction* over a coinciding index. Here, the tensor elements are depicted in arbitrary colours, while the colour for each tensor is usually adapted for highlighting purposes. This graphical representation allows representing complex equations of the underlying tensor algebra in visual diagrams. Once you get used to this notation, it surely is more convenient and intuitive than explicitly writing down the underlying rigorously long equations.

The implementation and development of Tensor Networks for solving cutting edge problems in research have become an interdisciplinary field combining elements from mathematics,
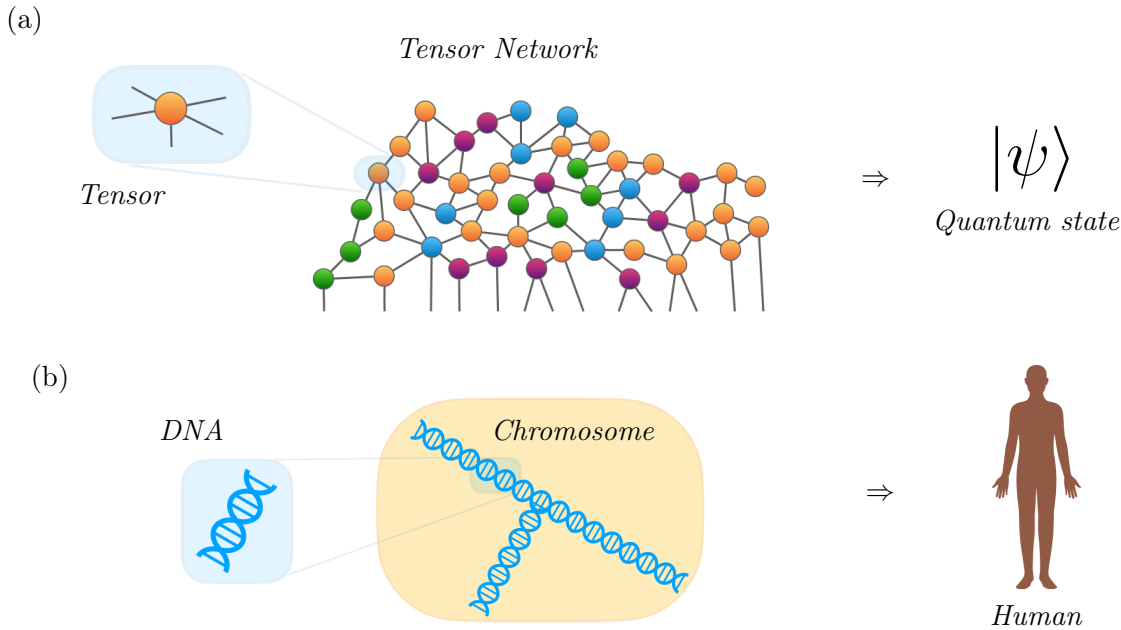
(a)

(b)

FIGURE 2.1:   (a) Graphical representation of an arbitrary Tensor Network state: A tensor (left inlet) is the fundamental building block for representing the quantum many-body state $|\psi\rangle$ via Tensor Networks. Each tensor is typically represented graphically by a circle-like object with several links which can be open or connected to other tensors and contains a certain amount of information about the state. (b) The DNA similarly constitutes the fundamental building stone of a chromosome, and ultimately of a human.

such as linear algebra and graph theory, with High-Performance computing and Quantum Information. In this context, the exploitation of symmetries has become a crucial, state-of-the-art add-on for many Tensor Network algorithms that are deployed nowadays to investigate quantum many-body Systems [59, 81]. This incorporation of symmetries reduces the numerical resources needed and thereby helps to push forward the limitations of simulating quantum systems on classical computers via Tensor Networks. For global symmetries, this can be realised independent of the network geometry by decomposing each tensor into smaller, independent subspaces [106, 107]. Furthermore, such symmetrically invariant Tensor Networks offer to efficiently investigate physical properties regarding certain selected symmetry sectors as towards the end of the chapter.

This chapter introduces selected topics on Tensor Network methods crucial for their theoretical understanding and further aims to grant a broad overview of the state-of-the-art Tensor Network states and algorithms. In particular, the chapter is divided into three main sections: The first section (Sec. 2.1) formally introduces the mathematical concepts behind Tensors and the important operations for Tensor Network algorithms. Afterwards, Sec. 2.2 is dedicated to the description of general Tensor Network states and algorithms. Therein, Tensor Networks and their mathematical properties are formally defined followed by a discussion on the importance for efficient Tensor Networks to *(i)* satisfy the same area-law of the quantum states they represent and *(ii)* to be efficiently contractable at the same time. Further, this section provides a general overview of different Tensor Network states and common Tensor Network algorithms. The last section (Sec. 2.3.1) illustrated the benefits of exploiting symmetries in Tensor Network simulations and presents how symmetries can be incorporated in Tensor Networks rounding up this introduction to Tensor Networks.

2.1 Tensors and Tensor Operations

In the application of Tensor Network algorithms, the Tensors making up the network are constantly manipulated in order to achieve the goal of the algorithm, such as finding the ground state of a Hamiltonian or investigating the dynamics of a quantum system. The most important numerical operations thereby are the *contraction* of two tensors, *reshaping* a tensor or performing a *factorisation*. In a sophisticated algorithm, these operations are executed several thousand or even millions of times. Thus, these operation form the low-level toolbox for Tensor Network algorithms and their numerical efficiency is key to a successful application of Tensor Networks.

This section formally introduces the tensor object together with a selected set of these operations which serve as the low-level toolbox of all Tensor Network algorithms. In practice, the tensors of a Tensor Network are complex (or in some cases real) numerical data stored in a computer memory device and manipulated on a processing unit. In the following, the mathematical tensor object is formally and graphically introduced followed by the basic operations commonly performed on these tensors.

It is worth mentioning that Ref. [24] and Ref. [108] offer a similar in-depth description with the former focusing in particular on the practical implementation of tensors and their operations.

2.1.1 Tensors

From a mathematical point of view, a single tensor $\mathcal{T}$ is a multidimensional array describing linear relations of several different vector spaces. It can be seen as the generalisation of the well-known vectors and matrices from linear algebra: While a vector is an element of a single vector spaces $\mathbb{V} \in \mathbb{C}^d$ and a matrix is assigned to the product space of two vector spaces $\mathbb{V}_1 \times \mathbb{V}_2$, a tensor represents information in a general space composed out of arbitrarily many vector spaces $\mathbb{V}_1 \times ... \times \mathbb{V}_r$.

Definition 2.1.1: *(Tensor).*

A complex tensor $\mathcal{T} \in \mathbb{C}^{d_1 \times d_2 \times ... \times d_R}$ is defined as a dense R-dimensional array of complex numbers mapping the R-dimensional manifold to a complex scalar:

$$\mathcal{T} : \mathbb{V}_1 \times ... \times \mathbb{V}_R \to \mathbb{C} \ .$$

Remarks.

- We also refer to $d_1 \times d_2 \times ... \times d_R$ as the *shape* of a tensor.

- The total number of vector spaces R of the tensor $\mathcal{T}$ is its *order*.

- Each element $\mathcal{T}_{i_1,...i_R}$ of the tensor can be addressed by a set of indices $i_r \in \{1, ..., d_r\}$ assigned to vector spaces $\mathbb{V}_r$ ($r \in \{1, ..., R\}$). The range of values $\{1, ..., d_r\}$ which the index i_r may take is defined as the *link* ν_r of the tensor assigned to the vector space $\mathbb{V}_r$.

- To simplify the notation, we often use the same notation exchangeable for the link ν_r and the corresponding index i_r within the same vector space $\mathbb{V}_r$; whether it is referred to the entire link or individual entries is easily distinguishable by
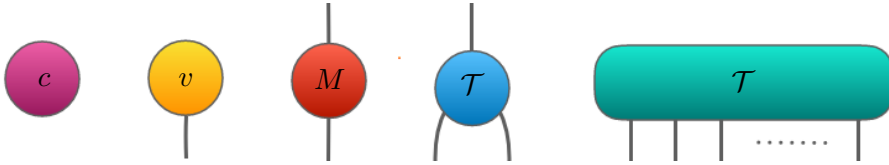
FIGURE 2.2: Graphical representation of (a) a scalar $c \in \mathbb{C}$ (i.e. zero-link tensor), (b) a one-link tensor (i.e. a vector) $v \in \mathbb{C}^{d_1}$, (c) a matrix $M \in \mathbb{C}^{d_1 \times d_2}$ (i.e. a two-leg tensor), (d) a three link tensor $\mathcal{T} \in \mathbb{C}^{d_1 \times d_2 \times d_3}$ and (e) an arbitrary order-R tensor. The tensor elements are depicted here by blue areas, while the open-ending links represent the indices of the tensor.

context. Thus, in order to explicitly indicate the specific shape of a tensor $\mathcal{T}$, we may write $\mathcal{T} = \mathcal{T}_{\nu_1,...,\nu_R}$ overloading the expression for addressing a certain element $\mathcal{T}_{i_1,...,i_R}$ in $\mathcal{T}$.

- Considering all d_r-dimensional spaces $\mathbb{V}_r$, a tensor consists of $\xi = \prod_{r=1}^{R} d_r$ elements.

- The norm of a tensor shall be the Frobenuius norm

$$||\mathcal{T}|| = ||\mathcal{T}||_F \equiv \sqrt{\sum_{i_1=1}^{d_1} \sum_{i_2=1}^{d_2} \cdots \sum_{i_R=1}^{d_R} ||\mathcal{T}_{i_1,...i_R}||^2}$$

- Similarly to this definition, we can define a real tensor $\mathcal{T} \in \mathbb{R}^{d_1 \times d_2 \times ... \times d_R}$ as a dense R-dimensional array of real numbers.

From Def. 2.1.1, it becomes evident that a vector is equivalent to an order-1 tensor with one corresponding link addressing a certain vector element by the index i_1. Further, a matrix forms an order-2 tensor with two indices i_1 and i_2 for its row and column respectively. Equivalently, you may imagine an order-3 tensor $\mathcal{T} \in \mathbb{V}_1 \times \mathbb{V}_2 \times \mathbb{V}_3$, as a three-dimensional cube with one index for each dimension. This conceptual procedure can be extended towards tensors with higher order r even though we as humans may lack the ability to imagine these objects geometrically. Fig. 2.2 illustrates the typical graphical representation for different tensors of various orders.

In practical applications, Tensor Networks exploit, or consist of, special tensors which obey additional conditions. The most important ones which are frequently used in different Tensor Network states and algorithms are *isometries* and *disentanglers*.

Isometry   The first special kind of tensors to be introduced as a vital tool in Tensor Networks are *isometries*. An isometry can be seen as the generalisation of a matrix U with orthonormal rows towards higher-order tensors. Similar to $UU^\dagger = \mathbb{1}$ for such a matrix U, an isometry $\mathcal{T}$ becomes an identity when contracting it with its complex conjugate over the proper links.
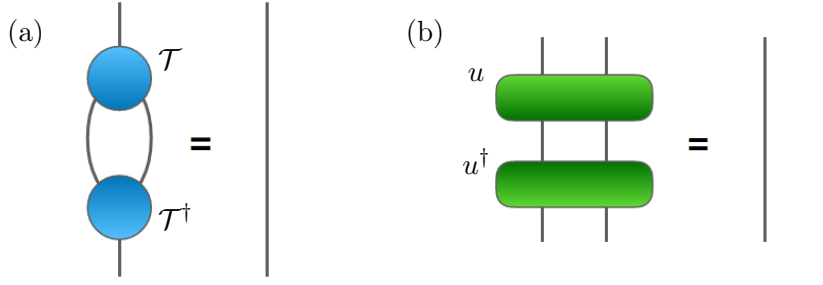
FIGURE 2.3: Graphical representation of the condition for an order-3 isometry (a) and a disentangler (b): When contracting one of them with their complex conjugate over the proper links, they become an identity indicated by a straight line.

> **Definition 2.1.2:** *(Isometry).*
>
> Let $\mathcal{T}$ be an order-$(R > 1)$ tensor with the links $(i_1, ..., i_R)$, then $\mathcal{T}$ is called an *isometry* with respect to the link i_k (with $k \in \{1, ..., R\}$) if it obeys the isometry condition
>
> $$\sum_{\{i_r\}\backslash i_k} (\mathcal{T})_{i_1,...,i_R} (\mathcal{T}^\dagger)_{i'_1,...,i'_R} = \delta_{i_k, i'_k}.$$

Fig. 2.3a shows this isometry condition for an order-3 tensor in the graphical notation. In general, a tensor can be *isometrised* by performing QR-decompositions as described in Sec. 2.1.2. An important example is the insometry $\mathcal{T}_{[i_1, i_2, i_3]}$ of order $R = 3$ which when it is isometrised with respect to i_3 performs a unitary renormalisation on the two subspaces (i_1, i_2). In Chapt. 3, it is described that this kind of tensor is crucial for an efficient implementation of a Tree Tensor Network. In particular, therein we will see that the Tree Tensor Network developed within this thesis is mainly composed out of order-3 isometries.

Disentanglers A disentangler can be seen as the generalisation of a unitary matrix U towards an order-4 tensor (see. Fig. 2.3b). Indeed, every disentangler u_k is a unitary when reshaping it into a matrix by fusing its first two and its last two indices respectively (see Sec. 2.1.2 for more details on the corresponding *link-fusion* operation).

> **Definition 2.1.3:** *(Disentangler).*
>
> Let u be an order-4 tensor with the links (i_1, i_2, i_3, i_4), then u is called a *disentangler* if it obeys the unitary condition
>
> $$\sum_{i_3, i_4} (u)^{i_1, i_2}_{i_3, i_4} (u^\dagger)^{i_3, i_4}_{i'_1, i'_2} = \delta_{i_1, i'_1} \delta_{i_2, i'_2}.$$

Hence, one disentangler u_k performs a unitary transformation on the two subspaces addressed by i_1 and i_2. In practical Tensor Network applications, this transformation aims to locally decouple - or *disentangle* - relevant degrees of freedom in a quantum many-body state (hence the name *disentangler*). This particular tensor was introduced as a key ingredient of the MERA geometry [73, 73, 74] which is described in more detail later on in Sec. 2.2.

Further, it plays a crucial role in the novel aTTN method [23] developed within the scope of this thesis (see Chap. 5 for an in-depth presentation of the aTTN).

2.1.2   Tensor Operations

Having formally introduced the tensor object, this section is dedicated to the description of selected tensor operations to manipulate tensor objects. These operations can be seen as a generalisation of linear algebra towards a *tensor algebra*. We will see throughout this section that many tensor operations have defining similarities with corresponding linear algebra operations and further that some tensor operations are numerically executed via linear algebra libraries, such as the *Basic Linear Algebra Subprograms* (BLAS) [109, 110] or the *Linear Algebra PACKage* (LAPACK) [111].

In general, the following tensor operations can be separated into two classes: *(i) Reshaping operations*, such as the tensor transposition or a link fusion, aiming to reshuffle the data of a tensor without changing its information content and *(ii) tensor algebra operations*, such as the tensor contraction or a tensor decomposition, performing actual floating-point operations on one or more tensor objects. It is worth mentioning that there are various established numerical libraries available able to perform all or a subset of existing tensor operations (for a detailed overview see Ref. [112]).

Tensor Transposition

Similar to the matrix transposition, the tensor transposition reshapes the entries of a tensor by reordering its indices. However, considering the higher order of a tensor, and thereby the several different indices, this idea generalises to a permutation of the tensor links.

Definition 2.1.4: *(Tensor Transposition).*

Let $\mathcal{T} \in \mathbb{C}^{d_1 \times ... \times d_R}$ be a tensor of order R with link indices $(i_1, ..., i_R)$ and further let π be a permutation of size R, then the transposed tensor $\mathcal{T}^\pi$ under the permutation π is defined as reordering its elements

$$\mathcal{T}_{i_1,...,i_R} \to \mathcal{T}^\pi_{i'_1,...,i'_R} = \mathcal{T}_{\pi(i_1,i_2,...,i_R)} \ ,$$

where the link indices $i_1, ..., i_R$ of $\mathcal{T}$ permute to $\pi(i_1, ..., i_R) = i_{\pi(1)}, ..., i_{\pi(R)}$.

Remark.

The permutation of the links leaves the order of the tensor invariant. We call the resulting tensor *transposed under the permutation* π. The numerical complexity of this operation is linear in the number of elements $\mathcal{O}\{\xi\}$.

As an example, Fig. 2.4a shows the transposition of an order-5 tensor $\mathcal{T}_{i_1,i_2,i_3,i_4,i_5}$ under the permutation $\pi = (1, 2, 4, 5, 3)$. The orange shade marks the effect of the permutation on the tensor T. After the operation, the resulting tensor has the elements $\mathcal{T}^\pi_{i_1,i_2,i_5,i_3,i_4} = \mathcal{T}_{i_1,i_2,i_3,i_4,i_5}$. Further, in agreement to linear algebra, the transposition of an order-2 tensor equals the matrix-transposition $M^T_{mn} = M_{nm}$.
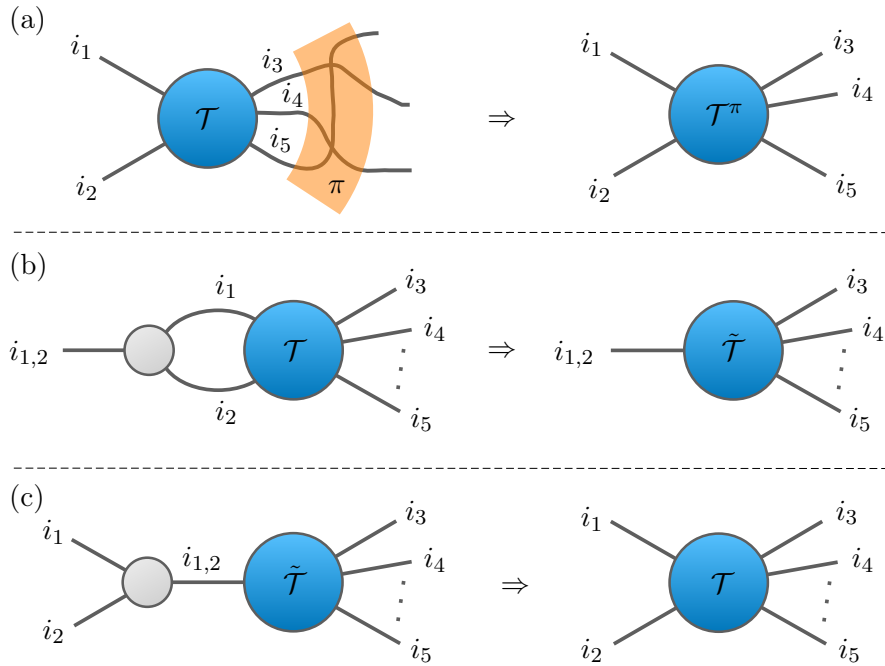
FIGURE 2.4: Different reshaping operations for tensors. (a) Transposition of the tensor $\mathcal{T} \rightarrow \mathcal{T}^\pi$ under the permutation $\pi = (1, 2, 4, 5, 3)$. Graphically, this transposition is illustrated by the permutation of the links within the orange-shaded area. (b) Fusing the links i_1 and i_2 of the tensor $\mathcal{T}$ resulting in the tensor $\tilde{\mathcal{T}}$ with link $i_{1,2}$. The gray tensor denotes the so called *fuse tensor*. (c) Splitting the link $i_{1,2}$ again into the links i_1 and i_2.

Fusing and splitting

Another form of reshaping a tensor without changing the information content it contains is the operation of fusing or splitting links. The former operation combines several adjacent links of a tensor into a single *fused link* while the latter describes the inverse operation which splits one single fused link again into its original decomposed links.

Definition 2.1.5: *(Link fusion and splitting).*

Let $(\nu_1, \nu_2, ..., \nu_k)$ be a tuple of links with corresponding indices $(i_1, i_2, ..., i_k)$, then the fusion ν of the links $(\nu_1, \nu_2, ..., \nu_k)$ is defined by the Cartesian product of all fused links

$$\nu = \nu_1 \times \nu_2 \times ... \times \nu_k \, ,$$

where the index i of the fused link runs over all possible combinations for $i \leftarrow (i_1, i_2, ..., i_k)$ within the spanned product space ν.

On the flip-side, we define the operation of *splitting links* as the inverted process of a link fusion, i.e. splitting one link i of a product space $i \in \left(\bigotimes_{\kappa=1}^{k} \mathbb{V}_\kappa \right)$ into k different links $i_1, i_2, ..., i_k$ inverting the process of fusing links.

Remark.

> The targeted fused link ν contains every possible index pair of the initial links and consequently exhibits the dimension $d_i = \prod_{\kappa=1}^{k} \dim\{i_\kappa\}$.
> The combination of spaces itself can be represented by a so-called *fuse-tensor* containing the links of $i_1, i_2, ..., i_k$ and i. To complete the fusion, the original tensor will be contracted with the fuse-tensor. The tensor rank will be reduced by this operation.

As an example, Fig. 2.4 illustrates both operations, the fusion and splitting of two links. The grey tensor therein is the fuse-tensor which in the case of fusing links and consequently splitting them again is equivalent for both operations.

As another crucial example, a tensor $\mathcal{T}_{i_1,i_2,...,i_R}$ of arbitrary order $R > 2$ can be converted into a matrix by bipartiting all links of the tensor and fusing both sets of links $i_1, ..., i_j$ and $i_{j+1}, ..., i_R$ respectively (with $0 < j < R$). This fusion results in a Tensor with two indices $m = (i_1, ..., i_j)$ and $n = (i_{j+1}, ..., i_R)$, or i.e. a matrix

$$M_{mn} = \mathcal{T}_{(i_1,...,i_j),(i_{j+1},...,i_R)} \ .$$

This fusion is an important part of many Tensor Network algorithms as this operation enables to apply matrix operations from classical linear algebra straightforwardly to tensors with more than two links.

Contraction

The contraction of two tensors generalises the matrix-matrix-multiplication towards the tensor algebra: It combines the two tensors A and B to a resulting third tensor $C = AB$ with different properties and elements while preserving the underlying information content. Differently from the matrix-multiplication, the operands can be of arbitrary order so that the tensor contraction can be performed over several coinciding links $m_1, ..., m_\mu$ (see Fig. 2.5).

> **Definition 2.1.6:** *(Tensor Contraction).*
>
> Let A and B be an order-r_a and an order-r_b tensor, respectively, with coinciding links $m_1, ..., m_\mu$, then a tensor contraction is defined as the summation over all indices $m_k \in \{m_1, ..., m_\mu\}$ of the coinciding links
>
> $$C_{l,n} = \sum_m A_{l,m} B_{m,n} \ ,$$
> $$\text{with} \quad l = (l_1, ..., l_\lambda), m = (m_1, ..., m_\mu)$$
> $$\text{and} \ n = (n_1, ..., n_\nu)$$
>
> where the links $l_k \in \{l_1, ..., l_\lambda\}$ and $n_k \in \{n_1, ..., n_\nu\}$ remain after the contraction as the links of the resulting tensor C. Consequently, the order of C is $r_c = \lambda + \nu = r_a + r_b - \mu$.

In this tensor algebra, the matrix-multiplication itself can be seen as a contraction of two order-2 tensors over the coinciding link m as illustrated in Fig. 2.5 showing the graphical representation of both, the matrix-multiplication as well as a general tensor contraction.

While matrix-matrix multiplications are typically carried out with optimised linear algebra libraries, such as GEMM [113], there exist three main ways for the computational
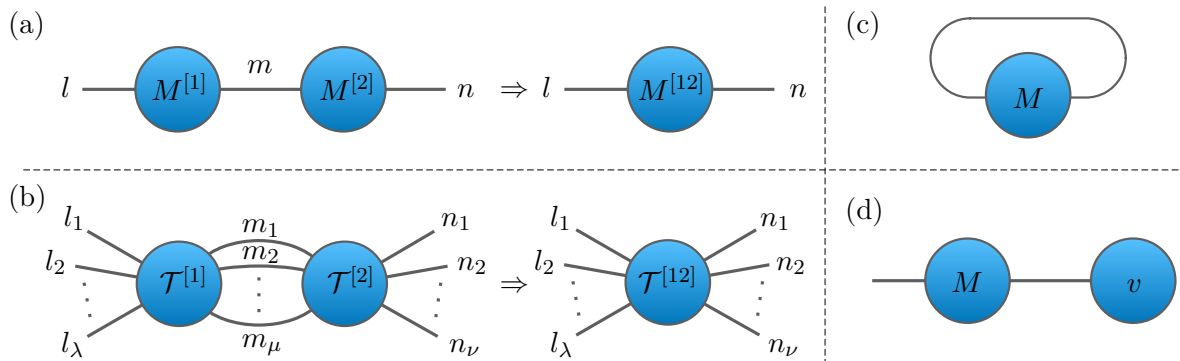
FIGURE 2.5: Graphical representation of different Tensor Contractions: A Matrix-Matrix-multiplication (a) can be generalised to a tensor contraction for two tensors $\mathcal{T}^{[1]}$ and $\mathcal{T}^{[2]}$ of arbitrary orders (b). A special kind of tensor contraction is the trace of an order-2 tensor, i.e. a matrix (c). As another example a vector-matrix multiplication can be written in the graphical tensor notation as well (d).

implementation of tensor contractions. The most prominent one is the *Transpose-Transpose-GEMM-Transpose* (TTGT) approach [24, 114], in which the two tensors are reshaped into matrix shape by proper transposition and link fusion, before performing the contraction as optimised matrix-matrix multiplication. Crucially at this step, the algorithm can exploit efficient linear algebra libraries, such as the GEMM routine from LAPACK [111]. Finally, the resulting matrix, i.e. order-2 tensor, is reshaped by splitting its links again into the original links of the now contracted tensors. As alternative methods, *Loops-over-GEMM* (LoG) [115–118] or *GEMM-Like Tensor-Tensor Contraction* (GETT) [118] are the potentially more efficient strategies trying to get rid of the pre- and post-processing steps of the common TTGT to solve the contraction without changing the shapes of the tensors. However, they are non-trivial to implement and yet to be implemented in a highly optimised way to compete with the TTGT strategy.

The computational complexity of the tensor contraction $\mathcal{O}\left(d_{l_1}...d_{l_\lambda} d_{m_1}...d_{m_\mu} d_{n_1}...d_{n_\nu}\right)$ scales with the dimensions of all links involved[1]. Due to this complexity increasing with the number of links, the contractions play a crucial role for the efficiency of Tensor Network algorithms. Thus, the contraction schemes have to be thought-through when designing and developing an efficient Tensor Network from the computational point of view.

Factorisation

Each tensor can be factorised by taking advantage of classical linear algebra by using matrix decompositions. Therefore, the tensor to be decomposed will be reshaped into an order-2 tensor by fusing the links of two different subsets as described in Sec. 2.1.2. After the factorisation, the links are properly split with respect to the two different sets of its links. The most prominent decompositions for tensors (applied in matrix-shape) are the *Singular Value Decomposition* (SVD) and the *reduced QR-decomposition*.

---

[1]when carried out as numerically optimised multiplication, the complexity can be reduced in practical applications

Definition 2.1.7: *(Reduced QR-Decomosition).*

Let $M \in \mathbb{C}^{m \times n}$ be a complex matrix. Then, the reduced QR-decomposition is defined as

$$M = QR$$

decomposing M into a $Q \in \mathbb{C}^{m \times \min\{m,n\}}$ featuring $Q^\dagger Q = \mathbb{1}$ and a matrix $R \in \mathbb{C}^{\min\{m,n\} \times m}$ consisting of non-vanishing entries only on and above its diagonal.

Remark.

For $m \leq n$, R becomes square and thus forms an upper triangular matrix, whereas in the case of $m \geq n$, Q equals a unitary matrix with $Q^\dagger Q = QQ^\dagger = \mathbb{1}$.
This factorisation refers to the *reduced* or *thin QR-decomposition* while by the conventional definition the full matrix Q_f is always a squared matrix, which can be achieved for $m < n$ by filling the missing entries in the reduced Q with orthonormal columns Q_c and R with zeros. Thus, it contains the same amount of information while being beneficial for computational purposes.

Definition 2.1.8: *(Singular Value Decomposition).*

Let $M \in \mathbb{C}^{m \times n}$ be a complex matrix, then the singular value decomposition is defined as

$$M = USV^\dagger$$

decomposing M into a matrix $U \in \mathbb{C}^{m \times \min\{m,n\}}$ with orthonormal columns, a diagonal matrix $S \in \mathbb{R}^{\min\{m,n\} \times \min\{m,n\}} \in \mathbb{C}^{\min\{m,n\} \times n}$ and a matrix $V^\dagger \in \mathbb{C}^{\min\{m,n\} \times n}$ consisting of orthonormal rows. Thus, U obeys $U^\dagger U = \mathbb{1}$ where its orthonormal colums are referred to as left singular vectors of M. Analogously, $V^\dagger V = \mathbb{1}$ holds true where the orthonormal rows of $V^\dagger$ are the right singular vectors of M. The diagonal matrix S contains only real, non-negative elements referred to as *singular values* of M. The number of non-zero singular values is the *rank* of the matrix M.

Fig. 2.6 illustrates the SVD (QR-decomposition) following Def. 2.1.8 (Def. 2.1.7) applied on an order-5 tensor $\mathcal{T}$: Firstly a direction of truncation has to be defined, i.e. which sets of links shall be fused before decomposing the tensor in the resulting matrix shape. After the decomposition, the first set, here (i_1, i_2), is attached to the U (Q) tensor, while the other set, here (i_3, i_4, i_5), is attached to the V (R) tensor. After splitting the links again, $\mathcal{U}$, $\mathcal{V}$ and $\mathcal{Q}$ are isometries obeying Def. 2.1.2.

In practical applications, the SVD is a key tool of Tensor Networks: On one hand, the SVD enables to establish a desired gauge in a Tensor Network (see Sec. 2.2 for more details on gauge-freedom in Tensor Networks) which in return can provide higher efficiency in certain computations, such as the calculation of observables (see Sec. 3.1.3). Additionally, the singular values can be used to measure entanglement properties within loop-less networks as described later on in more detail for Tree Tensor Networks (see Sec. 3.2.2). Furthermore, the SVD can be deployed to efficiently compress a tensor or even a complete Tensor Network:
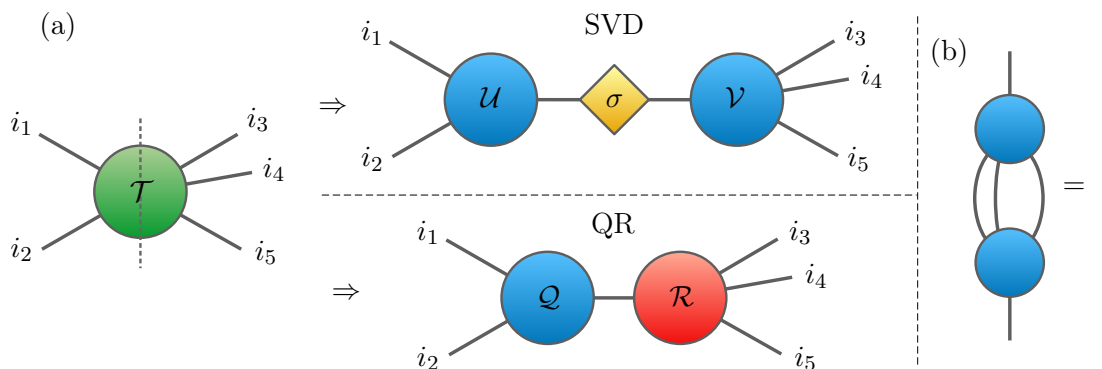
FIGURE 2.6: Graphical representation of tensor factorisations. (a) The tensor $\mathcal{T}$ is decomposed using a singular value decomposition (top) or using a QR-decomposition (bottom). The former decomposes $\mathcal{T}$ into three tensors: Two isometries $\mathcal{U}$ and $\mathcal{V}$ and the diagonal matrix σ containing the singular values. The latter decomposes $\mathcal{T}$ into an isometry $\mathcal{Q}$ and a second tensor $\mathcal{R}$. (b) All isometries (blue) obey the isometry condition of Def. 2.1.2 when contracted with their complex conjugate over the open links of the decomposition.

Lemma 1. *For a given matrix M, the SVD enables to approximate M by another matrix $\tilde{M}$ with lower rank $\tilde{r} < r$, i.e. living in a smaller space, with a minimal deviation introduced in its Frobenius norm $||M||_F^2 := \sum_{ij} |M_{ij}|^2 = \mathrm{tr}\{MM^\dagger\}$.*

Proof. The SVD on M yields

$$||M||_F^2 = \mathrm{tr}\{USV^\dagger VS^\dagger U^\dagger\} = \mathrm{tr}\{SS^\dagger\} = \sum_{i=1}^{r} \sigma_i^2 \ ,$$

where r is the Schmidt rank of the matrix M. Consequently, the deviation $\epsilon_{\tilde{r}}(M, \tilde{M}) = |\,||M||_F^2 - ||\tilde{M}||_F^2\,|$ due to the approximation $\tilde{M}$ is minimal for considering the highest $\tilde{r}$ singular values, together with their corresponding singular vectors. If all σ_i are ordered descendingly, the deviation $\epsilon_{\tilde{r}}$ introduced can be computed by the discarded singular values σ_k.

$$||\tilde{M}||_F^2 = \sum_{i=1}^{\tilde{r}} \sigma_i^2 = ||M||_F^2 - \underbrace{\sum_{k=\tilde{r}+1}^{r} \sigma_k^2}_{=\epsilon_{\tilde{r}}}$$

Thus, $\tilde{M}$ is the best possible approximation of M for a lower rank $\tilde{r}$ with respect to the Frobenius norm. *q.e.d.*

2.2   Tensor Network states and algorithms

(a)   (b)
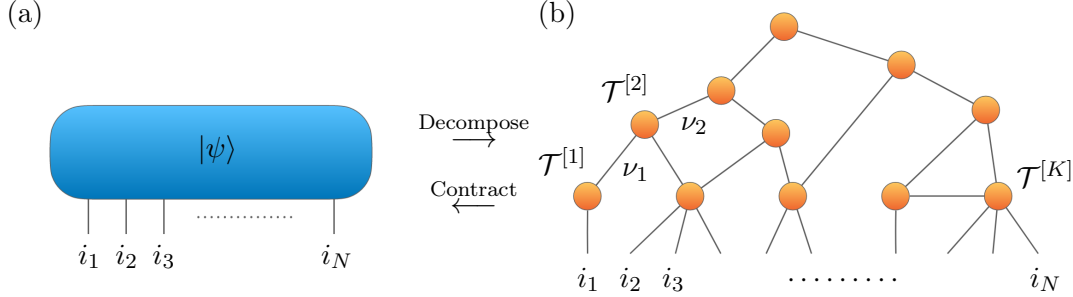

FIGURE 2.7:   (a) High-order tensor $|\psi\rangle$ describing a quantum many-body wavefunction exactly with all exponentially many coefficients. (b) A possible Tensor Network decomposition of the same wavefunction where the tensors are illustrated in different colours and connected with each other over *internal bond-links*. The dimension of the bond-links can be controlled to compress the information represented within the network.

In quantum mechanics, Tensor Network states represent an arbitrary pure quantum many-body state $|\Psi\rangle \in \mathscr{H}$ on a lattice $\mathscr{L}$ with corresponding Hilbert space $\mathscr{H}$. In general, the lattice $\mathscr{L}$ can be any D-dimensional lattice containing N sites, where each site $k \in [1, ..., N]$ is described by a local Hilbert space $\mathscr{H}_k$ with finite dimension d_k, so that the complete Hilbert space of $\mathscr{L}$ is spanned by $\mathscr{H} = \otimes_{k=1}^{N} \mathscr{H}_j$. The most general pure state living on such a lattice can be written as

$$|\psi\rangle = \sum_{i_1=1}^{d_1} \sum_{i_2=1}^{d_2} ... \sum_{i_N=1}^{d_N} \Psi_{i_1,...,i_N} |i_1\rangle \otimes |i_2\rangle \otimes ... \otimes |i_N\rangle \, , \qquad (2.1)$$

where $|i_k\rangle$ denotes the local state space of the site k. For sake of simplicity, we assume in the following that all local dimensions are of equal size ($\forall k : d_k = d$) which is frequently the case in various quantum systems of interest. The complete state vector seeks for an exponentially growing number of coefficients $\Psi_{i_1,...,i_N}$ with increasing system size, namely $\dim\{\Psi\} = \prod_{k=1}^{N} d_k \overset{d_k=d}{=} d^N$. Due to the separable Hilbert spaces $\mathscr{H}$, the wave function can be recast in an order-N tensor, where each link i_k of this tensor corresponds to a local Hilbert space $\mathscr{H}_k$ (see Fig. 2.7).

As denoted in the introduction of this chapter, the idea of Tensor Networks is to decompose this exponentially large tensor into a set of connected local tensors $\{\mathcal{T}^{[k]}\}$. Some of these tensors take on the links $\{i_k\}$ of the complete tensor as open links corresponding to the local Hilbert spaces $\mathscr{H}_k$ of the physical system. Thus, we can formally define a Tensor Network similar to a graph in *graph theory* with the addition of these physical links $\{i_k\}$ by Def. 2.2.1

Definition 2.2.1: *(Tensor Network).*

A Tensor Network shall be defined, as a tuple $\mathcal{G} = (\mathfrak{T}, \mathfrak{N})$ consisting of a set $\mathfrak{T} = \{\mathcal{T}\}$ of two or more tensors and a set $\mathfrak{N} = \{\nu\}$ of corresponding links. The set $\mathfrak{N}$ is further divided into *internal* links $\{\chi\}$ connecting the tensors with each other and *external* or *physical* links $\{i\}$ which are connected to one tensor only, thus having one open ending each.
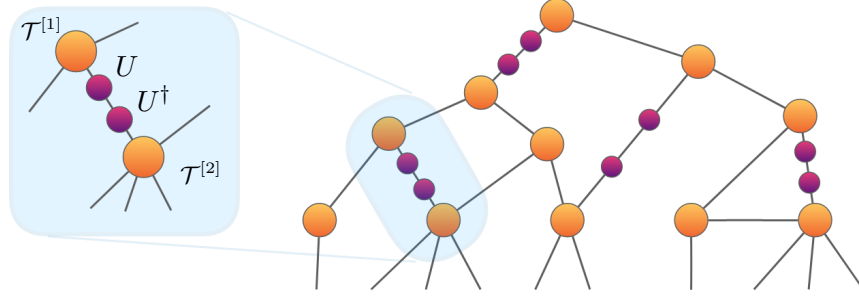
FIGURE 2.8: Gauge freedom in a Tensor Network: At every internal link an identity $\mathbb{1} = UU^{\dagger}$ can be inserted without changing the global state representation. When contracting U ($U^{\dagger}$) into the nearest tensor, the structure of the original Tensor Network is restored, giving an alternative description of the same information content.

Further, we denote:

- The *size* of the network K shall be defined as the total number of tensors within the tuple $\mathcal{G}$, i.e. the size $\mathcal{G}$ or $\mathfrak{T}$ respectively.

- The total number of physical links N, i.e. the size of $\{i\}$, is the *system size*.

- The typical literature introduces to $d_k = \dim\{i_k\}$ as the dimension of the physical link i_k. Depending on the network and description, the dimension of an internal link is typically referred to as χ, m, or D. With some exceptions, this thesis uses $m_k = \dim\{\chi_k\}$ describing the dimension of the internal link χ_k, i.e. its bond-dimension.

- Every Tensor Network should, in theory, allow to be contracted to a single tensor $\Psi_{\{i\}}$ when contracting the constituting tensors over all internal links $\{\chi\}$, following

$$\Psi_{i_1,\dots,i_N} = \sum_{\chi_1=1}^{m_1} \sum_{\chi_2=1}^{m_2} \cdots \sum_{\chi_N=1}^{m_N} \left(\prod_{k=1}^{K} \mathcal{T}^{[\kappa]}_{\{i\}_\kappa, \{\chi\}_\kappa} \right), \tag{2.2}$$

where $\mathcal{T}^{[k]}_{\{i\}_\kappa, \{\chi\}_\kappa}$ is the κ-th tensor in the set $\mathfrak{T}$ with the links $\{i\}_\kappa$ and $\{\chi\}_\kappa$ denoting the set of all physical and internal links respectively attached to the κ-th tensor. Noting that each internal link χ_k occurs twice within the set of tensors $\mathfrak{T}$ and shall be contracted over, while each physical link occurs only once within $\mathfrak{T}$.

- Finally, we define the maximum dimension of all internal links as the *bond-dimension* of the network $m = \max_{k=1,\dots,K}\{\dim\{\chi_k\}\}$.

Lemma 2. *Every Tensor Network has a gauge freedom.*

Proof. Let $\mathcal{T}^{[1]}$ and $\mathcal{T}^{[2]}$ be two tensors within a Tensor Network connected by the link ν of dimension n, such that the tensors can be contracted over ν to

$$\mathcal{T}^{[12]} = \mathcal{T}^{[1]} \mathcal{T}^{[2]} \ .$$

Further be $U \in \mathbb{C}^{n \times n}$ a unitary matrix, then

$$\mathcal{T}^{[12]} = \mathcal{T}^{[1]}\mathcal{T}^{[2]} = \underbrace{\mathcal{T}^{[1]}U}_{\equiv \tilde{\mathcal{T}}^{[1]}} \underbrace{U^{\dagger}\mathcal{T}^{[2]}}_{\equiv \tilde{\mathcal{T}}^{[2]}} = \tilde{\mathcal{T}}^{[1]}\tilde{\mathcal{T}}^{[2]}$$

holds true. This basis transformation can be applied to any link connecting two tensors within the complete Tensor Network as illustrated in Fig 2.8.   *q.e.d.*

Depending on the way of decomposing the complete state vector $\Psi_{i_1,\dots,i_N}$, several Tensor Network geometries have been established over the last decades. The most prominent of which are the *Matrix Product States* (MPS), the *Tree Tensor Networks* (TTN), the *Multiscale Entanglement Renormalisation Ansatz* (MERA) and the *Projected Entangled Pair States* (PEPS) which are discussed in more detail in the upcoming Sec. 2.2. However, before going to this section, let us first briefly raise the daring question: What is the best Tensor Network geometry? Or in other words, what makes a Tensor Network a *good* Tensor Network state? Well, that one is easy to answer theoretically: On the one hand, the Tensor Network should faithfully describe the properties of the quantum many-body state they represent. For such a faithful description, a *good* Tensor Network has to be able to accurately capture the underlying quantum entanglement of the system. In particular, they shall satisfy the same entanglement bounds under real-space bipartitions as the physical states they represent. On the other hand, since Tensor Networks are a numerical tool, a *good* Tensor Network shall be numerically efficient, thus, the underlying operations performed on the network shall include a minimum amount of FLoating point OPerations (FLOPs). In other words, the numerical complexity $\zeta = \mathcal{O}(poly(N)poly(m))$ shall scale as low as possible with respect to the system size N and the network's bond-dimension m.

The following section addresses the first point regarding the faithfulness of the representation by discussing the properties of the quantum entanglement of quantum many-body systems and how they translate into a general Tensor Network representation of a state. The subsequent section is then dedicated to the different Tensor Network geometries and further aims to illustrate that, in fact, there is no Tensor Network which generally is the *best* since the accuracy of each network highly depends on the system they investigate.

2.2.1   Area law and Entanglement Properties

Quantum entanglement is one of the most impactful properties separating quantum many-body states from classical many-body systems. It gives rise to the high potential of quantum systems in capturing an exponentially large amount of information. When formally introducing Tensor Networks, it already became evident that a general quantum wavefunction lives in a with the number of constituents exponentially growing vector space and therefore is able to capture such a large information content from a mathematical point of view. In this respect, when decomposing such a random state into a Tensor Network, this Tensor Network would require an exponentially large bond-dimension to represent the exponentially large wavefunction exactly. This would make the use of Tensor Networks as a quantum state representation practically unfeasible.

However, most *physical* quantum states obey certain entanglement bounds under real-space bipartitions, also known as *Area Laws* making it possible to decompose these states very efficiently into a Tensor Network structure, to a high degree even loss-less. The keypoint for such an efficient decomposition into a Tensor Network is the remarkable feature of Tensor Network states that their entanglement properties can be controlled by the network

topology itself. Therefore, it is crucial for an efficient numerical implementation that the Tensor Network topology satisfies the same entanglement bounds as the *physical states* they represent. In what follows, these entanglement bounds are discussed in more detail for the Tensor Network states.

In information theory, entropy as such is used as a measure of the information content inherent in a system [6, 119, 120]. This information content, or amount of quantum entanglement incorporated in a quantum state, can be measured utilizing the entanglement entropy S which describes the shared information between two bipartitions. The *Von Neumann entropy*, as a measure of the entanglement, can be assessed by the *Schmidt decomposition*, which formally decomposes the state $|\psi\rangle$ into two bipartitions $|\Psi^{[\mathcal{A}]}\rangle$ and $|\Psi^{[\mathcal{B}]}\rangle$ [6].

Theorem 2.2.1: *(Schmidt Decomposition).*

Every bipartite quantum state $|\psi\rangle \in \mathscr{H}_\mathcal{A} \otimes \mathscr{H}_\mathcal{B}$ can be decomposed as

$$|\Psi\rangle = \sum_{\alpha=1}^{\chi} \lambda_\alpha^{[\mathcal{A},\mathcal{B}]} |\Psi^{[\mathcal{A}]}\rangle \otimes |\Psi^{[\mathcal{B}]}\rangle \ ,$$

into two orthonormal bases $|\Psi^{[\mathcal{A}]}$ ($|\Psi^{[\mathcal{B}]}$), or *Schmidt bases*, corresponding to the Hilbert-space $\mathscr{H}_\mathcal{A}$ ($\mathscr{H}_\mathcal{B}$) of the bipartition $\mathcal{A}$ ($\mathcal{B}$), and the *Schmidt-coefficients* $\lambda_\alpha^{[\mathcal{A},\mathcal{B}]}$.

Remark.

In practical applications, this decomposition can be performed using an SVD where the state is first rewritten as a matrix $\Psi = \sum_{\{i\}} \psi_{(i_1,...,i_k),(i_{k+1},...,i_N)} |i_1,...,i_k\rangle_\mathcal{A} |i_{k+1},...,i_N\rangle_\mathcal{B}$ with respect to the two bipartitions which is consequently decomposed into $\Psi = U\sigma V^\dagger$. Here, U (V) with $U^\dagger U = \mathbb{1}$ ($VV^\dagger = \mathbb{1}$) span the orthonomal basis, while the non-zero, normalised singular values of σ are the Schmidt-coefficients $\lambda_\alpha^{[\mathcal{A},\mathcal{B}]}$.

The Schmidt decomposition is directly linked to the entanglement of the state Ψ over its Schmidt-coefficients: The entanglement entropy S between the bipartitions $\mathcal{A}$ and $\mathcal{B}$ equals

$$S = -\sum_{\alpha=1}^{\chi} |\lambda_\alpha^{[\mathcal{A},\mathcal{B}]}|^2 \log\{|\lambda_\alpha^{[\mathcal{A},\mathcal{B}]}|^2\} \ .$$

Consequently, the minimal entropy $S = 0$ is reached only when there is only one non-zero singular value, or Schmidt coefficient, $\lambda_1 = 1$ while $\lambda_{\alpha \neq 1} = 0$. In this case, the two bipartitions can be completely separated since they do not share any information, i.e. the bipartitions are not entangled. In contrast, a higher $S > 0$ indicates shared information between the two bipartitions, thus the state is increasingly entangled with growing S. Evidently, this Von Neumann entropy is bounded by $S \leq \log \chi$.

Applying this upper bound to a general random state obeying Eq. (2.1) leads to the so-called *volume law* stating that the entanglement entropy scales with the volume of the system: Considering for instance a D-dimensional system with equal system sizes L (such that the total number of sites $N = L^D$), this volume law would suggest that the entanglement entropy when bipartiting the system indeed scales extensively with $S \sim L^D$. However, in the recent year it has become evident that most *physical states* do not hold up to this amount of entanglement and rather exhibit significantly less entanglement following to the so-called *Area*
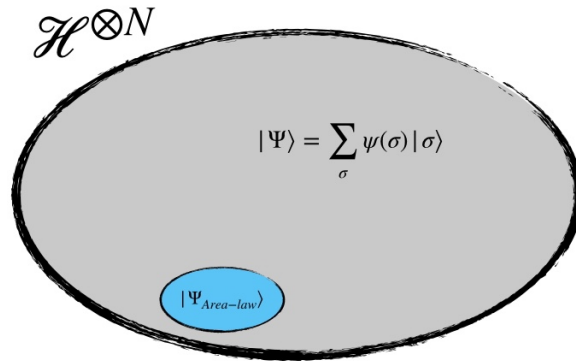
FIGURE 2.9: Quantum many-body Hilbert space spanning the complete manifold for pure states $|\psi\rangle$. The states $|\psi_{Area-Law}\rangle$ that obey an area law scaling of entanglement entropy occupy a tiny, distinct corner of the complete manifold.

law [45, 46, 121]. Thus, the physical states live in a tiny fraction of the complete Hilbert space as illustrated in Fig. 2.9 which can be ideally accessed by Tensor Network representations.

As we will see in the following, the family of Tensor Network states precisely targets this corner of relevant states in the complete Hilbert space [48, 121, 122]. Due to the renormalisation group nature of Tensor Networks, they are able to identify and keep track of the relevant degrees of freedom in a quantum many-body system.

Area laws in physical systems — The entanglement entropy of quantum states obeying area law is expected to increase proportionally to the boundary area of the bipartitions rather than the volume of the bipartitions:

$$S = \mathcal{O}(\partial\mathcal{A}) \left(\sim L^{D-1} \right)$$

The latter proportion in brackets describes the entanglement behavior for area law states in a D-dimensional system with each linear system dimension equaling L, i.e. $N = \prod_{i=1}^{D} L_i \stackrel{(L_i=L)}{=} L^D$. Such an area law behaviour of the entanglement properties of a state has been proven for various systems of different dimensions: In $D = 1$, it is well-known that all gaped models obey the area law with $S \sim \mathcal{O}(1)$ [48, 123–125] while only critical states corresponding to a conformal field theory may exhibit a logarithmic behaviour of the entropy $S \sim \log L$ [126–128]. Further, area laws of $S \sim L^{D-1}$ are proven for $D > 1$ among others for ground states of gaped models in the same phase as ones satisfying an area law [129, 130], for frustration-free spin models, in free gaped bosonic and fermionic systems [131, 132], and even in critical bosonic systems [133] while in the gap-less case of Fermions, for instance, the entropy scales slightly higher with $S \sim L^{D-1} \log L$ [133–135]. However, in general, it is expected that all gaped lattice models satisfy such a behaviour of the entanglement properties governed by the area law. For an in-depth review on entanglement properties in various quantum systems, see Ref. [121].

The entanglement properties of a quantum many-body state, and in particular the insight that most physical states obey an area law, has direct implications on the Tensor Networks designed to represent such a quantum state since they define an upper bound for the entanglement in the system.
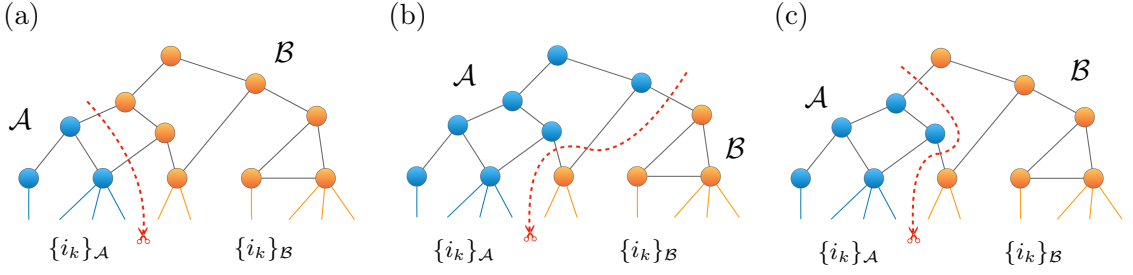
FIGURE 2.10: Possible Tensor Network bipartitions with the links $\{i_k\}_{\mathcal{A}}$ (blue) and $\{i_k\}_{\mathcal{B}}$ (orange) addressing the physical sites within the bipartitions $\mathcal{A}$ and $\mathcal{B}$, respectively. Each cut-through link gives an entropy contribution $\leq \log m_{\bowtie}$ upper-bounded by its bond-dimension $m_{\bowtie}$.

Entanglement bounds in Tensor Networks — Let Ψ be a Tensor Network representation as introduced in Def. 2.2.1 with physical links $\{i_k\}$. Further, among the physical links, the disjoint subsets $\{i_k\}_{\mathcal{A}}$ ($\{i_k\}_{\mathcal{B}}$) shall address the physical sites belonging to a bipartition $\mathcal{A}$ ($\mathcal{B}$) as illustrated in Fig. 2.10. Then, the network can be bipartitioned with respect to the two the physical partitions $\mathcal{A}$ and $\mathcal{B}$ by cutting through a set of internal links $\{\chi_{\bowtie}\}$. The Von Neuman entanglement entropy captured within the network for such a bipartition satisfies the inequality

$$S_{TN}^{[\mathcal{A},\mathcal{B}]} \leq \min_{\{\bowtie\}} \left[\sum_{\bowtie} \log m_{\bowtie} \right] , \tag{2.3}$$

where $m_{\bowtie}$ denotes the bond-dimension of the link $\chi_{\bowtie}$. Consequently, this minimum over all possible cuts through the network which result in the same bipartition gives an upper bound for the entanglement entropy the network is able to represent.

As an example, Fig. 2.10 illustrates different possibilities for such a bipartition of the Tensor Network representation. This bipartition cuts through the selected links and breaks the network into separate sub-(Tensor Networks). Crucially, every broken link can maximally give an entropy contribution of $\log m_{\bowtie}$ upper-bounded by its bond-dimension $m_{\bowtie}$. Thus, if we would have a network with bond-dimension $m = 1$, this upper-bound indicates that $S_{TN}^{[\mathcal{A},\mathcal{B}]} = 0$ independently on the sizes of $\mathcal{A}$ and $\mathcal{B}$. As a consequence, every Tensor Network state with $m = 1$ represents a wave function without entanglement, i.e. is a product state. Therefore, such a network coincides with the approach used in mean-field theory. However, increasing the bond-dimension m enables the Tensor Network to faithfully represent states obeying the same entanglement bounds as Eq. (2.3) suggests for the network. Therefore, either the bond-dimension m or the number of cut links should scale according to these entanglement bounds for a Tensor Network to be able to capture the entanglement properties of the underlying quantum system. On the other hand, a Tensor Network representation can be improved computationally by limiting m to the minimum value required to capture the correlations of the system. Finding this balance of computational efficiency, like in mean-field approaches, and faithfully representing the underlying quantum entanglement, as it is the case in the exact but exponentially large description, is in practice a highly non-trivial challenge for Tensor Network geometries in higher dimensions.

2.2.2 Tensor Network state representations

There are numerous ways to create a tensor network to represent a certain quantum many-body state. In particular, the complete, exponentially large state vector $\Psi_{i_1,\ldots,i_N}$ can be
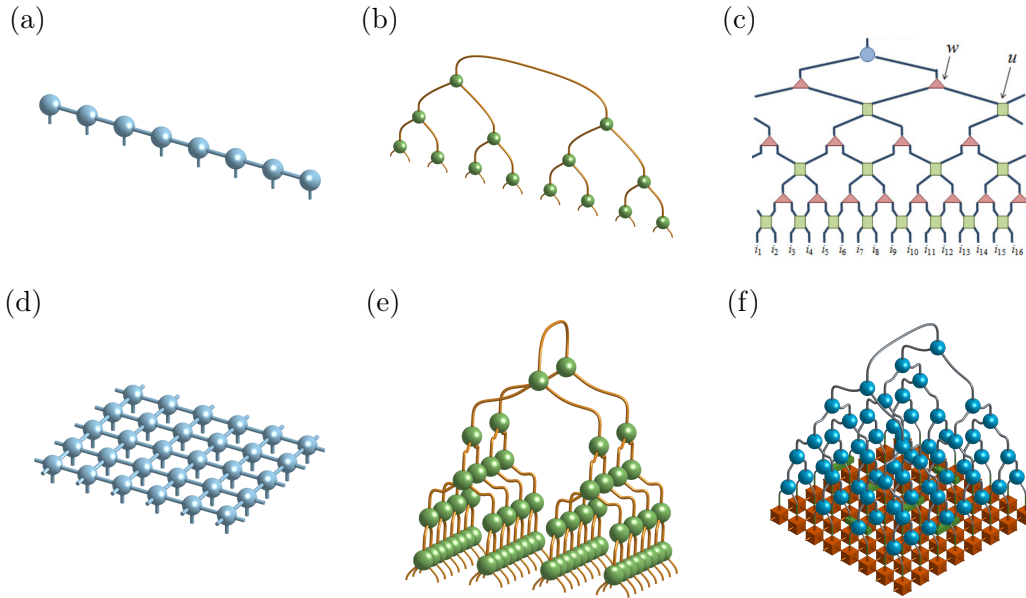
FIGURE 2.11:   Different Tensor Network representations for one- and two-dimensional quantum many-body wavefunctions: The MPS (a), the TTN (b) and the MERA (c) for 1D systems, and the PEPS (d), the TTN (e) and the aTTN (f) for 2D systems. Sub-picture (c) of MERA is taken from Ref. [74].

decomposed following different strategies. This section aims to present the most commonly used Tensor Networks in more detail discussing the strengths and weaknesses of each and giving an overview of typical problems the different networks are frequently deployed to solve. While there are some approaches designed to work directly in the thermodynamical limit, the following discussion focuses on Tensor Networks for the simulation of finite size systems.

In a nutshell, the Matrix Product States (MPS) are the established Tensor Network geometry for equilibrium, and in many cases even out-of-equilibrium problems of one-dimensional systems while MERA provides an efficient representation for critical systems in 1D [66, 74–76]. However, the development of algorithms for Tensor Networks in two-dimensions or higher is still in progress, as the combination of accuracy and scalability remains an open challenge for Tensor Network methods when simulating high-dimensional systems [21, 24]. The most prominent Tensor Network representations are the Projected Entangled Pair States (PEPS) [45, 71, 77–79] for a two-dimensional system only, and the Tree Tensor Networks (TTN) [22, 24, 35, 37, 67, 68, 80] which, in principle, can be defined for any system dimensionality. The following discussion is partly taken from the appendix of Ref. [37] which is fully incoperated in this thesis as Chapt. 4.

Matrix Product States —   The MPS [46, 48, 61, 136] are the first developed geometry in the age of Tensor Network. In the last two decades, various algorithms for MPS have been introduced making the MPS up to date the main working tool for one-dimensional problems in equilibrium [61, 137], and in many cases even out-of-equilibrium [62, 63]. The geometry of an MPS is suited for a one-dimensional system quite intuitively: As illustrated in Fig. 2.11a, each tensor $\mathcal{T}^{[k]}$ had one physical link attached to the physical site i_k and two internal links (except for the boundary tensors). With a numerical complexity of $\mathcal{O}(m^3)$ for prominent minimisation and time-evolution algorithms, it is the most efficient network from

a computational point of view while it as well fulfils the area law in 1D. Another established approach is to extend the MPS for two-dimensional systems, also known as 2D-DMRG, [72] which can as well be optimised with the same economic scaling of $\mathcal{O}\left(m^3\right)$. It works extremely well for cylindrical two-dimensional systems with limited system width reaching in practical applications bond-dimensions over $m \approx 4000$ when executed on a High-Performance cluster. However, for $D > 1$-dimensional systems, the MPS by design does not fulfil the underlying area law scaling; To compensate for this drawback, the MPS requires an exponentially large bond-dimension for higher dimensional systems and thus eventually fails to faithfully describe the quantum many-body state when scaling both system sizes equivalently.

Projected Entangled Pair States —   Comparable to the MPS in 1D, the PEPS approximates the complete state vector of a two-dimensional system by a Tensor Network representation with one tensor for each physical site. These tensors are then connected through a grid mimiking the physical lattice as shown in Fig. 2.11d, yielding a Tensor Network with *loops* (non-local gauge redundancies). Thus, the PEPS is the most intuitive - and potentially most powerful - representation of a two-dimensional quantum many-body wavefunction and satisfies the area laws of entanglement by its structure [45, 121]. However, the PEPS, it suffers from a high numerical complexity (typically with $\mathcal{O}(m^{10})$ for finite-sized PEPS [77, 138]) and it lacks an exact calculation of expectation values [20]. In fact, for a finite square lattice with $N = L \times L$ sites, the contraction of the complete PEPS to perform this calculation scales exponentially on average system length L [139]. Further, due to this computationally more expensive optimization of the PEPS ansatz, the typical bond-dimensions achieved are in the order of $m \sim 10$ which is sufficiently large for many spin systems with local dimension $d = 2$. However, two-dimensional systems with high local dimension $d > 10$, such as higher-order spin systems, Lattice Gauge Theories or fermionic Hubbard-like systems, raise a non-trivial challenge for the PEPS ansatz. Furthermore, the PEPS is strictly designed for two-dimensional problems with open boundary conditions while its application to systems with periodic boundary conditions or even of higher dimensionality is practically unfeasible.

Multi-scale Entanglement Renormalization Ansatz —   The MERA was first introduced in 2007 providing an efficient representation for analysing critical systems in 1D [66]. Thus, in contrast to MPS, it can handle the entanglement entropy in critical systems going beyond area law in 1D. Due to its advanced structure, consisting out of alternating layers of isometries and disentanglers (see Fig. 2.11c), it can be efficiently contracted, however, requires a complexity of at least $\mathcal{O}(m^7)$ or higher (depending on the exact structure) to be optimised [74].

The general idea of MERA can as well be extended towards two-dimensional systems [73]: The MERA in 2D, in contrast to the PEPS, is able to calculate expectation values exactly while satisfying area law [121]. However, it suffers from a critical high numerical complexity $\mathcal{O}\left(m^{16}\right)$ making it unfeasible in practical applications for high dimensional systems. Thus, the application of the network in 2D problems has decreased significantly in the last years with the increasing competition of PEPS and TTN. In principle, the MERA can be defined for systems with higher dimensionality as well, however, as already seen by going from one to two dimensions the numerical scaling increases unfavourably. Examples, where the MERA is currently playing an important role, can be found in quantum gravity [75, 76, 82, 140–142], where it is linked to the geometry of space, e.g. via Anit-de Sitter (AdS) or Conformal Field Theory (CFT).

Tree Tensor Network — As a good compromise, TTNs, with their hierarchical network geometry, benefit from a reasonably low numerical complexity and allow (polynomial-scaling) universal contraction schemes [24] to calculate expectation values exactly. Both, the exact contraction and the optimization algorithms scale with $\mathcal{O}(m^4)$, which in turn allows typical bond-dimensions to exceed $m \geq 1000$. Moreover, a TTN is fairly straight-forward to implement and not restricted to any dimensionality of the underlying system, thus the extension to 3D systems is theoretically straightforward. With its flexible geometry, the TTN has proven itself to be a stable alternative with its particular strong-points ranging from applications in gapped 1D systems with periodic boundary conditions [80, 143] over two-dimensional systems with large local dimensions d [35, 37] up to 3D systems [144]. However, TTNs fail to capture the area law for higher-dimensional systems and thus eventually fade with increasing system size [145].

In Chapt. 3, the TTN and its implementation for high-dimensional systems is presented in more details featuring a formal definition, typical structures, its optimisation, the calculation of observables and most critically the practical hurtles when extending this Tensor Network to higher dimensions.

Augmented Tree Tensor Network — Within the scope of this thesis, we introduce a novel Tensor Network geometry, which augments the TTN approach in a way that the area law is indeed encoded (see Chapt. 5). Thus this *augmented Tree Tensor Network* (aTTN) can be used to efficiently simulate quantum many-body problems in any spatial dimension with a numerical complexity of $\mathcal{O}(m^4)$ while encoding the entropic area law for the state ψ it represents.

2.2.3   Tensor Network operators

In the language of Tensor Networks, we typically describe operators which can be applied to a Tensor Network state as *Tensor Product Operators* (TPOs). In accordance with the Tensor Network representation of a quantum state, TPOs represent general operators $\mathcal{T}$ in a decomposed Tensor Network form following their introduction in Ref. [24, 146]:

$$(\mathcal{T})^{\{i_j\}}_{\{i'_j\}} = \sum_{\{\gamma_\nu\}} \prod_j (\mathrm{t}^{[j]})^{\{\gamma_{\nu'}\}}_{i_j, i'_j} \ . \tag{2.4}$$

Thereby, the j-th tensor $\mathrm{t}^{[j]}$ is acting locally on the site i_j and is connected by the links $\{\gamma_{\nu'}\}$ to other tensors within the TPO. Thus, the complete TPO acts on the physical sites $\{i_j\}$. Thereby, they are a Tensor Network following Def. 2.2.1, but instead of having one physical link for each physical site i_j as the Tensor Network states have, the TPOs have two physical links i_j and i'_j for each local site they address. In this TPO formalism, we can describe for instance a local observable $\mathcal{T}_i$, an interaction part $\mathcal{T}_{\{i_j\}} = \mathcal{H}_p$ of a Hamiltonian $\mathcal{H} = \sum_p \mathcal{H}_p$, a string observable $\mathcal{T}_{\{i_j\}}$ or more general structures like a *Matrix Product Operator* (MPO) [61, 147, 148] or a *Projected Entangled Pair Operator* (PEPO) [149–151]. Fig. 2.12 illustrates some of the different TPOs encountered in the Tensor Network analysis of typical quantum many-body systems. For the sake of computational efficiency, it is important to keep the number of links and their internal bond-dimension κ reasonably small.

Hamiltonian as a TPO — A TPO, or a set of TPOs, can for instance be used to represent the Hamiltonian $\mathcal{H} \in \mathscr{H}$ of a quantum system in a suitable way for Tensor Network
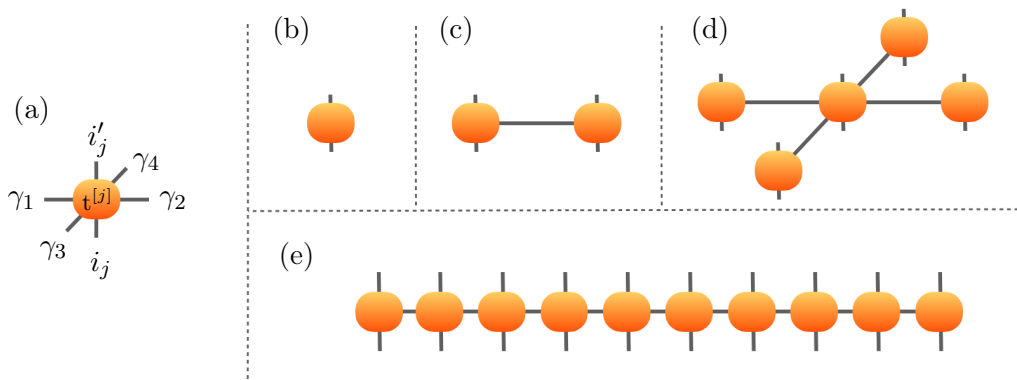
FIGURE 2.12: Different Tensor Network Operators (TPO) for typical quantum many-body systems in one and two dimensions. (a) Each TPO is composed out of local tensors t which are applied to different physical links i_j of a Tensor Network, i.e. to a different local Hilbert space. (b)-(e) Different examples of operators represented by a TPO: (b) a local operator, (c) a two-body operator, or correlator, (d) a star operator modelling a possible interaction of in a two-dimensional system, and (e) a Matrix Product Operator (MPO) which is frequently used for modelling an Hamiltonian in MPS analysis.

calculations. In fact, the computation of $\langle \mathcal{H} \rangle_\psi$ of the complete Hamiltonian $\mathcal{H} \in \mathscr{H}$ requires an exponentially costly computation, even though having a Tensor Network representation. However, the complete Hamiltonian is often a product of interactions $\mathcal{H} = \sum_p \mathcal{H}_p$. Thus, every interaction $\mathcal{H}_p$ can be described as a Tensor Product Operator (TPO). In Sec. 3.2.3, such a description of the Hamiltonian via a set of TPOs is illustrated in a hands-on example showing the optimisation procedure of a two-dimensional Tree Tensor Network.

2.2.4 Accessible information in Tensor Networks

While Tensor Network can efficiently represent a quantum state by exploiting the underlying entanglement structure, it may, for a general Tensor Network geometry, no be easy to extract information incorporated in the network. Indeed, the ability to access the information is a crucial problem for designing Tensor Network geometries. What use does a Tensor Network have which perfectly determines a desired quantum many-body state but does not allow to efficiently probe and characterise this exact state? In the end, we want to be able to extract information about the state from the Tensor Network! To address important this point, let us have a look at the typical quantities of interest for analysing quantum systems with Tensor Networks.

Expectation values $\langle \psi | O | \psi \rangle$ — Calculating expectation values of an observable O is arguably the most important way of extracting information about a quantum many-body state from a Tensor Network representation. Surely, the first important observable in the analysis of a quantum system would be its Hamiltonian $\mathcal{H}$. As introduced in Sec. 2.2.3, $\mathcal{H}$ is typically represented by *Tensor Network Operators* which enables to compute the energy expectation value $\langle \mathcal{H} \rangle_\psi$ in practical applications.

Besides $\mathcal{H}$, we are typically interested in three types of observables: *Local observables* $O = O^{[s]}$ acting on a single local site s and which are frequently used to identify local order parameters, *correlators* $O = O^{[s_1,s_2]}$ acting on two sites and which are crucial for investigating long-range order, and more general *string observables* $O = O^{[s_1,s_2,...,s_\omega]}$ acting ω sites of the

system which are important for instance to detect topological properties.

Depending on the geometry such expectation values can either be computed exactly or have to be approximated. A prominent example of the latter is the PEPS which is not *efficiently contractible*, i.e. the network cannot be contracted with itself in polynomial time as a function of its size.

State Overlap $\langle \psi^{[1]} | \psi^{[2]} \rangle$ — The overlap of two states $|\psi^{[1]}\rangle$ and $|\psi^{[2]}\rangle$, or their fidelity $|\langle \psi^{[1]} | \psi^{[2]} \rangle|^2$, are measures of interest when comparing two physical states. In practice, this computation is performed by contracting the two networks representing $|\psi^{[1]}\rangle$ and $|\psi^{[2]}\rangle$, respectively, over their coinciding physical links. Tensor Networks that are efficiently contractible can easily assess this information while for networks, such as PEPS, this contraction can become challenging with increasing system size.

Consequently, the norm of a state $\langle \psi | \psi \rangle$ can be computed in the same manner. Additionally, for performing this computation, the Tensor Network gauge freedom can often be exploited to further decrease the numerical costs [24, 61].

Entanglement properties — As discussed in Sec. 2.2.1, the entanglement plays a fundamental role in quantum systems and in the Tensor Network representation. Therein, it was elaborated that a bipartition of the system corresponds to a bipartition of the Tensor Network into two sub-networks which may be connected over several links. We can measure the amount of entanglement captured in a Tensor Network for any bipartition by contracting one of the corresponding two sub-networks. This contraction allows reconstructing the spectrum of the reduced density matrix and thereby enabling to measure the *Von Neumann entropy* as well as all *Rényi entropies* [22, 126, 152]. However, as the contractions above, depending on the Tensor Network geometry, such a contraction of an arbitrary sub-system can become infeasibly hard. In contrast, *loop-less* Tensor Networks flourish in this task (see Chapt. 3). By taking advantage of a proper isometrisation, they can directly be rewritten in terms of the Schmidt-decomposition making it computationally easy to extract the Von Neumann entropy (see Sec. 3.2.2).

2.2.5 Tensor Network algorithms

Tensor Networks have a broad range of applications in Quantum Physics and beyond. In what follows, the most prominent algorithms and problems tackled by employing Tensor Networks are described.

Energy minimisation

Finding the ground state of an interacting quantum many-body system is often a challenging task for numerical tools. For Tensor Networks, this can be done very efficiently when the network is properly designed to (i) obey the area law of the system it represents, and (ii) to provide a reasonably low numerical complexity which is mainly governed by the contractions required to minimise the energy. Given the Hamiltonian $\mathcal{H}$, we optimise the variational parameters of the Tensor Network wavefunction ψ in order to find the ground state of the system by minimising the energy

$$E = \langle \psi | \mathcal{H} | \psi \rangle \ . \tag{2.5}$$

Here, the common advantage of Tensor Networks is that all single tensors within the network can be treated as independent variables and thereby the key idea is to minimize one, or only a limited subset, of the tensors at a time. This procedure shifts the global, exponentially large, optimisation problem to a set of local minimisations in a drastically smaller subspace which is controlled in size by the chosen bond-dimension of the network. Still, based on this idea, we can either perform an eigenvalue minimisation, a gradient descent method or an imaginary time evolution. In Sec. 3.1.4, we will see a more technical illustration of this approach for Tree Tensor Networks.

Time Evolution

Especially in one dimensional systems, Tensor Networks are frequently deployed to investigate non-equilibrium dynamics by evolving a state in time under a unitary Hamiltonian evolution. In particular, there are two established methods for performing such a time evolution numerically. The first to be mentioned would be a class of algorithms, such as the Time-Evolving Block Decimation (TEBD) [62] and the tDMRG [63], which are based on decomposing the unitary evolution by means of a Suzuki-Trotter decomposition into a set of local operators to be applied to the network. Another remarkable method is the Time-Dependent Variational Principle (TDVP) [143, 153] which exploits mathematical concepts of differential calculus to compute a time evolution based on the geometry of the network. While these techniques have a broad range of applications such as performing various quenches [154–157], controlled dynamics [158] or, using imaginary time, quantum annealing, these applications are currently limited to either small system sizes in one-dimension or short time-scales.

Open quantum systems

Another remarkable application of Tensor Networks in quantum many-body physics is the representation of mixed states in open systems [64, 65, 159–161]. In fact, Tensor Networks can be deployed to investigate steady states or analyse dissipative time evolutions governed by a Lindblad master equation. The idea behind this application is, instead of representing the wave function of a pure state, to efficiently decompose the exponentially large density matrix of a mixed many-body state via Tensor Networks. The underlying master equation dynamics in then implemented by a direct integration [65, 160] or using stochastical methods [161]. This line of research is quite recently explored and might be a promising contender among others for simulating and validating the dissipation in novel quantum technologies such as quantum computers.

2.3   Symmetries in Tensor Networks

In the prior section, we got to know that by taking advantage of entanglement properties, Tensor Networks can efficiently compress the information content of a quantum many-body system. Even though a Tensor Network efficiently represents a quantum many-body state, the underlying simulations still rapidly reach the computational limits of current classical machines, especially for highly entangled systems. One of the main benefits for Tensor Network analyses which comes from present symmetries in a quantum system is that they allow restricting the degrees of freedom of the system as each symmetry obeys a well-defined internal structure. Thus, exploiting the symmetries of a quantum system is a powerful way to reduce the required numerical resources and helps significantly to push forward the capabilities of numerical methods for studying quantum many-body systems. Furthermore, incorporating symmetries enables the investigation of physical properties concerning certain selected symmetry condition, e.g. analysing a system with a fixed global charge or with a pre-defined number of particles. Consequently, the incorporation of symmetries in Tensor Networks has become an essential add-on for state-of-the-art computations of quantum many-body systems. In particular, Tensor Networks enables an efficient implementation of *pointwise symmetries*. However, before going into the more technical details, let us have a look at the importance of symmetries in general in quantum systems.

In fact, symmetries are omnipresent not only in quantum systems but furthermore in nature, ranging from the structure of molecules where the underlying symmetry determines all internal interactions [162–164], over the biology of butterflies which have a higher probability of reproduction when their wings are perfectly symmetric [165, 166], up to animal responses where both behavioural and neurophysiological studies confirm that humans and other animals have a high sensitivity to reflection symmetries [167]. In quantum mechanics, or even more general in theoretical physics, the key role of symmetries is well established: Every conserved physical quantity is linked to a present, underlying symmetry in the system which is also known as *Noether's theorem* [168]. Thus, in a quantum spin system in which for instance the total number of particles is conserved, there is a $U(1)$ symmetry present. Tab. 2.1 present some more examples of typical symmetries occurring in a quantum many-body system together with the corresponding conservation laws.

TABLE 2.1: Examples of Symmetries and corresponding Conservation laws in Physics

Conserved quantity		Symmetry group	
Parity	$\leftrightarrows$	$\mathbb{Z}_2$	(abelian)
Particle number	$\leftrightarrows$	$U(1)$	(abelian)
Spin or angular momentum	$\leftrightarrows$	$SU(2)$	(non-abelian)
Spin and particle number	$\leftrightarrows$	$SU(2) \otimes U(1)$	(non-abelian)
Flavour conservation	$\leftrightarrows$	$SU(3)$	(non-abelian)

When dealing with symmetries, the underlying group theory ultimately implies that the tensors within a Tensor Network can be decomposed into a *structural* and *degeneracy* part. The former is determined by the symmetry itself while the latter contains the remaining degrees of freedom unconstrained by the underlying symmetry. Therefore, the presence of

symmetries reduces the full problem into a set of decoupled, simpler sub-problems. Consequently, incorporating symmetries in Tensor Networks is a very deep, multidisciplinary topic and ranges from the basic concepts of group theory, over its application in quantum physics up to the technical numerical implementation in a Tensor Network.

The following sections give a general overview from the basics of symmetries in quantum systems to presenting the main mathematical concepts from representation theory which lead to the efficient representation of a symmetrically invariant Hamiltonian $\mathcal{H}$. Building on these fundamentals, the implementation of symmetries in Tensor Networks is described in more detail including the description of the main effects on the implementation. Most parts of this section are based on Ref. [108] which provides a more technical thesis on the incorporation of *non-Abelian* symmetries in Tensor Networks. Other comprehensive reviews on the subject of symmetries in Tensor Network include Ref. [24] for *Abelian* symmetries, Ref. [169] including the non-Abelian symmetry $SU(2)$, and the appendix of Ref. [81] for an in-depth description of $SU(N)$ symmetries in the context of Tensor Network simulations.

2.3.1 Symmetries in Quantum many-body systems

A present symmetry in a quantum many-body system manifests itself when the system is invariant under a set of transformations $\{W_g\}$. This set forms the symmetry group G where each group element W_g, i.e. each symmetry operation, can be represented by a matrix $\Gamma(W_g)$ which commutes with the Hamiltonian $\mathcal{H}$ of the symmetrically-invariant quantum system:

$$[\mathcal{H}, \Gamma(W_g)] = 0 \quad \forall W_g \in G \quad .$$

Given the present symmetry, the Hamiltonian can be transformed into a symmetric eigenbasis providing an internal structure that consequently decomposes the Hamiltonian in a beneficial block-diagonal shape to be exploited when analysing the system. Depending on whether all $W_g \in G$ commute with each other, the symmetry is called *Abelian* or *non-Abelian*:

$$\begin{array}{cc} \textit{Abelian} & \textit{non-Abelian} \\ [\Gamma(W_g), \Gamma(W_k)] = 0 \;\; \forall (W_g, W_k) \in G & \exists (W_g, W_k) \in G : \;\; [\Gamma(W_g), \Gamma(W_k)] \neq 0 \end{array}$$

This differentiation has an impact from a theoretical point of view as well as from a computational perspective, as we will see later on. In a nutshell: An *Abelian symmetry* is easier to handle with a straight-forward internal structure while a *non-Abelian symmetry* can become overwhelmingly complex with a non-trivial internal structure which on the other hand typically constrains a higher number of degrees of freedom and thus can offer a remarkable computational benefit when exploited.

Before moving forward with describing these typical symmetries in more detail, the following section will recall the crucial principles of representation theory, a crucial part for the application of group theory in mathematics and physics, and illustrate how the presence of symmetry transforms the Hamiltonian of a system.

From Representation theory to the structure of symmetric Hamiltonians

The idea in representation theory is to describe a group G by a homomorphic *representation* $\Gamma(G)$, i.e. by an element-wise one-to-one correspondence, which is suitable for linear algebra operations. It is convenient to choose the group representation as a set of matrices, where

each group element $W_g \in G$ is represented by a corresponding matrix $\Gamma(W_g)$ and the group multiplication is defined by the matrix multiplication.

Theorem 2.3.1: *(Reducibility).*

Every group representation $\Gamma(G)$ can either be decomposed by a transformation M into a block-diagonal form such that

$$M^{-1}\Gamma(W_g)M = \begin{pmatrix} \Gamma_1(W_g) & 0 & 0 \\ 0 & \Gamma_2(W_g) & 0 \\ 0 & 0 & \ddots \end{pmatrix} \quad \forall g \in G$$

where each set of *reduced* matrices $\Gamma_k(G)$ again forms a valid representation of the original symmetry group G, or if no such transformation exists, the representation Γ is called an *irreducible representation* (irrep). Taking this further, each *reducible representation* Γ can be decomposed into its irreps $\Gamma^{[j]}$:

$$\tilde{M}^{-1}\Gamma\tilde{M} = \bigoplus_j \left(\mathbb{1}_{d_j} \otimes \Gamma^{[j]} \right) \equiv \begin{pmatrix} \Gamma^{[1]} & & & & & \\ & \ddots & & & 0 & 0 \\ & & \underbrace{\Gamma^{[1]}}_{d_1} & & & \\ & & & \Gamma^{[2]} & & \\ 0 & & & & \ddots & 0 \\ & & & & & \underbrace{\Gamma^{[2]}}_{d_2} \\ 0 & & & 0 & & \ddots \end{pmatrix} \tag{2.6}$$

Consequently, when working in the irrep subspaces, each sector j describes an independent subspace, which can again be separated into the degeneracy space $\mathbb{D}_j$ and a structural space $\mathscr{H}_j$. Thus, the total vector space $\mathscr{H}$ naturally decomposes into the direct sum $\oplus$ of irreducible sub-spaces $\mathscr{H}_j$. Each irreducible subspace corresponds to an irrep $\Gamma^{[j]}$ which in return defines the dimension m_j of $\mathscr{H}_j$. The label j of the various sub-spaces is often referred to as *symmetry sector*, or in the framework of quantum mechanics as *primary quantum number*. The label d_j further indicates the degeneracy - also referred to as (outer) multiplicity - of the irrep $\Gamma^{[j]}$ within the block-diagonalised form of Γ.

In conclusion, we can describe the complete space $\mathscr{H}$ in a decomposed basis $|j, t_j, m_j\rangle = |j, t_j\rangle \otimes |j, m_j\rangle$ where the symmetry sector j labels the subspace $\mathscr{H}_j$, $t_j \in \{1, ..., d_j\}$ labels the degeneracy space $\mathbb{D}_j$, and m_j labels the internal states within the subspace $\mathscr{H}_j$. The labels in the new symmetry basis are typically assigned to quantum numbers of the system: As an example, the irrep label j corresponds to the spin j in case of spin conservation, or equally $SU(2)$ symmetry (with $j = n/2, n \in \mathbb{N}_0$).

Lemma 3 (Schur's Lemma). *Any matrix T, which commutes with all elements $\Gamma^{[j]}(W_g)$ of an irrep $\Gamma^{[j]}$, is proportional to the identity.*

$$T\Gamma^{[j]}(W_g) = \Gamma^{[j]}(W_g)T \quad \forall W_g \in G \quad \Leftrightarrow \quad T = c \cdot \mathbb{1}$$

Now, let us have a look at the transformation of the Hamiltonian $\mathcal{H} \to \tilde{\mathcal{H}} = M^{-1}\mathcal{H}M$ of a quantum system into its symmetry basis. In this basis, $\mathcal{H}$ maintains the permutation relations with the symmetry operations $\tilde{\Gamma} = \oplus_j \left(\mathbb{1} \otimes \Gamma^{[j]}\right)$ since

$$[\mathcal{H}, \Gamma] = 0 \;\; \Rightarrow \;\; \tilde{\mathcal{H}}\tilde{\Gamma} := M^{-1}\mathcal{H}MM^{-1}\Gamma M = M^{-1}\mathcal{H}\Gamma M = M^{-1}\Gamma \mathcal{H} M = \tilde{\Gamma}\tilde{\mathcal{H}} \; .$$

In order to fulfill this relation, the Hamiltonian has to exhibit the same block-diagonal shape, namely $\tilde{\mathcal{H}} = \oplus_j \left(\tilde{\mathcal{H}}_j \otimes I_j\right)$. Now, I_j has to commute with all elements within the irrep $\Gamma^{[j]}$ and therefore, due to *Schur's Lemma*, equals the identity $I_j = \mathbb{1}$, while $\tilde{\mathcal{H}}_j$ can be an arbitrary $d_j \times d_j$ matrix containing the remaining physical degrees of freedom which are unconstrained by symmetry. In conclusion the system's Hamiltonian can be decomposed into the sub-spaces $\mathscr{H}_j$ and the d_j-fold degeneracy spaces $\mathbb{D}_j$ corresponding to the $\tilde{\mathcal{H}}_j$s.

$$\tilde{\mathcal{H}} = \bigoplus_j \underbrace{\begin{pmatrix} H_{11}^j & H_{12}^j & \cdots & H_{1d_j}^j \\ H_{21}^j & H_{22}^j & \cdots & H_{2d_j}^j \\ \vdots & \vdots & \ddots & \vdots \\ H_{d_j1}^j & H_{2d_j}^j & \cdots & H_{d_jd_j}^j \end{pmatrix}}_{= \tilde{\mathcal{H}}_j \;\; \in \mathbb{D}_j} \otimes \underbrace{\begin{pmatrix} 1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix}}_{\mathscr{H}_j} \tag{2.7}$$

While the set of occurring symmetry sectors j and their degeneracy d_j is a property of the particular representation Γ, the irreps $\Gamma^{[j]}$ and thereby the irreducible sub-spaces $\mathscr{H}_j$ depends only on the symmetry group G itself. In the case of an Abelian symmetry, all irreps $\Gamma^{[j]}$ are scalars, leaving Γ in a dense block-diagonal form, whereas in the non-abelian case these irreps, in general, form quadratic matrices providing a further, non-trivial internal structure within each block j.

Representing symmetries in quantum systems

Very common symmetries in quantum mechanics are cyclic symmetry groups, in particular, $U(1)$, $SU(N)$ or $\mathcal{Z}_n$. The group elements $W(\mathbf{g})$ of such symmetries can be represented by an exponential map

$$W(\mathbf{g}) = e^{i\mathbf{g}\cdot\mathbf{J}} = \exp\left(i\sum_k g_k J_k\right) \tag{2.8}$$

where $\mathbf{J}$ denotes the so-called generators, a set of hermitian matrices, which together with a vector of parameters $\mathbf{g}$ create each group element $W(\mathbf{g}) \in G$. The corresponding group operation $W(\mathbf{h}) \otimes W(\mathbf{f})$ for this unitary representation is the matrix-multiplication $W(\mathbf{g}) = W(\mathbf{h}) \times W(\mathbf{f})$. In order for the multiplication of two group elements $W(\mathbf{h}) \in G$ and $W(\mathbf{f}) \in G$ to result in another element of the group $W(\mathbf{g}) \in G$, the generators shall satisfy the commutator relations

$$[J_k, J_l] = \sum_n f_{kln}J_n \; , \tag{2.9}$$

with f_{kln} being a structure constant defining the underlying algebra.

A special case of these groups are *Lie-Groups* which are used to describe continuous, unitary symmetries, such as $SU(2)$ in the case of spin conservation or $U(1)$ for a conserved

number of total particles. In these cases the number of group elements $W(\mathbf{g}) \in G$ is infinite, however, they can still be represented by a finite set of generators $\mathbf{J}$.

The following paragraphs present in more detail some of the most important Abelian and non-Abelian symmetry groups frequently encountered in quantum systems.

Abelian symmetries   Abelian symmetries can be represented by a single scalar generator J by using one group element parameter g (dim$\{\mathbf{g}\} = 1$). Thus, fulfilling the generator condition in Eq. (2.9) becomes trivial and following the chosen representation of Eq. (2.8), each element can be created by $W(g) = e^{iJg}$. The most important Abelian symmetries for quantum systems are the discrete symmetry group $\mathcal{Z}_n$ and the continuous rotation group $U(1)$:

(i) $\mathcal{Z}_n$ **(Discrete rotation)**. This cyclic group describes the rotation symmetry of a regular polygon with n edges. As example, $\mathcal{Z}_4$ denotes the symmetry of a square, $\mathcal{Z}_5$ of a pentagon, $\mathcal{Z}_6$ a hexagon, etc.. This group has n elements $W(g) \in \mathcal{Z}_n$, one for each rotation leaving the system invariant. Each element $W(g)$ can be created by the group parameters $g \in \{0, ..., n-1\}$ and the generator $J = 2\pi/n$ following the chosen representation of Eq. (2.8), i.e. $W(g) = e^{-i2\pi g/n}$. It is worth to mention that a valid, however non-unitary representation of this group would also be the function $W(g) = g$ under the group operation $g \otimes g' = (g + g') \mod n$. A particularly relevant case of this group in quantum mechanics is the omnipresent $\mathcal{Z}_2$, or *parity group*. It is either present directly as the single symmetry of a system as it is the case in the Ising model or in more advanced models of interacting Rydberg atoms (see Sec. 5.1), or it is implicitly present as a subgroup of higher order symmetry groups, such as $U(1)$ or $SU(2)$.

(ii) $U(1)$ **(Continuous rotation)**. This group describes the cyclic rotation symmetry of a circle. Thus, the number of group elements is infinite with its continuous group parameters $g \in [0, 1)$ and the generator $J = 2\pi$, i.e. $W(g) = e^{-i2\pi g}$. This Lie-group is connected to the conservation of integer quantities in physical systems, among others to particle conservation, and therefore present in systems such as Bose-Hubbard models [170]. Further, it occurs in various versions of the Heisenberg model (e.g. the XXZ model and the homogeneous Heisenberg model) as rotation invariance in the z-direction (see Sec. 5.1 or Sec. 6).

Non-Abelian symmetries   For non-Abelian symmetries the representation of Eq. (2.8) and the generator condition in Eq. (2.9) becomes non-trivial. The minimal choice of generators should together with the identity span the whole space $\mathscr{H}_j$ on which the corresponding symmetry operation acts ($W(\mathbf{g}) : \mathscr{H}_j \to \mathscr{H}_j$).

The most important non-Abelian symmetries for quantum systems are the special unitary groups $SU(N)$, in particular, the spin symmetry $SU(2)$ and the flavour symmetry $SU(3)$. All $SU(N)$ groups are Lie-Groups with the additional condition det$\{W(\mathbf{g})\} = 1$ which can be satisfied by choosing traceless generators J_k. For their comprehensive description, it is sufficient to define their *fundamental representation*, the simplest non-trivial irrep, since all other irreps can be created iteratively from there on by combining the given symmetry sectors [81, 169]. As an example: The fundamental representation of $SU(2)$ is the Spin-$\frac{1}{2}$ representation, while we can obtain e.g. a spin-1 by combining two Spin-$\frac{1}{2}$ into one Spin-1 and one Spin-0 or higher-order spin-representation by iteratively combining several spin-$\frac{1}{2}$ representations.

(i) $SU(2)$ **(Complex 2D-rotation)**. This Lie-group describes the rotation symmetry within a two-dimensional complex space, such as the angular momentum rotation symmetry of a quantum spin systems. It is a manifold out of two $SO(3)$ rotation groups, or "Sphere-symmetries", i.e. there is a surjective, two-to-one, homomorphism from $SU(2)$ onto $SO(3)$. This symmetry is connected to the spin conservation and occurs for instance in the homogeneous Heisenberg model (Sec. 5.1 or Sec. 6), the bilinear-quadratic Heisenberg model [171] or in Bose-Hubbard models [170].

The *defining representation* can be conveniently chosen to be the *Pauli matrices* σ:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \ , \quad \sigma_y = \begin{pmatrix} 0 & \text{-}i \\ i & 0 \end{pmatrix} \ , \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & \text{-}1 \end{pmatrix}$$

Obeying the commutator relation $[\sigma_k, \sigma_l] = \sum_n 2i\epsilon_{kln}\sigma_n$ with ϵ_{kln} being the Levi-Cevita constant, the Pauli matrices fulfill the condition of Eq. (2.9) and since they span the complete space of two-dimensional matrices together with the identity $\mathbb{1}_2$, they are commonly used as generators $\mathbf{J} = \frac{\sigma}{2}$. In quantum spin systems, this representation corresponds to the spin-$\frac{1}{2}$ algebra. Further, each irrep $\Gamma^{[j]}$ of the $SU(2)$ symmetry can be used to represent the $2j + 1$-dimensional space corresponding to the spin-j multiplet algebra.

(ii) $SU(3)$ **(Complex 3D-rotation)**. The $SU(3)$ group denotes the rotation symmetry within a three-dimensional complex space, such as the flavour symmetry in particle physics. It includes among others $U(1)$, $SU(2)$, $SO(2)$ and $SO(3)$ as subgroups. This symmetry is famously connected to particle physics and occurs in the Standard model [8] as well as in some spin model such as the bilinear-quadratic Heisenberg model at the Lai-Sutherland point [171].

For this symmetry, the *Gell matrices*

$$J_{x,\{1,2\}} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \ , \quad J_{x,\{2,3\}} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \ , \quad J_{x,\{1,3\}} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \ ,$$

$$J_{y,\{1,2\}} = \begin{pmatrix} 0 & \text{-}i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \ , \quad J_{y,\{2,3\}} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & \text{-}i \\ 0 & i & 0 \end{pmatrix} \ , \quad J_{y,\{1,3\}} = \begin{pmatrix} 0 & 0 & \text{-}i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix} \ ,$$

$$J_{z,1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \text{-}1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \ , \quad J_{z,2} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \text{-}2 \end{pmatrix}$$

offer a convenient defining representation, obeying the commutator relation of Eq. (2.9) and together with the identity $\mathbb{1}_3$, they span the complete space of three-dimensional matrices. In quantum spin systems, this representation corresponds to the flavour algebra. Crucially, in contrast to the $SU(2)$ symmetry, the $SU(3)$ symmetry exhibits an inner multiplicity for higher dimensional irreps, i.e. high order multiplet state spaces, making it quite challenging to implement [81, 108].

It is worth to mention that even more complex symmetries, such as $SO(6)$ may occur in spinor systems. However, incorporating such symmetries, in particular once with inner

multiplicity, is a highly challenging task in practical applications and further, the number of occurrences is typically limited in ordinary quantum many-body systems compared to the above-presented symmetry groups. Still, when having a Tensor Network framework capable of handling for instance $SU(2)$ symmetries, an $SO(6)$-symmetric system can be simulated more efficiently by exploiting its $SU(2)$ subgroup compared to the straightforward approach without respecting any symmetries.

Further, even though a non-Abelian symmetry might be present in the system to analyse, it can as well be convenient for the sake of further simplicity in the implementation to reduce the non-Abelian to an Abelian symmetry - or a set of multiple Abelian symmetry. As an example, a general $SU(N)$ symmetry can be conveniently reduced to a set of $N-1$ independent Abelian $U(1)$ symmetries.

2.3.2   Incorporating Symmetries in Tensor Networks

By taking advantage of entanglement properties, Tensor Networks can efficiently compress the information content of a quantum many-body system. However, the present Tensor Network simulations still rapidly reach the computational limits of current classical machines. Additionally exploiting symmetries of the underlying quantum many-body system helps to overcome this bottleneck towards a higher performance of Tensor Network algorithms and further, it offers to target specific symmetry sectors. In particular, Tensor Networks enables an efficient implementation of pointwise symmetries.

The following section will describe the effect of a present, homogeneous, pointwise symmetry on the Tensor Network states and how to implement these symmetries efficiently on the level of every single tensor.

Pointwise symmetries

Global symmetries which act separably on the local sites of the quantum many-body state are defined as *pointwise symmetries*. Thus, each global symmetry operation

$$W_g = \bigotimes_{i=1}^{N} W_g^i$$

consists of the terms W_g^i which locally represent the group element $W_g \in G$ at site i. Thus, each local symmetry operation acts only on the degrees of freedom of one physical site i. For *homogeneous pointwise symmetries*, these local terms W_g^i are even equivalent and do not depend on the site i on which they act. These types of symmetries are exceedingly relevant from the physical point of view as they describe extensive physical constraints defined by *quantum numbers* such as total particle conservation, total parity, spin conservation or total magnetisation.

Symmetric Tensor Network geometry

Let the Tensor Network $|\Psi\rangle$ represent the wavefunction of a quantum many-body system which is invariant under the pointwise symmetry G. Thus, the state is invariant under the action of each symmetry element $W_g = \bigotimes_{i=1}^{N} W_g^i$, i.e. it fulfills

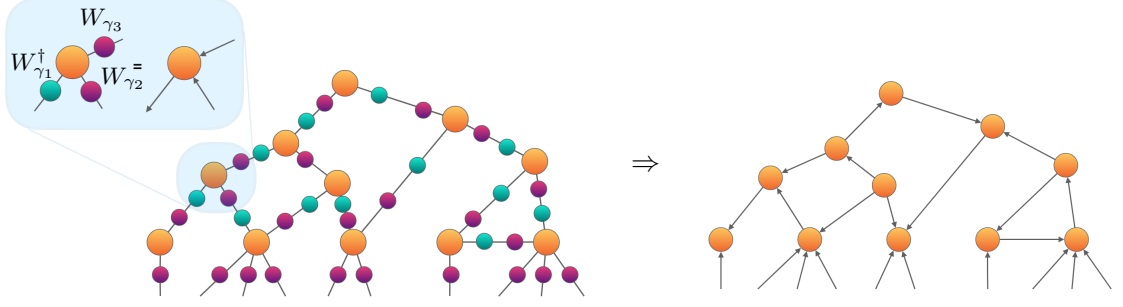$$W_g^{\otimes N}|\Psi\rangle = |\Psi\rangle \;\; \forall g \in G \;.$$

FIGURE 2.13:   Incorporation of Symmetries in a Tensor Network: A global pointwise symmetry operation $W_g = \bigotimes_{i=1}^N W_g^i$ is acting separately on the physical links (left: each local operation W_g^i illustrated in purple). Introducing a gauge freedom $W_\gamma W_\gamma^\dagger = \mathbb{1}$ on the internal links enables to transform each tensor into a symmetry basis. The contraction of W_γ (purple) or $W_\gamma^\dagger$ (turquoise) with the nearest tensor introduces a *direction* for the symmetry operations.

Now, the *gauge invariance* of the Tensor Network allows to transform every single tensor within the network into the symmetry basis of G: Introducing a unitary representation $\Gamma(W_g)$ together with its inverse $\Gamma^\dagger(W_g)$ at each internal link χ enables to map each tensor $\mathcal{T}$ into a *gauge-transformed* version leaving the global state unaltered (see Fig. 2.13). However, in this mapping, we introduce a clear orientation in the links depending on whether $\Gamma(W_g)$ or its inverse is applied to a certain tensor $\mathcal{T}$. Thus, for symmetrically invariant Tensor Networks, a link direction is defined for each tensor where a link is *incoming* when it is transformed into the symmetry space $\mathcal{D} \ni \Gamma(W_g)$ or is *outgoing* when acting on its duplex space $\mathcal{D}^\dagger \ni \Gamma^\dagger(W_g)$. For sake of compactness, we refer to the applied action on the tensor as X_g with

$$
X_g^{(k)} = \left\{
\begin{array}{lll}
\Gamma(W_g) & \Leftrightarrow & \mathfrak{D}_k = \text{``incoming''} \\[2ex]
\Gamma(W_g^{-1}) & \Leftrightarrow & \mathfrak{D}_k = \text{``outgoing''}
\end{array}
\right.
\tag{2.10}
$$

given that $\Gamma^\dagger(W_g) = \Gamma(W_{g'}) \in G$ is a valid representation of another group element g' of the symmetry G. Consequently, each tensor $\mathcal{T}$ of the Tensor Network $|\Psi\rangle$ is transformed into the symmetry basis of G, i.e. $\mathcal{T}$ is invariant under the symmetry operations of G:

$$
\sum_{i_1,\dots,i_r} \prod_{r=1}^R [X_g^{(r)}]_{i'_r,i_r} \mathcal{T}_{i_1,\dots,i_R} = \mathcal{T}_{i'_1,\dots,i'_R} \quad \forall g \in G \ .
$$

where $X_g = \bigotimes_{r=1}^R X_g^{(r)}$ with $X_g^{(r)} \in G$ represents the g-th element of the pointwise symmetry for the tensor $\mathcal{T}$. Thus, all symmetry actions X acting on each single tensor $\mathcal{T}$ can be interpreted as an independent symmetry group for the vector space spanned by $\mathcal{T}$.

Since the global gauge transformation $\Gamma^\dagger(W_g)\Gamma(W_g) = \mathbb{1}$ of the Tensor Network is performed locally on all internal link, each of the internal links implicitly contains information about the symmetry group. While W_g is fixed by the underlying physics of the many-body system, the representation $\Gamma(W_g)$ can be freely chosen for each internal link allowing to model $\Gamma(W_g)$ depending on the physical problem. In a simulation algorithm, this allows for instance to dynamically adapt the symmetry sectors for each internal link individually.

Consequently, a global symmetry can be exploited point-wise addressing each tensor separately. In the following, we describe the influences of present symmetric on each tensor and how the tensor operations presented in section 2.1 are influenced.

(a) Schur's Lemma

(b) Wigner-Eckart theorem
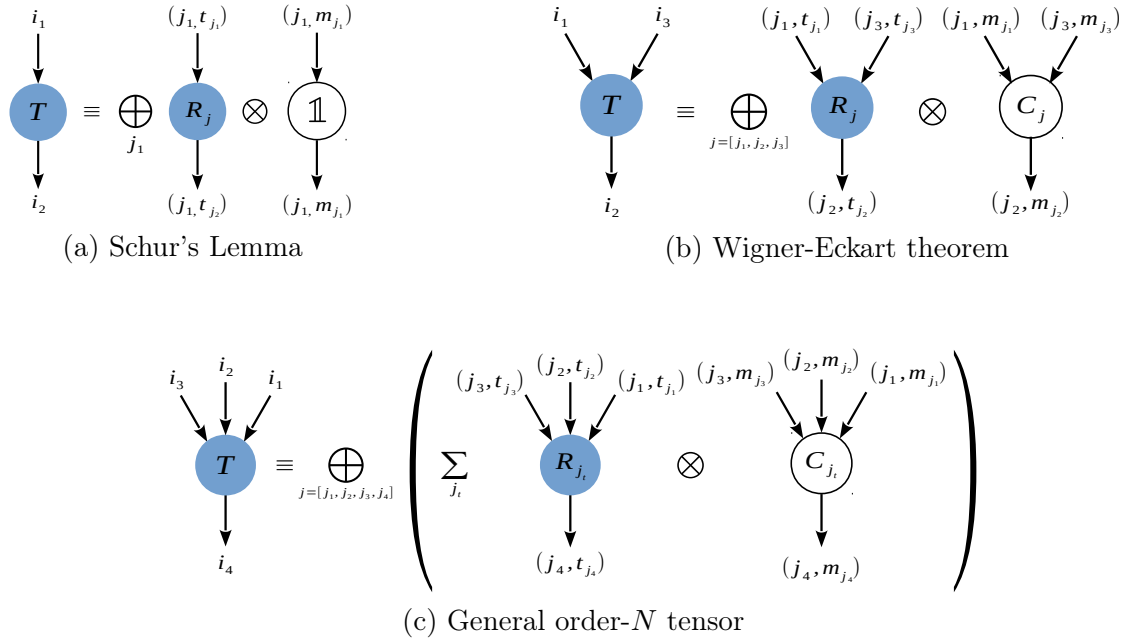
(c) General order-N tensor

FIGURE 2.14: The inner structure of symmetric tensors. An order-2 tensor obeys Schur's Lemma with its structural components being identities (a). The order-3 tensor obeys the Wigner-Eckart theorem, coupling only certain sectors $j = [j_1, j_2, j_3]$ allowed by symmetry with its structural tensor C_j given by the Clebsch-Gordan-coefficients (b). A general order-N tensor decomposes into a block-diagonal form with each block further splitting in one degeneracy tensor and one structural tensor (c). Illustrations taken from Ref. [108].

Symmetric Tensors

Similar to the decomposition of a quantum many-body Hamiltonian in the presence of symmetry, a symmetrically invariant tensor $\mathcal{T}$ can be block-diagonalised by the transformation into the symmetry basis. This underlying principle for the Hamiltonian which is governed by Schur's Lemma can be expanded to the tensor algebra for arbitrary order-R tensors. Ultimately, the symmetric tensor decomposes into a set of *structural tensors* R_{s_j} which are entirely determined by symmetry and set of *degeneracy tensors* C_{s_j} which contains the unconstrained degrees of freedom (see Fig. 2.14). Thus, we can reduce the *physical* degrees of freedom which are not constrained by symmetry by isolating them from the symmetry constrains. Consequently, a general order-r tensor decomposes into

$$\mathcal{T} = \bigoplus_{j=[j_1,\ldots,j_r]} \left(\sum_{s_j=1,\ldots,n_j} \left[R_{s_j} \otimes C_{s_j} \right] \right) , \tag{2.11}$$

where n_j characterises the number of structural tensors C_{s_j} arising for one specific configuration of coupling symmetry sectors $j = [j_1, j_2, \ldots, j_r]$. While in general the combination of the sectors $j = [j_1, j_2, \ldots, j_r]$ leads to n_j-degenerate structural tensors, this degeneracy does not apply in the most crucial cases which are presented later on [108, 169, 172].

Link fragmentation — In a symmetrically invariant Tensor Network, we assign a representation X_g of the symmetry group G to each internal link χ.

As described in Sec. 2.1, the space $\mathscr{H}_k$ of each link i_k can be transformed into the symmetry basis decomposing the total space $\mathscr{H}_k$ into a set of independent subspaces $\mathscr{H}_{j_k}$, or symmetry sectors. Therefore, the link index i_k splits according to the irrep subspace expansion into the three symmetry labels $i_k \to (j_k, t_{j_k}, m_{j_k})$. Thus, each symmetric link index is described by a tuple of the symmetry sectors j_k, the degeneracy index $t_{j_k} \in \mathbb{D}_{j_k}$ and the index $m_{j_k} \in \mathscr{H}_{j_k}$ for the irrep basis vector space. As mentioned above, the dimension of the irrep vector space $d_{s,j_k} = \dim\{\mathscr{H}_{j_k}\}$, i.e. the number of allowed values of m_{j_k} is determined by the symmetry sector j_k while the dimension of the degeneracy space $d_{r,j_k} = \dim\{\mathbb{D}_{j_k}\}$ can be adjusted by the user. Therefore, the total link dimension is given by summing over all irreps of G

$$D_\chi = \sum_{j_k} d_{r,j_k} d_{s,j_k} \ .$$

In practice, we keep only a finite amount of symmetry sectors with upper-bounded degeneracy $d_{r,j_k} \le c$ for each sector j_k to ensure a finite link dimension. Additionally, $d_{r,j_k} = 0$ implies that the possible irrep j_k is not occupied and therefore there is no need in storing it explicitly in the link description. Further, we recall that for Abelian symmetries $d_{s,j_k} = 1 \ \forall j_k$ holds true. Thus, for Abelian symmetries the complete tensor decomposes simply into a block-diagonal structure with C_j being a constant for all sectors j, while for non-Abelian symmetries each block further exhibits an internal structure governed by the symmetry.

Order-Two Tensor, i.e. matrix — Equivalently to the Hamiltonian decomposition described in section 2.3.1, the decomposition of an order-2 tensor $\mathcal{T}_{i,k}$ obeys Schur's Lemma. Thus, when written in the symmetry basis $|j\rangle\langle k| \to |j, t_j, m_j\rangle\langle k, t'_k, m'_k|$, the tensor with incoming link j and one outgoing link k has to preserve the symmetry sector as well as the irrep label m (see Fig 2.14a). Consequently, the structural tensors $C_j = \delta_{j,k}\delta_{m_j m'_k}$ all become an identity:

$$\mathcal{T} = \bigoplus_{j=[j_1,j_1]} R_j \otimes \mathbb{1}_j \ . \tag{2.12}$$

This structure is one of the key points for the drastic speed-up when incorporating symmetries in Tensor Networks. It becomes evident that employing symmetries reduces the number of variational parameters while maintaining entanglement features since the overall bond-dimension remains unchanged. Indeed, the number of free parameters within the tensor in the symmetrically invariant basis equals

$$\vartheta(\mathcal{T}) = \sum_j d_{r,j}^{[\alpha]} d_{r,j}^{[\beta]}$$

with $d_{r,j}^{[\alpha]}$ being the degeneracy of the j-th symmetry sector within the fragmented link α. Thus, in practice, the number of free parameters $\vartheta(\mathcal{T}) \sim D_\chi << D_\chi^{[\alpha]} D_\chi^{[\beta]} \sim D_\chi^2$ is heavily reduced, especially when several symmetry sectors or higher-dimensional irrep spaces are occupied.

It is worth mentioning that the structural tensor may couple the irrep labels m differently when considering a tensor with two unidirectional links, thus still obeying $C_j = \delta_{j,k}\tilde{C}_j$.

Order-three Tensor — An order-3 tensor with two incoming and one outgoing is connected on the level of group theory to the tensor product of two irreps, say $\Gamma^{[j_1]} \in G$ and

$\Gamma^{[j_2]} \in G$. This product of two irreps $\Gamma^{[j_1]} \otimes \Gamma^{[j_2]} \in G$ is a group element of G and can be decomposed following Eq. (2.6) into

$$\Gamma^{[j_1]} \otimes \Gamma^{[j_2]} = \bigoplus_j \left(\mathbb{1}_{d_j} \otimes \Gamma^{[j]} \right)$$

where the inner direct sum runs over the number times the same irrep $\Gamma^{[j]}$ occurs within the representation (see Fig. 2.14b). The degeneracy d_j in the occurrence of $\Gamma^{[j]}$ is as well referred to *inner multiplicity*. In every Abelian symmetry and in many non-Abelian symmetries relevant for quantum mechanics, such as the $SU(2)$, the symmetry group G is free of such inner multiplicities, i.e. the symmetry sector j only occurs once when coupling two irreps. Taking the $SU(2)$ symmetry as an example, we obtain only one spin-1 and only one spin-0 when coupling the irreps of two spin-$\frac{1}{2}$. Further, each combination of two arbitrary spin multiplet subspaces always results in a set of undegenerate subspaces.

In these cases where no inner multiplicity is present, the structure of an order-3 tensor is determined by the *Wigner-Eckart-theorem* and follows the decomposed block-structure

$$\mathcal{T}_{i_1,i_2,i_3} = \bigoplus_{j=[j_1,j_2,j_3]} R^{[j]}_{t_{j_1},t_{j_2},t_{j_3}} \otimes C^{[j]}_{m_{j_1},m_{j_2},m_{j_3}} \ .$$

The structural parts $C^{[j]}$ are given by the *Clebsch-Gordan-Coefficients* of G. For different link directions, the order-3 tensor decomposes into the same structure with different structural parts $C^{[j]}$. Thus, given a multiplicity-free symmetry, each combination of three labels j_1, j_2 and $j_3 \in \{j_1 \otimes j_2\}$ which is not forbidden by the symmetry uniquely corresponds to one specific structural tensor C_j. This does not apply in the case of inner multiplicity, where the tensor however still decomposes into the general form of a symmetrically invariant tensor given in Eq. (2.11).

Further, the number of free parameters within a symmetrically invariant tensor $\mathcal{T}$ is given by

$$\vartheta(\mathcal{T}) = \sum_{j_1} \sum_{j_2} \sum_{j_3 \in \{j_1 \oplus j_2\}} d^{[\alpha]}_{j_1} d^{[\beta]}_{j_2} d^{[\gamma]}_{j_3}$$

with $d^{[\alpha]}_j$ being the dimension of the j-th symmetry sector of the fragmented link α. Thus, the number of free parameters $\vartheta(\mathcal{T}) << D^3_\chi$ is considerably smaller compared to the total number of parameters D^3_χ of the original tensor.

Influence on Tensor Network algorithms

As mentioned above, exploiting the presence of symmetry is crucial for an efficient Tensor Network algorithm. Similar to the fact that using the structure of a block-diagonal or sparse matrix in linear algebra leads to drastic performance gains, the block-structure of a symmetrically invariant tensor, and further, the decomposition into a degeneracy and a structural space for each block can be used for a crucial computational benefit.

In the following, we will describe the most important operations explaining the efficient exploitation of symmetries in Tensor Networks. In particular, we present the factorisation of a tensor and the contraction of two tensors within a symmetric framework.

Tensor Factorisation —   Beyond the decomposed structure of a symmetric tensor $\mathcal{T}$, the tensor factorisation further benefits from Schur's Lemma as they are ultimately executed in a matrix shape, i.e. as an order-2 tensor. In this shape, the factorisation of $\mathcal{T}$ can be executed on its degeneracy tensors R_j only since all its structural tensors $C_j = c_j \mathbb{1}_j$ are proportional to the identity. As an example, the QR-decomposition of $\mathcal{T}$ within a symmetric Tensor Network framework is given by

$$\mathrm{QR}(\mathcal{T}) = \bigoplus_{j=[j_1,j_1]} \left(\underbrace{\mathrm{QR}(c_j R_j)}_{= R_j^Q R_j^R} \otimes \mathbb{1}_{m_j} \right) = \underbrace{\bigoplus_{j=[j_1,j_1]} R_j^Q \otimes \mathbb{1}_{m_j}}_{= \mathcal{T}^Q} \underbrace{\bigoplus_{j=[j_1,j_1]} R_j^R \otimes \mathbb{1}_{m_j}}_{= \mathcal{T}^R}$$

where we assume $\mathcal{T}$ to be already in an order-2 shape. Note, that the factorisation of the reduced tensors $R_j = R_j^Q R_j^R$ can even be performed in parallel over the different symmetry sectors j. Thus, considering that the tensor holds a $(d_{r,k} \times d_{r,k})$-dimensional reduced tensor and a $(d_{s,k} \times d_{s,k})$-fold structural tensor for each coupled symmetry sector $[j_k, j_k]$, the numerical complexity of this operation drops from $\zeta = \mathcal{O}(D^3)$ for the non-symmetric framework down to

$$\zeta_{\mathrm{QR}} = \mathcal{O}\left(\max_k \{d_{r,k}^3\} \right) \; . \tag{2.13}$$

Equivalently, a valid SVD of a symmetric tensor can be performed on the reduced tensors only. This factorisation, like the QR-decomposition, scales with the largest dimension of all reduced tensors

$$\zeta_{\mathrm{SVD}} = \mathcal{O}\left(\max_k \{d_{r,k}^3\} \right) \equiv \mathcal{O}\left(\max_k \{D^\eta\} \right) \; . \tag{2.14}$$

and can be executed in parallel as the different symmetry sectors j are independent of each other. Compared to the bond-dimension D of the complete tensor, the complexity of this symmetric SVD of Eq. (2.14) can be reduced with $\eta = 3 \log_D(\max_k \{d_{r,k}\})$. The exact value η highly depends on the occupation of the symmetry sectors j and can range from $\eta \to 0$ for a flat distribution of symmetry sectors with $\max_k d_{r,k} = 1$ or in the worst case to $\eta = 3$ if only one trivial symmetry sector is present such that $\max_k d_{r,k} = D$. For Abelian symmetries where only trivial symmetry sectors occur, this factor typically lays around $\eta \sim 2$ while for some non-Abelian systems it can be realistic to reach $\eta < 1$. The latter however depends on the order of the symmetry, i.e. the size of the defining representation: Since the dimension of the irrep spaces grows faster with increasing symmetry label for higher-order symmetries, η on average decreases with increasing order of the symmetry. An additional speed-up can be achieved by considering several symmetries at once, e.g. $SU(2) \otimes U(1)$.

Contraction —   The symmetric Tensor Network framework further impacts the contraction of two symmetrically invariant tensors $\mathcal{T}^{[1]}_{i_1,\ldots,i_r}$ and $\mathcal{T}^{[2]}_{i'_1,\ldots,i'_\rho}$. Firstly, this process of contracting the two tensors over coinciding links $\{i_c\}$ can be performed independently for each matching symmetry sectors $\{j_c\}$. Further, when contracting the matching symmetry blocks of the two tensors their reduced and structural part can be contracted independently of each other as well:

$$\mathcal{T}^{[1]}_{i_1,\dots,i_r} \overset{\{i_c\}}{\odot} \mathcal{T}^{[2]}_{i'_1,\dots,i'_\rho} = \left[\bigoplus_{j=[j_1,\dots,j_r]} \left(R^{[1]}_j \otimes C^{[1]}_j \right) \right] \overset{\{i_c\}}{\odot} \left[\bigoplus_{j'=[j'_1,\dots,j'_\rho]} \left(R^{[2]}_{j'} \otimes C^{[2]}_{j'} \right) \right]$$

$$= \bigoplus_{[j,j']\setminus\{j_c\}} \left(\left[R^{[1]}_j \otimes C^{[1]}_j \right] \overset{\{j_c\}}{\odot} \left[R^{[2]}_{j'} \otimes C^{[2]}_{j'} \right] \right)$$

$$= \bigoplus_{[j,j']\setminus\{j_c\}} \left(\left[R^{[1]}_j \overset{\{j_c\}}{\odot} R^{[2]}_{j'} \right] \otimes \left[C^{[1]}_j \overset{\{j_c\}}{\odot} C^{[2]}_{j'} \right] \right)$$

where the operator $\overset{\{i_c\}}{\odot}$ denotes a contraction over the coinciding links $\{i_c\}$. In case, the resulting coupling of the symmetry sectors $[j,j'] \setminus \{j_c\}$ is forbidden by symmetry, the contraction of the structural part naturally results in a zero tensor, i.e. a tensor with each element equal to zero, as the two structural tensors are by construction orthogonal to each other. Recalling that the complexity of the contraction for general tensors is given by $\zeta = \mathcal{O}(D_k D_c D_{k'})$ with

$$D_k = \prod_{i_k \notin \{i_c\}} \dim\{i_k\} , \quad D_c = \prod_{\{i_c\}} \dim\{i_c\} , \quad D_{k'} = \prod_{i_{k'} \notin \{i_c\}} \dim\{d_{k'}\}$$

its symmetric counterpart scales with

$$\zeta_{\text{sym}} = \mathcal{O}\left(\max_j \{ \underbrace{D_k^{[j,m_j]} D_c^{[j,m_j]} D_{k'}^{[j,m_j]}}_{\text{structural tensors}}, \underbrace{D_k^{[j,t_j]} D_c^{[j,t_j]} D_{k'}^{[j,t_j]}}_{\text{reduced tensors}} \} \right) , \tag{2.15}$$

where $D_k^{[j,m]} = \prod_{i_k \neq i_c} \dim\{m_{j,k}\}$, $D_c^{[j,m]} = \prod_{i_c} \dim\{m_{j,c}\}$, $D_{k'}^{[j]} = \prod_{i_{k'} \neq i_c} \dim\{m_{j,k'}\}$

and $D_k^{[j,t]} = \prod_{i_k \neq i_c} \dim\{t_{j,k}\}$, $D_c^{[j,t]} = \prod_{i_c} \dim\{t_{j,c}\}$, $D_{k'}^{[j]} = \prod_{i_{k'} \neq i_c} \dim\{t_{j,k'}\}$

Thus, the complexity of contracting two symmetric tensors scales with the most expensive sub-contraction to be executed. In contrast to the factorisation, the size of the structural tensors, i.e. the size of the irrep symmetry-subspace, cannot be neglected in the determination of the numerical complexity. Consequently, the actual speed-up highly depends on the occupation of the symmetry sectors and the symmetry itself.

2.3.3 Finite symmetry sectors in Tensor Network states

Next to the significant improvement in computational performance, the incorporation of symmetries has a second crucial benefit: A symmetric Tensor Network enables to select a specific symmetry sector of the underlying system for the numerical analysis. This enables us to study systems by implicitly fixing conserved quantities in the variational problem. As an example, a pointwise $SU(2)$ symmetric Hamiltonian preserves the total spin j, i.e. global many-body states always couple irrep subspaces $|j, t_j, m_j\rangle$ belonging to the same symmetry sector j. Here, as a valid physical question, we could be interested in the ground state properties of a spin-$\frac{1}{2}$ lattice with a finite global spin j, say $j = \frac{5}{2}$. Analogously, we might be interested in
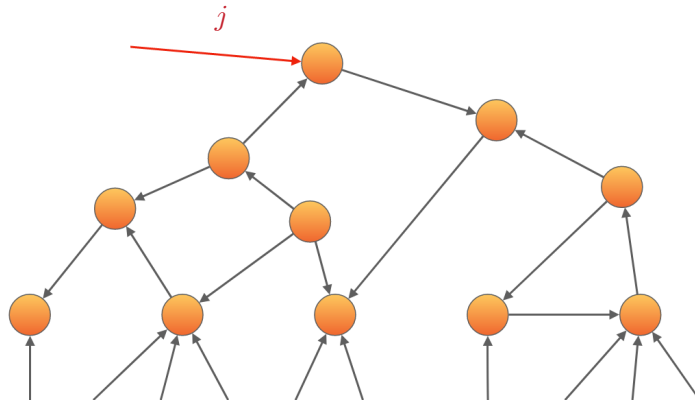
Figure 2.15: Finite symmetry sector in a Tensor Network state. The *selector link* (red) introduces the irrep $\Gamma^{[j]}$ of the desired global symmetry sector j.

describing the ground state, or even a time-evolution, for a $U(1)$ symmetrical Hamiltonian (corresponding to particle conservation) with a fixed number of particles in our system.

Achieving such a selection of a specific symmetry sector in numerical analysis is, in general, not a trivial task. Even when starting in the correct sector j and applying only symmetrically invariant operations, the simulation will ultimately introduce numerical errors which inevitably result in fluctuations exploring other symmetry sectors $j' \neq j$. Thus, the variational analysis has to be able to strictly enforce the symmetrically invariance in a way that is implicitly embedded in the numerical framework. This exactly is another strong point of incorporating symmetries in Tensor Networks. As illustrated in the following, we will be able to select specific symmetry sectors j for the global state which is implicitly respected within the symmetry Tensor Network framework.

Undoubtedly, the complete Tensor Network state $|\Psi\rangle$ is invariant under all action of the symmetry group G if all tensors of $|\Psi\rangle$ are symmetrically invariant. However, this state composed of a set of symmetrically invariant tensors would strictly belong to the global $j = 0$ sector of G, thus we refer to the state as *zero-sector* state $|\Psi^{[0]}\rangle$. Now, a *finite sector* state $|\Psi^{[j]}\rangle$ can be realised very easily by introducing an additional *selector link* into the zero sector state. This selector link shall hold only the desired global symmetry sector j with its irrep $\Gamma^{[j]}$. Attaching this selector link to a tensor within the directed, symmetrically invariant network $|\Psi^{[0]}\rangle$ enables to conveniently set the global symmetry sector: In this way, the network is strictly invariant under the symmetric transformation $W_g^{\otimes N}\Psi^{[j]}(\Gamma^{[j']})^{-1} = \Psi^{[j]}\delta_{j,j'}$ only if the sector j is applied [24, 172].

Tensor

auxiliary Bond-Link
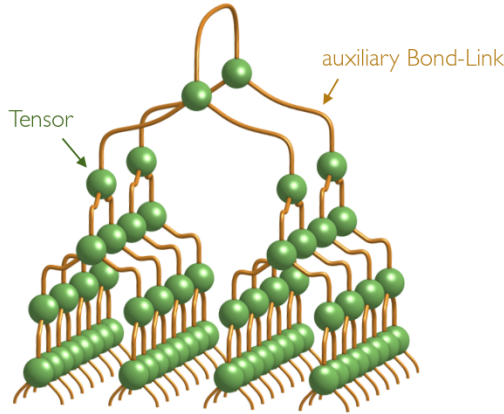
3

Tree Tensor Networks and their implementation for high-dimensional systems

The fundamental part of this thesis is the development of a Tree Tensor Network algorithm for investigating high-dimensional quantum many-body systems. Building on an existing library for Tensors invariant under Abelian symmetries, this algorithm includes the structure of the network as well as the main programs performing ground-state search and time evolution. In the last chapter which introduces different Tensor Network states and algorithms, it was stated that Tree Tensor Networks cannot efficiently encode the area law for quantum many-body systems in higher dimensions. Thus, the question arises whether they are actually a reasonable method for analysing such systems, especially compared to competing Tensor Networks, such as PEPS and MERA which can faithfully represent area law states, or even compared to other numerical methods, such as *Quantum Monte Carlo* (QMC).

In fact, QMC is one of the most prominent methods deployed for solving quantum many-body problems in higher dimensions systems [50–53]. However, there are a set of fundamental problems in which QMC simulations suffer from the sign problem, such as quantum electrodynamics, quantum chromodynamics at finite density, and fermionic systems. As we will see in Chapt. 4, these systems can be attacked by a Tensor Network approach. In particular, therein Tree Tensor Networks are successfully applied to Lattice Gauge Theories at finite density in two and even three dimensions. Studying these systems with alternative Tensor Network methods, such as PEPS, becomes, even in two dimensions, practically unfeasible since the local dimension $d \geq 35$ poses a non-trivial challenge; and for three dimensions, there are no known alternative Tensor Network ansätze. Thus, Tree Tensor Networks offer unique flexibility since they can be, in theory, easily adapted to simulate systems of any dimension. This flexibility makes Tree Tensor Networks an interesting candidate for the analysis of general high-dimensional quantum many-body systems. Additionally, they come with a significant computational advantage, as discussed in Sec. 2.2.2, which enables computations with higher bond-dimensions, thus to some extend compensating the lower entanglement scaling of its structure compared to *looped* Tensor Networks, such as PEPS or MERA [69]. Such increased bond-dimension further translates into a crucial benefit when exploiting symmetries since more

symmetry sectors can be explored on a global level while the incorporation of symmetries in general looped Tensor Networks might require an increased bond-dimensions [24, 107].

However, while a Tree Tensor Network performs extremely well for smaller system sizes of such systems, it lacks the scalability into higher dimensions, due to its lack of efficiently encoding the area law in its structure. This clearly denotes the main bottleneck of Tree Tensor Networks when further numerically investigating large-scale quantum many-body systems towards, for instance, promoting the understanding of our universe by studying non-Abelian high-dimensional Lattice Gauge Theories. However, as we will see in Chap. 5, the Tree Tensor Network can be *augmented* in a way to efficiently encode the area law by its structure offering better scalability then an ordinary TTN. Thus, this *augmented Tree Tensor Network* (aTTN) can help to overcome the current limitations of Tensor Networks in high dimensional systems opening the pathway to novel physics beyond the reach of every other numerical method.

In this respect, this chapter is dedicated to an in-depth discussion on the basics of Tree Tensor Networks. Starting in Sec. 3.1, the fundamental ideas of Tree Tensor Network states are presented. After introducing the Tree Tensor Network in a formal definition, we will see that these class of Tensor Network states naturally have a loop-less, hierarchical structure which greatly benefits the underlying algorithms for calculating observables and finding ground states of many-body Hamiltonians. In Sec. 3.2, this idea of Tree Tensor networks is generalised towards the study of higher-dimensional systems. Therein, we will see that the concrete structure of a Tree Tensor Network is crucial for simulating systems in 2D and beyond and how a system can be mapped ideally to suit the topology of a Tree Tensor Network. This mapping is further underlined with a practical step-by-step example for a 4×4-system with nearest-neighbour interactions. Further, we will see that this network itself (without the augmentation introduced in Chap. 5) does not satisfy the area law in higher dimensions, thus eventually fails to faithfully represent physical states for growing system sizes. The last section of this chapter, Sec. 3.2, is dedicated to an in-depth technical description for the numerical implementation of a Tree Tensor Network, featuring discussions on the underlying data structure, important modules and practical note on the High-Performance-Computing aspect of the implementation.

3.1 Tree Tensor Network states

Definition 3.1.1: *(Tree Tensor Network).*

A Tree Tensor Network, abbreviated by TTN, is a Tensor Network following Def. 2.2.1 in which any two tensors are connected by exactly one path.

Definition 3.1.2: *(Path).*

Let $\{\mathcal{T}\}$ be a set of distinct tensors $(\forall i \neq j : \mathcal{T}^{[i]} \neq \mathcal{T}^{[j]})$ and let $\{\nu\}$ be a set of links such that each link ν_i connects the tensors $\mathcal{T}^{[i]}$ and $\mathcal{T}^{[i+1]}$ for $1 \leq i < k$. Then, a *path* is defined as the ordered, alternating list $(\mathcal{T}^{[1]}, \nu_1, \mathcal{T}^{[2]}, \nu_2, ..., \nu_{k-1}, \mathcal{T}^{[k]})$.
Further, we define:

- The *length* of the path as the number of links in the set $\dim\{\nu\} = k - 1$.

- $\mathcal{T}^{[1]}$ and $\mathcal{T}^{[k]}$ as the endpoints of the path.

- A cycle as a path with an additional link connecting its endpoints.

Theorem 3.1.1: *(Tree Tensor Network properties).*

Let Ψ be a Tensor Network, then the following statements are equivalent:

- Ψ is a Tree Tensor Network.

- Ψ contains no cycles, i.e. Ψ is acyclic (loop-free), and the addition of any internal link results in the formation of cycles

- Ψ is connected and the removal of any internal link disconnects G.

Lemma 4. *Every Tensor Network without cycles is a Tree Tensor Network.*

3.1.1 Hirarchical Structure of Tree Tensor Networks

As illustrated in Fig. 3.1, each TTN can be arranged in a hierarchical structure with different layers Λ_ℓ when defining a *root link* which will be the topmost-link in the network. In this structure, the first layer Λ_1 consists of two tensors connected via the root link while all other neighbours of the two tensors from the layer Λ_2. Thus, further filling all other layers accordingly leaves us with a well defined hierarchical structure. In this structure, each tensor in layer $\Lambda_{\ell>1}$ has exactly one neighbor in layer $\Lambda_{\ell-1}$ which is defined as its *parent tensor*, and may have several neighbors in layer $\Lambda_{\ell+1}$ which are defined as its *child tensors*. For $\ell = 1$, we define the two tensors to be each other's parents (which might go against traditional biological intuition however ensures that every tensor within the network has exactly one parent tensor). Within this structure, the k-th tensor within the ℓ-th layer Λ_ℓ can be enumerated by $\mathcal{T}^{[\ell,k]}$. Further, the *depth* of a hierarchical TTN, is defined as its number of layers $\mathcal{L}$, i.e. the longest
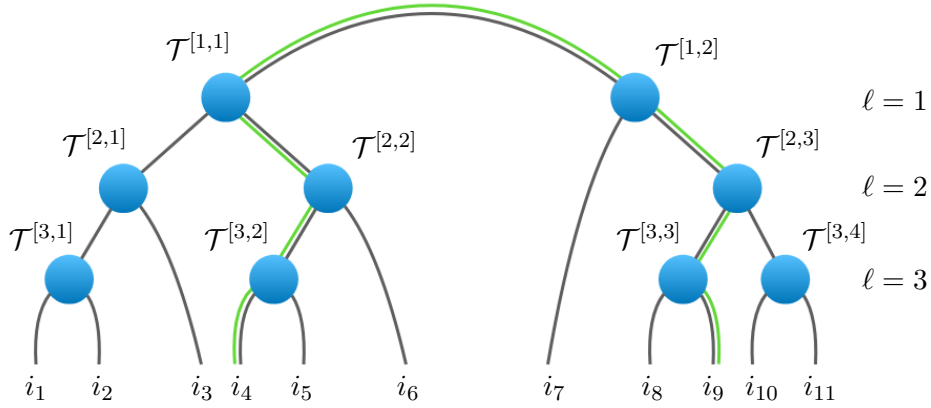
FIGURE 3.1: Structure of a general binary Tree Tensor Network (TTN). The green line indicates the only possible path from Tensor $\mathcal{T}^{[3,2]}$ to $\mathcal{T}^{[3,3]}$, or more concretely from site i_4 to site i_9. In a binary TTN, each tensor has three links while in a general TTN, the tensors may have different numbers of links higher than three.

path between a tensor and the root link. For sake of compactness, we always choose the root link to be an internal link which minimises the number of layers of the hierarchically-ordered tree.

   Given that the computational complexity of the contraction within a network scales with the product of the bond-dimensions of all links attached to the tensors, it is reasonable to minimise the number of links per tensor in an attempt to build an efficiently contractable Tensor Network. Therefore, a particular class of interest are the *binary TTNs* where each tensor can have three neighbours at most, i.e two children in the hierarchical structure. Further, all tensors within this network are of the order-3. Thus, when one tensor does not have two child tensors, it has a physical link attached. Based on the definitions above, we can further categorise a binary TTN into different subclasses:
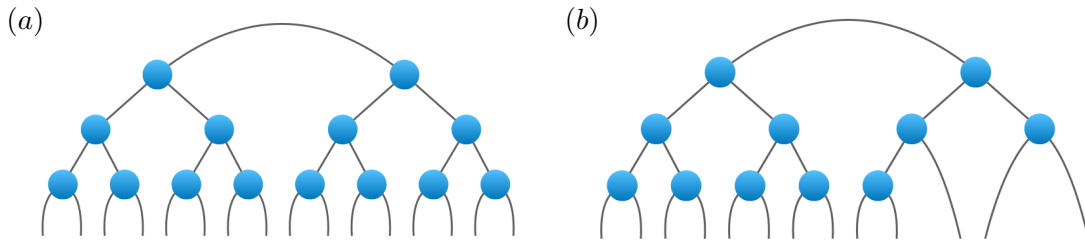

FIGURE 3.2: Structure of a perfect TTN (a) where each of the $\mathcal{L}$ layers is completely filled with tensors resulting in $N = 2^{\mathcal{L}}$ sites and of a complete TTN (b) where the last layer may not be filled entirely. The perfect TTN is a subclass of the complete TTN.

Perfect TTN — In a perfect binary TTN, each layer ℓ is completely filled with the maximum number of tensors (2^{ℓ}). Thus, each tensor within the layers $\Lambda_{l<\mathcal{L}}$ has two child tensors while each tensor $\mathcal{T}^{[\mathcal{L},k]}$ within the last layer $\Lambda_{\mathcal{L}}$ has one internal link upwards connecting it to its parent tensor and two physical links downwards addressing the local Hilbert spaces of the quantum many-body system (see Fig. 3.2a). Consequently, this TTN is only able to capture systems with $N = 2^{\mathcal{L}}$ sites.

Complete TTN —   Analogously to the perfect TTN, each layer within a complete TTN is completely filled, except for the last layer $\Lambda_{\mathcal{L}}$ which is filled from the left onward, i.e. all tensors within the last layer are as far left as possible (see Fig. 3.2b). Evidently, the *perfect TTN* is a subclass of a *complete TTN*. In contrast to the perfect TTN, however, this TTN is able to describe systems of arbitrary size.

Matrix Product States —   One particular subclass of binary TTNs is the MPS. In the framework of hierarchical trees, the MPS equals a binary TTN where each layer $\Lambda_{l>2}$ consists of one tensor only.
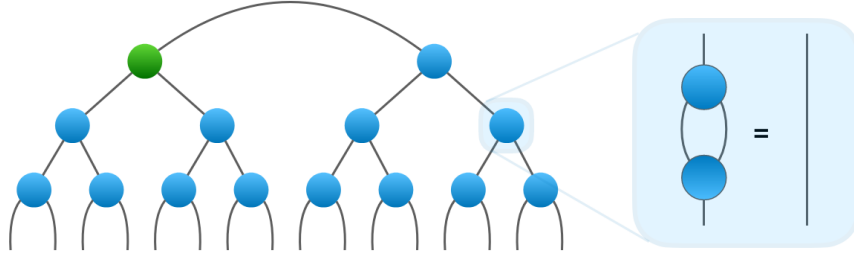
3.1.2   Isometry in a Tree Tensor Network

FIGURE 3.3: TTN isometrised towards the tensor $\mathcal{T}^{[1,1]}$. All other tensors (blue) are isometries as introduced in Def. 2.1.2 and obey the isometry condition (right). The tensor $\mathcal{T}^{[1,1]}$ (green) can be an arbitrary tensor.

One crucial property of TTNs is that they can be brought in a certain isometric form by exploiting the gauge freedom of Tensor Networks. In this isometric form, or *unitary gauge*, the TTN mainly consists of isometries following their introduction in Def. 2.1.2. As illustrated in Fig. 3.3, only one tensor within the network does not obey the isometry condition in the unitary gauge. Thus, this gauge is always defined with respect to a certain tensor $\mathcal{T}^{[l,k]}$ of the network. We call such a TTN *isometrised towards* $\mathcal{T}^{[l,k]}$.

The fact that most of the tensors are isometries and thus vanish when contracted with their complex conjugate over the proper links can be of great benefits in many numerical scenarios. The first benefit on hands is the computation of the norm $\langle\psi|\psi\rangle$ of the TTN state: In this computation, the entire TTN has to be contracted with its complex conjugate. Having the TTN isometrised to any tensor $\mathcal{T}^{[l,k]}$ reduces this computation to the contraction $\sum_{i_1,\dots,i_K}(\mathcal{T}^{[l,k]})_{i_1,\dots,i_K}(\mathcal{T}^{[l,k]})^{\dagger}_{i_1,\dots,i_K}$ of $\mathcal{T}^{[l,k]}$ with its own complex conjugate over all of its K links. In the next section (Sec. 3.1.3), for instance, we will see that, in the same way, the calculation of observables can be executed far more efficient when taking advantage of a properly isometrised TTN.

Having an isometrised TTN towards a given tensor $\mathcal{T}^{[l,k]}$, we can always change this gauge towards another tensor $\mathcal{T}^{[l',k']}$ in the network, or in other words *reisometrise* the TTN towards $\mathcal{T}^{[l',k']}$. For such a reisometrisation, we have to define the path $\Gamma_{l,k}^{l',k'}$ from the tensor $\mathcal{T}^{[l,k]}$ towards which the network is currently isometrised towards the target tensor $\mathcal{T}^{[l',k']}$. As first step in the procedure, the tensor $\mathcal{T}^{[l,k]}$ becomes decomposed via QR-decomposition into an isometry tensor $\mathcal{Q}$ and an order-2 tensor $\mathcal{R}$ along the direction defined by the path.
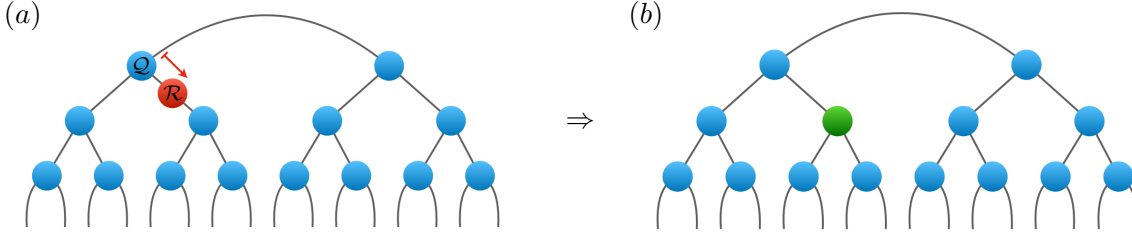
FIGURE 3.4: Algorithm to reisometrise a TTN from position tensor $\mathcal{T}^{[1,1]}$ towards tensor $\mathcal{T}^{[2,2]}$: (a) The tensor $\mathcal{T}^{[1,1]}$ was decomposed via QR-decomposition into an isometry $\mathcal{Q}$ and an order-2 tensor $\mathcal{R}$ (red). $\mathcal{Q}$ denotes the new tensor $\tilde{\mathcal{T}}^{[1,1]}$. (b) The tensor $\mathcal{R}$ is contracted with the tensor $\mathcal{T}^{[2,2]}$ resulting in a new tensor $\tilde{\mathcal{T}}^{[2,2]}$ which does not fulfil the isometry condition any more (marked in green). Consequently, the TTN is isometrised towards $\mathcal{T}^{[2,2]}$.

Subsequently, the tensor $\mathcal{T}^{[l,k]} \leftarrow \mathcal{Q}$ is replaced by the created isometry while the tensor $\mathcal{R}$ is contracted with the neighbouring tensor along the path $\Gamma_{l,k}^{l',k'}$. This first step already reisometrises the TTN towards the neighbouring tensor of $\mathcal{T}^{[l,k]}$. Consequently, performing this procedure iteratively along the entire path $\Gamma_{l,k}^{l',k'}$, reisometrises the TTN towards the desired tensor $\mathcal{T}^{[l',k']}$. As an example, Fig. 3.4 illustrates the reisometrisation from the tensor $\mathcal{T}^{[1,1]}$ towards its neighbouring tensor $\mathcal{T}^{[2,2]}$.

Since we create an isometry at each step of the reisometrisation, this procedure can be extended for establishing the unitary gauge within an arbitrary TTN by iteratively performing a QR-decomposition through the entire network. In practice, however, an isometrisation is typically created during the initialisation of a TTN as illustrated in Sec. 3.1.4. Once created, the executed TTN algorithms are usually designed to manage, keep track and take advantage of the isometrisation.

3.1.3 Calculation of Observables

In what follows, the computation of the expectation values of different observables O are presented for the TTN. Of particular interest in many simulations of quantum many-body systems are the calculations for *(i)* local observables, acting on one site s only, *(ii)* correlation functions, i.e. two-body operators, and *(iii)* general String observables.

Local observables

Let $O^{[s]}$ be a local observable acting on the site i_s of the TTN state ψ. In order to efficiently compute the expectation value $\langle O^{[s]} \rangle_\psi = \langle \psi | O^{[s]} | \psi \rangle$ the TTN can be isometrised towards the tensor $\mathcal{T}^{[\mathcal{L}, \lfloor (s+1)/2 \rfloor]}$ which acts on the physical site i_s ($\lfloor \cdot \rfloor$ denotes the floor function). Using this isometry in the TTN, the computation of $\langle O^{[s]} \rangle_\psi$ can be executed by contracting three tensors only with a complexity of $\mathcal{O}(d^3 m) + \mathcal{O}(d^2 m)$ where d is the local physical dimension of the system (see Fig. 3.5).

Correlators

Let $O^{[s_1,s_2]}$ be a correlator acting on the sites i_s and $i_{s'}$ of the TTN state ψ. In order to compute the expectation value $\langle O^{[s,s']} \rangle_\psi = \langle \psi | O^{[s,s']} | \psi \rangle$, the TTN can be isometrised towards
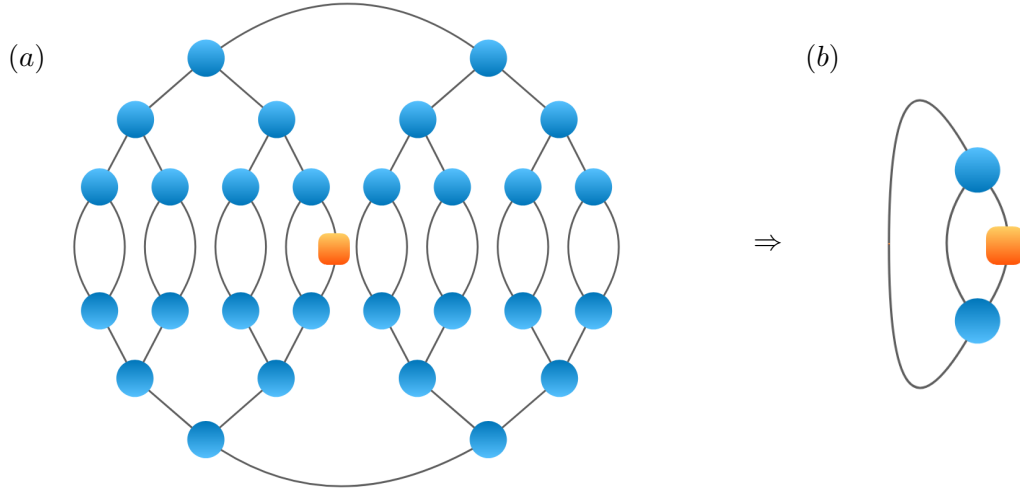
FIGURE 3.5: Contraction scheme for the calculation of the expectation value of a local observable (orange) acting on a single site of the TTN (a). Isometrising the network towards the tensor directly attached to the operator enables to reduce the computation to a contraction of three tensors since all other tensors are isometries and become an identity when contracted with their corresponding complex conjugate (b).
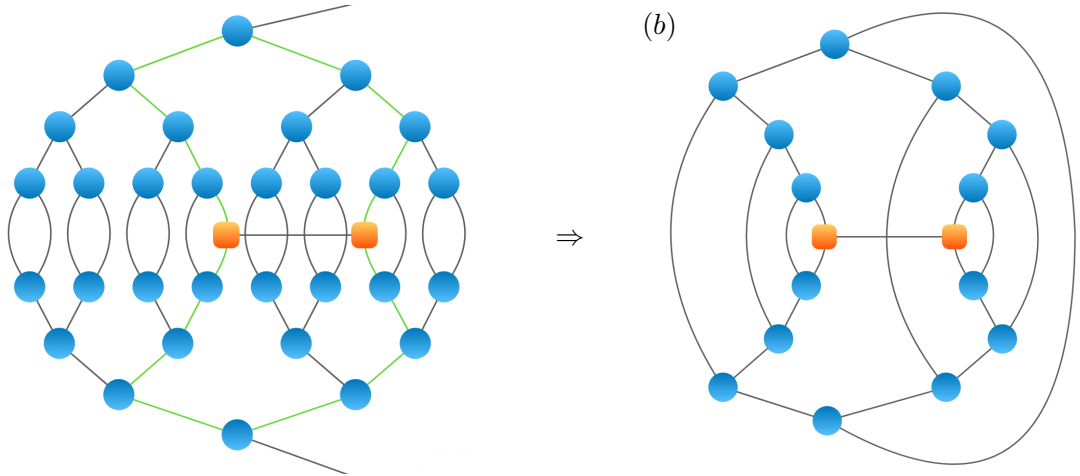

FIGURE 3.6: Contraction scheme for the calculation of the expectation value of a correlator (orange) acting on two sites of the TTN (a). The open links indicate that the total TTN may be larger than the illustrated subtree. The green line indicates the path combining the two tensors directly attached to the operator. Isometrising the network towards the anchor node of this path enables to reduce the computation to the contraction of all tensors along the path since all other tensors vanish in the contraction due to the isometry condition (b).

the anchor of the path from the physical site s to s'. Thereby, only the tensors within this path are to be contracted since every other tensor cancels out with its complex conjugate in the global contraction (see Fig. 3.6). This contraction can be performed with a worst-case complexity of $\mathcal{O}(m^4 \kappa \log N) + \mathcal{O}(m^4 \kappa \log (N-1)) + \mathcal{O}(m^3 \kappa)$ where κ denotes the internal dimension of the Tensor Product Operator (TPO) representing the correlator $O^{[s_1, s_2]}$ (for TPO see Sec. 2.2.3) and $\log N$ equals the number of layers within the TTN.

String observables

Let $O^{[s_1,s_2,...,s_\Omega]}$ be a string operator acting on the sites $s_1, s_2, ..., s_\Omega$ of the TTN state ψ. In order to compute the expectation value $\langle O^{[s_1,s_2,...,s_\Omega]} \rangle_\psi = \langle \psi | O^{[s_1,s_2,...,s_\Omega]} | \psi \rangle$, the TTN can be isometrised towards the anchor of the longest path of each combination of two the physical site $s \in \{s_1, s_2, ..., s_\Omega\}$ to $s' \in \{s_1, s_2, ..., s_\Omega\}$. In this scenario, only the tensors within all these different paths are to be contracted since every other tensor cancels out with its complex conjugate in the global contraction. In the most general case, the observable forms an MPO [61, 147, 148]. In this, numerically worst case, non of the tensors can be neglected in the contraction and the complexity of the underlying contraction becomes $\mathcal{O}(m^4\kappa^2(\kappa + 2)(N - 2)) + \mathcal{O}(m^2\kappa)$ where κ is the dimension of the TPO representing $O^{[s_1,s_2,...,s_\Omega]}$ and $N - 2$ equals the total number of tensors within the TTN.

3.1.4 Ground state computation

This section describes the optimisation algorithm for a TTN to compute the ground state of an interacting quantum many-body Hamiltonian. As introduced in Sec. 2.2.5, such a ground state search can be done efficiently with Tensor Networks by iteratively optimising a sub-set of its variational parameters minimising the energy $E = \langle \psi | \mathcal{H} | \psi \rangle$. In particular, this optimisation exploits the fact that all single tensors within the network can be treated as independent variables. Consequently, the key idea behind the optimisation of a Tensor Network state $|\psi\rangle$ is to optimise one, or only a limited subset, of the tensors at a time, thus reformulating the global, exponentially large, optimisation problem towards a set of local minimisations in a drastically smaller subspace which is controlled in size by the chosen bond-dimension of the network.

In the following, this optimisation procedure is illustrated in more detail starting from the global optimisation and how it can be reformulated as a local minimisation problem for either single tensors or a subset of locally connected tensors. Afterwards, the local optimisation problem is presented and how we can solve this minimisation in practical applications before finalising this section with a discussion on the importance of the initialisation of the TTN for an efficient simulation.

Solving the global optimisation problem

Let $\mathcal{H} \in \mathscr{H}$ be the Hamiltonian of a quantum many-body system and Ψ a TTN to represent a wave-function of the same Hilbert space. Then, the ground state of $\mathcal{H}$ is the solution to the minimisation problem

$$\min_\Psi \{E(\Psi)\} = \min_\Psi \{\langle \Psi | \mathcal{H} | \Psi \rangle\} \ .$$

To find the ground state of $\mathcal{H}$, the variational parameters of the TTN are optimised targeting the minimisation of the energy $E(\Psi)$ of the wave function it represents. Here, the TTN allows performing this optimisation problem by treating each tensor within the network as an independent set of variational parameters. Thus, the global optimization can be executed by *sweeping* through the network (see Fig. 3.7), where we subsequentially perform a local optimisation for each tensor. In order to perform such a local optimisation, the Tensor Network can be differentiated with respect to any single arbitrary tensor $\mathcal{T}^{[\kappa]}$ within the network:
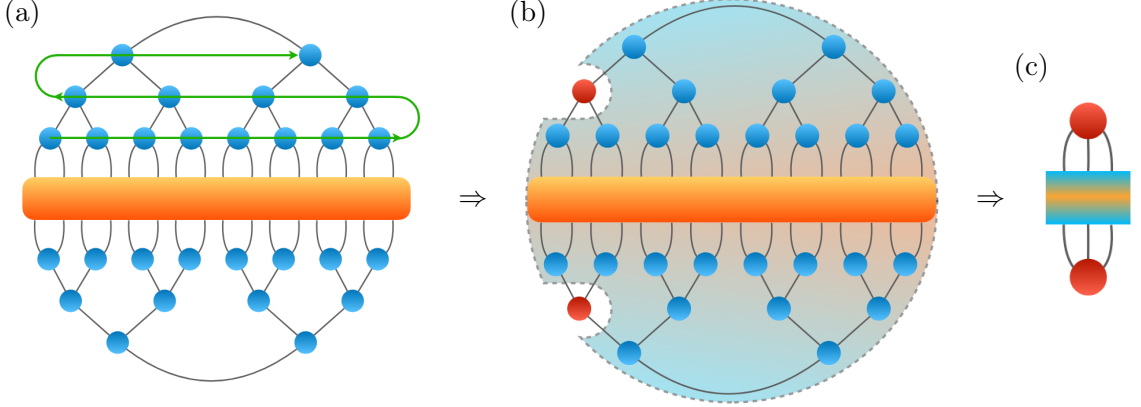
FIGURE 3.7: Optimisation algorithm for a Tree Tensor Network to find the ground-state of a quantum many-body system. (a) The current energy can be calculated by contracting the Hamiltonian (orange) with the TTN and its complex conjugate. The minimisation of the energy is executed by iteratively optimising the variational parameters of each tensor. Therefore, a sweeping sequence (green) is defined in which we optimise the local tensors. (b) To optimise a targeted tensor (red), its environment (shaded area with border in dashed lines) is contracted describing the *effective Hamiltonian* of the local problem. (c) The local optimisation problem has been boiled down to the tensor (red), its complex conjungate, and its *effective Hamiltonian* (orange/blue).

$$\frac{\partial \Psi}{\partial \mathcal{T}^{[\kappa]}} = \frac{\partial}{\partial \mathcal{T}^{[\kappa]}} \sum_{i_1, i_2, \ldots, i_N} \left(\prod_{k=1}^{K} \mathcal{T}^{[k]}_{\{i\}_k, \{\chi\}_k} \right)$$

$$= \sum_{i_1, i_2, \ldots, i_N} \left(\prod_{k=1}^{K} \frac{\partial}{\partial \mathcal{T}^{[k]}} \mathcal{T}^{[k]}_{\{i\}_k, \{\chi\}_k} \right)$$

$$= \sum_{i_1, i_2, \ldots, i_N} \left(\prod_{k \neq \kappa} \mathcal{T}^{[k]}_{\{i\}_k, \{\chi\}_k} \right) \equiv \Psi^{[\kappa]}_{env}$$

Thus, the derivation of any function linear in $\mathcal{T}^{[\kappa]}$ can be calculated numerically where we call the resulting $\Psi^{[\kappa]}_{env}$ the environment of the tensor $\mathcal{T}^{[\kappa]}$. With this definition, the energy expectation value can be reformulated with respect to a single tensor

$$E = \langle \Psi | \mathcal{H} | \Psi \rangle = \langle \mathcal{T}^{[\kappa]} | \underbrace{\Psi^{[\kappa]}_{env} \mathcal{H} \Psi^{[\kappa]}_{env}}_{\mathcal{H}^{[\kappa]}_{eff}}^{\dagger} | \mathcal{T}^{[\kappa]} \rangle$$

where we further define the *effective Hamiltonian* $\mathcal{H}^{[\kappa]}_{eff}$ for the tensor $\mathcal{T}^{[\kappa]}$. Consequently, the global Lagrangian functional $\mathcal{L}(\Psi, \Psi^{\dagger}) = \langle \Psi | \mathcal{H} | \Psi \rangle - \epsilon \langle \Psi | \Psi \rangle$ of the optimisation problem can be reformulated into a local Lagrangian functional $\mathcal{L}(\mathcal{T}^{[\kappa]}, \mathcal{T}^{[\kappa]\dagger}) = \langle \mathcal{T}^{[\kappa]} | \mathcal{H}^{[\kappa]}_{eff} - \epsilon \mathcal{N} | \mathcal{T}^{[\kappa]} \rangle$ with respect to $\mathcal{T}^{[\kappa]}$ which is reshaped as a vector $| \mathcal{T}^{[\kappa]} \rangle$. The *effective Hamiltonian* $\mathcal{H}^{[\kappa]}_{eff}$ hereby includes the physical Hamiltonian $\mathcal{H}$ and all tensors in the network surrounding $\mathcal{T}^{[\kappa]}$ while $\mathcal{N}$ denotes a positive operator for the effective square norm constraint. Thus, we can lower the global energy by optimising only the variational parameters within $\mathcal{T}^{[\kappa]}$. In practice, this reformulation is done by contracting $\mathcal{H}$ with the whole environemnt $\Psi^{[\kappa]}_{env}$, i.e. the complete

TTN Ψ and its complex conjugate except for the tensors $\mathcal{T}^{[\kappa]}$ and $\mathcal{T}^{[\kappa]\dagger}$. Fig. 3.7 illustrates such a contraction of the network towards the local tensor $\mathcal{T}^{[\kappa]}$ with $\kappa = (2, 1)$. Once, the Hamiltonian and the environment of $\mathcal{T}^{[\kappa]}$ is contracted into the effective Hamiltonian, the minimisation problem for $\mathcal{T}^{[\kappa]}$ lives in a space restricted by the chosen maximum bond-dimension m of the network, i.e. $\dim\{\mathcal{H}_{eff}^{[\kappa]}\} \leq m^3$. Typically this local optimisation problem can be solved directly by an eigenvalue decomposition $\mathcal{H}_{eff}^{[\kappa]} = Ev^{\dagger}\mathcal{H}_{eff}^{[\kappa]}v$ where we update the tensor $\mathcal{T}^{[\kappa]} \to v_0$ with the eigenvector of $\mathcal{H}_{eff}^{[\kappa]}$ corresponding to the lowest eigenvalue E_0. Thus, optimising the tensor $\mathcal{T}^{[\kappa]}$ within the TTN consists of three different steps:

1. Contract the Hamiltonian $\mathcal{H}$ with the environment $\Psi_{env}^{[\kappa]}$ to the effective Hamiltonian $\mathcal{H}_{eff}^{[\kappa]} = \Psi_{env}^{[\kappa]}\mathcal{H}\Psi_{env}^{[\kappa]}$

2. Compute the eigenvector v_0 with minimum eigenvalue for $\mathcal{H}_{eff}^{[\kappa]}$

3. Update the tensor $\mathcal{T}^{[\kappa]} \to v_0$

Moving on to solve the local optimisation problem iteratively for all tensors within the network gives a high likelihood to converge into the global minimisation problem for the Lagrangian functional $\mathcal{L}(\Psi, \Psi^{\dagger})$. In this way, the complexity of the global optimisation can be successfully reduced towards a significantly smaller space bound by the bond-dimension m of Ψ.

The exact *sweeping* sequence of the optimisation, in general, can be random which turns out to be beneficial in complex systems, such as the study of Lattice Gauge systems in higher dimensions (presented in Chapt. 4). Another successful sequence is going iteratively from the bottom left to the top right tensor, as indicated in Fig. 3.7a (green). In this way, the information of the system can be propagated upwards through the TTN starting from the tensors directly attached to the Hamiltonian while still keeping a relatively short path to complete one sweep. Iterating this sweeping procedure through the complete network leads to the convergence for the global TTN in terms of energy and selected observables. Equivalently this strategy can be extended to optimising not only one tensor at a time but a small subset of tensors.

Local optimisation

As discussed above, the global minimisation problem can be solved by iteratively solving a local eigenvalue problem for each local tensor $\mathcal{T}^{[\kappa]}$. For this optimisation, i.e. eigenvalue problem, we can exploit well-established algorithms from linear algebra. In particular, the Arnoldi algorithm [173, 174] implemented in the ARPACK library [175], a collection of Fortran77 subroutines designed to solve large scale eigenvalue problems, offers an efficient numerical solution. The Arnoldi algorithm solves the local eigenvalue problem for each tensor $\mathcal{T}^{[\kappa]}$ by iteratively diagonalising its corresponding effective Hamiltonian $H_{eff}^{[\kappa]}$. In the practical application, the algorithm only needs to know the action of $H_{eff}^{[\kappa]}$ on $\mathcal{T}^{[\kappa]}$ and thereby provides the lowest eigenpairs of $H_{eff}^{[\kappa]}$ within a predetermined accuracy ϵ (see Sec. 3.2.3 for a hands-on example). The complete numerical complexity for this operation is determined by the dimension of all links of $\mathcal{T}^{[\kappa]}$, upper-bounded by $\mathcal{O}(m^4)$. In practical applications, it might be wise to start with a low optimization precision ϵ of the Arnoldi algorithm and increasing it after each sweep such that the solution of the local eigenvalue problems becomes more and more accurate as the TTN comes closer to global convergence.

Interesting strategies are further to extend this single tensor update towards a two-site update, in which two neighbouring tensors $\mathcal{T}^{[\kappa]}$ and $\mathcal{T}^{[\kappa']}$ are optimised at the same time. This can be done analogously to the description above by contracting the two tensors and calculating the effective Hamiltonian, now with the environment $\Psi_{eff}^{[\kappa,\kappa']}$ with respect to the complete space of $\mathcal{T}^{[\kappa]}$ and $\mathcal{T}^{[\kappa']}$. However, since this would lead to a local minimisation problem for a contracted two-tensor space, it is more expensive numerically with a complexity of $\mathcal{O}(m^5)$. After the optimisation of this contracted tensor, it is decomposed again into its original form, i.e. two interconnected tensors, keeping the TTN structure alive. This strategy may increase the numerical complexity, but it allows to adapt the bond-link dynamically (including its dimension and for symmetric TTNs its symmetry sectors) and thereby may provide a better or faster convergence to the global minimum.

As a good compromise, the subspace-expansion [24] can be exploited for the local single tensor optimisations. This technique approximates a two-site update by *(i)* first increasing the bond-dimension of the link connecting the two tensors, *(ii)* then iteratively optimising the two local tensors separately before *(iii)* truncating the bond-link back to original dimension. In this way, the local optimisation exploits the best of both worlds: It keeps the favourable algorithmic complexity of $\mathcal{O}\left(m^4\right)$ while effectively performing a complete two-site optimisation and thus being able to adapt the bond-link dynamically. For more details on this technique, Ref. [24] offers an in-depth illustration and a comparison between the three strategies.
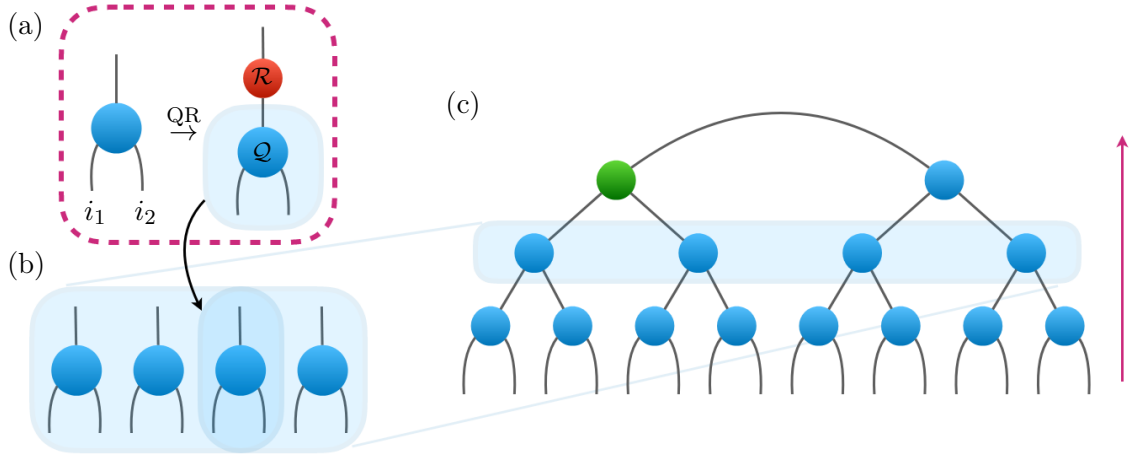
Initialisation

FIGURE 3.8: Construction of a Layer for the initialisation of a Tree Tensor Network. (a) Given two sites (i_1, i_2) of local state spaces as input, an order-3 tensor is initialised with random elements combining the local spaces (bottom left). Performing a QR-decomposition creates an order-3 isometry $\mathcal{Q}$ which is used as the initialised tensor within a TTN layer (b). For the next TTN layer, the procedure is repeated, now using the links from the just created tensors of the lower layer. (c) The TTN is created layer after layer from the bottom to the top.

For the ground-state search of a Hamiltonian $\mathcal{H}$, it is beneficial to start from a random guess $|\psi_{init}\rangle$ in order to maximise the probability of an overlap with the target state $|\langle\psi_{\text{init}}|\psi_{\text{target}}\rangle|^2 > 0$. Such a randomly initialised tree can be created with the following bottom-top strategy illustrated in Fig. 3.8.

Starting from the physical indices at the bottom of the TTN, the lowest layer $\Lambda_{l=\mathcal{L}}$ is built up by merging two neighbouring local sites into one internal link introducing a randomly initialized tensor. Subsequently, the next layer $\Lambda_{\mathcal{L}-1}$ can be constructed similarly by taking the resulting internal links from each tensor within the just constructed layer $\Lambda_{\mathcal{L}}$, and again merging two neighbouring internal links by introducing randomly initialized tensors. Following this procedure until $l = 1$, the complete TTN would be initialised with an exponentially increasing bond dimension for each layer $m_{\mathcal{L}-l} = d^l$. Thus, in case we reach the user-defined maximum bond dimension m of the network when constructing each layer from the bottom to the top, we truncate the bond dimension $m_{\mathcal{L}-l} = \min\{d^l, m\}$ accordingly.

Importantly, in the case of present symmetries, this truncation of the maximum allowed bond-dimension has to be done by truncating the single symmetry sectors within the internal link. Here, we can introduce further randomness in the state by truncating these occupied coupling sectors randomly down until reaching the maximum set bond-dimension for the link in question. Thereby, not only the tensors themselves will be randomly initialised but further the distribution of the coupling symmetry sectors within the complete network. However, to ensure optimal convergence during the optimisation, it is critical to occupy the correct symmetry sectors within the internal links corresponding to the symmetry sectors of the ground state. However, as mentioned before, the local bond-dimension can be dynamically increased during the optimisation algorithm in order to adapt the symmetry sectors within the TTN. Thus, before a symmetrically invariant TTN converges to the ground state, it first converges in the occupied symmetry sectors on each link within the network.
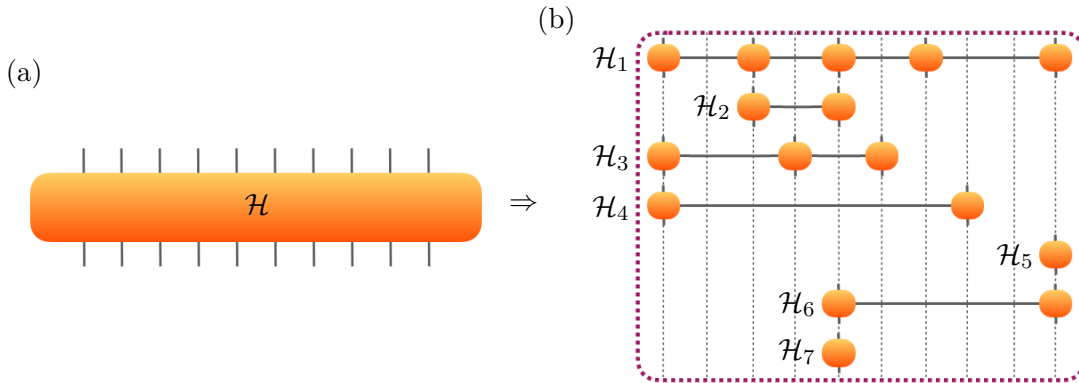
3.1.5 Efficient handling of the Hamiltonian

Figure 3.9: Quantum many-body Hamilton $\mathcal{H} = \sum_p \mathcal{H}_p$ in a Tensor Network framework. The complete Hamiltonian (a) is represented as a set of *Tensor Product Operators* (TPOs), one for each interaction $\mathcal{H}_p$ (b).

As described in Sec. 2.2.2, the complete many-body Hamiltonian $\mathcal{H} = \sum_p \mathcal{H}_p$ typically consists of several separate interactions $\mathcal{H}_p$ which can each be represented as a *Tensor Product Operator* (TPO). Fig. 3.9 illustrates an example of such a Tensor Network representation decomposing $\mathcal{H}$ into a set of TPOs. In the TTN analysis, these TPOs act on different physical links of the network while the complete Hamiltonian acts on all open links, i.e. physical links, of the TTN.
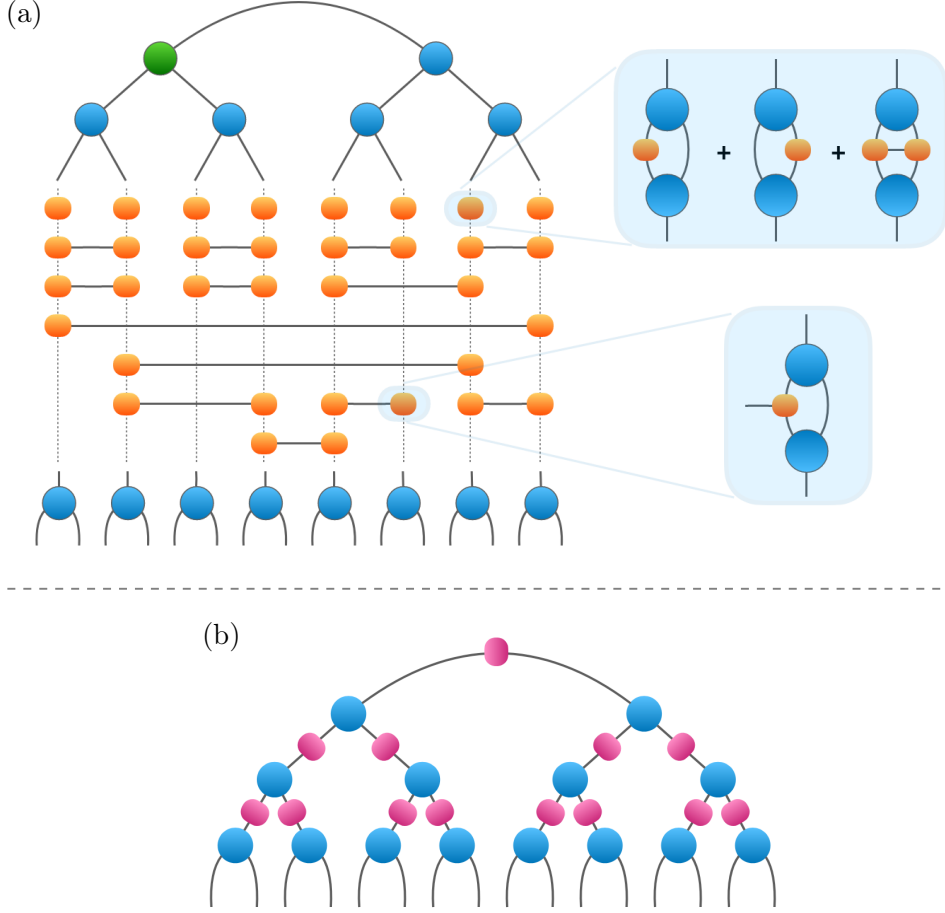
FIGURE 3.10:   Link-operators of a TTN for efficient computation of effective Hamiltonians $\mathcal{H}_{eff}$.
(a) Each tensor of the TPO representing the physical interactions $\mathcal{H}_p$ have been contracted with the
corresponding tensor in the lowest layer of the TTN and its complex conjugate (see inlet). (b) For
each link, we store the corresponding coarse-grained operator-tensors as link-operators (purple) to
reduce the required computations during the optimisation algorithm.

In the previous Sec. 3.1.4, it was illustrated that during the optimisation procedure, we
repeatedly need to compute the *effective Hamiltonian* $\mathcal{H}_{eff}$ for each tensor $\mathcal{T}^{[l,k]}$ within the
TTN. This computation can be performed by contracting all TPOs separately through the
network as indicated in Fig. 3.10: Each tensor of all the TPOs $\mathcal{H}_p$ are contracted with the
tensor $\mathcal{T}^{[l,k]}$ they act on and its complex conjugate $(\mathcal{T}^{[l,k]})^\dagger$. Thus, we perform a renormal-
isation procedure of the interactions for each layer in the TTN. This procedure maps the
single TPO parts onto the virtual subspace of the internal links of the TTN. In practical
applications, we store these renormalised operators at each internal link χ of the TTN as
link-operators. The explicit storing of the link-operators enables to improve the computation
for various TTN algorithms, such as time-evolution or ground state search, by reducing the
required computations during the algorithms. For further practical information, Sec. 3.2.3
illustrates a hands-on example including the initialisation of the link-operators and the un-
derlying contraction schemes to calculate the *effective Hamiltonian* $\mathcal{H}_{eff}$ for a targeted TTN
tensor $\mathcal{T}^{[l,k]}$.

3.2 Simulating high-dimensional systems with TTNs

Once a complete TTN is working for one-dimensional systems, the application to higher dimensions is in theory straightforward. However, there are still non-trivial practical hurdles to overcome in the implementation. The following section describes in more detail this application of a complete TTN to high-dimensional systems.

3.2.1 Two-dimensional TTN structure and Hamiltonian mapping

The first and obvious difference going from analysing one-dimensional systems to higher dimensions is the fact that we have to consider the additional dimensions when attaching the physical sites to the TTN structure. In other words, the question arises of how to map the high-dimensional lattice to the generic one-dimensional TTN structure. Indeed, while for a 1D system the physical site i is applied to the i-th external link of the TTN, each physical site at position $(x_1, x_2, ..., x_D)$ of a general D-dimensional system requires a well-defined mapping to one of the physical links of the TTN. In fact, this mapping is highly critical for the representation power and accuracy of a TTN approach. Thus, for an efficient mapping, we aim to take advantage of the physical principle that the entanglement typically decays with distance in quantum many-body systems. Therefore, the mapped TTN structure shall be arranged in a way that the high entanglement bipartitions correspond to the lower branches of the TTN. For these lower branches $\nu \in \Lambda_l$ with $l \to \log N$, the bond-dimension m_ν is sufficiently large to capture the area law entanglement - or even the complete state - accurately, especially for reasonably small local dimensions d. Instead, the bond-dimension in the higher layers of the TTN ($l \to 1$) suffers from the required exponentially large bond-dimension.

A straightforward way to implement such a mapping is to map the two-dimensional lattice in a zig-zag pattern towards the physical sites $j \in \{1, \ldots, N\}$ of the TTN. More concretely, each site (x, y) of the two-dimensional lattice with $x, y \in \{1, \ldots, L\}$ becomes mapped to an auxiliary site $j = x + L \cdot (y - 1)$ for the TTN simulation. In this way, the TTN performs a renormalises each "column" in x-direction of the system at its lower layers fist, before
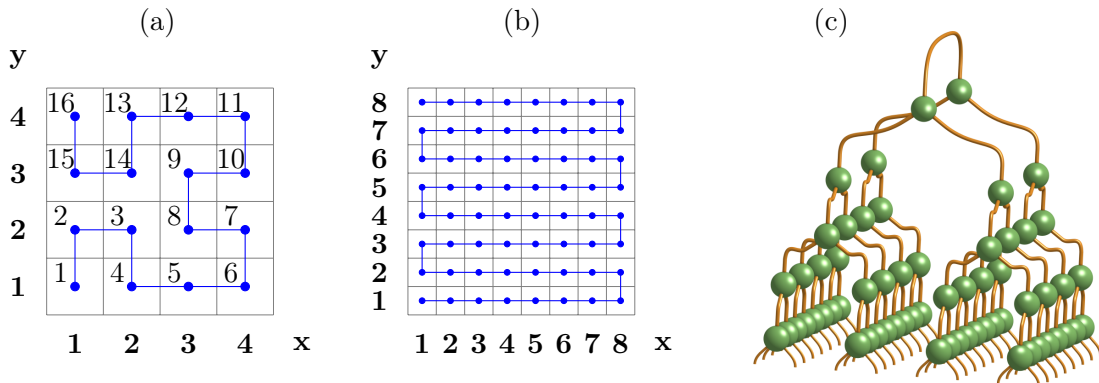

FIGURE 3.11: Possible structures of a TTN deployed for two-dimensional systems. The TTN can be aligned along different paths through the system, such as the Hilbert curvature (a) or a snake-like pattern (b). The chosen path defines the mapping for each local site $(x, y) \to i \in \{1, ..., N\}$ from the two-dimensional system to the sites i of the TTN giving rise to different structures for the two-dimensional TTN simulation. In (c), a binary TTN structure is defined such that its tensors (green) each merge two sites to one bond link (brown). The mapping here groups alternatingly in x- and y-direction from layer to layer starting from the physical sites of the 8×8 lattice.

(a)

(2D)
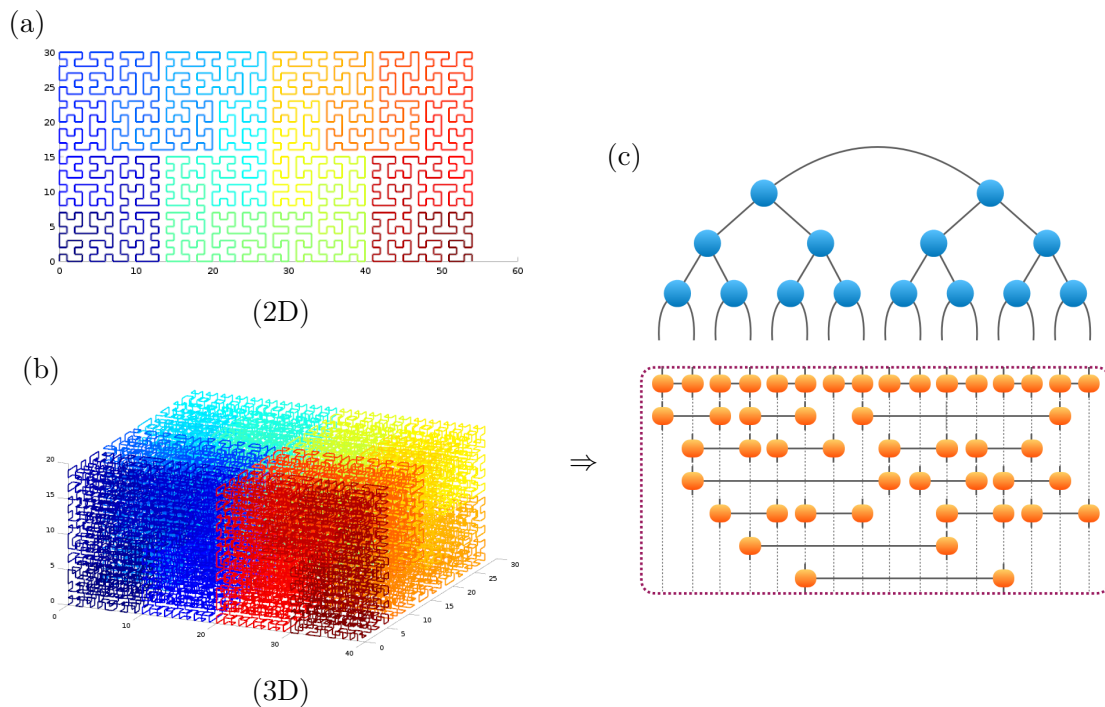
(b)

(3D)

(c)

$\Rightarrow$

FIGURE 3.12: Mapping of a high-dimensional system towards a generic TTN structure. (a)+(b) The generalised Hilbert-curvature defines the mapping for the local physical sites from the high-dimensional system to a one-dimensional path: (a) 2D mapping $(x, y) \rightarrow i \in \{1, ..., N\}$; (b) 3D mapping $(x, y, z) \rightarrow i \in \{1, ..., N\}$. The colour indicates the magnitude of i. (c) Mapping all interaction parts $\mathcal{H}_p$ of the high-dimensional Hamiltonian gives rise to a long-range interacting, one-dimensional system (purple-dashed box) in the generic basis of a TTN. Subfigures (a)+(b) taken from Ref. [176].

the coarse-grained "columns" become merged together in y-direction at the upper layers. Consequently, the resulting TTN topology is not well-suited for capturing correlations in y-directions and the numerical simulations will be biased in x-direction. Thus, this mapping is not ideal for a general TTN approach in two dimensions.

A well-established procedure is to construct the network in a way such that each layer of the TTN alternatingly groups together neighbouring sites in x- or y-direction respectively, as illustrated in Fig. 3.11c. Thereby, each tensor of a layer combines two sites to one coarse-grained virtual site, i.e. internal link, building up the hierarchical tree structure from the bottom to the top. In the case of a $L \times L$ system, this technique effectively coarse-grains the system in local plaquettes and enables to improve the representation power of correlations within these plaquettes in the TTN structure, resulting in a more precise representation of a quantum many-body state of the underlying system in general. However, this procedure by construction is only able to deal with binary-sized dimensions with $L_k = 2^n$.

For this problem, the *generalised Hilbert-curvature* introduced in Ref. [176, 177] offers an efficient compromise: Exploiting curvature allows to construct a precise TTN structure analysing high-dimensional systems while being able to keep the system linear dimensions L_k of arbitrary size. As illustrated in Fig. 3.12, generalised Hilbert-curvature defines the mapping from the high-dimensional system to one-dimensions. Accordingly, each interaction becomes mapped, in general, to an arbitrary one-dimensional long-range interaction. Consequently, the mapped one-dimensional system becomes analysed by the generic TTN structure. In fact, for

isotropic systems with nearest-neighbour interactions, the TTN construction described above allows creating a graph together with the physical interactions which is surjectively homomorphic to the resulting graph when mapping the system using the Hilbert curvature. In other words, for binary sized systems, the TTN arising from the x-/y-coarse-graining procedure is equivalent to the TTN using the Hilbert-curvature apart from a structural degeneracy.

A priori, the exact entanglement entropy within a quantum many-body state is usually unknown before analysing the system which makes it challenging to figure out the ideal Tensor Network structure for a given system. However, we know that the entanglement is governed by the interactions within the systems. Moreover, there is even a connection between the coupling strengths g of a given interaction and the amount of introduced entanglement for the ground state: For increasing g the interaction introduces entanglement accordingly with $g \to 0$ keeping the sited disentangled and $g \to \infty$ introduces completely entangles the sites. However, in both extremes, disentangled and completely entangled, the introduced entanglement entropy $S = 0$ vanishes. Thus, the entanglement entropy maximises for medium coupling strength g. Thus, knowing the interactions of the system can help to find an efficient high-dimensional TTN structure by optimising the TTN layout such that it minimises the average weighted distance between any interacting sites

$$\langle D \rangle = 1/N_p \sum_{p=1}^{N_p} |g_p| D_p$$

where D_p is the path distance of an interaction with corresponding interaction strength g_p and N_p the total number of interactions in the system.

An example illustrating the importance of this mapping can be found by looking at a two-dimensional isotropic model with nearest-neighbour interactions only and periodic boundary conditions. These Hamiltonians of the shape $H = \sum_p \mathcal{H}_p$ with $|g_p| \equiv ||H_p||$ are frequently used to investigate magnetisation, e.g. in the form of the Heisenberg model.

One-dimensional Systems   Before going to the analysis of the mapping in two dimensions, let us have a look at the one-dimensional case: As illustrated in Fig 3.13, each interaction is acting on two neighbouring sites of the TTN with the number of interaction $N_\mathcal{H} = L - 1$ proportional to the system length L for open boundary conditions ($N_\mathcal{H} = L$ for periodic boundary conditions). Thus, there are $\frac{L}{2}$ interactions with distance $D_p = 1$, i.e. both physical sites of the interactions are assigned to the physical links of the same tensor. Further, $\frac{L}{4}$ interactions are acting with a path distance $D_p = 3$ on the TTN, i.e. they are connected over a tensor in the second layer, $\frac{L}{8}$ interactions with $D_p = 5$, etc.. Following this patter up to the highest layer, the average distance equals

$$\langle D_{1D} \rangle = \frac{1}{L} \left(\frac{L}{2} \cdot 1 + \frac{L}{4} \cdot 3 + \frac{L}{8} \cdot 5 + ... \right) = \frac{1}{L} \sum_{k=1}^{\log L} \frac{L}{2^k} (2k-1) = \sum_{k=1}^{\log L} 2^{-k}(2k-1) \quad (3.1)$$

$$= 3 - \frac{1}{L}(2 \log L + 3) \quad (3.2)$$

which converges to $\langle D_{1D} \rangle \to 3$ for $L \to \infty$. Thus, the average path distance of a nearest-neighbour interaction within a TTN representing a one-dimensional system is finite in the thermodynamical limit.
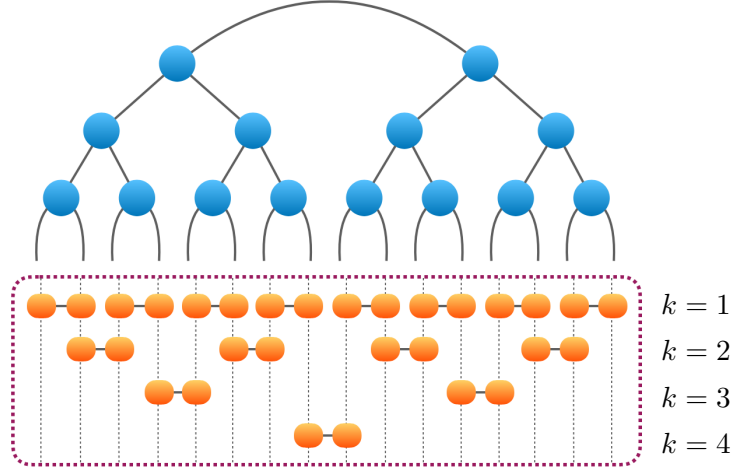
FIGURE 3.13: Nearest-neighbour interactions for a one-dimensional TTN analysis. The interactions are ordered with respect to their distance D_p: All interactions in the k-level for $k \in [1, ..., \log L - 1]$ are connected through a path of length $D_p = 2k - 1$ through the network.

Two-dimensional Systems   For the same analysis in two dimensions, the average path distance $\langle D_{2D} \rangle = \frac{1}{2}(\langle D_x \rangle + \langle D_y \rangle)$ can be separated into the average path distance of all interactions in x- and y-directions respectively. For the sake of simplicity, we may take a binary-sized $L \times L$ system with $L = 2^n$, thus, totalling L^2 interactions in each direction[1].

Snake strategy   Applying the Snake strategy to the system, the TTN first coarse grains in x-direction with each interaction in x-direction being connected in the lower half (compare Fig. 3.14b). Thus, we have L coarse-grained subbranches with each one following the one-dimensional case equivalently. Consequently, the computation of $< D_x >$ follows the same structure of Eq. (3.2) except for the sum running to $\frac{1}{2} \log N$ and that there are now L of these sums for each coarse-grained "system column":

$$\langle D_x \rangle = L/N \sum_{k=1}^{\log N/2} \frac{L}{2^k}(2k-1) \stackrel{(N=L^2)}{=} \sum_{k=1}^{\log L} 2^{-k}(2k-1)$$

For each interaction in y-direction, however, the complete lower half of the TTN with distance $\log N/2$ always has to be gone through in the corresponding paths. In other words, these interactions recombine only in the higher half of the TTN topology offsetting each path distance D_p by $2 \frac{1}{2} \log N = 2 \log L$ (once going up through the half-tree and once going down). Thus, there are $\frac{L}{2}L$ interactions with distance $D_p = 2 \log L + 1$, $\frac{L}{4}L$ with $D_p = 2 \log L + 3$, $\frac{L}{8}L$ interactions with $D_p = 2 \log L + 5$, etc.:

---

[1] L^2 for each direction with periodic boundary conditions and $(L^2 - L)$ for each direction in case of open boundary conditions
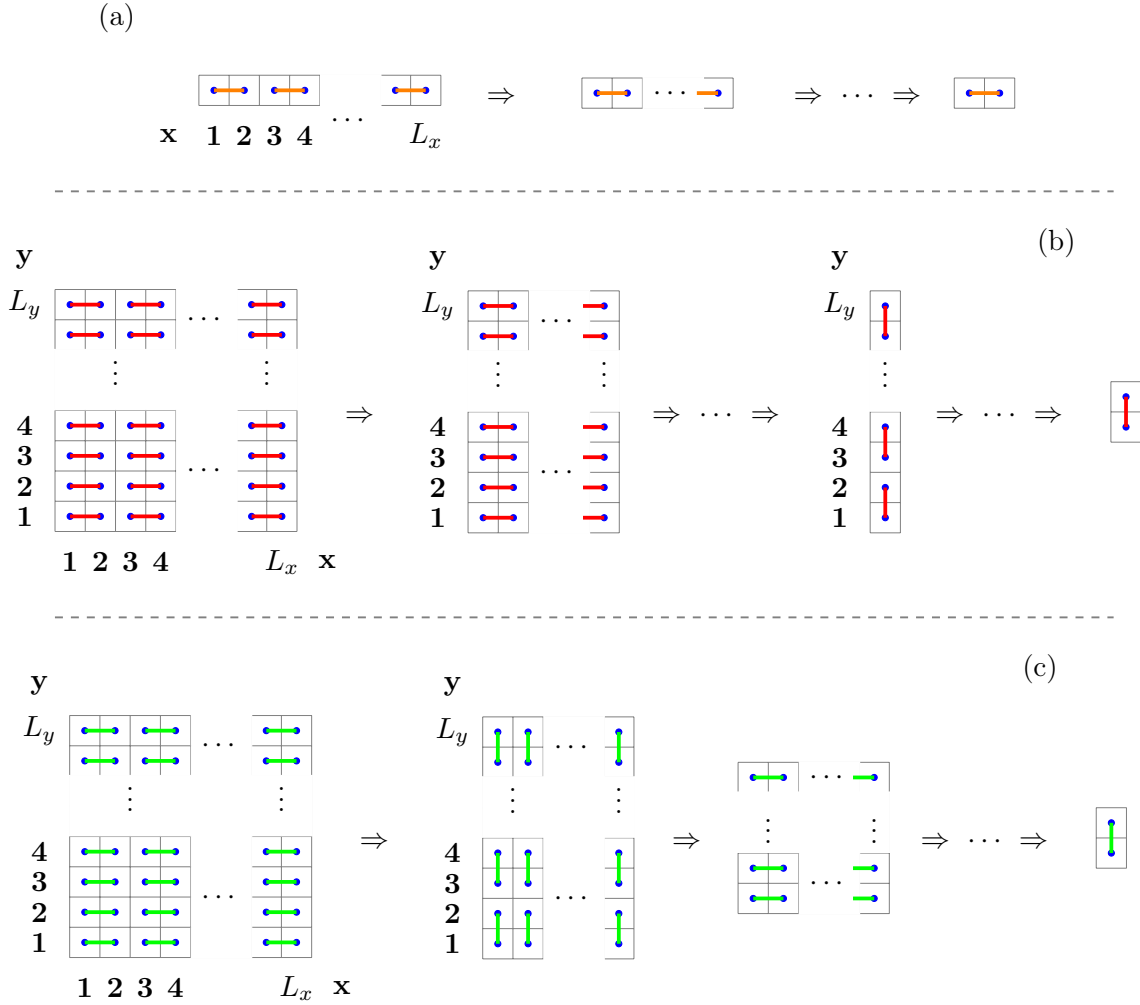
(a)

(b)

(c)

FIGURE 3.14:   Renormalisation procedures performed by each layer of the TTN for different underlying mappings. (a) Coarse-graining procedure of a one-dimensional system with L_x sites: Each tensor in the layer $\Lambda_{\mathcal{L}}$ of the TTN renormalises two neighbouring physical sites into an artificial site. The next higher layers continue merging neighbouring artificial sites until the last two sites are linked by the topmost link in the TTN. (b) Coarse-graining procedure of a two-dimensional system when exploiting the *Snake-like pattern* (Fig. 3.11b): The tensors in the layers Λ_l with $l = [\mathcal{L}, ..., \mathcal{L} - \log L_x]$ renormalise the entire first dimension of the system in x-direction, before the layers with $l = [\log L_y, ..., 1]$ coarse-grain the system in y-direction. (c) Renormalisation procedure of a two-dimensional system when exploiting the strategies of Fig. 3.11a,c: Each layer Λ_l coarse-grains the system alternatingly along the x- (for even l) and the y-direction (odd l).

$$\langle D_y \rangle = \frac{L}{N} \left(\frac{L}{2} \cdot (2 \log L + 1) + \frac{L}{4} \cdot (2 \log L + 3) + \frac{L}{8} \cdot (2 \log L + 5) + ... \right)$$

$$= \sum_{k=1}^{\log L} 2^{-k}(2 \log L + (2k - 1)) = 2 \log L \underbrace{\sum_{k=1}^{\log L} 2^{-k}}_{=\log L(1 - 1/L)} + \sum_{k=1}^{\log L} 2^{-k}(2k - 1)$$

Indeed, for the calculation of $\langle D_y \rangle$, this can be seen as equivalently to the one-dimensional

case with this offset in D_p and L times the interaction. These lower subbranches completely coarse-grain the system in x-direction reducing the problem to the y-direction at the layer $l = \log L$.

Consequently, the average interaction distance of the complete system with the Snake strategy equals

$$\langle D_{\text{Snake}} \rangle = \frac{\langle D_x \rangle + \langle D_y \rangle}{2} = \log L \underbrace{\sum_{k=1}^{\log L} 2^{-k}}_{=(1-1/L)} + \underbrace{\sum_{k=1}^{\log L} 2^{-k}(2k-1)}_{=3-\frac{1}{L}(2\log L+3)} = \log L + 3 - \frac{1}{L}(3\log L + 3)$$

which diverges logarithmic for $L \to \infty$. Thus, using these naive Snake patters leads to a logarithmic divergence of the average interaction distance and thereby to a with system size increasing incapability of the network to properly capture the entanglement properties of the underlying system.

Hilbert curvature The fundamental idea of the Hilbert-curvature is to minimise the increase of the spatial distance when mapping two arbitrary sites from a high-dimension to a one-dimensional system. This is done by iteratively merging plaquettes of the system. The average interaction distance in the TTN can be calculated by analysing the TTN illustrated in Fig. 3.11c which alternatingly groups sites in x- and y-direction at each layer within the TTN (compare Fig. 3.14c). Consequently, when starting coarse-graining in x-direction at the lowest layer, there are immediately $\frac{N}{2}$ interactions acting in x-direction being connected by a tensor in this layer, i.e. with path length $D_p = 1$. The next layer connects $\frac{N}{2}$ interactions in y-direction with a path distance of $D_p = 3$. Moving further upwards in the network, the $l = \log L - 2$-th layer combines $\frac{N}{4}$ interactions again acting in x-direction with $D_p = 5$, and the $l = \log L - 3$-th layer combines $\frac{N}{4}$ interactions in y-direction with $D_p = 7$. Following this pattern, the average path distance in x-direction reads

$$\langle D_x \rangle = \frac{1}{N}\left(\frac{N}{2} \cdot 1 + \frac{N}{4} \cdot 5 + \frac{N}{8} \cdot 9 + ...\right) = \sum_{k=1}^{\log N/2} 2^{-k}(4k-3) \stackrel{(N=L^2)}{=} 5 - \frac{1}{L}(4\log L + 5)$$

while the average path distance in y-direction equals

$$\langle D_y \rangle = \frac{1}{N}\left(\frac{N}{2} \cdot 3 + \frac{N}{4} \cdot 7 + \frac{N}{8} \cdot 11 + ...\right) = \sum_{k=1}^{\log N/2} 2^{-k}(4k-1) \stackrel{(N=L^2)}{=} 7 - \frac{1}{L}(4\log L + 5) \, .$$

Consequently, the average path distance $\langle D_{\text{2D}} \rangle = \frac{1}{2}(\langle D_x \rangle + \langle D_y \rangle) \to 6 \ \ (L \to \infty)$ for this TTN structure converges in the thermodynamical limit making this strategy a solid approach for capturing the entanglement properties of the underlying system via TTN.

3.2.2 Area law for Tree Tensor Networks

As discussed in Sec. 2.2.1, the von Neumann entanglement entropy for a bipartition of the system is directly linked to the bond-links of a Tensor Network. In this context, the loop-free nature of a TTN enables an efficient way of measuring the entropy captured in this approach of representing a quantum many-body wave function.
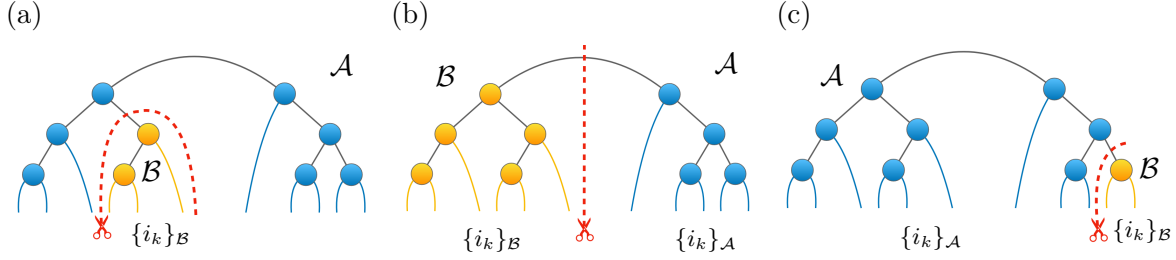
(a)   (b)   (c)

FIGURE 3.15: Possible TTN bipartitions with the links $\{i_k\}_\mathcal{A}$ (blue) and $\{i_k\}_\mathcal{B}$ (yellow) addressing the physical sites within the bipartitions $\mathcal{A}$ and $\mathcal{B}$, respectively.

In every TTN state $|\psi\rangle$, each internal link ν connecting two tensors bipartites the underlying lattice $\mathcal{L}$ into two subsystems $\mathcal{A}^{[\nu]}$ and $\mathcal{B}^{[\nu]}$ (see Fig. 3.15). This allows the TTN to be rewritten for each ν in the form of the Schmidt decomposition

$$|\psi\rangle = \sum_{\alpha=1}^{m_\nu} \lambda_\alpha^{[\mathcal{A},\mathcal{B}]} |\psi^{[\mathcal{A}]}\rangle \otimes |\psi^{[\mathcal{B}]}\rangle \qquad (3.3)$$

where m_ν is the bond-dimension of ν and $\lambda_\alpha^{[\mathcal{A},\mathcal{B}]}$ are the Schmidt-values. For each of the bipartitions, the Schmidt rank is upper-bounded by the bond dimension m_ν of the corresponding link ν. Consequently, the bond-dimension m_ν as well limits the amount of bipartite entanglement the TTN is able to capture.

If we assume the physical state ψ to fulfill the entropic area law, the entanglement entropy $S(\mathcal{A}^{[\nu]})$ of such a bipartition scales with the number of sites γ_ν forming the boundary ∂_ν of $\mathcal{A}^{[\nu]}$:

$$S(\mathcal{A}^{[\nu]}) \sim \gamma_\nu \qquad (3.4)$$

For a practical illustration regarding the TTN representation $|\psi\rangle$ of the state ψ, we consider $\mathcal{L}$ to be a two-dimensional system with $N = L \times L$ sites. In this case, the required bond dimension to encode the area law of the physical state ψ scales exponentially with L (see Sec. 2.2.1 for more details on entanglement bounds in Tensor Networks). To be even more accurate, the bond dimension m_ν of each link ν should scale with

$$m_\nu \approx e^{c\gamma_\nu} \qquad (3.5)$$

where c is a constant factor. This exponential scaling leads to the fact, that with increasing system size for a two-dimensional system, a TTN representation eventually fails to faithfully represent area law states as it rapidly becomes unfeasible to deal with such an exponentially large network especially with $\gamma_{\nu=1} \sim L$. And going beyond a 2D system to even higher dimensions, this scaling becomes all the more critical for the TTN.

3.2.3   A practical example: 4x4-System

To give a hands-on illustration, this section will describe the optimisation of a TTN for analysing high-dimensional systems in a step-by-step procedure for a two-dimensional isotropic Hamiltonian on a square 4×4-Lattice with open boundary conditions consisting of nearest-neighbour interactions $\mathcal{H}_p$. Therefore, we assume the complete Hamiltonian to be a product of interactions $\mathcal{H} = \sum_p \mathcal{H}_p$ with every interaction $\mathcal{H}_p$ being described as a Tensor Product Operator (TPO) (see Sec. 2.2.3). For sake of simplicity, we will restrict ourselves in this example to Hamiltonians $\mathcal{H}$ containing exclusively *(i)* local terms $\mathcal{H}_p = h_{i_p}^p$ (acting on the site
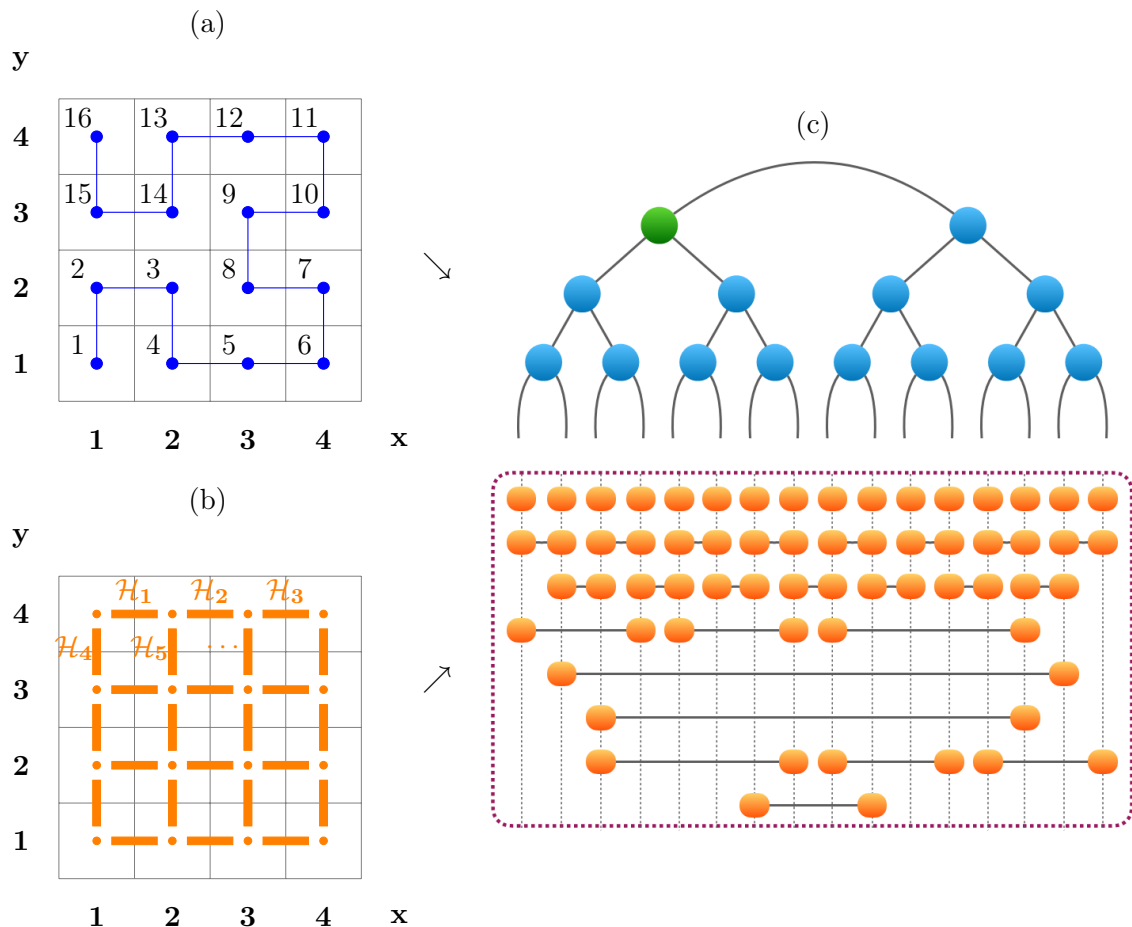
FIGURE 3.16:   Mapping of a two-dimensional 4×4 system with nearest-neighbour interactions and open boundary conditions to a one-dimensional TTN via Hilbert curvature. (a) The Hilbert curvature defines the mapping for the local sites $(x, y) \to i \in \{1, ..., 16\}$ from two into a one-dimensional path. (b) All interaction parts $\mathcal{H}_p$ of the two-dimensional Hamiltonian including local terms (orange dots) and nearest-neighbour interactions (orange rectangles), the latter enumerated from $\mathcal{H}_1$ to $\mathcal{H}_{24}$. (c) The operators in (b) are mapped by the Hilbert curvature in (a) towards the TTN. Each interaction $\mathcal{H}_p$, in general, becomes long-range interaction indicated as orange TPOs, making up the mapped one-dimensional model.

i_p) and *(ii)* two-body interaction $\mathcal{H}_p = h^{p,1}_{i_p} h^{p,2}_{i'_p}$ between the physical sites i_p and i'_p. However, the basic idea of the TTN approach can, in general, be applied to more complex Hamiltonians $\mathcal{H}$ and systems with periodic boundary conditions.

This optimisation procedure for the TTN consists of four main parts: *(i)* Mapping of the high-dimensional systems towards the TTN structure, *(ii)* initialising the TTN, *(iii)* initialising the TTN link-operators, and *(iv)* optimising the internal TTN with the mapped Hamiltonian.

Step 1 (Mapping) — Before starting the actual analysis, we map the system as described above in Sec. 3.2.1. For our 4×4-example, Fig. 3.16 illustrates the mapping of the two-dimensional system towards the one-dimensional TTN structure. The Hilbert curvature defines the mapping for the local sites $(x, y) \to i$ in the two-dimensional system to

their one-dimensional correspondence $i \in \{1, ..., 16\}$ used for the TTN analysis. Given this mapping, all interaction parts $\mathcal{H}_p$ of the two-dimensional Hamiltonian become transformed accordingly into the one-dimensional basis. Crucially, all nearest-neighbour interactions in the two-dimensional plane are mapped to general long-range interactions in the one-dimensional representation. Thus, the Hamiltonian which in 2D consists of nearest-neighbour interactions only becomes a one-dimensional long-range model.

Step 2 (Initialisation TTN) — Once the system is mapped and the physical sites are assigned to the physical links of the TTN, we start constructing the tree structure. This construction can be done layer-by-layer from the bottom beginning with the physical links up to the top following the instructions in Sec. 3.1.4. In this example, we would construct a perfect binary TTN while, in general, for a total number of sites which is not a power of 2, the initialisation of the lowest layer has to be adapted towards a complete TTN since not all physical sites will be attached to a tensor in this lowest layer. After initialising such a random tree which is isometrised towards the tensor $\mathcal{T}^{[1,1]}$, the link operators within the TTN are initialised.

Step 3 (Initialisation Link-operators) At this stage, the coarse-grained Hamiltonian parts $\mathcal{H}_p^{[\nu]}$, i.e. link-operators, are calculated for each link ν in the network. This is done in a bottom-top approach following the procedure presented in Sec. 3.1.5. Fig. 3.17 illustrates this procedure for our 4×4 example contracting for each layer the corresponding Hamiltonian parts with the proper tree tensor and its complex conjugate. Note, that after this initialisation, the effective Hamiltonian $\mathcal{H}_{eff}^{[1,1]}$ is already present for the tensor $\mathcal{T}^{[1,1]}$ towards which the TTN is isometrised.

Step 4 (Optimisation) After initialising all link-operators, the algorithm starts sweeping through the network optimising each tensor. The exact sweeping order can, in general, be chosen arbitrarily from the user. When selecting one tensor to optimise, say $\mathcal{T}^{[2,1]}$ as illustrated in Fig. 3.18, the effective Hamiltonian has to be calculated following the description of Sec. 3.1.4. Having all link operators initialised and the TTN in an isometric form, this can be done easily by properly contracting the effective Hamiltonian parts $(\mathcal{H}_{eff}^{[l,k]})_p$ with all the tensors along the path from the tensor the TTN is current isometrised towards, to the target tensor. While performing this contraction along the path, it might be wise to isometrise the TTN towards the target tensor along the way by performing a QR-decomposition. Once we isometrised the TTN towards our target tensor, the local optimisation solves the underlying eigenvalue problem as illustrated in Sec. 3.1.4. When executed via ARPACK implementation, the optimisation algorithm only requires the knowledge of the action of the effective Hamiltonian $\mathcal{H}_{eff}^{[l,k]}$ on the tensor to optimise $\mathcal{T}^{[l,k]}$. Thus, we iteratively need to perform the contraction of $\mathcal{H}_{eff}^{[l,k]}$ with $\mathcal{T}^{[l,k]}$ which can be done by contracting each parts $(\mathcal{H}_{eff}^{[l,k]})_p$ of the effective Hamiltonian separately. Afterwards, the eigenvalue solver returns the eigenvector with the lowest eigenvalue, i.e. energy, which we use to replace the old tensor in the network as a result of the local optimisation. After optimising a single tensor, we move on to the next tensor of the global optimisation sweep and after one sweep, we perform another one until the selected values, such as global energy, expectation values, or entanglement entropy, converge.
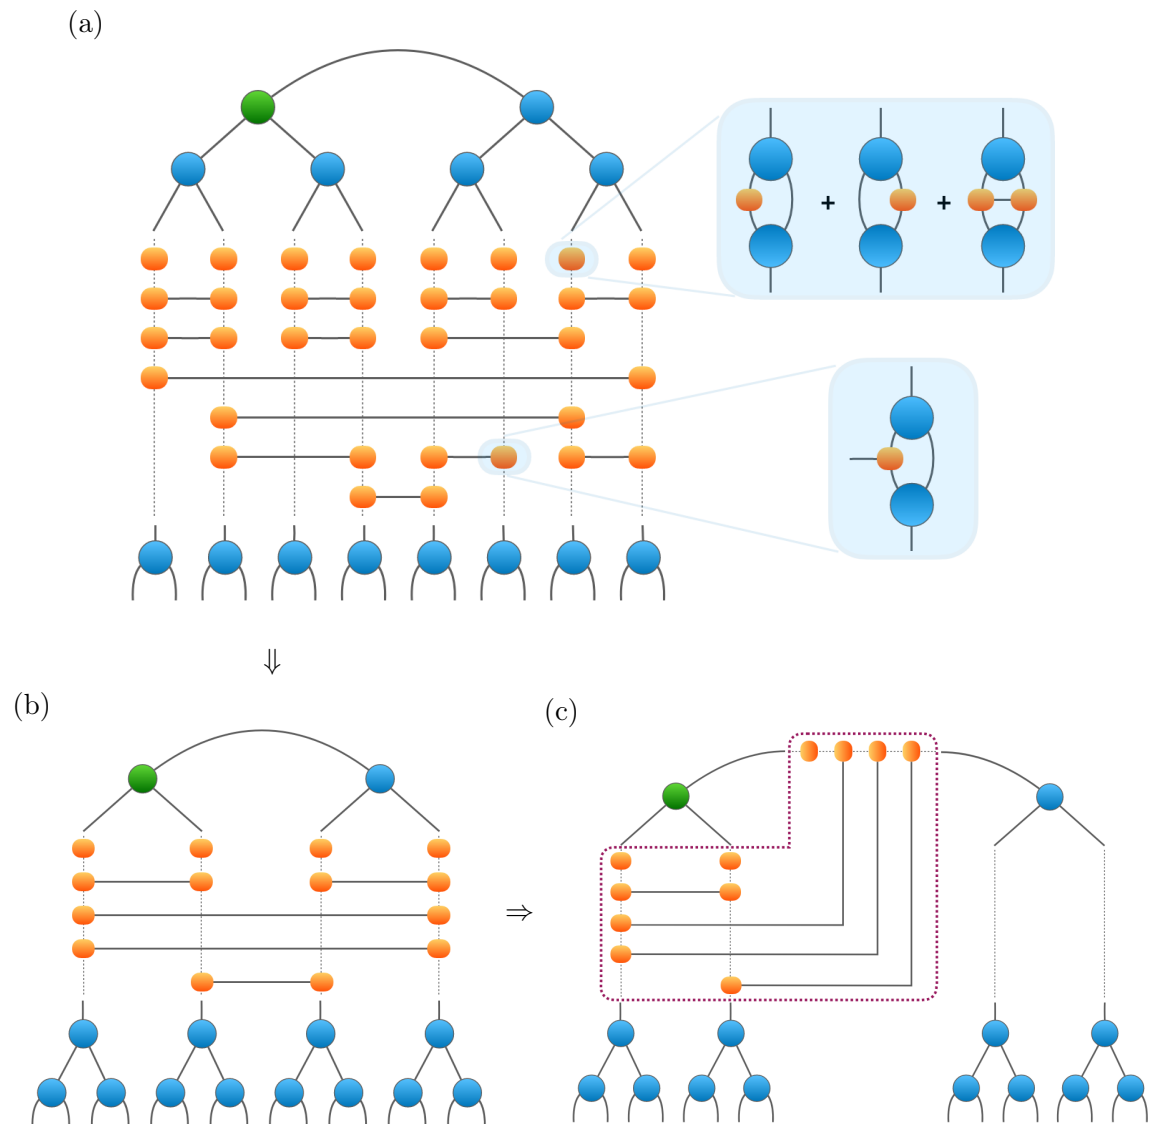
FIGURE 3.17:   Initialisation of the Hamiltonian parts of Fig. 3.16c within the isometrised TTN as link operators. (a) Each tensor of the TPO representing the physical interactions $\mathcal{H}_p$ have been contracted with the corresponding tensor in the lowest layer of the TTN and its complex conjugate (see inlet) making up the TPO tensors of the different link operators. (b) Analogously, the link operators of (a) are coarse-grained upwards with the next layer. (c) The last contraction is done with the tensor $\mathcal{T}[1,2]$ resulting in all the link operators (in purple box) which constitute the effective Hamiltonian $\mathcal{H}_{eff}^{[1,1]}$ for $\mathcal{T}^{[1,1]}$.

Add-on (Towards numerical complexity and periodic boundary conditions)   Having presented a 4×4 system with open boundary conditions you may ask about the feasibility of simulating periodic boundary conditions as well. In fact, the difference between periodic and open boundary computations in terms of the numerical complexity of the TTN is merely a matter of the number of Hamiltonian parts $N_{\mathcal{H}}$, that is, it impacts sub-leading terms in the overall algorithmic complexity of m^4. Indeed, for open boundary conditions $N_{\mathcal{H}}$ is lower then for periodic once by a margin scaling with system length $N_{\mathcal{H}}^{periodic} - N_{\mathcal{H}}^{open} \sim L$.
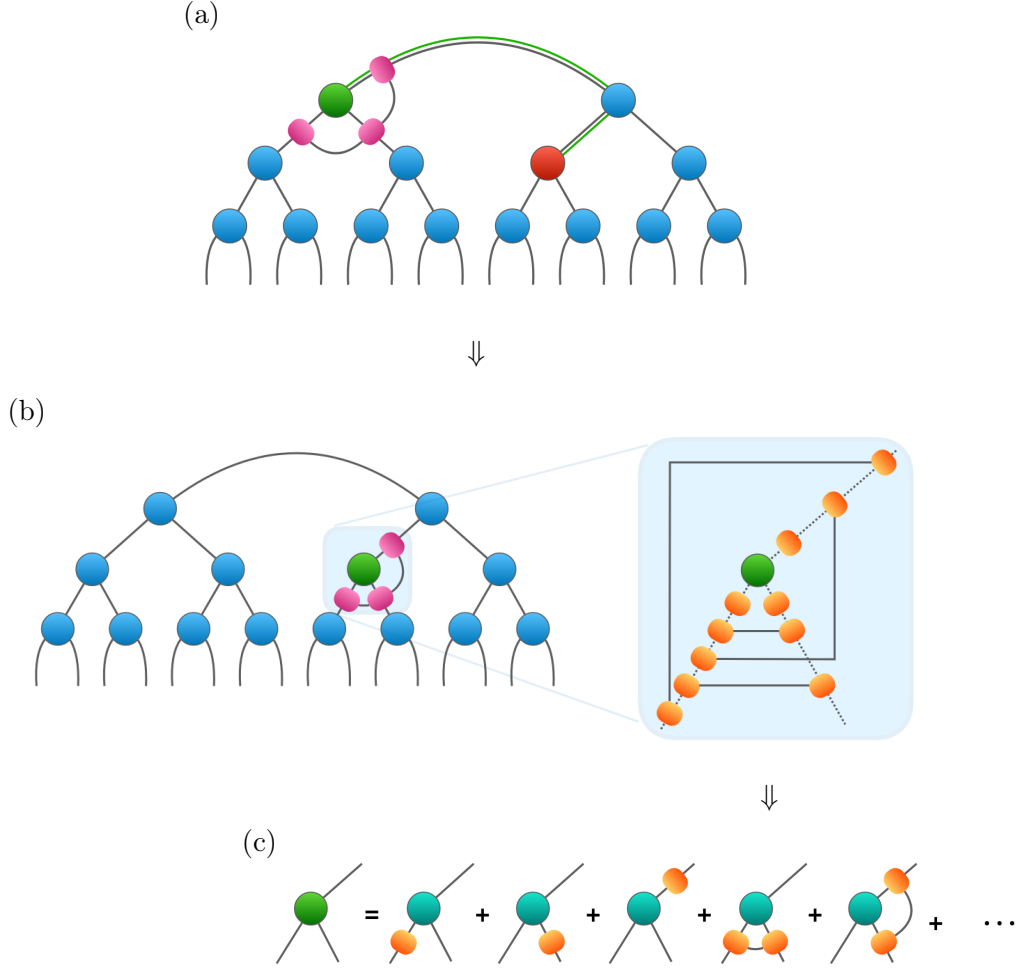
FIGURE 3.18: Optimisation of a tensor for the 4×4 TTN simulation. (a) The TTN is isometrised to $\mathcal{T}^{[1,1]}$ (green) and all link operators are initialised with the effective Hamiltonian $\mathcal{H}_{eff}^{[1,1]}$ (purple) containing all local and interacting TPOs illustrated in Fig. 3.17c. The target tensor $\mathcal{T}^{[2,3]}$ (red) to be optimised is connected to $\mathcal{T}^{[1,1]}$ via the green path. (b) Having isometrised the network towards the target tensor (now green) and contracted the link operators accordingly, the effective Hamiltonian $\mathcal{H}_{eff}^{[2,3]}$ (purple) arises for the optimisation consisting of different TPOs (see inlet). (c) For the optimisation, the action (turquoise) of the effective Hamiltonian on the target tensor (green) is calculated by summing the contractions of all TPOs in $\mathcal{H}_{eff}^{[2,3]}$ with $\mathcal{T}^{[2,3]}$.

The complete optimisation scales maximally with the complexity

$$\zeta_T = \mathcal{O}\left(m^4 d^2 (N_\mathcal{T} N_\mathcal{H} I_\mathcal{T} + \sum_k N_{\mathcal{H},\mathcal{T}^{[k(\Gamma)]}})\right)$$

(plus some sub-leading terms) where d is the physical dimension, $N_\mathcal{T}$ the number of tensors in the TTN, $N_{\mathcal{H},\mathcal{T}}$ the number of operators acting on the tensor $\mathcal{T}$, $I_\mathcal{T}$ the number of iterations optimising the tensor $\mathcal{T}$, Γ the path from one tensor to the next one in the sweeping optimisation with $\mathcal{T}^{[k(\Gamma)]}$ its k-th tensor. Now, going from open to periodic boundary conditions,

we get ΔN_H more Hamiltonian parts. Thus, the additional complexity would read

$$\zeta_{periodic} = \zeta_{open} + \mathcal{O}\left(m^4 d^2 (N_T \Delta N_{H,T} I_T + \sum_k \Delta N_{H,T^{[k(\Gamma)]}})\right) \ .$$

While $\Delta N_H \sim L$ holds true, the additional number of Hamiltonian parts $\Delta N_{H,T}$ affecting each tensor T highly depends on the mapping itself. Still, for the sake of simplicity, we may assume that $\Delta N_{H,T} \sim L$ then we get in all leading terms an increase in complexity $\Delta\zeta = \mathcal{O}(L)$ which scales with system length L.

In contrast, the implementation of periodic boundary conditions is by construction unfeasible for the finite PEPS and would result in an increasing complexity for a 2D-DMRG approach from $\mathcal{O}(m^3)$ to $\mathcal{O}(m^5)$.

To summarise, while for alternative methods the implementation of periodic boundary conditions might not be possible or at least require a change in the network structure resulting in a major increase in the scaling with bond-dimension, the aTTN can efficiently cope with the additional complexity of periodic boundary conditions without any fundamental changes where only the numerical complexity of sub-leading terms is affected.
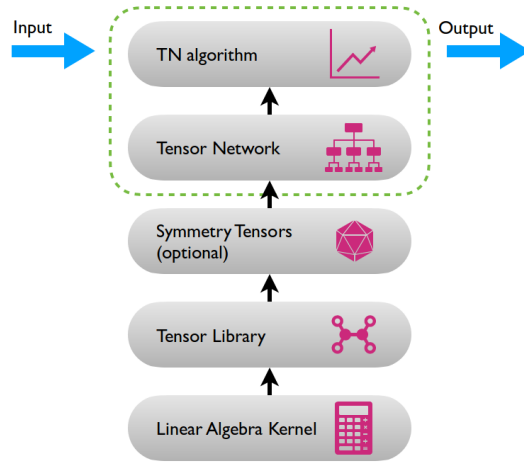
3.3   Technical Implementation

FIGURE 3.19:   Overview of a typical software structure of a generic Tensor Network code with
its different levels ranging from the *Linear Algebra Kernel* performing operations on the hardware
up to the main *TN algorithm* interfacing with the user. The main Tensor Network developments
performed within this thesis, build on a pre-existing *symmetry tensors* library and consists of the two
highest levels (green-dashed box), in particular the implementation of a TTN as *Tensor Network* and
a corresponding ground-state optimisation as *TN algorithm*. Fig. 3.20 illustrates both of the two levels
in more detail.

The general structure of the implementation is illustrated as a simplified overview in
Fig. 3.19. It shows the hierarchical structure of general Tensor Network algorithms in quantum
mechanics. The lowest level is a common linear algebra kernel, such as LAPACK or BLAS,
executing the core functionalities on the computational hardware. Using these mathematical
libraries, the tensor module extends the linear algebra to the tensor algebra described in
Sec. 2.1 including the numerical definition of a tensor type instance and its functions. As
described briefly in Sec. 2.3.2, we can include another level managing the incorporation of
symmetries in the code by, loosely speaking, defining the symmetrically invariant tensor as
a list of tensors, one for each symmetry subspace. This library includes the definition of
symmetric tensors including all operations one may perform to manipulate a symmetrically
invariant tensor (see Ref. [24, 108] for an in-depth description). For general Tensor Networks,
this symmetry layer in the structure is not required. However, it can drastically improve
the computation by taking advantage of the well-defined internal structure of symmetrically
invariant Tensor Networks on the level of the single tensors within the network. The actual
Tensor Network then utilises this library (or the tensor library itself in case of neglecting
symmetries) and the therein defined tensors as the building blocks for the complete network.
Here, the exact geometry is defined together with the important functions required, such
as calculation of observables, optimisation or time evolution. Finally, on top of everything,
the main algorithm for the Tensor Network itself is defined. Later on, the structure of the
implementation of such a Tensor Network level is exemplified in more detail for a Tree Tensor
Network (TTN). In what follows, each of the nodes for the implementation of the TTN is
described in more detail.

3.3.1   Data Structure of a Tree Tensor Network

In general, there are two main ways to implement a Tree Tensor Network for the simulation of quantum many-body systems. On one hand, the TTN can be implemented as a set of nodes where each node consists of a tensor and a list of references, i.e. pointer, to its neighbouring tensors. As an example a binary TTN can be described by the following object *node* written in pseudo-code:

```
1 struct node
2 {
3     tensor T;
4     struct node* left;
5     struct node* right;
6 };
```

In this way of implementation, the topology of the TTN is implicitly stored by the references for the links on each tensor. This concept in general allows very flexible TTN topologies which can be beneficial when the topology shall be adapted to the underlying problem or even on the fly in an optimisation algorithm. The benefit of adapting the network topology can be valuable for its applications in quantum chemistry where the interactions in the system are, in general, not homogeneous but rather determined by molecular structure [178]. Further, a dynamically adaptable topology might turn out the key to automatically finding the most efficient topology for a given problem which is a very promising concept that is not well studied so far and denotes an interesting direction for future research of Tensor Networks.

On the other hand, the TTN can be structured in a static hierarchical order with an array of tensors each being assigned a fixed position within the tree as illustrated in the following pseudo-code.

```
1 struct Tree
2 {
3     tensor[] T;
4     address[] x,y;
5     topology K;
6 };
```

Thus, in this concept, the topology of the TTN is explicitly defined in the data structure. While this static implementation may lack some flexibility in dynamically adapting the topology, it is very well suited for efficient TTN simulations with a well, predefined topology. Fixing the exact topology on a global level allows incorporating the knowledge of the topology in the algorithms manipulating the TTN, such as the optimisation, calculation of observables and computation of entropy, which potentially leads to a more efficient, well-structured implementation and allows for efficient parallelisation techniques on these global algorithms. Consequently, this data structure is well suited for implementing a *perfect TTN*, or a *complete TTN* for which the topology is static and well-defined with each tensor sitting at a predefined position within the tree.
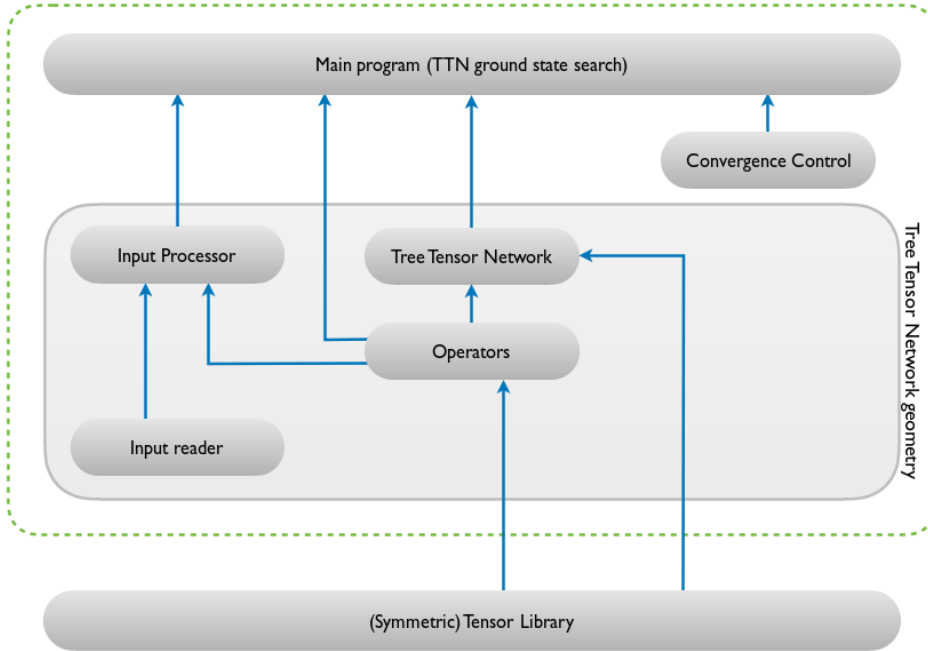
FIGURE 3.20: Software structure of a Tree Tensor Network code performing analyses such as ground-state search or time-evolution. Build on a library for (symmetric) tensors including their operations, a TTN geometry can be properly defined by the modules in the shaded box: The *input reader* gets the data from the user which is then pre-processed for the analysis by the *input processor*. The *Operators* module defines all required low-level data types while the TTN itself is defined by *Tree Tensor Network* module. The modules are interfaced with the *main program* controlling the execution of the specific Tensor Network algorithms with the aid of a *Convergence Control* module.

The code developed in the course of this thesis follows the latter strategy. In particular, it implements a complete TTN that can be applied for high-dimensional systems. This implementation is not restricted to represent a binary TTN but is further able to incorporate tensors of higher order within the complete TTN.

3.3.2 Tree Tensor Network Implementation

The following section is dedicated to the implementation of a TTN software capable of performing ground-state search. In particular, it describes the different modules shown in Fig. 3.20 in more detail and presents some hands-on examples taken in a simplified form from the FORTRAN90 code developed in the course of this thesis. These examples illustrate the most important data structures and most relevant functions for implementing a TTN algorithm for high-performance computations of high-dimensional quantum many-body systems.

Library for (Symmetric) Tensors

As mentioned at the beginning of Sec 3.3, The actual TTN code utilises a library in which a tensor instance is defined together with all operations one may perform to manipulate a tensor (see Sec. 2.1). This library may optionally feature the ability to deal with symmetries to improve computational time and to target a desired global symmetry sector (see Sec. 2.3.2 for more details). For the description of the TTN software described in this section, we will
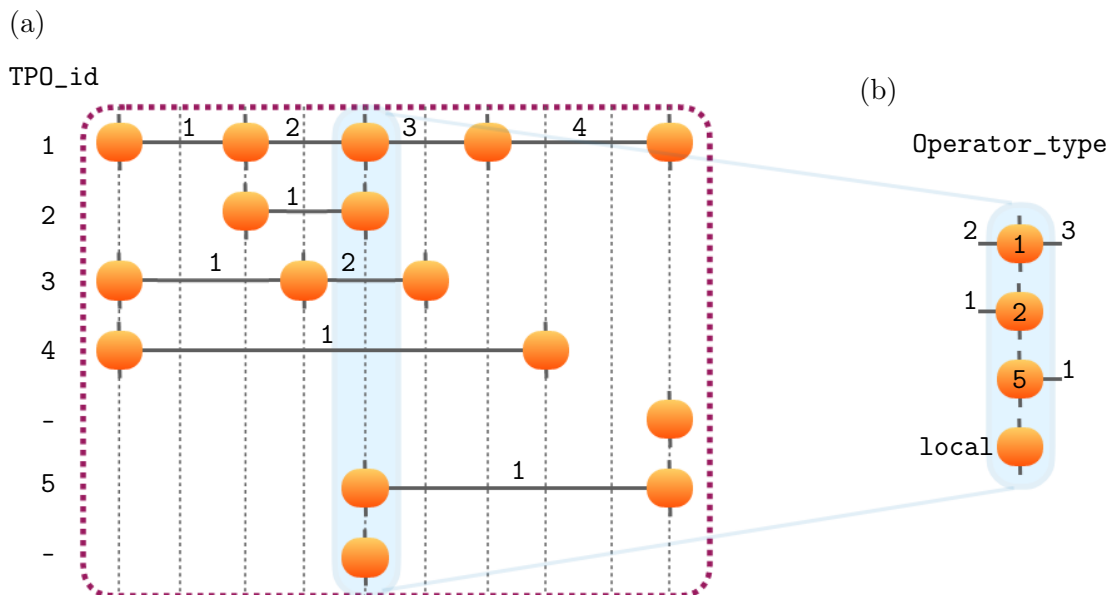
(a)

(b)

FIGURE 3.21: `MPOLayer_type` instance for storing a Hamiltonian $\mathcal{H} = \sum \mathcal{H}_p$ with different interactions $\mathcal{H}_p$. (a) All interactions are represented as TPO. Except for local operations, all TPOs are enumerated (`TPO_id`). Further, all links are enumerated for each TPO. (b) An `Operator_type` instance models all operations acting on the same site i: All local operations are accumulated into the tensor `local` while for the interaction parts, the single tensors are stored with the corresponding `TPO_id` and the link indices.

assume such a library as given and use it as the building block for the developed code. For an in-depth technical description of such a library, Ref. [24] offers profound literature on the implementation of general tensor libraries with Abelian symmetries and Ref. [81, 108, 169] further provide in-depth technicalities for incorporating non-Abelian symmetries onto an existing, ordinary tensor library.

Operators

Building on the tensor library, this module is mainly dedicated to the description of operators, and in particular of the Hamiltonian $\mathcal{H} = \sum_p \mathcal{H}_p$ of the underlying system to analyse. The complete operator is stored as an `MPOLayer_type` instance, which stores and keeps track of all single interactions $\mathcal{H}_p$ of the Hamiltonian and on a global level mimics an MPO. Thus, when using this `MPOLayer_type` instance it can be treated like an MPO while it internally manages the TPOs it consists of (where each TPO represents an interaction part $\mathcal{H}_p$). However, due to numerical efficiency in the global TTN algorithms, the `MPOLayer_type` instance is not decomposed directly into a plain list of TPOs, but rather into a list of `Operator_type` instances which for each site i of the MPO consists of all tensors acting thereon. Thus, effectively, the `MPOLayer_type` instance describes a layer of operators or, more generally, describes a list of lists of tensors.

Fig. 3.21 illustrates the way of storing an Hamiltonian as `MPOLayer_type` instance. In fact, with the illustration in mind, it becomes evident that there are two-dimensions in the data: The number of TPOs, i.e. interaction parts $\mathcal{H}_p$, and the number of physical sites N of the Hamiltonian. Thus, we order the data first in the direction of the latter dimension in the

`MPOLayer_type` instance and subsequently in the former dimension within the `Operator_type`.
Thus, each `Operator_type` consists of a list of tensors modelling the overall action of the
Hamiltonian on one single site i. Here, we can exploit another numerical benefit: All operator
acting locally on a single site can be added together. As example, we can take a Hamiltonian,
such as $\mathcal{H} = \sum_i \alpha_i \sigma_x^{[i]} + \beta_i \sigma_z^{[i]} + \sigma_x^{[i]} \sigma_x^{[i+1]}$, for which we can merge the local fields $\mathcal{H}_{i,\text{local}} =$
$\alpha_i \sigma_x^{[i]} + \beta_i \sigma_z^{[i]}$ into one local operator. Such a summation cannot be done locally, however,
on the latter interaction part $\mathcal{H}_i = \sigma_x^{[i]} \sigma_x^{[i+1]}$ of the Hamiltonian. Thus, the `Operator_type`
instance only needs to store one local tensor at most modelling the local single-site operators
of the Hamiltonian and a list of tensors which together with all other `Operator_type` instance
in an `MPOLayer_type` instance. The two instances might be implemented as follows.

`Operator_type` — An `Operator_type` instance can be implemented as an extension of a
general tensor list. This can be useful, in particular, since a general instance for a tensor list
can be exploited on numerous occasions within a TTN algorithm. As shown in the following
listing, the tensor list would model all tensors corresponding to a many-body interaction term
$\mathcal{H}_p$ while the extension *(i)* adds an optional tensor modelling all local single-site operators
of the Hamiltonian and *(ii)* includes an identification array `id` to keep track of the connec-
tions of each single interaction term $\mathcal{H}_p$ within the complete Hamiltonian, i.e. the global
`MPOLayer_type` instance.

```fortran
1   ! This is a general list of Tensors which can be extended for the
        Operator_type instance:
2   type, public :: TensorList_type
3     private
4     integer :: numtensors = 0
5     type(Tensor_type), dimension(:), allocatable :: tensor
6   end type
7
8   ! This is how an Operator_type instance could look like extending the
        TensorList_type instance above. In comments it can be useful to
        illustrate this instance as
9   !                                    |
10  !                                  --O--
11  !                                    |
12  type, extends(TensorList_type), public :: Operator_type
13    private
14    logical :: contains_local
15    type(Tensor_type), allocatable :: local
16    integer, dimension(:,:), allocatable :: id
17  end type
```

The attributes `contains_local` and `local` take care of the possible presence of a single-
site term in the Hamiltonian while `id` crucially stores all information about the interaction
structure within the Hamiltonian. In particular, this will become a 3×`numtensors` sized
array where each triplet entry `id(:,k)=[TPO_id, id_left_link, id_right_link]` contains
the all the required information regarding the `k`-th tensor in the list: As indicated in Fig. 3.21,
each TPO $\mathcal{H}_p$ becomes a unique identification number `TPO_id` assigned (such as the index p
of the interaction $\mathcal{H}_p$ it represents). Further, all tensors of one TPO are ordered with respect
to the physical sites they are applied to. In this order, all the internal links of the TPO

are enumerated from left to right. Consequently each tensor within the TPO may contain one link to the left and one to the the right indicated by `id_left_link` and `id_right_link`, respectively. Additionally allowing these values to be, e.g. `-1`, in case the tensor denotes the endpoint of a TPO, enables `id` to sufficiently capture all relevant information for a typical TTN analysis.

Note that this kind of implementation only allows representing interaction terms $\mathcal{H}_p$ as TPOs where each tensor has two internal links at most. Thus, the modelling of star-shaped interaction terms as illustrated in Fig. 2.12d is not possible. Furthermore, it is worth mentioning that the first index is the running index in FORTRAN while in other languages such as C++ this is not the case. Thus, for the latter scenario, it is numerically more efficient to exchange the indices.

`MPOLayer_type` — Having implemented the `Operator_type` instance, the data structure of an `MPOLayer_type` is straight-forwardly a list of these `Operator_type` instances as presented in the following.

```fortran
1   ! This could be an MPOlayer_type instance in FORTRAN which might look
        like this in the code commentation:
2   !
3   !          | | |      |
4   !        -O-O-O-  ~  -O-
5   !          | | |      |
6   !
7   ! Designed to be applied to the lower links of a TreeLayer_type
        instance which we will formally introduce later on. This could be an
        example of such an application:
8   !          ___|___    _|_   _|_   |   _|_
9   !         (_____)  (___) (___) | (___)
10  !          | | | |    | |   | |  |  | |
11  !        -O-O-O-O—  O-O—  O-O—O—O-O-
12  !          | | | |    | |   | |  |  | |
13  type, public :: MPOLayer_type
14     integer :: NumSites
15     type(Operator_type), dimension(:), allocatable :: site
16  end type
```
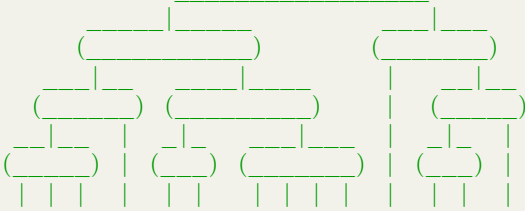
Tree Tensor Network

This code module describes a general TTN and contains all its relevant functions. In particular, this module includes the object definition for one layer of the TTN (`TreeLayer_type`), for a general path (`path_type`) which can be defined in the underlying tree structure, and ultimately the object definition modelling the TTN itself (`tree_type`). For the ground-state search, important functions to be included in this module for the `tree_type` instance would be in particular the isometrisation of the network towards a targeted tensor including the proper contraction of the effective Hamiltonian $\mathcal{H}_{eff}$ (see Sec. 3.1.4), and the local optimisation of a single tensor.

tree_type —   The tree_type instance is modelling a Tree Tensor Network with a strict
layer orientated topology as illustrated in Sec. 3.1 and following Def. 3.1.1. The complete tree
consists of several layers (TreeLayer_type, see below) with each layer consisting of several
tensors. In general, each tensor can have an arbitrary order and only one link per tensor is con-
nected with the next upper tree layer. A certain tensor within the TTN structure is addressed
by two indices, one layer_index for the layer and the second index tensor_index for the
tensor within the specific layer. Thus, a general tree_type instance might be implemented
as:

```fortran
1    ! This is a tree_type instance in FORTRAN which can be a general TTN
         structure like this one:
2    !
3    !    Layer                              _____
4    !                            _____|_____          ___|___
5    !       1                   (_____)         (_____)
6    !                     ___|___  ____|____    |     __|__
7    !       2            (_____) (_____)   |   (_____)
8    !                  __|__   |  _|_  ___|___   |  _|_  |
9    !       3         (_____)  | (___)(_____)  | (___) |
10   !                 | | |    |  | |  | | | |   |  | |  |
11   !
12   type, public :: tree_type
13      ! Tree topology attributes
14      integer :: NumSites, NumLayers
15      type(TreeLayer_type), dimension(:), allocatable :: Layer
16
17      ! Tree isometrisation
18      logical :: is_isometrised = .false.
19      integer, dimension(2) :: iso_position
20
21      ! Link Operators
22      type(MPOLayer_type), dimension(:), allocatable :: Operators
23   end type
24
25   interface tree_type
26      module procedure :: tree_construct
27   end interface
```

where the flag is_isometrised together with iso_position keeps track of the isometrisa-
tion of the TTN. Additionally, the attribute Operators can be used to store the link operators
as discussed in Sec. 3.1.5.

Crucially, the constructor tree_construct of this instance creates a TTN as described
in Sec. 3.1.4: It initialises a tree_type instance for a given number of physical sites N with
random tensors where the isometrisation of the TTN can be done on the fly.

TreeLayer_type —   The TreeLayer_type instance describes a tensor layer within a
Tree_type instance. Therefore, it consists of a list of tensors including the references
(upper_address) for each of the tensors in this layer to its parent tensor (the corresponding
tensor within the next-higher layer to which it is connected). To make this address system
consistent, each layer needs to keep track of its own sites which are pointing downwards to
child tensors. This track-keeping can be done by ordering and enumerating all the downwards

pointing sites of the layer, and then storing the first and the last site indices for each tensor within the layer (`site`). Thus, with the information of `upper_address` from a lower layer and `site` from the next-higher layer, we can reconstruct tensor are connected inter-layers.

```fortran
1    ! This is a TreeLayer_type instance which might look like this:
2    !
3    !            ___|___   |  ___|___   _|_   _|_   |   _|_   |
4    !           (_____)  | (_____) (___) (___)  |  (___)  |
5    !            | | | |   |  | | | |   | |   | |   |   | |   |
6    !   site:   1 2 3 4    5  6 7  ...
7    !
8    type, private :: TreeLayer_type
9       ! Tree Layer attributes
10      integer :: numtensors = 0
11      type(Tensor_type), dimension(:), allocatable :: tensor
12      integer :: cutoff
13
14      ! Topology attributes
15      integer, dimension(:,:), allocatable :: site
16      integer, dimension(:,:), allocatable :: upper_address
17
18      ! Singular values
19      type(real_vector), dimension(:), allocatable :: SingVals
20   end type
```

Additionally, with `sv_present` and `SingVals`, we can keep track of the singular values sitting on each of the link directed upwards in the TTN layer. This might be interesting since *(i)* the singular values are required to calculate the entropy following the description in Sec. 2.2.1, *(ii)* they are beside the global energy a good predictor for convergence of the state representation in a ground state search and *(iii)* can be a good stop criterion for time evolution in which the entanglement may grow beyond the faithful representation power of a TTN with limited bond-dimension m (which is here set as `cutoff`).

`path_type` — This auxiliary object follows Def. 3.1.2 (Path) and should, among others, help to propagate through the TTN during the optimisation procedure. Each path has an *anchor tensor* which is defined as the highest tensor within the hierarchical network structure. This anchor is stored redundantly by its position within the network `anchor_address` and its position `position_anchor` when traveling along the path. The path itself consists of an `4×length` array describing the `steps` of the path. The quartet `steps(:,k)` contains all necessary information for the `k`-th step of the path: *(i)+(ii)* The address [`layer_index`, `tensor_index`] of the `k+1`-th tensor as a tuple for the tree layer and the tensor position within the layer, *(iii)* the index for the link of the `k+1`-th tensor which is connected to the `k`-th tensor and accordingly *(iv)* the index for the same link on to the `k`-th tensor. Consequently, the instance can be implemented as shown in the following listing.

```fortran
1    type, public :: path_type
2       ! Path attributes
3       integer :: length, position_anchor
4       integer, dimension(2) :: anchor_address
```

```
5
6      ! Single steps of the path
7      integer, dimension(:,:), allocatable :: steps
8    end type
```

`tree_optimization_sweep` — Having encoded the TTN as a `tree_type` instance, we can implement an optimisation sweep as described in Sec. 3.1.4 simply by looping over all tensors within the TTN structure, for each one contracting the effective Hamiltonian towards the targeted tensor, isometrising the network accordingly and optimising this tensor.

```
1     ! This could be an optimisation sweep for a tree_type instance:
2     subroutine tree_optimization_sweep( tree, energy, arnoldi_tolerance )
3       class(tree_type), intent(inout) :: tree
4       double complex, intent(out) :: energy
5
6       double precision, intent(in), optional :: arnoldi_tolerance
7       integer :: ll, tt
8
9
10      do ll = 1, tree%numlayers
11        do tt = 1, tree%layer(ll)%numtensors
12          !
13          ! Isometrise the network towards the target tensor at position [
                 ll,tt]
14          !
15          call tree%isometrise_towards( [ll,tt] )
16
17          !
18          ! Optimise target tensor (one-tensor update)
19          !
20          call tree%optimise_tensor( energy, tolerance = arnoldi_tolerance
                 )
21        end if
22      end do
23    end subroutine
```

Note, that as mentioned in Sec. 3.2.3, for the sake of performance, we can implement one function `tree%isometrise_towards()` for performing both, the isometrisation and the contraction of the link operators at once, since they both manipulate tensors along the same path . Thus, in this way, the code will only iterate through the path from the last optimised tensor towards the target tensor once instead of twice. Finally, the function `tree%optimise_tensor()` performs the local optimisation as described in Sec. 3.1.4 and basically wraps the ARPACK solver straightforwardly into the code.

In the same manner as shown here, we can further include different optimisation techniques, such as the two-site optimisation, the space expansion technique, or even a gradient descent method, and a different sweeping sequence, for instance optimising in a random sweep order or performing a user-defined sweeping sequence.

Input reader

The input reader module reads in pre-defined files which among others model the Hamiltonian to analyse by defining the interactions and coupling-strengths or set different control parameters for the ground state search, such as the number of maximum sweep iterations or the detailed options of the local optimisation procedure. In particular, from the given interactions $\mathcal{H}_p$ of a generic Hamiltonian $\mathcal{H} = \sum_p \mathcal{H}_p$, the module constructs all interactions in the system as TPOs respecting the chosen boundary conditions. Additionally, we can read single operators which for instance shall be measured as observables during the optimisation procedure.

Input processor

Out of the interactions defined by the input reader module, this module defines the `MPOLayer_type` structure used to model the Hamiltonian. Further, at this level, we can take care of the exact mapping of the operators from the high-dimensional space towards the one-dimensional `tree_type` instance as illustrated in Sec. 3.2. Further, this module is the interface keeping track of the observables to be measured and, in particular, their mapping, thus where they shall be applied in the TTN structure.

Convergence control

This module is independent of the Tensor Network structure itself and designed to track and control the convergence for a Tensor Network simulation. Thus, this module is to be incorporated in the main program and is an interface for managing simulation parameters and for enhanced stopping criteria based on variables (i.e. the current energy) of the simulation.

On one hand, the module defines all the important simulation parameters for a ground-state search, such as tolerance of the local Arnoldi solver, or potentially even parameters for time evolution. The module further tracks the number of iterations, the energy at each iteration as well as the desired numerical precision of the objectives of the final target state.

At the end of each iteration, the current energy has to be provided as an interface for the main program. Based on the number of iterations and the development of the energy, the module can adapt and return important simulation parameters, such as the suggested tolerance in solving the eigenvalue problem. For this particular parameter, it turns out empirically that the estimate

```
tol = max(tol_min, min(tol_max, abs( (energy-energy_prev)/energy )/Ntensors ) )
```

, i.e. the relative energy difference after one sweep normalised by the number of tensors `Ntensors` in the network and upper- (lower-)bounded by predefined values `tol_max` (`tol_min`), provides an efficient value for the balance between fast convergence and avoiding to get trapped in local minima.

TTN algorithm

Based on the code structure illustrated above, the algorithm for searching the ground state of a quantum many-body system can be implemented following the description of Sec. 3.1.4 which can be implemented in the following way.

```fortran
!{{{ subroutine start_simulation ()
subroutine start_simulation ()
  type ( Link_type ), dimension (:) , allocatable :: physical_links
  type ( MPOLayer_type ) :: Hamiltonian

  type ( tree_type ) :: tree
  double precision :: current_energy
  double precision :: tol

  integer :: ii
  character :: converged

  !
  ! Reading the Hamiltonian from defined input files
  !
  tree = read_physical_system ( Hamiltonian , physical_links )

  !
  ! Initialising the Tree Tensor Network
  !
  tree = tree_type ( Nsites , physical_links , cutoff , .true.)

  !
  ! Initialising the Link Operators
  !
  call tree%create_operators ( Hamiltonian )

  !
  ! Calculating energy and observables (optional)
  !
  current_energy = tree%calculate_energy ()
  call measure_observables ()

  !##################################################################
  !    Start Iterating
  !##################################################################

  do ii = 1, control%get_max_number_of_iterations ()
    !
    ! Start optimization
    !
    call control%get_iteration_parameters ( arnoldi_tolerance = tol )

    call tree%optimization_sweep ( energy = current_energy , &
        arnoldi_tolerance = tol )

    converged = control%check_convergence ( current_energy )

    select case ( converged )
      case ('f')

      case ('a')
        write (*,*) '_____',
        write (*,*) 'Converged on absolute deviation'
        print *, ii
        exit
      case ('r')
```

```
58            write (*,*) '————————————————————————————————',
59            write (*,*) 'Converged_on_relative_deviation'
60            print *, ii
61            exit
62         case ('d')
63            write (*,*) '————————————————————————————————',
64            write (*,*) 'Failed_converging'
65            exit
66         case default
67            write (*,*) '————————————————————————————————',
68            write (*,*) 'Error_in_checking_convergence'
69            stop
70      end select
71
72      call measure_observables ()
73   end do
74 end subroutine
```

The simulation is set up by reading the systems Hamiltonian together with information about its physical sites, initialising the `tree_type` instance based on the physical system, initialising the link operators within the TTN, and optionally calculating observables for the randomly initialised network. Afterwards, the simulation starts by iterating through the number of sweeps we want to perform. As described above, the control module takes care of the hyperparameters for each sweep, such as the Arnoldi tolerance, and tracks the global energy E after each sweep. Based on this tracking, the convergence is checked using the value `converged`. In practice, it is feasible to track the relative value $|\frac{\delta E}{E} < \epsilon_{rel}|$ and the absolute convergence $|\delta E| < \epsilon_{abs}$. Further, at each iteration, we can measure and keep track of the observables of interest allowing us to follow how they develop with convergence. This might be interesting, in particular, to investigate whether the global optimisation has reached a global minimum or is stuck in a local minimum (see Appendix of the publication in Chapt. 4).

3.3.3 Towards High-Performance Computing

The final part of this chapter shall provide a general discussion on the performance of a Tree Tensor Network code which is implemented in the way described in Sec. 3.3.2. In particular, it should illustrate the perspectives concerning the parallelisation capabilities of a TTN simulation performed on High-Performance Computing (HPC) systems.

There are several computations within a TTN algorithm that can be parallelised in different ways. On the lowest level, we can parallelise the tensor contractions, i.e. the LAPACK routines performing the underlying matrix-multiplications. This can be done efficiently using GPUs. However, it is crucial to respect the sizes of the tensors to be contracted, since there can be a detrimental overhead moving the data to the GPU and back for small tensors. Having an educated heuristic, the incorporation of GPUs or even GPU clusters might be an interesting pathway for further development which has not been exploited within the scope of this thesis.

Secondly, the computations performed on symmetrically invariant tensor can be done to a large extent independently for the different symmetry sectors (see Sec. 2.3). Since these tensors can be seen numerically as block-diagonal instances, we could parallelise computations, such as tensor decompositions or tensor contractions, similarly to the parallelisation of a matrix decomposition or matrix multiplication, respectively, for a block-diagonal matrix. However,

the benefit of this parallelisation highly depends on the distribution of symmetry sectors, i.e.
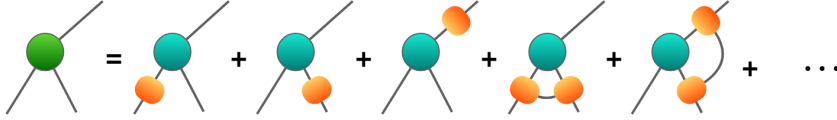the number and sizes of different blocks in the tensor.

FIGURE 3.22:   Contraction scheme for a local optimisation. During the optimisation, the action
of the effective Hamiltonian $\mathcal{H}_{eff}$ on a local tensor (green) is calculated by summing the contractions
of all TPOs in $\mathcal{H}_{eff}$ with $\mathcal{T}$. In practice, these contractions make up for about 75% in a TTN ground
search algorithm when performed on one CPU while they can be efficiently performed in parallel.

Furthermore, as described in Sec. 3.1.4, the local optimisation of a single tensor within the
TTN can be parallelised since we need to sum the contraction of several Hamiltonian parts
with a single tensor (see Fig. 3.22). Each of these contractions can be performed independently
and since they all act on the same space, the computational time should be of the same order
for each contraction. In turns out that in the ground state search this parallelisation idea is
the most efficient one when using a limited amount of CPUs only, especially since depending
on the bond-dimension about 75% of the serial CPU time is spent performing this set of
contractions. However, it clearly has an upper limit of efficiency, namely the number of TPO
operators that have to be contracted in the local optimisation which highly depends on the
physical system itself. Nevertheless, with this optimisation, we can reach reasonable physical
results for systems up to $N = 16 \times 16$ (containing in total 2^{256} degrees of freedom) when
running 3 days on standard HPC systems, such as CINECA [179]. In this case, we perform
the parallelisation with OpenMP and reach typically a ground state with a precision in the
energy density of the order $\sim 10^{-5}$, which is sufficiently accurate to obtain valuable insights
into the physics behind. However, this parallelisation technique does not scale sufficiently well
with system size N. Thus, since the number of tensors to be manipulated in the optimisation
scales with the system size N we investigate, it is quite challenging to obtain high-precision
results for higher system sizes like $N = 32 \times 32$ or even larger (obviously depending on the
system as well).

Another possible form of parallelising a TTN ground-state search would be an MPI paral-
lelisation of the global optimisation technique which could indeed offer the desired scalability
with system size N. This parallelisation aims to perform the global optimisation algorithm
by optimising all local tensors in parallel. The success of this idea has already been shown
in Ref. [180] where it was introduced for an MPS. The main principle therein is to bring the
network into the form of the Schmidt-decomposition obtaining the singular values, or Schmidt
values, on an internal link. Now, we contract the singular values into both sides of the network
keeping the inverse of them on the link. In this way, each side of the network is gauged so
that it can perform the local optimisations independently. Applying this idea to a TTN, we
can divide the tree into different sub-branches containing a fixed number of tensors. All the
sub-branches will be optimised locally on different computational nodes in parallel via MPI.
After the optimisation, the information affecting neighbouring branches will be passed to the
corresponding nodes.

While at the final stage of this thesis, the latter option has not been implemented, this
might be a further major step forward in the development of Tensor Network codes, and
thereby the investigations of high-dimensional quantum many-body systems since it can offer
an ideal scalability of the code with system size N.
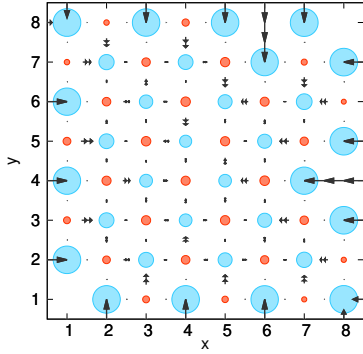
4

Applications to Lattice Gauge Theory

This chapter focuses on the application of Tensor Network algorithms on models in the field of Lattice Gauge Theory (LGT). Lattice Gauge Theories are rich formulations of quantum field theories with gauge symmetries on a discretised space-time lattice. The numerous applications of these theories range from low-energy physics, describing high-temperature superconductivity [181–183] or spin liquids [184, 185], all the way to high-energy particle physics, where they model the fundamental building blocks in the Standard Model [186, 187]. Indeed, since the fundamental forces in the Standard model are mediated through gauge fields, Lattice Gauge Theories are successfully deployed to understand and unveil microscopic processes ruling the dynamics of elementary particles. Thus, being able to capture their complex many-body behaviour is not only a cornerstone of contemporary scientific research but further provides us with a microscopical description of the universe. In particular in high-energy physics, the most prominent Lattice Gauge Theories are found in the study of quantum chromodynamics (QCD) which incorporates non-Abelian gauge symmetries or its counterpart for Abelian gauge symmetries quantum electrodynamics (QED) [188–191].

Solving such interacting quantum field theories, from quantum electrodynamics to the Standard model, is a fundamental goal of modern physics. However, the complexity of describing the underlying many-body phenomena from these strongly-interacting field theories stretches beyond the power of the most advanced analytical and numerical tool available. This challenge makes Lattice Gauge Theories a particularly interesting problem for Tensor Networks, not at least because one of its most competitive methods, the Monte-Carlo simulations, suffers here from the well-known sign-problem in regimes crucial for further novel physical insights. In the last decade, Tensor Networks have already shown great potential in addressing quantum gauge theories on one-dimensional lattices. However, in higher dimensions, these Lattice Gauge systems become highly non-trivial for Tensor Network approaches due to the Jordan-Wigner strings in the underlying fermionic system. Introducing a mapping to a computational basis that aims to get rid of the Jordan-Wigner strings enables the simulation of Lattice Gauge Theory with Tensor Networks. Still, the approach poses, in particular, two technical challenges: *(i)* The Hamiltonian in these systems is not homogeneous, i.e. the physical sites and their dimension varies depending on even and odd sites, with comparably

high local dimension *d*. *(ii)* Additionally to the Hamiltonian terms, a penalty term has to be implemented following the aforementioned mapping.

This thesis incorporates two manuscripts that illustrate this highly non-trivial extension of Tensor Networks to study LGTs in higher dimensions and present the first-ever publications of such a Tensor Network analysis in two- and three dimensions at finite density. In particular, we exploit a Tree Tensor Network to study a Lattice Gauge Theory model with discretised Abelian gauge field which can be linked to quantum electrodynamics (QED).

The first manuscript is presented in this chapter. Therein, we demonstrate the successful applicability of Tensor Network methods for two-dimensional gauge theories. We describe in more detail how to account for the presence of matter beside the quantum gauge fields and how we are efficiently able to allow a finite charge density, a scenario that challenges the most advanced numerical tools, including Monte Carlo with its infamous sign problem. This first application studying a "simpler" (though still highly non-trivial), Abelian Lattice Gauge Theory model in two dimensions significantly reduces the gap between the current numerical capabilities and the ultimate goal of fully understanding the standard model in three dimensions including non-Abelian gauge symmetries. The manuscript was published in *Physical Review X*, **10** [37].

The second manuscript of interest in this context is presented in App. A. This work builds on the former publication extending the Tree Tensor Network analysis to a three-dimensional lattice gauge problem. Due to the fundamental challenge in representing high-dimensional systems with Tensor Networks (as discussed in Sec. 2.2), this publications was the first attempt to simulate a Lattice Gauge Theory in 3+1 dimensions via Tensor Networks. Thus, it shows that Tensor Networks, and in particular Tree Tensor Networks, can be a fruitful direction to solve fundamental long-standing problems. Additionally, Tensor Networks provide an ideal cross-verification tool for near-future quantum simulations of Lattice Gauge Theories available shortly (see Chapt. 5). At the state of submission of this thesis, this second manuscript is accepted for publication in *Nature Communications* and in the process of being published. In the meantime, it is available on arXiv [144].

Two-Dimensional Quantum-Link Lattice Quantum Electrodynamics at Finite Density

Timo Felser[1,2,3] Pietro Silvi,[4,5] Mario Collura[1,2,6] and Simone Montangero[2,3]

[1]*Theoretische Physik, Universität des Saarlandes, D-66123 Saarbrücken, Germany*
[2]*Dipartimento di Fisica e Astronomia "G. Galilei," Università di Padova, I-35131 Padova, Italy*
[3]*Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Padova, I-35131 Padova, Italy*
[4]*Center for Quantum Physics, and Institute for Experimental Physics,*
University of Innsbruck, A-6020 Innsbruck, Austria
[5]*Institute for Quantum Optics and Quantum Information,*
Austrian Academy of Sciences, A-6020 Innsbruck, Austria
[6]*SISSA-International School for Advanced Studies, I-34136 Trieste, Italy*

We present an unconstrained tree-tensor-network approach to the study of lattice gauge theories in two spatial dimensions, showing how to perform numerical simulations of theories in the presence of fermionic matter and four-body magnetic terms, at zero and finite density, with periodic and open boundary conditions. We exploit the quantum-link representation of the gauge fields and demonstrate that a fermionic rishon representation of the quantum links allows us to efficiently handle the fermionic matter while finite densities are naturally enclosed in the tensor network description. We explicitly perform calculations for quantum electrodynamics in the spin-one quantum-link representation on lattice sizes of up to 16×16 sites, detecting and characterizing different quantum regimes. In particular, at finite density, we detect signatures of a phase separation as a function of the bare mass values at different filling densities. The presented approach can be extended straightforwardly to three spatial dimensions.

I. INTRODUCTION

Recent progress in quantum simulations is paving the way for the possibility of studying high-energy physics phenomena with tools developed in low-energy quantum physics [1–13]. In the Standard Model, forces are mediated through gauge fields; thus, gauge-invariant field theories— e.g., quantum electrodynamics (QED) for the Abelian case or quantum chromodynamics (QCD) for the non-Abelian scenario—are fundamental building blocks to our understanding of all microscopic processes ruling the dynamics of elementary particles [14,15]. When discretizing the gauge theories, the dynamical gauge variables obey a lattice formulation of the original quantum field theory, which is referred to as a lattice gauge theory (LGT) [16,17]. LGTs encode many-body interactions satisfying exact constraints, encoding a lattice-discretized version of the local gauge invariance, e.g., in QED, the Gauss law $\nabla \cdot E = 4\pi\rho$. Many of the collective phenomena arising from these theories, including the phase diagram, have yet to be fully

characterized [18], especially for higher spatial dimensions at finite charge density.

Possibly the most successful tools to investigate LGTs are Monte Carlo simulations based on lattice formulations [16,19–23]. However, the Monte Carlo approach suffers from the infamous sign problem for complex actions, e.g., at finite fermion density (matter-antimatter unbalance), which naturally arises in LGTs [13,24]. Another very promising alternative to simulate lattice gauge theories is based on tensor network (TN) methods. They have already shown significant capabilities in describing many condensed matter and chemistry problems and for studying lattice gauge theories in one spatial dimension [12,25–31, 31–41]. So far, very few attempts have been made to capture the phase properties (e.g., at zero temperature) of a lattice analogue of an Abelian gauge theory in higher spatial dimensions [4,42–50], none of them in the presence of fermionic matter at finite density.

In this work, we fill this gap and develop a computationally tractable Hamiltonian formulation of low-energy QED in two spatial dimensions. We show that TN states allow for an accurate representation of its many-body ground state, thus allowing us to identify the different regimes and effectively test the response of the system to a finite density of charge. The study of lattice gauge Hamiltonians at finite chemical potential is, in general,

out of reach for Monte Carlo-based techniques [13,24]: Here, we show that by using an unconstrained tree tensor network (TTN) [51] and the quantum-link formalism of lattice gauge theories [2,52–55], we can face this highly nontrivial setup. The techniques developed in this paper not only provide the basic ingredients for an efficient calculation of the phase diagram of simple lattice gauge models, but they can also be extended to more complex theories and higher dimensions.

We demonstrate the effectiveness of the presented approach by focusing on the low-energy properties, both at zero and finite charge density, of a two-dimensional lattice quantum-link theory with $U(1)$ gauge symmetry. Specifically, we investigate a model involving (spinless, flavorless) Kogut-Susskind matter fermions [16,21] and $U(1)$ electromagnetic gauge fields, truncated to a spin-S compact representation. Hereafter, we set $S = 1$, the smallest representation where all Hamiltonian terms are nontrivial. The calculations for higher spin representations are numerically demanding but straightforward. We investigate the (zero-temperature) phase diagram in the zero global charge scenario without and with finite magnetic coupling. We observe that both magnetic and electric Hamiltonian terms, separately, hinder the creation of a charge-crystal configuration, which emerges at large negative bare masses. However, when electric and magnetic terms are mutually frustrated, the charge crystal is restored. Moreover, we study the ground state in the presence of a finite charge density, which we can directly control in the TN ansatz state. Small charge densities impact the zero-charge phases as follows: In the vacuum regime, charges aggregate at the system (open) boundaries, suggesting the existence of a spatial phase separation between the bulk and the boundaries; this scenario is reminiscent of the classical electrodynamics properties of a perfect conductor, where $\nabla \cdot E = 0$ in the bulk and the excess of charge is redistributed on the outer surface of the conductor. On the contrary, the charge-crystal regime, which is full of matter or antimatter, is characterized by a homogeneous delocalization of the charge hole, resulting in a quasiflat charge distribution in the bulk and therefore reminiscent of a plasma phase [56].

Finally, we stress that the quantum-link formulation provides the ideal tools to establish a connection between LGTs and atomic lattice experiments [57,58]. In this framework, the dynamical gauge fields are usually represented by spin degrees of freedom, which have a natural mapping to typical condensed-matter models, like Hubbard Hamiltonians or locally constrained Ising-like Hamiltonians. These models can be engineered with cold atoms in optical lattices [7,11], or within the very promising experimental setups involving Rydberg atom chains [59,60], and they can be straightforwardly numerically simulated with the presented techniques to verify and benchmark the experimental results and to carefully

and quantitatively compare the limits, the precision, and the efficiencies of the classical and quantum simulations.

The paper is structured as follows: In Sec. II, we present the 2D lattice gauge Hamiltonian and its quantum-link formulation in terms of the gauge-field spin-1 compact representation. We also give some technical details of the tensor network numerical simulations. In Sec. III, we focus on the ground-state properties in the zero-charge sector: We explore the phase space of the model by varying the mass and the electric coupling; we then analyze the effect of a finite magnetic coupling. Section IV is devoted to studying the equilibrium properties at finite charge density. We exploit TTN techniques to investigate how the charges redistribute all over the lattice, depending on the Hamiltonian couplings. Finally, we draw our conclusions in Sec. V and give additional supplementary technical details in the Appendixes.

II. MODEL AND METHODS

We consider a field theory on a 2D square lattice with $U(1)$ local gauge symmetry. The sites of a finite $L \times L$ square lattice host the matter field, while the quantum gauge field lives on the lattice links, with open boundary conditions. Following the Kogut-Susskind (staggered) formulation [16,21], the discretization of the matter field is performed by introducing a staggered fermionic field, whose positive energy solutions lie on the even sites and whose negative ones lie on the odd sites. The matter field is thus described by spinless, flavorless Dirac fermions, whose operator algebra satisfies the usual canonical anticommutation relations $\{\hat{\psi}_x, \hat{\psi}_{x'}^\dagger\} = \delta_{x,x'}$. In particular, in the even sublattice, particles represent fermions with electric charge $+q$ ("positrons"), while in the odd sublattice, holes represent antifermions with electric charge $-q$ ("electrons"). Here, a lattice site x labels a 2D coordinate $x \equiv (i, j)$, and the parity $p_x \equiv (-1)^x = (-1)^{i+j}$ of a site, distinguishing the two sublattices, is well defined on the square lattice (see Fig. 1).

The gauge field is defined on the lattice links, and its algebra is constructed by the electric-field operator $\hat{E}_{x,\mu} = \hat{E}_{x,\mu}^\dagger$ and its associated parallel transporter $\hat{U}_{x,\mu}$, which is unitary, $U_{x,\mu}U_{x,\mu}^\dagger = \mathbb{1}$, and satisfies $[\hat{E}_{x,\mu}, \hat{U}_{y,\nu}] = \delta_{x,y}\delta_{\mu,\nu}\hat{U}_{x,\mu}$. Here, μ (ν) represents the positive unit lattice vector in one of the two orthogonal directions, namely, $\mu_x \equiv (1,0)$ and $\mu_y \equiv (0,1)$; thus, (x, μ) uniquely defines a link. For comfort of notation, we also allow (technically redundant) negative unit lattice vectors $-\mu_x$ and $-\mu_y$, with the convention $\hat{E}_{x+\mu,-\mu} = -\hat{E}_{x,\mu}$, and in turn $\hat{U}_{x+\mu,-\mu} = \hat{U}_{x,\mu}^\dagger$. With such definitions, apart from a rescaling due to the lattice spacing regularization [16,21], the two-dimensional lattice QED Hamiltonian, including a magnetic plaquette term, reads
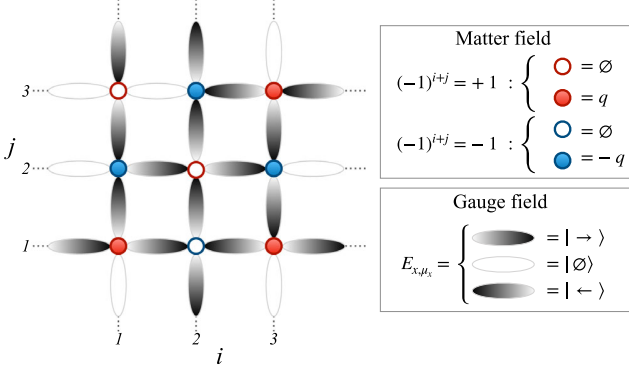
FIG. 1. Sketch of the 2D lattice gauge theory in the spin-1 representation. The cartoon of the square lattice shows a specific gauge-invariant configuration of the matter and gauge fields with zero total charge (three particles and three antiparticles). Staggered fermions represent matter and antimatter fields on a lattice bipartition: On the even (odd) bipartition, a full red (blue) site represents a particle (antiparticle). The gauge-field points in the positive direction of the color gradient.

$$\hat{H} = -t \sum_{x,\mu} (\hat{\psi}_x^\dagger \hat{U}_{x,\mu} \hat{\psi}_{x+\mu} + \mathrm{H.c.})$$

$$+ m \sum_x (-1)^x \hat{\psi}_x^\dagger \hat{\psi}_x + \frac{g_e^2}{2} \sum_{x,\mu} \hat{E}_{x,\mu}^2$$

$$- \frac{g_m^2}{2} \sum_x (\hat{U}_{x,\mu_x} \hat{U}_{x+\mu_x,\mu_y} \hat{U}_{x+\mu_y,\mu_x}^\dagger \hat{U}_{x,\mu_y}^\dagger + \mathrm{H.c.}), \quad (1)$$

where μ in $\{\mu_x, \mu_y\}$ are the unit vectors of the square lattice. The first term in Eq. (1) provides the minimal coupling between gauge and matter fields associated with the coupling strength t. It describes a process of particle-antiparticle pair creation or annihilation, where the parallel transporter operator guarantees that the local gauge symmetries are not violated. The second term in the Hamiltonian represents the energy associated with the fermionic bare mass, and it appears as a staggered chemical potential according to the Kogut-Susskind prescription. For numerical purposes, it has been redefined by adding an overall constant $mL^2/2$, thus replacing $(-1)^x \hat{\psi}_x^\dagger \hat{\psi}_x \to \delta_{x,e} \hat{\psi}_x^\dagger \hat{\psi}_x + \delta_{x,o} \hat{\psi}_x \hat{\psi}_x^\dagger$ (see Appendix C). Thus, a filled local state in the even sublattice costs positive energy m and carries charge q; otherwise, when an odd site is empty, the energy cost is still m, but it corresponds to having an antiparticle (a hole) with charge $-q$. The last two terms contribute to the gauge-field dynamics: The electric part, with coupling g_e, is completely local. The magnetic part, with coupling g_m instead, is constructed by considering the smallest Wilson loop—the product of the parallel transporter $\hat{U}_{x,\mu}$ in a closed loop—the size of a plaquette. Its name is related to the fact that it generates the magnetic contribution to the energy density in the continuum limit.

The LGT Hamiltonian $\hat{H}$ commutes with the local Gauss law generators (in units of q)

$$\hat{G}_x = \hat{\psi}_x^\dagger \hat{\psi}_x - \frac{1 - p_x}{2} - \sum_\mu \hat{E}_{x,\mu}, \quad (2)$$

where the unit lattice vector μ in the sum runs in $\{\pm\mu_x, \pm\mu_y\}$, while $p_x = (-1)^x$ is, again, the lattice site parity. In addition, the model exhibits a $U(1)$ global symmetry—namely, the conservation of the total charge $\hat{Q} = \sum_x [\hat{\psi}_x^\dagger \hat{\psi}_x - (1 - p_x/2)] = -(L^2/2) + \hat{N}$, equivalent (apart from a constant) to the number conservation $\hat{N} = \sum_x \hat{\psi}_x^\dagger \hat{\psi}_x$ of Kogut-Susskind matter fermions. As a consequence of the convention, using $\hat{E}_{x,-\mu} = -\hat{E}_{x-\mu,\mu}$, the sum of all four terms of the gauge field around the lattice site x corresponds to the outgoing electric flux, i.e., $\sum_\mu \hat{E}_{x,\mu} = E_{x,\mu_x} + E_{x,\mu_y} - E_{x-\mu_x,\mu_x} - E_{x-\mu_y,\mu_y}$. The gauge-invariant Hilbert space is thus given by all states $|\Phi\rangle$ satisfying $\hat{G}_x|\Phi\rangle = 0$ at every site x. As each electric-field degree of freedom is shared by two Gauss generators G_x, the generators themselves overlap, and projecting onto the gauge-invariant subspace becomes a nonlocal operation. Only for 1D lattice QED, or the lattice Schwinger model [19], is it possible to integrate out the gauge variables and work with the matter field only (albeit with long-range interactions) [61]. However, in two dimensions, a given (integer occupation) realization of the matter fermions does not fix a unique gauge-field configuration, thus requiring explicit treatment of the gauge fields as quantum variables. A numerically relevant complication, related to the standard Wilson formulation of lattice gauge theories, arises from the gauge-field algebra, $[\hat{E}, \hat{U}] = \hat{U}$, with $\hat{E} = \hat{E}^\dagger$ and $\hat{U}\hat{U}^\dagger = \hat{U}^\dagger\hat{U} = \mathbb{1}$, whose representations are always infinite dimensional. Simply put, if a representation contains the gauge-field state $|\alpha\rangle$, such that $\hat{E}|\alpha\rangle = \alpha|\alpha\rangle$ with $\alpha \in \mathbb{R}$, then the states $|\alpha \pm 1\rangle = \hat{U}^{\pm 1}|\alpha\rangle$ belong to the representation as well. By induction, the representation must contain all the states $|\alpha + \mathbb{N}\rangle$, which are mutually orthogonal as distinct eigenstates of $\hat{E}$; thus, the representation space dimension is at least countably infinite.

In order to make the Hamiltonian numerically tractable via tensor network methods, we need to truncate the local gauge-field space to a finite dimension. For bosonic models, this truncation is typically done by introducing an energy cutoff and eliminating states with single-body energy density beyond it, while checking *a posteriori* the introduced approximation. Similarly, for $U(1)$ lattice gauge theories, we truncate the electric field according to the quantum-link model formulation. Specifically, the gauge fields are substituted by spin operators, namely, $\hat{E}_{x,\mu} = (\hat{S}_{x,\mu}^z + \alpha)$ and $\hat{U}_{x,\mu} = \hat{S}_{x,\mu}^+/s$, such that $\hat{E}$ is still Hermitian and the commutation relation $[\hat{E}_{x,\mu}, \hat{U}_{y,\nu}] = \delta_{x,y}\delta_{\mu,\nu}\hat{U}_{x,\mu}$ is preserved [2]; however, $\hat{U}$ is no longer unitary for any *finite*

spin-s representation $|\hat{S}|^2 = s(s+1)\mathbb{1}$. The original algebra is then restored in the large spin limit $s \to \infty$, for any *background* field $\alpha \in \mathbb{R}$. Similar truncation strategies, based on group representations, can be applied to non-Abelian gauge theories as well [35,62]. In the following, we make use of the spin-1 representation ($s = 1$), under zero background field $\alpha = 0$, which captures reasonably well the low-energy physics of the theory, especially in the parameter regions wherein the ground state is characterized by small fluctuations above the bare vacuum. Here, $s = 1$ is the smallest spin representation exhibiting a nontrivial electric energy contribution. In fact, for $s = 1/2$, we have that $\hat{E}_{x,\mu}^2 \propto (\sigma_{x,\mu}^z)^2 = \mathbb{1}$ is simply a constant in the Hamiltonian; thus, g_e^2 plays no role. In 1D, evidence suggest that truncated gauge representations converge rapidly to the continuum theory, e.g., in the Schwinger model [37,63,64], reinforcing the quantitative validity of the results obtained in the simplified model. Deviations between the truncated and the full-fledged lattice theory are expected to arise when g_m^2 is the dominant coupling, as we show for a 2×2 example in Appendix I.

Let us mention that, in the formulation of the lattice QED implemented on our numerical algorithms, Eq. (1), we consider the respective couplings of the various Hamiltonian terms, namely, t, m, g_e^2, and g_m^2, as independent, dimensionless parameters. In this way we have a practical advantage in our numerical interface which allows us to treat the Hamiltonian terms on equal footage, and, in what follows, we set the energy scale as $t = 1$. However, we stress that in the original Hamiltonian formulation of lattice QED [16,21], these couplings are mutually related as $t = (1/a)$, $m = m_0$, $g_e^2 = (g^2/a)$, $g_m^2 = (8/g^2 a)$, where g is the coupling constant of QED, m_0 is the matter-field bare mass, and a is the lattice spacing of the lattice discretization. In this sense, *physical* realizations of lattice QED only depend on two actual parameters: $m' = m_0 a > 0$ and $g^2 > 0$. Nevertheless, in this work, we aim to highlight that our numerical simulations are not limited to these physical scenarios, and we keep our effective couplings independent and not bound to positive values. We leave a more detailed convergence analysis along the physical regimes to future work, along the lines of similar studies already presented for 1D systems [13].

A. Spin-1 compact representation of $U(1)$

In the spin-1 representation, the electric-field operator allows three orthogonal states for the electric flux (in units of the charge q), graphically represented in Fig. 1. For a horizontal link (x, μ_x), we write the eigenbasis of E_{x,μ_x} as

$$\hat{E}_{x,\mu_x}|\rightarrow\rangle = +|\rightarrow\rangle, \quad \hat{E}_{x,\mu_x}|\varnothing\rangle = 0, \quad \hat{E}_{x,\mu_x}|\leftarrow\rangle = -|\leftarrow\rangle,$$

on which the parallel transporter acts as $\hat{U}_{x,\mu_x}|\rightarrow\rangle = 0$, $\hat{U}_{x,\mu_x}|\varnothing\rangle = |\rightarrow\rangle$ and $\hat{U}_{x,\mu_x}|\leftarrow\rangle = |\varnothing\rangle$, and analogously

$\hat{E}_{x+\mu_x,-\mu_x}|\rightarrow\rangle = -|\rightarrow\rangle$. A similar set of states can be defined in the vertical links (x, μ_y), such that $\hat{E}_{x,\mu_y}|\uparrow\rangle = |\uparrow\rangle$, $\hat{E}_{x,\mu_y}|\varnothing\rangle = 0$, and $\hat{E}_{x,\mu_y}|\downarrow\rangle = -|\downarrow\rangle$.

In this work, we introduce an algebraic technique, similar to the rishon representations common in quantum-link models [2,52–55], which has the advantage of automatically accounting for the Gauss law, while carefully reproducing the anticommutation relations of the matter fermions without resorting to Jordan-Wigner string terms (see next section). This strategy relies on splitting the gauge-field space on each link (x, μ) into a *pair* of three hardcore fermionic modes, defined later. We say that each mode in this pair "belongs" to either of the sites sharing the link, in this case, x and $x + \mu$.

Thus, we write $\hat{U}_{x,\mu} = \hat{\eta}_{x,\mu}\hat{\eta}_{x+\mu,-\mu}^\dagger$, where the three hardcore fermionic operators $\hat{\eta}$ satisfy $\hat{\eta}_{x,\mu}^3 = 0$ (while $\hat{\eta}_{x,\mu}^2 \neq 0$) and anticommute at different positions $\{\hat{\eta}_{x,\mu}, \hat{\eta}_{y,\nu}^{(\dagger)}\} = 0$, for $x \neq y$ or $\mu \neq \nu$. Moreover, these new modes obey anticommutation relations with the matter field as well, i.e., $\{\hat{\eta}_{x,\mu}^{(\dagger)}, \hat{\psi}_y^{(\dagger)}\} = 0$. To explicitly build this three hardcore fermionic mode $\hat{\eta}_{x,\mu}$, we use two subspecies of Dirac fermions $\hat{a}_{x,\mu}$ and $\hat{b}_{x,\mu}$, such that

$$\hat{\eta}_{x,\mu}^\dagger = \hat{n}_{x,\mu}^a \hat{b}_{x,\mu}^\dagger + (1 - \hat{n}_{x,\mu}^b)\hat{a}_{x,\mu}^\dagger, \quad (\hat{\eta}_{x,\mu}^\dagger)^2 = \hat{b}_{x,\mu}^\dagger \hat{a}_{x,\mu}^\dagger, \quad (3)$$

where $\hat{n}_{x,\mu}^a = \hat{a}_{x,\mu}^\dagger \hat{a}_{x,\mu}$ and $\hat{n}_{x,\mu}^b = \hat{b}_{x,\mu}^\dagger \hat{b}_{x,\mu}$ are the occupation number operators for each subspecies. This construction provides the local algebra

$$[\hat{\eta}_{x,\mu}, \hat{\eta}_{x,\mu}^\dagger] = 1 - \hat{n}_{x,\mu}^a - \hat{n}_{x,\mu}^b \quad (4)$$

and only grants access to the three-dimensional subspaces for each three hardcore fermion mode, spanned by the following three states:

$$|0\rangle_{x,\mu}, \quad |1\rangle_{x,\mu} = \hat{a}_{x,\mu}^\dagger |0\rangle_{x,\mu}, \quad |2\rangle_{x,\mu} = \hat{b}_{x,\mu}^\dagger \hat{a}_{x,\mu}^\dagger |0\rangle_{x,\mu}, \quad (5)$$

where $|0\rangle_{x,\mu}$ is the Dirac vacuum of both subspecies, i.e., $a_{x,\mu}|0\rangle_{x,\mu} = b_{x,\mu}|0\rangle_{x,\mu} = 0$. The state $\hat{b}_{x,\mu}^\dagger |0\rangle_{x,\mu}$ is disconnected from the other three and thus projected away. Such a "half-link" local subspace is joined with a similar construction $\hat{a}_{x+\mu,-\mu}^\dagger$ and $\hat{b}_{x+\mu,-\mu}^\dagger$ on the other half of the link; thus, the pair defines the link space, and $\hat{E}_{x,\mu}$ will be diagonal in the occupation basis. While, in principle, a full link space is $3 \times 3 = 9$ dimensional, we can now exploit the following symmetry:

$$\hat{L}_{x,\mu} = \hat{n}_{x,\mu}^a + \hat{n}_{x,\mu}^b + \hat{n}_{x+\mu,-\mu}^a + \hat{n}_{x+\mu,-\mu}^b, \quad (6)$$

which counts the total number of fermions in each link and is a conserved quantity since $[\hat{L}_{x,\mu}, \hat{E}_{x,\mu}] = [\hat{L}_{x,\mu}, \hat{U}_{x,\mu}] = 0$. By working in the subspace with two fermions per link,
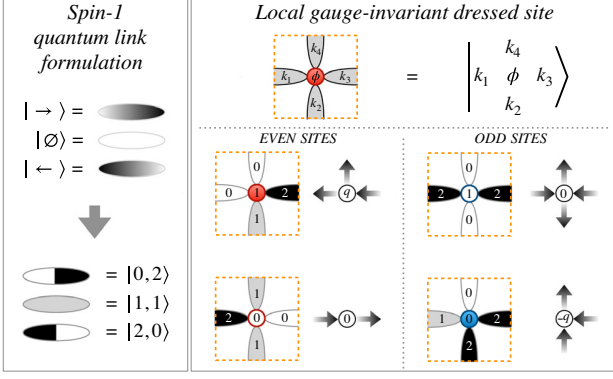
FIG. 2. Mapping of the gauge-field states in the spin-1 representation to the fermionic Fock states. Each half-link is constructed by employing two species of Dirac fermions. The link symmetry formally reduces the total number of states to only the three allowed states with two fermions. We therefore construct the local gauge-invariant dressed site by gluing each single matter site together with its neighboring half-links. In the four examples, notice how the quantum number $\phi = 1$ represents the presence (absence) of a charge (anticharge) in the even (odd) sites.

$\hat{L}_{x,\mu} = 2$, we reduce the link space to dimension three, and we can restore the desired algebra. First, we write the occupation basis (see also Fig. 2) as

$$
\begin{aligned}
|\rightarrow\rangle &= -|0,2\rangle = \hat{a}^\dagger_{x+\mu,-\mu}\hat{b}^\dagger_{x+\mu,-\mu}|0\rangle_{x,\mu}|0\rangle_{x+\mu,-\mu}, \\
|\varnothing\rangle &= |1,1\rangle = \hat{a}^\dagger_{x,\mu}\hat{a}^\dagger_{x+\mu,-\mu}|0\rangle_{x,\mu}|0\rangle_{x+\mu,-\mu}, \\
|\leftarrow\rangle &= |2,0\rangle = \hat{b}^\dagger_{x,\mu}\hat{a}^\dagger_{x,\mu}|0\rangle_{x,\mu}|0\rangle_{x+\mu,-\mu},
\end{aligned}
\tag{7}
$$

so that $\hat{U}_{x,\mu}$ acts correctly, i.e., where the minus sign in the first equation ensures $\hat{U}_{x,\mu}|\varnothing\rangle = |\rightarrow\rangle$ and $\hat{U}^\dagger_{x,\mu}|\rightarrow\rangle = |\varnothing\rangle$. Then, we express the electric-field operator as the unbalance of fermions between the two halves of a link, precisely

$$
\hat{E}_{x,\mu} = \frac{1}{2}(\hat{n}^a_{x+\mu,-\mu} + \hat{n}^b_{x+\mu,-\mu} - \hat{n}^a_{x,\mu} - \hat{n}^b_{x,\mu}), \tag{8}
$$

implementing the correct action of $E_{x,\mu}$. It is worth mentioning that this formulation can be extended to higher spin-j representations, where each link becomes a pair of $(2j + 1)$ hardcore fermions.

B. Local gauge-invariant dressed sites

One of the common issues while working with a lattice gauge theory, even in the compact representation of the electric field, is to properly identify the gauge-invariant Hilbert space. Because of the overlapping of the Gauss law generators $\hat{G}_x$, the identification of the correct local basis is highly nontrivial, especially for dimensions higher than one [26].

Using the three-hardcore-fermion pairs language gives us a shortcut to this issue. In fact, we are able to recast the Gauss law generators as nonoverlapping operators, at the price of enforcing the link constraint $(\hat{L}_{x,\mu} - 2)|\Phi\rangle = 0$. Using this constraint, we can rewrite the electric-field operator in Eq. (8), taking only the fermionic operators into account, which act on the half-link connected to x, i.e.,

$$
\hat{E}_{x,\mu} = 1 - \hat{n}^a_{x,\mu} - \hat{n}^b_{x,\mu} = [\hat{\eta}_{x,\mu}, \hat{\eta}^\dagger_{x,\mu}], \tag{9}
$$

which is valid in the link-symmetry invariant space $(\hat{L}_{x,\mu} - 2)|\Phi\rangle = 0$. As a consequence, in the Hilbert space with two Dirac fermions per link, the Gauss law generators become strictly local, i.e., containing quantum variables belonging solely to site x; they read

$$
\hat{G}_x = \hat{\psi}^\dagger_x\hat{\psi}_x - \frac{1-p_x}{2} - \sum_\mu(1 - \hat{n}^a_{x,\mu} - \hat{n}^b_{x,\mu}). \tag{10}
$$

Within this picture, it is easy to identify a local gauge-invariant basis for the *dressed* site

$$
\left|\begin{matrix} & k_4 & \\ k_1 & \phi & k_3 \\ & k_2 & \end{matrix}\right\rangle_x = (-1)^{\delta_{k_1,2}+\delta_{k_2,2}}
$$
$$
\times |\phi\rangle_x|k_1\rangle_{x,-\mu_x}|k_2\rangle_{x,-\mu_y}|k_3\rangle_{x,\mu_x}|k_4\rangle_{x,\mu_y}, \tag{11}
$$

where $|\phi\rangle_x = (\hat{\psi}^\dagger_x)^\phi|0\rangle$ for $\phi \in \{0,1\}$ describes the matter content, while $k_j \in \{0,1,2\}$ selects a state, from those in Eq. (5), for each respective half-link. The factor $(-1)^{\delta_{k_1,2}+\delta_{k_2,2}}$ accommodates the sign in Eq. (7). In this language, the Gauss law, cast as Eq. (10), simplifies to

$$
\phi + \sum_{j=1}^4 k_j = 4 + \frac{1-p_x}{2}, \tag{12}
$$

which fixes the total number of fermions in each dressed site, specifically four in the even sites and five in the odd sites. Equation (12) actually reduces the Hilbert space dimension of each dressed site from 162 to 35, and we use these 35 states as a computational basis for tensor network algorithms.

A fundamental feature of this language is that, since the total number of fermions at each dressed site is conserved, their parity is conserved as well; thus, the gauge-invariant model will not exhibit any Jordan-Wigner strings (outside the dressed sites) in the computational basis. An operative way to show this property is to consider that the Hamiltonian term $\hat{\psi}^\dagger_x\hat{U}_{x,\mu}\hat{\psi}_{x+\mu}$ decomposes as the product of $\hat{\psi}^\dagger_x\hat{\eta}_{x,\mu}$ and $\hat{\eta}^\dagger_{x+\mu,-\mu}\hat{\psi}_{x+\mu}$: Each of these two factors is local (acts on a single dressed site) and commutes with the

algebra of other dressed sites. The same applies to the magnetic plaquette term. In conclusion, by working on the dressed-site computational basis, we can employ standard (spin-model-like) tensor network techniques, without the requirement of keeping track of fermionic parity at each site [65–71]. Notice that this construction can also be exploited to perform quantum computations of two-dimensional LGTs.

C. Tensor network for 2D lattice gauge simulations

In order to numerically simulate the quantum system, we use a two-dimensional TTN state to represent the many-body wave function [51,72,73]. We work in the computational 35-dimensional local basis for each dressed site, defined in the previous section, which automatically encodes the Gauss law. Operators appearing in the Hamiltonian (1) can be cast in this basis, either as local operators or by acting on a pair or plaquette of neighboring dressed sites (see Appendix C for the explicit construction). The extra link symmetry $\hat{L}_{x,\mu} = 2$ must be enforced at every pair of neighboring sites. We do so by introducing an energy penalty for all states violating the link constraints. This penalty term is included in the optimization by a driven penalty method—similar to an augmented Lagrangian method—which is described in more detail in Appendix E. Under all other aspects, the TTN algorithm employed here for finding the many-body ground state follows the prescriptions of Ref. [74].

In the numerical simulations, we fix the energy scale by setting the coupling strength $t = -1$. Furthermore, we work within a sector with a fixed total charge $\hat{Q}$, by using standard techniques for global symmetry conservation in TNs [74–76]. We thus characterize the ground-state properties as a function of the mass m, the electric coupling g_e, and the magnetic coupling g_m.

In order to exploit the best performances of our TTN algorithm, we run simulations on square lattices $L \times L$, with the linear length L being a binary power; in particular, we consider $L = 4$, 8, and 16, and vary the TN auxiliary dimension (or bond dimension) up to $\chi \sim 300$. Depending on L and the physical parameters, we obtain a convergence precision between about 10^{-2} and 10^{-5}, sufficient to characterize the ground-state properties.

III. ZERO CHARGE DENSITY SECTOR

In this section, we focus on the zero charge density sector $\rho \equiv \langle \hat{Q} \rangle / L^2 = 0$, where there is a balance between matter and antimatter, and we analyze the ground state of Hamiltonian (1) within this subspace. Unless otherwise stated, we consider periodic boundary conditions. We characterize the ground state of the Hamiltonian by looking at the energy density $\langle \hat{H} \rangle / L^2$ and the particle density $\langle \hat{n} \rangle = (1/L^2) \sum_x \langle \hat{n}_x \rangle$, where $\hat{n}_x = (\delta_{x,e} \hat{\psi}_x^\dagger \hat{\psi}_x + \delta_{x,o} \hat{\psi}_x \hat{\psi}_x^\dagger)$

counts how many charges are in the system, both positive and negative, i.e., fermions in even sites plus holes in odd sites. We start our analysis by first focusing on the case in which the magnetic coupling has been set to zero, $g_m = 0$. Before detailing the numerical results, some analytically solvable limit cases should be considered. For large positive values of the bare mass $m \gg t$, the fluctuations above the bare vacuum are highly suppressed; the system exhibits a unique behavior since there is no competition between the matter term and the electric-field term in the Hamiltonian. Indeed, to construct particle-antiparticle pairs, the matter energy and the electric-field energy both contribute to an overall increase of the ground-state energy. In order to explore more interesting phenomena, we allow the mass coupling to reach negative values. By doing so, we can identify two different regions, depending on the competition between the electric coupling $g_e^2/2$ and the values of the mass $m < 0$:

(i) For $g_e^2/2 \gg 2|m|$, we still have a vacuumlike behavior, where we expect a unique nondegenerate ground state with small particle-density fluctuations. This regime exists, no matter the value of the mass, as long as the energy cost to turn on a nonvanishing electric field on a single link overcomes the gain in creating the associated particle-antiparticle pairs. Indeed, for any value of the mass and $g_e^2/2 \to \infty$, or for $g_e^2/2 \neq 0$ and $m \to \infty$, the presence of a finite electric field, or finite particle density, is strictly forbidden, and the ground state flows toward the only admissible configuration, namely, the bare vacuum.

(ii) For $-2m \gg g_e^2/2 > 0$, the system is characterized by slightly deformed particle-antiparticle dimers; this regime of course only exists for negative values of the mass and represents the region wherein the energy gain for creating a particle-antiparticle couple largely overcomes the associated electric-field energy cost. Here, the ground state remains highly degenerate as long as the kinetic energy coupling $|t|$ is much smaller than all the other energy scales (with the degeneracy being lifted only at the fourth order in t). In particular, for $g_e^2/2 \neq 0$ and $m \to -\infty$, the ground state reduces to a completely filled state. In order to minimize the electric-field energy, particles and antiparticles are arranged in $L^2/2$ pairs (where we are assuming L even), sharing a single electric flux in between. All of these configurations are energetically equivalent, and their degeneracy corresponds to the number of ways in which a finite quadratic lattice (with open or periodic boundary conditions) can be fully covered with given numbers of "horizontal" and "vertical" dimers. This number scales exponentially with the system size as $\exp(L^2 C/\pi)$ for $L \to \infty$, with $C \simeq 0.915966$, the Catalan's constant [77]. For the sake of clarity, we stress that such dimers are not entangled
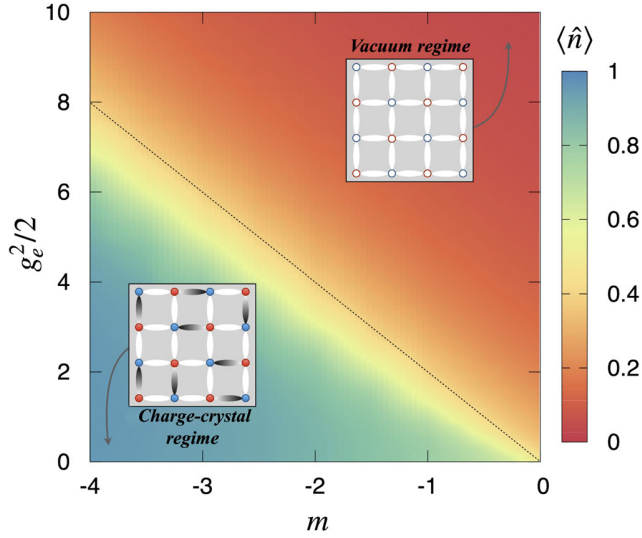
FIG. 3. Color density plot for $m < 0$ obtained from the evaluation of the density of matter in the TTN ground state for an 8×8 lattice system with periodic boundary conditions. The insets are schematic representations of the ground state deep in the two regimes: the bare vacuum for $g_e^2/2 \gg 2|m|$ and a typical dimer configuration for $g_e^2/2 \ll -2m$. The dashed line is located at the classical ($t = 0$) transition $g_e^2/2 = -2m$.

clusters of matter and gauge fields; they are roughly product states.

Let us mention that the case $g_e = 0$ with $m \to \infty$ ($m \to -\infty$) is more pathological since any gauge-field configuration compatible with the vacuum (dimerized)

state is admissible, provided the Gauss law is fulfilled. In practice, we may draw a generic closed loop with finite electric flux on top of the vacuum state without modifying its energy; similar gauge loops may be realized on top of the dimerized state, provided it is compatible with the occupied links, without changing its energy. All of these configurations are gauge invariant by construction and increase the degeneracy of the ground-state energy sector.

Our numerical results confirm and extend this picture, as can easily be seen in the phase diagram displayed in Fig. 3, obtained from TTN simulations in an 8×8 system. The matter density is roughly zero in the vacuum regime; otherwise, it takes on a finite value whenever the system exhibits "dimerization", i.e., in the charge-crystal regime. We check that the numerical data, both the ground-state energy density and the particle density, show an asymptotic tendency toward the perturbative estimates. Interestingly, the particle density experiences an abrupt change mainly in a narrowed region around $m \simeq -g_e^2/4$, where the local slope becomes steeper as the electric coupling (and the mass) approaches zero (see left panel in Fig. 4), as roughly predicted by perturbation theory and supported by the exact results in the 2×2 case (see Appendixes H and I).

The two regimes exhibit opposite long-range ordering, identified by the order parameter $\hat{\mathcal{O}}_x = (-1)^x (2\psi_x^\dagger \psi_x - 1)$, and are mutually frustrated in the proximity of the transition. As a confirmation of this property, we can track, for instance, the *full* density-density correlation function
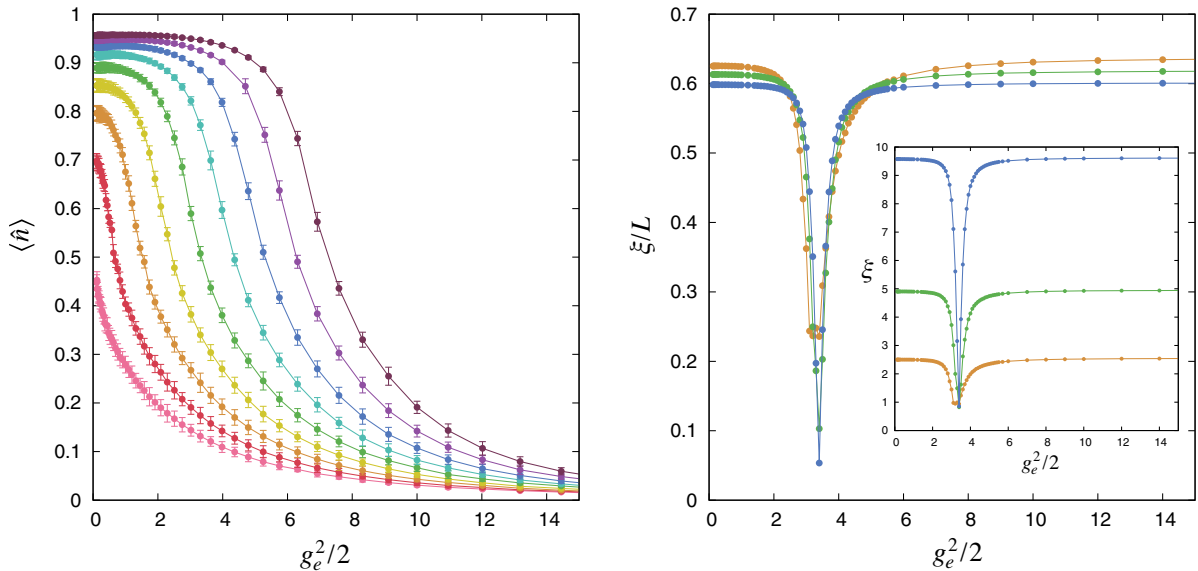

FIG. 4. Left panel: profiles of the matter density as a function of the electric-field coupling for different values of negative mass obtained by vertically cutting the color plot in Fig. 3. From bottom (pink circles) to top (purple circles), the mass takes the values $m \in \{-0.01, -0.5, -1, -1.5, -2, -2.5, -3, -3.5, -4\}$. Right panel: correlation length, in the ground state for $m = -2$ and varying the electric coupling, extracted from the density-density correlation function as explained in the main text. Different colors represent different sizes: 4×4 (orange), 8×8 (green), and 16×16 (blue).

$\mathcal{C}_{x,x'} = \langle \hat{\mathcal{O}}_x \hat{\mathcal{O}}_{x'} \rangle$ (as compared to its connected component $\mathcal{C}^0_{x,x'} = \mathcal{C}_{x,x'} - \langle \hat{\mathcal{O}}_x \rangle \langle \hat{\mathcal{O}}_{x'} \rangle$). We expect both regimes to exhibit an extensive (linear with L) full correlation length, while a sudden drop in such a quantity identifies frustration and helps us locate the transition with high precision. Such behavior is shown in the right panel of Fig. 4, where we quantify the full correlation length via the estimator $\xi^2_{\text{est}} = \sum_v D^2(v) \bar{\mathcal{C}}(v) / \sum_v \bar{\mathcal{C}}(v)$, which uses the spatially averaged correlation function $\bar{\mathcal{C}}(v) = L^{-2} \sum_x \mathcal{C}_{x,x+v}$ and the Euclidean metric $D^2(v) = v_x^2 + v_y^2$. Such an estimator effectively calculates the two-dimensional variance of $\bar{\mathcal{C}}(v)$ meant as a distribution (further discussed in Appendix F). By testing sizes up to 16×16, we observe that the actual transition point is slightly below the classical (i.e., $t = 0$) position $g_e^2/2 = -2m$. Such an outcome confirms our predictions that particle-antiparticle fluctuations, induced by a finite value of the hopping amplitude t, naturally discourage the charge-crystal order.

This effect emerges already at the second order in the perturbation theory treatment (see Appendix H), where the crossing point of the two different ground-state energies, E_v (vacuum) and E_d (dimer), slightly shifts toward the dimerized configuration.

A relevant physical question is whether the system undergoes an actual quantum phase transition across the two regions. Exactly at $t = 0$, when m crosses the critical value $-g_e^2/4$, the ground state exhibits an exact level crossing, passing from the bare vacuum to the charge-crystal energy sector. In this limit case, the system experiences a trivial first-order phase transition since the gauge-field energy term and the matter-field mass term commute between each other. However, if we tune the mass at the classical critical value $m = -g_e^2/4$, a small hopping amplitude $t \neq 0$ is already sufficient to remove such a degeneracy: Namely, the bare-vacuum energy and the charge-crystal energy get modified in a different way so that a gap opens between the two sectors. At the critical value of the mass, creation or annihilation of particle-antiparticle pairs has no energetic cost, and the ground-state energy sector is characterized by all possible states with any number of dimers; however, creating a pair in the vicinity of the bare vacuum is more favorable than annihilating a pair on top of a fully dimerized state; this is at least true for any finite L (see Appendix H for details).

A crucial insight comes from the features of the overlap between the exact ground state $|GS\rangle$ and the unique bare vacuum $|\Omega\rangle$ (see Appendix I for exact results in the 2×2 case). Indeed, for the $t = 0$ trivial case, it experiences a discontinuous transition when passing from the vacuum sector to the full dimerized sector, suddenly jumping from one to zero. Interestingly, for fixed system size L, we may evaluate such an overlap in the approximate ground state $|GS^{(k)}\rangle$ at a given order k in perturbation theory. The resulting perturbative expansion of the square of the overlap $|\langle \Omega | GS^{(k)} \rangle|^2$ changes continuously in the vacuum regime, while it remains identically vanishing, i.e., $|\langle \Omega | GS^{(k)} \rangle|^2 = 0$, when correcting the fully dimerized state up to $k < L^2/2$. We thus expect the exact overlap to be continuous at the transition point and identically vanish in the thermodynamic limit $L \to \infty$. As a consequence, for any finite t, no first-order phase transition occurs, and we may have a second-order phase transition. Let us stress that, although the perturbative expansion of the fully dimerized state does not produce any change in the overlap with the bare vacuum for $k < L^2/2$, local observables do experience perturbative modifications, simply because the state by itself gets modified. In particular, as a consequence of the Hellmann-Feynman theorem, the particle density as a function of the mass coupling m coincides with the derivative of the ground-state energy density, $\langle GS(m) | \hat{n} | GS(m) \rangle = \partial_m E_{GS}(m)/L^2$. A second-order phase transition will thus imply continuous profiles of the particle density, with a discontinuous or diverging derivative at the transition point. In fact, we have numerical evidence that the matter density changes continuously when going from one phase to the other (see Fig. 4); however, it remains very hard to infer its derivative at the transition point.

A. Finite magnetic-coupling effects

We now analyze the case of nonvanishing magnetic coupling g_m, especially focusing on how it impacts the many-body quantum features at zero temperature.

In Fig. 5, we show the field-plot representations of the ground-state typical configurations for an 8×8 system in the presence of magnetic couplings. For the sake of visibility, we only plot a 4×4 subsystem out of the complete 8×8 lattice simulated with periodic boundary conditions. Both mass m and electric coupling g_e have been chosen so that the system is deep within the two different regimes (left panels represent the vacuum regime; right panels represent the charge-crystal regime). As the magnetic coupling g_m increases to commensurate values (bottom panels), we see negligible changes affecting the vacuum configuration. By contrast, in the charge-crystal regime, the nonvanishing magnetic coupling introduces a nontrivial reorganization of the electric fields.

Such an effect can be well understood in terms of perturbation theory: (i) In the vacuum region, the ground state is not degenerate, and the first nontrivial corrections are given by coupling such a state with all the states with a single flux loop over a single plaquette (whose energy is therefore $2g_e^2$). In this regime, the flux loop state has high electric-field energy; thus, it will only slightly impact the global features of the state. The first-order correction to the ground-state energy will be quadratic in the magnetic coupling, i.e., proportional to g_m^4/g_e^2 (see Fig. 6, left panel). Let us stress that, even though the state may experience an
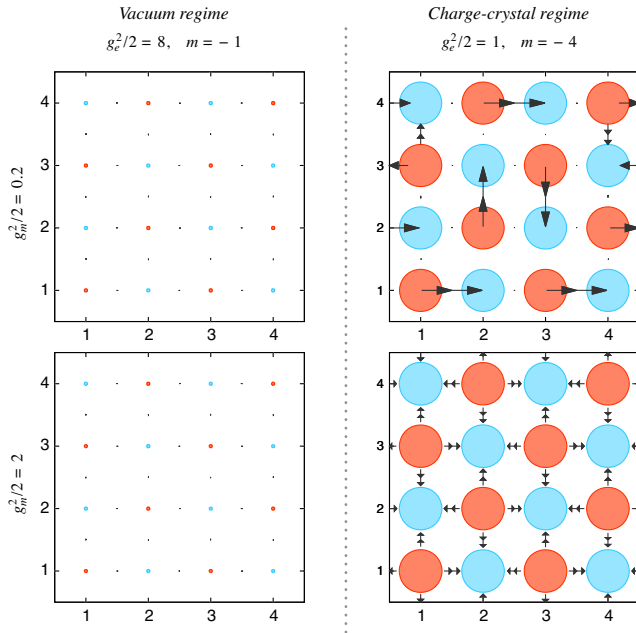
FIG. 5. Numerical results obtained via TTN simulations. The field plots reproduce the matter and gauge configurations for a 4×4 subsystem embedded in an 8×8 lattice with periodic boundary conditions. Red (blue) circles represent particles (antiparticles): Their diameter indicates the average density, from 0 (empty sites) to 1 (completely filled). The arrows in between represent the electric field: Larger arrows indicate greater electric flux.

electric-field reconfiguration due to the "field-loop" superpositions, the fact that in this regime the electric field is almost zero causes no visible effect on its expectation value; this is pretty clear from the left column of Fig. 5:

When passing from a small ($g_m^2/2 = 0.2$) to a slightly bigger value ($g_m^2/2 = 2$) of the magnetic coupling, the changes in the expectation value of the gauge field are negligible. (ii) In the charge-crystal configuration, the effect of the magnetic interaction is nontrivial. Indeed, the ground-state energy sector in this regime is highly degenerate, and the magnetic field contributes to lift such a degeneracy (at much lower order than t). The magnetic coupling introduces first-order transitions between different gauge-field configurations; therefore, its first contribution to the ground-state energy is scales like order g_m^2/g_e^2. Actually, a sufficiently large value of the magnetic coupling g_m helps the TTN wave function to restore the square lattice symmetry by introducing a gap between the actual ground state and all the energetically unfavorable configurations. This property is noticeable in Fig. 5 (bottom-right panel), where for $g_m^2/2 = 2$, the gauge-field distribution becomes uniform (on average) in the bulk. In this scenario, the charge crystal does not encourage the formation of dimers but instead a global entangled state of gauge fields.

The previous considerations are supported by the behavior of the ground-state energy and of the particle density, as a function of g_m. In Fig. 6, we plot the numerical results obtained via TTN simulations in an 8×8 system. We fix the value of the mass to $m = -2$ and explore the behavior for $g_e^2/2 = 2$ and 6, which are slightly below or above the classical transition point $g_e^2/2 = -2m$. We vary the magnetic coupling in a rather big interval $g_m^2/2 \in [0, 8]$. In the first case, i.e., when the system is initially in the charge-crystal configuration, we expect linear corrections to the ground-state energy as a function of g_m^2; this is pretty clear from Fig. 6, where, however, some deviations are visible

FIG. 6. Numerical results for the ground-state energy $\langle \hat{H} \rangle / L^2$ density and the particle density $\langle \hat{n} \rangle$ as a function of the magnetic coupling for an 8×8 system. The mass coupling has been fixed to $m = -2$. The energies have been renormalized to their absolute value at $g_m = 0$; dashed lines are guides for the eye. Data have been obtained by extrapolating from TTN simulations with different auxiliary dimensions.

because of the vicinity of the phase boundary. In the second case, i.e., starting from a quasivacuum, a quadratic deviation of the ground state is clearly visible (see Fig. 6).

Interestingly, in the parameter region we are exploring, the magnetic coupling enhances the production of particles, thus increasing the average matter density, even though the magnetic term does not directly couple to matter. Such an emergent behavior is physically relevant since it also arises when performing phase diagram simulations along a physical line of the QED problem. Specifically, setting $g_e g_m = 8t^2$, we realize the physical scenario of QED. Figure 11 in Appendix A shows a growing charge density at smaller QED couplings g, even when the (negative) bare mass is small.

In practice, the magnetic coupling creates resonating configurations of the gauge fields in the crystal charge regime, thus decreasing the electromagnetic energy density of the state itself, which in turn favors the crystal charge configuration in the proximity of the phase boundary. Hence, small g_m values effectively enlarge the charge-crystal regime. However, in the spin-1 representation of the gauge field, the dimerized configuration is not stable under an arbitrary large value of the magnetic coupling, and we expect $\langle \hat{n}_x \rangle = 0$ when $g_m \gg g_e$. This case can be easily understood at the classical level ($t = 0$), comparing the effect of a Wilson loop operator in the zero-matter (vacuum) sector and in the full-matter sector: In the former case, each single plaquette resonates between three different diagonal gauge-field configurations $\{|\circlearrowleft\rangle, |\varnothing\rangle, |\circlearrowright\rangle\}$; in the last case, only two configurations are resonating, e.g., $\{|\uparrow\downarrow\rangle, |\leftrightarrows\rangle\}$, since constructing a clockwise (anticlockwise) electric loop $\circlearrowright$ ($\circlearrowleft$) on top of the first (second) state is forbidden by the spin-1 finite representation. This leads to a contribution of the magnetic coupling to the energy, which is proportional to $-\sqrt{2}g_m^2$ for a plaquette in the vacuum, while it is only $-g_m^2$ for a dimerized plaquette. In practice, such a difference remains at the many-body level as well, and, therefore, although for $-2m \gg g_e^2/2$ the dimerized configuration represents the lower energy state at $g_m = 0$, it will become energetically unfavorable for sufficiently strong magnetic couplings.

We have further analytical confirmation of such behavior from the exact diagonalization of 2×2 systems (see Appendix I for details): For a single plaquette system, the first visible effect of a nonvanishing magnetic coupling is to mix up the two dimerized states into two different superpositions with different energies. The transition between the vacuum state toward the lower energetic charge-crystal state is therefore sharpened, and its position is shifted as well in $g_e^2/4 + m \simeq (g_e^2 + g_m^2/2 - \sqrt{g_e^4 + g_m^4/2})/4$. Interestingly, depending on the values of g_e, this shifting is not monotonous in g_m, producing an initial increase in the particle density followed by a definitive decrease toward zero (cf. Fig. 19) and thus confirming the previous heuristic argument based on perturbation theory. Again, this shifting is a strictly finite-spin representation effect, and it disappears

as the spin gets larger, as shown by analyzing the behavior for the single plaquette in the spin-2 compact representation of the gauge field (see Appendix I).

IV. FINITE CHARGE DENSITY SECTOR

One of the most intriguing phenomena we observed in our numerical simulations relies on the possibility to create a charge imbalance in the system. This scenario is challenging for Monte Carlo techniques as it produces the sign problem [13,24]. Instead, our gauge-invariant tensor network approach is very well suited to overcome such difficulty: The fact that the global $U(1)$ symmetry has been explicitly embedded in the tensor network ansatz [74] allows us to work exactly within each sector with fixed total charge. In the following, we only consider $g_m = 0$. Moreover, in this setup, because of the finite net electric flux coming out of the entire system, we have to work with open boundary conditions. In this geometry, the dressed sites at the boundary are now characterized by one outgoing half-link (two in the corners), which can support the electric field to allow the existence of a nonvanishing total outgoing flux.

When a finite density of charge $\rho \equiv \langle \hat{Q} \rangle / L^2 \in \{-1/2, 1/2\}$ is injected into the system, we expect a different behavior, depending on which part of the phase diagram the ground state belongs in. Indeed, when the ground state is very close to the bare vacuum, any charge created on top of it is forced to reach the boundaries so as to minimize the total energy; this is easily understood already with the classical ($t = 0$) Hamiltonian, and there are no fluctuations of the gauge fields. In this case, a classical configuration with a single charge located at distance ℓ from the boundary costs at least $\ell g^2/2$ more than the optimal configuration where the same charge is located at the surface (see Fig. 7). In this regime, the diagonal energy term gets modified as $E_v/L^2 = (g_e^2/4 + m)\rho$, as long as $\langle \hat{Q} \rangle \leq 2(L-1)$, i.e., whenever the total excess of charge is lower than the number of allowed free sites at the boundaries. When the total charge increases, deeper sites start to be filled; e.g., for $2(L-1) < \langle \hat{Q} \rangle \leq 4(L-2)$, one starts filling the next-neighboring sites to the surface (e.g., Fig. 8). Overall, this argument supports the existence of a phase separation between a boundary region attached to the surface, or strip, where charges aggregate and a bulk region expelling charges and electric fields. In the picture where the gauge field is not truncated ($S \to \infty$), both regions will scale as a surface. In practice, defining $\rho_\ell \equiv 2\ell(L-\ell)/L^2$ as the maximum amount of charge density the system can store within a strip of extension ℓ from the surface, we have a sharp discontinuity in local charge densities, at the smallest ℓ^* such that $\rho \leq \rho_{\ell^*}$, between a finite-charge region (for $j < \ell^*$) and a zero-charge region (for $j > \ell^*$). In particular, in the thermodynamic limit, we obtain $\ell^*/L = (1 - \sqrt{1 - 2|\rho|})/2$. In other words, the width ℓ^* of the surface strip where all charges are localized varies
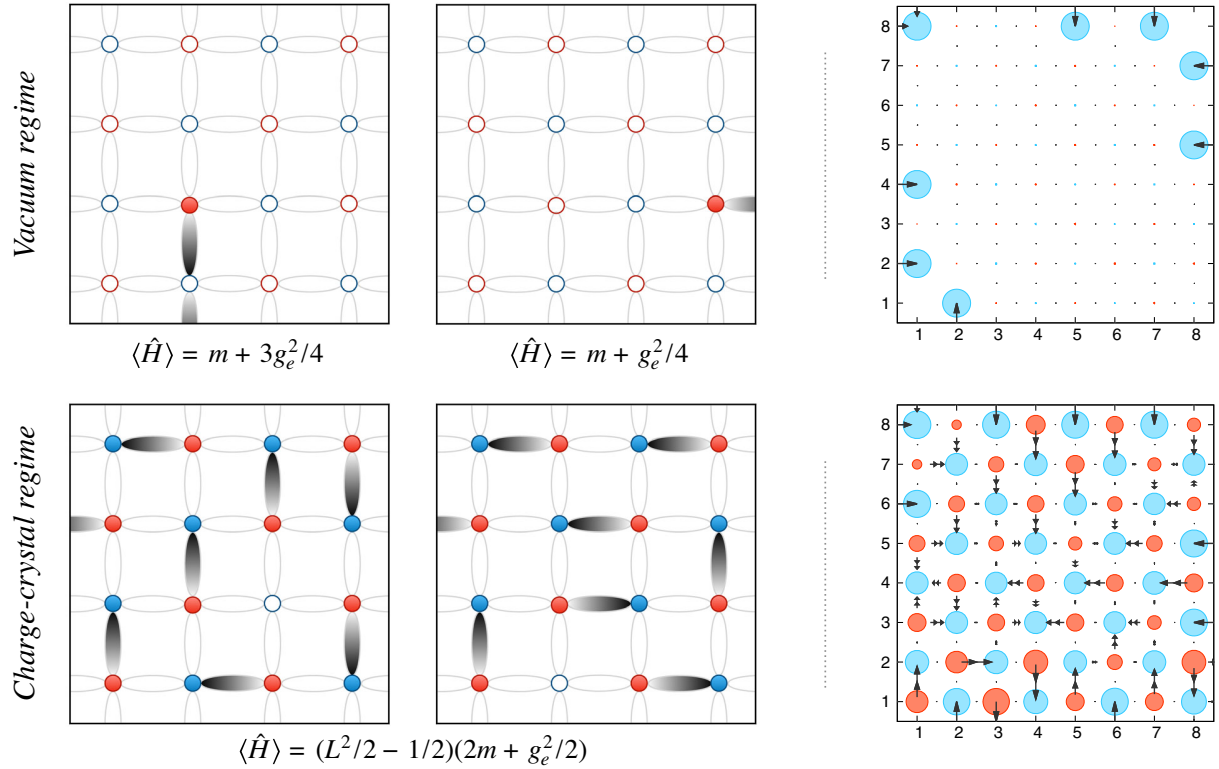
FIG. 7. Field plots in the finite charge density sectors. The top row refers to the vacuum regime and the bottom row to the charge-crystal regime. On the left of the figure, the four panels represent a sketch of the classical configurations (i.e., $t = 0$) for a 4×4 system with open boundary conditions in the $Q = 1$ charge sector. Gauge fields can now exit the system at the energy cost of a half-link each the system at the cost of the half-link. The panels on the right of the figure are the field plots obtained by numerical TTN simulations in 8×8 systems for the two different regimes, namely, $g_e^2/2 = 2$ and $m = 4$ (top) or $m = -2$ (bottom), and charge sector $Q = -8$. In the vacuum regime, the excess of charge prefers to be localized at the boundaries since such configurations are more energetically favorable. In the dimerized regime, the holes may occupy any position since the system can reconfigure the pairs of dimers so as to always use the same amount of energy. However, because of the very high degeneracy of the low-energy sector, the TTN simulations may get stuck in a slightly asymmetric configuration.

smoothly in $[0, L/2]$ as $|\rho|$ varies in $[0, 1/2]$. Quantitatively, both the depth of the surface strip (ℓ^*) and the diameter of the bulk region $(L/2 - \ell^*)$ scale linearly with L; thus, the phase
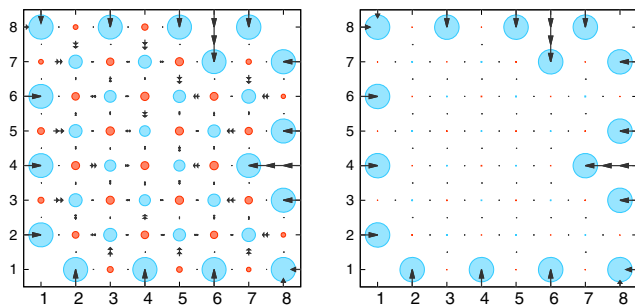

FIG. 8. Field plots from the TTN numerical simulations of 8×8 systems in the $Q = -16$ charge sector. The couplings are tuned in such a way that the system is deep in the vacuum regime (right panel, $g_e^2/2 = 2, m = 5$) or near the critical region (left panel, $g_e^2/2 = 2, m = -1$). Notice how the excess of charge is larger than the allowed antimatter sites at the surface: The system has to allocate two extra charges in the next-neighboring sites to the surface.

separation argument applies when approaching the thermodynamic limit, as long as the gauge field is unconstrained and the lattice spacing stays finite. In practice, as long as the average charge density is finite, we always have an extensive region in the bulk of the system whose linear dimension scales as $L\sqrt{1 - 2|\rho|}/2$, which exhibits no charges. However, we stress that when introducing a fixed truncation of the gauge field (to any spin S), the amount of total charge that can be injected in the system is limited to a linear scaling in L since, due to the Gauss law, the total electric flux at the boundary must match the total charge. Therefore, in order to approach the thermodynamical limit at finite charge density, one needs to increase the truncation or introduce static background fields.

We expect this picture to be slightly modified at finite hopping coupling $|t|$ but to remain valid as long as the system belongs to the vacuum regime. In practice, a finite tunneling amplitude introduces a small homogeneous particle density, thus slightly increasing ℓ^* and also building up a finite charge penetration length scaling linearly with $|t|$ such that the phase separation is
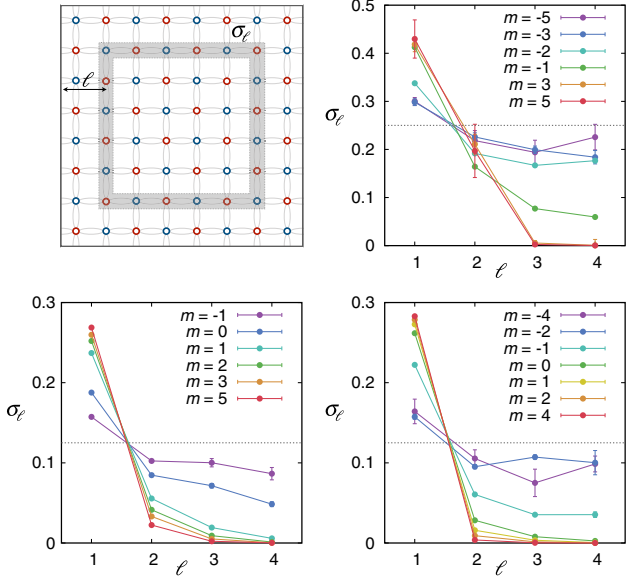
FIG. 9. Surface charge density evaluated in an 8×8 system as sketched in the top-left image. The shaded region represents the domain $\mathcal{D}_\ell$ defined in the main text. TTN simulations have been performed for different charge sectors and electric couplings; in clockwise order, $(Q = -16, g_e^2/2 = 2)$, $(Q = -8, g_e^2/2 = 2)$, and $(Q = -8, g_e^2/2 = 1/2)$.

smoothened out with an exponentially small density-charge tail penetrating into the bulk. The overall scenario is confirmed by the field plots in Figs. 7 and 8, where it is pretty clear that when the couplings are tuned in order for the system to be deep in the vacuum regime, the excess charges stick to the boundary so as to minimize the length of the attached electric strings. In principle, all possible configurations with all charges at the boundaries are energetically equivalent. However, the TTN many-body wave function spontaneously breaks such symmetry and picks up a single, specific

configuration, as is usually the case with DMRG-like algorithms. We stress that such phase separation, where both bulk and boundary regions scale extensively, is likely an artifact of the lattice discretization, where the amount of local charge density is bound.

When the state belongs to the charge-crystal regime, a finite positive (negative) charge density is mainly generated by creating holes in the odd (even) sublattice; namely, negative (positive) charges are removed from the fully dimerized state. In order to minimize the energy, the holes can now be fully delocalized: A hole in the bulk, or at the boundary, generates a reconfiguration of the charge-crystal state in such a way that it always requires the same amount of energy and guarantees the expected total outgoing electric flux. In this regime, the zero-order energy term gets modified as $E_d/L^2 = (g_e^2/4 + m)(1 - \rho)$. The entire system is now characterized by a unique spatial phase where we expect a uniform average charge density and finite electric field in the bulk. Let us mention that, for any finite value of the hopping amplitude, we still expect a similar behavior, where the transition toward the phase-separated phase will be driven by the competition between the mass and the electric coupling.

In order to highlight the different features of the low-energy state at finite chemical potential, we analyze the behavior of the surface charge density,

$$\sigma_\ell \equiv \frac{1}{\dim \mathcal{D}_\ell} \sum_{x \in \mathcal{D}_\ell} \langle \hat{\psi}_x^\dagger \hat{\psi}_x \rangle, \quad (13)$$

where $\mathcal{D}_\ell$ is a square that counts $\dim \mathcal{D}_\ell = 4(L + 1 - 2\ell)$ lattice sites as sketched in Fig. 9. Here, $\ell \in \{1, 2, ..., L/2\}$ represents the distance of the domain $\mathcal{D}_\ell$ from the external surface: Namely, as ℓ grows, we select domains deeper into the bulk.

In Fig. 9, we plot the surface charge density σ_ℓ as a function of ℓ for different points in the coupling-parameter space. As long as the Hamiltonian is tuned into the vacuum
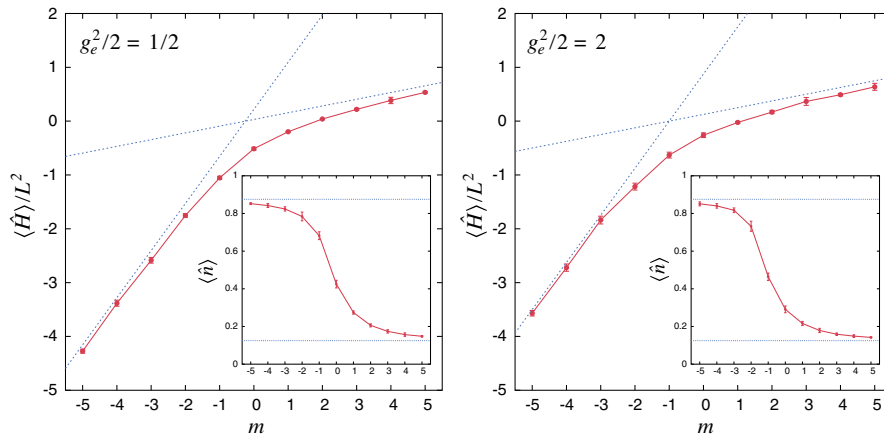

FIG. 10. Ground-state energy density and particle density (insets) as a function of the bare mass m for an 8×8 lattice with $\rho = -1/8$, i.e., in the $Q = -8$ charge sector. In the left panel, the transition between the two regimes occurs at $m \sim -1/4$, while in the right panel, it is located at $m \sim -1$. Dashed lines are the asymptotic values in the zero-order perturbative approximation.

regime, the surface charge suddenly drops when getting into the bulk of the system. As expected, for finite value of the couplings, when approaching the critical region, the bulk charge density is enhanced; finally, once the system reaches the charge-crystal regime, σ_ℓ acquires a loosely uniform shape.

Finally, we carefully check the ground-state energy density and the particle density, which are plotted in Fig. 10 as a function of the mass for two different values of the electric coupling. Notice how, for sufficiently large positive (negative) values of the mass, the data get closer to the perturbative predictions. The intermediate region, at $m \sim g_e^2/4$, is characterized by stronger quantum fluctuations and thus exhibit a smooth transition between uniform and nonuniform charge distribution in space.

As a concluding remark, we stress that, while in this section we consider the boundary conditions to be completely free, setting a specific set of boundary conditions for the problem of electrodynamics is not conceptually or numerically difficult. Typical boundary conditions are realized by means of a boundary Hamiltonian H_b to be added to the bulk Hamiltonian H from Eq. (1) (see Appendix G for details).

V. CONCLUSIONS

In this work, we demonstrated a novel, efficient, tensor network approach to the study of two-dimensional lattice gauge theories. By exploiting the quantum-link formulation of LGT, the fermionic rishon representation of quantum links, and unconstrained tree tensor networks, we investigated the equilibrium properties of a two-dimensional lattice QED within its first compact spin representation. We present results for lattice size up to 16×16, whose Hilbert space dimension is approximately equivalent to that of a system composed of spins-$1/2$ on a square lattice with edges of about 80 lattice sites. Whenever possible, we confirmed our results with perturbative analysis and small-scale exact simulations.

In particular, we identified different regimes at zero chemical potential, a vacuum state and a charge-density one, that reproduce what has been found in the one-dimensional case, and we investigated the effects of a magnetic term uniquely present in two dimensions. Finally, we explored the finite density scenario and individuated two distinct behaviors corresponding to the vacuum and charge-density configurations: In the former case, the excess charges accumulate on the boundaries. This configuration minimizes the electric energy density, and is analogous to how charges distribute in classical conductors. In the latter, the excess charge is distributed uniformly in the bulk and boundaries.

In conclusion, we have shown that unconstrained tree tensor networks are a powerful tool to obtain a non-perturbative description of a lattice gauge theory in two dimensions. We stress that these simulations have been obtained on standard clusters without exploiting heavy parallelization and with simulations lasting only a few days. Despite the fact that the presented results are not yet able to allow physical predictions in the continuous limit for the system we study, we foresee that upgrading the current software to exploit the full power of high performance computing—without major changes in the algorithms—larger system sizes, an additional dimension, the continuum and large-S limit, and more complex Abelian and non-Abelian lattice gauge theories will be in the range of the approach presented here, as already shown for the one-dimensional case [31,61]. Our proposed architecture is perfectly tailored to accommodate advanced strategies of diagnostics, including elaborate string order parameters capable of detecting deconfined phases and topological order [78–80].

APPENDIX A: PHYSICAL QED SCENARIO

Complementing the discussions in Sec. III A, we present two physical lines of phase diagram simulations of the QED problem with $g_e g_m = 8t^2$. Figure 11 shows a growing charge density at smaller QED couplings g, even when the (negative) bare mass is small.
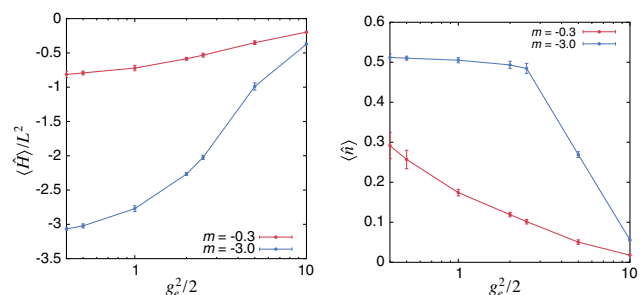

FIG. 11. Numerical results for the ground-state energy density $\langle \hat{H} \rangle / L^2$ and the particle density $\langle \hat{n} \rangle$ of the original QED Hamiltonian formulation as a function of the electric coupling g_e for 8×8 systems. The lattice spacing is set to $a = 1/t = 1$, and the magnetic coupling is tuned with respect to $g_m^2 = 8/(g_e^2 a^2)$. The mass coupling has been set to $m_0 = \{-0.3, -3.0\}$, respectively. The data have been obtained by extrapolating from TTN simulations with different auxiliary dimensions.

APPENDIX B: CHARGE SCREENING

Here, we briefly address the problem of detecting confinement. A natural way of exhibiting confinement in our lattice scenario is to show that electric-field lines do not extend over infinite lengths in the full-fledged theory (where $g_e g_m = 2\sqrt{2}t$), even when we enforce the presence of charges at specific locations. In this sense, the ground QED solution adjusts the mobile charges (and anticharges) to screen the pinned ones. We can insert such pinned charges by tuning a local chemical potential term $\tilde{m}_x(-1)^x \hat{\psi}_x^\dagger \hat{\psi}_x$, which shifts the mass at site x to strongly negative values ($\tilde{m}_x + m \ll -1$), thus favoring the presence of a full charge at x in the ground state.

Adding a single pinned charge has the effect of creating a local excitation in the vacuumlike regime. Figure 12(a) shows a scenario with global zero charge, and a single pinned charge, under periodic boundaries. While the crystal-charge regime is mostly unaffected, in the vacuum-like regime, opposite-sign charges are attracted around the pinned one, so as to form an almost perfect meson, carrying an electric-charge quadrupole. Field lines propagating from this configuration are very short ranged, thus supporting confinement.

It is important to mention that confinement can be further corroborated by string-breaking analysis, where an initial (high-energy) configuration with a long field-line string breaks down to multiple localized mesons. Configurations with long (extensive) field lines can be engineered either by field linking a bulk charge to a point in the boundary, in the sector $Q_{\text{tot}} = 1$, or by setting two pinned charges far apart

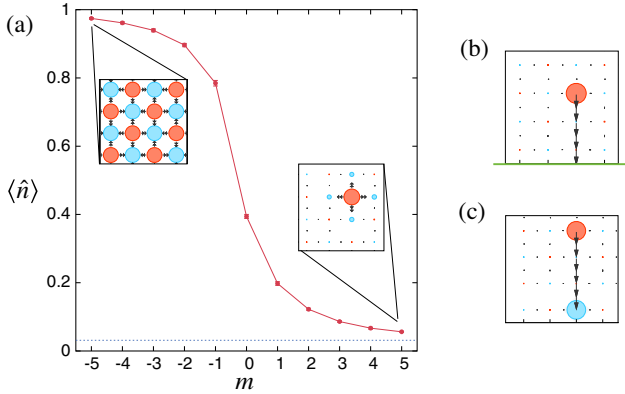FIG. 12. Occupation for pinned charges in an 8×8 system of the original QED Hamiltonian with $g_e = 1$ and $a = 1/t = 1$. (a) Particle density for one charge pinned in the zero-charge sector $Q_{\text{tot}} = 0$ with respect to the bare mass m for periodic boundaries. The system transitions from the completely filled charge-crystal phase to the pinned charge screening. (b) Field plot for a system with total charge $Q_{\text{tot}} = 1$ with open boundaries (green line). (c) Field plot for two pinned charges in the $Q_{\text{tot}} = 0$ symmetry sector. All of the field plots are 4×4 subsystems embedded in an 8×8 simulation. For panels (b) and (c), the mass term is set to $m = 4$.

in the $Q_{\text{tot}} = 0$ sector. When the field lines (strings) scale in length with the system size, we expect them to be broken by the appearance of screening charges around the pinned ones, in the thermodynamic limit. By contrast, for finite-length strings, it is possible to set the bare mass sufficiently large so that they will remain unbroken, as shown in Figs. 12(b) and 12(c).

APPENDIX C: CONSTRUCTING THE COMPUTATIONAL HAMILTONIAN

In this section, we sketch the steps needed to obtain the operator matrices, and their elements, which appear in the computational formulation of the quantum-link QED model. In particular, we stress how to construct building-block operators $A_j^{(\alpha)}$, each acting on a single dressed site j, which are genuinely local, in the sense that they commute, by construction, with every other building-block operator at another site: $[A_j^{(\alpha)}, A_{j' \neq j}^{(\alpha')}] = 0$. The electric-field term and the bare mass term are diagonal in the occupation basis of fermions and rishons, as in Eq. (11), and thus trivially obtained. The nondiagonal terms are decomposed as follows:

Matter-field coupling terms.—Matter-field terms decompose naturally as $\psi_x^\dagger U_{x,x+\mu_x} \psi_{x+\mu_x} = A_x^{(1)\dagger} A_{x+\mu_x}^{(3)}$ (and its Hermitian conjugate) for horizontal "hopping" terms, and $\psi_x^\dagger U_{x,x+\mu_y} \psi_{x+\mu_y} = A_x^{(2)\dagger} A_{x+\mu_y}^{(4)}$ for vertical hopping. The decomposition into building blocks is based upon

$$
\begin{aligned}
\psi_x^\dagger U_{x,x+\mu_x} \psi_{x+\mu_x} &= \psi_x^\dagger \eta_{x,\mu_x} \eta_{x+\mu_x,-\mu_x}^\dagger \psi_{x+\mu_x} \\
&= (\eta_{x,\mu_x}^\dagger \psi_x)^\dagger (\eta_{x+\mu_x,-\mu_x}^\dagger \psi_{x+\mu_x}) \\
&= A_x^{(1)\dagger} A_{x+\mu_x}^{(3)},
\end{aligned} \tag{C1}
$$

where $\eta_{x,\mu}$ are the three hardcore fermionic operators defined in Eq. (3). Both $A_x^{(1)}$ and $A_x^{(3)}$ are built on an even number of fermionic operators; therefore, they commute with any operator that does not act on site x and are thus genuinely local. The vertical hopping term is similarly decomposed into building-block operators.

Magnetic terms.—The magnetic (or plaquette) term decomposes into building-block operators, acting on the four dressed sites at the corners of a plaquette. Specifically, we have

$$
\begin{aligned}
U_{x,x+\mu_x} & U_{x+\mu_x,x+\mu_x+\mu_y} U_{x+\mu_y,x+\mu_x+\mu_y}^\dagger U_{x,x+\mu_y}^\dagger \\
&= \eta_{x,\mu_x} \eta_{x+\mu_x,-\mu_x}^\dagger \eta_{x+\mu_x,\mu_y} \eta_{x+\mu_x+\mu_y,-\mu_y}^\dagger \\
&\times (\eta_{x+\mu_y,\mu_x} \eta_{x+\mu_x+\mu_y,-\mu_x}^\dagger)^\dagger (\eta_{x,\mu_y} \eta_{x+\mu_y,-\mu_y}^\dagger)^\dagger \\
&= -(\eta_{x,\mu_y}^\dagger \eta_{x,\mu_x})(\eta_{x+\mu_x,-\mu_x}^\dagger \eta_{x+\mu_x,\mu_y}) \\
&\times (\eta_{x+\mu_x+\mu_y,-\mu_y}^\dagger \eta_{x+\mu_x+\mu_y,-\mu_x})(\eta_{x+\mu_y,\mu_x}^\dagger \eta_{x+\mu_y,-\mu_y}) \\
&\equiv -A_x^{(5)} A_{x+\mu_x}^{(6)} A_{x+\mu_x+\mu_y}^{(7)} A_{x+\mu_y}^{(8)},
\end{aligned} \tag{C2}
$$

to be added, in the Hamiltonian, to its Hermitian conjugate. All operators in this decomposition are local and ready to use for TTN algorithms.

APPENDIX D: TENSOR NETWORKS

In what follows, we describe the background and main principles of tensor networks and, in particular, the TTN ansatz considered in this work. For a more in-depth description of TNs, we refer to more technical reviews and textbooks [32,74,81,82].

TNs are used to efficiently represent (pure) quantum many-body wave functions $|\psi\rangle$, which live in the tensor product $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \mathcal{H}_N$ of N local Hilbert spaces $\mathcal{H}_k$, each assumed to be of finite dimension d. Expressing such a state in the real-space product basis means decomposing the wave function as

$$|\psi\rangle = \sum_{i_1,\ldots,i_L=1}^{d} c_{i_1,\ldots,i_L} |i_1\rangle_1 \otimes |i_2\rangle_2 \otimes \ldots \otimes |i_L\rangle_L, \quad \text{(D1)}$$

where $\{|i\rangle_k\}_i$ is the canonical basis of site k, spanning $\mathcal{H}_k$. Describing such a general state by all possible combinations of local states requires d^N coefficients $c_{i_1,\ldots,i_N}$. Thus, we have exponential growth with the system size N in the exact representation of the wave function. For physical states, which satisfy certain entanglement bounds under real-space bipartitions (area laws) [83,84], tensor networks offer a more efficient representation. This representation is given by decomposing the complete rank-N tensor into a set of local tensors with smaller rank, connected with auxiliary indices. We control the dimension of the auxiliary indices with the bond dimension χ and thereby the amount of captured information. Thus, tuning this parameter χ, TNs interpolate between a product state, where quantum correlations are neglected, and the exact, but inefficient representation. The most prominent TN representations are the matrix product states (MPS) for 1D systems [82,84,85] and their higher-dimension variant, the projected entangled pair states (PEPS) [83,86,87], while TTN [51,72,88,89] (as well as multi-scale entanglement renormalization Ansatz [90,91]) can, in principle, be defined in any lattice dimension.

Algorithms for MPS have been developed for over 20 years and have established the MPS ansatz as the primary workhorse for equilibrium problems in 1D [82,92] and, in many cases, even out of equilibrium [93,94]. By contrast, the development of TN algorithms, which are both quantitatively accurate and polynomially scalable, for two-dimensional lattices is still ongoing. Currently, we are still facing the open question of which tensor network geometry is generally best suited for 2D simulations. The PEPS approximates the complete rank-N tensor by a decomposition with one tensor for each physical site. These tensors are then connected through a grid analogous to the lattice, resulting in a TN with "loops" (nonlocal gauge redundancies). On the other hand, TTNs represent the wave function with a network geometry without loops, thus allowing (polynomially scaling) universal contraction schemes [74].

By its structure, the PEPS is the intuitive (and potentially more powerful) representation of a two-dimensional quantum many-body wave function satisfying the area laws of entanglement. However, in general, it lacks an exact calculation of expectation values. In fact, for a finite square lattice with $N = L \times L$ sites, the contraction of the complete PEPS to perform this calculation scales exponentially on an average system length L [95]. Additionally, the optimization of the PEPS ansatz has a higher numerical complexity $\mathcal{O}(\chi^{10})$ with the bond dimension, so the typical bond dimensions achieved are on the order of $\chi \sim 10$, which is sufficiently large for many spin systems with local dimension $d = 2$. For the 2D LGT simulations presented in this work, however, we have to deal with a local dimension of $d = 35$, which raises a nontrivial challenge for the PEPS ansatz. Furthermore, this local dimension increases for 3D systems or a higher representation for the discretization of the electric field.

The TTN, on the other hand, offers a more favorable computational scaling with bond dimension: Both exact full contraction and optimization algorithms scale with $\mathcal{O}(\chi^4)$, which in turn allows typical bond dimensions to even exceed $\chi \geq 1000$. Moreover, a TTN is fairly straightforward to implement and not restricted to any dimensionality of the underlying system; thus, the extension to 3D systems is straightforward. On the other hand, TTNs have been shown to poorly embed the area laws in two or higher dimensions [96]. Eventually, when increasing the system size, $N = L \times L$, the TTN may fail to accurately describe the quantum wave function. Thus, even though the TTN is a powerful tool to tackle systems in one, two, and three dimensions, further development and improvement is needed for reaching a scalable algorithm for higher system sizes. However, since this is a variational ansatz with increasing precision for increasing bond dimension, we can always give an estimate of the total error of our simulation results.

Our TTN algorithm implemented for finding the many-body ground state in this LGT analysis follows the prescriptions of Ref. [74]. In the numerical implementation, we exploit the $U(1)$ symmetry corresponding to the conservation of total charge Q using common techniques for global symmetry conservation in TNs [76]. We construct the tree starting from the physical indices at the bottom by iteratively merging two local sites into one by a randomly initialized tensor coarse-grained site. In case we reach the maximum bond dimension for the coarse-grained space, we truncate the coupling symmetry sectors randomly in order to keep the bond dimension. Thereby, we randomly initialize not only the tensors themselves but the
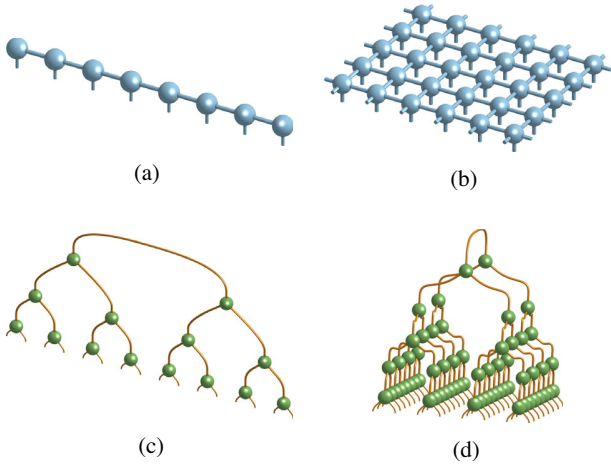
FIG. 13. Tensor network representations for a quantum many-body wave function: the matrix product states (MPS) (a) and the tree tensor network (TTN) (c) on the left hand side are for 1D systems, while the projected entangled pair states (PEPS) (b) and the TTN (d) on the right hand side represent 2D systems.

distribution of the coupling symmetry sectors within the tensors as well. In order to ensure convergence during the optimization, we dynamically increase the bond dimension locally, allowing us to adapt the symmetry sectors within the tree. In particular, we exploit the single tensor optimization with subspace expansion presented in Ref. [74], which approximates a two-site update by expanding the connecting link and iteratively optimizing the two local tensors separately. Thereby, we maintain the beneficial numerical complexity of $\mathcal{O}(\chi^4)$ instead of a heavier scaling of $\mathcal{O}(\chi^6)$ for the complete two-site optimization. For the single tensor optimization, we exploit the Arnoldi algorithm implemented in the ARPACK library. In this algorithm, the local eigenvalue problem is solved by iteratively diagonalizing the effective Hamiltonian H_{eff} for the single tensor. It delivers the lowest eigenpairs of H_{eff} up to a predefined precision ϵ by requiring only knowledge of the action of the operator H_{eff}. In the global optimization, we sweep through the TTN from the bottom to the top, performing the subspace expansion from each tensor towards its "parent" tensor [the one located directly above in the geometry of Fig. 13(d)]. After one complete sweep, we start over, iterating until global convergence (in terms of energy and selected observables) is achieved. As we come closer to convergence with each sweep, we also drive the optimization precision ϵ of the Arnoldi algorithm, such that we become more and more accurate in solving the local eigenvalue problems.

The computations with TTN presented in this work were run on different HPC clusters (the BwUniCluster and CINECA), where a single simulation of, e.g., an 8×8 system can last up to three weeks until final convergence, depending on the system parameters. Here, we point out that we can still improve the efficiency of the code and can have the potential to heavily parallelize our TTN to decrease the computational effort.

APPENDIX E: TENSOR NETWORK SIMULATIONS FOR LATTICE GAUGE THEORIES

In this section, we describe the TN approach for LGT in greater technical detail. As mentioned in Sec. II C, we already fulfill the Gauss law by choosing the local gauge-invariant states (Sec. II B) as the logical basis in the TN simulations. In particular, we use an unconstrained TTN to represent the many-body wave function [51]. We adapt the TTN structure for the 2D system as shown in Refs. [72,73]. Following the description in Ref. [74], we additionally exploit the Abelian $U(1)$ symmetry, which corresponds to the total charge $\hat{Q}$ for the TTN representation. In this way, we keep the total charge $\hat{Q}$ fixed for each simulation by choosing the proper global symmetry sector.

As discussed in Sec. II A, the chosen local basis does not naturally respect the extra link symmetry arising from the division of the Hilbert space for each link into two half-links. Thus, in additional to the LGT Hamiltonian $\hat{H}$ Eq. (1), we include a term to penalize the states violating the link constraint during the simulation. In conclusion, we simulate the Hamiltonian

$$\hat{H}_{\text{sim}} = \hat{H} + \nu \sum_{x,\mu} (1 - \delta_{2,\hat{L}_{x,\mu}}), \tag{E1}$$

with $\mu \in \{\mu_x, \mu_y\}$, where the penalty term vanishes when the link symmetry is respected and increases the energy for a state breaking the symmetry. Let us mention that this additional term translates to a nearest-neighbor interaction term in the TN simulations.

In theory, the penalty factor ν should be chosen as large as possible, as the link symmetry is strictly enforced for $\nu \to \infty$. But choosing a too-large ν leads to the optimization focusing on this penalty term only and fails to optimize for the physical quantities. Depending on the physical simulation parameter, t, m, g_e, and g_m, the penalty factor ν has to be chosen in a balanced way, such that we are able to optimize for the physical quantities as much as for the link constraint. In fact, when choosing ν too low, we end up with a result where the state does not strictly obey the link symmetry. If ν is too large, artifacts can appear in the proposed ground state, as the penalty term can introduce local minima and thus freeze the state in the optimization. These artifacts can either be a matter-antimatter pair for the vacuum regime or, as shown on the left in Fig. 14, a matter-antimatter hole for the charge-crystal regime. As the total charge $\hat{Q}$ is strictly conserved by the chosen symmetry sector during the simulation, there are only two ways to get rid of such an artifact. The optimizer has to either locally violate the link symmetry or change the state at the neighboring sites together with the artifact—both of which would increase the energy in the simulation, given a large value for ν.
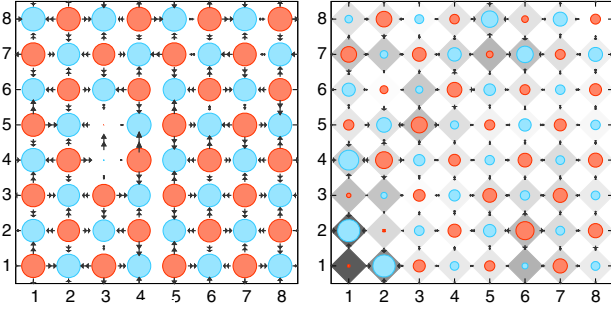
FIG. 14. Field plots from a TTN numerical simulation of 8×8 systems. The left panel depicts a configuration corresponding to a local minimum in the total energy in which a simulation got stuck because of a poor choice of the penalty parameter ν. The right panel shows the field plot for a typical randomly initialized state. Note that in this case, the link symmetry is not respected. The gray diamonds in the background of each site signal the violation of this constraint. The darkness of the gray color corresponds to the contribution of the penalty term in Eq. (E1) for the site positioned at $x \equiv (i, j)$.

In order to improve this approach, we exploit two different methods. First, we start with a random state, which, in general, violates the link symmetry. One example for this random initialization is reported on the right side of Fig. 14. Second, we drive the penalty term by increasing ν after every optimization sweep. In particular, we start by linearly increasing ν, until we observe an increment in the energy, which signals that the penalty term becomes significant for the optimization. Consequently, we switch to a quadratic tuning of ν such that, in the following few iterations, we increase ν more slowly than in the linear regime. Finally, we also set a maximum value for ν, at which we stay for the rest of the optimization. The three different regimes of driving the penalty parameter ν are depicted in Fig. 15, showing the energy difference δe to a higher bond dimension, together with ν with respect to the iterations for an exemplifying simulation.

With this driving, we optimize the random initial state in the first phase without focusing too strictly on obeying the link symmetry. This method flattens the local minima arising from including the penalty in the Hamiltonian and thereby helps converge to the global minimum. When choosing the linear tuning correctly, most physical observables are qualitatively already captured at the end of the first driving phase without strictly obeying the link symmetry. Thus, the second phase enforces the link symmetry, while the last phase—with a constant ν—optimizes the state for the final quantitative ground state.

Although introducing the driven penalty drastically decreases the number of simulations that are stuck in artificial configurations, this cannot be completely avoided. Therefore, we simulate several samples with different random initial states. From these samples, we perform a
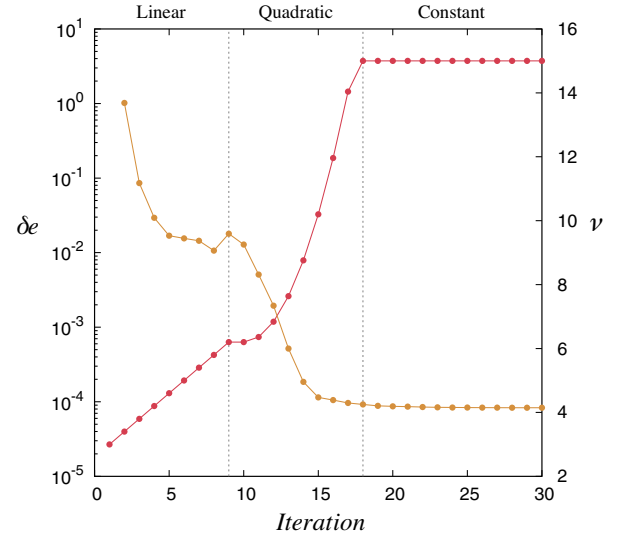

FIG. 15. Penalty parameter ν (red) and energy (yellow) with respect to the number of iterations for a typical LGT simulation. The energy is plotted here as a deviation δe from the ground-state energy obtained with highest bond dimension available. We start with linearly increasing ν. When the energy increases, we change to a quadratic driving regime with zero gradient at the transition point. Finally, we reach a predefined maximum value for ν.

postselection and check whether the obtained wave functions are indeed physically correct ground states. We also observe the typical convergence for TN with increasing bond dimension when we discard the results with artifacts. From the different samples and the convergence in bond dimension, we can estimate the relative error in the energy, which, depending on the physical parameters, typically lies in the range of about 10^{-2}–10^{-4} for an 8×8 system.

APPENDIX F: INTEGRAL ESTIMATORS FOR THE CORRELATION LENGTH

Here, we briefly discuss a strategy to estimate correlation lengths based on integrals of the correlation functions. The obvious advantage is that this strategy employs all of the data within the correlation function itself while requiring no data regression. Therefore, it can be easily automatized and needs no careful initialization of the fit parameters for data regression; at the same time, it is a reliable, only slightly biased, estimator for correlation lengths [97].

While in the main text we applied an analogous estimator to the full correlation function, in this section, we perform an estimator analysis on the connected component $C^0_{x,x'} = \langle \hat{\mathcal{O}}_x \hat{\mathcal{O}}_{x'} \rangle - \langle \hat{\mathcal{O}}_x \rangle \langle \hat{\mathcal{O}}_{x'} \rangle$ of the correlation function, which we spatially average to $\bar{C}^0_v = L^{-2} \sum_x C^0_{x,x+v}$. In the absence of strong quantum correlations, lattice systems at

low temperatures typically exhibit $\bar{C}^0(v) \simeq \alpha_0 \exp(-|v|/\xi)$ exponentially decaying in the relative coordinate modulus $|v| = \sqrt{v_x^2 + v_y^2}$, where ξ is the actual correlation length and α_0 is a (not interesting) prefactor. Here, we construct an integral estimator ξ_{est} for (connected) correlation lengths and show that, on the exponentially decaying class, it returns ξ to an acceptable precision.

In deriving these expressions, we assume that the system is much larger than the correlation length $L \gg \xi$ to avoid observing finite-size or boundary effects (we effectively approximate the lattice to $\mathbb{Z}^2$). For a 2D square lattice, we consider the following estimator:

$$\xi_{\text{est}}^2 = \frac{\sum_{v_x, v_y \in \mathbb{Z}} |v|^2 \bar{C}^0(v)}{6 \sum_{v_x, v_y \in \mathbb{Z}} \bar{C}^0(v)}, \qquad (\text{F1})$$

which, apart from the $1/6$ prefactor, is the (Euclidean) variance of $P(v) = \bar{C}^0(v)[\sum_{v'} \bar{C}^0(v')]^{-1}$, the correlation function normalized to a probability distribution over $\mathbb{Z}^2$ [assuming $\bar{C}^0(v)$ is symmetric, $\bar{C}^0(v) = \bar{C}^0(-v)$].

In the limit of correlation lengths that are large compared to the lattice spacing, $\xi \gg 1$, the discrete sums in Eq. (F2) converge to Riemann integrals,

$$\xi_{\text{est}}^2 = \frac{\int_{\mathbb{R}^2} |v|^2 \bar{C}^0(v) d^2 v}{6 \int_{\mathbb{R}^2} \bar{C}^0(v) d^2 v}, \qquad (\text{F2})$$

yielding an unbiased estimator $\xi_{\text{est}}^2 = \xi^2$ for the family of correlation functions $\bar{C}^0(v) \simeq \alpha_0 \exp(|v|/\xi)$. For correlation lengths comparable in magnitude to the lattice spacing, the finite sum in Eq. (F2) can produce a bias

$B(\xi) = \xi^2 - \xi_{\text{est}}^2 \geq 0$ in the estimator. Unfortunately, $B(\xi)$ is not an analytic function and thus cannot be removed altogether. However, we numerically verified that $B(\xi)$ is upper bounded by $1/17$, for any $\xi \in \mathbb{R}$, which makes Eq. (F2) a satisfactory estimator for the purposes of identifying phases and transitions.

APPENDIX G: BOUNDARY HAMILTONIAN AND TYPICAL BOUNDARY CONDITIONS

In this section, we discuss strategies to realize a specific set of (open) boundary conditions for problems of equilibrium electrodynamics. These strategies present an extension to the simulations realized in this work, which assume the boundary conditions to be either free (for finite charge density) or periodic (for zero charge density).

Von Neumann boundary conditions.—In this simple scenario, the outgoing electric flux at each boundary site is fixed and defined by the user. To realize this boundary condition, start the TTN algorithm from a product state that has the desired configuration of electric fluxes at the open boundary, and then simply carry out the optimization algorithm (without a boundary Hamiltonian, i.e., $H_b = 0$). The algorithm has no means of changing the electric fluxes at the boundaries and will converge to the bulk ground state, given that specific boundary flux configuration.

Dirichelet boundary conditions.—To model the scenario where the boundaries are a perfect conductor, we actually assume the boundaries to be superconductive and expel magnetic fields by displaying huge magnetic couplings at the boundary. This model requires the usage of a magnetic boundary Hamiltonian

$$
\begin{aligned}
H_b = J_b \Bigg(& \sum_{j=1}^{L-1} (\hat{U}_{(1,j),-\mu_x}^\dagger \hat{U}_{(1,j),\mu_y} \hat{U}_{(1,j+1),-\mu_x} + \hat{U}_{(j,L),\mu_y}^\dagger \hat{U}_{(j,L),\mu_x} \hat{U}_{(j+1,L),\mu_y} + \hat{U}_{(L,j+1),\mu_x}^\dagger \hat{U}_{(L,j),\mu_y}^\dagger \hat{U}_{(L,j),\mu_x} \\
& + \hat{U}_{(j+1,1),-\mu_y}^\dagger \hat{U}_{(j,1),\mu_x}^\dagger \hat{U}_{(j,1),-\mu_y}) + U_{(1,L),-\mu_x}^\dagger \hat{U}_{(1,L),\mu_y} + U_{(L,L),\mu_y}^\dagger \hat{U}_{(L,L),\mu_x} + U_{(L,1),\mu_x} \hat{U}_{(L,1),-\mu_y}^\dagger \\
& + U_{(1,1),-\mu_y}^\dagger \hat{U}_{(1,1),-\mu_x} + \text{H.c.} \Bigg),
\end{aligned}
\qquad (\text{G1})
$$

which contains both edge terms (top rows) and corner terms (bottom rows). To address the problem of electrodynamics, the ground-state algorithm is carried out while setting $J_b \gg \max\{|t|, |m|, g_e^2, g_m^2\}$, ensuring that the magnetic fields will approach a constant value (equal to zero) at the boundary, once converged.

APPENDIX H: PERTURBATION THEORY

Here, we describe the corrections to the ground state in both regimes outlined in Sec. III. Let us start by

considering particle fluctuations due to the presence of a small tunneling $|t|$. The system has periodic boundary conditions.

Perturbation around the vacuum state.—For $m \gg |t|$, the vacuum state (with zero energy) is corrected by strictly local particle-antiparticle fluctuations. The first nontrivial contribution comes from a local dimer excitation as depicted in Fig. 16, whose average energy is $2m + g_e^2/2$. The truncated Hamiltonian reads (apart from the sign of the tunneling coupling, which, however, does not affect the results)

$$H_v = \begin{bmatrix} 0 & t & \cdots & t \\ t & 2m + g_e^2/2 & & \\ \vdots & & \ddots & \\ t & & & 2m + g_e^2/2 \end{bmatrix}, \quad \text{(H1)}$$

which is a $(1 + 2L^2) \times (1 + 2L^2)$ matrix. The correction to the vacuum energy is therefore

$$E_v = \left[\frac{g_e^2}{4} + m - \sqrt{\left(\frac{g_e^2}{4} + m \right)^2 + 2L^2 t^2} \right]. \quad \text{(H2)}$$

Perturbation around the dimer state.—Small-order tunneling perturbations on top of the fully dimerized states are not sufficient to remove their degeneracy. The ground-state energy sector remains degenerate up to the fourth order in perturbation theory. Here, we focus on the smallest order energy corrections for one specific dimerized configuration and consider the possible excitations as depicted in Fig. 16. We now have two different excitation sectors, depending on where we remove a particle-antiparticle pair: When the pair is annihilated on top of a dimer, the energy cost is $2m + g_e^2/2$; otherwise, when we remove a pair in between two dimers, it costs $2m - g_e^2/2$. The number of possible configurations of the first type coincides with the number of dimers, i.e., $L^2/2$; in the other case, we have $3L^2/2$ different possibilities. The full truncated Hamiltonian is still a $(1 + 2L^2) \times (1 + 2L^2)$ matrix, which now reads [apart from the overall extensive constant $E_N \equiv (2m + g_e^2/2)L^2/2$]
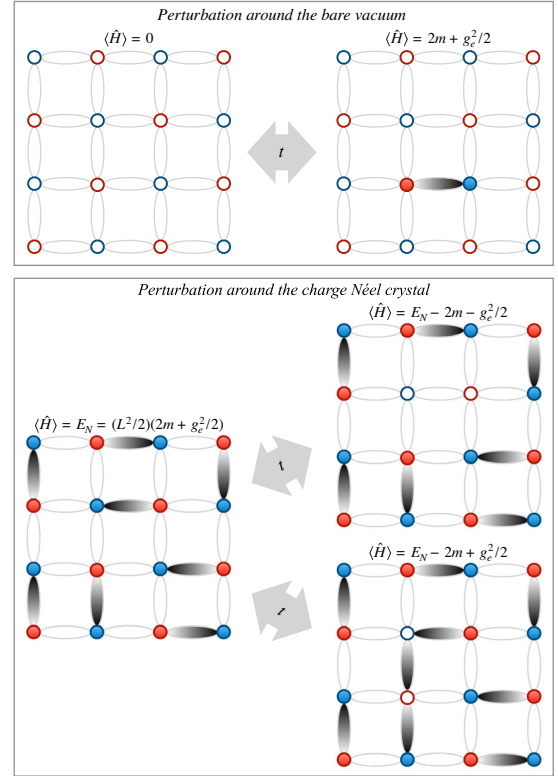

FIG. 16. Example of excited states coupled to the vacuum (top) or to the fully dimerized state (bottom) at the lowest order in perturbation theory in the tunneling coupling t, as described in Appendix H.

$$H_d = \begin{bmatrix} 0 & t & \cdots & t & t & \cdots & t \\ t & -2m - g_e^2/2 & & & & & \\ \vdots & & \ddots & & & & \\ t & & & -2m - g_e^2/2 & & & \\ t & & & & -2m + g_e^2/2 & & \\ \vdots & & & & & \ddots & \\ t & & & & & & -2m + g_e^2/2 \end{bmatrix}, \quad \text{(H3)}$$

where, also in this case, the sign of t does not affect the results. The correction to the vacuum energy can be evaluated as well by solving $\det(H_d - \varepsilon) = 0$; indeed, because of the structure of the matrix, and thanks to the properties of the determinant, we find

$$E_d = \frac{L^2}{2} \left(\frac{g_e^2}{2} + 2m \right) + \varepsilon_-, \quad \text{(H4)}$$

where ε_- is the negative solution of

$$\varepsilon \left[\frac{g_e^4}{4} - (2m + \varepsilon)^2 \right] + 2L^2 t^2 \left(\frac{g_e^2}{4} + 2m + \varepsilon \right) = 0. \quad \text{(H5)}$$

Now, it is clear that, when m is approaching the value $-g_e^2/4$, the biggest corrections, at the lower order in t, solely come from the sector that is quasidegenerate with the
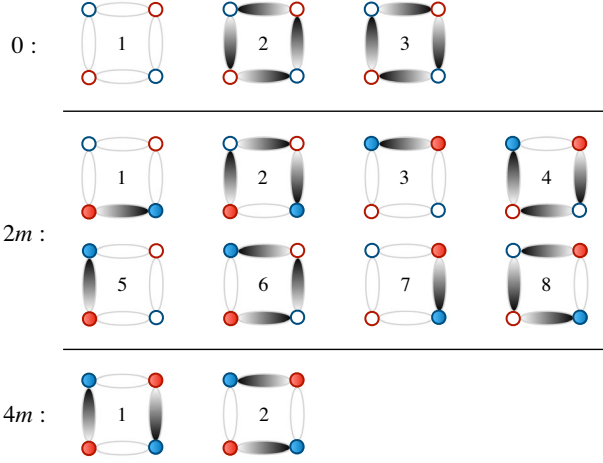
FIG. 17. Graphic representation of the basis vectors, in each mass sector, used to build up the 2×2 LGT Hamiltonian in the $S = 1$ representation, as outlined in Appendix I.

classical dimerized configuration. The finite-size scaling depends on whether the mass is approaching from above (i.e., from the vacuum) or form below (i.e., from the dimerized configuration): In the first case, $2L^2$ states contribute to the energy corrections; in the second case, if $g_e > 0$, only $L^2/2$ states get involved. Here, an energy gap, $|E_v - E_d| \sim Lt/\sqrt{2}$, opens. Notice that, in the pathological situation where $g_e = 0$ as well, there is no gap opening at the second order in t, and therefore a sharper transition is expected.

Let us mention that the correction to the ground-state energy coincides, as it should, with the second-order degenerate perturbation theory. In practice, if $\hat{\mathcal{Q}}$ is the projector into the classical charge-crystal sector and $\hat{\mathcal{P}} = 1 - \hat{\mathcal{Q}}$ the projector into the complementary sector, then we may split the eigenvectors into two contributions: $|E_k\rangle = |\phi_k\rangle + |\varphi_k\rangle$, where $|\phi_k\rangle \equiv \hat{\mathcal{Q}}|E_k\rangle$ and $|\varphi_k\rangle \equiv \hat{\mathcal{P}}|E_k\rangle$. The eigenvalue equation $(\hat{H}_0 - t\hat{V})|E_k\rangle = E_k|E_k\rangle$ therefore splits into two coupled equations:

$$-t\hat{\mathcal{P}}\,\hat{V}\,|\phi_k\rangle = (E_k - \hat{H}_0 + t\hat{\mathcal{P}}\,\hat{V}\,\hat{\mathcal{P}})|\varphi_k\rangle, \quad \text{(H6)}$$

$$-t\hat{\mathcal{Q}}\,\hat{V}\,|\varphi_k\rangle = (E_k - E_N)|\phi_k\rangle, \quad \text{(H7)}$$

where we used the fact that, in our case, $\hat{\mathcal{Q}}\,\hat{V}\,\hat{\mathcal{Q}} = 0$. Corrections within the degenerate subsector are thus given by recursively solving the following equation:

$$(E_k - E_N)|\phi_k\rangle = \hat{\mathcal{Q}}\,\hat{V}\,\hat{\mathcal{P}}\frac{t^2}{E_k - \hat{H}_0 + t\hat{\mathcal{P}}\,\hat{V}\,\hat{\mathcal{P}}}\hat{\mathcal{P}}\,\hat{V}\,|\phi_k\rangle. \quad \text{(H8)}$$

At the second order in the tunneling, the dimerized subsector degeneracy is not lifted, and the energy changes according to Eq. (H5). Let us stress that, deep in the

charge-crystal regime, these are the dominant corrections. However, close to the classical transition, the creation or annihilation of a particle-antiparticle is energetically favorable, and nontrivial corrections to the degeneracy of the ground-state energy sector are induced by fourth-order tunneling transitions: Two different classical dimerized states are coupled whenever they share at least one "resonating" *plaquette*, which consists in two neighboring horizontal or vertical dimers (see the $4m$ mass sector in Fig. 17). This effect partially removes the ground-state degeneracy, making a specific superposition of different dimer states energetically favorable. Incidentally, let us mention that, in the thermodynamic limit, there exist classical dimer configurations, e.g., the state where dimers are all vertically (horizontally) aligned with all local electric fluxes pointing in the same direction, which are not resonating with any other fully dimerized state at any order in perturbation theory.

APPENDIX I: EXACT RESULTS OF THE 2 × 2 SYSTEM

In the zero charge density sector, the single plaquette system, i.e., 2×2, admits only 13 gauge-invariant diagonal configurations in the spin-1 compact representation of the electric field. The full Hamiltonian can be easily constructed by considering each mass sector $\{0, 2m, 4m\}$ independently, and it acquires the following block structure,

$$H_{2\times 2} = \begin{bmatrix} D_0 & T_{02} & \varnothing \\ T_{20} & D_2 & T_{24} \\ \varnothing & T_{42} & D_4 \end{bmatrix}, \quad \text{(I1)}$$

where $D_j = D_j^\dagger$, $T_{20} = T_{02}^\dagger$, $T_{42} = T_{24}^\dagger$, and all matrix entries are real. To construct each block, we use the gauge-invariant eigenstates of the electric field $\hat{E}_{x,\mu}$ and particle number $\hat{n}_x$, as listed in Fig. 17.

The diagonal blocks read

$$D_0 = \begin{pmatrix} 0 & -g_m^2/2 & -g_m^2/2 \\ -g_m^2/2 & 2g_e^2 & 0 \\ -g_m^2/2 & 0 & 2g_e^2 \end{pmatrix}, \quad \text{(I2)}$$

$$D_2 = \mathbb{I}_4 \otimes \begin{pmatrix} 2m + g_e^2/2 & -g_m^2/2 \\ -g_m^2/2 & 2m + 3g_e^2/2 \end{pmatrix}, \quad \text{(I3)}$$

$$D_4 = \begin{pmatrix} 4m + g_e^2 & -g_m^2/2 \\ -g_m^2/2 & 4m + g_e^2 \end{pmatrix}, \quad \text{(I4)}$$

where $\mathbb{I}_4$ is a 4×4 identity matrix. The out-diagonal blocks are responsible for creation or annihilation of particle-antiparticle pairs and are given by
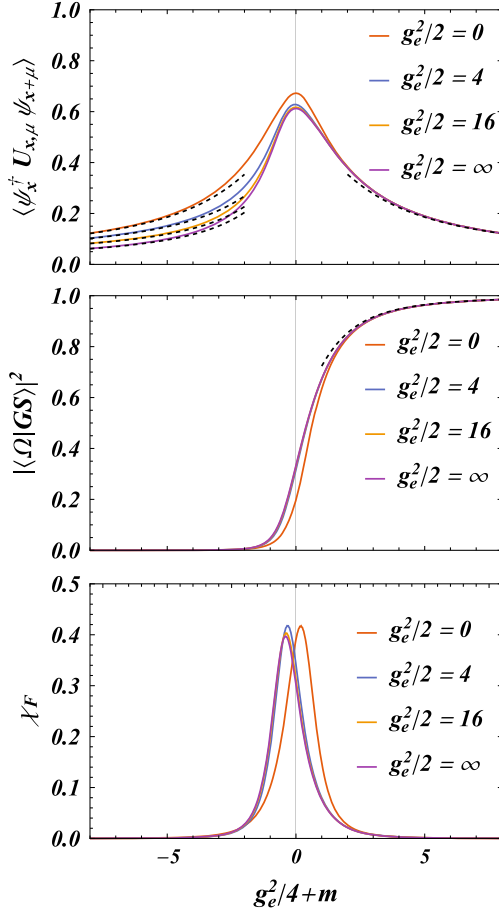
FIG. 18. Top panel: expectation value of the tunneling Hamiltonian as a function of the distance from the classical transition for a 2×2 system. Center panel: module square of the overlap between the exact 2×2 ground state and the vacuum state when varying the coupling across the classical transition point. Dashed lines are the perturbative predictions. Bottom panel: fidelity susceptibility of the ground state as defined in the main text.

$$T_{02} = \begin{pmatrix} -t & 0 & -t & 0 & t & 0 & -t & 0 \\ 0 & 0 & 0 & 0 & 0 & t & 0 & -t \\ 0 & -t & 0 & -t & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (\mathrm{I}5)$$

$$T_{42} = \begin{pmatrix} 0 & -t & 0 & -t & -t & 0 & -t & 0 \\ -t & 0 & -t & 0 & 0 & -t & 0 & -t \end{pmatrix}. \quad (\mathrm{I}6)$$

The exact diagonalization of the Hamiltonian $\hat{H}_{2\times2}$ allows us to explore the behavior of the ground state in the vicinity of the transition $m \simeq g_e^2/4$. As expected from the enhancement of quantum fluctuations, the gauge-invariant hopping term gets picked at the transition (Fig. 18, top panel). The overlap of the ground state with the bare vacuum as a function of m for different values of the electric coupling is analyzed as well (central panel in Fig. 18). Exact curves are compared with first-order perturbative results.
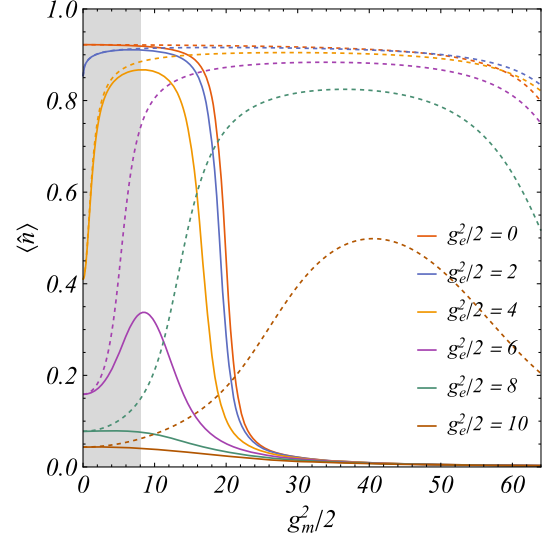

FIG. 19. Behavior of the particle density vs the magnetic coupling in the 2×2 system for $m = -2$ and different electric couplings $g_e^2/2$. The shaded gray area represents the region explored in Fig. 6. Solid lines are the $S = 1$ results; dashed lines are the $S = 2$ results.

In order to more carefully explore the transition region, we look at the fidelity susceptibility of the ground state [98–101], $\chi_F(m) \equiv \langle \partial_m GS(m) | \partial_m GS(m) \rangle - |\langle GS(m) | \partial_m GS(m) \rangle|^2$, which gives the leading contribution to the ground-state fidelity $|\langle GS(m) | GS(m+\delta) \rangle| = 1 - \delta^2 \chi_F(m)/2 + o(\delta^2)$, since the linear contribution in δ vanishes due to the normalization condition $\langle GS(m) | GS(m) \rangle = 1$. This quantity is the perfect indicator of a change in the geometrical properties of the ground state when varying the couplings. Moreover, from perturbation theory, it can be easily shown that $\chi_F(m) \leq [\langle GS(m) | (\sum_x \hat{n}_x)^2 | GS(m) \rangle - \langle GS(m) | \sum_x \hat{n}_x | GS(m) \rangle^2]/\Delta^2$, where Δ is the energy gap between the ground state and the lower excitations. In practice, the fidelity susceptibility of the ground state is bounded from above by the number of particle fluctuations (which is an extensive quantity) divided by the gap. Whenever $\chi_F(m)$ shows a superextensive behavior, the ground state of the system should be gapless. From the numerical data, we have confirmation that $\chi_F(m)$ is enhanced in the vicinity of the transition between the two regions, as depicted in the bottom panel of Fig. 18.

In Fig. 19, we reproduce the behavior of the matter density as a function of the magnetic coupling, for different values of the electric-field couplings. As explained in the main text and confirmed by these exact results in the 2×2 plaquette, the local density gets enhanced by applying a small magnetic coupling; however, when $g_m^2 \simeq g_e^2$, the particle density starts decreasing and eventually vanishes for $g_m^2 \gg g_e^2$. Let us stress that this phenomenon is strictly due to the finite compact representation of the gauge field.

Indeed, when gauge-field fluctuations are very strong, we may expect deviations in the observables due to the
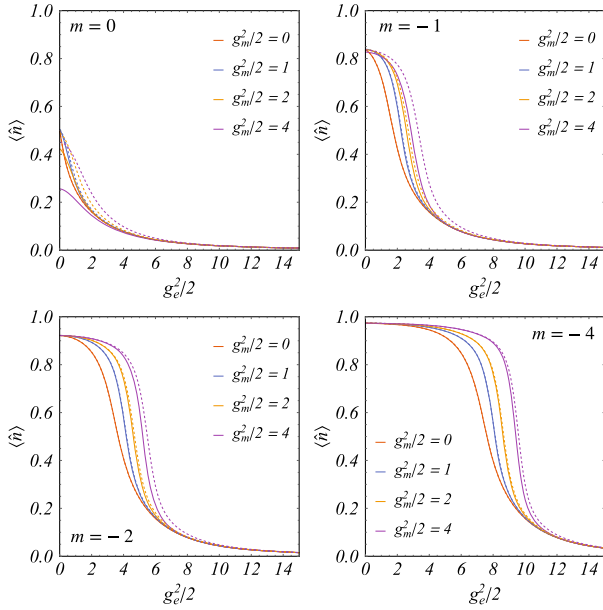
FIG. 20. Particle density as a function of the electric coupling g_e for different values of the magnetic coupling g_m. The four panels represent different bare masses. Solid lines are the $S = 1$ results; dashed lines are the $S = 2$ results.

finite spin representation of the electric field; in order to have an estimate of the finite-S representation accuracy, we further analyze the 2×2 plaquette system in the $S = 2$ compact representation, namely, when the electric field (in units of flux) can have the values $\{-2, -1, 0, 1, 2\}$. The full Hamiltonian still preserves the block structure in Eq. (I1), where now each mass sector acquires further gauge-invariant states, for a total of five states in the zero-mass sector; 16 states in the $2m$-mass sector; four states in the $4m$-mass sector.

As a matter of fact, when $S = 2$, the phenomenon of density suppression depicted in Fig. 19 occurs for much larger values of the magnetic couplings, thus disappearing in the limit $S \to \infty$.

In Fig. 20, we compare the matter density for the two compact representations $S = 1, 2$ and different values of the couplings. As expected, for $g_e^2 \gg g_m^2$, the two representations are equivalent; moreover, if $-m \gg 1$ (i.e., very negative) the diagonal configurations are more energetically favorable, and even for small electric coupling and a finite value of g_m^2, the truncation of the gauge-field representation does not affect the results too much ($S = 1$ and $S = 2$ are almost identical indeed); of course, for $m \geq 0$, this is not the case, and we need $g_e^2 \gg g_m^2$. Notice that, in the actual QED, this condition is satisfied as long as the electric coupling is sufficiently large since $g_e^2 \sim g_m^{-2} \sim g^2$.

[1] E. Zohar, J. I. Cirac, and B. Reznik, *Simulating Compact Quantum Electrodynamics with Ultracold Atoms: Probing Confinement and Nonperturbative Effects*, Phys. Rev. Lett. **109,** 125302 (2012).

[2] D. Banerjee, M. Dalmonte, M. Müller, E. Rico, P. Stebler, U.-J. Wiese, and P. Zoller, *Atomic Quantum Simulation of Dynamical Gauge Fields Coupled to Fermionic Matter: From String Breaking to Evolution after a Quench*, Phys. Rev. Lett. **109,** 175302 (2012).

[3] S. P. Jordan, K. S. M. Lee, and J. Preskill, *Quantum Algorithms for Quantum Field Theories*, Science **336,** 1130 (2012).

[4] L. Tagliacozzo, A. Celi, A. Zamora, and M. Lewenstein, *Optical Abelian Lattice Gauge Theories*, Ann. Phys. (Amsterdam) **330,** 160 (2013).

[5] D. Banerjee, M. Bögli, M. Dalmonte, E. Rico, P. Stebler, U.-J. Wiese, and P. Zoller, *Atomic Quantum Simulation of* $\mathbf{U}(n)$ *and* $\mathrm{SU}(n)$ *Non-Abelian Lattice Gauge Theories*, Phys. Rev. Lett. **110,** 125303 (2013).

[6] E. Zohar, J. I. Cirac, and B. Reznik, *Cold-Atom Quantum Simulator for SU(2) Yang-Mills Lattice Gauge Theory*, Phys. Rev. Lett. **110,** 125304 (2013).

[7] L. Tagliacozzo, A. Celi, P. Orland, M. W. Mitchell, and M. Lewenstein, *Simulation of Non-Abelian Gauge Theories with Optical Lattices*, Nat. Commun. **4,** 2615 (2013).

[8] A. Mezzacapo, E. Rico, C. Sabín, I. L. Egusquiza, L. Lamata, and E. Solano, *Non-Abelian SU(2) Lattice Gauge Theories in Superconducting Circuits*, Phys. Rev. Lett. **115,** 240502 (2015).

[9] U.-J. Wiese, *Ultracold Quantum Gases and Lattice Systems: Quantum Simulation of Lattice Gauge Theories*, Ann. Phys. (N.Y.) **525,** 777 (2013).

[10] E. Zohar, J. I. Cirac, and B. Reznik, *Quantum Simulations of Lattice Gauge Theories Using Ultracold Atoms in Optical Lattices*, Rep. Prog. Phys. **79,** 014401 (2015).

[11] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller, and R. Blatt, *Real-Time Dynamics of Lattice Gauge Theories with a Few-Qubit Quantum Computer*, Nature (London) **534,** 516 (2016).

[12] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. V. Acoleyen, F. Verstraete, U. J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, *Simulating Lattice Gauge Theories within Quantum Technologies*, arXiv:1911.00003.

[13] M. C. Bañuls and K. Cichy, *Review on Novel Methods for Lattice Gauge Theories*, arXiv:1910.00257.

[14] I. J. R. Aitchison and A. J. G. Hey, *Gauge Theories in Particle Physics: A Practical Introduction. Vol. 1: From Relativistic Quantum Mechanics to QED* (CRC Press, Bristol, England, 2003).

[15] T. Cheng and L. Li, *Gauge Theory of Elementary Particle Physics* (Oxford University, New York, 2006).

[16] J. Kogut and L. Susskind, *Hamiltonian Formulation of Wilson's Lattice Gauge Theories*, Phys. Rev. D **11,** 395 (1975).

[17] E. Fradkin, *Field Theories of Condensed Matter Physics*, 2nd ed. (Cambridge University Press, Cambridge, England, 2013).

[18] N. Brambilla *et al.*, *QCD and Strongly Coupled Gauge Theories: Challenges and Perspectives*, Eur. Phys. J. C **74**, 2981 (2014).

[19] J. Schwinger, *Gauge Invariance and Mass. II*, Phys. Rev. **128**, 2425 (1962).

[20] K. G. Wilson, *Confinement of Quarks*, Phys. Rev. D **10**, 2445 (1974).

[21] J. B. Kogut, *An Introduction to Lattice Gauge Theory and Spin Systems*, Rev. Mod. Phys. **51**, 659 (1979).

[22] J. B. Kogut, *The Lattice Gauge Theory Approach to Quantum Chromodynamics*, Rev. Mod. Phys. **55**, 775 (1983).

[23] C. J. Hamer, Z. Weihong, and J. Oitmaa, *Series Expansions for the Massive Schwinger Model in Hamiltonian Lattice Theory*, Phys. Rev. D **56**, 55 (1997).

[24] M. Troyer and U.-J. Wiese, *Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations*, Phys. Rev. Lett. **94**, 170201 (2005).

[25] T. M. R. Byrnes, P. Sriganesh, R. J. Bursill, and C. J. Hamer, *Density Matrix Renormalization Group Approach to the Massive Schwinger Model*, Phys. Rev. D **66**, 013002 (2002).

[26] P. Silvi, E. Rico, T. Calarco, and S. Montangero, *Lattice Gauge Tensor Networks*, New J. Phys. **16**, 103015 (2014).

[27] T. Pichler, M. Dalmonte, E. Rico, P. Zoller, and S. Montangero, *Real-Time Dynamics in U(1) Lattice Gauge Theories with Tensor Networks*, Phys. Rev. X **6**, 011023 (2016).

[28] M. Dalmonte and S. Montangero, *Lattice Gauge Theory Simulations in the Quantum Information Era*, Contemp. Phys. **57**, 388 (2016).

[29] M. C. Bañuls, K. Cichy, J. I. Cirac, K. Jansen, and H. Saito, *Matrix Product States for Lattice Field Theories*, Proc. Sci. LATTICE 2013 (2013).

[30] M. C. Bañuls, K. Cichy, J. I. Cirac, K. Jansen, and S. Kühn, *Tensor Networks and their Use for Lattice Gauge Theories*, Proc. Sci. LATTICE 2018 (2018).

[31] P. Silvi, Y. Sauer, F. Tschirsich, and S. Montangero, *Tensor Network Simulation of an SU(3) Lattice Gauge Theory in 1D*, Phys. Rev. D **100**, 074512 (2019).

[32] S. Montangero, *Introduction to Tensor Network Methods* (Springer, Cham, 2018).

[33] L. Funcke, K. Jansen, and S. Kühn, *Topological Vacuum Structure of the Schwinger Model with Matrix Product States*, Phys. Rev. D **101**, 054507 (2020).

[34] F. Bruckmann, K. Jansen, and S. Kühn, *O(3) Nonlinear Sigma Model in $1+1$ Dimensions with Matrix Product States*, Phys. Rev. D **99**, 074501 (2019).

[35] L. Tagliacozzo, A. Celi, and M. Lewenstein, *Tensor Networks for Lattice Gauge Theories with Continuous Groups*, Phys. Rev. X **4**, 041024 (2014).

[36] B. Buyens, J. Haegeman, K. Van Acoleyen, H. Verschelde, and F. Verstraete, *Matrix Product States for Gauge Field Theories*, Phys. Rev. Lett. **113**, 091601 (2014).

[37] E. Rico, T. Pichler, M. Dalmonte, P. Zoller, and S. Montangero, *Tensor Networks for Lattice Gauge Theories and Atomic Quantum Simulation*, Phys. Rev. Lett. **112**, 201601 (2014).

[38] J. Haegeman, K. Van Acoleyen, N. Schuch, J. I. Cirac, and F. Verstraete, *Gauging Quantum States: From Global to Local Symmetries in Many-Body Systems*, Phys. Rev. X **5**, 011024 (2015).

[39] P. Silvi, E. Rico, M. Dalmonte, F. Tschirsich, and S. Montangero, *Finite-Density Phase Diagram of a $(1+1)-d$ Non-Abelian Lattice Gauge Theory with Tensor Networks*, Quantum **1**, 9 (2017).

[40] M. C. Bañuls, K. Cichy, J. I. Cirac, K. Jansen, and S. Kühn, *Density Induced Phase Transitions in the Schwinger Model: A Study with Matrix Product States*, Phys. Rev. Lett. **118**, 071601 (2017).

[41] P. Sala, T. Shi, S. Kühn, M. C. Bañuls, E. Demler, and J. I. Cirac, *Gaussian States for the Variational Study of $(1+1)$-Dimensional Lattice Gauge Models*," arXiv:1811.04899.

[42] A. Polyakov, *Quark Confinement and Topology of Gauge Theories*, Nucl. Phys. **B120**, 429 (1977).

[43] T. Banks, R. Myerson, and J. Kogut, *Phase Transitions in Abelian Lattice Gauge Theories*, Nucl. Phys. **B129**, 493 (1977).

[44] S. D. Drell, H. R. Quinn, B. Svetitsky, and M. Weinstein, *Quantum Electrodynamics on a Lattice: A Hamiltonian Variational Approach to the Physics of the Weak-Coupling Region*, Phys. Rev. D **19**, 619 (1979).

[45] S. Ben-Menahem, *Confinement in Compact QED for Low Couplings*, Phys. Rev. D **20**, 1923 (1979).

[46] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler, *A Rydberg Quantum Simulator*, Nat. Phys. **6**, 382 (2010).

[47] E. Zohar and J. I. Cirac, *Combining Tensor Networks with Monte Carlo Methods for Lattice Gauge Theories*, Phys. Rev. D **97**, 034510 (2018).

[48] F. Tschirsich, S. Montangero, and M. Dalmonte, *Phase Diagram and Conformal String Excitations of Square Ice using Gauge Invariant Matrix Product States*, SciPost Phys. **6**, 28 (2019).

[49] A. Celi, B. Vermersch, O. Viyuela, H. Pichler, M. D. Lukin, and P. Zoller, *Emerging 2D Gauge Theories in Rydberg Configurable Arrays*, Phys. Rev. X **10**, 021057 (2020).

[50] Y.-P. Huang, D. Banerjee, and M. Heyl, *Dynamical Quantum Phase Transitions in U(1) Quantum Link Models*, Phys. Rev. Lett. **122**, 250401 (2019).

[51] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero, *Unconstrained Tree Tensor Network: An Adaptive Gauge Picture for Enhanced Performance*, Phys. Rev. B **90**, 125154 (2014).

[52] D. Horn, *Finite Matrix Models with Continuous Local Gauge Invariance*, Phys. Lett. B **100**, 149 (1981).

[53] P. Orland and D. Rohrlich, *Lattice Gauge Magnets: Local Isospin From Spin*, Nucl. Phys. **B338**, 647 (1990).

[54] S. Chandrasekharan and U.-J. Wiese, *Quantum Link Models: A Discrete Approach to Gauge Theories*, Nucl. Phys. **B492**, 455 (1997).

[55] R. Brower, S. Chandrasekharan, and U.-J. Wiese, *QCD as a Quantum Link Model*, Phys. Rev. D **60**, 094502 (1999).

[56] R. J. Goldston and P. H. Rutherford, *Introduction to Plasma Physics* (CRC Press: Taylor & Francis, London, 1995).

[57] I. Bloch, J. Dalibard, and W. Zwerger, *Many-Body Physics with Ultracold Gases*, Rev. Mod. Phys. **80**, 885 (2008).

[58] T. Hartung and K. Jansen, *Zeta-Regularized Vacuum Expectation Values*, J. Math. Phys. (N.Y.) **60**, 093504 (2019).

[59] F. M. Surace, P. P. Mazza, G. Giudici, A. Lerose, A. Gambassi, and M. Dalmonte, *Lattice Gauge Theories and String Dynamics in Rydberg Atom Quantum Simulators*, Phys. Rev. X **10**, 021041 (2020).

[60] S. Notarnicola, M. Collura, and S. Montangero, *Real Time Dynamics Quantum Simulation of (1 + 1)-D Lattice QED with Rydberg Atoms*, Phys. Rev. Research **2**, 013288 (2020).

[61] M. Bañuls, K. Cichy, J. Cirac, and K. Jansen, *The Mass Spectrum of the Schwinger Model with Matrix Product States*, J. High Energy Phys. 11 (2013) 158.

[62] E. Zohar and M. Burrello, *Formulation of Lattice Gauge Theories for Quantum Simulations*, Phys. Rev. D **91**, 054506 (2015).

[63] S. Notarnicola, E. Ercolessi, P. Facchi, G. Marmo, S. Pascazio, and F. V. Pepe, *Discrete Abelian Gauge Theories for Quantum Simulations of QED*, J. Phys. A **48**, 30FT01 (2015).

[64] E. Ercolessi, P. Facchi, G. Magnifico, S. Pascazio, and F. V. Pepe, *Phase Transitions in Z_n Gauge Models: Towards Quantum Simulations of the Schwinger-Weyl QED*, Phys. Rev. D **98**, 074503 (2018).

[65] P. Corboz, G. Evenbly, F. Verstraete, and G. Vidal, *Simulation of Interacting Fermions with Entanglement Renormalization*, Phys. Rev. A **81**, 010303(R) (2010).

[66] P. Corboz, R. Orús, B. Bauer, and G. Vidal, *Simulation of Strongly Correlated Fermions in Two Spatial Dimensions with Fermionic Projected Entangled-Pair States*, Phys. Rev. B **81**, 165104 (2010).

[67] C. V. Kraus, N. Schuch, F. Verstraete, and J. I. Cirac, *Fermionic Projected Entangled Pair States*, Phys. Rev. A **81**, 052338 (2010).

[68] I. Pižorn and F. Verstraete, *Fermionic Implementation of Projected Entangled Pair States Algorithm*, Phys. Rev. B **81**, 245110 (2010).

[69] Q.-Q. Shi, S.-H. Li, J.-H. Zhao, and H.-Q. Zhou, *Graded Projected Entangled-Pair State Representations and an Algorithm for Translationally Invariant Strongly Correlated Electronic Systems on Infinite-Size Lattices in Two Spatial Dimensions*, arXiv:0907.5520.

[70] T. Barthel, C. Pineda, and J. Eisert, *Contraction of Fermionic Operator Circuits and the Simulation of Strongly Correlated Fermions*, Phys. Rev. A **80**, 042333 (2009).

[71] C. Pineda, T. Barthel, and J. Eisert, *Unitary Circuits for Strongly Correlated Fermions*, Phys. Rev. A **81**, 050303 (R) (2010).

[72] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte, and S. Montangero, *Fractional Quantum Hall Effect in the Interacting Hofstadter Model via Tensor Networks*, Phys. Rev. B **96**, 195123 (2017).

[73] M. Collura, L. Dell'Anna, T. Felser, and S. Montangero, *On the Descriptive Power of Neural-Networks as Constrained Tensor Networks with Exponentially Large Bond Dimension*, arXiv:1905.11351.

[74] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, *The Tensor Networks Anthology: Simulation Techniques for Many-Body Quantum Lattice Systems*, SciPost Phys. Lect. Notes 8 (2019).

[75] I. P. McCulloch, *From Density-Matrix Renormalization Group to Matrix Product States*, J. Stat. Mech. (2007) P10014.

[76] S. Singh and G. Vidal, *Global Symmetries in Tensor Network States: Symmetric Tensors Versus Minimal Bond Dimension*, Phys. Rev. B **88**, 115147 (2013).

[77] P. Kasteleyn, *The Statistics of Dimers on a Lattice: I. The Number of Dimer Arrangements on a Quadratic Lattice*, Physica **27**, 1209 (1961).

[78] K. Fredenhagen and M. Marcu, *Confinement Criterion for QCD with Dynamical Quarks*, Phys. Rev. Lett. **56**, 223 (1986).

[79] K. Gregor, D. A. Huse, and R. Moessner, and S. L. Sondhi, *Diagnosing Deconfinement and Topological Order*, New J. Phys. **13**, 025009 (2011).

[80] E. Zohar, M. Burrello, T. B. Wahl, and J. I. Cirac, *Fermionic Projected Entangled Pair States and Local U(1) Gauge Theories*, Ann. Phys. (Amsterdam) **363**, 385 (2015).

[81] R. Orús, *Tensor Networks for Complex Quantum Systems*, Nat. Rev. Phys. **1**, 538 (2019).

[82] U. Schollwöck, *The Density-Matrix Renormalization Group in the Age of Matrix Product States*, Ann. Phys. (Amsterdam) **326**, 96 (2011).

[83] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, *Criticality, the Area Law, and the Computational Power of Projected Entangled Pair States*, Phys. Rev. Lett. **96**, 220601 (2006).

[84] F. Verstraete and J. I. Cirac, *Matrix Product States Represent Ground States Faithfully*, Phys. Rev. B **73**, 094423 (2006).

[85] S. Östlund and S. Rommer, *Thermodynamic Limit of Density Matrix Renormalization*, Phys. Rev. Lett. **75**, 3537 (1995).

[86] F. Verstraete and J. I. Cirac, *Valence-Bond States for Quantum Computation*, Phys. Rev. A **70**, 060302(R) (2004).

[87] R. Orús, *A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States*, Ann. Phys. (Amsterdam) **349**, 117 (2014).

[88] Y.-Y. Shi, L.-M. Duan, and G. Vidal, *Classical Simulation of Quantum Many-Body Systems with a Tree Tensor Network*, Phys. Rev. A **74**, 022320 (2006).

[89] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio, *Homogeneous Binary Trees as Ground States of Quantum Critical Hamiltonians*, Phys. Rev. A **81**, 062335 (2010).

[90] G. Vidal, *Entanglement Renormalization*, Phys. Rev. Lett. **99**, 220405 (2007).

[91] G. Evenbly and G. Vidal, *Entanglement Renormalization in Two Spatial Dimensions*, Phys. Rev. Lett. **102**, 180406 (2009).

[92] S. R. White, *Density Matrix Formulation for Quantum Renormalization Groups*, Phys. Rev. Lett. **69**, 2863 (1992).

[93] S. R. White and A. E. Feiguin, *Real-Time Evolution Using the Density Matrix Renormalization Group*, Phys. Rev. Lett. **93**, 076401 (2004).

[94] G. Vidal, *Efficient Classical Simulation of Slightly Entangled Quantum Computations*, Phys. Rev. Lett. **91**, 147902 (2003).

[95] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, *Contracting Projected Entangled Pair States is Average-Case Hard*, Phys. Rev. Research **2**, 013010 (2020).

[96] A. J. Ferris, *Area Law and Real-Space Renormalization*, Phys. Rev. B **87**, 125139 (2013).

[97] T. D. Kühner, S. R. White, and H. Monien, *One-Dimensional Bose-Hubbard Model with Nearest-Neighbor Interaction*, Phys. Rev. B **61**, 12474 (2000).

[98] S.-J. Gu, *Fidelity Approach to Quantum Phase Transitions*, Int. J. Mod. Phys. B **24**, 4371 (2010).

[99] B. Damski, *Fidelity Susceptibility of the Quantum Ising Model in a Transverse Field: The Exact Solution*, Phys. Rev. E **87**, 052131 (2013).

[100] D. Braun, G. Adesso, F. Benatti, R. Floreanini, U. Marzolino, M. W. Mitchell, and S. Pirandola, *Quantum-Enhanced Measurements without Entanglement*, Rev. Mod. Phys. **90**, 035006 (2018).

[101] D. Rossini and E. Vicari, *Ground-State Fidelity at First-Order Quantum Transitions*, Phys. Rev. E **98**, 062137 (2018).
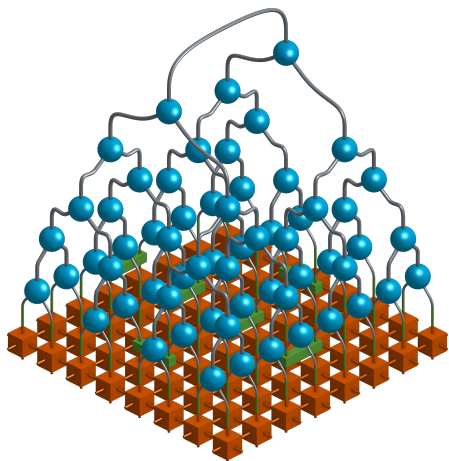
5

augmented Tree Tensor Network

This chapter describes a novel Tensor Network geometry, the augmented Tree Tensor Network (aTTN), which was developed in the scope of this thesis to efficiently study high-dimensional quantum many-body systems. In particular, the developed Tensor Network was introduced as a way to numerical benchmarks of near-future quantum simulators and quantum computers, showing the fact that Tensor Networks might have a relevant impact not only within the vast-growing Tensor-Network-community but further on experimental groups working on the implementation of quantum simulation and quantum computing protocols. In fact, very recent experiments have proven the possibility to realize quantum many-body states on a variety of platforms such as superconducting qubits, ultracold atoms, trapped ions and Rydberg-atom lattices with high dimensions at unprecedented sizes [26, 29, 31, 192]. These experimental setups pave the way for future quantum technologies, however, require a robust numerical apparatus able to provide proper validation and benchmarking.

As discussed in the previous Chapt. 2, Tensor Networks have been developed over the last three decades to simulate such quantum many-body systems on classical computers. The Matrix Product State (MPS) [46, 48, 61–63, 136, 137] has been established as the main workhorse for one-dimensional systems, not only against other Tensor Network approaches, but against other variational ansätze as well (see Chapt. 6). However, the development of Tensor Network algorithms for studying quantum many-body behaviour in two- or even higher-dimensional systems is still ongoing, since combining accuracy and scalability in simulating these systems remains a difficult challenge for Tensor Network methods [21, 24]. Different Tensor Network geometries have been brought forward and studied extensively to tackle this challenge with the most prominent being the Projected Entangled Pair States (PEPS) [45, 71, 77–79] for two-dimensional systems only, and the Tree Tensor Networks (TTN) [22, 35, 37, 67, 68, 80], as well as MERA [66, 73–75], which, in principle, can be defined for any system dimensionality. However, all of these Tensor Network geometries have fundamental limitations in representing the underlying entanglement properties of a high-dimensional quantum state and are yet to overcome the challenge to efficiently combine accuracy and scalability (see Sec. 2.2 and Chapt. 3).

At this point, the augmented Tree Tensor Network (aTTN) provides a novel tool that can capture complex entanglement properties of high-dimensional systems, while keeping a favourable computational complexity compared to the competing Tensor Network geometries, such as PEPS and MERA. The aTTN combines the advantages of MERA and TTN, making it possible to simulate systems in and out of criticality at unprecedented system sizes, up to $N = 64 \times 64$.

In the following, we firstly provide the manuscript *Efficient Tensor Network ansatz for high-dimensional quantum many-body problems* published in the journal *Physical Review Letters*, **126** [23]. Therein, we propose the new Tensor Network ansatz in the context of paving the way to numerical benchmarks of near-future quantum simulators and quantum computers. We further show that our ansatz outperforms most of the competing techniques, which are currently limited to smaller system sizes. As an experimental case of study, we investigate a two-dimensional Rydberg-atom lattice, one of the most promising platforms for quantum simulation and computation, and compute its ground-state phase diagram, predicting unexplored phases of matter and quantum phase transitions as they would be observed in real experiments.

Afterwards, the aTTN is described in more technical detail, providing a comprehensive overview of our novel Tensor Network method. In particular, this addendum explains how the structure of the aTTN compensates the lack in faithfully capturing the area law of the TTN in high-dimensional systems while maintaining its main advantages: *(i)* the low scaling with the bond dimension m compared to both MERA and PEPS, and *(ii)* the ability to contract the network exactly. Further, the section provides a detailed description of the optimisation technique of the aTTN used for the ground-state search of quantum many-body systems and additionally address the procedure on how to efficiently compute observables. Finally, it compares the aTTN in more detail with the TTN in an extensive study on the critical two-dimensional Heisenberg model that *(i)* demonstrates the higher precision of the aTTN in describing two-body correlations $\langle \sigma_{i,j}^{\gamma} \sigma_{i',j'}^{\gamma} \rangle$, *(ii)* provides an analysis on the engineering of the disentangler positions including their effects on the numerical results, and *(iii)* illustrates the more favourable CPU time of the aTTN.

5.1 Efficient Tensor Network ansatz for high-dimensional quantum many-body problems

Efficient Tensor Network *Ansatz* for High-Dimensional Quantum Many-Body Problems

Timo Felser[1,2,3] Simone Notarnicola,[2,3,4] and Simone Montangero[2,3,4]

[1]*Theoretische Physik, Universität des Saarlandes, D-66123 Saarbrücken, Germany*
[2]*Dipartimento di Fisica e Astronomia "G. Galilei", Università di Padova, I-35131 Padova, Italy*
[3]*INFN, Sezione di Padova, Via Marzolo 8, I-35131 Padova, Italy*
[4]*Padua Quantum Technology Research Center, Università di Padova, I-35131 Padova, Italy*

We introduce a novel tensor network structure augmenting the well-established tree tensor network representation of a quantum many-body wave function. The new structure satisfies the area law in high dimensions remaining efficiently manipulatable and scalable. We benchmark this novel approach against paradigmatic two-dimensional spin models demonstrating unprecedented precision and system sizes. Finally, we compute the ground state phase diagram of two-dimensional lattice Rydberg atoms in optical tweezers observing nontrivial phases and quantum phase transitions, providing realistic benchmarks for current and future two-dimensional quantum simulations.

Recent experiments investigated one- and two-dimensional lattice quantum many-body systems at unprecedented sizes, calling for a continuous search of numerical techniques to provide accurate benchmarking and verification of future quantum simulations [1–9]. In particular, Rydberg atoms in optical tweezers are one of the most promising platforms for the study of quantum phase transitions, quantum simulation and computation [10–19]. In the last decades, Monte Carlo and Tensor Networks (TN) algorithms have been employed widely to study quantum many-body systems, and they are routinely used to benchmark quantum simulation results [20–30]. However, Monte Carlo methods are limited by the sign problem [31], while combining accuracy and scalability in simulating high-dimensional systems still represent an open challenge for TN methods [32,33]. Here, we introduce a novel TN variational *Ansatz*, able to encode the area law of quantum many-body states in any spatial dimension by keeping a low algorithmic complexity with respect to standard algorithms (see Fig. 1), thus opening a pathway towards the application of TN to high-dimensional systems. Hereafter, we benchmark this approach against spin models up to sizes of $N = 64 \times 64$, in and out of criticality. Finally, we simulate 2D Rydberg-atom lattices at sizes of up to ~ 1000, demonstrating the ability of providing the missing benchmarks for very recent quantum simulation experiments [34–37]: nontrivial phase transitions are characterized, in agreement with those experimentally observed in Ref. [34].

In the last three decades, TN have been developed and applied to classically simulate quantum many-body systems, representing the exponentially large wave function with a set of local tensors connected via auxiliary indices with a *bond-dimension m*. The bond dimension m allows

us to control the amount of information in the TN, interpolating between mean field ($m = 1$) and the exact but inefficient representation. While for one-dimensional (1D) systems the matrix product states (MPS) are the established TN geometry for equilibrium and out-of-equilibrium problems with open boundary conditions, the development of TN algorithms for 2D or 3D systems
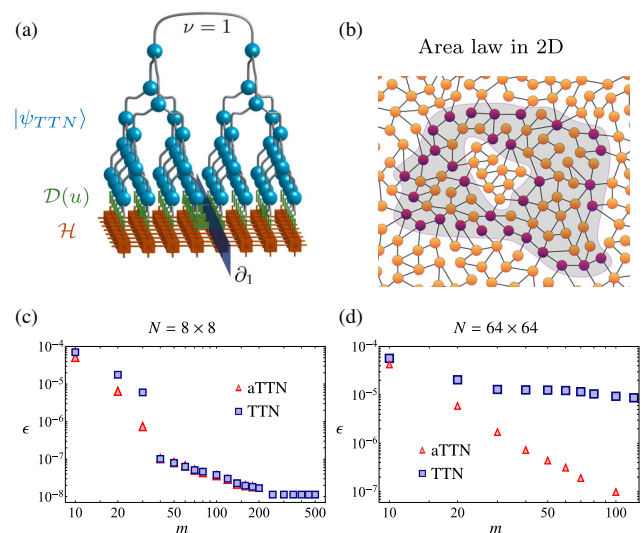

FIG. 1. (a) An aTTN for a 8×8 2D system: The disentanglers in $\mathcal{D}(u)$ are applied to the TTN state $|\psi_{\text{TTN}}\rangle$ across the boundaries ∂_ν of each link ν, in order to fulfill the area law depicted (b) for a sublattice $\mathcal{A}$ (shaded region) and its boundary $\partial\mathcal{A}$ (purple dots). (c),(d) Relative error of the Ising model ground state energy computed with the aTTNs and the TTNs. While for $L = 8$ the precision achieved with the two methods is the same, a clear improvement emerges for $L = 64$.

is still ongoing [38–43]. The most successful TN representations are the projected entangled pair states (PEPS) [44–47] and the tree tensor networks (TTN) [48–51], as well as the multiscale entanglement renormalization *Ansatz* (MERA) [52–54]. The PEPS flourishes for (infinite) 2D systems with open boundary condition and small physical dimension d and MERA provides an efficient representation in critical systems. The TTN offers a very flexible geometry which has proven to be a valid alternative with its particular strong points ranging from applications in gapped 1D systems with periodic boundary conditions [28,50], 2D systems with large local dimension d [24,51] to 3D systems [55].

TNs shall satisfy entanglement bounds under real-space bipartitions, known as area laws, of the *physical states* they represent [39,45,56]. The PEPS is the potentially most powerful TN *Ansatz* and by construction satisfies the area laws of entanglement [56]. However, it suffers from a high algorithmical complexity [typically $\mathcal{O}(m^{10})$ for finite-sized PEPS [46,57,58]] and lacks an exact calculation of expectation values. Indeed, the exact contraction a finite square lattice of the complete PEPS scales exponentially with the system linear dimension L and sophisticated numerical methods shall be introduced to mitigate this unfavorable scaling [58–63]. On the contrary, the MERA in two dimensions is able to calculate expectation values exactly while satisfying area law but suffers from an even higher algorithmical complexity [at least $\mathcal{O}(m^{16})$] [56]. Another well-established approach is to extend the MPS for 2D systems [64]: this approach has a very low algorithmic complexity [$\mathcal{O}(m^3)$], however, it is limited by an exponential scaling of the required bond dimension $m \sim e^{L_k}$ with the system minimal linear size $L_k \equiv \min\{L_x, L_y\}$. As a compromise, TTNs are equally scalable in both system dimensions while still benefiting from a low numerical complexity, $\mathcal{O}(m^4)$ but may fail to satisfy the area law for large systems sizes in higher dimensions [32,65].

Hereafter, we introduce a novel *Ansatz* which augments the TTN and show that is able to encode the area law keeping constant the algorithmic complexity to [$\mathcal{O}(m^4)$] in any physical dimension. As numerically demonstrated hereafter, the augmented tree tensor network (aTTN) allows us to efficiently tackle open challenges in two- and three-dimensional systems at sizes inaccessible before.

Augmented tree tensor network.—The aTTN *Ansatz* $|\psi_{\text{aTTN}}\rangle = \mathcal{D}^\dagger(u)|\psi_{\text{TTN}}\rangle$ is based on a TTN wave function (a binary tree) $|\psi_{\text{TTN}}\rangle \in \otimes_i^N \mathcal{H}_i$, with $\mathcal{H}_i = \mathbb{C}^d$, with an additional sparse layer $\mathcal{D}(u) = \prod_k u_k$ of two-site unitary operators $\{u_k\}$ acting on (some of) the physical links of the TTN (see Fig. 1). The additional layer $\mathcal{D}(u)$ contains N_D independent nonoverlapping (i.e., acting on different couples of sites of the lattice $\mathcal{L}$) and thus commuting disentanglers $\{u_k\}$. In this way, $\mathcal{D}(u)$ describes a unitary mapping of the Hamiltonian $\mathcal{H}$ to an auxiliary Hamiltonian $\mathcal{H}_{\text{aux}} = \mathcal{D}(u)\mathcal{H}\mathcal{D}^\dagger(u)$. Each local

transformation u_k aims to decouple—or *disentangle* in the spirit of the MERA language [54]—entangled degrees of freedom in the quantum many-body state, that are then trivially included in the TTN layer. As described in the following, $\mathcal{D}(u)$ modifies the TTN in such a way that the aTTN satisfies the area law while keeping the complexity for the optimisation at $\mathcal{O}(m^4)$. Thus, the aTTN overcomes the drawback of the TTN while maintaining its main advantages: (i) the low scaling with the bond dimension m compared to both MERA and PEPS, and (ii) the ability to contract the network exactly. We stress that the aTTN can be applied straighforwardly to a general D-dimensional system. Finally, we notice that the aTTN is effectively a particular subclass of a MERA, where the structure scale invariance is traded for efficiency, as the scale invariance is not necessary to ensure the area law at the tensor structure level. Figure 1(a) reports an illustrative example of an aTTN for a two-dimensional 8×8 system with the $\mathcal{D}(u)$ layer composed by 6 disentanglers u_k (green). Notice that not every physical site j is addressed by a disentangler, a key property for preserving numerical efficiency. Indeed, the disentangler positioning is critical in order to (i) keep an optimal numerical complexity for the optimization and (ii) efficiently encode an area law in the TN.

Area law in aTTN.— Hereafter, we specialize the discussion for the case of a two-dimensional square lattice $\mathcal{L}$ with $N = L \times L$ sites, and $L = 2^n$. Moreover, we consider a binary TTN, where the tree tensors coarse grain neighboring sites for each layer Λ_l alternately along the x (for even l) and the y direction (odd l) with l going from $l = 1$ addressing the topmost layer to $l = \log L$ for the lowest layer [Fig. 1(a)]. Each link ν of the tree bipartites the whole system $\mathcal{L}$ into two subsystems $\mathcal{A}^{[\nu]}$ and $\mathcal{B}^{[\nu]}$, separated by the boundary ∂_ν with length γ_ν. The area law implies that the entanglement entropy of the bipartition $S(\mathcal{A}^{[\nu]})$ (or $\mathcal{B}^{[\nu]}$, respectively) scales with γ_ν. Thus, in order to faithfully represent the area law, the bond dimension m_ν of each link ν should scale with $m_\nu \approx e^{c\gamma_\nu}$, where c is a constant factor. This scaling argument implies that for two dimensions the TTN *Ansatz* requires an exponentially large bond dimension m within the topmost layers, for which $\gamma_\nu \sim L$. In conclusion, with increasing L a TTN representation eventually fails to capture area law states' properties as it becomes exponentially inefficient. This necessary exponential scaling of the bond dimension can be prevented by inserting the tensors layer $\mathcal{D}(u)$ that augments the TTN with $N_D = \sum_\nu K_\nu$ disentanglers, where K_ν is the number of disentanglers along the boundary ∂_ν for each link ν. More precisely, each disentangler u_k is positioned such that it acts on one physical site in the subsystems $\mathcal{A}^{[\nu]}$ and the other in subsystems $\mathcal{B}^{[\nu]}$. Thus, each disentangler can maximally assess information in a d^2-dimensional space belonging to two local Hilbert spaces, reducing the entanglement for the TTN up to the order of d^2. As a result, all the K_ν disentanglers support the TTN link ν by

disentangling information on the order of $m_{\nu,\text{aux}} \approx (d^2)^{K_\nu}$. Therefore, when applying $\mathcal{D}(u)$, the information assessed by the aTTN for the bipartition defined by each link ν scales with $m_{\nu,\text{eff}} \approx m_{\nu,\text{aux}} m_\nu = d^{2K_\nu + \xi_\nu}$, where we introduced the parameter $\xi_\nu \equiv \log_d m_\nu$ describing the contribution of the TTN bond dimension m_ν. If we now impose $K_\nu \sim \gamma_\nu$, we obtain the exponential scaling required to encode the area law for the two-dimensional aTTN state. Notice that the number of disentanglers $\Gamma_l = \sum_{\nu \in \Lambda_l} K_\nu$ for each layer shall be directly proportional to $\sum_{\nu \in \Lambda_l} \gamma_\nu \sim L$. However, placing exactly L disentanglers for each layer of the tree may lead to an unfavorable, L-dependent scaling of $\mathcal{O}(m^4 d^L)$ for the computational complexity. Thus, a careful balance between the position of the disentanglers and their density has to be found. This balance can be found as when no couple of disentanglers is directly connected by a Hamiltonian interaction term, the algorithmic scaling remains of the order $\mathcal{O}(m^4 d^2)$. Moreover, the area law is still satisfied, removing the disentanglers crossing the boundaries of the bipartitions ∂_ν corresponding to the lower layers of the tree ($l \to \log L$). On the contrary, one shall keep the maximal allowed number of disentanglers (i.e., not connected by Hamiltonian terms) to support the boundaries corresponding to the higher branches ($l \to 1$). Indeed, for $\nu \in \Lambda_l$ with $l \to \log L$, the TTN bond-dimension m_ν is sufficiently large to capture the area law entanglement—or even the complete state—accurately, especially for reasonably small local dimensions d. Instead, the contribution ξ_ν of the TTN is negligibly small for $\nu \in \Lambda_l$ with $l \to 1$ compared to the required exponentially large bond dimension, calling for the support of the disentanglers.

In conclusion, different disentangler configurations $\mathcal{D}(u)$ exist, matching the aforementioned criteria, computational efficiency and the area law. In our numerical simulations, the final resulting precision was not significantly affected by the particularly chosen configuration (for details see Supplemental Material [66]).

Ising model.—We first benchmark the aTTN *Ansatz* against the ordinary TTN via a ground state search on the ferromagnetic 2D Ising model with periodic boundary conditions (BC). We consider a $L \times L$ lattice with $L = \{8, 16, 32, 64\}$ and the Ising Hamiltonian $\mathcal{H} = \sum_{i,j=1}^{L} \sigma_{i,j}^x \sigma_{i+1,j}^x + \sigma_{i,j}^x \sigma_{i,j+1}^x + \sum_{i,j=1}^{L} \sigma_{i,j}^z$, where $\sigma_{i,j}^\gamma$ (with $\gamma \in \{x, y, z\}$) denote the are Pauli matrices acting on the site (i, j). For small system sizes ($L = 8$ and $L = 16$) both the TTN and the aTTN reach the chosen machine precision of `1E-8` with high bond dimension. However, as expected, for larger sizes we find a significant improvement in the precision of the aTTN simulations. Indeed, the different performances become evident for $L = 32$ and $L = 64$, as the aTTN and the TTN converge with increasing bond dimension to different values for the energy Fig. 1 reports the relative error $\epsilon_m = |(\langle \mathcal{H} \rangle_m - E_{ex})/E_{ex}|$ for increasing bond dimension m with respect to the energy E_{ex} obtained by extrapolating the results of the aTTN for
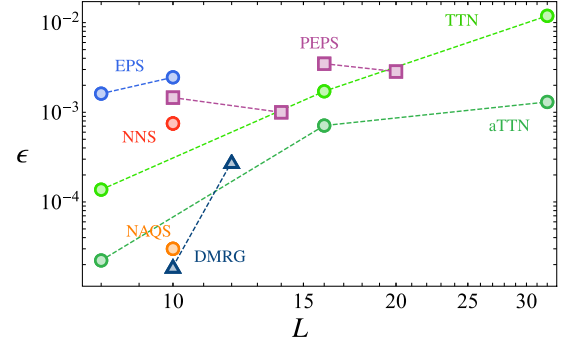

FIG. 2. Relative error ϵ of the 2D Heisenberg ground-state energy as a function of the system linear size L for the TTN, aTTN, NNS [68], NAQS [69], EPS [70], PEPS [58], 2D-DMRG [64] (circles, squares and triangles indicate open BC, periodic BC and cylindrical BC, respectively) each compared with the best available estimates obtained by MC with the same BC (for pbc [22], for obc [58] obtained via ALPS library [75–77]).

$L = 8$ and $L = 64$ (For the $L = 16, 32$ results see Fig. 3 in the Supplemental Material [66]).

Heisenberg model.—We now analyze the more challenging critical antiferromagnetic two dimensional Heisenberg model $\mathcal{H} = \sum_{i,j=1}^{L} \sum_{\gamma \in \{x,y,z\}} \sigma_{i,j}^\gamma \sigma_{i+1,j}^\gamma + \sigma_{i,j}^\gamma \sigma_{i,j+1}^\gamma$, with periodic BC. In Fig. 2 we compare the estimated energy
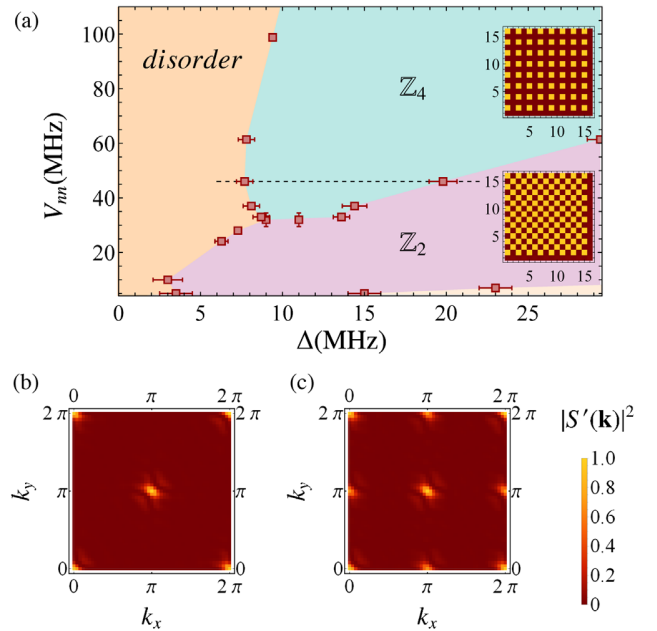

FIG. 3. Up: Phase diagram as a function of the detuning Δ and the nearest-neighbors interaction energy V_{nn}. The disordered phase is characterized by a substantially uniform distribution of the excitations, while in the phases $\mathbb{Z}_2$ and $\mathbb{Z}_4$ the excitations are distributed as shown in the upper ($\mathbb{Z}_4$) and lower ($\mathbb{Z}_2$) insets. Down: Renormalized structure factor $S'(\mathbf{k}) = S(\mathbf{k})/S(\mathbf{0})$ for $V_{nn} = 46$ MHz and (a) $\Delta = 28$ MHz ($\mathbb{Z}_2$ phase) and (b) $\Delta = 12$ MHz ($\mathbb{Z}_4$ phase). Other parameters: $\Omega = 4$ MHz.

density obtained by extrapolating the results from the TTN and the aTTN at $m \to \infty$ with previous results from different variational *Ansätze* obtained by means of ordinary cluster resources [67]. In particular, we plot the relative error obtained by the different tensor network *Ansätze* and the best known results, obtained via quantum Monte Carlo simulations [22]. Differently from the Ising model, we find the aTTN to be more accurate than the TTN even at lower system sizes, such as $L = 8$, 16. Interestingly, the aTTN for $L = 16$ obtains an even more precise ground state energy density compared to most of the alternative variational *Ansätze* at lower finite system size of $L = 10$, such as neural network states (NNS), neural autoregressive quantum state, entangled plaquette states (EPS), or PEPS [58,68–70]. We mention that, while the PEPS is very efficient with its ability to work directly in the thermodynamic limit in describing infinite systems as iPEPS [71,72], the PEPS analysis for finite sizes are, for now, limited to $N = 20 \times 20$ systems. It turns out that for this model a very competitive variational approach is the 2D-DMRG, which outperforms the alternative methods for finite sizes with open or cylindrical BC up to the system size $L = 12$, but struggles with periodic BC and with increasing both system sizes $L \gtrsim 12$ [64]. Finally, we extended our analysis to $L = 32$: In this case no public result is available for periodic BC and thus we estimated the error by extrapolating the value of the finite size scaling of Monte Carlo simulations [22]. In recent works, heavy parallelization over large high performance computing systems have been exploited to study the Heisenberg model with open BC combining the PEPS structure with MC techniques, reaching a precision of $\sim 10^{-4}$ at 32×32 [73,74].

We point out that the here performed aTTN simulations (as well as the TTN simulations) exploit a $U(1)$ symmetry. However, for this model, we could further drastically improve the performance of the aTTN by incorporating the present SU(2) symmetry in the simulation framework [32,78,79].

Interacting Rydberg atoms.—We now present new physical results, on a long-range interacting system by studying the zero-temperature phase diagram of an interacting Rydberg atoms two-dimensional lattice [4], described by the Hamiltonian $\mathcal{H}_{\mathrm{ryd}} = \sum_{\mathbf{r}} [(\Omega/2)\sigma_{\mathbf{r}}^x - \Delta n_{\mathbf{r}} + \frac{1}{2}\sum_{\mathbf{s}} V(|\mathbf{r} - \mathbf{s}|) n_{\mathbf{r}} n_{\mathbf{s}}]$, where the Rabi frequency Ω couples the ground $|g\rangle_{\mathbf{r}}$ and the excited Ryderg state $|r\rangle_{\mathbf{r}}$ and $n_{\mathbf{r}} = |r\rangle\langle r|_{\mathbf{r}}$. Δ is the detuning and $V(|\mathbf{r} - \mathbf{s}|) = c_6/|\mathbf{r} - \mathbf{s}|^6$ is the interaction strength between two excited atoms placed at sites $\mathbf{r}$ and $\mathbf{s}$. We keep the interaction terms up to the fourth-nearest neighbor and set the Rabi frequency $\Omega = 4 \mathrm{MHz}$, while the interaction parameters refer to $^{87}\mathrm{Rb}$ atoms excited to the state $|70S_{1/2}\rangle$, for which $c_6 = 863 \, \mathrm{GHz} \, \mu m^6$.

The interactions limit the maximum excitation density according to the Rydberg blockade radius r^*—the minimum distance at which two atoms can be simultaneously excited—defined by the relation $V(r^*) = \Omega$. The

competition between the interactions strength and Δ generates nontrivial phases characterized by regular spatial excitation-density distributions. Figure 3(a) shows the phase diagram of the system as a function of the detuning and the nearest-neighbor interaction energy V_{nn}, obtained via aTTN simulations with $L = 4$, 8, 16, 32 with open BC.

For low values of the detuning Δ, the system exhibits a disordered phase characterized by the absence of excitations while, increasing Δ, excitations are energetically favored and the interactions determine their spatial arrangement. In the limit of $V_{nn} \to 0$, or $a \to \infty$, the atoms are noninteracting and the expectation value $\langle n_{\mathbf{r}} \rangle \to 1$ for $\Delta \ll \Omega$. At larger values of V_{nn}, corresponding to $r^*/\sqrt{2} < a < r^*$, nearest neighbor atoms cannot be simultaneously excited, giving rise to the $\mathbb{Z}_2$ phase [4,80] with a two-degenerate ground state with the excitations distributed in a chess board like configuration, as shown in Fig. 3(a). Nevertheless, the $\mathbb{Z}_2$ disappears at low values of V_{nn} and large detuning, as all the atoms are excited (light orange, right-bottom region of the phase diagram). The spatial distribution of the excitations in the orderded phase is well captured by the peaks of the static structure factor $S(\mathbf{k}) = (1/N^2) \sum_{\mathbf{r},\mathbf{s}} e^{-i\mathbf{k}\cdot(\mathbf{r}-\mathbf{s})} \langle n_{\mathbf{r}} n_{\mathbf{s}} \rangle$. In particular, the phase $\mathbb{Z}_2$ exhibits a peak in (π, π), as shown in Fig. 3(b). The transition from the disordered to the $\mathbb{Z}_2$ phase is a second-order one, as it emerges by computing the second derivative of the energy with respect to Δ (see Supplemental Material [66]). In order to determine the critical line separating the two phases we define the non-local order parameter $O_{\mathbf{r}}^{(2)} = (n_{r_x,r_y} - n_{r_x+1,r_y} - n_{r_x,r_y+1} + n_{r_x+1,r_y+1})/4$ and perform a finite-size scaling analysis of $\langle O_{\mathbf{r}}^{(2)\dagger} O_{\mathbf{r}}^{(2)} \rangle$ vs Δ, where $\langle O_{\mathbf{r}}^{(2)\dagger} O_{\mathbf{r}}^{(2)} \rangle$ is estimated by $S(\pi, \pi)$ [81,82] (see Supplemental Material [66]). By further reducing a, the blockade radius prevents diagonal-adjacent atoms to be excited. As a consequence, each one of the $\mathbb{Z}_2$ ground states breaks into two different states, giving rise to the four-degenerate phase $\mathbb{Z}_4$: In each one of the ground states of this phase, each excited atom is surrounded by atoms in their ground states [see upper inset in Fig. 3(a)]. We observe a second-order phase transition in V_{nn} for $\Delta \simeq 10$ MHz from the $\mathbb{Z}_2$ to the $\mathbb{Z}_4$ phase at $V_{nn}^c = 32 \pm 2.5$ MHz (or equivalently $a = r^*/\sqrt{2}$). The static structure factor exhibits four additional peaks in the points such as $(0, \pi)$ as shown in Fig. 3(c). As in the $\mathbb{Z}_2$ case, a second-order phase transition occurs between the disordered phase to the $\mathbb{Z}_4$ by changing Δ at a fixed V_{nn}. We determine the critical line by introducing the order parameter $O_{\mathbf{r}}^{(4)} = (n_{r_x,r_y} + i n_{r_x+1,r_y} - i n_{r_x,r_y+1} - n_{r_x+1,r_y+1})/4$, defined such that the value of $\langle O_{\mathbf{r}}^{(4)\dagger} O_{\mathbf{r}}^{(4)} \rangle$ equals $S(0, \pi)$ in the $\mathbb{Z}_4$ phase. Remarkably, we find that another second-order phase transition occurs by further increasing Δ, leading the system from the $\mathbb{Z}_4$ to the $\mathbb{Z}_2$ phase. We expect that at

larger values of V_{nn} new phases would emerge and accordingly, new phase transitions would occur by changing Δ.

Conclusions.—We have augmented the well-established TTN geometry with a new *Ansatz* which reproduces area law for high dimensional quantum many-body systems. The efficiency of aTTNs allowed us to reach large sizes (32×32) in the study of critical system, going beyond the current possibilities of standard PEPS and DMRG, and therefore set new benchmarks for future numerical simulations [83]. As a first application of aTTNs, we have characterized the phase diagram of two-dimensional Rydberg atoms in optical tweezers, with atoms number of the order of current and near future experiments [2]. Further applications include the study of systems which cannot be studied via Monte Carlo simulations due to the sign problem [31,84], such as Abelian and non-Abelian lattice gauge theories at finite densities [24,55,85–88]. Such an application enables the study of the continuum limit at higher dimensions, paving the way to unveil novel insights into our understanding of the fundamental constituents of our universe [89,90]. Finally, the aTTN *Ansatz* can support known TN algorithms [28,91–96] to investigate nonequilibrium dynamics in open and closed high-dimensional systems, including annealing, quenches, or controlled dynamics.

In conclusion, the aTTN *Ansatz* introduced here provides a novel powerful tool for simulating quantum systems in two or higher dimensions, which, beyond many interesting physical applications will provide benchmark near-future quantum simulations and computations on different platforms, as we have demonstrated for Rydberg atoms in optical tweezers.

[1] S. de Léséleuc, V. Lienhard, P. Scholl, D. Barredo, S. Weber, N. Lang, H. P. Büchler, T. Lahaye, and A. Browaeys, Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms, Science **365,** 775 (2019).

[2] A. Browaeys and T. Lahaye, Many-body physics with individually controlled Rydberg atoms, Nat. Phys. **16,** 132 (2020).

[3] J. Rui, D. Wei, A. Rubio-Abadal, S. Hollerith, J. Zeiher, D. M. Stamper-Kurn, C. Gross, and I. Bloch, A subradiant optical mirror formed by a single structured atomic layer, Nature (London) **583,** 369 (2020).

[4] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, Nature (London) **551,** 579 (2017).

[5] P. Schauß, M. Cheneau, M. Endres, T. Fukuhara, S. Hild, A. Omran, T. Pohl, C. Gross, S. Kuhr, and I. Bloch, Observation of spatially ordered structures in a two-dimensional Rydberg gas, Nature (London) **491,** 87 (2012).

[6] A. Elben, B. Vermersch, R. van Bijnen, C. Kokail, T. Brydges, C. Maier, M. K. Joshi, R. Blatt, C. F. Roos, and P. Zoller, Cross-Platform Verification of Intermediate Scale Quantum Devices, Phys. Rev. Lett. **124,** 010504 (2020).

[7] T. Olsacher, L. Postler, P. Schindler, T. Monz, P. Zoller, and L. M. Sieberer, Scalable and parallel tweezer gates for quantum computing with long ion strings, PRX Quantum **1,** 020316 (2020).

[8] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, Self-verifying variational quantum simulation of lattice models, Nature (London) **569,** 355 (2019).

[9] A. Celi, B. Vermersch, O. Viyuela, H. Pichler, M. D. Lukin, and P. Zoller, Emerging Two-Dimensional Gauge Theories in Rydberg Configurable Arrays, Phys. Rev. X **10,** 021057 (2020).

[10] H. Labuhn, D. Barredo, S. Ravets, S. de Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models, Nature (London) **534,** 667 (2016).

[11] P. Schauss, Quantum simulation of transverse Ising models with Rydberg atoms, Quantum Sci. Technol. **3,** 023001 (2018).

[12] S. Notarnicola, M. Collura, and S. Montangero, Real-time-dynamics quantum simulation of (1+1)-dimensional lattice QED with Rydberg atoms, Phys. Rev. Research **2,** 013288 (2020).

[13] F. M. Surace, P. P. Mazza, G. Giudici, A. Lerose, A. Gambassi, and M. Dalmonte, Lattice Gauge Theories and String Dynamics in Rydberg Atom Quantum Simulators, Phys. Rev. X **10,** 021041 (2020).

[14] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, Synthetic three-dimensional atomic structures assembled atom by atom, Nature (London) **561,** 79 (2018).

[15] D. Barredo, V. Lienhard, P. Scholl, S. de Léséleuc, T. Boulier, A. Browaeys, and T. Lahaye, Three-Dimensional Trapping of Individual Rydberg Atoms in Ponderomotive Bottle Beam Traps, Phys. Rev. Lett. **124,** 023201 (2020).

[16] V. Lienhard, S. de Léséleuc, D. Barredo, T. Lahaye, A. Browaeys, M. Schuler, L.-P. Henry, and A. M. Läuchli, Observing the Space- and Time-Dependent Growth of Correlations in Dynamically Tuned Synthetic Ising Models with Antiferromagnetic Interactions, Phys. Rev. X **8,** 021070 (2018).

[17] A. Omran *et al.*, Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, Science **365,** 570 (2019).

[18] H. Levine *et al.*, Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms, Phys. Rev. Lett. **123,** 170503 (2019).

[19] M. Morgado and S. Whitlock, Quantum simulation and computing with Rydberg-interacting qubits, arXiv:2011.03031.

[20] J. Kolorenč and L. Mitas, Applications of quantum Monte Carlo methods in condensed systems, Rep. Prog. Phys. **74**, 026502 (2011).

[21] P. H. Acioli, Review of quantum Monte Carlo methods and their applications, J. Mol. Struct. Theochem **394**, 75 (1997).

[22] A. W. Sandvik, Finite-size scaling of the ground-state parameters of the two-dimensional Heisenberg model, Phys. Rev. B **56**, 11678 (1997).

[23] S. Montangero, *Introduction to Tensor Network Methods* (Springer International Publishing, Cham, 2018).

[24] T. Felser, P. Silvi, M. Collura, and S. Montangero, Two-dimensional quantum-link lattice Quantum Electrodynamics at finite density, arXiv:1911.09693.

[25] R. Orús, Tensor networks for complex quantum systems, Nat. Rev. Phys. **1**, 538 (2019).

[26] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).

[27] Bañuls *et al.*, Simulating lattice gauge theories within quantum technologies, Eur. Phys. J. D **74**, 165 (2020).

[28] L. Kohn, P. Silvi, M. Gerster, M. Keck, R. Fazio, G. E. Santoro, and S. Montangero, Superfluid-to-Mott transition in a Bose-Hubbard ring: Persistent currents and defect formation, Phys. Rev. A **101**, 023617 (2020).

[29] E. Macaluso, T. Comparin, R. O. Umucalilar, M. Gerster, S. Montangero, M. Rizzi, and I. Carusotto, Charge and statistics of lattice quasiholes from density measurements: A tree tensor network study, Phys. Rev. Research **2**, 013145 (2020).

[30] P. Emonts, M. C. Bañuls, I. Cirac, and E. Zohar, Variational Monte Carlo simulation with tensor networks of a pure $\mathbb{Z}_3$ gauge theory in (2+1)D, Phys. Rev. D **102**, 074501 (2020).

[31] M. Troyer and U.-J. Wiese, Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations, Phys. Rev. Lett. **94**, 170201 (2005).

[32] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, The Tensor Networks Anthology: Simulation techniques for many-body quantum lattice systems, SciPost Phys. Lect. Notes **8**, 3 (2019).

[33] S.-J. Ran, E. Tirrito, C. Peng, X. Chen, L. Tagliacozzo, G. Su, and M. Lewenstein, *Two-dimensional tensor networks and contraction algorithms, in Tensor Network Contractions: Methods and Applications to Quantum Many-Body Systems* (Springer International Publishing, Cham, 2020), pp. 63–86.

[34] S. Ebadi *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, arXiv:2012.12281.

[35] D. Bluvstein *et al.*, Controlling many-body dynamics with driven quantum scars in Rydberg atom arrays, arXiv:2012.12276.

[36] P. Scholl *et al.*, Programmable quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, arXiv:2012.12268.

[37] R. Samajdar, W. W. Ho, H. Pichler, M. D. Lukin, and S. Sachdev, Complex Density Wave Orders and Quantum Phase Transitions in a Model of Square-Lattice Rydberg Atom Arrays, Phys. Rev. Lett. **124**, 103601 (2020).

[38] S. Östlund and S. Rommer, Thermodynamic Limit of Density Matrix Renormalization, Phys. Rev. Lett. **75**, 3537 (1995).

[39] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, Phys. Rev. B **73**, 094423 (2006).

[40] S. R. White, Density matrix formulation for quantum renormalization groups, Phys. Rev. Lett. **69**, 2863 (1992).

[41] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, Ann. Phys. (Amsterdam) **326**, 96 (2011).

[42] S. R. White and A. E. Feiguin, Real-Time Evolution Using the Density Matrix Renormalization Group, Phys. Rev. Lett. **93**, 076401 (2004).

[43] G. Vidal, Efficient Classical Simulation of Slightly Entangled Quantum Computations, Phys. Rev. Lett. **91**, 147902 (2003).

[44] F. Verstraete and J. I. Cirac, Valence-bond states for quantum computation, Phys. Rev. A **70**, 060302(R) (2004).

[45] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of Projected Entangled Pair States, Phys. Rev. Lett. **96**, 220601 (2006).

[46] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, Ann. Phys. (Amsterdam) **349**, 117 (2014).

[47] F. Verstraete and J. I. Cirac, Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions, arXiv:cond-mat/0407066.

[48] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, Phys. Rev. A **74**, 022320 (2006).

[49] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio, Homogeneous binary trees as ground states of quantum critical Hamiltonians, Phys. Rev. A **81**, 062335 (2010).

[50] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero, Unconstrained tree tensor network: An adaptive gauge picture for enhanced performance, Phys. Rev. B **90**, 125154 (2014).

[51] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte, and S. Montangero, Fractional quantum Hall effect in the interacting Hofstadter model via tensor networks, Phys. Rev. B **96**, 195123 (2017).

[52] G. Vidal, Entanglement Renormalization, Phys. Rev. Lett. **99**, 220405 (2007).

[53] G. Evenbly and G. Vidal, Entanglement Renormalization in Two Spatial Dimensions, Phys. Rev. Lett. **102**, 180406 (2009).

[54] G. Evenbly and G. Vidal, Algorithms for entanglement renormalization, Phys. Rev. B **79**, 144108 (2009).

[55] G. Magnifico, T. Felser, P. Silvi, and S. Montangero, Lattice Quantum Electrodynamics in (3+1)-dimensions at finite density with Tensor Networks, arXiv:2011.10658.

[56] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, Rev. Mod. Phys. **82**, 277 (2010).

[57] P. Czarnik and J. Dziarmaga, Variational approach to projected entangled pair states at finite temperature, Phys. Rev. B **92**, 035152 (2015).

[58] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, Phys. Rev. B **90,** 064425 (2014).

[59] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, Contracting projected entangled pair states is average-case hard, Phys. Rev. Research **2,** 013010 (2020).

[60] S.-J. Dong, C. Wang, Y. Han, G.-c. Guo, and L. He, Gradient optimization of fermionic projected entangled pair states on directed lattices, Phys. Rev. B **99,** 195153 (2019).

[61] H. Ueda, K. Okunishi, S. Yunoki, and T. Nishino, Corner transfer matrix renormalization group analysis of the two-dimensional dodecahedron model, Phys. Rev. E **102,** 032130 (2020).

[62] T. Nishino and K. Okunishi, Corner transfer matrix renormalization group method, J. Phys. Soc. Jpn. **65,** 891 (1996).

[63] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, Phys. Rev. Lett. **98,** 140506 (2007).

[64] E. Stoudenmire and S. R. White, Studying two-dimensional systems with the density matrix renormalization group, Annu. Rev. Condens. Matter Phys. **3,** 111 (2012).

[65] A. J. Ferris, Area law and real-space renormalization, Phys. Rev. B **87,** 125139 (2013).

[66] See Supplemental Material at http://link.aps.org/supplemental/ 10.1103/PhysRevLett.126.170603 for a more technical description of the aTTN Ansatz and a deeper analysis of the 2D quantum Ising model, critical Heisemberg and Rydberg atom lattice simulations.

[67] The simulations presented here have been performed on a single cluster node.

[68] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, Science **355,** 602 (2017).

[69] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Deep Autoregressive Models for the Efficient Variational Simulation of Many-Body Quantum Systems, Phys. Rev. Lett. **124,** 020503 (2020).

[70] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, Ground-state properties of quantum many-body systems: entangled-plaquette states and variational Monte Carlo, New J. Phys. **11,** 083026 (2009).

[71] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions, Phys. Rev. Lett. **101,** 250602 (2008).

[72] R. Orús and G. Vidal, Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction, Phys. Rev. B **80,** 094403 (2009).

[73] L. He, H. An, C. Yang, F. Wang, J. Chen, C. Wang, W. Liang, S. Dong, Q. Sun, W. Han, W. Liu, Y. Han, and W. Yao, PEPS++: Towards extreme-scale simulations of strongly correlated quantum many-particle models on sunway taihulight, IEEE Trans. Parallel Distributed Syst. **29,** 2838 (2018).

[74] W.-Y. Liu, Y.-Z. Huang, S.-S. Gong, and Z.-C. Gu, Accurate simulation for finite projected entangled pair states in two dimensions, arXiv:1908.09359.

[75] S. Todo and K. Kato, Cluster Algorithms for General-*S* Quantum Spin Systems, Phys. Rev. Lett. **87,** 047203 (2001).

[76] A. Albuquerque, F. Alet, P. Corboz, P. Dayal, A. Feiguin, S. Fuchs, L. Gamper, E. Gull, S. Gürtler, A. Honecker *et al.,* The ALPS project release 1.3: Open-source software for strongly correlated systems, J. Magn. Magn. Mater. **310,** 1187 (2007).

[77] B. Bauer *et al.,* The ALPS project release 2.0: open source software for strongly correlated systems, J. Stat. Mech. (2011) P05001.

[78] S. Singh and G. Vidal, Global symmetries in tensor network states: Symmetric tensors versus minimal bond dimension, Phys. Rev. B **88,** 115147 (2013).

[79] S. Singh, Tensor network states and algorithms in the presence of abelian and non-abelian symmetries, arXiv:1203.2222.

[80] P. Schauß, J. Zeiher, T. Fukuhara, S. Hild, M. Cheneau, T. Macrì, T. Pohl, I. Bloch, and C. Gross, Crystallization in Ising quantum magnets, Science **347,** 1455 (2015).

[81] K. Binder, Finite size scaling analysis of ising model block distribution functions, Z. Phys. B **43,** 119 (1981).

[82] G. Giudici, A. Angelone, G. Magnifico, Z. Zeng, G. Giudice, T. Mendes-Santos, and M. Dalmonte, Diagnosing Potts criticality and two-stage melting in one-dimensional hard-core boson models, Phys. Rev. B **99,** 094434 (2019).

[83] For a discussion concerning the possibility to improve the efficiency our simulations see the Supplemental Material [66] and Refs. [97,98].

[84] M. C. Bañuls and K. Cichy, Review on novel methods for lattice gauge theories, Rep. Prog. Phys. **83,** 024401 (2020).

[85] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. V. Acoleyen, F. Verstraete, U. J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, Simulating lattice gauge theories within quantum technologies, arXiv:1911.00003.

[86] M. Dalmonte and S. Montangero, Lattice gauge theory simulations in the quantum information era, Contemp. Phys. **57,** 388 (2016).

[87] T. M. R. Byrnes, P. Sriganesh, R. J. Bursill, and C. J. Hamer, Density matrix renormalization group approach to the massive Schwinger model, Phys. Rev. D **66,** 013002 (2002).

[88] P. Silvi, E. Rico, T. Calarco, and S. Montangero, Lattice gauge tensor networks, New J. Phys. **16,** 103015 (2014).

[89] I. J. R. Aitchison and A. J. G. Hey, *Gauge Theories in Particle Physics: A Practical Introduction. Vol. 1: From Relativistic Quantum Mechanics to QED* (CRC Press, Bristol, UK, 2003).

[90] T. Cheng and L. Li, *Gauge Theory of Elementary Particle Physics* (Oxford University Press, New York, 2006).

[91] M. Rizzi, S. Montangero, and G. Vidal, Simulation of time evolution with multiscale entanglement renormalization ansatz, Phys. Rev. A **77,** 052328 (2008).

[92] L. Arceci, P. Silvi, and S. Montangero, Entanglement of formation of mixed many-body quantum states via Tree Tensor Operators, arXiv:2011.01247.

[93] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-Dependent Variational

Principle for Quantum Lattices, Phys. Rev. Lett. **107,** 070601 (2011).

[94] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, Phys. Rev. B **94,** 165116 (2016).

[95] C. Lubich, I. V. Oseledets, and B. Vandereycken, Time integration of tensor trains, SIAM J. Numer. Anal. **53,** 917 (2015).

[96] D. Bauernfeind and M. Aichhorn, Time dependent variational principle for tree Tensor Networks, SciPost Phys. **8,** 24 (2020).

[97] N. Ossi, L. Bitton, D. B. Gutman, and A. Frydman, Zero-bias anomaly in a two-dimensional granular insulator, Phys. Rev. B **87,** 115137 (2013).

[98] A. Milsted, M. Ganahl, S. Leichenauer, J. Hidary, and G. Vidal, TensorNetwork on TensorFlow: A spin chain application using tree Tensor Networks, arXiv:1905.01331.

5.2 Technical description of an aTTN

In this section, we describe the aTTN in more technical detail, providing a comprehensive overview of our novel Tensor Network method. We begin by explaining the structure of the aTTN in more detail and how it compensates the drawbacks of the TTN while maintaining its main advantages: *(i)* the low scaling with the bond dimension m compared to both MERA and PEPS, and *(ii)* the ability to contract the network exactly. Further, we provide a detailed description on the optimisation technique of the aTTN used for the ground-state search of quantum many-body systems and additionally address the procedure on how to efficiently compute observables. Finally, we compare the aTTN in more detail with the TTN in an extensive study on the critical two-dimensional Heisenberg model, demonstrating among others the higher precision of the aTTN in describing two-body correlations $\langle \sigma_{i,j}^{\gamma} \sigma_{i',j'}^{\gamma} \rangle$, an analysis on the engineering of the disentangler positions including their effects on the numerical results, and the more favourable CPU time of the aTTN.

Since the following builds on the presented publication and aims to provide a deeper understanding, this chapter has to be understood in this context only and includes selected parts and material of the original publication without explicit citation.

5.2.1 aTTN geometry

The augmented Tree Tensor Network (aTTN) represents a wave-function $|\Psi\rangle \in \mathscr{H}$ on a lattice $\mathscr{L}$ with the Hilbert space $\mathscr{H}$. In general, the lattice $\mathscr{L}$ can be an D-dimensional lattice containing $N = \prod_{i=1}^{D} L_i$ sites, where each site $j \in \mathscr{L}$ is described by a local Hilbert space $\mathscr{H}_j$ with finite dimension d_j, so that its complete Hilbert space is spanned by $\mathscr{H} = \otimes_j \mathscr{H}_j$.

The geometry of the aTTN is based on a TTN wave-function with an additional layer of disentanglers $\mathcal{D}(u)$ attached to the outgoing links on the bottom of the TTN. Thus the aTTN representation of a pure state $|\psi\rangle \in \otimes_i^N \mathscr{H}_i$ on the lattice $\mathscr{L}$ is given by

$$|\psi_{\mathrm{aTTN}}\rangle = \mathcal{D}^{\dagger}(u)|\psi_{\mathrm{TTN}}\rangle \tag{5.1}$$

with $|\psi_{\mathrm{TTN}}\rangle$ describing the wave-function parametrised by the internal TTN. Therein, the appended layer $\mathcal{D}(u)$ contains N_D disentanglers $\{u_k\}$, which all act independently of each other on different sites of the lattice $\mathscr{L}$. Each of the disentanglers u_k is a unitary when fusing its first two and its last two indices respectively, thus obeying the isometry condition

$$\sum_{k_3,k_4} (u_k)_{k_3,k_4}^{k_1,k_2} (u_k^{\dagger})_{k_1',k_2'}^{k_3,k_4} = \delta_{k_1,k_1'} \delta_{k_2,k_2'}. \tag{5.2}$$

Hence, one disentangler u_k performs a unitary transformation on two physical sites $(i_1^{[k]}, i_2^{[k]})$ towards the attached TTN. This local transformation aims to decouple - or *disentangle* - relevant degrees of freedom in the quantum many-body state which consequently disappear for the TTN. Thus, the complete layer $\mathcal{D}(u)$ maps a pure state ψ of the lattice $\mathscr{L}$ to another pure state ψ_{aux} within the same Hilbert space $\mathscr{H}$ by applying all of its disentanglers u_k:

$$\mathcal{D}(u): \ \mathscr{H} \to \mathscr{H} \tag{5.3}$$

$$\mathcal{D}(u)|\psi\rangle = u_1 u_2 ... u_K |\psi\rangle = |\psi_{\mathrm{aux}}\rangle \tag{5.4}$$

Note, that the different disentanglers u_k commute with each other, as they all act on different spaces $\mathscr{H}_{k_1} \otimes \mathscr{H}_{k_2}$. In this manner, $\mathcal{D}(u)$ can be as well seen as a unitary mapping for a

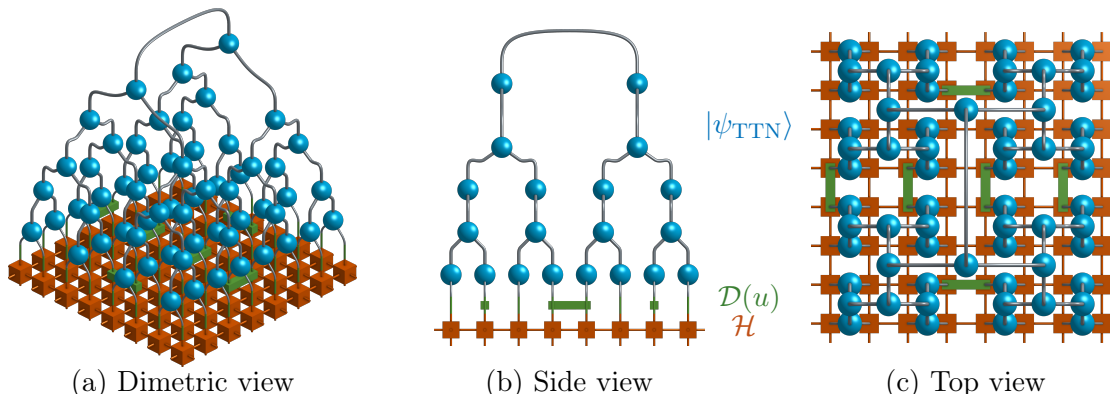(a) Dimetric view   (b) Side view   (c) Top view

FIGURE 5.1: Example of an aTTN for a 2D system with $N = 8 \times 8$ sites in different views (a)-(c). The aTTN contains an internal TTN ψ_{TTN} (illustrated with tensors in blue and links in gray) with a layer of disentanglers $\mathcal{D}(u)$ (green) attached to the outgoing links of the TTN. For illustration, we indicated a Hamiltonian $\mathcal{H}$ (orange) acting on the physical sites of the aTTN. (c) The disentangler layer $\mathcal{D}(u)$ supports the most critical links of the internal TTN with a number of disentanglers K_ν connecting the bipartitions introduced by the link ν. To encode the area law, K_ν scales with the boundary length γ_ν.

given physical Hamiltonian $\mathcal{H} \in \mathscr{H}$ to an auxiliary Hamiltonian $\mathcal{H}_{\mathrm{aux}} = \mathcal{D}(u)\mathcal{H}\mathcal{D}^\dagger(u)$ towards the TTN within the aTTN. This preconditioning of the Hamiltonian $\mathcal{H}$ for the internal TTN can be performed in a way, such that it introduces an area law for higher-dimensional systems of the complete network while keeping the complexity for the optimisation at $\mathcal{O}\left(m^4\right)$. Thus the aTTN avoids the weakness of a TTN - being the lack of an area law - and still maintains its main advantages, namely *(i)* the reasonably low scaling with bond dimension m compared to both MERA and PEPS, and *(ii)* the ability to contract the network exactly, which in general is not guaranteed for a PEPS (see also Tab. 5.1 for comparison). Let us point out as well, that the aTTN is not restricted to a certain dimensionality of the underlying system, but can be applied for a general D-dimensional system.

TABLE 5.1: Comparison of the most prominent Tensor Networks discussed in the main text: Numerical complexity as a function of the bond-dimension m, obeying the entanglement area law, the typical bond-dimension to be used in current high performance simulations and the calculation of expectation values, i.e. the exact contractability.

Tensor Network	Complexity	Area law in 2D	Typical Bond dimensions	Exact contractable
MPS / DMRG	$\mathcal{O}\left\{m^3\right\}$	No (Only in 1D)	> 10.000	Yes ($\mathcal{O}\left\{m^3\right\}$)
TTN	$\mathcal{O}\left\{m^4\right\}$	No (Only in 1D)	$\approx 1.000 - 2.000$	Yes ($\mathcal{O}\left\{m^4\right\}$)
PEPS	$\mathcal{O}\left\{m^{10}\right\}$	Yes	~ 10	No ($\mathcal{O}\left\{m^L\right\}$)
MERA	$\mathcal{O}\left\{m^8\right\}$ (1D), $\mathcal{O}\left\{m^{16}\right\}$ (2D)	Yes	~ 10	Yes ($\mathcal{O}\left\{m^8\right\}$)
aTTN	$\mathcal{O}\left\{m^4\right\}$	Yes*	≈ 500	Yes ($\mathcal{O}\left\{m^4\right\}$)

In Fig. 5.1 we give an illustrative example of an aTTN for a two-dimensional 8×8 system with its disentangler layer $\mathcal{D}(u)$ consisting of 6 different disentanglers u_k (green). As shown therein, not every physical site j is addressed by a disentangler, which - as we will explain later on - is key for a better numerical complexity. Thus the positioning of the disentanglers u_k for a general aTTN is critical in order to *(i)* keep an optimal numerical complexity for the optimisation and *(ii)* efficiently encode an area law in the Tensor Network.

Resuming, the total number of parameters for an aTTN state scales with $\mathcal{O}(Nm^3 + N_D d^4)$ where m is the bond dimension of the TTN within, and N_D is the number of disentanglers u_k in the disentangler layer $\mathcal{D}(u)$.

5.2.2 Area Law in aTTN

For the description of the area-law captured by the aTTN, we illustrate the case of a two-dimensional square lattice $\mathscr{L}$ with $N = L \times L$ sites, where furthermore $L = 2^n$. The fundamental idea anyhow holds true for an D-dimensional lattice structure with arbitrary dimensions L_i. Furthermore, we here assume a binary TTN, which is arranged so that the tensors within the tree alternatingly in x- and y-direction coarse-grain neighboring sites going from layer to layer, as it is the case for the internal TTN in Fig. 5.1 (the internal TTN is illustrated by its tensors in blue and its links in gray). The topmost link of such a tree bipartites the whole system $\mathscr{L}$ into the two equally $(L \times L/2)$-sized subsystems $\mathcal{A}$ and $\mathcal{B}$. Going from the topmost link downwards in the TTN, each link within a layer further divides a smaller $l_x \times l_y$-dimensional sublattice.

	$\gamma_{\nu[1]}$	$\gamma_{\nu[2]}$		$\gamma_{\nu[3]}$		$\gamma_{\nu[4]}$	$\cdots$
Open BC	L	L		$\frac{5}{4}L$ or $\frac{3}{4}L$			
cylindrical BC	L	$\frac{3}{2}L$		$\frac{3}{2}L$ or L			
toroidal BC	$2L$	$\xrightarrow{\quad}$ $2L$	$\xrightarrow{\cdot 3/4}$	$\frac{3}{2}L$	$\xrightarrow{\cdot 2/3}$	$\frac{3}{2}L$	$\xrightarrow{\cdot 3/4}$

TABLE 5.2: Size of boundaries of different bipartitions in a 2D system $\mathscr{L}$ for different boundary conditions of $\mathscr{L}$

When we now consider an area-law state ψ, the bipartition entanglement scales with the boundary ∂_ν of the divided subsystem, as stated in Eq. (3.4). In Tab. 5.2 this area is presented depending on the boundary conditions of $\mathscr{L}$. Therein, we introduced the notation $\gamma_{\nu[k]}$ for the boundary of a subsystem $\mathcal{A}^{[\nu]}$ emerging from the partitions introduced by a link ν within the k-th layer of the TTN (with subsystem size $N/2^k$). As mentioned above in Eq. 3.5, the TTN itself would require an in γ_ν exponentially large bond dimension m_ν in order to capture the area-law entanglement for the state ψ. In order to prevent this scaling, we now place K_ν disentanglers on the bottom of the TTN connecting the subsystems $\mathcal{A}^{[\nu]}$ and $\mathcal{B}^{[\nu]}$ given by the bipartition introduced by the link ν. More precise, we position the disentanglers u_k such that one of its physical sites $i_1^{[k]}$ or $i_2^{[k]}$ belongs to the subsystem $\mathcal{A}^{[\nu]}$, while the other one corresponds to $\mathcal{B}^{[\nu]}$. Each one of the disentanglers can maximally asses information in a d^2-dimensional space belonging to two local Hilbert spaces. Thus it can reduce the entanglement for the TTN up to the order of d^2. Consequently, all the K_ν disentanglers together can support the TTN link ν by disentangling information in the order of

$$m_{\nu,\text{aux}} \approx (d^2)^{K_\nu} \tag{5.5}$$

When we now scale the number of disentanglers K_ν according to the boundary γ_l of the bipartition introduced by the link ν, the information captured within the complete aTTN for the cut through the network introduced by ν scales with

$$m_{\nu,\text{eff}} \approx m_{\text{aux}} m_\nu = d^{2K_\nu + \xi_\nu} \tag{5.6}$$

Thus, the proper positioning of the disentanglers results in an exponential scaling of the effective bond dimension where we defined the parameter $\xi_n u \equiv \log_d m_\nu$ reflecting the contribution of the TTN bond dimension m_ν. In order to encode the area law for the two-dimensional aTTN state according to Eq. (3.5), the number of disentanglers K_ν connecting the different bipartitions introduced by the links ν has to be directly proportional to the boundaries γ_ν and thereby for $\nu = 1$ to the length L of the system.

In practice we know, that ξ_ν in Eq. (5.6) is sufficiently large for the lower branches of the tree in order to capture the area law entanglement - or even the complete state - accurately, especially for reasonably small local dimensions d. Thus we waive disentanglers within the smaller plaquettes of the TTN. Following the same idea, we focus on supporting the links higher up in the TTN with disentanglers $K_\nu \sim \gamma_\nu$, as the contribution of ξ_ν compared to the required exponentially large bond dimension becomes negligibly small. Equivalently, while going to lower layers, the contribution of ξ_ν weights increasingly stronger and thus we do not need to strictly enforce $K_\nu \sim \gamma_\nu$ in order to capture area law within the complete network.

5.2.3   Connection to other Tensor Networks

In the prior sections, we described the connection of the aTTN to the well-established TTN geometry in detail. In what follows, we will elaborate on the similarity of the aTTN to the MERA and the connections to a PEPS.

MERA —   The aTTN can be seen as a particular subclass of a MERA, where we sacrifice the disentanglers at higher renormalisation levels, and thereby the scale-invariance, for a higher numerical efficiency. In particular, the scale-invariance is not strictly necessary to ensure the area law. Indeed, in the aTTN we shift the problem of encoding the area law to a non-trivial positioning problem of the disentanglers at the bottom of the aTTN which we can solve with several approaches as discussed later on. Thus, the aTTN does not necessarily have the predefined, straight-forward scale-invariance structure of the aTTN but, with a clever positioning of the disentanglers, it is indeed able to encode area law with a leading complexity of $\mathcal{O}\left(m^4\right)$ instead of $\mathcal{O}\left(m^{16}\right)$ in case of the MERA.

PEPS —   It has already been shown that the MERA can directly be mapped to a PEPS [193]. Following this procedure, and keeping in mind that the aTTN is a subclass of the MERA, we in principle can map the aTTN to a PEPS with in general unisotropic bond dimensions. In fact, the bond-dimensions of the resulting PEPS would highly depend on the positioning of the disentanglers and the internal TTN bond-dimensions indicating that this interplay between disentanglers at the bottom and the internal TTN is the crucial key to the success of the aTTN. We mention, however, that performing this mapping numerically might be not beneficial from a computational point of view.

5.3   Optimisation of the aTTN

For the optimisation of the aTTN, we assume the complete Hamiltonian to be a product of interactions $\mathcal{H} = \sum_p \mathcal{H}_p$. Thus every interaction $\mathcal{H}_p$ can be described as a Tensor Product Operator (TPO). For sake of simplicity, we will restrict ourselves to Hamiltonians $\mathcal{H}$ containing exclusively *(i)* local terms $\mathcal{H}_p = h_{i_p}^p$ (acting on the site i_p) and *(ii)* two-body interaction $\mathcal{H}_p = h_{i_p}^{p,1} h_{i'_p}^{p,2}$ between the physical sites i_p and i'_p. Note, that the basic idea of the aTTN

approach can, in general, be applied to more complex Hamiltonians $\mathcal{H}$, but depending on the Tensor Network representation of $\mathcal{H}$ the computational complexity in secondary terms may increase. Furthermore the numerical implementation becomes more challenging as well when going beyond two-site interactions.

Given the Hamiltonian $\mathcal{H}$ we optimise the variational parameters of the aTTN wavefunction ψ in order to find the ground state of the system by minimising the energy

$$E = \langle\psi|\mathcal{H}|\psi\rangle \ . \tag{5.7}$$

This optimisation procedure for the aTTN consists of three different parts: *(i)* The optimisation of the disentanglers u_k, *(ii)* the mapping of the Hamiltonian $\mathcal{H} \xrightarrow{\mathcal{D}(u)} \mathcal{H}_{\mathrm{aux}}$, and *(iii)* the optimisation of the internal TTN with the auxiliary Hamiltonian $\mathcal{H}_{\mathrm{aux}}$.

Below we will describe each part of the optimisation in more detail, followed by some general remarks towards the practical application of the optimisation.

Disentangler Layer $\mathcal{D}(u)$

We optimise the layer $\mathcal{D}(u)$ of an aTTN by optimising all of the disentanglers forming $\mathcal{D}(u)$ one-by-one individually. In Fig. 5.2 we illustrate the procedure for one disentangler u_k to optimise. The fundamental idea hereby origins from the general MERA optimisation [74] with some minor adaptions for the aTTN geometry. Thus for the disentangler u_k the energy $E = \langle\Psi|H|\Psi\rangle$ to minimise depends bilinearly on u_k and $u_k^\dagger$,

$$E(u_k) = \underbrace{\mathrm{tr}\left\{\sum_p u_k M_p u_k^\dagger N_p\right\}}_{\equiv E_k(u_k)} + c_k \ , \tag{5.8}$$

where M_p and N_p are two sets of matrices corresponding to different contractions of the environment of u_k and $u_k^\dagger$ with the Hamiltonian part $\mathcal{H}_p$. Note that p only runs over the Hamiltonian parts $\mathcal{H}_p$ which act on one or both sites of the disentangler to optimise. All the other Hamiltonian parts $\mathcal{H}_{p'}$ contribute to the overall energy independently of the disentangler u_k and thus are all included in the constant c_k for the Energy $E(u_k)$. Consequently, for the optimisation of u_k, we minimise the first part $E_k(u_k)$ of Eq. (5.8) only and neglect the constant contributions in c_k. The sum over the different Hamiltonian parts p thereby is illustrated in Fig. 5.2 in terms of the Tensor Network representation. We point out, that in Eq. (5.8), the disentangler u_k and its complex conjugate $u_k^\dagger$ are reshaped as unitary matrices by fusing their legs respectively.

In order to contract the complete network towards the matrices M_p and N_p, we can exploit the isometry of the internal TTN. As introduced for the MERA structure [74], the positions of the disentangler u_k and the interaction part $\mathcal{H}_p$ for the contraction determine a *causal cone* of tensors. We hereby define the *anchor* of a causal cone as the topmost tensor (with respect to the hierarchical TTN structure), which is still included in the causal cone. Thus, isometrising the internal TTN towards the anchor node of the causal cone for u_k and $\mathcal{H}_p$ results in all tensors outside of the causal cone vanishing to identities due to their enforced isometry. Consequently, the complete contraction of the network can be reduced to the tensors within the causal cone only.

Anyhow, after contracting, the environment resulting in the matrices M_p and N_p, there is no generic algorithm that solves the bilinear problem of Eq. (5.8) while additionally fulfilling
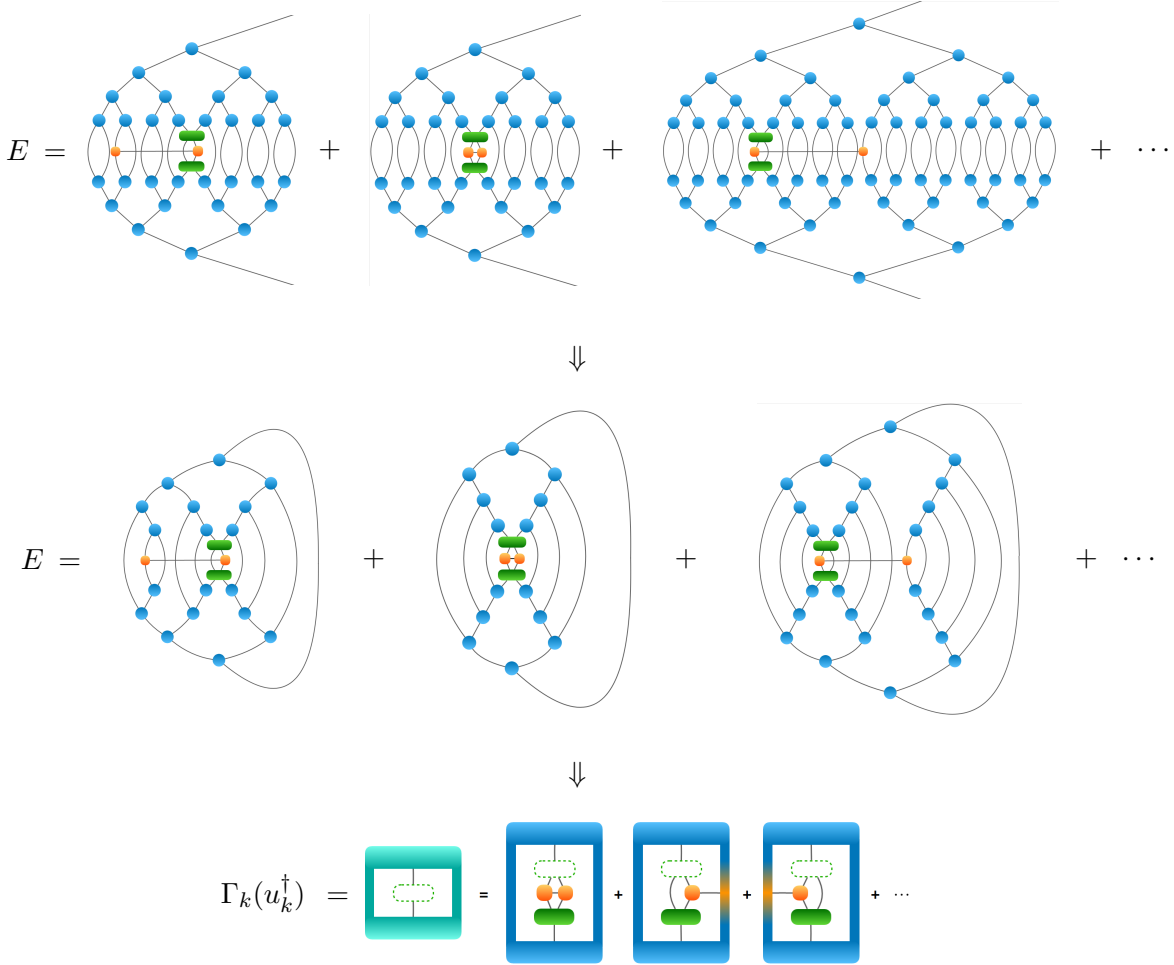
FIGURE 5.2:   Optimisation procedure for one disentangler u_k (green) of the aTTN. [top] Energy expectation value $E(u_k)$ of the aTTN. Each Hamiltonian part $\mathcal{H}_p$ (orange) is contracted with the complete aTTN and adding up to $E(u_k)$. For the optimisation of u_k only the parts $\mathcal{H}_p$ attached to u_k have to be considered. [middle] Contraction for $E(u_k)$ when isometrising the internal TTN towards the anchor node of the causal cone [74]. All tensors outside annihilate to identities. [bottom] Environment Γ_k for u_k including all the relevant contributions $\mathcal{H}_p$. By construction $E = \mathrm{tr}\left\{u_k\Gamma_k\right\} + c_k$ holds true. By decomposing $\Gamma_k = U\sigma V^\dagger$, we optimise the disentangler u_k with $u_k \equiv -VU^\dagger$.

the isometry constraint in Eq. (5.2). Thus, we solve this optimisation problem iteratively by linearising the function for the Energy $E_k(u_k)$ with respect to the disentangler u_k following the idea in Ref. [74]. Thereby, u_k and $u_k^\dagger$ are temporarily considered as independent tensors within one iteration step. This allows to optimise the Energy $E_k(u_k)$ with respect to u_k while keeping its complex conjugate $u_k^\dagger$ constant. For this linearised problem, the energy functional now reads

$$E_k(u_k) \;=\; \mathrm{tr}\left\{u_k\Gamma_k(u_k^\dagger)\right\}\,, \tag{5.9}$$

$$\text{with} \qquad \Gamma_k(u_k^\dagger) = \sum_p M_p u_k^\dagger N_p \tag{5.10}$$

where Γ_k is a matrix obtained by the contraction of the complete environment around

the disentangler u_k, as it is illustrated in Fig. 5.2. Thus, each iteration step, we optimise $E_k(u_k)$ with respect to u_k, while including $u_k^\dagger$ in the environment Γ_k. Thus, assuming Γ_k to be independent of u_k, the energy $E_k(u_k)$ is minimised by setting $u_k = -VU^\dagger$ where U and V are obtained via singular value decomposition of $\Gamma_k = U\sigma V^\dagger$. Thus by construction u'_k obeys the isometry constraint of Eq. (5.2) as both U and V are unitary matrices. The minimised energy for this iteration step becomes

$$\tilde{E}_{k,min} = \text{tr}\,\{u_k\Gamma_k\} = \text{tr}\left\{-VU^\dagger U\sigma V^\dagger\right\} = -\sum_i \sigma_i \tag{5.11}$$

Note, that the total energy is $E(u_k) = E_{k,min} + c_k$, as we neglected the constant c_k for the minimisation.

Concluding one iteration step, the complete procedure will be repeated with the new disentangler u'_k, until the singular values σ_i of the SVD convergence which equals a convergence in the minimisation of the Energy to E_{min}.

Summarising the optimisation of the disentangler u_k:

(i) Contract the environment matrices M_p and N_p for all Hamiltonian parts $\mathcal{H}_p$ addressed by the disentangler u_k

(ii) Compute the environment Γ_k of the u_k for one iteration

(iii) Decompose $\Gamma_k = U\sigma V^\dagger$ via SVD

(iv) Update the disentangler $u_k \rightarrow u'_k = -VU^\dagger$

(v) Restart from (ii) until all σ_i converge.

Once the disentangler u_k is optimised we move on the next disentangler, until all the disentanglers within the layer $\mathcal{D}(u)$ have been optimised. Due to the disentanglers being positioned in a way, that they don't share interaction parts $\mathcal{H}_p$ of the physical Hamiltonian $\mathcal{H}$, the optimisation of each disentangler u_k is completely independent of all the other disentanglers $u_{k'}$. Thus it is sufficient to optimise each disentangler just once to obtain the optimised disentangler layer $\mathcal{D}(u)$ - and furthermore, all optimisations for each disentangler can be fully parallalised.

The complete optimisation of one disentangler can be done with the complexity $\mathcal{O}(m^4d^2 + m^3d^4 + d^6)$, where the contractions for part (i) of the summary above scale with $\mathcal{O}(m^4d^2 + m^3d^4)$, for (ii) with $\mathcal{O}(d^6)$, and the decomposition in (iii) as well with $\mathcal{O}(d^6)$. Note that the most expensive part (i) can be done once for the complete optimisation of one disentangler u_k and is not required to be recomputed during the iterative procedure.

Hamiltonian Mapping

After each optimisation of the disentangler layer $\mathcal{D}(u)$, the physical Hamiltonian $\mathcal{H}$ mapped by the newly optimised disentangler layer $\mathcal{D}(u)$ for the subsequent TTN optimisation. For the aTTN, we consider the systems Hamiltonian $\mathcal{H} \in \mathscr{H}$ to be a product of interactions $\mathcal{H} = \sum_p \mathcal{H}_p$. As mentioned above, the aTTN introduces a layer of disentanglers $\mathcal{D}(u)$ which
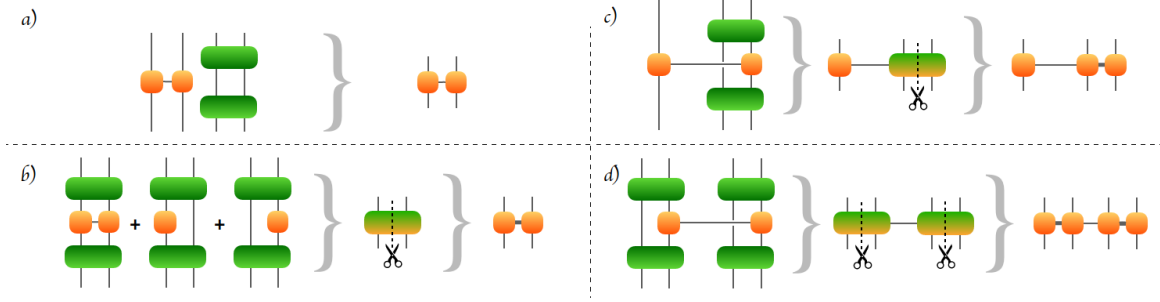
FIGURE 5.3:   (a) Trivial mapping of an operator that is not attached to a disentangler. (b) Mapping of several Hamiltonian parts $\mathcal{H}_p$ acting exclusively on the subspace addressed by the disentangler u_k. After contracting each operator $\mathcal{H}_p$ individually, the resulting tensors act on the same space and can be summed up. The obtained tensor is decomposed concluding the mapping. (c) Mapping of an operator with one site addressed by the disentangler u_k and one site free of disentanglers. After contracting the operator with u_k and $u_k^\dagger$, the contracted tensor is decomposed resulting in the mapped TPO. (d) Mapping of an operator with both site addressed by different disentangler u_k and u_k'. After contracting both tensors, the contracted tensors are decomposed resulting in the mapped TPO.

maps $\mathcal{H}$ to an auxiliary Hamiltonian $\mathcal{H}_{\mathrm{aux}} \equiv \mathcal{D}(u)\mathcal{H}\mathcal{D}^\dagger(u)$ as preconditioning for the TTN. The energy expectation value

$$\langle\psi_{\mathrm{aTTN}}|\mathcal{H}|\psi_{\mathrm{aTTN}}\rangle = \langle\psi_{\mathrm{TTN}}|\underbrace{\mathcal{D}(u)\mathcal{H}\mathcal{D}^\dagger(u)}_{\equiv\mathcal{H}_{\mathrm{aux}}}|\psi_{\mathrm{TTN}}\rangle \qquad (5.12)$$

equals the expectation value of the mapped Hamiltonian for the internal TTN wavefunction. This transformation of the complete Hamiltonian $\mathcal{H} \xrightarrow{\mathcal{D}(u)} \mathcal{H}_{\mathrm{aux}}$ is done by contracting each interaction part $\mathcal{H}_p$ separately with $\mathcal{D}(u)$.

In what follows, we illustrate this mapping for a general operator $\mathcal{T}$, which we describe as a Tensor Product Operator (TPO)

$$(\mathcal{T})_{\{i_j'\}}^{\{i_j\}} = \sum_{\{\gamma_\nu\}}\prod_j (\mathrm{t}^{[j]})_{i_j,i_j'}^{\{\gamma_{\nu'}\}} \qquad (5.13)$$

following the definition in Ref. [24]. Thereby, the j-th tensor $\mathrm{t}^{[j]}$ is acting locally on the site i_j and is connected by the links $\{\gamma_{\nu'}\}$ to other tensors within the TPO. Thus, the complete TPO acts on the physical sites $\{i_j\}$. In this TPO formalism, we can describe for instance a local observable $\mathcal{T}_i$, an interaction part $\mathcal{H}_p$ of the Hamiltonian $\mathcal{H}$, a string observable $\mathcal{T}_{\{i_j\}}$ or more general structures like an MPO.

Consequently, the mapping $\mathcal{T} \xrightarrow{\mathcal{D}(u)} \mathcal{T}_{aux}$ is done by contracting the TPO $\mathcal{T}$ with the proper disentanglers of layer $\mathcal{D}(u)$. This mapping can transform the operator $\mathcal{T}$ according to four different cases depicted in Fig. 5.3.

a) Trivial mapping —   Let there be no disentangler u_k in $\mathcal{D}(u)$ applied to neither one of the physical sites $\{i_j\}$ of the operator $\mathcal{T}$ ($\nexists u_k : i_k \in \{i_j\} \vee i_{k'} \in \{i_j\}$). Therefore, $\mathcal{T}$ is unaffected by the disentangler layer $\mathcal{D}(u)$ and the mapping effectively becomes an identity resutling in $\mathcal{T}_{\mathrm{aux}} = \mathcal{T}$.

b) Mapping to two-site TPO — Let the disentangler u_k be applied to all the sites $\{i_j\}$ on which the operator $\mathcal{T}$ is acting ($\forall j : i_j \in \{k_1, k_2\}$). In this case, $\mathcal{T} \equiv \mathcal{T}_i$ is either a local operator acting on the single site $i_1 \in \{k_1, k_2\}$ only or a two-site operator $\mathcal{T} \equiv \mathcal{T}_{k_1,k_2}$ acting on the exact same sites $\{k_1, k_2\}$ addressed by the disentangler u_k.

For the first step of the mapping, we contract the TPO with u_k and $u_k^\dagger$ to

$$(\mathcal{T}_{\text{contr}})^{k_1,k_2}_{k_1',k_2'} = \sum_{i_1,i_2,i_1',i_2'} (u_k)^{k_1,k_2}_{i_1,i_2} (\mathcal{T})^{\{i_j\}}_{\{i_j'\}} (u_k^\dagger)^{i_1',i_2'}_{k_1',k_2'} . \tag{5.14}$$

Afterwards, we decompose the resulting tensor $\mathcal{T}_{\text{contr}}$ via singular value decomposition obtaining two tensors - one for each site addressed by u_k - forming the mapped TPO $\mathcal{T}_{\text{aux}}$. Note, that if the original TPO was a two-site operator with an internal bond dimension κ, this mapping may lead to a different bond dimension for $\kappa' \neq \kappa$ within the mapped operator $\mathcal{T}_{\text{aux}}$. Anyhow, this new bond dimension is restricted by $\kappa' \leq d^2$.

Let us mention as well, that in case of mapping the Hamiltonian $\mathcal{H}$, we can have several interaction parts $\mathcal{H}_p$ which are all exclusively acting within the subspace $\mathscr{H}_{k_1} \otimes \mathscr{H}_{k_1}$ addressed by the disentangler u_k. In Fig. 5.3 we illustrate such an example with one two-site interaction and two local Hamiltonian parts, all to be contracted with the same disentangler u_k. Here it is worth to point out, that after the contraction of Eq. (5.14) all the interaction parts $(\mathcal{H}_p)_{\text{contr}}$ can be added up, as now all of them act as one tensor on the same subspace. The summed tensor will subsequently be decomposed resulting in the mapped two-site TPO for the auxiliary Hamiltonian $\mathcal{H}_{aux}$.

c) Mapping with one disentangler — Let there be only one disentangler u_k addressing at least one but not all of the sites $\{i_j\}$ ($(\exists i_k : i_k \in i_j) \wedge (\exists i_m \in i_j : i_m \notin \{i_k\})$). For the mapping we here proceed in a similar way as in the prior scenario.

First, we contract the operator $\mathcal{T}$ over the connected links with the diesentangler u_k and its c.c $u_k^\dagger$:

$$(\mathcal{T}_{\text{contr}})^{k_1,k_2,\{i_m\}}_{k_1,k_2,\{i_m'\}} = \sum_{i_1,i_2,i_1',i_2'} (u_k)^{k_1,k_2}_{i_1,i_2} (\mathcal{T})^{\{i_j\}}_{\{i_j'\}} (u_k^\dagger)^{i_1',i_2'}_{k_1',k_2'} . \tag{5.15}$$

After the contraction, we decompose the tensor $\mathcal{T}_{\text{contr}}$ in order to keep a TPO structure where each tensor within corresponds to one physical site. Thereby, we obtain the mapped TPO which now contains either the same number of tensors (when u_k is contracted with two tensors of the TPO) or one additional tensor (u_k addresses only one tensor of $\mathcal{T}$). Note, that here again the bond dimension for $\kappa' \neq \kappa$ within the TPO might increase but still is restricted by $\kappa' \leq d^2$ for H_p being a two-site TPO or to $\kappa' \leq \kappa d^2$ for some cases with larger TPO sizes respectively.

d) Mapping with two or more disentanglers — Let there be K disentanglers $\{u_k\}$ be applied for the sites $\{i_j\}$. In this most general case, the operator $\mathcal{T}$ becomes mapped highly non-trivial. Here, the contraction with the disentangler layer is done by contracting the operator $\mathcal{T}$ with all the K disentanglers attached:

$$(\mathcal{T}_{\text{contr}})^{\{i_m\}}_{\{i_m\}} = \sum u_{k_1} u_{k_2} ... u_{k_K} \mathcal{T}^{\{i_j\}}_{\{i_j'\}} u_{k_1}^\dagger ... u_{k_2}^\dagger u_{k_1}^\dagger . \tag{5.16}$$

For the sake of compactness, we dropped the indices for the disentanglers. Furthermore, we introduced the set of indices $\{i_m\}$ containing all the indices i_k of all K disentanglers $\{u_k\}$

and additionally all the physical legs i_j of the original operator $\mathcal{T}$ over which we did not contract. Note as well that the order of application of the disentanglers u_k is not important, as all u_k act on different subspaces of the complete Hilbert space $\mathscr{H}$ and thus commute with each other.

Here again the resulting operator $\mathcal{T}_{\mathrm{contr}}$ is decomposed back into a TPO structure with its tensors acting on one single site of the TTN. For a general TPO which may even contain loops in its structure, like e.g. a PEPO [149–151], this decomposition is a highly non-trivial task. Furthermore, for loopless TPOs, this mapping might again result in an increased internal bond dimension within the TPO.

Considering the Hamiltonian $\mathcal{H}_{\mathrm{aux}}$ is used to optimise the TTN, an arbitrary loopless interaction $\mathcal{H}_p$ which becomes mapped to a general MPO might in the worst case still increase the complexity of the overall optimisation to a with system size exponential scaling $\mathcal{O}(m^4 d^L)$. But if we restrict ourselves to all Hamiltonian parts $\mathcal{H}_p = h_{p,1} h_{p,2}$ being nearest-neighbor interactions only, the additional cost is restricted to a worst-case prefactor of d^4. In order to go even further and prevent this prefactor, we can constraint the aTTN in its construction to prevent an interaction part $\mathcal{H}_p$ being addressed by two - or more - disentanglers. Thus, for nearest-neighbor interactions, we strictly forbid to attach two disentanglers u_k and u'_k onto neighboring sites, as can be seen in Fig. 5.1. This not only decreases the second-order scaling of the complete algorithm but additionally simplifies the numerical implementation and introduces a higher potential of parallelisation, as now all the disentanglers can be treated completely independent within the optimisation.

Tree Tensor Network

The optimisation of the TTN follows the prescriptions of Chapt. 3. Each pair of connected tensors within the TTN is optimised via the subspace-expansion technique which approximates a two-site update [24]. This technique allows to keep the favourable scaling of $\mathcal{O}\left(m^4\right)$ compared to the higher numerical complexity of $\mathcal{O}\left(m^6\right)$ in case of the direct two-site optimisation (see Sec. 3.1.4). The local optimisations of each single tensor optimisation are performed via Arnoldi algorithm as discussed in Sec. 3.1.4 which solves the local eigenvalue problem by iteratively diagonalising the corresponding effective Hamiltonian H_{eff}, returning the lowest eigenvalues of H_{eff} within a predetermined accuracy ϵ.

We optimise the global TTN state by performing these local optimisations including the space-expansion technique step-by-step from the bottom of the TTN to its top. The space expansion, in particular, is always performed for the local target tensor together with its *parent* tensor as space expansion partner.

General remarks

In practice, we start the complete optimisation of the aTTN with a randomly initialised TTN. We first iterate a few times threw the TTN with the physical Hamiltonian $\mathcal{H}$ disregarding the disentangler layer $\mathcal{D}(u)$. In this part, we use a relatively large space expansion and a low precision for solving the local eigenvalue problems for each tensor. Thereby, we aim to efficiently *(i)* adapt the randomly initialised symmetry sectors within the TTN for a qualitative description of the ground state and *(ii)* use this qualitative TTN wavefunction for an advanced the initialisation of the aTTN. Consequently, we start with the aTTN consisting of the resulting TTN with a set of identities $\{u_k\} \equiv \{\mathbb{1}\}$ attached at the sites we aim to disentangle forming the initial disentangler layer $\mathcal{D}(u)$. We subsequently optimise the aTTN

by optimising the disentaglers once, mapping $\mathcal{H}$ to $\mathcal{H}_{\text{aux}}$ and performing n steps of optimising the TTN, where we usually choose $n \in \{1, 2, 3\}$.

Furthermore, in the results obtained for this paper, we improved the TTN optimisation by parallelising the contractions over different Hamiltonian parts $\mathcal{H}_p$ within the optimisation procedure. Additionally, we want to mention, that the computational time for optimisation of the disentanglers and the mapping of the Hamiltonian was negligible compared to the TTN optimisation in our practical applications. As of last we point out, that in the aTTN approach there is still more room for improvement by parallelisation, as the TTN optimisation can be further parallelised in a similar way as shown for parallel-DMRG [180] as well as the optimisation of the disentanglers.

5.4 Computation of observables

5.4.1 Local observables

Let $O^{[s]}$ be a local observable acting on the site s of the aTTN state ψ. In order to compute the expectation value $\langle O^{[s]} \rangle_\psi = \langle \psi | O^{[s]} | \psi \rangle$ we distinguish the two different cases:

(i) A disentangler u_k is placed at site s - thus connecting it with the site $\tilde{s}$, or

(ii) there is no disentangler attached to the physical site s.

In the latter case, the one-site operator $O^{[s]}$ is simply passed to the TTN and calculated by the standard computation of a local observable in TTN, as $O^{[s]}$ is unaffected by the disentangler layer $\mathcal{D}(u)$ and the isometry property $\mathcal{D}(u)\mathcal{D}^\dagger(u) = \mathbb{1}$ holds true. Using the proper isometry in the TTN, this calculation can be done by contracting three tensors only [24].

In case *(i)* however, the local operator becomes mapped to a two-site operator by the disentangler layer $\mathcal{D}(u)$, as it is contracted with u_k and $u_k^\dagger$ to

$$(O^{[s]}_{aux})^{k_1,k_2}_{k_1',k_2'} = \sum_{s,s',\tilde{s}} (u_k)^{k_1,k_2}_{s,\tilde{s}} (O^{[s]})^s_{s'} (u_k^\dagger)^{s',\tilde{s}}_{k_1',k_2'} \ . \tag{5.17}$$

Consequently, in this case, the computation of $\langle O^{[s]} \rangle_\psi$ effectively becomes a calculation of the expectation value of the correlator $\langle (O^{[s]}_{aux})^{k_1,k_2}_{k_1',k_2'} \rangle_\psi$ for the TTN. The worst case complexity of the underlying contraction in the TTN is $\mathcal{O}(m^4 \kappa \log N) + \mathcal{O}(m^4 \kappa \log N - 1) + \mathcal{O}(m^3 \kappa)$.

5.4.2 Correlators

Let $O^{[s_1,s_2]}$ be a correlator acting on the sites s and s' of the aTTN state ψ. In order to compute the expectation value $\langle O^{[s,s']} \rangle_\psi = \langle \psi | O^{[s,s']} | \psi \rangle$ we distinguish the following four different cases for the mapping $\mathcal{D}(u)$ - some of which are similar or even identical to the mapping described in section 5.3:

Trivial mapping Let there be no disentangler u_k in the layer $\mathcal{D}(u)$ attached to neither one of the physical sites s_1 or s_2 of the correlator. Therefore, the correlator $O^{[s]}$ can - as described for the local observable - simply passed threw to the TTN and calculated by the standard computation of a correlator in TTN, as here again $O^{[s_1,s_2]}$ is unaffected by the disentangler layer $\mathcal{D}(u)$ and the mapping effectively becomes an identity. This subsequent contraction of the TTN can be done with the worst-case complexity of $\mathcal{O}(m^4 \log N) + \mathcal{O}(m^4 \log N - 1) + \mathcal{O}(m^3)$.

Mapping to two-site TPO   Let the disentangler u_k connect both sites s_1 and s_2 on which the the correlator $O^{[s_1,s_2]}$ acts. In this case, the correlator becomes mapped to another two-site operator by the disentangler layer $\mathcal{D}(u)$, as it is contracted with u_k and $u_k^\dagger$ to

$$(O_{aux}^{[s_1,s_2]})_{k_1',k_2'}^{k_1,k_2} = \sum_{s_1,s_2,s_1',s_2'} (u_k)_{s_1,s_2}^{k_1,k_2} (O^{[s_1,s_2]})_{s_1',s_2'}^{s_1,s_2} (u_k^\dagger)_{k_1',k_2'}^{s_1',s_2'} \ . \tag{5.18}$$

Therefore, the computation of $\langle O^{[s_1,s_2]} \rangle_\psi$ becomes a calculation of the expectation value of the correlator $\langle O_{aux}^{[s_1,s_2]} \rangle_\psi$ for the TTN. The worst-case complexity of the underlying contraction in the TTN is $\mathcal{O}(m^4\kappa \log N) + \mathcal{O}(m^4\kappa \log N - 1) + \mathcal{O}(m^3\kappa)$ where κ describes the internal bond dimension of the TPO after the mapping.

Mapping to a three-site TPO   Let there be only one disentangler u_k addressing at either site s_1 or s_2, and connecting it with the site $s_3 \notin \{s_1, s_2\}$. Here the layer $\mathcal{D}(u)$ maps the correlator to a three-site operator, as it is contracted to

$$(O_{aux}^{[s_1,s_2]})_{k_1',k_2',s}^{k_1,k_2,s} = \sum_{s_c,\tilde{s},s_c'} (u_k)_{s_c,\tilde{s}}^{k_1,k_2} (O^{[s_c,s]})_{s_c',s'}^{s_c,s} (u_k^\dagger)_{k_1',k_2'}^{s_c',\tilde{s}} \ . \tag{5.19}$$

Note that we here changed the notation of $O^{[s_1,s_2]}$ to $O^{[s_c,s]}$ without loss of generality where s_c indicates the link of the correlator attached to the disentangler u_k and thus is to be contracted over while s indicates the remaining link for the resulting three-site operator. After the mapping the computation equals a calculation of the expectation value $\langle O_{aux}^{[s_1,s_2]} \rangle_\psi$ for the TTN - describing a string observable which in this case scales with a worst-case complexity of $\mathcal{O}(m^4\kappa \log N) + \mathcal{O}(m^4\kappa \log N - 1) + \mathcal{O}(m^3\kappa)$.

Mapping to a four-site TPO   Let the disentanglers u_{k_1} and u_{k_2} be placed at the sites s_1 or s_2, and connecting them with the sites $s_3 \notin \{s_1, s_2\}$ and $s_4 \notin \{s_1, s_2\}$ respectively. In this most complex case, the correlator becomes mapped to a four-site operator, the contraction reads

$$(O_{aux}^{[s_1,s_2]})_{k_1',k_2',k_3',k_4'}^{k_1,k_2,k_3,k_4} = \sum_{s_1,s_2,k',k''} (u_{k_1})_{k',s_1}^{k_1,k_2} (u_{k_2})_{k'',s_2}^{k_3,k_4} (O^{[s_1,s_2]})_{s_1',s_2'}^{s_1,s_2} (u_{k_2}^\dagger)_{k_3',k_4'}^{k'',s_2'} (u_{k_1}^\dagger)_{k_1',k_2'}^{k',s_1'} \ .$$

After this mapping the computation again equals a calculation of the expectation value of a string observable for the TTN, which in this case exhibits a worst-case complexity of $\mathcal{O}(m^4\kappa^2 \log N) + \mathcal{O}(m^4\kappa^2 \log N - 1) + \mathcal{O}(m^3\kappa)$.

5.5 Numerical studies

In order to benchmark the aTTN algorithm, we analyze the antiferromagnetic two dimensional Heisenberg model

$$\mathcal{H} = \sum_{i,j=1}^{L} \sum_{\gamma \in \{x,y,z\}} \sigma_{i,j}^{\gamma} \sigma_{i+1,j}^{\gamma} + \sigma_{i,j}^{\gamma} \sigma_{i,j+1}^{\gamma},$$

with $\sigma_{i,j}^{\gamma}$ ($\gamma \in \{x,y,z\}$) the Pauli matrices acting on the site (i,j). We consider a $L \times L$ lattice with system sizes $L = \{8, 16, 32\}$ and set the boundary terms with $\sigma_{i,N+1}^{x} = \sigma_{i,1}^{x}$ and $\sigma_{N+1,j}^{x} = \sigma_{1,j}^{x}$ as periodic boundary conditions. The ground state of this Heisenberg Hamiltonian is characterised by power-law decaying correlations and has been established by several numerical approaches as as challenging benchmark.

5.5.1 Energy density calculation
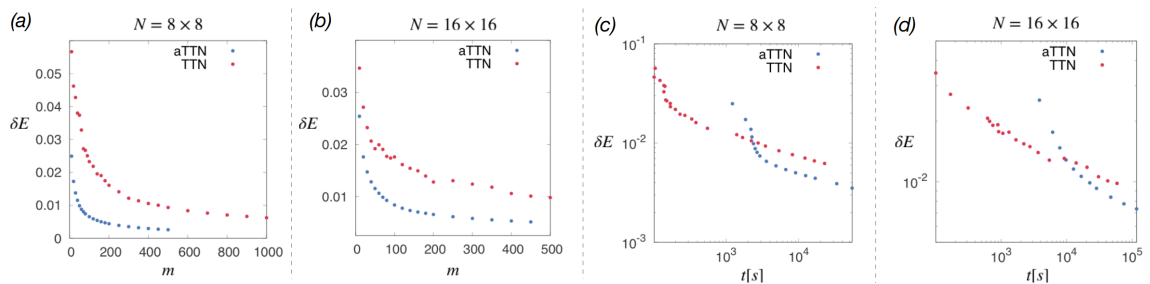
FIGURE 5.4: Relative error δE of the 2D Heisenberg ground-state energy compared with the best available estimates obtained via Quantum Monte Carlo [52] as a function of (a) the bond dimension m at $L = 8$, (b) the bond dimension m at $L = 16$, (c) the CPU time $t[s]$ at $L = 8$ and (d) CPU time $t[s]$ at $L = 16$. The TTN results in red, aTTN in blue.

In Fig. 5.4 we compare the energy density obtained by the TTN with the ones obtained via aTTN with increasing bond-dimension m. Evidently, the aTTN is more accurate than the TTN when compared at same bond-dimension. However, due to a constant prefactor c in the numerical complexity, the CPU time $t \approx \mathcal{O}\left(cm^4\right)$ for the aTTN is higher at same bond-dimension. Therefore, we compute the relative error obtained by the different ansätze and the best known results, obtained via Quantum Monte Carlo [52]. As expected, when keeping the bond dimension equal for both approaches, the aTTN obtains a higher precision for the energy density compared to the TTN but requires more CPU time for the complete optimisation. However, when we look at constant CPU times, the aTTN eventually still outperforms the TTN in terms of accuracy in the energy density for sufficiently large bond dimension m. Thus, it is evidently the more efficient Tensor Network for these two-dimensional simulations. In Fig. 5.4 c and d we compare the CPU time of each calculation against the accuracy as a figure of merit for the efficiency of the ansatz, i.e. which of the Tensor Networks provide a lower relative error at constant CPU time. We see, that for low CPU time (consequently, low bond dimensions) it is more efficient to scale up the bond-dimension of the TTN, while eventually the aTTN becomes the more efficient ansatz when increasing the simulation time, i.e. the bond-dimension. We mention that this prefactor c depends on the number of Hamiltonian parts $\mathcal{H}_p$ addressed by the disentanglers of the aTTN.

When we extrapolate with the bond dimension, we further confirm that the aTTN is more accurate compared to the TTN for both system sizes with a deviation of $\epsilon_8 = 2.22 \cdot 10^{-5}$ and $\epsilon_{16} = 7.11 \cdot 10^{-4}$ in the energy density for both sizes $L = 8$ and $L = 16$ (with respect to Monte Carlo). The TTN, in contrast, performs reasonably well in the 8×8 case but struggles in precision when going to $L = 16$. Interestingly, the aTTN for $L = 16$ obtains an even more precise ground state energy density compared to alternative variational ansätze at lower finite system size of $L = 10$, such as Neural Network states [97] ($\epsilon_{10} \approx 7.5 \cdot 10^{-4}$), Entangled Pair States [194] ($\epsilon_{10} = 2.5 \cdot 10^{-3}$) or PEPS [138] ($\epsilon_{10} = 1.5 \cdot 10^{-3}$). It turns out that for this model a very competitive variational approach is the 2D-DMRG [72], which outperforms the alternative methods for finite size with open or cylindrical boundary conditions up to the system size $L = 10$ ($\epsilon_{10} \approx 2 \cdot 10^{-5}$), but struggles with periodic boundary conditions and with increasing system sizes $L \gtrsim 12$. We further mention that, PEPS is very efficient with its ability to work directly in the thermodynamic limit in describing infinite systems as iPEPS [195, 196]. We as well point out that further more complex Neural Network states, such as the Neural Autoregressive Quantum States [197], obtain comparable results for $L = 10$ with $\epsilon_{10} \approx 3.5 \times 10^{-5}$.

Furthermore, we extended our analysis to reach the system size of $N = 32 \times 32$ for which, to the best of our knowledge, no public result of alternative methods is yet available as a comparison. In consistency with the finite size scaling analysis of Ref. [52], we obtain an energy density $\epsilon_{32} = -0.669(1)$ as extrapolated value $m \to \infty$. We mention that there exist techniques combining Tensor Networks with Variational Monte Carlo, shown with the PEPS++ [198, 199] providing a high potential with strong results at system sizes up to $N = 32 \times 32$. However due to the underlying PEPS geometry in this case, the approach still suffers from an unappealing numerical scaling of $\mathcal{O}\left(m^{10}\right)$ and requires an tremendous computational effort for obtaining these results.

5.5.2   Correlations

Next to the energy density, we compared the correlations captured within both approaches, the TTN and the aTTN. In particular, we computed the two-body correlations $\langle \sigma_{i,j}^z \sigma_{i',j'}^z \rangle$. In Fig. 5.5 we show the correlations for the 8×8 size (a) and 16×16 size (b) for different bond dimensions. It is evident that the computed correlations of the ansätze are growing and getting more accurate with an increasing number of variational parameters, i.e. with increasing bond dimension m. Indeed, when comparing both approaches at equal bond dimension, the obtained correlations of the aTTN simulations are clearly more accurate, demonstrating the higher precision of the aTTN in describing two-body correlations. Interestingly, for lower distances $r(i,j) < 5$ in the $L = 16$ case, some of the correlations computed via TTN are more precise then the aTTN correlations at equal bond-dimension, while the TTN correlations clearly decrease faster for higher distances $r(i,j) < 5$. At the most far distant meassurement at $r(8,8) = 8\sqrt{2}$, we observe that the TTN correlations for $m_{\text{TTN}} = 120$ even end up at the same precision as the aTTN correlations for lower bond dimension $m_{\text{aTTN}} = 80$ and for $m_{\text{TTN}} = 80$ clearly below the aTTN correlations for $m_{\text{aTTN}} = 40$.

We point out that the here performed aTTN simulations (as well as the TTN simulations) exploit a $U(1)$ symmetry. However, for this model, we could further drastically improve the performance of the aTTN by incorporating the present $SU(2)$ symmetry in the simulation framework [24, 169, 200]. In particular, this would *(i)* drastically improve the connected correlations since we enforce $\langle \hat{\sigma}_j^\alpha \rangle = 0$, as well as $\langle \hat{\sigma}_i^\alpha \hat{\sigma}_j^\alpha \rangle_c$ to be equivalent independent
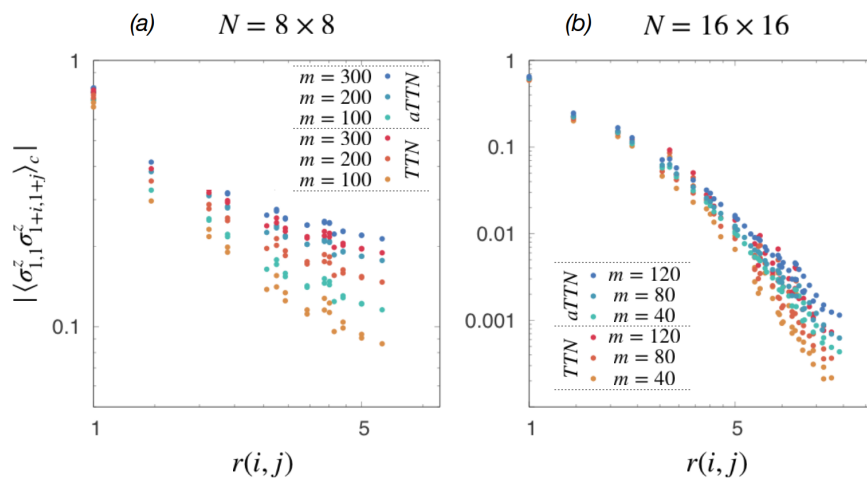
FIGURE 5.5:   Connected correlation function in the TTN (in shades of red for different bond dimensions) and the aTTN (in shades of blue) representation of the ground state of the $L = 8$ (a) and $L = 16$ (b) Heisenberg Hamiltonian plotted against the distance $r(i,j) \equiv [i^2 + j^2]^{1/2}$, where $i \leq j \in \{0, L/2\} \times \{0, L/2\}$.

on $\alpha \in \{x, y, z\}$, and *(ii)* reduce the effective bond dimension and thereby the computational time since for non-abelian symmetry we may only work within the symmetry multiplet spaces. Both advantages are expected to lead to dramatically increase the accuracy in the estimated energy.

5.5.3   Disentangler engineering

Additionally, we performed an analysis on the positioning of the disentanglers in the aTTN based on the Heisenberg model. In particular, we perform the simulation with three different geometries (see Fig. 5.6) investigating the influence of the aTTN geometry on the overall precision for $L = 8, 16$:

(a) In this strategy, we start placing as many disentanglers as possible to reinforce the top-most link (its bipartition indicated with red-dotted lines in Fig. 5.6a), then place as many disentangler as possible to reinforce the second highest links (indicated in blue), and so on. Following this strategy for the 8×8 system, we position all $N_D = K_1 = 8$ disentanglers supporting the highest TTN layer with no place left for disentanglers supporting the second layer.

(b) Starting from strategy (a), we now sacrifice disentanglers reinforce the top-most link in order to place some disentanglers for the second highest tree layer (its bipartition indicated with blue-dotted lines in Fig. 5.6b). In this way we are able to increase the density of disentanglers applied. For the 8×8 system, we position $K_1 = 6$ disentanglers supporting the highest TTN layer and $K_2 = 4$ disentanglers for the second highest TTN layer, totalling $N_D = 10$ disentanglers.

(c) In the last strategy, we change the structure of the internal TTN: Now, instead of coarse-graining alternatingly in $x-$ and $y-$direction, now we coarse-grain the complete y-direction first, before we go on to coarse-grain in x-direction when initialising the
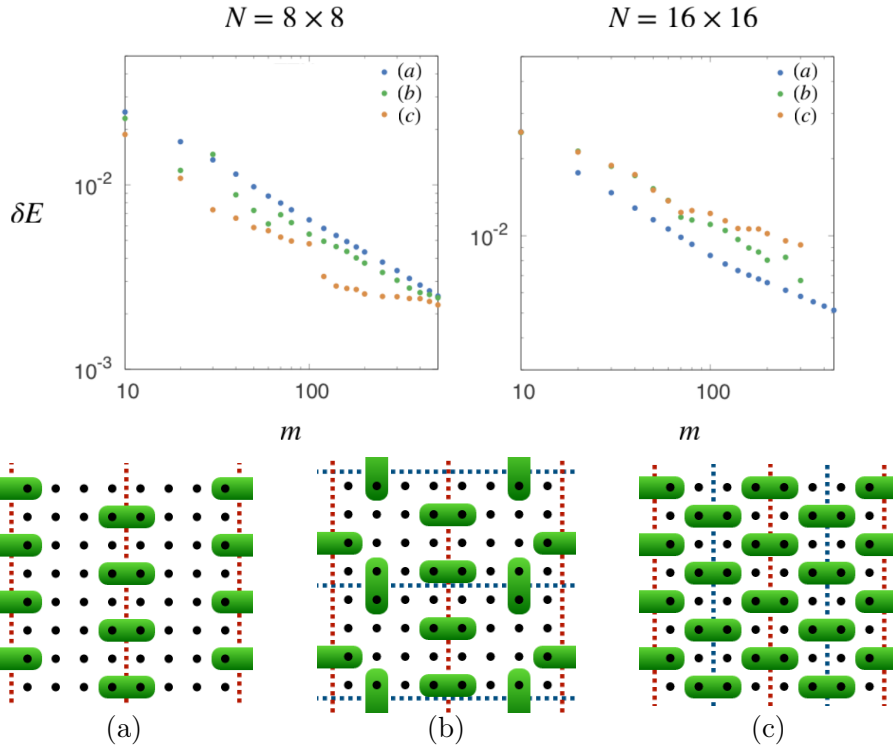
FIGURE 5.6: (top) Relative error ϵ of the 2D Heisenberg ground-state energy for different geometrical configurations (exemplified for $L = 8$ at the bottom) of the aTTN compared with the best available estimates obtained via Quantum Monte Carlo [52] as a function of the bond dimension m at $L = 8$. (bottom) The green shapes represent the disentanglers, dashed lines indicate the bipartitions introduced by the top-most (red) and the second highest (blue) links within the tree.

tree from the bottom to the top. Thus, as indicated by the dashed lines in Fig. 5.6c, the topmost links bipartite the system in x-direction. In this way, we can place the maximum amount of disentanglers to reinforce the two top-most links, namely $N_D = 16$ with $K_1 = K_2 = 8$. Thereby, in the $L = 8$ case, the lowest 3 layers will coarse-grain each column of the system with each of the lowest branches grouping the 8 sites in y-direction together while for $L = 16$, each column of 16 sites will be joined in the lowest 4 layers.

In Fig. 5.6 we present the deviation in the energy density for the different geometries with respect to Quantum Monte Carlo [52]. In the 8×8 analysis, we see that all of the geometries give very similar results along the complete range of m. Surprisingly, reshaping the entire tree structure, i.e. geometry (c), offers the best results out of the three configurations, while the initial strategy (a) which was used in the previous analysis comes out worst. When moving to the $L = 16$ case, however, geometry (c) clearly fails to keep up with the other two and strategy (a) comes out ahead. Indeed this can be traced back to the same problem of 2D-DMRG for $L \gtrsim 12$: Keeping L reasonably small (e.g. $L = 8$), the lower branches of the internal TTN are able to efficiently capture the necessary information to faithfully represent the ground state. However, moving to higher system sizes, the internal TTN struggles to keep up efficiently joining the 16 "column" sites together. Therefore, comparable to the 2D-DMRG, the strategy (c) should offer the best results for an aTTN simulation for limited second dimension L_y and might be a good compromise for non-squared systems ($L_x > L_y$) with periodic boundary conditions. On the other hand, we find that the most efficient strategy

for general $L \times L$ systems indeed is strategy (a). However, we point out that finding the best possible geometry is a highly non-trivial engineering task and depends on many factors, such as the system size, the interactions within the system and the boundary conditions. The optimisation of the Tensor Network structure towards the particular problem might be an interesting problem to solve in future work.
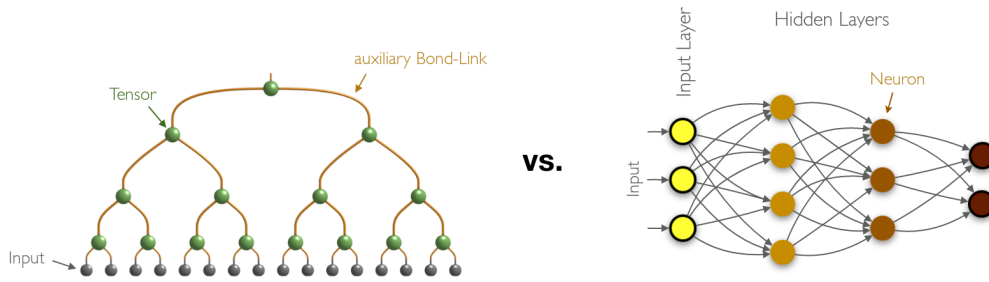
Comparing Quantum State Representations

This Chapter aims to compare the Tensor Network representation of quantum many-body states to a class of alternative variational ansätze for representing a quantum state, namely the class of Neural Network states. The underlying Artificial Neural Networks by design mimic the neural networks of the human brain and are famously applied in the context of data science and artificial intelligence [201–203]. They are arguably the most prominent numerical technique for performing Machine Learning tasks and have been deployed in a broad range of applications in research as well as industry.

In physics, Machine Learning has already shown to be very useful for a broad range of applications in condensed matter, quantum chemistry, material science and statistical physics [204–207]. In the most recent years, further connections between Quantum Information and Machine Learning have been shown and continue to be uncovered giving among others raise to technologies such as Quantum Machine Learning [208–211]. This interplay between Machine Learning and Quantum Information has brought to light the aforementioned application of artificial Neural Networks to describe the behaviour of complex quantum many-body systems as an alternative to Tensor Network.

Indeed, Neural-Network- and Tensor-Network-based algorithms are tightly bound to a number of fundamental issues in statistical mechanics [212]. Their descriptive power and numerical efficiency strongly affect the ability to explore the quantum properties of many-body quantum systems. In fact, in many cases, Neural Networks can even be mapped to a Tensor Network showing the close connection between the two numerical methods [213]. In particular, recent progress in the application of Neural Network techniques in quantum physics established a functional representation to characterise a quantum many-body wave function. This quantum state representation was introduced in the form of a Restricted Boltzmann Machine with the intention of representing highly correlated quantum states whose entanglement content scales with the system volume [97, 214–216]. Therefore, with this approach of Neural Networks becoming a more and more prominent technique to analyse complex quantum systems, the comparing of this ansatz with the well-established Tensor Networks naturally arises as one of the most interesting numerical questions to be investigated.

In the following manuscript, we compare the Tensor-Network-based algorithms against Neural-Network-based algorithms in two exemplary benchmarks, the one-dimensional ferromagnetic Ising model and the two-dimensional antiferromagnetic Heisenberg model. This manuscript was published in *SciPost Physics Core*, **4** [98]. Therein, we present a new mapping for the Neural-Network-based, unrestricted Boltzmann Machine (uRBM), both in 1D and 2D, to constrained Tensor Network representations, namely the constrained MPS (coMPS) and constrained PEPS (coPEPS) respectively. We unveil fundamental aspects of this important connection illustrating that the actual descriptive power of the Neural Network ansatz as constrained Tensor Networks with exponentially large bond dimension should be taken with a grain of salt as it fails to fulfil the aforementioned *volume-law*.

Therefore, this chapter gives new insights on the connection between Neural Networks and Tensor Networks, as well as illustrating current limitations of the representation power of Neural Network states. However, this encouraging field of research poses the open, highly nontrivial question on how to properly estimates the information encoded into a Neural Network and in this context on how to fairly compare different numerical approaches for quantum state representations to be addressed by future systematic research. Further, with the introduced mapping, we open a pathway to go beyond Variational Montecarlo optimisation techniques enabling the development of new methods, based on intuitions of Tensor Network techniques such as DMRG, TEBD and TDVP, for the optimisation of Neural Network states.

On the descriptive power of Neural Networks as constrained Tensor Networks with exponentially large bond dimension

Mario Collura[1,2,3], Luca Dell'Anna[2,4], Timo Felser[1,2,5] and Simone Montangero[2,4,5]

1 Theoretische Physik, Universität des Saarlandes, D-66123 Saarbrücken, Germany.
2 Dipartimento di Fisica e Astronomia, "G. Galilei",
Università di Padova, I-35131 Padova, Italy.
3 SISSA – International School for Advanced Studies, I-34136 Trieste, Italy.
4 Padua Quantum Technologies Research Center, Università di Padova, I-35131 Padova, Italy.
5 INFN, Sezione di Padova, Via Marzolo 8, I-35131 Padova, Italy.

Abstract

In many cases, neural networks can be mapped into tensor networks with an exponentially large bond dimension. Here, we compare different sub-classes of neural network states, with their mapped tensor network counterpart for studying the ground state of short-range Hamiltonians. We show that when mapping a neural network, the resulting tensor network is highly constrained and thus the neural network states do in general not deliver the naive expected drastic improvement against the state-of-the-art tensor network methods. We explicitly show this result in two paradigmatic examples, the 1D ferromagnetic Ising model and the 2D antiferromagnetic Heisenberg model, addressing the lack of a detailed comparison of the expressiveness of these increasingly popular, variational ansätze.

Contents

1 Introduction

Artificial Neural Networks (NN) have increasingly taken hold in various research fields and technology [1–3]. Their power in recognizing special patterns behind a huge amount of raw data allowed a revolutionary change in our approach towards deep learning [4, 5]. Taking inspiration from biological neural networks, NNs can be a good framework to process big sets of data. As a matter of fact, NNs can be seen as special functional mappings of many variables (physical and hidden), which can be trained by specific algorithms and applied to a very broad spectrum of applications in different fields, one of which is statistical physics [2, 6–9]. A powerful example is the Restricted Boltzmann Machine (RBM) which has been largely employed to mimic the behaviour of complex quantum systems [10–17]. Essentially, RBM is a type of artificial neural network which, in interacting quantum systems, can be understood as a particular variational ansatz for the many-body wave function.

Another very successful class of wave-function variational ansatz that has been widely exploited are Tensor Network (TN) states [18–34]. They are based on the replacement of the non-local rank-N tensor representing the N-body wave function, with $O(N)$ local tensors with smaller rank, connected via auxiliary indexes. Such ansatz interpolates between the mean-field approach, where quantum correlations are completely neglected, and the exact (but inefficient) representation of the state. The interpolation is governed by the dimension χ of the auxiliary indexes connecting local tensors. In one dimension, a very successful tensor network representation is the so-called Matrix Product States (MPS) [18, 19, 25, 33].

When applied to the study of many-body quantum systems, these two very powerful approaches reduce the exponentially large Hilbert space dimension by optimally tuning a number of parameters which scales polynomially with N. In particular, the number of free parameters scales as $O(NM)$ for a fully connected RBM (where M is the number of hidden variables), while is $O(N\chi^2)$ for an MPS. Recently, a strong connection between NN and TN has been pointed out [40–46]: among others, it has been shown that the fully connected RBM can be explicitly rewritten as a MPS with an exponentially large auxiliary dimension, i.e. $\chi = 2^M$ (see Fig. 1). Considering the fact that the bipartite entanglement entropy in an MPS is proportional to $\log \chi$, suggests that RBMs may provide a way to represent highly correlated quantum states whose entanglement content scales with the system volume [47], thus going much beyond the MPS descriptive power that is limited to area-law states [48, 49].

Here, present a systematic comparison between constraint TN representations of NNs and the unconstrained counterparts aiming to investigating the actual descriptive power of the NN ansatz as constrained Tensor Networks with exponentially large bond dimension. With this comparison we further aim to address the lack of a detailed comparison of the expressiveness of these various ansatz considering the increasing popularity of such variational states and encourage further work in this direction. In the following section, we unveil fundamental aspects of this important connection between TN and NN states and along the way we introduce a new mapping between NN and TN, valid also in two-dimensions. This mapping can be exploited to introduce more efficient strategies to optimize NN states. Finally, we present two paradigmatic examples, the study of one-dimensional and two-dimensional ground states of many-body quantum Hamiltonians and show that a TN with moderate bond dimensions can match – if not overcome – the prediction of the correspondent NN, that were expected to deliver results equivalent to those obtained via TN with an in practice unreachable bond dimension. In particular, we compare the prediction of NN and TN, and show that the latter are able to deliver higher precision in local quantities and in correlation functions with respect to the equivalent NN. We further clarify why these results have to be expected.

In one dimension, these results are based on the fact that beyond the naive expectations, the detailed aspects in the relationship between RBM and MPS shall be considered. In partic-
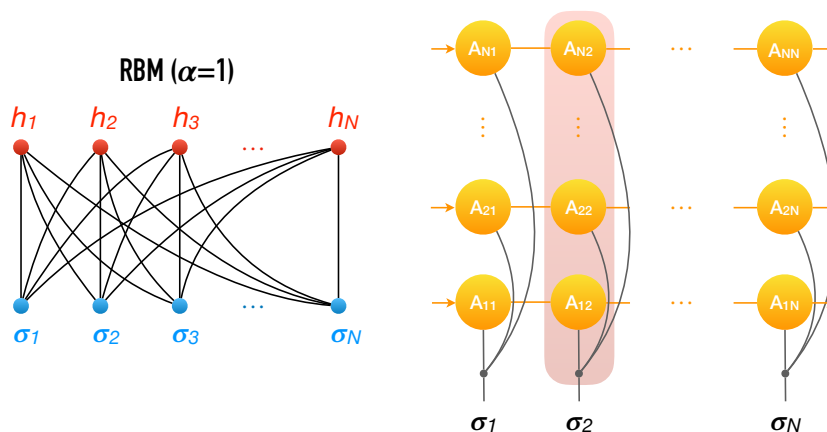
Figure 1: The RBM for $\alpha = 1$ ($M = N$) (left) can be mapped to an exponentially large coMPS (right). The pink-shaded region represents the $2^N \times 2^N$ local matrix $\mathbf{\Sigma}_j^{\sigma_j}$ which is constructed as a tensor product of N matrices (yellow circles) with bond dimension $\chi = 2$ (see main text for details). Arrows indicate periodic boundary conditions.

ular, *(i)* different topologies of the NN may result in MPSs with very different auxiliary dimensions; *(ii)* the emerging "constrained" MPS (coMPS) is highly constrained due to the mapping itself. In practice, even though the formal mapping reveals an exponentially large auxiliary dimension of the related coMPS, the number of independent entries of the tensors scales polynomially with the size of the artificial neural network. This makes the coMPS representation of the RBM inefficient.

The first point above raises a very delicate question regarding the efficiency of the NN *tout court*. Physical many-body states are usually eigenstates of short-range interacting Hamiltonians, which naturally introduce the notion of distance between lattice sites, suggesting that a valuable approach should encode this information *ab initio*. Therefore, a useful modification of the RBM is obtained by introducing proper intra-layer connections in the wave function ansatz; the resulting unRestricted Boltzmann Machine (uRBM) has been recently employed in order to describe the ground state of the ferromagnetic Ising quantum chain [50]. A simple structure with one layer of hidden variables (see Fig. 2) explicitly encodes the underlying Hamiltonian geometry and makes it possible to obtain an increased accuracy with respect to the corresponding RBM (i.e. with $\alpha = 1$), with a much smaller number of free parameters. The surprising effect of this result will become evident in the following, when the RBM and the uRBM will be compared at the level of the mapping to the corresponding coMPS.

2 Constrained Matrix Product States

Here we introduce the "constrained" Matrix Product State, highlighting the differences between RBM and uRBM. In the following, periodic boundary conditions are assumed, and results are valid for $N > 2$.

<u>RBM</u>: An RBM is defined as follows: given a set $\boldsymbol{\sigma} = \{\sigma_1, \sigma_2, \dots, \sigma_N\}$ of N physical binary variables (e.g. the eigenvalues $\sigma_j = \pm 1$ of a spin-1/2), one introduces an extra set of "unphysical" hidden variables $\boldsymbol{h} = \{h_1, h_2, \dots, h_M\}$ (with density $\alpha \equiv M/N$) such that the unnormalised many-body wave function of the RBM type is obtained from a full Boltzmann distribution by

tracing out the hidden set,

$$\Psi_\alpha(\boldsymbol{\sigma}) = \sum_{\boldsymbol{h}} \exp\left[-\psi_\alpha(\boldsymbol{\sigma},\boldsymbol{h})\right], \tag{1}$$

where ψ_α is the RBM's functional

$$\psi_\alpha(\boldsymbol{\sigma},\boldsymbol{h}) = \sum_{j=1}^{N} a_j \sigma_j + \sum_{j=1}^{M} b_j h_j + \sum_{i,j=1}^{M,N} \Gamma_{ij} h_i \sigma_j, \tag{2}$$

and $\alpha = M/N$ is assumed to be a positive integer for convenience. Parameters a_j and b_j are local biases (in the spin-1/2 picture, they represent local magnetic fields) applied to the variables, whilst Γ_{ij} are couplings between the physical and the hidden variables (see Fig. 1).

As stated in the introduction, in the previous prescription there is no direct connections within the same set of variables while the two sets are fully connected, thus there is no notion of a distance. Nonetheless, correlations between physical variables may be mediated by their fictitious interactions with the hidden variables. A priori the RBM variational ansatz is well suited to work in any dimension. As a matter of fact, it is a promising tool to describe many-body quantum states [35–39]; recently convolutional NNs have been employed to improve the level of accuracy of the shallow RBMs in order to deal with frustrated 2D lattice models [51].

RBM wave functions can be rewritten as a coMPS with periodic boundary conditions. The absence of intra-layer couplings allows to easily take the sum over the hidden variables in Eq. (1), thus obtaining [40]

$$\Psi_\alpha(\boldsymbol{\sigma}) = e^{-\sum_{j=1}^{N} a_j \sigma_j} \prod_{i=1}^{M} 2\cosh\left(b_i + \sum_{j=1}^{N} \Gamma_{ij}\sigma_j\right) = \mathrm{Tr}\left[\prod_{j=1}^{N} \Sigma_j^{\sigma_j}\right] \tag{3}$$

in terms of $2^M \times 2^M$ real diagonal matrices (see Fig. 1) of the form

$$\Sigma_j^\sigma = e^{-a_j \sigma} \bigotimes_{i=1}^{M} \begin{pmatrix} e^{-b_i/N - \Gamma_{ij}\sigma} & 0 \\ 0 & e^{b_i/N + \Gamma_{ij}\sigma} \end{pmatrix}. \tag{4}$$

Notice that, if the RBM wave function describes a translational invariant quantum state, we should have $\Psi_\alpha(\boldsymbol{\sigma}) = \Psi_\alpha(\boldsymbol{\sigma}')$, where $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}'$ differ for an arbitrary cyclic permutation of local spin variables: thus, all local tensors $\Sigma_j^{\sigma_j}$ can be set to be equal and independent of the lattice site j, reducing the number of free parameters to $2M + 1$. Let us mention that when the local biases are set to zero, the wave function becomes spin-flip invariant as well.

1D – uRBM: Let us now turn our attention to the unrestricted Boltzmann machine for 1D systems. When explicitly encoding the geometry of the underlying model into the artificial neural network, we can describe the many-body quantum state via the following uRBM un-normalised wave function

$$\Phi_\ell(\boldsymbol{\sigma}) = \sum_{\boldsymbol{h}} \exp\left[-\phi_\ell(\boldsymbol{\sigma},\boldsymbol{h})\right], \tag{5}$$

where now the hidden variables are labeled according to $\boldsymbol{h} = \{h_j^\gamma\}$ with $j \in [1,N]$, and $\gamma \in [1,\ell]$ denoting the different layers. The uRBM's functional is now given by

$$\begin{aligned} \phi_\ell(\boldsymbol{\sigma},\boldsymbol{h}) &= \sum_{j=1}^{N}\Big[K_j^0 \sigma_j \sigma_{j+1} + \sum_{\gamma=1}^{\ell} K_j^\gamma h_j^\gamma h_{j+1}^\gamma \\ &+ J_j^1 \sigma_j h_j^1 + \sum_{\gamma=2}^{\ell} J_j^\gamma h_j^{\gamma-1} h_j^\gamma \Big]. \end{aligned} \tag{6}$$
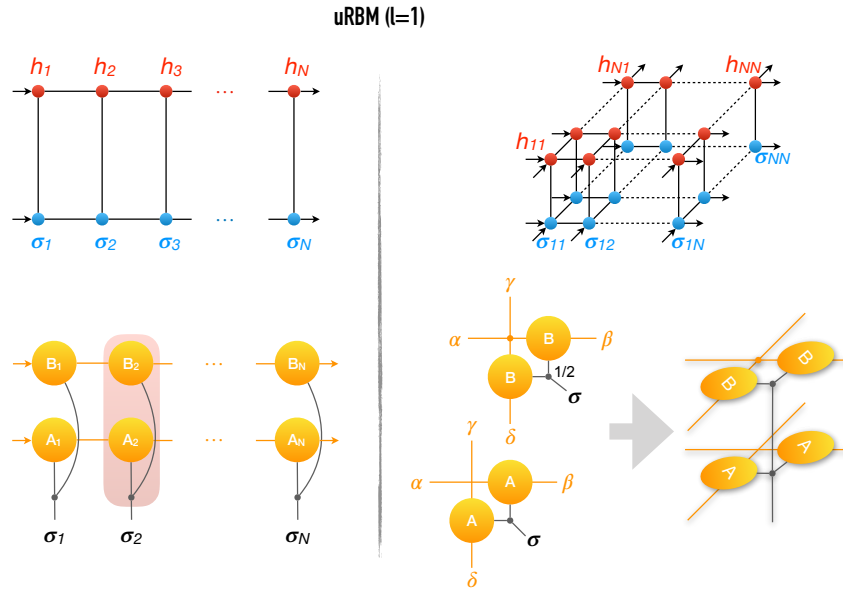
uRBM (l=1)

Figure 2: The uRBM with one layer of hidden variables in one (two) dimension is mapped to the corresponding coMPS (coPEPS). In the 2D case we only draw the local building block of the full tensor network. Here in general crossing lines are independent, except when a dot fixes them to be equals. The $1/2$ in one of the dot means that the $\boldsymbol{B}$ matrices have to be evaluated at $\sigma/2$ (see main text for details). Arrows indicate periodic boundary conditions.

Interestingly, the state described by the uRBM wave function enforces the spin-flip invariance of the Ising Hamiltonian, namely $\Phi_\ell(-\boldsymbol{\sigma}) = \Phi_\ell(\boldsymbol{\sigma})$; moreover, it is also invariant under the transformation $J_j^\gamma \rightarrow -J_j^\gamma$, for arbitrary γ. Although we cannot analytically sum over the hidden variables as before, it is still possible to trace-out the hidden variables recasting Eq. (5) in a simple coMPS form. Exploiting the transfer matrix approach for evaluating the partial partition function, we obtain (e.g. for $\ell = 1$)

$$\Phi_1(\boldsymbol{\sigma}) = \text{Tr}\left[\prod_{j=1}^N (A_j^{\sigma_j} \otimes B_j^{\sigma_j})\right], \tag{7}$$

where

$$A_j^\sigma = \begin{pmatrix} \cosh(K_j^0) & -\sinh(K_j^0)\sigma \\ \cosh(K_j^0)\sigma & -\sinh(K_j^0) \end{pmatrix}, \tag{8}$$

$$B_j^\sigma = \begin{pmatrix} e^{-K_j^1 - J_j^1\sigma} & e^{K_j^1 - J_j^1\sigma} \\ e^{K_j^1 + J_j^1\sigma} & e^{-K_j^1 + J_j^1\sigma} \end{pmatrix}. \tag{9}$$

Also, in this case, translational invariance of the many-body state can be exploited reducing the number of free parameters by a factor N. Interestingly, the spin-flip symmetry of the state is reflected in the coMPS representation as the local invariance $A_j^{-\sigma} = \hat{\sigma}^z A_j^\sigma \hat{\sigma}^z$, $B_j^{-\sigma} = \hat{\sigma}^x B_j^\sigma \hat{\sigma}^x$. Finally, the mapping can be extended to an arbitrary number of additional hidden layers which results in a coMPS with auxiliary dimension $\chi_\ell = 2^{\ell+1}$.

<u>2D – uRBM</u>: The geometry encoded in the uRBM may affect the tensor network representation of the many body wave function which, in the 2D cases, can be written as a coPEPS

(constrained Projected Entangled Pair State). For the sake of simplicity, we focus only on the translational invariant case where the 2D uRBM wave function reads

$$\Phi_\ell^{2D}(\boldsymbol{\sigma}) = \sum_{\boldsymbol{h}} \exp\left[-\phi_\ell^{2D}(\boldsymbol{\sigma}, \boldsymbol{h})\right], \tag{10}$$

with

$$
\begin{aligned}
\phi_\ell^{2D}(\boldsymbol{\sigma}, \boldsymbol{h}) &= \sum_{i,j=1}^{N} \left[K^0(\sigma_{i,j}\sigma_{i,j+1} + \sigma_{i,j}\sigma_{i+1,j}) \right. \\
&+ \sum_{\gamma=1}^{\ell} K^\gamma (h_{i,j}^\gamma h_{i,j+1}^\gamma + h_{i,j}^\gamma h_{i+1,j}^\gamma) \\
&+ \left. J^1 \sigma_{i,j} h_{i,j}^1 + \sum_{\gamma=2}^{\ell} J^\gamma h_{i,j}^{\gamma-1} h_{i,j}^\gamma \right]. \tag{11}
\end{aligned}
$$

Summing over the hidden variables, the wave function can be rewritten as a translational invariant coPEPS built from the local tensors (see Fig. 2)

$$
\begin{aligned}
\mathbb{A}_{\alpha\beta\gamma\delta}^\sigma &= (\boldsymbol{A}^\sigma)_{\alpha\beta}(\boldsymbol{A}^\sigma)_{\gamma\delta}, \\
\mathbb{B}_{\alpha'\beta'\gamma'\delta'}^\sigma &= \delta_{\alpha'\gamma'}(\boldsymbol{B}^{\sigma/2})_{\alpha'\beta'}(\boldsymbol{B}^{\sigma/2})_{\gamma'\delta'}, \tag{12}
\end{aligned}
$$

with matrices $\boldsymbol{A}^\sigma$ and $\boldsymbol{B}^\sigma$ given by Eqs. (8) and (9), where we discarded the label j due to the translational invariance. The local building block for the coPEPS is obtained by index fusion, paring each couple of indices to a single index which spans a four-dimensional auxiliary space, i.e. $\boldsymbol{\alpha} = (\alpha, \alpha')$, getting $\mathbb{C}_{\boldsymbol{\alpha\beta\gamma\delta}}^\sigma = \mathbb{A}_{\alpha\beta\gamma\delta}^\sigma \mathbb{B}_{\alpha'\beta'\gamma'\delta'}^\sigma$. Again, in this case, the extension to an arbitrary number of hidden layers is straightforward, and the coPEPS auxiliary dimension is the same as in the 1D case, namely $\chi_\ell = 2^{\ell+1}$.

The coMPS (coPEPS) mapping of the uRBM variational ansatz can be proficiently used together with Monte Carlo techniques in order to avoid the full sampling over the set of hidden variables. Indeed, those representations are a practical way to explicitly trace out the full set of the hidden variables.

3 Numerical results

In what follows we investigate how the different ansätze are able to describe the ground state of critical Hamiltonians (both in 1D and 2D), where bipartite entanglement entropy scales logarithmically with the system size, and the correlation functions decay algebraically.

3.1 The Ising quantum chain

We start our analysis with the ferromagnetic Ising quantum chain, whose Hamiltonian, for N lattice sites and periodic boundary conditions, is given by

$$H_I = -\sum_{j=1}^{N} \hat{\sigma}_j^z \hat{\sigma}_{j+1}^z - \lambda \sum_{j=1}^{N} \hat{\sigma}_j^x, \tag{13}$$

where $\hat{\sigma}_j^\gamma$ (for $\gamma \in \{x, y, z\}$) are Pauli matrices acting on the site j, and $\hat{\sigma}_{N+1}^\gamma = \hat{\sigma}_1^\gamma$. The transverse field λ drives the ground state from a ferromagnetic region ($\lambda < 1$) to a paramagnetic region ($\lambda > 1$) across a quantum critical point.
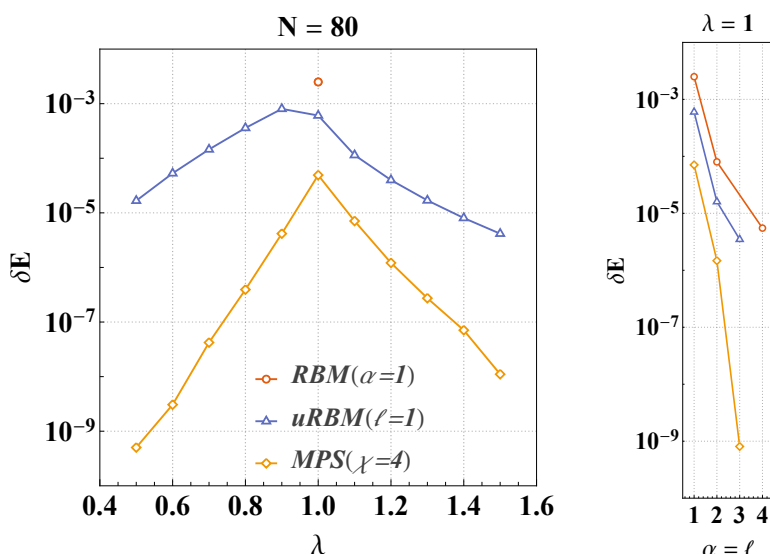
Figure 3: Relative error in the ground-state energy estimate for different many-body wave function representations. (left panel) We compare the uRBM with $\ell = 1$ with respect to the canonical MPS with the same bond dimension ($\chi = 2^{\ell+1} = 4$) as a function of the transverse field λ. (right panel) Scaling analysis of the energy error at the critical point, as a function of the hidden variable density $\alpha = \ell$ (for RBM and uRBM) and analogous bond dimension $\chi = 2^{\ell+1}$ (for the MPS). RBM (uRBM) representation reaches (overtakes) the accuracy of the MPS with $\chi = 4$ only for $\alpha = 2$ ($\ell = 2$); however, they remain above the estimate obtained with a canonical MPS with the same auxiliary dimension of the uRBM, i.e. $\chi = 2^{\ell+1} = 8$.

Exploiting the coMPS mapping of the uRBM, we are able to optimize the many-body wave function very efficiently. We consider a chain with periodic boundary conditions and mainly focus on the one layer case ($\ell = 1$), thus reducing the number of variational parameters to 3. Due to the coMPS representation of the variational wave function in Eq. (7) we are able to evaluate the Hamiltonian expectation value exactly. Thus, we improve the accuracy and the computational time compared to what has been recently found for the ground state energy with an uRBM in Ref. [50] via Monte Carlo methods. For sake of clarity, we point out that this approach drastically improves the evaluation of expectation values only and does not affect the time required for the optimisation of the uRBM wave-function.

In the left panel of Fig. 3 we report the relative error of the best estimate of the ground-state energy with respect to the exact value, namely $\delta E = |(\langle H_I \rangle - E_{ex})/E_{ex}|$, for a system of size $N = 80$ and varying the transverse field $\lambda \in [0.5, 1.5]$. We compare the results of the uRBM with $\ell = 1$ against the data obtained with a traditional MPS-based algorithm [52] with the same auxiliary dimension $\chi = 2^{\ell+1} = 4$. At the critical point, we also report the result obtained in Ref. [13] with the RBM variational ansatz and the *same number of hidden variables* (i.e. $\alpha = 1$). We confirm that appropriate physical insights about the model under investigation not only reduce the computational effort of the algorithm (from $2N + 1$ parameters in the RBM to 3 parameters in the uRBM), but results in higher precision. However, we notice that results based on the canonical MPS representation are order of magnitudes more accurate than those based on the corresponding uRBM representation.

Of course, it must be said that the canonical MPS representation contains more variational parameters than the uRBM with the same bond dimension; we would expect all represen-
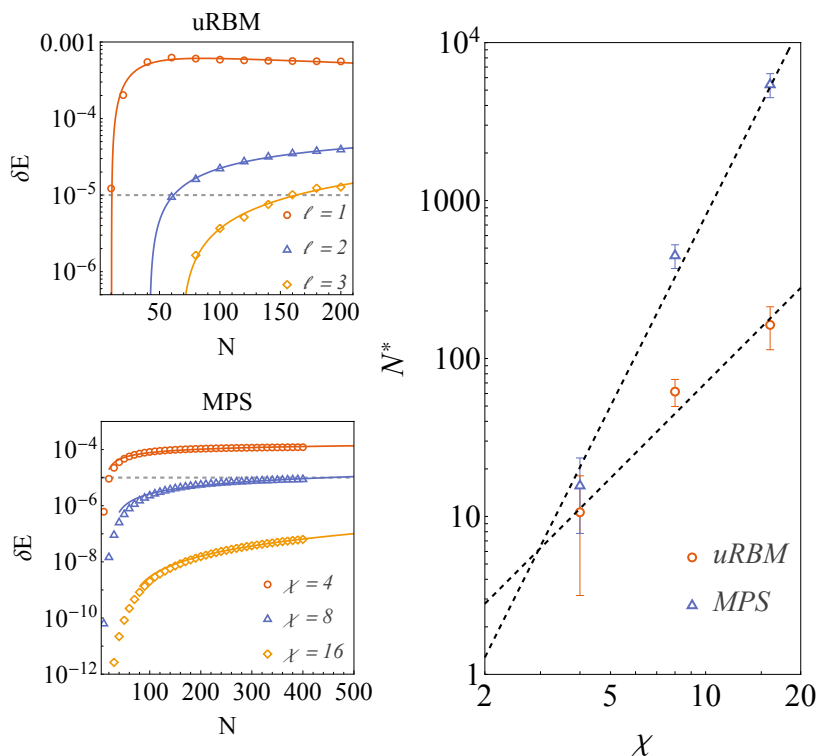
Figure 4: (left panels) Finite size scaling analysis of the relative error in the ground-state energy of the critical Ising chain for different many-body wave function ansatz (uRBM and MPS); the grey dashed lines correspond to the accuracy goal 10^{-5}. (right panel) Scaling of the largest system size N^* which can be described with an energy accuracy $\leq 10^{-5}$ for the two different ansätze. The black dashed lines are power-law fits as reported in the main text.

tations becoming better by increasing the number of variational parameters. Therefore, we investigate this aspect at the critical point (the more computational demanding case), where we perform a scaling analysis of the accuracy in the energy estimation for the different wave function representations. It turns out that, for equal hidden variable density $\alpha = \ell$, the uRBM overtakes the RBM; however, the canonical MPS with the same bond dimension $\chi = 2^{\ell+1}$ of the uRBM remains highly more accurate than any NN representation (see Fig. 3 right panel). For example, for $\alpha = \ell = 2$, we obtain $\delta E_{RBM} \simeq 0.8 \cdot 10^{-4}$, $\delta E_{uRBM} \simeq 0.16 \cdot 10^{-4}$, whilst the MPS with $\chi = 2^{\ell+1} = 8$ gives $\delta E_{MPS} \simeq 0.15 \cdot 10^{-5}$.

Once established that the uRBM with the Ising-like geometry gives better estimates of the ground-state Ising energy with respect to the RBM variational wave function, we may now concentrate on a more systematic Finite Size Scaling (FSS) comparison between uRBM and MPS. Indeed, in order to infer about a possible definition of the *descriptive power* of a given many-body wave-function ansatz, we decided to proceed in the following way: *(i)* We fixed the level of accuracy to be 10^{-5} so as to have a good interpolation of the uRBM data; *(ii)* we extract the largest system size N^* whose ground-state energy can be estimated within that accuracy goal; *(iii)* we analyse the behaviour of N^* as a function the (effective) bond dimension χ, which indeed gives the algorithmic complexity of the energy minimisation procedure for both ansätze. Once again, we concentrate the FSS analysis to the more computational demanding critical chain (i.e. $\lambda = 1$). We perform numerical simulations with sizes $N \in [10, 200]$ for the uRBM with $\ell \in \{1, 2, 3\}$, and $N \in [10, 400]$ for the MPS with $\chi \in \{4, 8, 16\}$.
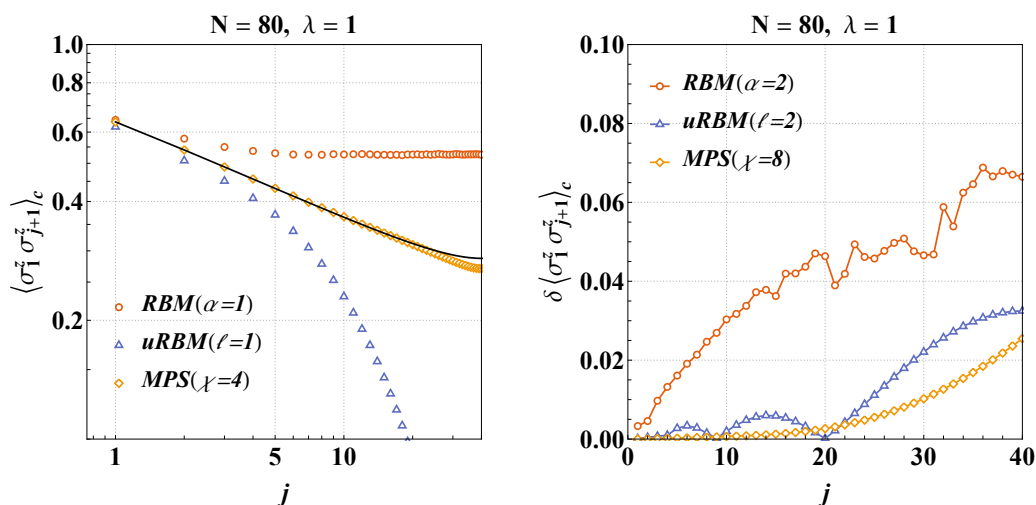
Figure 5: (**left**) Two-point connected correlation function in log-log scale at the critical point for different variational ansatz and smaller description, i.e. $\alpha = \ell = 1$ and $\chi = 2^{\ell+1} = 4$. Black full line are the exact analytical results. (**right**) Relative error from the exact data when larger NN representations are considered; here we compare RBM/uRBM with $\alpha = \ell = 2$ with the canonical MPS with $\chi = 2^{\ell+1} = 8$. All the data for RBM has been obtained by using the optimised wave functions in Ref. [13].

In the left panels of Fig. 4 we report the data (symbols) together with the best power-law fits $\delta E = a + bN^c$ (full lines). From the best-fit parameters and their errors we obtain the estimate of N^*, where the effective bond-dimension of the uRBM is $\chi = 2^{\ell+1}$. A different scaling of the *descriptive power* of the two ansätze is reasonably clear from the right panel of Fig. 4: The black dashed lines indeed show a fair agreement with $N^* \sim \chi^{4.2}$ for the MPS variational wave-function, while only $N^* \sim \chi^{1.9}$ for the uRBM one.

Let us stress once more that, the performances of the uRBM to characterise the low-energy properties of the Ising chain are strictly related to the fact that such variational ansatz somehow encode the Ising geometry. In principle, we do not expect the same degree of accuracy when analysing different 1D critical models; this is what happen, for example, when trying to characterise the ground-state energy of the XXZ spin-1/2 chain in the gapless phase, where the performances of both the uRBM and the MPS are one/two orders of magnitude worst.

Even though the different variational ansätze may give reasonable estimates of the ground-state energy, it is worth investigating the large-distance behaviour of correlation functions. Indeed, at the critical point, we expect a power-law decay of the two-point connected correlation function $\langle \sigma_1^z \sigma_{j+1}^z \rangle_c = \langle \sigma_1^z \sigma_{j+1}^z \rangle - \langle \sigma_1^z \rangle \langle \sigma_{j+1}^z \rangle$, as far as $j \ll N$. However, the MPS structure of the variational ansatz introduces an unavoidable fictitious correlation length. Moreover, the fact that the uRBM energy estimate is better than the RBM estimate (see Ref. [13,50] for a comparison), implies that the uRBM may give a better estimate at the level of the correlation functions as well. With this respect, in the left panel of Fig. 5, we compare the connected two-point function $\langle \sigma_1^z \sigma_{j+1}^z \rangle_c$ evaluated in the optimised uRBM with $\ell = 1$ against the same two-point function evaluated in the unconstrained MPS with auxiliary dimension $\chi = 4$. In order to have the same number of hidden variables in both NN representations, we show the RBM correlations with $\alpha = 1$, which have been obtained by sampling the optimised wave function in Ref. [13] over 10^6 configurations. We focus our analysis to the critical point, where a larger deviation from the exact data is expected. From the figure it is clear that, the canonical

MPS is largely better than the neural-network representation.

In a way, the RBM suffers from a sort of over-estimation of the long-range correlations due to the presence of unphysical long-range couplings between hidden and physical variables; on the contrary, the over-constrained structure of the coMPS representation of the uRBM reflects into a stronger exponential decay of the two-point correlations. However, there is the possibility for those functions obtained by optimised NNs to be improved by the inclusion of further layers in the ansatz, which can increase the degree of correlations.

Indeed, when the number of hidden variables is increased to $\alpha = \ell = 2$ in such a way to reach the same energy accuracy of the MPS with $\chi = 4$ (see Fig. 3, right panel), the corresponding NN description of the correlation function improves as well and reaches that of the MPS with $\chi = 4$. However, it still remains less accurate with respect to the MPS representation with an auxiliary dimension $\chi = 2^{\ell+1} = 8$ (see right panel in Fig. 5). In particular, for distances $j \lesssim 20$ lattice sites, the RBM relative error remains $\lesssim 5\%$, the uRBM gets an error $\lesssim 0.6\%$ whilst finally the MPS reaches a better accuracy with an error $\lesssim 0.16\%$.

From this point of view a structured neural network states, namely a uRBM, seem better tailored than an RBM, at least, to deal with short-range one-dimensional systems. Once again, we may stress that when simulating with a uRBM, the coMPS representation should be used to calculate energy and further expectation values for higher accuracy and lower computational time.

3.2 The 2D Heisenberg model

We now extend our analysis to two-dimensional systems as well, by considering the 2D Heisenberg antiferromagnetic model, whose Hamiltonian is

$$H_H = \sum_{i,j=1}^{N} \sum_{\gamma} \left(\hat{\sigma}_{i,j}^{\gamma} \hat{\sigma}_{i,j+1}^{\gamma} + \hat{\sigma}_{i,j}^{\gamma} \hat{\sigma}_{i+1,j}^{\gamma} \right), \tag{14}$$

where γ runs over $\{x, y, z\}$. We assume here periodic boundary conditions. The ground state of Heisenberg Hamiltonian is characterised by power-law decaying correlations, thus being a perfect two-dimensional benchmark.

As already stressed, an RBM is characterised by an exponentially large bond dimension and seems to work pretty well in the presence of long-range interactions and correlations [47]. Here, in order to avoid the expensive optimisation procedure for a PEPS wave function, we only consider the state-of-the-art RBM variational results in Ref. [13] and compare its descriptive power against a Tree Tensor Network (TTN) [34, 54–57] representation for the 2D Heisenberg ground-state. A TTN is a loop-free Tensor Network which can be efficiently contracted and it is characterised by a finite bond dimension χ which enforces the maximum amount of entanglement for any bipartition of the 2D state to be finite.

In our TTN simulations, we consider both 8×8 and 10×10 sizes; the former is more suitable for a TTN algorithm since it can fully exploit the binary-tree structure. We compare the estimated energy density with the best known results obtained via Quantum Monte Carlo [53] and show the relative deviation in Fig. 6 for both system sizes. In both cases (for $L = 8$ and $L = 10$) the errors of the Quantum Monte Carlo results are below 10^{-5} and therefore negligible compared to both, the TTN as well as the RBM.

The 10×10 case presents a geometry which is much less easy to adapt for TTN computations; nevertheless we are able to reach the same accuracy of an RBM with $\alpha = 1$ by only keeping $\chi = 340$ states (see Fig. 6), which should be compared with the RBM equivalent bond dimension $\chi_{RBM} \sim 2^{\alpha N^2} = 2^{100}$. Let us point out that in the 8×8 case, when a better-suited TTN geometry can be used, we are able to reach one order of magnitude better precision. With a relatively small bond dimension $\chi = 700$ we already almost meet the RBM results with
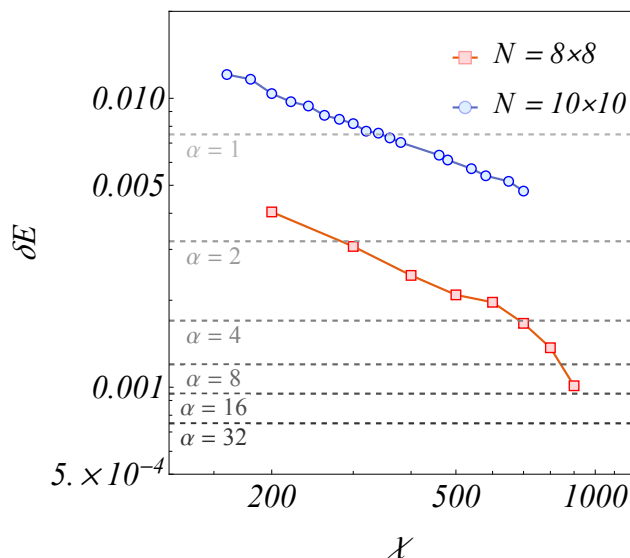
Figure 6: Relative error of the 2D Heisenberg ground-state energy compared with the best available estimates obtained by the finite-size Quantum Monte Carlo analysis in Ref. [53]. Symbols are TTN results for two different system sizes as function of the maximum bond dimension. Dashed lines represent RBM accuracy from Ref. [13] for the 10×10 system and different hidden variable densities $\alpha \in [1, 32]$.

$\alpha = 16$ (whose bond-dimension scale as $\chi_{RBM} \sim 2^{16 \cdot 64}$). Let us stress that, this huge difference in the bond dimension scaling, suggests that such parameter is not a sensible measure of complexity for a Neural Network.

At this point, we may wonder how well different representations reproduce two-point correlation functions. Due to the $SU(2)$ symmetry of the Heisenberg Hamiltonian, we expect all correlations $\langle \hat{\sigma}_i^\gamma \hat{\sigma}_j^\gamma \rangle_c$ being independent of γ when evaluated in the exact ground state. In the TTN framework, we enforced $U(1)$ symmetry along the $\hat{z}$ axis which provides $\langle \hat{\sigma}_j^{x,y} \rangle = 0$ thus the connected correlations are more accurate in the $\hat{x}$-$\hat{y}$ plane; we therefore compute correlations along the $\hat{x}$ axis. In the RBM case, we considered correlations in the $\hat{z}$ axis, since by construction they are more accurate and easier to measure here; again in this case, the RBM correlations have been obtained by sapling over 10^6 configurations the optimised wave function in Ref. [13]. With this prescription we are sure to compare the best estimates in both representations.

In Fig. 7 we show the TTN correlations for the 10×10 size and different bond dimensions, and compare them to the RBM with $\alpha = 1$ and 2. It is clear that correlations are growing and getting better with an increasing number of variational parameters. However, the insight from this comparison is twofold: (i) the exponentially large bond-dimension of the RBM is not a guarantee for this representation to be able to encode power-law correlations in critical 2D short-range interacting models and thus to overtake a Tensor Network representation based on finite bond-dimension; (ii) when both TTN and RBM get the same accuracy in the energy (i.e. $\alpha = 1$ and $\chi = 340$ for the 10×10 size), TTN is more accurate for characterising the correlation functions. However, regarding the latter, we mention that this finding and especially the magnitude of the difference in characterising the correlation functions may as well be dependent on the model of investigation. Thus, the generalisation of this statement has to be investigated in future work.
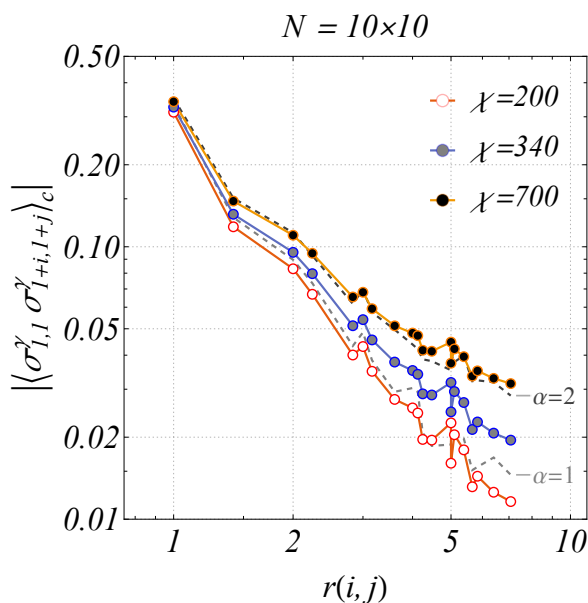
$$N = 10 \times 10$$

Figure 7: Connected correlation function in the TTN representation of the ground state of the 10×10 Heisenberg Hamiltonian for different bond dimensions (symbols) vs the distance $r(i,j) \equiv [i^2 + j^2]^{1/2}$, where $i \leq j \in \{0,5\} \times \{0,5\}$. The dashed grey lines represent correlations obtained by sampling the RBM optimised wave function in Ref. [13].

Let us further point out that in this particular benchmark we may as well exploit the $SU(2)$ symmetry in the TTN simulations thus allowing us to: (1) dramatically increase the accuracy in the estimated energy; (2) reduce the effective bond dimension and thereby the computational time since for non-abelian symmetry we may only work within the symmetry multiplet spaces; (3) drastically improve the connected correlations since we enforce $\langle \hat{\sigma}_j^\gamma \rangle = 0$, as well as $\langle \hat{\sigma}_i^\gamma \hat{\sigma}_j^\gamma \rangle_c$ to be equivalent independent on γ. While (3) can be achieved as well for the RBM when encoding symmetries, (2) does not apply for RBMs. Thus, for RBMs, as for the TTN, we would expect to gain a higher precision in the final results when exploiting the $SU(2)$ symmetry as a result of (3). However, we would not expect such a dramatic increase of the computational time as in the case of the TTN, which in return enables the TTN to achieve higher bond dimensions and thereby to further increase the accuracy additionally to point (3).

4 Discussions

We investigated the efficiency of Neural Network quantum states with respect to Tensor Network quantum states when used to describe the ground-state of critical short-range interacting Hamiltonians both in one and two dimensions.

We pointed out and exploited the "constrained" Tensor Network State (coMPS/coPEPS in 1D/2D) representation of the neural-network wave function. As a matter of fact, RBM and uRBM have very different representations. Even though the coMPS associated to the RBM has an auxiliary dimension which scales exponentially with the number of hidden variables, it still struggles to properly describe the ground-state correlation functions of critical Hamiltonians.

Indeed, in the 1D Ising case, a much smaller coMPS dimension associated to a much more constrained uRBM wave function gives a more accurate description if compared to the RBM

parametrisation. However, it turns out that, for equal auxiliary dimension, standard MPS algorithms give a more acurate description of the many-body ground state. In addition, since the performances of bothe, the uRBM (in the coMPS representation) and the MPS algorithm, scale with the (effective) auxiliary dimension χ, an accurate FSS analysis suggests a possible definition of the *descriptive power* of a variational wave-function: namely, the largest system size which can be faithfully (i.e. within an accuracy goal) described by a given ansatz. As a matter of fact, the descriptive power of the un-constrained MPS outperforms the uRBM descriptive power, when both are used to approximate the ground-state of a critical Ising chain.

In this sense, the exponentially large auxiliary dimension of the coMPS associated to a generic RBM seems not enough to provide a good characterisation of the long-range correlations in the critical Ising quantum chain. In order to obtain more accurate estimates, system-dependent deep neural network states, namely a uRBM with $\ell \gg 1$, have to be properly optimised. Our explicit coMPS representation of the uRMB variational ansatz can be eventually combined with Monte Carlo techniques thus overcoming the limitation, pointed out in Ref. [50], of sampling over the hidden-variable configurations; thus making the Monte Carlo approach also effective for Hamiltonians where the sign-problem occurs. Moreover, an optimised uRBM can be used to optimally initialise convolutional NN algorithms, so as to speed up the computations.

In 2D we compared the RBM representation against the TTN representation when both are used to approximate the ground-state many-body wave function of the two-dimensional Heisenberg Hamiltonian. As expected, both methods are well suited to describe the 2D many-body quantum system. However, when reaching the same level of accuracy in the energy estimate, a TTN is more precise in characterising long-range correlations, even though they are employing a strictly finite bond dimension far below the mapped counterpart of the RBM representation. This, from the one side, leaves no doubt that the exponentially large auxiliary dimension of the RBM does not ensure an adequate descriptive power; from the other side, it leaves us with the open question on how to properly estimate the information which is encoded in a NN, so as to define a proper *measure of complexity* for a neural network quantum state.

5 Acknowledgments

Author contributions. — All authors agreed on the approach to pursue, and contributed to the interpretation of the results and to the writing of the manuscript. L. D. conceived the work. M. C. and T. F. performed the numerical simulations. S. M. supervised and merged the research. The authors declare no competing interests.

A 1D numerical simulations

The numerical simulations for the ground-state optimisation in the Ising quantum chain have been performed by means of different approaches. In particular, the optimisation of the canonical MPS has been done by using the well established DMRG algorithm [25]. In our algorithm,

we fixed the auxiliary dimension χ to remain constant. A preliminary "infinite" size procedure enlarges the system up to the desired linear dimension $N = 80$. Thereafter, the usual "sweeps" procedure locally optimises the MPS wave function. The algorithm is stopped when the energy difference between two consecutive sweeps is less that the machine precision.

In the uRBM approach, we exploited the coMPS representation of the ansatz for the wave-function so as to get very accurate results. If $\mathbf{M}^\sigma$ is the local tensor depending on $2\ell + 1$ real variational parameters $\vec{K}$, we introduced the local operator-dependent transfer-matrix

$$\mathbf{T}_{\hat{O}} = \sum_{\sigma,\sigma'} \langle \sigma' | \hat{O} | \sigma \rangle \, (\mathbf{M}^*)^{\sigma'} \otimes \mathbf{M}^\sigma, \tag{15}$$

which implicitly depends on the variational parameters $\vec{K}$, and globally minimised the energy density of the Ising quantum chain

$$\varepsilon[\vec{K}] = -\frac{\mathrm{Tr}[\mathbf{T}_{\hat{\sigma}^z}\mathbf{T}_{\hat{\sigma}^z}\mathbf{T}_{\hat{I}}^{N-2}] + \lambda\,\mathrm{Tr}[\mathbf{T}_{\hat{\sigma}^x}\mathbf{T}_{\hat{I}}^{N-1}]}{\mathrm{Tr}[\mathbf{T}_{\hat{I}}^N]}. \tag{16}$$

In the simplest case of $\ell = 1$ we used the Mathematica builtin routine `NMinimize` which turns out to be stable and it efficiently converges to the global minimum. However, for larger parameter spaces, namely $\ell = 2$ and 3, the Mathematica routine does not give the expected improvement in the energy minimisation, and it gets stacked on some local minimum. We thus improved the global minimisation by randomly reducing the dimension of the parameter space wherein `NMinimize` has to look for a global minimum. In practice, we proceed in the following way:

1. We randomly initialise a real vector $\vec{K}$ which contains the $2\ell + 1$ variational parameters.

2. We randomly construct a $(2\ell+1)\times(2\ell+1)$ orthogonal matrix $\mathbb{R}$ and define the variational parameter vector in the new basis $\vec{K}' = \mathbb{R}^T \vec{K}$.

3. We pick up three components of the vector $\vec{K}'$ and promote them as variational variables, thus defining the 3-variable dependent vector $\vec{K}'(x,y,z)$ where $\{x,y,z\}$ is a three dimensional subset of variational parameters. We thus transform back the vector to the original basis so as to have $\vec{K}(x,y,z) = \mathbb{R}\vec{K}'(x,y,z)$. We minimise $\varepsilon[\vec{K}(x,y,z)]$ with respect to $\{x,y,z\}$ by using `NMinimize`, thus finding the best parameters $\{x^*,y^*,z^*\}$. If the new optimised energy density is lower than the actual best estimate, we upgrade the solution $\vec{K} = \vec{K}(x^*,y^*,z^*)$.

4. We repeat point 3 for all possible different way of taking three components of the vector $\vec{K}'$, i.e. $\binom{2\ell+1}{3}$. Thereafter, we go back to point 2 and repeat the procedure.

5. We stop the recipe when the difference in the two best energy estimates is less than 10^{-9}.

B 2D Numerical simulations

The simulations for the ground-state computation of the isotropic 2D Heisenberg model have been done using a binary Tree Tensor Networks (TTN). In this approach each tensor within the Network combines two sites to one coarse-grained virtual site (or bond link), resulting in the hierarchical tree structure. The optimisation of the TTN, as well as the calculation of the observables for the optimised ground-state, were obtained following the description for loopless Networks in Ref. [34]. For all simulation the $U(1)$ symmetry has been exploited. Furthermore,
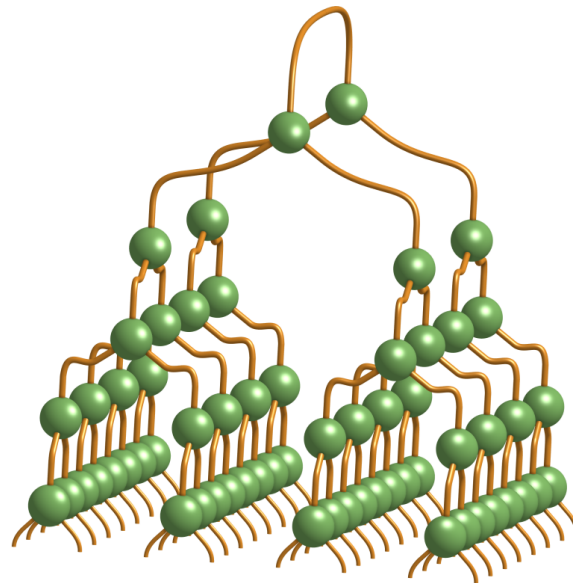
Figure 8: Binary TTN structure for the 8×8 simulations. The tensors (green) each merges two sites of the lower layer to one bond link (brown). The mapping here groups alternatingly in x- and y-direction from layer to layer starting from the physical sites of the 8×8 lattice.

for each simulation the Network was randomly initialised within the zero-magnetisation symmetry sector and within the given bond dimension χ.

For the 10×10 simulations, the physical sites $j \in \{1, \ldots, 100\}$ of the TTN were assigned in a zig-zag pattern to the two-dimensional lattice, such that the lattice site (x, y) (with $x, y \in \{1, \ldots, 10\}$) is mapped to the TTN site $j = x + 10 \cdot (y - 1)$. Thereby, the system is coarse-grains in x-direction first at the lower layers of the tree, and afterwards at the upper layers in y-direction. Thus the simulation is biased towards the x-direction as the topology of the TTN is not well suited to capture correlations in y-directions. This makes the 10×10 system size in general not ideal for a TTN approach.

In the case of the 8×8 system size, the TTN was arranged, such that the grouping within the network is done in an alternating form from layer to layer, as depicted in Fig. 8. Thus the Tensors in the TTN is coarse graining the system in local plaquettes and thereby better capture the correlations within these plaquettes. This mapping leads to a more precise description, which can be observed in the energy being an order of magnitude more accurate (see Fig. 6).

References

[1] C. M. Bishop, *Neural networks for pattern recognition*, Clarendon Press, Oxford, NewYork (1996).

[2] S. S. Haykin, *Neural networks and learning machines*, Vol. 3, Pearson Upper Saddle River, NJ, USA (2009).

[3] M. A. Nielsen, *Neural networks and deep learning*, Determination Press (2015).

[4] Y. LeCun, Y. Bengio and G. Hinton, *Deep learning*, Nature **521**, 436 (2015), doi:10.1038/nature14539.

[5] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*, MIT Press (2016).

[6] D. J. Amit, H. Gutfreund and H. Sompolinsky, *Spin-glass models of neural networks*, Phys. Rev. A **32**, 1007 (1985), doi:10.1103/PhysRevA.32.1007.

[7] P. Mehta and D. J. Schwab, *An exact mapping between the Variational Renormalization Group and Deep Learning* (2014), arXiv:1410.3831.

[8] H. W. Lin, M. Tegmark and D. Rolnick, *Why does deep and cheap learning work so well?*, J. Stat. Phys. **168**, 1223 (2017), doi:10.1007/s10955-017-1836-5.

[9] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto and L. Zdeborová, *Machine learning and the physical sciences*, Rev. Mod. Phys. **91**, 045002 (2019), doi:10.1103/RevModPhys.91.045002.

[10] G. Torlai and R. G. Melko, *Learning thermodynamics with Boltzmann machines*, Phys. Rev. B **94**, 165134 (2016), doi:10.1103/PhysRevB.94.165134.

[11] L. Huang and L. Wang, *Accelerated Monte Carlo simulations with restricted Boltzmann machines*, Phys. Rev. B **95**, 035105 (2017), doi:10.1103/PhysRevB.95.035105.

[12] J. Liu, Y. Qi, Z. Yang Meng and L. Fu, *Self-learning Monte Carlo method*, Phys. Rev. B **95**, 041101 (2017), doi:10.1103/PhysRevB.95.041101.

[13] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, Science **355**, 602 (2017), doi:10.1126/science.aag2302.

[14] D.-L. Deng, X. Li and S. Das Sarma, *Machine learning topological states*, Phys. Rev. B **96**, 195145 (2017), doi:10.1103/PhysRevB.96.195145.

[15] Y. Nomura, A. S. Darmawan, Y. Yamaji and M. Imada, *Restricted Boltzmann machine learning for solving strongly correlated quantum systems*, Phys. Rev. B **96**, 205152 (2017), doi:10.1103/PhysRevB.96.205152.

[16] G. Fabiani and J. Mentink, *Investigating ultrafast quantum magnetism with machine learning*, SciPost Phys. **7**, 004 (2019), doi:10.21468/SciPostPhys.7.1.004.

[17] S. Pilati, E. M. Inack and P. Pieri, *Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines*, Phys. Rev. E **100**, 043301 (2019), doi:10.1103/PhysRevE.100.043301.

[18] D. Perez-Garcia, F. Verstraete, M. M. Wolf and J. I. Cirac, *Matrix product state representations* (2006), arXiv:quant-ph/0608197.

[19] R. Orús, *A practical introduction to tensor networks: Matrix product states and projected entangled pair states*, Ann. Phys. **349**, 117 (2014), doi:10.1016/j.aop.2014.06.013.

[20] R. Orús, *Advances on tensor network theory: symmetries, fermions, entanglement, and holography*, Eur. Phys. J. B **87**, 280 (2014), doi:10.1140/epjb/e2014-50502-9.

[21] G. Vidal, *Entanglement renormalization: an introduction*, in Understanding Quantum Phase Transitions, edited by Lincoln D. Carr, Taylor & Francis, Boca Raton (2010).

[22] F. Verstraete, V. Murg and J. I. Cirac, *Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems*, Adv. Phys. **57**, 143 (2008), doi:10.1080/14789940801912366.

[23] J. I. Cirac and F. Verstraete, *Renormalization and tensor product states in spin chains and lattices*, J. Phys. A: Math. Theor. **42**, 504004 (2009), doi:10.1088/1751-8113/42/50/504004.

[24] M. J. Hartmann, J. Prior, S. R. Clark and M. B. Plenio, *Density matrix renormalization group in the Heisenberg picture*, Phys. Rev. Lett. **102**, 057202 (2009), doi:10.1103/PhysRevLett.102.057202.

[25] U. Schollwöck, *The density-matrix renormalization group in the age of matrix product states*, Ann. Phys. **326**, 96 (2011), doi:10.1016/j.aop.2010.09.012.

[26] U. Schollwöck, *The density-matrix renormalization group: a short introduction*, Phil. Trans. R. Soc. A. **369**, 2643 (2011), doi:10.1098/rsta.2010.0382.

[27] S. Sachdev, *Tensor networks—a new tool for old problems*, Physics **2**, 90 (2009), doi:10.1103/Physics.2.90.

[28] J. Eisert, *Entanglement and tensor network states* (2013), arXiv:1308.3318.

[29] G. Evenbly and G. Vidal, *Tensor network states and geometry*, J. Stat. Phys. **145**, 891 (2011), doi:10.1007/s10955-011-0237-4.

[30] J. C. Bridgeman and C. T. Chubb, *Hand-waving and interpretive dance: an introductory course on tensor networks*, J. Phys. A: Math. Theor. **50**, 223001 (2017), doi:10.1088/1751-8121/aa6dc3.

[31] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao and D. P. Mandic, *Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions*, Found. Trends Mach. Learn, **9**, 249 (2016), doi:10.1561/2200000059.

[32] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, M. Sugiyama and D. P. Mandic, *Tensor networks for dimensionality reduction and large-scale optimization: Part 2 Applications and future perspectives*, Found. Trends Mach. Learn, **9**, 249 (2017), doi:10.1561/2200000067.

[33] S. Montangero, *Introduction to tensor network methods*, Springer International Publishing (2018).

[34] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi and S. Montangero, *The tensor networks anthology: Simulation techniques for many-body quantum lattice systems*, SciPost Phys. Lect. Notes 8 (2019), doi:10.21468/SciPostPhysLectNotes.8.

[35] A. N. Kolmogorov, *On representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables*, Dokl. Akad. Nauk SSSR **108**, 179 (1956).

[36] K. Hornik, *Approximation capabilities of multilayer feedforward networks*, Neural Netw. **4**, 251 (1991), doi:10.1016/0893-6080(91)90009-T.

[37] L. Younes, *Synchronous Boltzmann machines can be universal approximators*, Appl. Math. Lett. **9**, 109 (1996), doi:10.1016/0893-9659(96)00041-9.

[38] N. Le Roux and Y. Bengio, *Representational power of restricted Boltzmann machines and deep belief networks*, Neural Comput. **20**, 1631 (2008), doi:10.1162/neco.2008.04-07-510.

[39] G. Montufar and N. Ay, *Refinements of universal approximation results for deep belief networks and restricted Boltzmann machines* (2010), arXiv:1005.1593.

[40] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez and J. Ignacio Cirac, *Neural-network quantum states, string-bond states, and chiral topological states*, Phys. Rev. X **8**, 011006 (2018), doi:10.1103/PhysRevX.8.011006.

[41] J. Chen, S. Cheng, H. Xie, L. Wang and T. Xiang, *Equivalence of restricted Boltzmann machines and tensor network states*, Phys. Rev. B **97**, 085104 (2018), doi:10.1103/PhysRevB.97.085104.

[42] S. R. Clark, *Unifying neural-network quantum states and correlator product states via tensor networks*, J. Phys. A: Math. Theor. **51**, 135301 (2018), doi:10.1088/1751-8121/aaaaf2.

[43] W. Huggins, P. Patil, B. Mitchell, K. Birgitta Whaley and E. Miles Stoudenmire, *Towards quantum machine learning with tensor networks*, Quantum Sci. Technol. **4**, 024001 (2019), doi:10.1088/2058-9565/aaea94.

[44] L. Pastori, R. Kaubruegger and J. Carl Budich, *Generalized transfer matrix states from artificial neural networks*, Phys. Rev. B **99**, 165123 (2019), doi:10.1103/PhysRevB.99.165123.

[45] E. Robeva and A. Seigal, *Duality of graphical models and tensor networks* (2017), arXiv:1710.01437.

[46] I. Glasser, N. Pancotti and J. Ignacio Cirac, *From probabilistic graphical models to generalized tensor networks for supervised learning* (2018), arXiv:1806.05964.

[47] D.-L. Deng, X. Li and S. Das Sarma, *Quantum entanglement in neural network states*, Phys. Rev. X **7**, 021021 (2017), doi:10.1103/PhysRevX.7.021021.

[48] M. B. Hastings, *An area law for one-dimensional quantum systems*, J. Stat. Mech. P08024 (2007), doi:10.1088/1742-5468/2007/08/P08024.

[49] J. Eisert, M. Cramer and M. B. Plenio, *Colloquium: Area laws for the entanglement entropy*, Rev. Mod. Phys. **82**, 277 (2010), doi:10.1103/RevModPhys.82.277.

[50] E. M. Inack, G. E. Santoro, L. Dell'Anna and S. Pilati, *Projective quantum Monte Carlo simulations guided by unrestricted neural network states*, Phys. Rev. B **98**, 235145 (2018), doi:10.1103/PhysRevB.98.235145.

[51] K. Choo, T. Neupert and G. Carleo, *Two-dimensional frustrated J1-J2 model studied with neural network quantum states*, Phys. Rev. B **100**, 125124 (2019), doi:10.1103/PhysRevB.100.125124.

[52] G. Vidal, *Efficient simulation of one-dimensional quantum many-body systems*, Phys. Rev. Lett. **93**, 040502 (2004), doi:10.1103/PhysRevLett.93.040502.

[53] A. W. Sandvik, *Finite-size scaling of the ground-state parameters of the two-dimensional Heisenberg model*, Phys. Rev. B **56**, 11678 (1997), doi:10.1103/PhysRevB.56.11678.

[54] Y.-Y. Shi, L.-M. Duan and G. Vidal, *Classical simulation of quantum many-body systems with a tree tensor network*, Phys. Rev. A **74**, 022320 (2006), doi:10.1103/PhysRevA.74.022320.

[55] L. Tagliacozzo, G. Evenbly and G. Vidal, *Simulation of two-dimensional quantum systems using a tree tensor network that exploits the entropic area law*, Phys. Rev. B **80**, 235127 (2009), doi:10.1103/PhysRevB.80.235127.

[56] V. Murg, F. Verstraete, Ö. Legeza and R. M. Noack, *Simulating strongly correlated quantum systems with tree tensor networks*, Phys. Rev. B **82**, 205105 (2010), doi:10.1103/PhysRevB.82.205105.

[57] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte and S. Montangero, *Fractional quantum Hall effect in the interacting Hofstadter model via tensor networks*, Phys. Rev. B **96**, 195123 (2017), doi:10.1103/PhysRevB.96.195123.

7

Conclusion and Outlook

This thesis focused on the efficient simulation of quantum many-body systems in higher dimensions. In particular, a Tree Tensor Network (TTN) has been developed capable of solving equilibrium problems and applied to various systems benchmarking its correctness and further unveiling novel physical insights unexplored before.

In this context, Chapt. 2 provided an overview of selected topics around Tensor Network states and algorithms building the foundation of the presented development. It mathematically introduced the fundamental building blocks of a Tensor Network, namely the tensors, together with their basic operations, before formally introducing the concept of Tensor Network states itself. In particular, this chapter highlighted that the developments of Tensor Network geometries for two- or even higher dimensional systems is still ongoing with the TTN and PEPS being the most prominent candidates. The main challenge thereby lays in the combination of accuracy and scalability required for efficiently simulating quantum many-body systems with Tensor Networks in higher dimensions. To overcome this limitation, Tensor Network states shall by construction satisfy the same entanglement bounds under real-space bi-partitions as the physical states they represent and at the same time, their algorithmic complexity shall be as low as possible. In Sec. 2.2, it was illustrated that finding this balance in higher dimensions between sufficiently satisfying the entanglement bounds and keeping the numerical complexity feasibly low poses an infeasible challenge to the current state-of-the-art Tensor Networks.

With this paradigm in mind, Chapt. 3 formally and technically discussed the TTN geometry and its implementation for high-dimensional quantum systems. In this regard, Sec. 3.2 elaborated on how the structure of a TTN is key for a successful analysis of systems in two spatial dimensions and beyond. At this point, the thesis introduced a technically non-trivial mapping enabling an efficient representation of quantum states in higher dimensions via TTN and further provided a hands-on illustration in a step-by-step example for a ground-state search of a (4×4)-system.

In Chapt. 4, this TTN has been applied to study a Lattice Gauge Theory (LGT) approximating the low-energy behaviour of quantum electrodynamics in two dimensions. Therein, we demonstrated the successful application of TTNs on LGTs at finite density, which belongs

to the set of fundamental problems in which Quantum Monte Carlo (QMC) simulations suffer from the sign problem. Thus, this chapter has shown the potential of TTNs to accurately describe systems in High-Energy physics which in their full extend are used to promote the understanding of fundamental elements of the universe. Further, this application illustrated that TTNs are, in general, a highly promising candidate for the analysis of high-dimensional quantum many-body systems. However, as Sec. 3.2.2 highlighted, the TTN does not capture the Area-Law by its structure for two, or even higher, dimensions, even with the aforementioned mapping, and thus it eventually fails to faithfully represent a complex high-dimensional quantum many-body state.

At this point, we introduced a novel Tensor Network which *augments* the TTN with a layer of disentanglers, in a way that it is indeed able to encode the area law by construction. Thus, this *augmented Tree Tensor Network* (aTTN) offers better scalability and thereby overcomes the fundamental limitations of TTNs for high-dimensional systems. In particular, in Chapt. 5, we illustrated that this novel technique outperforms competing methods and reaches unprecedented system sizes in two-dimensions making it the ideal candidate to, among other applications, benchmark and validate state-of-the-art quantum technologies, such as quantum simulators or quantum computers.

Finally, Chapt. 6 presented a publication in which we compared different variational methods used as state representation. In particular, we analysed one-dimensional and two-dimensional critical systems with different Tensor Network approaches and Neural Network (NN) approaches and on top introduced a novel mapping that illustrates the mathematical connection between the networks.

In a scientific perspective, the developments achieved in the course of this thesis represent an encouraging step for efficiently simulating the quantum many-body systems in high-dimensions. The results presented with the introduction of the aTTN, in particular, show that Tensor Networks are indeed able to efficiently capture the underlying entanglement properties of the quantum states in two dimensional systems. This key step provides the basis for future Tensor Network analysis to investigate quantum many-body phenomena inaccessible before. Since the thesis touches on several important areas, the following outlook is categorised into different sections.

Augmented Tree Tensor Network

We introduced the aTTN as a novel Tensor Network geometry to successfully tackle the quantum many-body problem in higher dimensions. In particular, we augmented the well-established Tree Tensor Network to reproduce area law for the underlying system. We illustrated the technical insights into the aTTN required for its implementation, including the structure of the network, its optimisation and the calculation of observables. We further proved that this *augmented TTN* outperforms the ordinary TTN with increasing system size for two-dimensional systems, showing its power to represent the ground-state of critical quantum many-body systems in unprecedented sizes up to 32×32. Due to the efficient scaling for the aTTN of $\mathcal{O}\left(m^4\right)$, we can reach a higher bond dimension m compared to alternative methods for 2D, such as PEPS and MERA, and thereby thanks to the augmented layer of properly engineered disentanglers, reach higher precision and larger system sizes.

In fact, the most crucial point of this technique is to find the best network structure for a given Hamiltonian, i.e. the right number of disentanglers and their exact position within the system. In Chapt. 5, the thesis presented a solid, general and scalable procedure

of positioning the disentanglers which is applicable for arbitrary systems in two and three dimensions. However, even though this procedure is a generic strategy that, as proven within the thesis, leads to great improvements compared to previous approaches, it might not be the best one for every individual system. This is particularly the case for applications where one could exploit further physical information from the system, such as interaction positions and coupling strengths. Therefore, a promising line of future developments is to flexibly adapt the network structure, and in particular the positioning of the disentanglers, based on the interactions of the system to analyse. As an example, this might turn out to be beneficial for analysing phenomena in quantum chemistry where the interactions are governed by the molecular structure [178].

Additionally, this novel Tensor Network can be adapted in future work to explore aspects of quantum many-body systems beyond ground state properties, such as time evolution and finite temperature behaviour, which presents a highly promising potential of the aTTN yet to be unveiled. In particular, real and imaginary time evolution can be implemented either along the line of Ref. [217] following the time-evolution procedure for MERA or of Ref. [143] for the TTN. Further, the aTTN may be extended to an *aTTN Density Operator* (aTTN-DO) following the concept of Ref. [159] towards simulating finite temperature states and dynamics in high-dimensional open quantum systems. These implementations are technically and numerically challenging but almost straightforward generalizations of the original ideas presented in the previously mentioned references.

Comparing Quantum State representations

Chapt. 6 presented a comparison between Tensor Networks and Neural Network states as representations of a quantum many-body wave function. In many cases, it has already been shown that Neural Network states can be mapped into a Tensor Network with exponential bond-dimension [213]. On the other hand, we illustrated a mathematical connection between the networks and introduced a novel mapping. We further pointed out that comparing quantum state representations in higher dimensions, in general, is a non-trivial task for which sound measures have to be found in the future in order to provide a fair comparison. In particular, many difficulties arise when naively comparing neural network states to exponential large Tensor networks since the actual descriptive power of those methods may not always be straightforwardly measurable.

Since both are variational methods, one may argue that they can easily be compared by the obtained energy with a fixed number of variational parameters. However, the number of variational parameters in both approaches does not reflect the actual number of free degrees of freedom [218]. This is a well-known problem already discussed in quantum physics when comparing the two methods: The number of variational parameters may even be far from being a valid indicator of the actual degrees of freedom or an indicator of the accuracy in such a comparison. As an example, the MPS gives an exact representation of the ground-state of the AKLT model with bond dimension $m = 2$. However, a fully connected *Restricted Boltzmann Machine* (RBM), i.e. a Neural Network-based approach, would require $\mathcal{O}(N^2)$ hidden units to find this exact representation [219]. The Laughlin wave-function, on the other hand, can be efficiently represented by an RBM using $N(N-1)/2$ hidden units [219, 220] while an MPS would requires a quasi-exponentiallylarge bond-dimension of $\chi \sim e^{acL}(cL)^{-1/2}$

to simulate the system at fixed accuracy [221]. Thus, comparing the number of variational (or trainable) parameters may not lead to a proper comparison of their descriptive power in general.

This issue presents an important point for future work which is not fully addressed yet. In this comparison, the open question is how to properly estimate the information which is encoded in a Neural Network, or to define a proper *measure of complexity* for a Neural Network quantum state. However, the underlying lack of a profound measure to compare variational methods as state representations extends to all methods including Tensor Networks and QMC. In fact, by scaling up the computational resources the results can be improved for all of the variational approaches, thus putting a high focus on the actual implementation and the used computational resources. Consequently, it is, in general, not trivial to tell from a pure information-theory-based point-of-view which approach is the best-suited for a certain problem when going beyond Tensor Networks.

Lattice Gauge Theory

A further highly encouraging future direction of research unveiled by our work is to be found in the analysis of Lattice Gauge Theories. As demonstrated in Chapt. 4, Tensor Network can be successfully applied to the study of two-dimensional Lattice Gauge Theories. In particular, the developed concept takes advantage of the quantum link formulation of Lattice Gauge Theories and the fermionic rishon representation of the quantum links. Thereby, the TTN algorithm is, among others, capable to find equilibrium properties of a lattice QED system within its first compact spin representation $S = 1$ for system sizes up to 16×16. By exploring systems at finite density, the TTN has proven itself to be a powerful tool to obtain a non-perturbative description of a Lattice Gauge Theory in regimes where QMC suffers from the sign-problem. This capability of successfully simulating Lattice Gauge Systems of the TTN in higher dimensions is further demonstrated in App. A which presents an LGT analysis for three-dimensional systems.

The potential future development of Tensor Networks for this application is very encouraging for theoretical and experimental physics. On the one hand, exploiting the full potential of High-Performance Computing, and the aTTN approach, could enable to extend the presented analyses towards *(i)* the simulation of larger system sizes allowing physical predictions in the spatial continuous limit, *(ii)* the analysis of higher-order link representations $S \to \infty$ to investigate the continuum field or large-S limit, and *(iii)* more complex Abelian and non-Abelian Lattice Gauge Theories, such as quantum chromodynamics or, as a particular case, the Standard model. Some of those studies have already been performed for one-dimensional systems [222, 223] while crucially, these advanced studies for high-dimensional systems could even turn out to be useful for promoting our theoretical understanding of the universe.

On the other hand, quantum link model formulations are frequently studied in experimental physics with the intention to simulate LGTs on quantum hardware [224]. Such experimental quantum simulations have recently been realised for one-dimensional LGTs by the means of trapped ions, ultracold atoms, or Rydberg atoms in optical lattices [225–229]. As mentioned above for the aTTN, the development of these quantum simulators seeks stable benchmarking tools, such as the developed Tensor Network approach of this thesis, in order to faithfully verify the experimentally realised simulators.

Hybrid quantum technologies

As mentioned before, Tensor Networks offer a powerful tool to cross-validate quantum technologies since they can simulate the underlying quantum system on a classical computer. The connection between Tensor Networks and quantum technologies, in particular quantum computers, is even more profound than offering a benchmarking possibility. Tensor Networks can as well be interpreted as a particular class of quantum circuits [100], thus opening the pathway to more connected hybrid approaches between Tensor Network-based quantum circuits and quantum computation. Such a classical-quantum hybrid approach has already been proposed as a tool for solving supervised learning problems in Machine Learning where a Tensor Network on classical computers is used for the training process while the classification is realised by a quantum system [100]. Using an MPS as Tensor Network and single photons as a quantum device for the prediction, this approach has already been implemented experimentally proving the feasibility of such hybrid systems [230]. However, while these references provide encouraging results and illustrate a very interesting field of research, their approach by design restricts the information content of the MPS, and thereby its accuracy, which sets a clear limitation for the practical application to general problems in Machine Learning. Using more advanced Tensor Network methods, such as the aTTN, and further investigating the fundamental connections between Tensor Networks and quantum circuits can overcome this limitation in future applications and potentially lead to efficient hybrid quantum technologies.

Such hybrid approaches can further be brought in the more general context of solving classical optimisation problems. As an example, App. B presents an application of Tensor Networks within the optimisation procedure of the dose distribution for an IMRT cancer treatment. Foreseeing new applications of quantum-inspired techniques to the solution of classical optimization problems, this study may open the pathway to the application of quantum technologies to cancer treatment. As an example, this can be realised through the application of hybrid quantum-classical optimization algorithms [101], once such a quantum computer hardware can be scaled up correspondingly. Thus, this might be another promising direction in which the developed TTN approach might play a role to further enhance the current capabilities of fighting cancer in the future.

Technical advancement

Last but not least, the presented Tensor Network implementation itself still has a high potential to be heavily improved by exploiting the full capabilities of High-Performance Programming and Computing. So far, the TTN and aTTN code is only parallelised via OpenMP for the local optimisation procedure, which is not scalable with system size for HPC. However, there are possibilities to improve this underlying scalability with further implementations aiming to exploit the full potential of HPC systems. As discussed in Sec. 3.3.3, the parallel DMRG algorithm proposed in Ref. [180] can be adapted towards the TTN structure and implemented accordingly. Following this path, the code could be parallelised using MPI and OpenMP to run efficiently on several nodes. This implementation would enable to scale the parallelisation with system size and drastically reduces the CPU time for large system sizes. Another technical step to improve the TTN simulations is the implementation of GPUs since, at the current stage, the development is limited to CPUs. Ref. [231] shows that such implementation at the low level of tensor contractions can provide a speedup by a factor of up to two orders of magnitude. Thus, implementing and testing both improvements, the advanced parallelisation and the usage of GPUs, helps to overcome the current technical limitations and

enables to efficiently investigate high-dimensional quantum many-body systems. Thereby, we will be able to set new standards in the numerical simulation of these systems and reveal new insights in the affected areas, including Quantum Information, Condensed Matter Physics, High Energy Physics, Quantum Computing, and Machine Learning.

A

Applications for three-dimensional Lattice Gauge Systems

As referred to in Chapt. 4, the following manuscript appends the application of Tensor Network algorithms on models in the field of Lattice Gauge Theory (LGT). In particular, this work extends the publication presented in Chapt. 4 analysing three-dimensional lattice gauge problems with a Tree Tensor Network. Due to the fundamental challenge in representing high-dimensional systems with Tensor Networks (as discussed in Sec. 2.2), this publications was the first attempt to simulate a Lattice Gauge Theory in (3+1) dimensions via Tensor Networks. Thus, it shows that Tensor Networks, and in particular Tree Tensor Networks, can be a fruitful direction to solve fundamental long-standing problems. Additionally, Tensor Networks provide an ideal cross-verification tool for near-future quantum simulations of lattice gauge theories available shortly (see Chapt. 5). At the state of submission of this thesis, this second manuscript is accepted for publication in *Nature Communications* and in the process of being published. In the meantime, it is available on arXiv [144].

Lattice Quantum Electrodynamics in (3+1)-dimensions at finite density with Tensor Networks

Giuseppe Magnifico,[1,2] Timo Felser,[1,2,3] Pietro Silvi,[4,5] and Simone Montangero[1,2]

[1]*Dipartimento di Fisica e Astronomia "G. Galilei", Università di Padova, I-35131 Padova, Italy.*
[2]*Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Padova, I-35131 Padova, Italy.*
[3]*Theoretische Physik, Universität des Saarlandes, D-66123 Saarbrücken, Germany.*
[4]*Center for Quantum Physics, and Institute for Experimental Physics,*
University of Innsbruck, A-6020 Innsbruck, Austria.
[5]*Institute for Quantum Optics and Quantum Information,*
Austrian Academy of Sciences, A-6020 Innsbruck, Austria.
(Dated: November 18, 2020)

Gauge theories are of paramount importance in our understanding of fundamental constituents of matter and their interactions. However, the complete characterization of their phase diagrams and the full understanding of non-perturbative effects are still debated, especially at finite charge density, mostly due to the sign-problem affecting Monte Carlo numerical simulations. Here, we report the Tensor Network simulation of a three dimensional lattice gauge theory in the Hamiltonian formulation including dynamical matter: Using this sign-problem-free method, we simulate the ground states of a compact Quantum Electrodynamics at zero and finite charge densities, and address fundamental questions such as collective phases, confinement, and charge screening.

Ranging from high-energy particle physics (Standard Model) [1–3] to low-temperature condensed matter physics (spin liquids, quantum Hall, high-T_c superconductivity) [4, 5], gauge theories constitute the baseline in our microscopical description of the universe and are a cornerstone of contemporary scientific research. Yet, capturing their many-body body behavior beyond perturbative regimes, a mandatory step before experimentally validating these theories, often eludes us [6]. One for all, the quark confinement mechanism in Quantum Chromodynamics (QCD), a founding pillar of the Standard Model which have been studied for almost a century, is still at the center of current research efforts [7–12]. Indeed, a powerful numerical workhorse such as Monte-Carlo simulations [13–15], capable of addressing discretized lattice formulations of gauge theories [11, 16], struggles in highly interesting regimes, where matter fermions and excess of charge are concerned, due to the infamous sign problem [17]. In recent years, a complementary numerical approach, Tensor Networks (TN) methods, have found increasing applications for studying low-dimensional Lattice Gauge Theories (LGT) in the Hamiltonian formulation [18, 19]. As tailored many-body quantum state ansätze, TNs are an efficient approximate entanglement-based representation of physical states, capable of efficiently describe equilibrium properties and real-time dynamics of systems described by complex actions, where Monte Carlo simulations fail to efficiently converge [20]. TN methods have proven remarkable success in simulating LGTs in (1+1) dimensions [21–28], and very recently they have shown potential in (2+1) dimensions [29–33]. To date, due to the lack of efficient numerical algorithms to describe high-dimensional systems via TNs, no results are available regarding the realistic scenario of LGTs in three spatial dimensions.

Here, we bridge this gap by numerically simulating the Hamiltonian formulation of (3+1) Quantum Electrodynamics (QED) at zero temperature, via TN ansatz states. We show that, by using the quantum link formalism (QLM) of LGTs [34, 35] and an unconstrained Tree Tensor Network (TTN), we can access multiple equilibrium regimes of the model, including finite charge densities. Precisely, we analyze the ground state properties of quantum-link QED in (3+1)D for intermediate system sizes, up to 512 lattice sites. The matter is discretized as a staggered spinless fermion field on a cubic lattice [16], while the electromagnetic gauge fields are represented on lattice links, and truncated to a compact representation of spin-s. Here we present results from a non-trivial representation for lattice gauge fields (the spin-1 case), with possible generalizations to higher spin requiring only a polynomial overhead in s. Our picture can be similarly adapted to embed non-Abelian gauge symmetries, such as they appear in QCD [28]. Finally, we stress that the truncation of the gauge field is a common step in quantum simulations and computations [36–42], making the presented numerical approach a landmark benchmarking and cross-verification tool for current and future experiments.

Hereafter, by variationally approximating the lattice QED ground state with a TTN, we address a variety of regimes and questions inaccessible before. In the scenario with zero excess charge density, we observe that the transition between the vacuum phase and the charge-crystal phase is compatible with a second-order quantum phase transition [32]. In the limit of zero magnetic coupling, this transition occurs at negative bare masses m_0, but as the coupling is activated, the critical point is shifted to larger, and even positive, m_0 values. To investigate field-screening properties, we also consider the case where two parallel charged plates are placed at a distance (a capacitor). By studying the polarization of the vacuum in the inner volume, we observe an equilibrium string-breaking effect akin to the Schwinger mechanism. Furthermore,
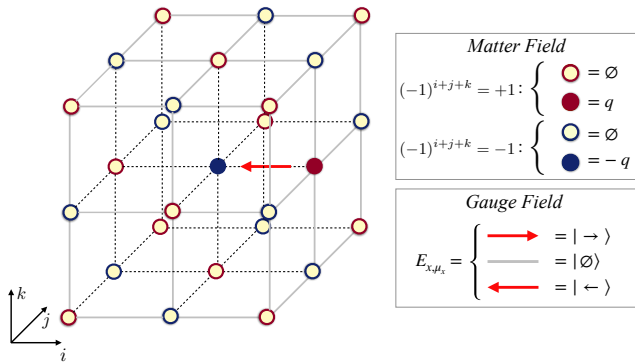
Figure 1: Scheme of the three-dimensional LGT with three electric field levels (spin-1 compact representation). Fermionic degrees of freedom are represented by staggered fermions on sites with different parity: on the even (odd) sites, a full red (blue) circle corresponds to a particle (antiparticle) with positive (negative) charge. As an illustrative example, it is shown a gauge-invariant configuration of matter and gauge fields with one particle and one antiparticle in the sector of zero total charge.

we address the confinement problem by evaluating the binding energies of charged particle pairs pinned at specified distances. Finally, we consider the scenario with a charge imbalance into the system, i.e. at finite charge density, and we characterize a regime where charges accumulate at the surface of our finite sample, analogously to a classic perfect conductor.

I. THE MODEL

Hereafter, we numerically simulate, at zero temperature, the Hamiltonian of $U(1)$ quantum electrodynamics on a finite $L \times L \times L$ three-dimensional simple cubic lattice [16]:

$$\hat{H} = -t \sum_{x,\mu} \left(\hat{\psi}_x^\dagger \hat{U}_{x,\mu} \, \hat{\psi}_{x+\mu} + \text{H.c.} \right) \tag{1a}$$

$$+ m \sum_x (-1)^x \hat{\psi}_x^\dagger \hat{\psi}_x + \frac{g_e^2}{2} \sum_{x,\mu} \hat{E}_{x,\mu}^2 \tag{1b}$$

$$- \frac{g_m^2}{2} \sum_x \left(\Box_{\mu_x,\mu_y} + \Box_{\mu_x,\mu_z} + \Box_{\mu_y,\mu_z} + \text{H.c.} \right) \tag{1c}$$

with $x \equiv (i,j,k)$ for $0 \leq i,j,k \leq L-1$ labelling the sites of the lattice and $\Box_{\mu_\alpha,\mu_\beta} = \hat{U}_{x,\mu_\alpha} \hat{U}_{x+\mu_\alpha,\mu_\beta} \hat{U}_{x+\mu_\beta,\mu_\alpha}^\dagger \hat{U}_{x,\mu_\beta}^\dagger$. Here we adopted the Kogut-Susskind formulation [16], representing fermionic degrees of freedom with a staggered spinless fermion field $\{\hat{\psi}_x, \hat{\psi}_{x'}^\dagger\} = \delta_{x,x'}$ on lattice sites. Their bare mass $m_x = (-1)^x m$ is staggered, as tracked by the site parity $(-1)^x = (-1)^{i+j+k}$, so that fermions on even sites represent particles with positive electric charge $+q$, while holes on odd sites represent anti-particles with negative charge $-q$, as shown in Fig. 1. Charge $\hat{Q}$ conservation is thus expressed as global fermion number $\hat{N}$ conservation, since $\hat{Q} = \sum_x \left(\hat{\psi}_x^\dagger \hat{\psi}_x - \frac{1-(-1)^x}{2} \right) = \hat{N} - L^3/2$.

The links of the 3D lattice are uniquely identified by the couple of parameters (x,μ) where x is any site, μ is one of the three positive lattice unit vectors $\mu_x \equiv (1,0,0)$, $\mu_y \equiv (0,1,0)$, $\mu_z \equiv (0,0,1)$. The gauge fields are defined on lattice links through the pair of operators $\hat{E}_{x,\mu}$ (electric field) and $\hat{U}_{x,\mu}$ (unitary comparator) that satisfy the commutation relation

$$[\hat{E}_{x,\mu}, \hat{U}_{x',\mu'}] = \delta_{x,x'} \delta_{\mu,\mu'} \hat{U}_{x,\mu}. \tag{2}$$

For comfort of notation, we can extend the definition to negative lattice unit vectors via $\hat{E}_{x+\mu,-\mu} = -\hat{E}_{x,\mu}$ and $\hat{U}_{x+\mu,-\mu} = \hat{U}_{x,\mu}^\dagger$.

The Hamiltonian of Eq. (1) consists of four terms: the parallel transporter (1a) describes creation and annihilation of a particle-antiparticle pair, shifting the gauge field in-between to preserve local gauge symmetries. The staggered mass and the electric energy density (1b) are completely local. Finally, the plaquette terms (1c) capture the magnetic energy density, and are related to the smallest Wilson loops along the closed plaquettes along the three planes $x-y$, $x-z$, $y-z$ of the lattice. In dimensionless units ($\hbar = c = 1$), the couplings in Eq. (1) are not independent: They can be expressed as $t = 1/a$, $m = m_0$, $g_e^2 = g^2/a$, $g_m^2 = 8/(g^2 a)$, where a is the lattice spacing, g is the coupling constant of QED and m_0 is the bare mass of particles/antiparticles. The numerical setup allows us to consider the couplings (t, m, g_e, g_m) as mutually independent. We then recover the physical regime of QED by enforcing $g_e g_m = 2\sqrt{2} t$. We also fix the energy scale by setting $t = 1$.

The local $U(1)$ gauge symmetry of the theory is encoded in Gauss's law, whose generators

$$\hat{G}_x = \hat{\psi}_x^\dagger \hat{\psi}_x - \frac{1-(-1)^x}{2} - \sum_\mu \hat{E}_{x,\mu}, \tag{3}$$

are defined around each lattice site x. The sum in Eq. (3) involves the six electric field operators on the links identified by $\pm\mu_x$, $\pm\mu_y$, $\pm\mu_z$. Each $\hat{G}_x$ commutes with the Hamiltonian $\hat{H}$ and the gauge invariant Hilbert space consists of physical many-body quantum states $|\Phi\rangle$ satisfying $\hat{G}_x |\Phi\rangle = 0$ at every site x.

As stressed in the standard Wilson's formulation of lattice QED [11], faithful representations of the $(\hat{E}, \hat{U})$ algebra are infinite-dimensional. A truncation to a finite dimension becomes therefore necessary for numerical simulations with TN methods, which require a finite effective Hilbert dimension at each lattice site. We use the quantum link model (QLM) approach in which the gauge field algebra is replaced by $SU(2)$ spin algebra, i.e. $\hat{E}_{x,\mu} \equiv \hat{S}_{x,\mu}^z$ and $\hat{U}_{x,\mu} \equiv \hat{S}_{x,\mu}^+/s$ for a spin-s representation. This substitution keeps the electric field operator
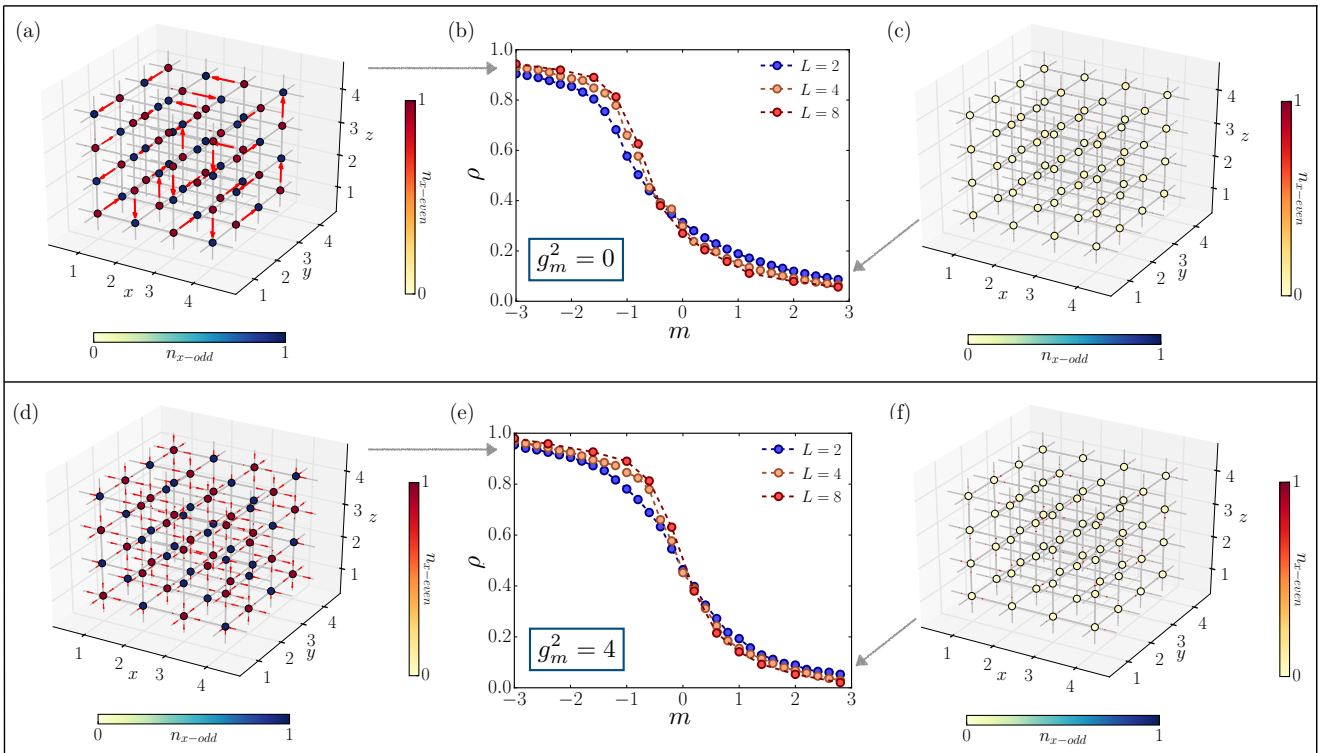
Figure 2: Ground state charge occupation and electric field on links for $m = -3.0$ (a) and $m = 3.0$ (c) and $g_m^2 = 0$. (b) Particle density as a function of m, for different system size L and $g_m^2 = 0$. Ground state charge occupation and electric field on links for $m = -3.0$ (d) and $m = 3.0$ (f) in the presence of magnetic interactions with $g_m^2 = 8/g_e^2 = 4$. (e) Particle density as a function of m, for different system size L and $g_m^2 = 8/g_e^2 = 4$.

hermitian and preserves Eq. (2), but $\hat{U}$ is no longer unitary. Throughout this work, we will select $s = 1$, the smallest representation ensuring a nontrivial contribution of all the terms in the Hamiltonian (see also Fig. 1). This truncation introduces a local energy cutoff based on g_e^2, which in turn requires larger spin s to accurately represent weaker coupling regimes, still potentially accessible via TNs [24].

II. TRANSITION AT ZERO CHARGE

We focus on the zero charge sector, i.e. $\sum_x \psi_x^\dagger \psi_x = \frac{L^3}{2}$, and Periodic Boundary Conditions (PBC). As shown in Fig. 2 (upper panel), for $g_m^2 = 0$ the system undergoes a transition between two regimes, analogously to the (1+1)D and (2+1)D cases [22, 25, 32]: for large positive masses, the system approaches the bare vacuum, while for large negative masses, the system is arranged into a crystal of charges, a highly degenerate state in the semiclassical limit ($t \to 0$) due to the exponential number of electric field configurations allowed. We track this transition by monitoring the average matter density $\rho = \frac{1}{L^3} \sum_x \langle GS| \hat{n}_x |GS \rangle$ where $\hat{n}_x = \frac{1+(-1)^x}{2} - (-1)^x \psi_x^\dagger \psi_x$ is the matter occupation operator and the many-body ground state $|GS\rangle$ has been computed by TTN algorithm

(see Appendices A, B, C for details). Fig. 2(b) displays the result for different sizes L (and $g_e^2/2 = t = 1$), portraying the transition. Panels (a) and (c) display local configurations of matter $\langle \hat{n}_x \rangle$ and gauge sites $\langle \hat{E}_{x,\mu} \rangle$ for $m = -3.0$ and $m = +3.0$ respectively. In the former regime, the algorithm seems to favor a single allowed configuration of gauge fields rather than a superposition of many configuations: This is due to the fact that, when $g_m^2 = 0$, the matrix element that rearranges the configurations occurs at very high perturbative order in $|t/m|$, and is numerically neglected. A finite-size scaling analysis of the transition (see Appendix D) yields results compatible with a II-order phase transition, with the critical point occurring at negative bare masses m.

The same transition appears to be more interesting when we 'activate' the magnetic coupling, by setting $g_m^2 = 8t^2/g_e^2 = 4$ (physical line). The phase at large negative m now appears to be a genuine superposition of many configurations of the electric field, as they are coupled by matrix elements of the order $\sim g_m^2$, kept as numerically relevant by the algorithm. Moreover, the transition is still compatible with a II-order phase transition, and the critical point is shifted to larger m values. This can lead to a critical bare mass m_c that is positive (as we observed $m_c \approx +0.22$ for the case $g_e^2/2 = t = 1$), ultimately making the transition physically relevant.
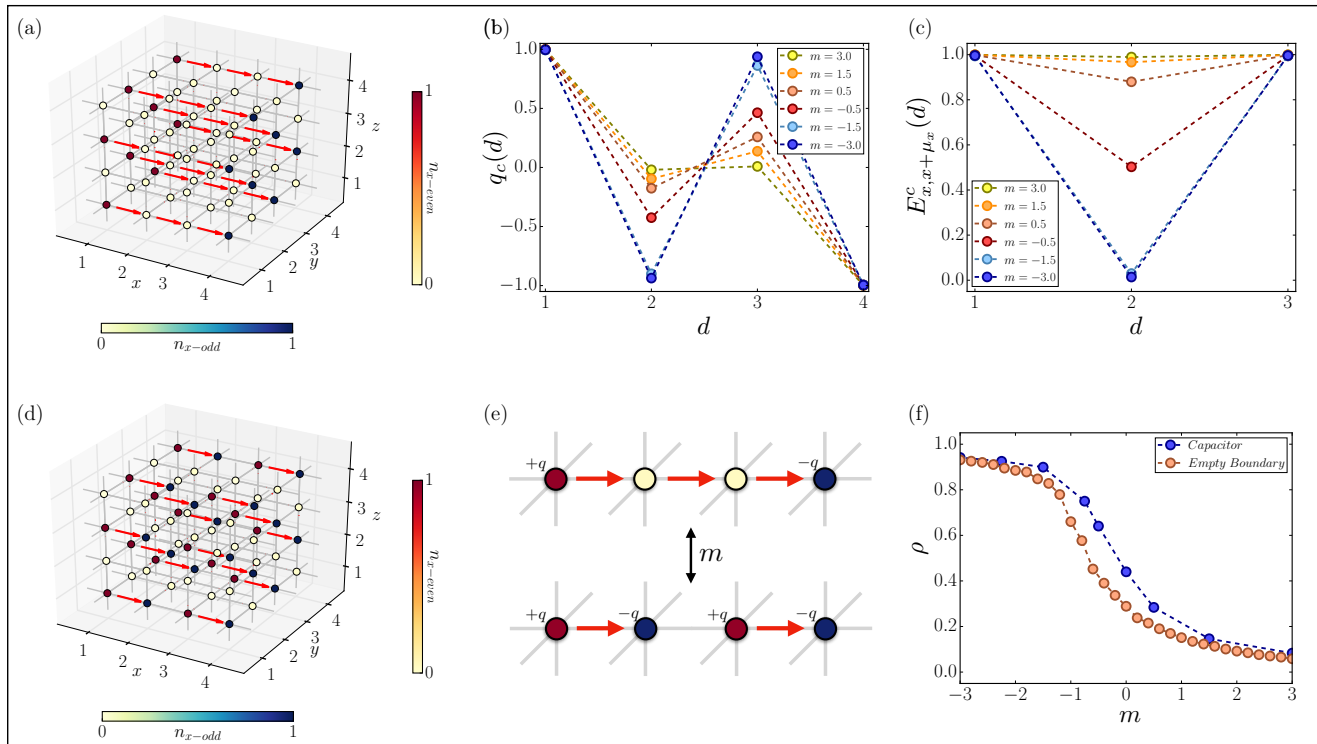
Figure 3: (a) Ground state configuration of the quantum capacitor for $m = 3.0$. (b) Mean charge density on the sites along the transverse direction for different values of m. (c) Mean value of the electric field on the transverse links for different values of m. (d) Ground state configuration of the quantum capacitor for $m = -3.0$. (e) Illustration of the creation of a particle-antiparticle pair along the transverse direction, starting from the initial electric field string generated by the boundary charges. (f) Particle density as a function of m, with a comparison to the case with no boundary charges.

III. QUANTUM CAPACITOR

To investigate field-screening and equilibrium string-breaking properties, we analyze the scenario where two charged plates (an electric capacitor) are placed at the opposite faces of a volume, with open boundary conditions (OBC). In our simulations, we achieve this regime by setting large local chemical potentials on the two boundaries. We expect that for small positive masses m, the vacuum inside the plates will spontaneously polarize to an effective dielectric, by creating particle and antiparticle pairs to screen the electric field from the plates, into an energetically-favorable configuration.

We observe this phenomenon by monitoring the charge density function along the direction μ_x orthogonal to the plates $q_c(d) = \frac{2}{L^2} \sum_{j,k=1}^{L} \langle GS| (-1)^x \hat{\psi}_{(d,j,k)}^{\dagger} \hat{\psi}_{(d,j,k)} |GS\rangle$ as well as the electric field amplitude along μ_x, $E_{x,x+\mu_x}^c(d) = \frac{2}{L^2} \sum_{j,k=1}^{L} \langle GS| \hat{E}_{(d,j,k),(d+1,j,k)} |GS\rangle$, as presented in Fig. 3.

A transition from a vacuum regime to a string-breaking dielectric regime is observed, when driving m from negative to positive. However, here the critical point occurs at positive masses ($m_c > 0$) even at zero magnetic cou-

pling $g_m^2 = 0$, analogously to the (1+1)D case [22]. In conclusion, the charged capacitor can make the phase transition *physical* even when g can not be tuned.

The observed behaviour can be interpreted as an equilibrium counterpart to the *Schwinger mechanism*, a real-time dynamical phenomenon in which the spontaneous creation of electron-positron pairs out of the vacuum is stimulated by a strong external electric field [43]. This could either be potentially verified in experiments or quantum simulations, by means of adiabatic quenches, ramping up the capacitor voltage.

IV. CONFINEMENT PROPERTIES

The (3+1)-dimensional pure compact lattice QED predicts a confining phase at large coupling g [11, 44–47]. This phase, where the magnetic coupling is negligible, is characterized by the presence of a linear potential between static test charges, and is expected to survive at the continuum limit. By decreasing g, the system undergoes a phase transition to the Coulomb phase where the magnetic terms are not negligible and the static charges interact through the $1/r$ Coulomb potential at distance r

[48]. When the gauge field is coupled to dynamical matter ($t \neq 0$ and finite m), new possible scenarios emerge, such as the string-breaking mechanism. Nevertheless, the transition between confined and deconfined phases is still expected to occur [49].

We can investigate this specific scenario with our TN method: we consider a $16 \times 4 \times 4$ lattice and pin two opposite charges via large local chemical potentials at distance r along direction μ_x. The energy $E(r) = V(r) - V(\infty) + 2\epsilon_1 + E_0$ of this ground state comprises: the work $V(r) - V(\infty)$ needed to bring two charges from infinity to distance r, plus twice the excitation energy ϵ_1 of an isolated pinned charge, on top of the dressed-vacuum energy E_0. Therefore we can estimate the interaction potential as $V(r) = E(r) - E_0 + \xi$ where the additive constant ξ does not scale with the volume (while $E(r)$ and E_0 separately do).

The presence of dynamical matter heavily impacts the strong-coupling picture ($g_m^2 \sim 0$), as it can be extrapolated in the semiclassical limit ($t \sim 0$). Here, a particle-antiparticle pair at distance r with, a field-string between them, has an energy

$$E(r) - E_0 = 2m + \frac{g^2}{2}r. \qquad (4)$$

that scales linearly with r. On the contrary, two mesons (neighboring particle-antiparticle pairs) have a flat energy profile

$$E_{pairs} - E_0 = 4m + g^2. \qquad (5)$$

Thus, for any mass m, there is critical distance r_0 above which the string is broken, and formation of two mesons is energetically favorable.

We observe this transition at finite t, as shown in Fig. 4 (bottom panel, $g^2 = 4$). The crossover from the short-range to long-range behavior is still relatively sharp, and the distance r_c at which it occurs strongly depends on the bare mass m. This is in contrast to the weak-coupling regime (top panel, $g^2 = 1/4$), where the potential profile $V(r)$ is smoothly increasing with r, and its slope at short distances disagrees with the string tension ansatz $rg^2/2 +$ const.. Thus our simulations highlight visibly different features between confined and deconfined regimes, even with dynamical matter.

V. FINITE DENSITY

One of the most important features of our numerical approach is the possibility to tackle finite charge-density regimes. In fact, by exploiting the global $U(1)$ fermion-number symmetry, implemented in our TTN algorithms, we can inject any desired charge imbalance into the system, while working under OBC. Fig. 5 shows the results for charge density $\rho = Q/L^3 = 1/4$. In the *vacuum phase* ($m \gg g_e^2/2 \approx t$), we obtain configurations as displayed in panel (a), where the charges are expelled from the bulk,
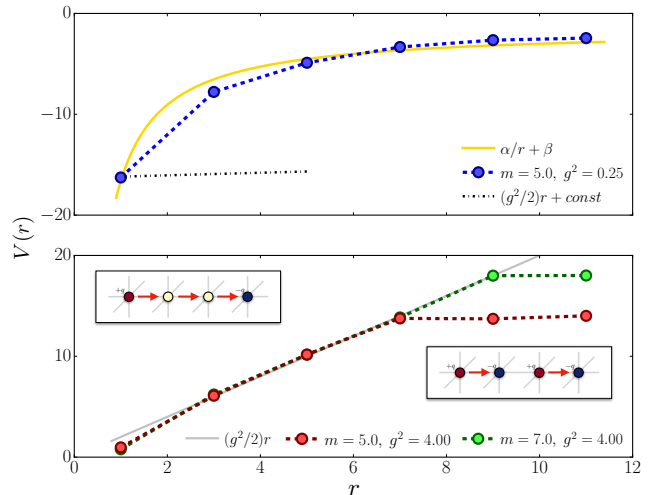

Figure 4: Interaction potential $V(r)$ between two charges of opposite sign as a function of their distance r in the (upper panel) weak coupling regime $g \ll 1$ and (lower panel) strong coupling regime $g \gg 1$.

and stick to the boundaries to minimize the electric field energy of the outcoming fields. To quantify this effect, which can also be interpreted as a field-screening phenomenon, we introduce the surface charge density

$$\sigma(l) = \frac{1}{A(l)} \sum_{x \in A(l)} \langle \psi_x^\dagger \psi_x \rangle \qquad (6)$$

where $A(l)$ contains only sites sitting at lattice distance l from the closest boundary. The deeper we are in the *vacuum phase*, the faster the surface charge decays to zero away from the boundary ($l = 1$). By contrast, close to the transition, the spontaneous creation of charge-anticharge pairs determines a finite charge density of the bulk. Finally, for large negative m, the charge distribution is roughly uniform.

VI. OUTLOOK

We have shown that TN methods can simulate LGT in three spatial dimensions, in the presence of matter and charge imbalance, ultimately exploring those regimes where other known numerical strategies struggle. We have investigated collective phenomena of lattice QED which stand at the forefront of the current research efforts, including quantum phase diagrams, confinement issues, and the string breaking mechanism at equilibrium. We envision the possibility of including more sophisticated diagnostic tools, such as the 't Hooft operators [50] which nicely fit TNs designs, to provide more quantitatively precise answers to the aforementioned open problems.
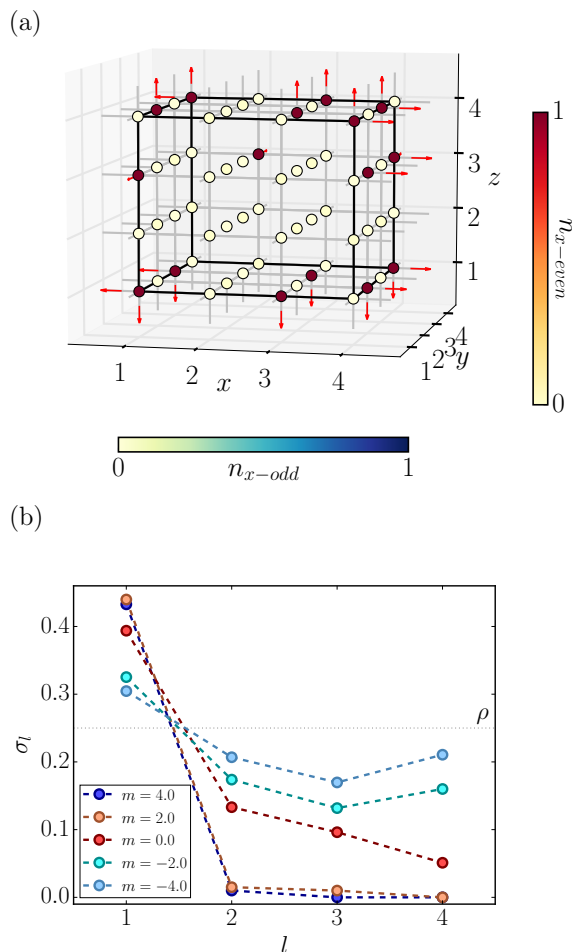
Figure 5: (a) Ground state configuration for $m = 4.0$ at finite charge density $\rho = Q/L = 1/4$. The system is in the global symmetry sector with $Q = 16$ positive charges on the lattice with linear size $L = 4$. (b) Surface charge density σ_l on a cube whose faces are at distance l from the boundaries of the lattice with linear size $L = 8$. The system is in the global symmetry sector with $Q = 128$ positive charges (finite density $\rho = 1/4$).

experimental quantum simulations of one-dimensional LGTs [36–40]. In this framework, numerical methods capable of accessing intermediate sizes, such as TNs, play a fundamental role as a cross-verification toolbox.

VII.   ACKNOWLEDGMENTS

From a theoretical standpoint, our work corroborates the long-term perspective to employ TN methods to efficiently tackle non-perturbative phenomena of LGTs, in high dimensions and in regimes that are out of reach for other numerical techniques. As ansatz states with a refinement parameter chosen by the user, the bond dimension, TTNs are automatically equipped with a self-validation tool: convergence of each quantity with the bond dimension can be verified in polynomial time.

From an experimental point of view, quantum link model formulations are the most studied pathway towards the simulation of LGTs on quantum hardware [41]. The recent developments in low-temperature physics and control techniques, for trapped ions, ultracold atoms and Rydberg atoms in optical lattices, have led to the the first

Appendix A: Fermionic compact representation of local gauge-invariant site

TN simulations of LGT, both Abelian and non-Abelian, are enabled by appropriate finite-dimensional representations of matter and gauge degrees of freedom [51–59], a necessary approach also towards quantum simulation of LGT models [60–63], and potentially useful for MonteCarlo as well [64]. This requirement applies to LGT descending from high-energy quantum field theories [65–68], as well as condensed matter models with emergent gauge fields [69, 70].

In this section we present the construction of the QED gauge-invariant configurations for the local sites that we exploit as computational basis in our TN algorithm.

The use of the spin-1 representation implies that the gauge degrees of freedom on each link of the lattice are represented by three orthogonal eigenstates of the electric field operator:

$$\hat{E}_{x,\mu}|1\rangle = |1\rangle \,, \; \hat{E}_{x,\mu}|0\rangle = 0, \; \hat{E}_{x,\mu}|-1\rangle = -|-1\rangle \,. \text{(A1)}$$

The parallel transporter, that in the spin language corresponds to the raising operator, acts on these states as

$$\hat{U}_{x,\mu}|1\rangle = 0, \; \hat{U}_{x,\mu}|0\rangle = |1\rangle \,, \; \hat{U}_{x,\mu}|-1\rangle = -|0\rangle \,. \text{ (A2)}$$

In the following, in order to obtain a representation of the gauge degrees of freedom that will be useful for constructing our TN ansatz, we employ the local mapping presented in Ref. [32] (see also [71, 72]), generalizing it to the case with three spatial dimensions. This technique is related to the standard rishon formulation of QLM [73–75] and allows us to encode the Gauss's law taking into account the anticommutation relations of the fermionic particles on the lattice.

Let us consider a generic link of the lattice (x,μ) between the two sites x and $x+\mu$: the starting point is the splitting of the gauge field of this link into a pair of *rishon* modes, so that each mode belongs to either one of the two sites. For the $s=1$ case, we can set each rishon mode (or half-link) to be a 3-hardcore fermionic field $\hat{\eta}_{x,\mu}$. Such lattice quantum fields satisfy $\hat{\eta}_{x,\mu}^2 \neq 0$ and $\hat{\eta}_{x,\mu}^3 = 0$. They mutually anticommute at different spatial positions, i.e. $\left\{\hat{\eta}_{x,\mu}, \hat{\eta}_{x',\mu'}^{(\dagger)}\right\} = 0$ for $x \neq x'$ or $\mu \neq \mu'$, and also anticommute with the staggered matter fermionic fields $\left\{\hat{\eta}_{x,\mu}, \hat{\psi}_{x'}^{(\dagger)}\right\} = 0$ [76, 77]. Then, we express the comparator on the link as $\hat{U}_{x,\mu} = \hat{\eta}_{x,\mu}\hat{\eta}_{x+\mu,-\mu}^{\dagger}$. To explicitly build these 3-hardcore fermions for each half-link, we consider two species of standard Dirac fermions $\hat{a}_{x,\mu}$ and $\hat{b}_{x,\mu}$ and we use the following relation:

$$\hat{\eta}_{x,\mu}^{\dagger} = \hat{n}_{x,\mu}^a b_{x,\mu}^{\dagger} + (1 - \hat{n}_{x,\mu}^b)a_{x,\mu}^{\dagger} \quad\quad \text{(A3)}$$

where $\hat{n}_{x,\mu}^a$ and $\hat{n}_{x,\mu}^b$ are the occupation number operators for each species, i.e., $\hat{n}_{x,\mu}^a = \hat{a}_{x,\mu}^{\dagger}\hat{a}_{x,\mu}$ and the same for $\hat{n}_{x,\mu}^b$. For each 3-hardcore mode, these operators act on a three-dimensional local Hilbert space with basis $|0\rangle_{x,\mu}$, $|1\rangle_{x,\mu} = a_{x,\mu}^{\dagger}|0\rangle_{x,\mu}$, $|2\rangle_{x,\mu} = b_{x,\mu}^{\dagger}a_{x,\mu}^{\dagger}|0\rangle_{x,\mu}$. In fact, due to the definition in Eq. (A3), the algebra of the operators $\hat{\eta}_{x,\mu}$ never accesses the fourth state obtained as $b_{x,\mu}^{\dagger}|0\rangle_{x,\mu}$. By using the same representation on the other half-link through the Dirac operators $\hat{a}_{x+\mu,-\mu}^{\dagger}$ and $\hat{b}_{x+\mu,-\mu}^{\dagger}$, we would obtain for the complete link a local space of dimension 9. However, the operator that counts the total number of fermions on the complete link as

$$\hat{L}_{x,\mu} = \hat{n}_{x,\mu}^a + \hat{n}_{x,\mu}^b + \hat{n}_{x+\mu,-\mu}^a + \hat{n}_{x+\mu,-\mu}^b, \quad \text{(A4)}$$

defines a symmetry of the Hamiltonian since it commutes with the operators $\hat{E}_{x,\mu}$ and $\hat{U}_{x,\mu}$. Thus, we can select the sector with $\hat{L}_{x,\mu} = 2$ (two rishons on each full link), reducing the link space to dimension 3 with the basis

$$|\rightarrow\rangle = -|0,2\rangle = \hat{a}_{x+\mu,-\mu}^{\dagger}\hat{b}_{x+\mu,-\mu}^{\dagger}|0\rangle_{x,\mu}|0\rangle_{x+\mu,-\mu},$$
$$|\emptyset\rangle = |1,1\rangle = \hat{a}_{x,\mu}^{\dagger}\hat{a}_{x+\mu,-\mu}^{\dagger}|0\rangle_{x,\mu}|0\rangle_{x+\mu,-\mu}, \quad \text{(A5)}$$
$$|\leftarrow\rangle = |2,0\rangle = \hat{b}_{x,\mu}^{\dagger}\hat{a}_{x,\mu}^{\dagger}|0\rangle_{x,\mu}|0\rangle_{x+\mu,-\mu},$$

where the minus sign in the first element allows the operator $\hat{U}_{x,\mu}$ to act correctly following the properties of Eq. (A2). By using this representation, the electric field finally corresponds to the imbalance of Dirac fermions between the two halves of the link, so that:

$$\hat{E}_{x,\mu} = \frac{1}{2}\left(\hat{n}_{x+\mu,-\mu}^a + \hat{n}_{x+\mu,-\mu}^b - \hat{n}_{x,\mu}^a - \hat{n}_{x,\mu}^b\right). \text{ (A6)}$$

This construction in terms of 3-hardcore fermions allows us to define, for each lattice site, a local basis that directly incorporates the Gauss's law, by constraining in this way the dynamics to the physical states only. This is a crucial point for both numerical and quantum simulations since non-physical states determine an exponential increase in the complexity of the problem.

From the definition of the link basis states of Eq. (A5), it follows that, within the sector with the link-symmetry constraint $\hat{L}_{x,\mu} = 2$, the electric field operator is uniquely identified by taking only the half-link fermionic configuration, namely:

$$\hat{E}_{x,\mu} = 1 - \hat{n}_{x,\mu}^a - \hat{n}_{x,\mu}^b. \quad\quad \text{(A7)}$$

In this way, the generators of the Gauss's law of Eq. (3) are transformed into completely local operators acting on the site x only:
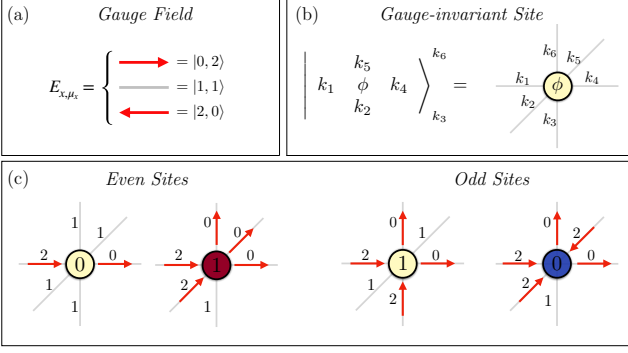
Figure 6: (a) Representation of the gauge field in terms of two species of Dirac modes in the sector with a total number of fermions equal to two. b) Generic state of the local site composed by the the matter degrees of freedom and six half-links along the three spatial directions. On each half-link the coefficients $k_j \in \{0,1,2\}$ define the fermionic modes. (c) Examples of gauge-invariant configurations for even and odd sites. Due to the use of staggered-fermions, the presence/absence of a fermion in an even/odd site represents the presence of a charge/anti-charge.

$$\hat{G}_x = \hat{\psi}_x^\dagger \hat{\psi}_x - \frac{1-(-1)^x}{2} - \sum_\mu \left(1 - \hat{n}_{x,\mu}^a - \hat{n}_{x,\mu}^b\right) \quad \text{(A8)}$$

Taking into account this property, it is possible to construct the gauge-invariant basis for the local site x, that is composed by the lattice site and the six half-links along the directions $\pm\mu_x, \pm\mu_y, \pm\mu_z$ (see Fig. 6):

$$\left| \begin{matrix} & k_5 & & k_6 \\ k_1 & \phi & k_4 \\ & k_2 & \end{matrix} \right\rangle_{k_3} = (-1)^{\delta_{k_1,2}+\delta_{k_2,2}+\delta_{k_3,2}} |\phi\rangle_x \quad \text{(A9)}$$
$$\times |k_1\rangle_{x,-\mu_x} |k_2\rangle_{x,-\mu_y} |k_3\rangle_{x,-\mu_z}$$
$$\times |k_4\rangle_{x,\mu_x} |k_5\rangle_{x,\mu_y} |k_6\rangle_{x,\mu_z}$$

where $|\phi\rangle_x = (\hat{\psi}_x^\dagger)^\phi |0\rangle$ with $\phi = 0,1$ describes the presence or the absence of the matter/antimatter particles. The indices k_j run over $\{0,1,2\}$ selecting a configuration of the 3-hardcore modes for each respective half-link. The presence of the factor $(-1)^{\delta_{k_1,2}+\delta_{k_2,2}+\delta_{k_3,2}}$ allows us to satisfy the anticommutation relations of the fermionic representation recovering the correct signs of Eq. (A5). The occupation numbers ϕ and k_j are not independent due to the constraint imposed by the Gauss's law

$$\hat{G}_x \left| \begin{matrix} & k_5 & & k_6 \\ k_1 & \phi & k_4 \\ & k_2 & \end{matrix} \right\rangle_{k_3} = 0. \quad \text{(A10)}$$

This equation, in the new language of matter fermions and rishons, reads

$$\phi + \sum_{j=1}^6 k_j = 6 + \frac{1-(-1)^x}{2}. \quad \text{(A11)}$$

where the factor 6 is indeed the coordination number of the cubic lattice. Thus, the gauge invariant configurations of the local basis are obtained by applying this constraint, effectively reducing the 'dressed-site' (matter and 6 rishon modes) dimension from $2 \cdot 3^6 = 1458$ to merely 267. We encode these states as building blocks of our computational representation for the TN algorithms. In Fig. 6 we show some examples of gauge-invariant configurations for even and odd sites.

The construction of the gauge-invariant local sites is particularly advantageous for our numerical purposes: in fact, it is now possible to express all the terms in the Hamiltonian of Eq. (1) of the main text as product of completely local operators that commute on different sites. Let us consider the kinetic term of the Hamiltonian and apply the representation of the gauge field in terms of the 3-hardcore fermionic modes:

$$\hat{\psi}_x^\dagger \hat{U}_{x,\mu} \hat{\psi}_{x+\mu} = \hat{\psi}_x^\dagger \hat{\eta}_{x,\mu} \hat{\eta}_{x+\mu,-\mu}^\dagger \hat{\psi}_{x+\mu}$$
$$= \left(\hat{\eta}_{x,\mu}^\dagger \hat{\psi}_x\right)^\dagger \left(\hat{\eta}_{x+\mu,-\mu}^\dagger \hat{\psi}_{x+\mu}\right)$$
$$= M_x^{(\alpha)\dagger} M_{x+\mu}^{\alpha'} \quad \text{(A12)}$$

where the indices α and α' select the right operators depending on the different directions in which the hopping process takes place. The operators $M_{x,\mu}^\alpha$ are genuinely local (i.e. they commute with operators acting elsewhere) as they are always quadratic in the fermionic operators (ψ and/or η). The same argument applies to the magnetic (plaquette) terms in the Hamiltonian

$$\Box_{\mu_x,\mu_y} = U_{x,x+\mu_x} U_{x+\mu_x,\mu_y} U_{x+\mu_y,\mu_x}^\dagger U_{x,\mu_y}^\dagger =$$
$$= \eta_{x,\mu_x} \eta_{x+\mu_x,-\mu_x}^\dagger \eta_{x+\mu_x,\mu_y} \eta_{x+\mu_x+\mu_y,-\mu_y}^\dagger$$
$$\times \left(\eta_{x+\mu_y,\mu_x} \eta_{x+\mu_x+\mu_y,-\mu_x}^\dagger\right)^\dagger \left(\eta_{x,\mu_y} \eta_{x+\mu_y,-\mu_y}^\dagger\right)^\dagger$$
$$= -\left(\eta_{x,\mu_y}^\dagger \eta_{x,\mu_x}\right)\left(\eta_{x+\mu_x,-\mu_x}^\dagger \eta_{x+\mu_x,\mu_y}\right)$$
$$\times \left(\eta_{x+\mu_x+\mu_y,-\mu_y}^\dagger \eta_{x+\mu_x+\mu_y,-\mu_x}\right)\left(\eta_{x+\mu_y,\mu_x}^\dagger \eta_{x+\mu_y,-\mu_y}\right)$$
$$\equiv -C_x^{(\alpha)} C_{x+\mu_x}^{(\alpha')} C_{x+\mu_x+\mu_y}^{(\alpha'')} C_{x+\mu_y}^{(\alpha''')}, \quad \text{(A13)}$$

where the indices $\alpha, \alpha', \alpha'', \alpha'''$ depend on the plane of the plaquette (in this case $x-y$) and the links involved into the loop. The operators C_x^α are genuinely local and act on the four sites at the corners of the plaquette. The decomposition is the same for the other plaquettes in the

planes $x - z$ and $y - z$. The present construction ensures that they can be treated as spin (or bosonic) operators [71, 72], so we can exploit standard TN algorithms, without the need of explicitly implementing the fermionic parity at each site [78–80].

The mass term and the electric field energy in the Hamiltonian of Eq. (1) of the main text are diagonal in the gauge-invariant basis with the rishon representation and so it is trivial to express them as local operators. These operators include the local chemical potential terms, which we use to pin charges in order to study confinement properties [81, 82]. In conclusion, *all* the operators we employ in the TTN algorithms (see Appendix B) are genuinely local. In order to get an idea on the numerical complexity, we emphasize that the dimension of these matrices acting on the local gauge-invariant basis is 267×267.

Appendix B: Tensor Networks

In this section we present the main concepts of Tensor Networks (TNs) with a particular focus on the Tree Tensor Network (TTN) ansatz that we exploit in this work [83]. For a detailed and exhaustive description of the subject, please see the technical reviews and textbooks [19, 84, 85].

Let us consider a generic quantum system composed by N lattice sites, each of which described by a local Hilbert space H_k of finite dimension d and equipped with a local basis $\{|i\rangle_k\}_{1 \leq i \leq d}$. The whole Hilbert space of the system will be generated by the tensor product of the local Hilbert spaces, that is, $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \mathcal{H}_N$, with a resulting dimension equal to d^N. Thus, a generic pure quantum state of the system $|\psi\rangle$ can be expressed as a linear combination of the basis elements of $\mathcal{H}$, i.e.,

$$|\psi\rangle = \sum_{i_1,...,i_N=1}^{d} c_{i_1,...,i_N} |i_1\rangle_1 \otimes |i_2\rangle_2 \otimes ... \otimes |i_N\rangle_N. \text{(B1)}$$

In principle, the coefficients $c_{i_1,...,i_N}$ are d^N complex numbers. As a consequence, this exact representation of the quantum state is completely inefficient from a computational point of view, since it scales exponentially with the system size N. In other words, the amount of information that we would need to store in memory for a computational representation of the generic quantum state of the system is exponentially large in the number of degrees of freedom.

However, if we are concerned with *local* Hamiltonians, which means that a lattice site interacts only with a finite set of neighboring sites and not with all sites of the lattice, it is possibile to exploit rigorous results on the scaling of entanglement under a bipartition (*area law*) [86, 87] in order to obtain an efficient representation of the states in the low-energy sectors of such Hamiltonians, e.g. ground-states and first excited states. Tensor

Networks provide a natural language for this representation [88, 89] by decomposing the complete rank-N tensor $c_{i_1,...,i_N}$ in Eq. (B1) into a network of smaller-rank local tensors interconnected with auxiliary indices (bond indices). If we control the dimension of the bond indices with a parameter χ, called the bond dimension, the number of coefficients in the TN is of the order $O(poly(N)poly(\chi))$, allowing an efficient representation of the information encoded in the quantum state. Furthermore, the bond dimension χ is a quantitative estimate of the amount of quantum correlations and entanglement present in the TN. In fact, by varying χ, TNs interpolate between a product state ($\chi = 1$) and the exact, inefficient, representation of the considered quantum state ($\chi \approx d^N$).

Matrix product states (MPS) for 1D systems [90–92], Projected Entangled Pair State (PEPS) for 2D and 3D systems [89, 93, 94], Multiscale Entanglement Renormalization Ansatz (MERA) [95, 96] and Tree Tensor Networks (TTN), that can be defined in any dimension, [83, 97, 98] are all important examples of efficient representations based on TNs.

MPS algorithms, such as the Density Matrix Renormalization Group (DMRG) [99], represent the state-of-the-art technique for the numerical simulation of manybody systems in 1D. MPS satisfy area-law and are extremely powerful since they allow to compute scalar products between two wave functions and local observables in an exact and efficient way. This property does not hold true for higher-dimensional generalizations, such as PEPS, and the development of TN algorithms, for accurate and efficiently scalable computations, is at the center of current research efforts.

In particular, one of the main problems is related to the choice of the TN geometry for simulating higher-dimensional systems. PEPS intuitively reproduce the structure of the lattice with one tensor for each physical site and the bond indices directly follow the lattice grid. The resulting TN follows the area-law of entanglement but it contains "loops", making the contractions for computing expectation values exponentially hard [100]. Furthermore, the computational cost for performing the variational optimization of PEPS, as for instance in the ground state searching, scales as $O(\chi^{10})$ as a function of the bond dimension. This severely limits the possibility of reaching high values of χ, especially for large system sizes (typical values are $\chi \approx 10$ for spin systems). For our purpose of simulating LGT in three-spatial dimensions this represents a crucial problem since the local dimension of our model is extremely high, i.e., $d = 267$, and so, it becomes necessary to be able to handle high values of χ in order to reach the numerical convergence.

Alternative ansatzs for simulating quantum manybody systems are the TTNs, that decompose the wave function into a network of tensors without loops, allowing efficient contraction algorithms with a polynomial scaling as a function of the system size. In Fig. 7, we show the typical TTN ansazts for 1D and 2D systems and our
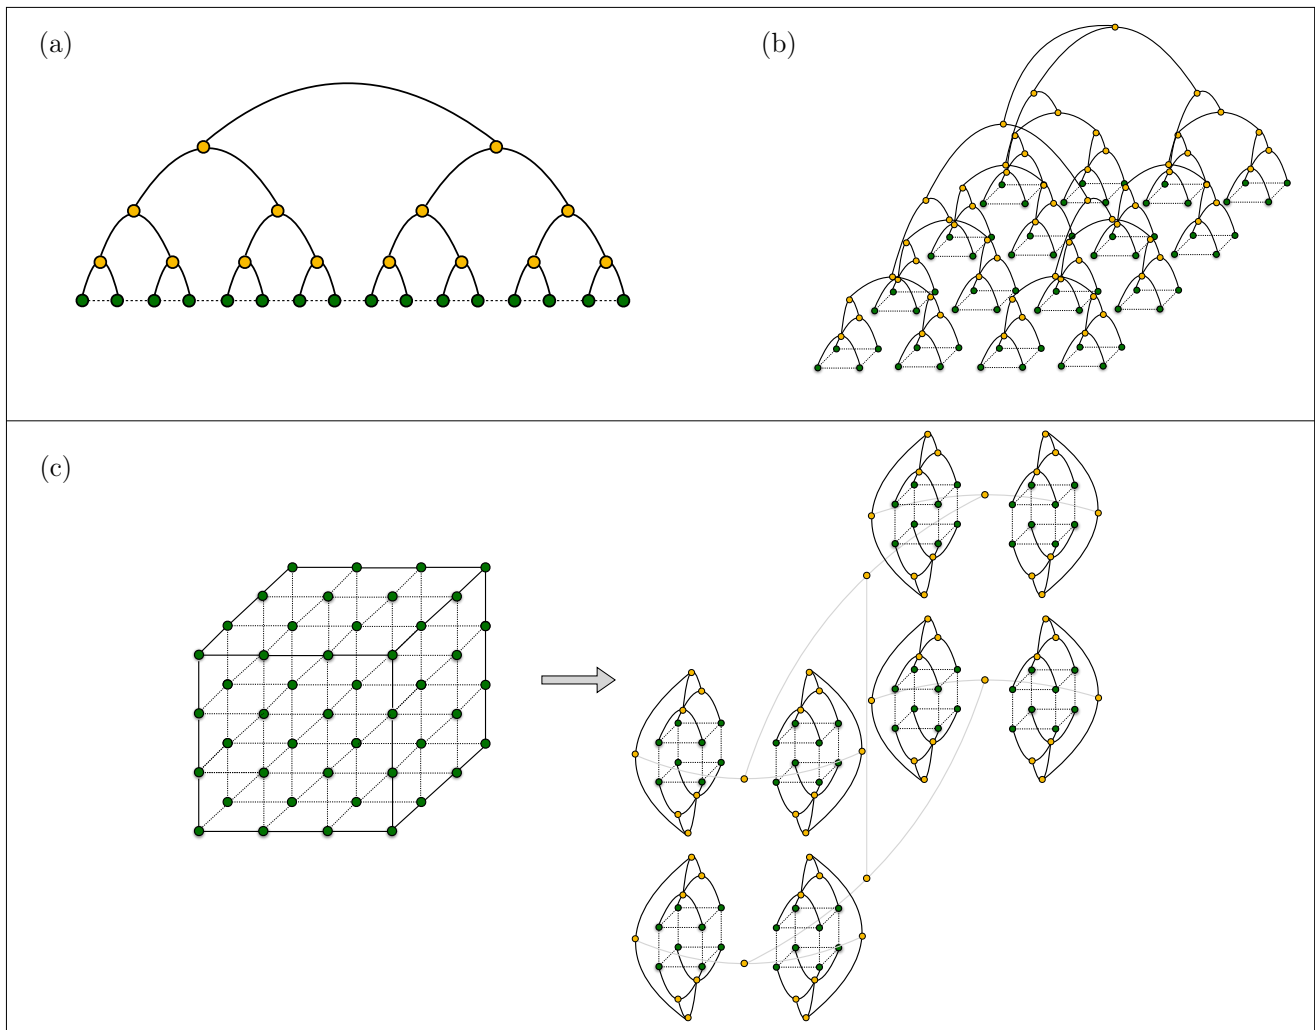
Figure 7: TTN representations for (a) 1D lattice and (b) 2D square lattice. Green circles indicate the sites of the lattice connected to the physical indices of the tree, whereas the yellow circles are the tensors making up the TTN. In (c) we shown our generalization to the 3D cubic lattice that we use for the numerical simulations of the LGT. The different colours of the bond indices are just for a better visualization of the tree structure.

generalization to 3D lattice. TTNs offer more tractable computational costs since the complete contraction and the variational optimization algorithms scale as $O(\chi^4)$, making it easier to reach high values of the bond dimension (up to $\chi \approx 1000$). The price to pay for using the loopless structure is related to the area law that TTNs may not explicitly reproduce in dimensions higher than one [101]. Nevertheless, we use the TTN ansatz in a variational optimization, so we can improve the precision by using increasing values of χ, providing in this way a careful control over the convergence of our numerical results.

The TTN algorithm for the ground state computation of our LTG model follows the technical implementation described in [85] and it takes into account the conservation of the total charge through the definition of global $U(1)$ symmetry sectors encoded in the TTN. In this way

we can easily access finite charge-density regimes, with any imbalance between charges and anticharges.

Our TTN for the 3D lattice is composed entirely of tensors with three links (this structure is usually called *binary tree*). The construction of the TTN starts from merging the physical indices at the bottom, that represent two neighboring lattice sites along the x-direction, into one tensor. Then, these tensors are connected along the y-direction through new tensors in an upper layer. The tensors in this layer are then connected along the z-direction through a new layer of tensors. Thus, this procedure is iteratively repeated by properly setting the connections along the three spatial directions in the upper layers of the tree. At the beginning of the simulation, we randomly initialize all the tensors in the network and the distribution of the global symmetry sec-
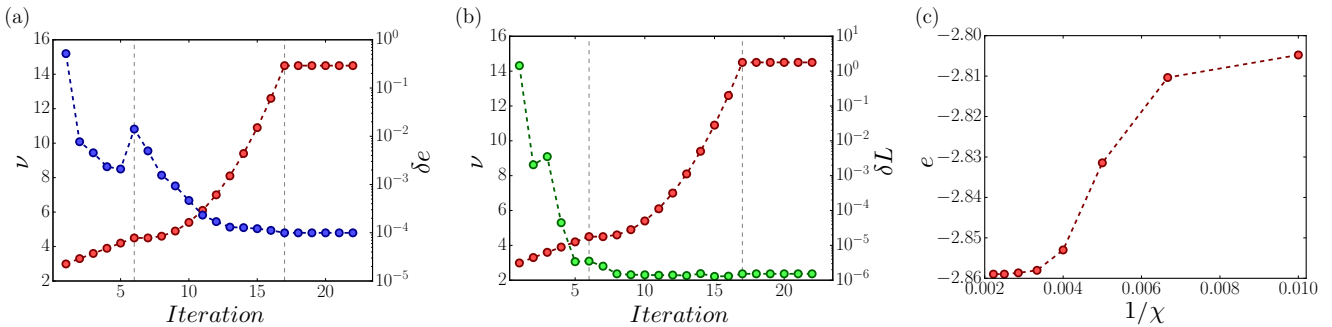
Figure 8: (a) Driven optimization (in three steps: linear, quadratic, constant) of the penalty coefficient ν (red) and behavior of the energy (blue) as a function of the iterations for an exemplifying simulation. The energy is reported as the difference with the lowest final energy that we reach. (b) Driven optimization of the penalty coefficient ν (red) and global error δL (green) with respect to the link symmetry during the optimization steps. (c) Scaling of the energy density as a function of the inverse of the bond dimension $1/\chi$. The bond dimension χ is in the range $[100, 450]$.

tors. During the variational optimization stage, in order to improve the convergence, we perform the single-tensor optimization with subspace-expansion technique, i.e., allowing a dynamical increase of the local bond dimension and adapting the symmetry sectors [85]. This scheme has a global computational cost of the order $O(\chi^4)$. The single tensor optimization is implemented in three steps: (i) the effective Hamiltonian H_{eff} for the tensor is obtained by contracting the complete Hamiltonian of the system with all the remaining tensors of the tree; (ii) the local eigenvalue problem for H_{eff} is solved by using the Arnoldi method of the ARPACK library; (iii) the tensor is updated by the eigenvector of H_{eff} corresponding to the lowest eigenvalue. This procedure is iterated by sweeping through the TTN from the lowest to the highest layers, gradually reducing the energy expectation value. After completing the whole sweep, the procedure is iterated again and again, until the desidered convergence in the energy is reached. The precision of the Arnoldi algorithm is increased in each sweep, for gaining more accuracy in solving the local eigenvalue problems as we approach the final convergence.

TTN computations presented in this work are extremely challenging due to the complexity of LGTs in the three-dimensional scenario. They were performed on different HPC-clusters (CloudVeneto, CINECA, BwUni-Cluster and ATOS Bull): a single simulation for the maximum size that we reached, a $8 \times 8 \times 8$ lattice, can last up to five weeks until final convergence, depending on the different regimes of the model and the control parameters of the algorithms.

Appendix C: Numerical Convergence

With our numerical simulations we characterize the properties of the ground state of the system as a function of the parameters in the Hamiltonian of Eq. (1) of the

main text. We fix the energy scale by setting the hopping coefficient $t = 1$ and we access several regimes of the mass m, the electric g_e and the magnetic coupling g_m. We consider simple cubic lattices $L \times L \times L$ with the linear size L being a binary power; in particular, we simulate the case with $L = 2, 4, 8$, that is, up to 512 lattice sites.

As explained in Appendix A, in order to obtain the right representation of the electric field operators, we have to enforce the extra link symmetry constraint $\hat{L}_{x,\mu} = 2$ at every pair of neighboring sites. For this reason, we include in the Hamiltonian additional terms that energetically penalise all the states with a number of hardcore fermions per link different from two, namely:

$$H_{pen} = \nu \sum_{x,\mu} \left(1 - \delta_{2,\hat{L}_{x,\mu}} \right) \tag{C1}$$

where $\nu > 0$ is the penalty coefficient and $\delta_{2,\hat{L}_{x,\mu}}$ are the projectors on the states that satisfy the extra link constraint. In this way, the penalty terms vanish when the link symmetry is satisfied and raise the energy of the states violating the constraint. In principle, the link symmetry is rigorously satisfied for $\nu \to \infty$. At numerical level, this limit translates into choosing ν much larger than the other simulation parameters of the Hamiltonian, i.e., $\nu \gg max\{|t|, |m|, |g_{el}|, |g_m|\}$. However, setting ν too large in the first optimisation steps could lead to local minima or non-physical states, since the variational algorithm would focus only on the penalty terms more than the physical ones. In order to avoid this problem and reach the convergence, we adopt a *driven optimization*, by varying the penalty coefficient ν in three steps: (i) starting from a very small value of ν and from a random state of the TTN, that in general does not respect the extra link symmetry, we drive the penalty term with a linear growth of ν during the first optimization sweeps. In this stage, the optimization will focus mainly on the physical quantities, until we notice a slight rise of the energy: this effect signals that the global optmiza-
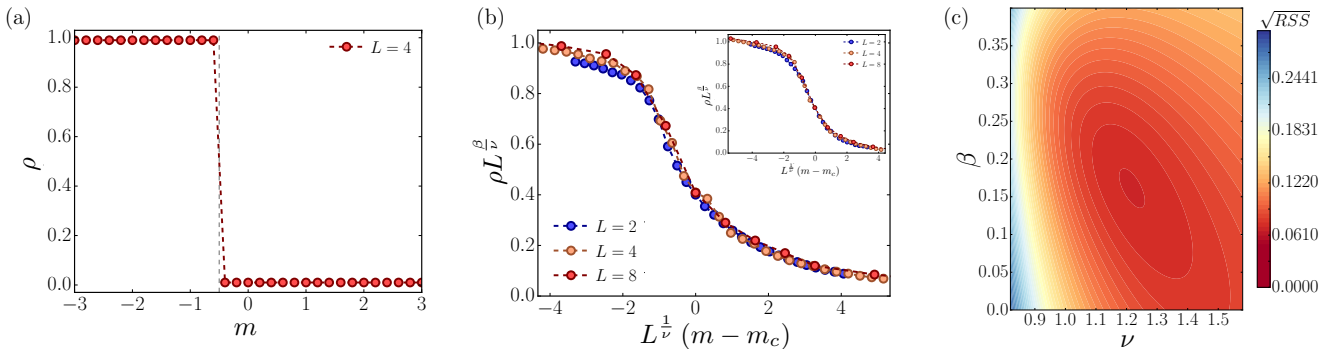
Figure 9: (a) Particle density as a function of m, for $t = 0$, $g_m^2 = 0$ and $L = 4$. (b) Universal scaling function $\lambda(x)$ close to the transition point $m_c \approx -0.39$ for $g_m^2 = 0$ with critical exponents $\beta \approx 0.16$ and $\nu \approx 1.22$. The inset shows the same universal behavior close to the transition point $m_c \approx 0.22$ in the presence of magnetic interactions with $g_m^2 = 8/g_e^2 = 4$ and the same critical exponents $\beta \approx 0.16$ and $\nu \approx 1.22$. (c) Contour plot of the square root of the residual sum of squares in the (ν, β) plane for the best-fitting values of the critical exponents.

tion procedure of the TTN become significantly sensitive to the penalty terms. (ii) Thus, we impose a quadratic growth of ν so that, in the immediately following sweeps, the penalty is increased at a slower rate with respect to the linear regime. (iii) After reaching the maximum desidered value of ν, that is an input parameter of the simulation, we keep it fixed, performing the last sweeps in order to ensure the convergence of the energy. This driven optimization strategy is summarized in Fig. 8(a) where we show the three different stages of the penalty coefficient ν and typical behavior of the energy difference δe, computed with respect to the lowest final energy that we reach, as a function of the iterations.

We can also quantify the global error with respect to the link symmetry during the driven optimization sweeps, by defining:

$$\delta L = \sum_{x,\mu} |\langle GS| (L_{x,\mu} - 2) |GS\rangle| \qquad (C2)$$

i.e., the sum of the deviations from the exact link constraint $L_{x,\mu} = 2$, computed over all the links of the lattice on the ground state. The typical behavior of this quantity is shown in Fig. 8(b): at the end of the optimization procedure, the global error results of the order of 10^{-6}. We also check the convergence of our TTN algorithms as a function of the bond dimension χ, by using $\chi = 450$ at most to ensure stability of our findings. Depending from the different system sizes and regimes of physical parameters, we estimate the relative error of the energy in the range $[10^{-2}, 10^{-4}]$. A typical scaling of the energy density as a function of the inverse of the bond dimension $1/\chi$ is shown in Fig. 8(c).

Appendix D: Critical points: scaling analysis

In this section we show the finite-size scaling analysis for detecting the phase transition separating the *charge-crystal phase* and the *vacuum phase* and the related location of the critical points.

At $t = 0$ and neglecting the magnetic interactions, i.e., for $g_m^2 = 0$, the Hamiltonian of Eq. (1) results diagonal in the local basis described in Appendix A and it is trivial to prove that the system undergoes a first-order phase transition between the *bare vacuum*, with energy $E_v = -m\frac{L^3}{2}$ and the *charge-crystal phase*, with energy $E_{ch} = (m + \frac{g_e^2}{2})\frac{L^3}{2}$. The ground-state exhibits a level-crossing at the critical value $m_c^{(0)} = -\frac{g_e^2}{4} = -\frac{1}{2}$ that is obtained at $E_v = E_{ch}$. This behaviour is clearly seen in Fig. 9(a), showing a discontinuous transition between the two configurations.

In order to understand the behavior of the system for finite $t = 1$ and $g_m^2 = 0$, we observe that the density, plotted in Fig. 2(a) of the main text, changes continuously as a function of the mass parameter and we might have a second-order phase transition. Finite-size scaling theory [102] implies that the behavior of the system close to a critical point, i.e. for $m \approx m_c$, can be described in terms of a universal function $\lambda(x)$ such that for our observable:

$$\rho L^{\frac{\beta}{\nu}} = \lambda \left(L^{\frac{1}{\nu}}(m - m_c) \right) \qquad (D1)$$

where β and ν are critical exponents. In particular this relation implies that for $m \approx m_c$, the value of $\rho L^{\frac{\beta}{\nu}}$ is independent of the size of the system. We use this property to get an estimate of the values of m_c, β, ν. In particular, we consider a grid of values for these parameter and for each set of values we fit our points $\rho L^{\frac{\beta}{\nu}}$ with an high-degree polynomial $f\left(L^{\frac{1}{\nu}}(m - m_c)\right)$. We compute the residual sum of squares (RSS) and we select the set of values which minimize this quantity, producing the best data collapse. We get for the critical point $m_c \approx -0.39$ and for the critical exponents $\beta \approx 0.16$ and $\nu \approx 1.22$. In

Fig. 9(b) we show the collapse of our numerical results onto the same universal function $\lambda(x)$ and in Fig. 9(c) a contour plot of the square root of the residual sum of squares in the (ν, β) plane for the best-fitting values.

By extending the previous considerations and the finite-size scaling analysis to the case with magnetic interactions with $g_m^2 = 8/g_e^2 = 4$, we check again the presence of a critical point and the values of critical exponents through the formula of Eq. (D1). We obtain a universal scaling function for $m_c \approx 0.22$ and the same critical exponents $\beta \approx 0.16$, $\nu \approx 1.22$, as reported in the inset of Fig. 9(b). Thus, while the transition and its univer-

sality remain unchanged in the presence of the magnetic coupling, the critical point is shifted toward positive values of the mass parameter, signaling that the magnetic interactions determine a visible enhancement of the production of charges and anticharges out of the vacuum.

Although a more precise determination of the numerical values of the critical exponents would require additional extensive analysis that results beyond the scope of this paper, our findings strongly indicate the presence of a phase transition at finite m for the three-dimensional lattice model of QED (compare with other previously investigated transitions in lattice QED, e.g. Ref. [103]).

[1] M. E. Peskin and D. V. Schroeder, *An introduction to quantum field theory* (Westview, Boulder, CO, 1995).
[2] T. Cheng and L. Li, *Gauge theory of elementary particle physics* (Oxford University Press, 2006).
[3] M. D. Schwartz, *Quantum Field Theory and the Standard Model* (Cambridge University Press, 2014).
[4] E. Fradkin, *Field Theories of Condensed Matter Physics*, 2nd ed. (Cambridge University Press, 2013).
[5] Y. Zhou, K. Kanoda, and T.-K. Ng, Rev. Mod. Phys. **89**, 025003 (2017).
[6] N. Brambilla *et al.*, The European Physical Journal C **74**, 2981 (2014).
[7] H. R. Fiebig and R. M. Woloshyn, Phys. Rev. D **42**, 3520 (1990).
[8] I. F. Herbut and B. H. Seradjeh, Phys. Rev. Lett. **91**, 171601 (2003).
[9] F. S. Nogueira and H. Kleinert, Phys. Rev. B **77**, 045107 (2008).
[10] X. Y. Xu, Y. Qi, L. Zhang, F. F. Assaad, C. Xu, and Z. Y. Meng, Phys. Rev. X **9**, 021022 (2019).
[11] K. G. Wilson, Phys. Rev. D **10**, 2445 (1974).
[12] R. Alkofer and J. Greensite, Journal of Physics G: Nuclear and Particle Physics **34**, S3 (2007).
[13] M. Creutz, Nuclear Physics B - Proceedings Supplements **10**, 1 (1989).
[14] A. Bazavov, D. Toussaint, C. Bernard, J. Laiho, C. DeTar, L. Levkova, M. B. Oktay, S. Gottlieb, U. M. Heller, J. E. Hetrick, P. B. Mackenzie, R. Sugar, and R. S. Van de Water, Rev. Mod. Phys. **82**, 1349 (2010).
[15] M. Tanabashi *et al.* (Particle Data Group), Phys. Rev. D **98**, 030001 (2018).
[16] J. Kogut and L. Susskind, Phys. Rev. D **11**, 395 (1975).
[17] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).
[18] F. Verstraete, V. Murg, and J. I. Cirac, Advances in Physics **57**, 143 (2008).
[19] R. Orús, Annals of Physics **349**, 117 (2014).
[20] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Phys. Rev. B **94**, 165116 (2016).
[21] T. M. R. Byrnes, P. Sriganesh, R. J. Bursill, and C. J. Hamer, Phys. Rev. D **66**, 013002 (2002).
[22] E. Rico, T. Pichler, M. Dalmonte, P. Zoller, and S. Montangero, Phys. Rev. Lett. **112**, 201601 (2014).
[23] J. Haegeman, K. Van Acoleyen, N. Schuch, J. I. Cirac, and F. Verstraete, Phys. Rev. X **5**, 011024 (2015).
[24] B. Buyens, S. Montangero, J. Haegeman, F. Verstraete,

[25] and K. Van Acoleyen, Phys. Rev. D **95**, 094509 (2017).
[25] E. Ercolessi, P. Facchi, G. Magnifico, S. Pascazio, and F. V. Pepe, Phys. Rev. D **98**, 074503 (2018).
[26] G. Magnifico, D. Vodola, E. Ercolessi, S. P. Kumar, M. Müller, and A. Bermudez, Phys. Rev. D **99**, 014503 (2019).
[27] G. Magnifico, M. Dalmonte, P. Facchi, S. Pascazio, F. V. Pepe, and E. Ercolessi, Quantum **4**, 281 (2020).
[28] M. C. Bañuls, K. Cichy, J. I. Cirac, K. Jansen, and S. Kühn, Phys. Rev. X **7**, 041046 (2017).
[29] L. Tagliacozzo, A. Celi, A. Zamora, and M. Lewenstein, Annals of Physics **330**, 160 (2013).
[30] L. Tagliacozzo, A. Celi, and M. Lewenstein, Phys. Rev. X **4**, 041024 (2014).
[31] E. Zohar, M. Burrello, T. B. Wahl, and J. I. Cirac, Annals of Physics **363**, 385 (2015).
[32] T. Felser, P. Silvi, M. Collura, and S. Montangero, arXiv e-prints , arXiv:1911.09693 (2019), arXiv:1911.09693 [quant-ph].
[33] P. Emonts, M. C. Bañuls, I. Cirac, and E. Zohar, Phys. Rev. D **102**, 074501 (2020).
[34] D. Horn, Physics Letters B **100**, 149 (1981).
[35] S. Chandrasekharan and U.-J. Wiese, Nuclear Physics B **492**, 455 (1997).
[36] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller, and R. Blatt, Nature **534**, 516 (2016).
[37] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. J. Savage, Phys. Rev. A **98**, 032331 (2018).
[38] A. Mil, T. V. Zache, A. Hegde, A. Xia, R. P. Bhatt, M. K. Oberthaler, P. Hauke, J. Berges, and F. Jendrzejewski, Science **367**, 1128 (2020).
[39] B. Yang, H. Sun, R. Ott, H.-Y. Wang, T. V. Zache, J. C. Halimeh, Z.-S. Yuan, P. Hauke, and J.-W. Pan, (2020), arXiv:2003.08945 [cond-mat.quant-gas].
[40] C. Schweizer, F. Grusdt, M. Berngruber, L. Barbiero, E. Demler, N. Goldman, I. Bloch, and M. Aidelsburger, Science **367**, 1128 (2020).
[41] S. V. Mathis, G. Mazzola, and I. Tavernelli, Phys. Rev. D **102**, 094501 (2020).
[42] S. P. Jordan, K. S. M. Lee, and J. Preskill, Science **336**, 1130 (2012).
[43] J. Schwinger, Phys. Rev. **82**, 664 (1951).
[44] A. Polyakov, Physics Letters B **59**, 82 (1975).
[45] T. Banks, R. Myerson, and J. Kogut, Nuclear Physics

B **129**, 493 (1977).

[46] S. D. Drell, H. R. Quinn, B. Svetitsky, and M. Weinstein, Phys. Rev. D **19**, 619 (1979).

[47] J. Jersák, T. Neuhaus, and P. Zerwas, Physics Letters B **133**, 103 (1983).

[48] A. H. Guth, Phys. Rev. D **21**, 2291 (1980).

[49] K.-I. Kondo, Phys. Rev. D **58**, 085013 (1998).

[50] G. Hooft, Nuclear Physics: B **138**, 1 (1978).

[51] P. Silvi, E. Rico, T. Calarco, and S. Montangero, New Journal of Physics **16**, 103015 (2014).

[52] P. Silvi, Y. Sauer, F. Tschirsich, and S. Montangero, Phys. Rev. D **100**, 074512 (2019).

[53] Y.-P. Huang, D. Banerjee, and M. Heyl, Phys. Rev. Lett. **122**, 250401 (2019).

[54] T. Pichler, M. Dalmonte, E. Rico, P. Zoller, and S. Montangero, Phys. Rev. X **6**, 011023 (2016).

[55] P. Silvi, E. Rico, M. Dalmonte, F. Tschirsich, and S. Montangero, Quantum **1**, 9 (2017).

[56] P. Sala, T. Shi, S. Kühn, M. C. Banuls, E. Demler, and J. I. Cirac, in *The 36th Annual International Symposium on Lattice Field Theory. 22-28 July* (2018) p. 230.

[57] B. Buyens, J. Haegeman, K. Van Acoleyen, H. Verschelde, and F. Verstraete, Phys. Rev. Lett. **113**, 091601 (2014).

[58] L. Funcke, K. Jansen, and S. Kühn, Phys. Rev. D **101**, 054507 (2020).

[59] M. C. Bañuls, K. Cichy, J. I. Cirac, K. Jansen, and S. Kühn, Phys. Rev. Lett. **118**, 071601 (2017).

[60] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).

[61] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, *et al.*, Nature **569**, 355 (2019).

[62] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. Van Acoleyen, F. Verstraete, U.-J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, European Physical Journal D **74**, 165 (2020).

[63] D. Paulson, L. Dellantonio, J. F. Haase, A. Celi, A. Kan, A. Jena, C. Kokail, R. van Bijnen, K. Jansen, P. Zoller, *et al.*, arXiv preprint arXiv:2008.09252 (2020).

[64] J. F. Haase, L. Dellantonio, A. Celi, D. Paulson, A. Kan, K. Jansen, and C. A. Muschik, arXiv preprint arXiv:2006.14160 (2020).

[65] C. N. Yang and R. L. Mills, Phys. Rev. **96**, 191 (1954).

[66] J. Goldstone, A. Salam, and S. Weinberg, Phys. Rev. **127**, 965 (1962).

[67] F. Englert and R. Brout, Phys. Rev. Lett. **13**, 321 (1964).

[68] P. W. Higgs, Phys. Rev. Lett. **13**, 508 (1964).

[69] G. Baskaran and P. W. Anderson, Phys. Rev. B **37**, 580 (1988).

[70] H. Kleinert, *Gauge Fields in Condensed Matter* (World Scientific, 1989).

[71] E. Zohar and J. I. Cirac, Phys. Rev. B **98**, 075119 (2018).

[72] E. Zohar and J. I. Cirac, Phys. Rev. D **99**, 114511 (2019).

[73] R. Brower, S. Chandrasekharan, and U.-J. Wiese, Phys. Rev. D **60**, 094502 (1999).

[74] R. Brower, S. Chandrasekharan, S. Riederer, and U.-J. Wiese, Nuclear Physics B **693**, 149 (2004).

[75] D. Banerjee, M. Bögli, M. Dalmonte, E. Rico, P. Stebler, U.-J. Wiese, and P. Zoller, Phys. Rev. Lett. **110**, 125303 (2013).

[76] J. B. Kogut, Rev. Mod. Phys. **51**, 659 (1979).

[77] L. Susskind, Phys. Rev. D **16**, 3031 (1977).

[78] C. V. Kraus, N. Schuch, F. Verstraete, and J. I. Cirac, Phys. Rev. A **81**, 052338 (2010).

[79] P. Corboz and G. Vidal, Phys. Rev. B **80**, 165129 (2009).

[80] T. Barthel, C. Pineda, and J. Eisert, Phys. Rev. A **80**, 042333 (2009).

[81] A. Polyakov, Nuclear Physics B **120**, 429 (1977).

[82] J. Greensite, *An introduction to the confinement problem*, Vol. 821 (2011).

[83] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero, Phys. Rev. B **90**, 125154 (2014).

[84] S. Montangero, *Introduction to Tensor Network Methods* (Springer International Publishing, 2018).

[85] P. Silvi, F. Tschirsich, M. Gerster, J. Junemann, D. Jaschke, M. Rizzi, and S. Montangero, SciPost Phys. Lect. Notes , 8 (2019).

[86] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, Rev. Mod. Phys. **80**, 517 (2008).

[87] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).

[88] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).

[89] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Phys. Rev. Lett. **96**, 220601 (2006).

[90] M. Fannes, B. Nachtergaele, and R. F. Werner, Communications in Mathematical Physics **144**, 443 (1992).

[91] A. Klumper, A. Schadschneider, and J. Zittartz, Journal of Physics A Mathematical General **24**, L955 (1991).

[92] A. Klümper, A. Schadschneider, and J. Zittartz, EPL (Europhysics Letters) **24**, 293 (1993).

[93] F. Verstraete and J. I. Cirac, (2004), arXiv:cond-mat/0407066 [cond-mat.str-el].

[94] M. Tepaske and D. J. Luitz, (2020), arXiv:2005.13592 [cond-mat.str-el].

[95] G. Vidal, Phys. Rev. Lett. **99**, 220405 (2007).

[96] G. Evenbly and G. Vidal, Phys. Rev. Lett. **102**, 180406 (2009).

[97] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Phys. Rev. A **74**, 022320 (2006).

[98] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio, Phys. Rev. A **81**, 062335 (2010).

[99] U. Schollwöck, Annals of Physics **326**, 96 (2011), january 2011 Special Issue.

[100] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, Phys. Rev. Research **2**, 013010 (2020).

[101] A. J. Ferris, Phys. Rev. B **87**, 125139 (2013).

[102] J. Cardy, *Scaling and Renormalization in Statistical Physics*, Cambridge Lecture Notes in Physics (Cambridge University Press, 1996).

[103] M. Göckeler, R. Horsley, E. Laermann, P. Rakow, G. Schierholz, R. Sommer, and U.-J. Wiese, Nuclear Physics B **334**, 527 (1990).

B

Solving classical optimisation problems with Tensor Networks

After the study of Quantum many-body systems via Tensor Networks, this appendix illustrates the analysis of classical systems in form of a manuscript which at the stage of submission of this thesis is peer-reviewed and available on arXiv [99].

In particular, therein we show how to solve classical optimisation problems with Tensor Networks with the following example of optimising therapy plans in cancer treatment. In this problem, we aim to optimise the radioactive dose distribution inserted in a patient in order to kill the cancer tissue to be treated while sparing surrounding healthy organs. With the aid of this exemplifying problem, we show how to solve classical optimisation problems with Tensor Networks in three steps: *(i)* Mapping the classical problem to a ground-state search of a many-body Ising-like spin-glass Hamiltonian, *(ii)* finding the ground-state via a quantum Tensor Network optimisation, and *(iii)* mapping the obtained ground-state wavefunction back giving us the solution of the classical problem.

With this particular Tensor Networks application of optimizing radiotherapy plans for cancer treatment, we further aim to build solid foundations for future applications of quantum science and technologies to solve life-changing problems in medicine. Optimizing the precise radiotherapy plans is one of the most complex computational tasks in cancer treatment: A radiation beam with hundreds of optimization variables is targeting a tumour inside a human being while healthy organs have to be avoided. We believe that the availability of large-scale quantum technology will strongly impact this area in the future, enhancing our resources in the fight against cancer, both in terms of sparing crucial time and increasing treatment effectiveness. Consequently, more people can be treated in less time and on the other hand, this will allow developing highly personalized treatments, increasing the probability to save human lives.

This work may pave the way for quantum technologies in the medical area of cancer treatment. We present and discuss the very first successful application of Tensor Networks to optimise such a cancer treatment plan. In particular, we use a Tree Tensor Network to solve the beam fluence optimization problem in Intensity Modulated Radiation Therapy (IMRT). We define a robust and clear strategy to map the classical IMRT problem to quantum-like

hardware and show some successful results of the application of TTN to scenarios of increasing complexity. Our TN approach can not only be used as a benchmark for the development of quantum algorithms for radiotherapy but further as a tool for partially overcoming the lack of scalable quantum hardware. Along the way, we provide a novel opportunity to develop strategies to map sophisticated classical problems to quantum hardware and thereby sets new standards for the development of hybrid classical-quantum optimization strategies.

Consequently, the following work not only gives an insight about how to solve classical optimisation problems with Tensor Networks but shows three further strong-points to build on in future research: *(i)* It will impact the quantum computing community, as it introduces a novel and innovative method to develop and test strategies for solving complex classical optimization problems on future quantum computers; *(ii)* it creates a direct connection between quantum computing and radiotherapy, laying the foundation for future application of quantum technologies to medicine; *(iii)* it creates a connection between the quantum science and technologies, such as TNs, to relevant problems of high impact beyond the quantum community, such as fighting cancer.

Optimizing Radiotherapy Plans for Cancer Treatment with Tensor Networks

Samuele Cavinato,[1] Timo Felser,[1, 2, 3] Marco Fusella,[4] Marta Paiusco,[4] and Simone Montangero[1, 2]

[1] *Dipartimento di Fisica e Astronomia "G. Galilei", Università di Padova, I-35131 Padova, Italy*
[2] *INFN, Sezione di Padova, I-35131 Padova, Italy.*
[3] *Theoretische Physik, Universität des Saarlandes, D-66123 Saarbrücken, Germany.*
[4] *IOV-IRCCS, I-35128 Padova, Italy*
(Dated: October 19, 2020)

We present a novel application of Tensor Network methods in cancer treatment as a potential tool to solve the dose optimization problem in radiotherapy. In particular, the Intensity-Modulated Radiation Therapy (IMRT) technique – that allows treating irregular and inhomogeneous tumors while reducing the radiation toxicity on healthy organs – is based on the optimization of the radiation beamlets intensities. The optimization aims to maximize the delivery of the therapy dose to cancer while avoiding the organs at risk to prevent their damage by the radiation. Here, we map the dose optimization problem into the search of the ground state of an Ising-like Hamiltonian, describing a system of long-range interacting qubits. Finally, we apply a Tree Tensor Network algorithm to find the ground-state of the Hamiltonian. In particular, we present an anatomical scenario exemplifying a prostate cancer treatment. A similar approach can be applied to future hybrid classical-quantum algorithms, paving the way for the use of quantum technologies in future medical treatments.

I. INTRODUCTION

Radiotherapy is one of the techniques used to treat solid tumors by means of a ionizing radiation. The radiation dose released into the cancer tissue damages the DNA of the tumor cells leading to their death or slowing down the growth of the tumor [1, 2]. Radiotherapy is often used in combination with other therapies like surgery, chemotherapy or immunotherapy to improve their global efficacy [3–6].

One of the hardest challenges encountered while treating patients with ionizing radiations is to deliver an optimal dose to the targeted tumor while keeping the radiation as low as possible in the surrounding healthy tissues. One of the most frequently used techniques nowadays is the *Intensity-Modulated Radiation Therapy* (IMRT) [7–10]. In IMRT, the applied radiation is modulated to reach an optimal dose distribution inside the patient. This optimal modulation is obtained by solving a highly nontrivial numerical optimization problem with a high number of optimization parameters and numerous constraints on the final radiation dose distribution. Over the past years, several numerical techniques have been developed to address this challenge [8, 11–14]. In 2015, Nazareth and Spaans proposed to solve the IMRT beam fluence optimization problem using the D-Wave annealer [15], and more recently El Naqa *et al.* proposed an approach exploiting the simulated quantum tunnelling effect [16].

Despite these remarkable novel approaches paved the way for future applications of quantum computation in medicine, their application is still limited. On the one hand, this is due to the lack of scalable quantum hardware. On the other hand, it is not straightforward to extend what is done on classical computers to quantum ones due to the lack of a robust and clear strategy to map the classical IMRT problem to quantum hardware. To try

overcome these limitations as well as to further investigate the applicability of quantum-inspired techniques to the solution of classical optimization problems and foster future applications of quantum technologies to medicine, here we apply Tensor Networks (TNs) to an IMRT dose optimization problem.

TNs are one of the most successful algorithms for simulating quantum many-body systems on classical computers. Indeed, whenever possible, they efficiently represent quantum many-body wavefunctions in a compact form on classical computers [17–21]. Il the last few decades, TNs have proven their effectiveness in the research and analysis of quantum many-body systems, especially for low-dimensional ground-state [22–29]. In addition to that, thanks to the properties they share with quantum hardware, TNs may play the role of test benches for the development of quantum algorithms [30–32].

Hereafter, we show how to solve an IMRT optimization problem with TNs. We first introduce how the classical cost function is mapped into an Ising-like Hamiltonian, where the optimization variables are represented as a set of long-range interacting spins. Finally, we solve the classical optimization problem by finding the ground-state for this Hamiltonian using TNs. We present the application of TNs to two different toy models and to a more realistic anatomical scenario simulating a prostate cancer treatment. We show that TNs results are compatible with other classical techniques, Quadratic Programming (QP) and Simulated Annealing (SA). Our results pave the way to the application of TNs methods to more complex and realistic clinical scenarios, contributing to building solid foundations for future applications of quantum computation to medicine. In the midterm, we foresee the development and application of classical TN methods to the solution of the IMRT problem.

The manuscript is structured as follows: In Section II we provide a brief introduction to the dose optimization problem in IMRT together with its basic mathematical
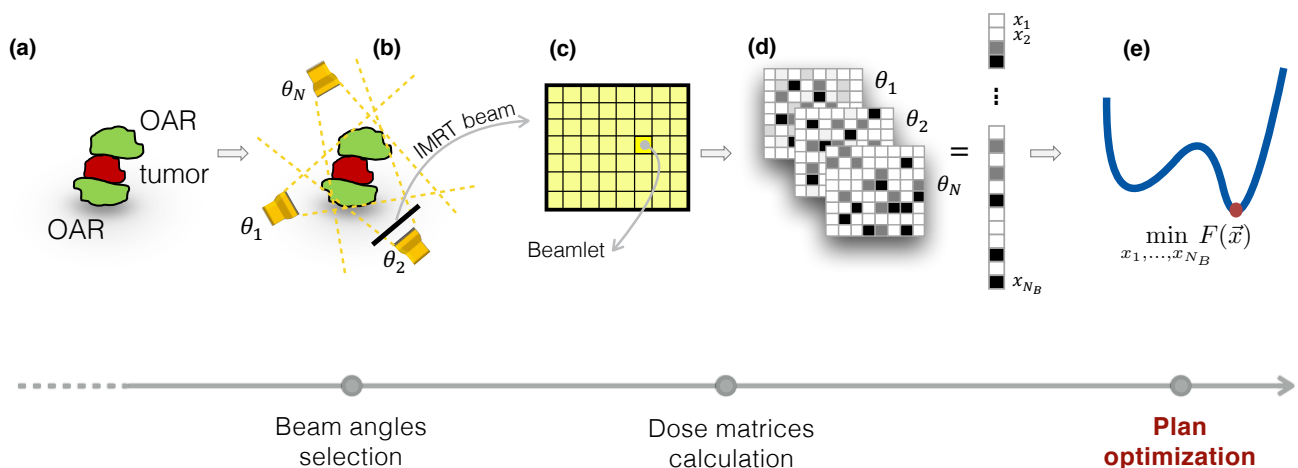
FIG. 1: Typical radiation therapy treatment procedure for the IMRT. The tumor (red) lies close to vital organs at risk (OARs, green) (a). In IMRT treatments the tumor is irradiated from different angles $(\theta_1, \ldots, \theta_N)$ using photon beams (b). Each beam is subdivided into a grid of smaller pencil beams called beamlets in order to modulate the local beam fluence (c). A weight $x_j \geq 0$ is associated to the j-th beamlet in order to quantify its contribution to the whole beam (d). The x_js becomes the optimization variables of a cost function $F(x_1, \ldots, x_{N_B})$ which expresses the distance between the desired dose and the delivered dose and the final aim of the RT planning procedure is to find the configuration $\vec{x}$ of the beamlets with minimizes the cost function (e).

description. In Section III the mapping procedure of the classical cost function to the Ising-type Hamiltonian is described in more detail. Furthermore, this section includes a brief introduction to TNs as well. Finally, we present and discuss the main results obtained from this study in Section IV.

II. RADIOTHERAPY OPTIMIZATION PROBLEM

In radiotherapy, cancer cells are treated by releasing a certain amount of radioactive dose inside the tumor. Modern radiotherapy offers various techniques for treating tumors [33–37] and the choice of one of them depends on factors like the site of the disease, type of cancer and overall patient's conditions [38–40]. One of the most impacting techniques in terms of improvement of treatments quality is the IMRT which became clinical available in its first implementations in the early 2000s [8, 41].

The goal of IMRT is to create a personalized dose distribution for each patient's anatomy that ensures the appropriate dose for the tumor while saving the Organs At Risk (OARs) as much as possible. The choice of IMRT treatment is nowadays mandatory in very challenging cases where a high dose is required for very irregular tumor shapes surrounded by critical OARs. In order to effectively treat this kind of diseases, the IMRT benefits from a non-uniform intensity distribution of the radiation beams. In particular, the radiation beams are modulated by dividing the fluence of each treatment beam into a certain number of smaller pencil-beams called *beamlets* that

can be delivered through the movement of the Multileaf Collimator (MLC).

However, finding the ideal intensity of every single beamlet for a desired treatment which optimises the relation between fields arrangement and dose distribution inside the patient is a highly complex problem. This optimisation problem is typically solved as an *inverse problem* which is encoded into a cost function to be minimized in order to find the optimal beamlet configuration.

In the past years, different approaches have been proposed, among them, the analytic transform method [42], algebraic solutions [13], gradient descent [11], genetic algorithms [14] and simulated annealing [12]. Although these powerful methods have been used in everyday clinical practice with increasing success, finding patient-specific plans is still an open problem since it should account for many factors, increasing the complexity of the problem. Therefore, new mathematical and physical solutions need to be developed in order both to spare precious computational time and to improve the quality of the treatments delivered, with the aim of enhancing our capability to save human lives in the fight againts cancer.

In what follows, we present the main elements of radiotherapy, explaining the general planning procedure for an IMRT treatment, focusing on the underlying numerical optimization problem.

Fig. 1 illustrates a typical treatment procedure. The targeted tumor (red), located between two OARs (green), is irradiated by several photon beams from different angles θ_k (see Fig. 1b). The underlying geometry is defined in the dosimetric plans created by the medical physicist.

The desired dose to the target and the OARs are defined in the treatment plan, following the goal of to spare healthy organs while the proper dose is delivered to the target.

To achieve this treatment goal, the IMRT optimization process begins with the splitting of each beam fluence into a grid of beamlets as shown in Fig. 1c. Only beamlets that traverse the target can be optimized, and the intensity of each j-th beamlet can be modulated independently with a weight x_j. (Fig. 1d). The volume of the patient considered for the optimization can be divided in 3D finite-size elements called *voxels*. The contribution of each beamlet to each voxel depends on their intensity as well as on the geometry and physical properties of each beam and patient's characteristics.

All these information are gathered in the *influence matrix* A which maps every beamlet to each voxel. The influence matrix is usually provided by the Treatment Planning System (TPS) used in the clinical environment, and a complete description of its computation goes far beyond the aim of this work.

The sum of all beamlets produces the total dose within a given voxel i and contributes to the definition of a global dose distribution $D(\mathbf{x})$. Thus, the total dose delivered to the voxel i can be expressed as:

$$D_i(\mathbf{x}) \equiv D_i(x_1, \ldots, x_{N_B}) = \sum_{j=1}^{N_B} a_{ij} x_j \, , \qquad (1)$$

where N_B is the total number of beamlets and a_{ij} describes the so-called *influence matrix*, A, giving the unmodulated contribution of the j-th beamlet to the i-th voxel.

The IMRT planning procedure is then solved through an iterative inverse planning process: the intensity of the applied photon beams are optimized towards the prescribed dose distribution $D^{(P)}$ inside the patient. The dose criteria are typically defined using Dose-Volume Histograms (DVH). In particular, we minimize the distance between the delivered dose $D(\mathbf{x})$ and the desired dose in the patient $D^{(P)}$ in the discretized volume with the beamlet weights $\mathbf{x}$ being the optimization variables (see in Fig. 1e). We will describe the process through the following quadratic cost function:

$$F(\mathbf{x}) = \sum_{r=0}^{R} \sum_{i=1}^{\mathcal{V}_r} \gamma_i \left[D_i(\mathbf{x}) - D_i^{(P)} \right]^2 \qquad (2)$$

where r is an index running over all the volumes (the targeted tumor and the OARs), and γ_i a weight assigned to the i^{th} voxel in order to prioritize certain volumes during the treatment. $\mathcal{V}_r$ gives the total number of voxels belonging to the object r.

III. SOLVING CLASSICAL PROBLEMS WITH TENSOR NETWORKS

A. Mapping the problem to a classical Hamiltonian

In the following, we describe how to rewrite the optimization problem in Eq. (2) into a ground-state search of a quantum many-body Hamiltonian. In particular, we here propose a procedure based on the binary-decimal conversion to map the cost function $F(\mathbf{x})$ into an Ising-type Hamiltonian, $\mathcal{H}_{IMRT}$. The procedure is summarized in Fig. 2.

We first discretize the weights x_j for each beamlet by a set of N_Q bits. N_Q is called the *bit-depth*. Thus, we represent

$$x_j \approx \frac{1}{(B/2)} \sum_{n=1}^{N_Q} 2^{n-1} b_n^{(j)} \qquad (3)$$

with the bits $b_n^{(j)} = \{0, 1\}$, and introducing a normalization constant B to set the range of the beamlets weights such that $x_j \in \left[0; \frac{2^{N_Q}-1}{(B/2)} \right]$. With increasing number of bits N_Q we can increase the resolution of the discretization. Then, we map the binary values $b_n^{(j)}$ into spin variables $s_n^{(j)} = \{-1; +1\} = 2b_n^{(j)} - 1$ for each site j. Consequently, we construct a $N_B \times N_Q$-dimensional many-body Hamiltonian with the first dimension running over the different beamlets and the second representing the discretized space for each beamlet. Inserting the spin variables together with Eq. (3) into the cost function of Eq. (2), we obtain this Ising-type Hamiltonian with the following general expression:

$$\mathcal{H}_{IMRT} = \mathcal{H}_{SP} + \mathcal{H}_{INT}^{(a)} + \mathcal{H}_{INT}^{(b)} \qquad (4)$$

where the first term describes the single particle terms with

$$\mathcal{H}_{SP} = \sum_{j,n} \left[\sum_i \frac{\gamma_i}{B} \left(a_{ij} \sum_k a_{ik} - 2D_i^{(P)} a_{ij} \right) 2^{n-1} \right] s_n^{(j)} \, , \qquad (5)$$

the second term captures spins interacting in the same beamlet

$$\mathcal{H}_{INT}^{(a)} = \sum_j \sum_{m \neq n} \left[\sum_i \gamma_i \frac{a_{ij}^2}{B^2} 2^{n-1} 2^{m-1} \right] s_n^{(j)} s_m^{(j)} \, , \qquad (6)$$

and the last term represents the interactions between different beamlets

$$\mathcal{H}_{INT}^{(b)} = \sum_{j \neq k} \sum_{n,m} \left[\sum_i \gamma_i \frac{a_{ij} a_{ik}}{B^2} 2^{n-1} 2^{m-1} \right] s_n^{(j)} s_m^{(k)} \, . \qquad (7)$$

We point out, that this Hamiltonian describes a two-dimensional fully-connected lattice of long-range interacting spins with N_B sites on one direction and N_Q on
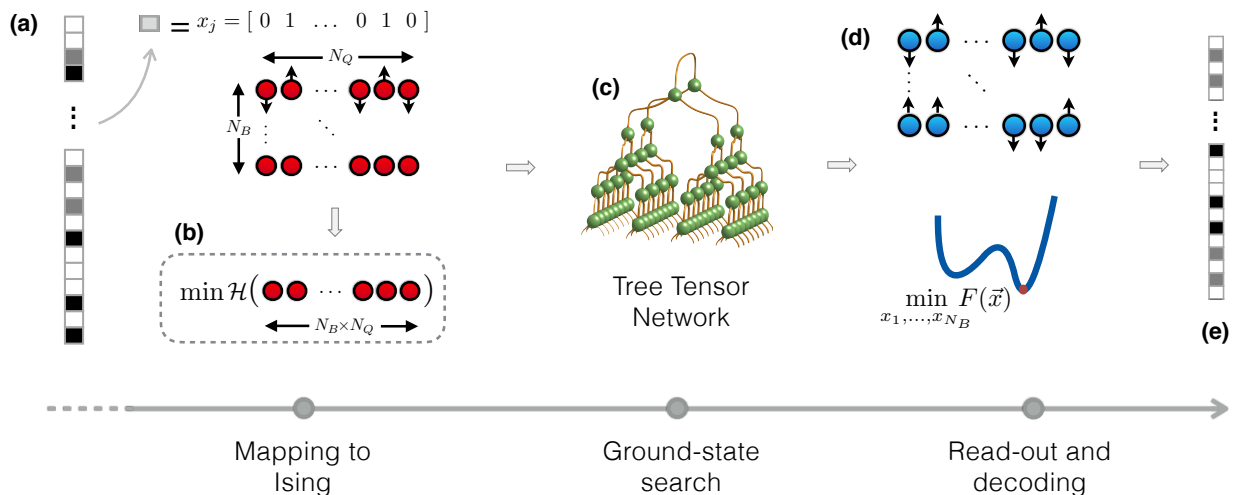
FIG. 2: Solving the classical IMRT optimization problem as a quantum Hamiltonian with TNs. The beamlet weights, x_j, are represented as a set of long-range pairwise interacting spins (a). The initial problem of minimizing a cost function is thus mapped into a ground-state search problem for the Ising-like Hamiltonian (b) which can be efficiently solved using the Tree Tensor Network algorithm (c). After the minimization the final spin configuration we read out the results by reconstructing the optimal values for the beamlet weights (d, e).

the other. At this point, we can solve the initial problem of minimizing Eq. (2) by finding the ground-state for the classical Ising-type Hamiltonian in Eq. (4).

The ground state of the system is given by the classical spin-configuration which provides the lowest energy E_0, that is the configuration that minimizes Eq. (2). The case $E_0 = 0$ corresponds to the optimal beamlet setup which results in exactly the desired dose distribution within the patient. Anyhow, in practice, the exact desired dose distribution is not always achievable for the optimization. Thus, we can find the solution of the initial optimization problem by going through all possible spin-configurations of the Ising-type Hamiltonian. Anyhow, this classical search rapidly becomes unfeasible, as the number of spin-configurations grows exponentially with increasing system size. And since the number of beamlets required in a radiotherapy treatment can easily get in the order of a few thousand, the number of spins required to represent them is practically too high to solve the problem efficiently in the classical regime. For this reason, in the following section we introduce an approach based on TNs to address this complex optimization task.

B. Description of TTN for solving quantum many-body systems

Considering the classical Hamiltonian $\mathcal{H}_{IMRT}$, we further allow each spin to be a quantum variable by representing $s_n^{(j)}$ with the Pauli matrix σ_z. In this way, the quantum Ising-type Hamitonian $\mathcal{H}_{IMRT}$ is a diagonal matrix with its entries corresponding to the energies of all possible spin combinations of the classical Hamil-

tonian. Here, TNs are a vital tool for finding ground states and their physical properties despite the exponentially growing Hilbert space [17, 19, 43]. In the following, we introduce the idea of TNs to address this challenging task of investigating complex quantum many-body systems. For a more in depth introduction on TNs we refer to sophisticated literature [17–19, 21]

TNs are used to efficiently represent quantum many-body wavefunctions $|\psi\rangle$, which live in the tensor product $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \mathcal{H}_N$ of N local Hilbert spaces $\mathcal{H}_k$, each assumed to be of finite dimension d. Expressing such a state in real-space product basis means decomposing the wavefunction as

$$|\psi\rangle = \sum_{i_1,\dots i_L=1}^{d} c_{i_1,\dots,i_L} |i_1\rangle_1 \otimes |i_2\rangle_2 \otimes \dots \otimes |i_L\rangle_L , \quad (8)$$

where $\{|i\rangle_k\}_i$ is the canonical basis of site k, spanning $\mathcal{H}_k$. The exact description of such a general state by all possible combinations of local states requires d^N coefficients $c_{i_1,\dots,i_N}$. Thus, the number of coefficients increases exponentially with the system size N in the exact representation of the wave-function.

TNs offer a more efficient representation by decomposing the complete rank-N tensor (containing all d^N coefficients) into a set of local tensors with smaller rank, connected with auxiliary indices. We control the dimension of the auxiliary indices with the bond-dimension χ and thereby the amount of captured information. Thus, tuning this parameter χ, TNs interpolate between a product state, where quantum correlations are neglected, and the exact, but inefficient representation.

The decomposition of the complete rank-N tensor can

be executed in several ways giving rise to different Tensor Network geometries. The most prominent TN representations are the Matrix Product States (MPS) for 1D systems [17, 44, 45], which addresses all sites with one corresponding tensor and their two-dimensional variant, the Projected Entangled Pair States (PEPS) [21, 46, 47]. Tree Tensor Networks (TTN) [26, 48–50] with their hierarchical structure can in principle be defined in any lattice dimension. As an illustration of the representation power, the MPS, for instance, reduces the number of parameters to the upper bound of $Nd\chi^2$ controlled by the chosen bond-dimension χ, leading to a linear dependence on the system size N rather than an exponential one.

In our analysis, we use the latter, a TTN. The TTN offers better connectivity between long-range interactions (with a logarithmic distance threw the network), while for the simpler MPS the distance by connecting tensors within the network is linear. In contrast to the PEPS, the TTN can be optimized with a lower computational complexity as well ($\mathcal{O}\left(\chi^4\right)$ for the TTN vs. $\mathcal{O}\left(\chi^{10}\right)$ for the PEPS)

Due to the bond-dimension $\chi > 1$, we perform a quantum ground state search within the subspace limited by χ. Thus, in contrast to the classical optimization routines, we explore within one optimization step several classical solutions as they are superposed in the quantum representation of the TTN. This allows us further to tunnel through higher, but reasonably thin, potentials within the optimization landscape.

After converging to the quantum ground state, we in general still obtain a superposition of classical solutions. As we know, that all solutions to the problem are classical, in theory, when the TTN algorithm is fully converged to the ground state, each of the superposed classical solutions separately has the same ground state energy E_0. Due to this possible degeneracy, we can select one classical solution from the TTN in the following way: We truncate the bond-dimension down to $\chi = 1$, leaving us with a separable, mean-field solution

$$|\Psi_{\chi=1}\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes ... \otimes |\psi_L\rangle \ ,$$

in which the only superposition can be local (such as a local site being $|\psi_n\rangle = 1/\sqrt{2}(|\uparrow\rangle + |\downarrow\rangle)$). From here on, we measure the quantum observable $\langle\sigma_n^z\rangle$ for each site n, resulting in $\langle\sigma_n^z\rangle = -1$ for a spin down, $\langle\sigma_n^z\rangle = 1$ for a spin up and in between those in case of a local superposition. In the case of the latter, we project the spin to the classical one with the highest probability by using the sign. The resulting spin configuration is further mapped back to the binary encoded solution for each voxel as described in the previous section.

In the case of the IMRT optimization problem, we are dealing with a non-trivial quantum Spin-Glass Hamiltonian type with over 32000 long-range interactions. Thus, the optimization is a highly non-trivial task and for the TTN-algorithm can be quite sensitive to the initialization procedure. Therefore, for each run, we randomly initialize several samples of the TTN from which we start the optimization. In the end, we can verify the best simulation by comparing the resulting energies.

IV. ANALYSIS

In this section, we compare the cancer treatment optimization performed with the TTN approach against QP and SA. In particular, we show the applications for two different toy models to validate our approach followed by a more realistic anatomical scenario simulating a prostate IMRT treatment.

A. Methods

We first compare the TTN algorithm with the analytical solution for a 3D box toy-model scenario. For the second, the two-sphere model and the more realistic IMRT phantom, we lack a general analytical solution. Therefore, we evaluate the results of our TTN approach by comparing it with two different optimization methods, QP and SA. In particular, QP refers to a set of methods and algorithms used to solve quadratic optimization problems subject to linear constraints and it was exploited to address the initial optimization problem in Eq. (2). On the other hand, the problem in the Ising-type formulation in Eq. (4) is addressed using both TTN and SA. QP and SA are used to validate the results obtained with the TTN approach. For further details on the QP and SA, we refer to App. A.

We point out, that both SA and TNs algorithms contain elements of randomness: samples of independent and randomly initialized simulations are always collected and the best solution considered. This also allows us to calculate the standard deviation of the samples and have a quantitative idea of the general behaviour of the algorithms.

B. Toy Models

Bipartite box. In this section we describe a simplified analytical model used to validate the correctness of the mapping discussed in Sec. III A and further our TN approach.

The model consists of a 3D box subdivided into two different regions for which we assign specific dose prescriptions as shown in Fig. 3. The red number on the front of the box corresponds to the desired beamlet weights for each bipartition. We radiate the cube with two beams from two opposite directions (i.e. $\theta_1 = 0°$, here represented as two rectangles, with a variable number of beamlets N_B for each beam, two in this example. The radiation beam is modelled as an *ideal beam* which releases the same amount of dose to each voxel. Moreover, we
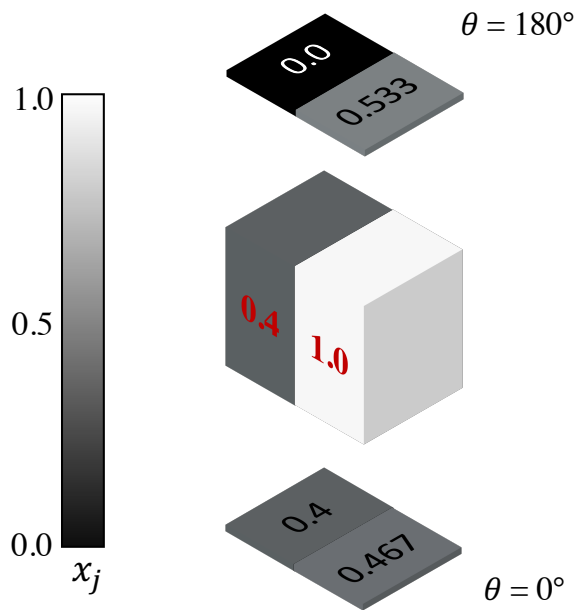
FIG. 3: **Optimization on the 3D box toy model.** The box is radiated from two opposite angles $\theta = 0°$ and $\theta = 180°$. The red numbers on the box are the desired beamlet weights; the number on the upper and lower rectangules are the beamlet weights obtained with TTN. Their sum for each partition equals the desired values.
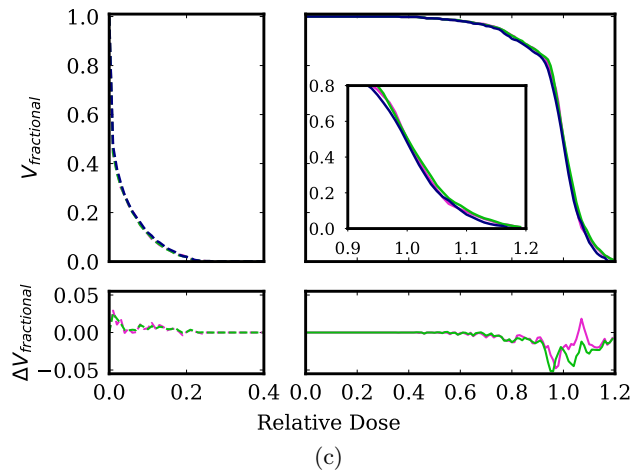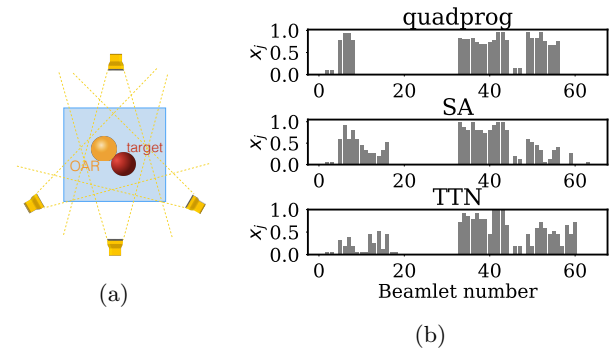

FIG. 4: **Optimization on the sphere Toy Model.** Schematic representation of the model with the the red sphere as target and the orange sphere as OAR;(b) best final beamlet configuration for the three algorithms; (c, top) Cumulative DVHs obtained with the three algorithms: quadprog (pink), SA (green) and TTN (blue). On the left panel the orange sphere (dashed line); on the right panel the red sphere (solid line). (c, bottom) Difference in the relative volume coverage between TTN and quadprog (pink line), and TTN and SA (green line).

neglect scattering effects limiting the interactions among different beamlets to those acting geometrically on the same voxels only. The dose prescriptions and the number of voxels are chosen in a way to ensure the exact ground state with energy $E_0 = 0$. Therefore, for instance, we choose dose values which are compatible with the number of discretization levels 2^{N_Q} used in Eq. (3).

We point out, that these conditions are not fulfilled in general cases. Anyhow, we introduce them to obtain an exactly solvable model for the sake of validation. Later on, we present a more realistic anatomical scenario of cancer treatment. The model is then described by the influence matrices A_r for each region r, the priorities γ_i and the dose prescriptions $D_i^{(P)}$ for each voxel. These are the same information which can be extracted from a real therapy planning system, as we will see also in the following.

The number of bits used to represent each beamlet is set to $N_Q = 4$. The target dose prescriptions for the two partitions are $D_{left}^{(P)} = 6.0$ and $D_{right}^{(P)} = 15.0$, which are arbitrary number chosen according to the number of discretization levels for the x_j, 16 in this case. The influence matrix, mapping the beamlets intensities to the voxels, is defined to have a uniform dose release into the box. Thereby, as we simplified physical effects of the beam, we end up with the influence matrix consisting of

either zero and otherwise constant entries mapping the beamlet intensities to the voxels.

The values contained inside the upper and lower and rectangles correspond to the beamlet weights obtained with the TTN algorithm. Wo observe that their sum equals the desired values for each bipartition, showing a perfect agreement between the analytical and numerical solution. This proves the well functioning of the algorithm as well as the correctness of the mapping procedure.

We point out, that, in general, there can be more than one configuration satisfying the constraints: indeed, depending on the system parameters, the ground state of the Ising-Hamiltonian can be degenerate. Anyhow, for the treatment, we are satisfied obtaining anyone of the degenerate ground states.

Sphere. In this paragraph, we illustrate the application of the TTN algorithm to a more realistic clinical scenario. In this case, instead of assuming an *ideal beam* we model the photon beam using the Matlab-based software CERR [51]. This software generates radiotherapy plans and can be used to obtain the influence matrixes a_{ij}. It allows working with physical effects introduced by the beams such as scattering and further allows for higher freedom in choosing the geometry, the number of beams and other typical model parameters. Thereby, the underlying medical images are in the standardised DICOM format.

In Fig. 4a we illustrate the model analysed in this paragraph. The model described consists of a cubic box dimensions $(50.0 \times 50.0 \times 100.0)$ cm^3 filled with water in which two spherical regions, shown in red and in orange respectively, of diameter $d = 3$ cm placed, one considered as the targeted tumor and the other one as an organ at risk. We irradiate the box using four beams at $0°$, $120°$, $180°$ and $240°$ and a total amount of $N_B = 64$ beamlets (16 beamlets/beam). The influence matrices (A_{red} and A_{orange}) are obtained using the CERR's dose calculation algorithm QIB [52] with the default settings; the dimensions of the beamlets are set (1.0×1.0) cm^2.

The optimization goals are set to $D_{red}^{(P)} = 50.0$ Gy for the red sphere and $D_{orange}^{(P)} = 0.0$ Gy for the orange one, considering the first as the targeted tumor and the second as an OAR. Each sphere is weighted equally in the cost function with $\gamma = 1.0$. During the optimization procedure, the influence matrixes are always normalized to keep the final beamlets intensities in the interval $[0, 1]$.

In the mapping to the discrete problem, the bit-depth was fixed to $N_Q = 4$ bits, resulting in a fully-connected lattice of 256 sites for the underlying Hamiltonian. The total number of non-zero interaction terms was $n_{int} = 32640$. Thus, the underlying quantum many-body system is a challenging long-range spin-glass Ising model to be solved with the TTN algorithm.

In practice, it is unfeasible to obtain $E_0 = 0$ and thereby to reach exactly the prescribed dose distribution. Thus, in this example, the optimization balances the different goals for each organ according to their priorities γ. For this reason, to evaluate the quality of the results returned by the TTN algorithm, the same optimization task was attacked using the Matlab build-in function *quadprog*, which exploits QP to optimize the cost function, and SA. We recall that the optimization with *quadprog* was directly performed on the function in Eq. (2), while SA and the TTN were applied to the discretized problem in Eq. (4).

A standard method used for plans quality evaluation is the cumulative DVH histogram, which shows a 2D projection of the 3D dose distribution inside a given volume. It represents the fractional volume receiving at least a given value of dose. Given a generic volume, r, we can easily build the dose vector $\mathbf{D}_r(\mathbf{x})$ as described in Eq. (1) by applying its influence matrix, A_r, to the beamlets vector $\mathbf{x}$. The resulting vector contains the total dose de-

livered to each one of the voxels in the volume r. By subdividing the dose interval $[0, D_i^{(max)}]$ into n_b (dose) bins, for each of them we can count how many voxels receive a dose greater or equal than the corresponding dose value. In other words, the number of entries in the k-th bin indicates the number of voxels receiving at least the corresponding dose. The obtained distribution results in a cumulative DVH with the fractional volume represented on the y-axis and the dose values on the x-axis. We point out, that in this representation of the dose distribution we lose the spatial information of the problem.

On the top panels of Figure 4c, we show the DVHs obtained with the three methods. It is clearly visible, that the three methods show a very good qualitative agreement in the resulting DVHs for each organ. This agreement is further quantitatively confirmed by the energy E_0 - or cost - after the minimization: Within the statistical uncertainty, all three methods result in a final energy $E_0 = 0.0181(8)$. By looking at the bottom panels of Figure 4c we see that the difference between TTN and the other two methods in the relative volume coverage is globally very close to zero, with only a few peaks at about 2-5 %

In Figure 4b, we present the final beamlets configuration for each of the three methods. Despite their global consistence, we observe that local differences arise. First, this is due to the fact the final configurations have slightly different energies, despite are all consistent. However, in general, there may exist more than one configuration satisfying the constraints and minimizing the energy E_0 for the underlying system and this effect is further amplified when comparing the optimization on the discrete space (SA and TTN) to that on the continuous one (quadrog). This happens because the energy landscape may be altered by the dicretization procedure. Thus, in this case, we have that the ground state of the underlying Hamiltonian is either degenerate or its energy gap is reasonably small. In fact, the underlying quantum spin-glass Hamiltonian has a highly non-trivial spectrum with many local optima and depending on the system parameters degenerate ground-states.

C. Prostate cancer treatment with TG119 IMRT phantom

We now show the results obtained on a standard IMRT phantom provided by the American Association of Physicists in Medicine Task Group 119 for use in institutional IMRT commissioning [53, 54]. This dataset contains several segmented structures and we chose the following with the aim of simulating a prostate cancer cases: prostate as the targeted tumor, bladder and rectum as the OARs.

The geometry we used is characterized by two beams (33 and 31 beamlets) placed at $90°$, $270°$ as shown in Fig. 5a, resulting in a total amount of $N_B = 64$ beamlets. The dimensions of each beamlet were (0.9×0.9)
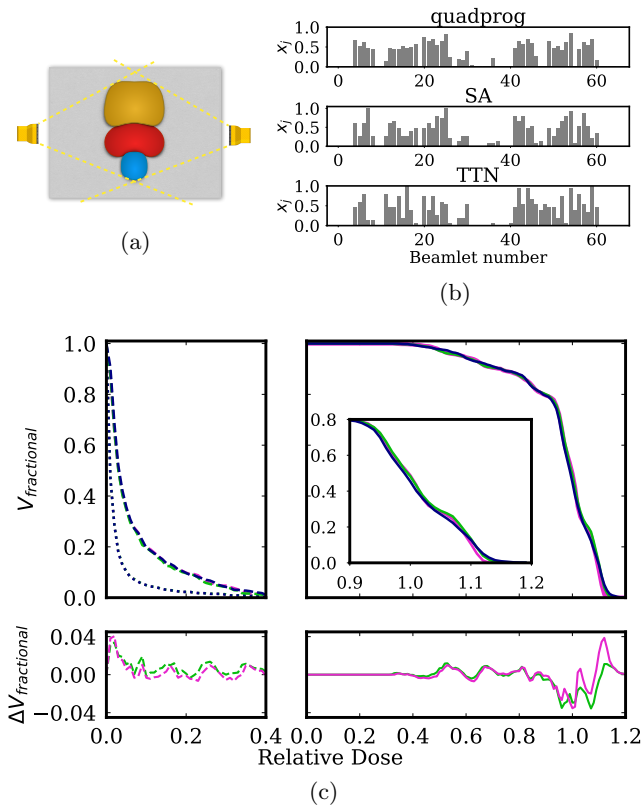
(a)

(b)

(c)

FIG. 5: **Optimization on the TG119 IMRT phantom.**(a) Schematic representation of the model with the the prostate as target (red) and the bladder (brown) and the rectum (blue) as OARs;(b) Best final beamlets configuration for the three algorithms; (c, top) Cumulative DVHs obtained with the three algorithms: quadprog (pink), SA (green) and TTN (blue). On the left panel the two OARs: bladder (dotted line) and rectum (dashed line). On the right panel the prostate (solid line). (c, bottom) Difference in the relative volume coverage for the rectum (dashed line) and the prostate (solid line) between TTN and quadprog (pink line), and TTN and SA (green line).

cm^2. The dose prescriptions were set to $D_{prostate}^{(P)} = 50$ Gy and $D_{bladder}^{(P)} = D_{rectum}^{(P)} = 0.0$ Gy and the priority assigned to the different structures was $\gamma = 1.0$ for all of them. The dose calculation algorithm used was CERR's QIB algorithm in the default settings. The fraction of non-zero elements in the influence matrixes was 0.55 for the $A_{prostate}$, 0.28 for $A_{bladder}$ and 0.45 for A_{rectum}. For the discrete problem, the system obtained is again characterized by 256 fully-connected long-range interacting spins and a total amount of $n_{int} = 32640$ interaction terms.

The top panels of Fig. 5c shows the comparison between the DVHs for the three optimization methods. The bottom panels show that the difference between TTN and the other two methods in the volume coverage for the rec-

tum and the prostate is within the 4%, proving a very good quantitative agreement between them. The results for the bladder are not shown since the differences were negligible. This agreement is additionally confirmed by the obtained energy E_0 - or cost - after the minimization: within the statistical uncertainty, the three methods result in a final energy $E_0 = 0.043(1)$. This result further confirms what found from the study on the toy models.

V. CONCLUSIONS

In this manuscript, we presented a new approach based on TNs to optimize the dose distribution for an IMRT cancer treatment. We showed a feasibility study on three different cancer treatment scenarios. First, we provided a proof-of-principle Tensor Network analysis by successfully investigating an analytically solvable toy-model of a radiated box. Then, we compared the TNs with the classical approaches of QP and SA in the case of a spherical cancer and a spherical organ at risk. Finally, we illustrated the successful application of the TNs approach to a more realistic anatomical scenario simulating a prostate cancer.

The main goal of our work was to show the applicability of TNs to the IMRT optimization problem, fostering new applications of quantum-inspired techniques to the solution of classical optimization problems. Along the road, we defined a clear strategy to map the classical problem to simulated quantum-like hardware.

Our results indicate that the TN approach can achieve results compatible with other optimization techniques such as QP and SA. We stress that for this feasibility study we (i) used a reduced number of beamlets to reduce the complexity in the models, (ii) kept the cost function convex for sake of simplicity and (iii) used the TNs code "out-of-the-box" originally engineered for typical quantum systems with significantly fewer interactions but higher entanglement. Further software developments will allow to address these three points, increasing the TN approach efficiency. In particular, extending this study to non-convex and non-differentiable functions and by further specializing and parallelizing the TNs code for this particular field of application, further significant steps forward can be made towards real-world scenarios and their use in every-day medical care.

Finally, we point out that TNs are particular examples of quantum circuits, thus, this study opens the way to the application of quantum computation to cancer treatment, for example through the application of hybrid quantum-classical optimization algorithms [55]. Once quantum computer hardware will be scaled up, one could replace the TNs simulation with actual quantum computation, possibly further enhancing our capabilities of fighting cancer via IMRT.

VI. ACKNOWLEDGMENTS

Appendix A: Classical optimization

1. Quadratic programming

Quadratic programming (QP) refers to a set of widespread methods for solving (non-linear) quadratic optimization problems subject to linear constraints. A general QP problem can be formulated as follows:

$$\min_x \ \mathbf{q}^T\mathbf{x} + \frac{1}{2}\mathbf{x}^T Q\mathbf{x} \quad s.\ t. \quad \begin{cases} A\mathbf{x} = \mathbf{a}, \\ B\mathbf{x} \leq \mathbf{b} \\ \mathbf{x} \geq \mathbf{0} \end{cases} \quad \text{(A1)}$$

where $f(\mathbf{x})$ is the target function, $\mathbf{q}$ its gradient and Q its Hessian matrix. The equation $A\mathbf{x} = \mathbf{a}$ contains all the equality constraints, while $B\mathbf{x} \leq \mathbf{b}$ all the inequality constraints.

In the QP formulation, the initial IMRT optimization problem in Eq. (2) can be written as follows:

$$\min_{\mathbf{x}} \ 2\sum_r \left(\mathbf{x}^T \tilde{A}_r^T \tilde{A}_r \mathbf{x} - \mathbf{D}_r^{(P)} \tilde{A}_r \mathbf{x} \right), \ 0 \leq \mathbf{x} \leq 1 \quad \text{(A2)}$$

where r is an index running over the different volumes or organs (the targeted tumor and the OARs), $\tilde{A}_r$ is the influence matrix for the r-th volume whose entries are weighted by the pre-assigned priorities $\gamma_{i(r)}$ to each voxel, i, and $\mathbf{D}_r^{(P)}$ a vector containing the dose prescriptions for the voxels in the r-th volume.

Matlab's *Optimization Toolbox*™ provides the function *quadprog* which exploits different solvers to attack a wide class QP problems. In this work, this methods is used to solve the initial quadratic problem and produce results to be compared to those obtained with TTN.

2. Simulated annealing

Simulated annealing (SA) is a widespread combinatorial optimization method based on randomization techniques [56], in particular based on the Metropolis-Hastings algorithm. By varying a control parameter T, called *temperature*, it is possible to explore the landscape of a target free-energy function in order to find its global minima. SA is typically applied to large-scale non-convex optimization problems where the number of local minima in the energy landscape is very high. The algorithm always requires a starting point which in practice is either a random one or the best one known for a specific problem. From the initial starting point, a rule to generate new configurations is given: Configurations with an energy lower than the previous configuration are always accepted (with probability one). On the other hand, a move towards configurations with a higher energy is accepted with a certain probability which significantly reduces the risk of getting stuck into a local minimum. In this case, the acceptance probability decreases dynamically with the temperature T, which itself decreases during the optimization from a given value T_{max} to T_{min} according to a pre-defined annealing schedule.

In this work, SA is applied to find the ground state of the classical Ising spin-glass problem representing the initial IMRT optimization problem. We always start the SA from a random configuration for the spins in the lattice. New configurations are generated by flipping a randomly chosen spin in the prior configuration. Due to the intrinsic stochastic nature of SA, we perform a statistical sampling of N_{runs} independent and randomly-initialized optimization runs for the same problem. For both, the sphere and the prostate cancer case, N_{runs} is set to 100.

The code used in this study to perform SA is based on the Python library in [57].

Appendix B: Ground-state search via TTN
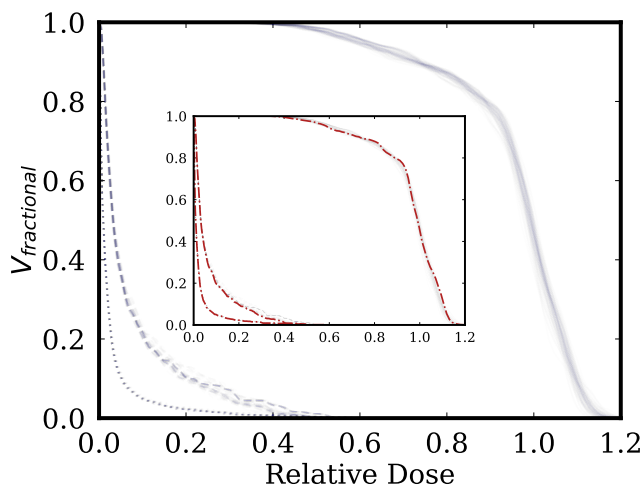
1. Trend of the solutions

The TTN algorithm explores a corner of the full many-body Hilbert space $\mathcal{H}$ which becomes wider as the bond dimension χ increases. In this work, the bond dimension was fixed to $\chi = 5$ throughout the whole study. Since we are dealing a classical problem we know that the final solution is not entangled. However, the introduction of a bond dimension $X > 1$ increases the probability for the algorithm to converge to the global minimum, as it thereby explores a greater solution space.

In Sec. IV we've shown the best results obtained with the TTN both for the prostate and the sphere over samples of $N_{runs} = 100$ runs. We recall that the tensors entries are randomly initialized at the beginning of each new optimization run.

In order to see which is the general behaviour of the algorithm throughout the whole sampling, we compare in Fig. 6a and 6b the results over all $N_{runs} = 100$ runs. We observe that they are globally very closely distributed around the best solution and this proves the precision of the TTN algorithm. However, making precise clinical considerations about the consistency between the different solutions goes far beyond the purpose of this feasibility study since many different factors should be considered depending on the specific case.
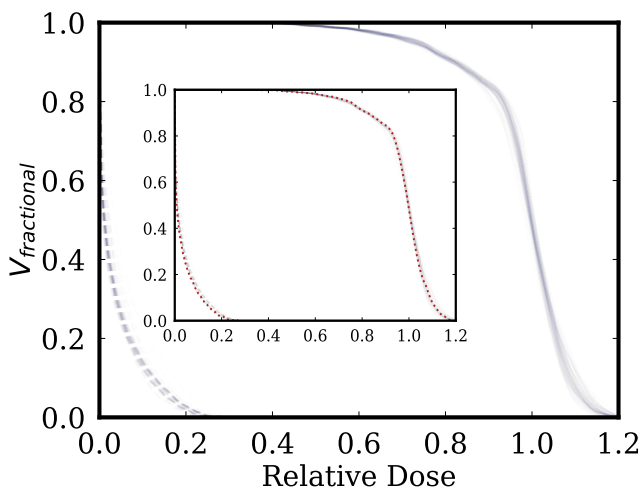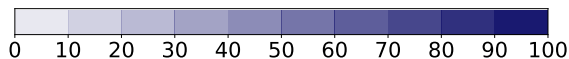
(a)

(b)

FIG. 6: **Trend of the 100 solutions obtained with the TTN algorithm.** (a) Sphere toy-model. (b) Prostate cancer case. The intensity of the blue lines is proportional to the number of superimposed solutions, according to the colorbar. The innermost panels show the comparison between the 100 solutions and the best one.

2. Reduction of interaction coefficients.

The model we have considered so far is a fully-connected lattice system of N interacting spins, with the number of interactions terms given by $N \times (N - 1)/2$. The density of the interaction scheme has an impact on the computational time required to the TTN to converge. Thereby, it is interesting to investigate the behaviour of
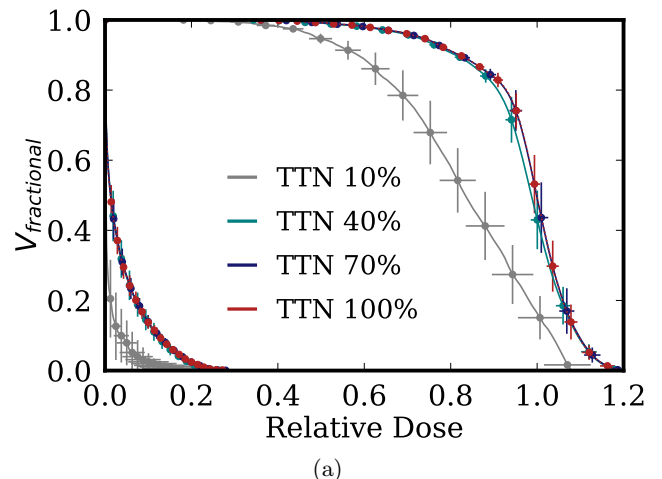
(a)
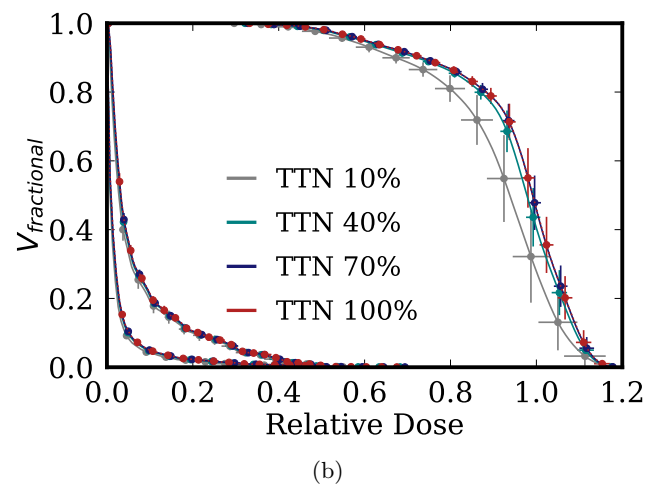
(b)

FIG. 7: **Reduction of the number of interaction coefficients.** (a) Sphere toy-model. (b) Prostate cancer.

the accuracy of the solution when neglecting some of the coefficients. In particular, we fix a threshold, λ_{min} such that $(|J_{ij}| < \lambda_{min}) = 0$, with J_{ij} the pairwise interaction term between the i-th and the j-th spin in the lattice, to cut away the interaction terms with lowest coefficients.

In the following, we considered four different cuts, keeping respectively the 10%, 40%, 70% and 100% of the coefficients. The result for the sphere cancer is shown in Fig. 7a. For each of the four cases, samples of $N_{runs} = 100$ runs were collected. Each histogram shows the average position of each bin (we set $N_{bin} = 100$ bins) with the associated standard deviation both on the relative dose and the factional volume axes. For the red sphere, (the most-right lines) 70%- and 100%- results are very well superimposed, while the 40%-line lays within the standard deviation of the former ones. Bigger differences arise between these three and the 10% line. No significative differences can be found for the orange sphere (most-left lines) between the 100%, 70% and 40% lines.
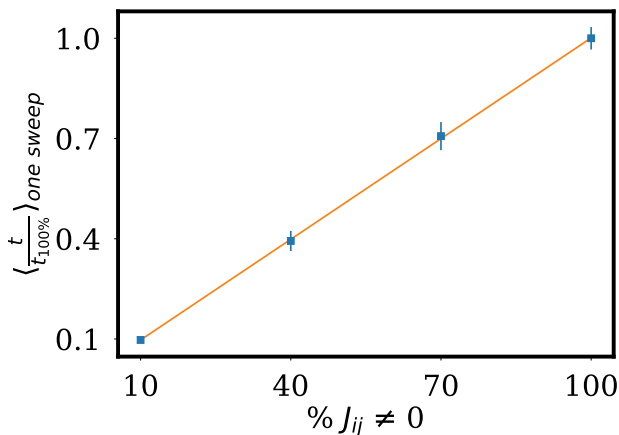
FIG. 8: **Scaling of the computational time.** The figures shows the average computational time to perform one sweep when different fractions of non-zero coefficients are considered. Each point is shown with the associated standard deviation. Each sample was made up of about 700 sweeps. Each time point is normalized by the average time to perform one sweep when the 100% of coefficients is considered. The orange line is the results of a linear fit.

Even in this case, the main differences arise between the 10% line and the others.

Interestingly, these results confirm that accurate solutions can be obtained even considering only 40% of the coefficients. The study on the prostate cancer case (see Fig. 7b) further confirms this insight. However, in this case, the 10% line is closer to the others than in the previous case.

This result becomes even more interesting if we con-sider the scaling of the computational time when reducing the number of interaction coefficients. Fig. 8 shows the result for the prostate cancer case. We observe that the scaling is linear. In conclusion, using this procedure we should be able to obtain very accurate solutions spending only the 40% of the initial computational time.

Consequently, one question arises naturally: How to find the optimal balance between speeding up the simulations by cutting terms versus keeping an adequate information to accurately describe the system? We therefore propose the following heuristic procedure:

- Build the fully-connected interaction matrix of the initial problem, J_{ij}^0 and compute its eigenvalues s_0.

- Choose a cut in order to keep the $\eta\%$ of the coefficients and compute the new eigenvalues s_{red} of the reduced interaction matrix, J_{ij}^{red};

- Check the statistical difference between the two samples of eigenvalues s_0 and s_{red} (before and after the cut). This can be performed using of an hypothesis testing procedure, where the null-hypothesis H_0 is that the two sets of eigenvalues are sampled from the same population, hence no statistical differences arise between them. Within this framework, the p-value of the test can be used to predict the accuracy of the optimization. In particular, the smaller the p-value is, the more the two sets of eigenvalues are likely to come from different populations and thus the accuracy of the optimization on the reduced model to be poor. We point out, that the choice of the statistical test in this procedure is not unique. Our trials exploited the Wilcoxon non-parametric test for dependent samples.

[1] R. Baskar, J. Dai, N. Wenlong, R. Yeo, and K.-W. Yeoh, Frontiers in Molecular Biosciences **1**, 24 (2014).
[2] J.-s. Wang, H.-j. Wang, and H.-l. Qian, Military Medical Research **5**, 20 (2018).
[3] Y. Ma, O. Kepp, F. Ghiringhelli, L. Apetoh, L. Aymeric, C. Locher, A. Tesniere, I. Martins, A. Ly, N. M. Haynes, M. J. Smyth, G. Kroemer, and L. Zitvogel, Seminars in Immunology **22**, 113 (2010), cancer Vaccines: The state of the art.
[4] S. C. Formenti and S. Demaria, JNCI: Journal of the National Cancer Institute **105**, 256 (2013).
[5] T. B. Brunner, Best Practice and Research Clinical Gastroenterology **30**, 515 (2016), chemoradiotherapy for gastrointestinal cancer.
[6] Y. Wang, W. Deng, N. Li, S. Neri, A. Sharma, W. Jiang, and S. H. Lin, Frontiers in pharmacology **9**, 185 (2018).
[7] A. Taylor and M. E. B. Powell, Cancer imaging : the official publication of the International Cancer Imaging Society **4**, 68 (2004).
[8] T. Bortfeld, Physics in Medicine and Biology **51**, R363 (2006).
[9] J. ur Rehman, Zahra, N. Ahmad, M. Khalid, H. N. ul Huda Khan Asghar, Z. A. Gilani, I. Ullah, G. Nasar, M. M. Akhtar, and M. N. Usmani, Journal of Radiation Research and Applied Sciences **11**, 361 (2018).
[10] C. Elith, S. E. Dempsey, N. Findlay, and H. M. Warren-Forward, *Journal of Medical Imaging and Radiation Sciences*, Journal of Medical Imaging and Radiation Sciences **42**, 37 (2011).
[11] C. Cotrutz, M. Lahanas, C. Kappas, and D. Baltas, Physics in Medicine and Biology **46**, 2161 (2001).
[12] S. Webb, International Journal of Imaging Systems and Technology **6**, 71 (1995).
[13] Y. Censor and J. Unkelbach, Physica Medica **28**, 109 (2012).
[14] G. A. Ezzell, Medical Physics **23**, 293 (1996).
[15] D. P. Nazareth and J. D. Spaans, *Physics in Medicine and Biology*, **60**, 4137 (2015).
[16] J. M. Pakela, H.-H. Tseng, M. M. Matuszak, R. K. Ten Haken, D. L. McShan, and I. El Naqa, *Medical*

Physics, Medical Physics **47**, 5 (2020).

[17] U. Schollwck, Annals of Physics **326**, 96192 (2011).

[18] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, SciPost Phys. Lect. Notes , 8 (2019).

[19] S. Montangero, *Introduction to Tensor Network Methods* (Springer International Publishing, 2018).

[20] G. Evenbly and G. Vidal, Phys. Rev. B **79**, 144108 (2009).

[21] R. Ors, Annals of Physics **349**, 117158 (2014).

[22] T. Felser, S. Notarnicola, and M. S., to be published.

[23] I. P. McCulloch, Journal of Statistical Mechanics: Theory and Experiment **2007**, P10014 (2007).

[24] S. Singh and G. Vidal, Phys. Rev. B **88**, 115147 (2013).

[25] M. Dalmonte and S. Montangero, Contemporary Physics **57**, 388 (2016).

[26] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte, and S. Montangero, Phys. Rev. B **96**, 195123 (2017).

[27] M. C. Bauls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. V. Acoleyen, F. Verstraete, U. J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, "Simulating lattice gauge theories within quantum technologies," (2019), arXiv:1911.00003 [quant-ph].

[28] T. Felser, P. Silvi, M. Collura, and S. Montangero, "Two-dimensional quantum-link lattice quantum electrodynamics at finite density," (2019), arXiv:1911.09693 [quant-ph].

[29] M. C. Bañuls and K. Cichy, Reports on Progress in Physics **83**, 024401 (2020).

[30] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, Quantum Science and Technology **4**, 024001 (2019).

[31] I. H. Kim and B. Swingle, "Robust entanglement renormalization on a noisy quantum computer," (2017), arXiv:1711.07500 [quant-ph].

[32] Y. Zhou, E. M. Stoudenmire, and X. Waintal, "What limits the simulation of quantum computers?" (2020), arXiv:2002.07730 [quant-ph].

[33] F. Alongi, S. Arcangeli, A. R. Filippi, U. Ricardi, and M. Scorsetti, The oncologist **17**, 1100 (2012).

[34] M. Teoh, C. H. Clark, K. Wood, S. Whitaker, and A. Nisbet, The British journal of radiology **84**, 967 (2011).

[35] D. A. Jaffray, Nature Reviews Clinical Oncology **9**, 688 (2012).

[36] C. Chargari, E. Deutsch, P. Blanchard, S. Gouy, H. Martelli, F. Gurin, I. Dumas, A. Bossi, P. Morice, A. N. Viswanathan, and C. Haie-Meder, CA: A Cancer Journal for Clinicians **69**, 386 (2019).

[37] R. Baskar, K. A. Lee, R. Yeo, and K.-W. Yeoh, International journal of medical sciences **9**, 193 (2012).

[38] S. Adeberg, S. B. Harrabi, N. Bougatf, V. Verma, P. Windisch, D. Bernhardt, S. E. Combs, K. Herfarth, J. Debus, and S. Rieken, Cancers **10** (2018), 10.3390/cancers10110401.

[39] Q. Zhang, J. Liu, N. Ao, H. Yu, Y. Peng, L. Ou, and S. Zhang, Scientific Reports **10**, 1220 (2020).

[40] A. R. Delaney, M. Dahele, J. P. Tol, I. T. Kuijper, B. J. Slotman, and W. F. A. R. Verbakel, *Radiotherapy and Oncology*, Radiotherapy and Oncology **124**, 263 (2017).

[41] *International Journal of Radiation Oncology, Biology, Physics*, International Journal of Radiation Oncology, Biology, Physics **51**, 880 (2001).

[42] T. R. Bortfeld and A. L. Boyer, International Journal of Imaging Systems and Technology **6**, 62 (1995).

[43] M. Nielsen and I. Chuang, eds., *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).

[44] S. stlund and S. Rommer, Physical Review Letters **75**, 35373540 (1995).

[45] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).

[46] F. Verstraete and J. I. Cirac, Phys. Rev. A **70**, 060302 (2004).

[47] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Phys. Rev. Lett. **96**, 220601 (2006).

[48] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Phys. Rev. A **74**, 022320 (2006).

[49] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio, Phys. Rev. A **81**, 062335 (2010).

[50] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero, Phys. Rev. B **90**, 125154 (2014).

[51] A. Apte, D. Khullar, J. Alaly, and J. O. Deasy, "Cerr," (2010).

[52] Y. Hardiyanti, M. Haekal, A. Waris, and F. Haryanto, Journal of Physics: Conference Series **739**, 012144 (2016).

[53] D. Craft, M. Bangert, T. Long, D. Papp, and J. Unkelbach, GigaScience **3**, 37 (2014).

[54] E. Villaggi, V. Hernandez, M. Fusella, E. Moretti, S. Russo, E. M. L. Vaccara, B. Nardiello, M. Esposito, J. Saez, S. Cilla, C. Marino, M. Stasi, and P. Mancosu, Physica Medica **62**, 73 (2019).

[55] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, New Journal of Physics **18**, 023023 (2016).

[56] E. H. L. Aarts and P. J. M. van Laarhoven, in *Pattern Recognition Theory and Applications*, edited by P. A. Devijver and J. Kittler (Springer Berlin Heidelberg, Berlin, Heidelberg, 1987) pp. 179–192.

[57] R. J. Wagner, "Python module for simulated annealing," (2019).

Bibliography

[1] A. G. J. MacFarlane, J. P. Dowling, and G. J. Milburn. *Quantum technology: the second quantum revolution.* Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **361**(1809), 1655 (2003). https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2003.1227, URL https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2003.1227. 1

[2] G. Kurizki, P. Bertet, Y. Kubo, K. Mølmer, D. Petrosyan, P. Rabl, and J. Schmiedmayer. *Quantum technologies with hybrid systems.* National Academy of Science of the United State of America **112**, 3866 (2015). DOI:10.1073/pnas.1419326112. 1

[3] S. Sachdev. *Quantum Phase Transitions* (Cambridge University Press, 1999). 1

[4] J. Bardeen, L. Cooper, and J. Schrieffer. *Theory of superconductivity.* Physical Review **108**, 1175 (1957). DOI:10.1103/PhysRev.108.1175. 1

[5] P. Lee, N. Nagaosa, and X.-G. Wen. *Doping a mott insulator: Physics of high temperature superconductivity.* Reviews of Modern Physics **78**, 17 (2006). DOI:10.1103/RevModPhys.78.17. 1

[6] M. Nielsen and I. Chuang, eds. *Quantum Computation and Quantum Information* (Cambridge University Press, 2000). 1, 25

[7] T. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. O'Brien. *Quantum computers.* Nature **464.4**, 45 (2010). DOI:10.1038/nature08812. 1

[8] W. Buchmüller and C. Lüdeling. *Field theory and standard model.* arXiv:hep-ph/0609174 (2006). 1, 39

[9] R. Brandenberger. *Quantum field theory methods and inflationary universe models.* Reviews of Modern Physics **57**, 1 (1985). DOI:10.1103/RevModPhys.57.1. 1

[10] A. Szabo and N. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, 1989). 1

[11] P. Atkins and R. Friedman. *Molecular Quantum Mechanics* (Oxford University Press, 2011). Original version published in 1997. 1

[12] A. Kitaev. *Fault-tolerant quantum computation by anyons.* Annals of Physics **303**(1), 2–30 (2003). URL http://dx.doi.org/10.1016/S0003-4916(02)00018-0. 1

[13] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma. *Non-abelian anyons and topological quantum computation.* Reviews of Modern Physics **80**(3), 1083–1159 (2008). URL http://dx.doi.org/10.1103/RevModPhys.80.1083.

[14] J. Clarke and F. Wilhel. *Superconducting quantum bits.* Nature **453.19**, 1031 (2008). DOI:10.1038/nature07128.

[15] R. Blatt and C. Roos. *Quantum simulations with trapped ions.* Nature **8**, 277 (2012). DOI:10.1038/NPHYS2252. 1

[16] M. Sarovar, A. Ishizaki, G. Fleming, and B. Whaley. *Quantum entanglement in photosynthetic light-harvesting complexes.* Nature **6**, 462–467 (2010). DOI:10.1038/nphys1652. 1

[17] W. P. Schleich, K. S. Ranade, C. Anton, M. Arndt, M. Aspelmeyer, M. Bayer, G. Berg, T. Calarco, H. Fuchs, E. Giacobino, M. Grassl, P. Hänggi, W. M. Heckl, I.-V. Hertel, S. Huelga, F. Jelezko, B. Keimer, J. P. Kotthaus, G. Leuchs, N. Lütkenhaus, U. Maurer, T. Pfau, M. B. Plenio, E. M. Rasel, O. Renn, C. Silberhorn, J. Schiedmayer, D. Schmitt-Landsiedel, K. Schönhammer, A. Ustinov, P. Walther, H. Weinfurter, E. Welzl, R. Wiesendanger, S. Wolf, A. Zeilinger, and P. Zoller. *Quantum technology: from research to application.* Applied Physics B **122**(5), 130 (2016). URL https://doi.org/10.1007/s00340-016-6353-8. 1

[18] U. Schollwöck. *The density-matrix renormalization group in the age of matrix product states.* Annals of Physics **326.11**, 96 (2011). 1

[19] M. Troyer and U.-J. Wiese. *Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations.* Phys. Rev. Lett. **94**, 170201 (2005). URL https://link.aps.org/doi/10.1103/PhysRevLett.94.170201. 1, 2, 3

[20] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac. *Computational complexity of projected entangled pair states.* Phys. Rev. Lett. **98**, 140506 (2007). URL https://link.aps.org/doi/10.1103/PhysRevLett.98.140506. 2, 29

[21] S.-J. Ran, E. Tirrito, C. Peng, X. Chen, L. Tagliacozzo, G. Su, and M. Lewenstein. *Two-Dimensional Tensor Networks and Contraction Algorithms*, pp. 63–86 (Springer International Publishing, Cham, 2020). URL https://doi.org/10.1007/978-3-030-34489-4_3. 4, 11, 28, 115

[22] L. Tagliacozzo, G. Evenbly, and G. Vidal. *Simulation of two-dimensional quantum systems using a tree tensor network that exploits the entropic area law.* Physical Review B **80**(23) (2009). URL http://dx.doi.org/10.1103/PhysRevB.80.235127. 4, 5, 28, 32, 115

[23] T. Felser, S. Notarnicola, and S. Montangero. *Efficient tensor network ansatz for high-dimensional quantum many-body problems* (2021). URL https://link.aps.org/doi/10.1103/PhysRevLett.126.170603. iii, 2, 8, 16, 116

[24] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero. *The Tensor Networks Anthology: Simulation techniques for many-body quantum lattice systems.* SciPost Phys. Lect. Notes p. 8 (2019). URL https://scipost.org/10.21468/SciPostPhysLectNotes.8. v, 2, 4, 5, 11, 13, 19, 28, 30, 32, 35, 47, 50, 59, 74, 77, 115, 132, 134, 135, 138

[25] S. de Léséleuc, V. Lienhard, P. Scholl, D. Barredo, S. Weber, N. Lang, H. P. Büchler, T. Lahaye, and A. Browaeys. *Observation of a symmetry-protected topological phase of interacting bosons with rydberg atoms*. Science **365**(6455), 775 (2019). https://science.sciencemag.org/content/365/6455/775.full.pdf, URL https://science.sciencemag.org/content/365/6455/775. 2

[26] A. Browaeys and T. Lahaye. *Many-body physics with individually controlled rydberg atoms*. Nature Physics **16**(2), 132 (2020). URL https://doi.org/10.1038/s41567-019-0733-z. 115

[27] J. Rui, D. Wei, A. Rubio-Abadal, S. Hollerith, J. Zeiher, D. M. Stamper-Kurn, C. Gross, and I. Bloch. *A subradiant optical mirror formed by a single structured atomic layer*. Nature **583**(7816), 369 (2020). URL https://doi.org/10.1038/s41586-020-2463-x.

[28] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin. *Probing many-body dynamics on a 51-atom quantum simulator*. Nature **551**(7682), 579 (2017). URL https://doi.org/10.1038/nature24622.

[29] P. Schauß, M. Cheneau, M. Endres, T. Fukuhara, S. Hild, A. Omran, T. Pohl, C. Gross, S. Kuhr, and I. Bloch. *Observation of spatially ordered structures in a two-dimensional rydberg gas*. Nature **491**(7422), 87 (2012). URL https://doi.org/10.1038/nature11596. 115

[30] A. Elben, B. Vermersch, R. van Bijnen, C. Kokail, T. Brydges, C. Maier, M. K. Joshi, R. Blatt, C. F. Roos, and P. Zoller. *Cross-platform verification of intermediate scale quantum devices*. Phys. Rev. Lett. **124**, 010504 (2020). URL https://link.aps.org/doi/10.1103/PhysRevLett.124.010504.

[31] T. Olsacher, L. Postler, P. Schindler, T. Monz, P. Zoller, and L. M. Sieberer. *Scalable and parallel tweezer gates for quantum computing with long ion strings*. PRX Quantum **1**(2) (2020). URL http://dx.doi.org/10.1103/PRXQuantum.1.020316. 115

[32] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller. *Self-verifying variational quantum simulation of lattice models*. Nature **569**(7756), 355 (2019). URL https://doi.org/10.1038/s41586-019-1177-4.

[33] A. Celi, B. Vermersch, O. Viyuela, H. Pichler, M. D. Lukin, and P. Zoller. *Emerging two-dimensional gauge theories in rydberg configurable arrays*. Phys. Rev. X **10**, 021057 (2020). URL https://link.aps.org/doi/10.1103/PhysRevX.10.021057. 2

[34] X.-G. Wen. *Quantum Field Theory of Many-Body Systems* (Oxford University Press, 2004). 2

[35] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte, and S. Montangero. *Fractional quantum hall effect in the interacting hofstadter model via tensor networks*. Phys. Rev. B **96**, 195123 (2017). URL https://link.aps.org/doi/10.1103/PhysRevB.96.195123. 2, 5, 28, 30, 115

[36] X.-G. Wen. *Colloquium : Zoo of quantum-topological phases of matter*. Reviews of Modern Physics **89**(4) (2017). URL http://dx.doi.org/10.1103/RevModPhys.89.041004. 2

[37] T. Felser, P. Silvi, M. Collura, and S. Montangero. *Two-dimensional quantum-link lattice quantum electrodynamics at finite density*. Phys. Rev. X **10**, 041040 (2020). URL https://link.aps.org/doi/10.1103/PhysRevX.10.041040. iii, 2, 5, 8, 11, 28, 30, 88, 115

[38] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. Van Acoleyen, F. Verstraete, U.-J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller. *Simulating lattice gauge theories within quantum technologies*. The European Physical Journal D **74**(8), 165 (2020). URL https://doi.org/10.1140/epjd/e2020-100571-8. 11

[39] M. C. Bañuls and K. Cichy. *Review on novel methods for lattice gauge theories*. Reports on Progress in Physics **83**(2), 024401 (2020). URL https://doi.org/10.1088%2F1361-6633%2Fab6311. 2

[40] H. Weimer, A. Kshetrimayum, and R. Orús. *Simulation methods for open quantum many-body systems* (2021). URL http://dx.doi.org/10.1103/RevModPhys.93.015008. 2

[41] Y. Ashida. *Quantum Many-Body Physics in Open Systems: Measurement and Strong Correlations* (Springer Singapore, 2020). 2

[42] N. Brambilla, S. Eidelman, P. Foka, S. Gardner, A. S. Kronfeld, M. G. Alford, R. Alkofer, M. Butenschoen, T. D. Cohen, J. Erdmenger, and et al. *Qcd and strongly coupled gauge theories: challenges and perspectives*. The European Physical Journal C **74**(10) (2014). URL http://dx.doi.org/10.1140/epjc/s10052-014-2981-5. 2

[43] P. A. M. Dirac and R. H. Fowler. *Quantum mechanics of many-electron systems*. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character **123**(792), 714 (1929). https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1929.0094, URL https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1929.0094. 3

[44] A. Klümper, A. Schadschneider, and J. Zittartz. *Matrix product ground states for one-dimensional spin-1 quantum antiferromagnets*. Europhysics Letters (EPL) **24**(4), 293–297 (1993). URL http://dx.doi.org/10.1209/0295-5075/24/4/010. 3, 4

[45] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac. *Criticality, the area law, and the computational power of projected entangled pair states*. Phys. Rev. Lett. **96**, 220601 (2006). URL https://link.aps.org/doi/10.1103/PhysRevLett.96.220601. 5, 26, 28, 29, 115

[46] F. Verstraete and J. I. Cirac. *Matrix product states represent ground states faithfully*. Phys. Rev. B **73**, 094423 (2006). URL https://link.aps.org/doi/10.1103/PhysRevB.73.094423. 26, 28, 115

[47] J. Eisert, M. Cramer, and M. Plenio. *Colloquium: Area laws for the entanglement entropy*. Reviews of Modern Physics **349**, 117 (2013). Eprint:1306.2164.

[48] M. B. Hastings. *Solving gapped hamiltonians locally*. Phys. Rev. B **73**, 085115 (2006). URL https://link.aps.org/doi/10.1103/PhysRevB.73.085115. 3, 4, 26, 28, 115

[49] R. B. Laughlin. *Anomalous quantum hall effect: An incompressible quantum fluid with fractionally charged excitations.* Phys. Rev. Lett. **50**, 1395 (1983). URL https://link.aps.org/doi/10.1103/PhysRevLett.50.1395. 3

[50] J. Kolorenč and L. Mitas. *Applications of quantum monte carlo methods in condensed systems.* Reports on Progress in Physics **74**(2), 026502 (2011). URL http://dx.doi.org/10.1088/0034-4885/74/2/026502. 3, 49

[51] P. H. Acioli. *Review of quantum monte carlo methods and their applications.* Journal of Molecular Structure: THEOCHEM **394**(2), 75 (1997). Proceedings of the Eighth Brazilian Symposium of Theoretical Chemistry, URL http://www.sciencedirect.com/science/article/pii/S016612809604821X.

[52] A. W. Sandvik. *Finite-size scaling of the ground-state parameters of the two-dimensional heisenberg model.* Physical Review B **56**(18), 11678–11690 (1997). URL http://dx.doi.org/10.1103/PhysRevB.56.11678. 137, 138, 140

[53] P. Emonts, M. C. Bañuls, I. Cirac, and E. Zohar. *Variational monte carlo simulation with tensor networks of a pure z_3 gauge theory in $(2 + 1)$D.* Phys. Rev. D **102**, 074501 (2020). URL https://link.aps.org/doi/10.1103/PhysRevD.102.074501. 3, 49

[54] W. H. and G. C. YOUNG. *Die fundamentalen physikalischen eigenschaften der krystalle in elementarer darstellung.* Nature **58**(1492), 99 (1898). URL https://doi.org/10.1038/058099a0. 3

[55] S. White. *Density matrix formulation for quantum renormalization groups.* Physical Review Letters **69.19**, 2863 (1992). 3

[56] S. White. *Numerical renormalization-group study of low-lying eigenstates of the antiferromagnetic s=1 heisenberg chain.* Physical Review B **48.6**, 3844 (1993). 3

[57] K. G. Wilson. *Renormalization group and critical phenomena. i. renormalization group and the kadanoff scaling picture.* Phys. Rev. B **4**, 3174 (1971). URL https://link.aps.org/doi/10.1103/PhysRevB.4.3174. 4

[58] K. G. Wilson. *The renormalization group: Critical phenomena and the kondo problem.* Rev. Mod. Phys. **47**, 773 (1975). URL https://link.aps.org/doi/10.1103/RevModPhys.47.773. 4

[59] I. P. McCulloch. *From density-matrix renormalization group to matrix product states.* Journal of Statistical Mechanics: Theory and Experiment **2007**(10), P10014 (2007). URL https://doi.org/10.1088/1742-5468/2007/10/p10014. 4, 5, 12

[60] S. Oestlund and S. Rommer. *Thermodynamical limit of density matrix renormalization.* Physical Review Letters **75.19**, 3537 (1995). 4

[61] U. Schollwöck. *The density-matrix renormalization group in the age of matrix product states.* Annals of Physics **326**(1), 96 (2011). URL https://www.sciencedirect.com/science/article/abs/pii/S0003491610001752. 4, 11, 28, 30, 32, 56, 115

[62] G. Vidal. *Efficient classical simulation of slightly entangled quantum computations.* Phys. Rev. Lett. **91**, 147902 (2003). URL https://link.aps.org/doi/10.1103/PhysRevLett.91.147902. 4, 28, 33

[63] S. R. White and A. E. Feiguin. *Real-time evolution using the density matrix renormalization group.* Phys. Rev. Lett. **93**, 076401 (2004). URL https://link.aps.org/doi/10.1103/PhysRevLett.93.076401. 4, 28, 33, 115

[64] M. Zwolak and G. Vidal. *Mixed-state dynamics in one-dimensional quantum lattice systems: A time-dependent superoperator renormalization algorithm.* Phys. Rev. Lett. **93**, 207205 (2004). URL https://link.aps.org/doi/10.1103/PhysRevLett.93.207205. 4, 33

[65] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac. *Matrix product density operators: Simulation of finite-temperature and dissipative systems.* Phys. Rev. Lett. **93**, 207204 (2004). URL https://link.aps.org/doi/10.1103/PhysRevLett.93.207204. 4, 33

[66] G. Vidal. *Entanglement renormalization.* Phys. Rev. Lett. **99**, 220405 (2007). URL https://link.aps.org/doi/10.1103/PhysRevLett.99.220405. 4, 5, 28, 29, 115

[67] Y.-Y. Shi, L.-M. Duan, and G. Vidal. *Classical simulation of quantum many-body systems with a tree tensor network.* Phys. Rev. A **74**, 022320 (2006). URL https://link.aps.org/doi/10.1103/PhysRevA.74.022320. 4, 5, 28, 115

[68] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio. *Homogeneous binary trees as ground states of quantum critical hamiltonians.* Phys. Rev. A **81**, 062335 (2010). URL https://link.aps.org/doi/10.1103/PhysRevA.81.062335. 4, 5, 28, 115

[69] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero. *Unconstrained tree tensor network: An adaptive gauge picture for enhanced performance.* Physical Review B **90**, 125154 (2014). DOI:10.1103/PhysRevB.90.125154. 4, 49

[70] M. Gerster, M. Rizzi, F. Tschirsich, P. Silvi, R. Fazio, and S. Montangero. *Superfluid density and quasi-long-range order in the one-dimensional disordered bose-hubbard model.* New Journal of Physics **18**(1), 015015 (2016). 4

[71] F. Verstraete and J. I. Cirac. *Valence-bond states for quantum computation.* Phys. Rev. A **70**, 060302 (2004). URL https://link.aps.org/doi/10.1103/PhysRevA.70.060302. 4, 5, 28, 115

[72] E. Stoudenmire and S. R. White. *Studying two-dimensional systems with the density matrix renormalization group.* Annual Review of Condensed Matter Physics **3**(1), 111–128 (2012). URL http://dx.doi.org/10.1146/annurev-conmatphys-020911-125018. 4, 29, 138

[73] G. Evenbly and G. Vidal. *Entanglement renormalization in two spatial dimensions.* Phys. Rev. Lett. **102**, 180406 (2009). URL https://link.aps.org/doi/10.1103/PhysRevLett.102.180406. 4, 15, 29, 115

[74] G. Evenbly and G. Vidal. *Algorithms for entanglement renormalization.* Phys. Rev. B **79**, 144108 (2009). URL https://link.aps.org/doi/10.1103/PhysRevB.79.144108. 5, 15, 28, 29, 129, 130

[75] B. Swingle. *Entanglement renormalization and holography.* Phys. Rev. D **86**, 065007 (2012). URL https://link.aps.org/doi/10.1103/PhysRevD.86.065007. 6, 29, 115

[76] G. Evenbly and G. Vidal. *Tensor network states and geometry.* Journal of Statistical Physics **145**, 891 (2011). DOI:10.1007/s10955-011-0237-4. 5, 28, 29

[77] R. Orús. *A practical introduction to tensor networks: Matrix product states and projected entangled pair states.* Annals of Physics **349**, 117 (2014). URL https://www.sciencedirect.com/science/article/abs/pii/S0003491614001596. v, 5, 28, 29, 115

[78] F. Verstraete and J. I. Cirac. *Renormalization algorithms for quantum-many body systems in two and higher dimensions* (2004). arXiv:0407066[cond-mat.str-el].

[79] D. Perez-Garcia, F. Verstraete, J. I. Cirac, and M. M. Wolf. *Peps as unique ground states of local hamiltonians* (2007). arXiv:0707.2260[quant-ph]. 5, 28, 115

[80] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero. *Unconstrained tree tensor network: An adaptive gauge picture for enhanced performance.* Phys. Rev. B **90**, 125154 (2014). URL https://link.aps.org/doi/10.1103/PhysRevB.90.125154. 5, 28, 30, 115

[81] A. Weichselbaum. *Non-abelian symmetries in tensor networks: A quantum symmetry space approach.* Annals of Physics **327**, 2972 (2012). 5, 12, 35, 38, 39, 77

[82] M. Nozaki, S. Ryu, and T. Takayanagi. *Holographic geometry of entanglement renormalization in quantum field theories.* Journal of High Energy Physics **2012**(10), 193 (2012). URL https://doi.org/10.1007/JHEP10(2012)193. 6, 29

[83] L. Susskind. *The Black Hole War: My Battle with Stephen Hawking to Make the World Safe for Quantum Mechanics* (Little, Brown and Company, 2008). 6

[84] B. Zwiebach. *A First Course in String Theory* (Cambridge University Press, 2009). 6

[85] P. Kharyuk, D. Nazarenko, and I. Oseledets. *Comparative study of discrete wavelet transforms and wavelet tensor train decomposition to feature extraction of ftir data of medicinal plants* (2018). arXiv:1807.07099[eess.SP]. 6

[86] W. Wang, V. Aggarwal, and S. Aeron. *Principal component analysis with tensor train subspace* (2018). arXiv:1803.05026[cs.LG].

[87] Q. Zhang, L. T. Yang, Z. Chen, and P. Li. *A tensor-train deep computation model for industry informatics big data feature learning.* IEEE Transactions on Industrial Informatics **14**(7), 3197 (2018). 6

[88] I. V. Oseledets. *Tensor-train decomposition.* SIAM Journal on Scientific Computing **33**, 2295–2317 (2011). 6

[89] W. Hackbusch and S. Kühn. *A new scheme for the tensor representation.* Journal of Fourier Analysis and Applications **15**, 706 (2009). URL https://doi.org/10.1007/s00041-009-9094-9. 6

[90] T. Felser, M. Trenti, L. Sestini, A. Gianelle, D. Zuliani, D. Lucchesi, and S. Montangero. *Quantum-inspired machine learning on high-energy physics data* (2020). arXiv:2004.13747[stat.ML]. 6

[91] E. Stoudenmire and D. J. Schwab. *Supervised learning with tensor networks*. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, eds., *Advances in Neural Information Processing Systems 29*, pp. 4799–4807 (Curran Associates, Inc., 2016). URL http://papers.nips.cc/paper/6211-supervised-learning-with-tensor-networks.pdf. 6

[92] E. M. Stoudenmire. *Learning relevant features of data with multi-scale tensor networks*. Quantum Science and Technology **3**(3), 034003 (2018). URL https://doi.org/10.1088/2058-9565/aaba1a.

[93] I. Glasser, N. Pancotti, and J. I. Cirac. *From probabilistic graphical models to generalized tensor networks for supervised learning* (2018). arXiv:1806.05964[quant-ph]. 6

[94] Raeky. *Top antitop quark events* (2009). URL https://commons.wikimedia.org/wiki/File:Top_antitop_quark_event.svg. 6

[95] A. Amini, A. Soleimany, S. Karaman, and D. Rus. *Spatial uncertainty sampling for end-to-end control* (2019). arXiv:1805.04829[cs.AI]. 6

[96] J. Steppan. *Mnist examples* (2017). URL https://commons.wikimedia.org/wiki/File:MnistExamples.png. 6

[97] G. Carleo and M. Troyer. *Solving the quantum many-body problem with artificial neural networks*. Science **355**(6325), 602–606 (2017). URL http://dx.doi.org/10.1126/science.aag2302. 9, 138, 143

[98] M. Collura, L. Dell'Anna, T. Felser, and S. Montangero. *On the descriptive power of neural-networks as constrained tensor networks with exponentially large bond dimension*. SciPost Physics Core **4**(1) (2021). URL http://dx.doi.org/10.21468/SciPostPhysCore.4.1.001. iv, 9, 144

[99] S. Cavinato, T. Felser, M. Fusella, M. Paiusco, and S. Montangero. *Optimizing radiotherapy plans for cancer treatment with tensor networks* (2020). arXiv:2010.09552[physics.med-ph]. v, 9, 187

[100] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire. *Towards quantum machine learning with tensor networks*. Quantum Science and Technology **4**(2), 024001 (2019). URL http://dx.doi.org/10.1088/2058-9565/aaea94. 9, 169

[101] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik. *The theory of variational hybrid quantum-classical algorithms*. New Journal of Physics **18**(2), 023023 (2016). URL https://doi.org/10.1088%2F1367-2630%2F18%2F2%2F023023. 9, 169

[102] S. Montangero. *Introduction to Tensor Network Methods* (Springer International Publishing, 2018). 11

[103] R. Orús. *Tensor networks for complex quantum systems*. Nature Reviews Physics **1**(9), 538 (2019). URL https://doi.org/10.1038/s42254-019-0086-7.

[104] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, and et al. *Simulating lattice gauge theories within quantum technologies* (2020). URL http://dx.doi.org/10.1140/epjd/e2020-100571-8.

[105] R. Augusiak, F. M. Cucchietti, and M. Lewenstein. *Many-body physics from a quantum information perspective.* Lecture Notes in Physics p. 245–294 (2012). URL http://dx.doi.org/10.1007/978-3-642-10449-7_6. 11

[106] S. Singh, H.-Q. Zhou, and G. Vidal. *Simulation of one-dimensional quantum systems with a global SU(2) symmetry.* New Journal of Physics **12**(3), 033029 (2010). URL https://doi.org/10.1088/1367-2630/12/3/033029. 12

[107] S. Singh and G. Vidal. *Global symmetries in tensor network states: Symmetric tensors versus minimal bond dimension.* Physical Review B **88.11** (2013). DOI:10.1103/PhysRevB.88.115147. 12, 50

[108] T. Felser. *Tensor network analysis on non-abelien quantum many-body systems.* Master thesis at the University of Ulm (Institute for Complex Quantum Systems) (2017). v, 13, 35, 39, 42, 74, 77

[109] Z. Xianyi, W. Qian, and W. Saar. *Openblas* (2017). URL http://www.openblas.net/. 16

[110] A. Petitet. *Pblas* (1995). URL http://www.netlib.org/utk/papers/scalapack/node9.html. 16

[111] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide* (Society for Industrial and Applied Mathematics, 1999), third ed. https://epubs.siam.org/doi/pdf/10.1137/1.9780898719604, URL https://epubs.siam.org/doi/abs/10.1137/1.9780898719604. 16, 19

[112] Flatiron Institute. URL https://tensornetwork.org/software/. 16

[113] B. Kågström, P. Ling, and C. van Loan. *Gemm-based level 3 blas: High-performance model implementations and performance evaluation benchmark.* ACM Trans. Math. Softw. **24**(3), 268–302 (1998). URL https://doi.org/10.1145/292395.292412. 18

[114] S. Hirata. *Tensor contraction engine: abstraction and automated parallel implementation of configuration-interaction, coupled-cluster, and many-body perturbation theories.* The Journal of Physical Chemistry A **107**(46), 9887 (2003). URL https://doi.org/10.1021/jp034596z. 19

[115] E. Di Napoli, D. Fabregat-Traver, G. Quintana-Ortí, and P. Bientinesi. *Towards an efficient use of the blas library for multilinear tensor contractions.* Applied Mathematics and Computation **235**, 454 (2014). URL https://www.sciencedirect.com/science/article/pii/S0096300314002902. 19

[116] J. Li, C. Battaglino, I. Perros, J. Sun, and R. Vuduc. *An input-adaptive and in-place approach to dense tensor-times-matrix multiply.* In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '15 (Association for Computing Machinery, New York, NY, USA, 2015). URL https://doi.org/10.1145/2807591.2807671.

[117] Y. Shi, U. N. Niranjan, A. Anandkumar, and C. Cecka. *Tensor contractions with extended blas kernels on cpu and gpu.* 2016 IEEE 23rd International Conference on

High Performance Computing (HiPC) (2016). URL http://dx.doi.org/10.1109/HiPC.2016.031.

[118] P. Springer. *High-performance tensor operations : tensor transpositions, spin summations, and tensor contractions.* Dissertation, RWTH Aachen University, Aachen (2019). Veröffentlicht auf dem Publikationsserver der RWTH Aachen University; Dissertation, RWTH Aachen University, 2019, URL https://publications.rwth-aachen.de/record/755345. 19

[119] C. E. Shannon. *A mathematical theory of communication.* The Bell System Technical Journal **27**(3), 379 (1948). 25

[120] C. E. Shannon. *A mathematical theory of communication.* The Bell System Technical Journal **27**(4), 623 (1948). 25

[121] J. Eisert, M. Cramer, and M. B. Plenio. *Colloquium: Area laws for the entanglement entropy.* Reviews of Modern Physics **82**(1), 277–306 (2010). URL http://dx.doi.org/10.1103/RevModPhys.82.277. 26, 29

[122] M. M. Wolf, F. Verstraete, M. B. Hastings, and J. I. Cirac. *Area laws in quantum systems: Mutual information and correlations.* Phys. Rev. Lett. **100**, 070502 (2008). URL https://link.aps.org/doi/10.1103/PhysRevLett.100.070502. 26

[123] K. Audenaert, J. Eisert, M. B. Plenio, and R. F. Werner. *Entanglement properties of the harmonic chain.* Phys. Rev. A **66**, 042327 (2002). URL https://link.aps.org/doi/10.1103/PhysRevA.66.042327. 26

[124] I. Arad, Z. Landau, and U. Vazirani. *Improved one-dimensional area law for frustration-free systems.* Phys. Rev. B **85**, 195145 (2012). URL https://link.aps.org/doi/10.1103/PhysRevB.85.195145.

[125] F. G. S. L. Brandão and M. Horodecki. *Exponential decay of correlations implies area law.* Communications in Mathematical Physics **333**(2), 761–798 (2014). URL http://dx.doi.org/10.1007/s00220-014-2213-8. 26

[126] P. Calabrese and J. Cardy. *Entanglement entropy and quantum field theory.* Journal of Statistical Mechanics: Theory and Experiment **2004**(06), P06002 (2004). URL https://doi.org/10.1088/1742-5468/2004/06/p06002. 26, 32

[127] G. Vidal, J. I. Latorre, E. Rico, and A. Kitaev. *Entanglement in quantum critical phenomena.* Phys. Rev. Lett. **90**, 227902 (2003). URL https://link.aps.org/doi/10.1103/PhysRevLett.90.227902.

[128] C. Holzhey, F. Larsen, and F. Wilczek. *Geometric and renormalized entropy in conformal field theory.* Nuclear Physics B **424**(3), 443 (1994). URL https://www.sciencedirect.com/science/article/pii/0550321394904022. 26

[129] K. Van Acoleyen, M. Mariën, and F. Verstraete. *Entanglement rates and area laws.* Physical Review Letters **111**(17) (2013). URL http://dx.doi.org/10.1103/PhysRevLett.111.170501. 26

[130] S. Michalakis. *Stability of the area law for the entropy of entanglement* (2012). arXiv:1206.6900[quant-ph]. 26

[131] M. B. Plenio, J. Eisert, J. Dreißig, and M. Cramer. *Entropy, entanglement, and area: Analytical results for harmonic lattice systems.* Phys. Rev. Lett. **94**, 060503 (2005). URL https://link.aps.org/doi/10.1103/PhysRevLett.94.060503. 26

[132] M. Cramer and J. Eisert. *Correlations, spectral gap and entanglement in harmonic quantum systems on generic lattices.* New Journal of Physics **8**(5), 71 (2006). URL https://doi.org/10.1088/1367-2630/8/5/071. 26

[133] M. Cramer, J. Eisert, and M. B. Plenio. *Statistics dependence of the entanglement entropy.* Phys. Rev. Lett. **98**, 220603 (2007). URL https://link.aps.org/doi/10.1103/PhysRevLett.98.220603. 26

[134] M. M. Wolf. *Violation of the entropic area law for fermions.* Phys. Rev. Lett. **96**, 010404 (2006). URL https://link.aps.org/doi/10.1103/PhysRevLett.96.010404.

[135] D. Gioev and I. Klich. *Entanglement entropy of fermions in any dimension and the widom conjecture.* Phys. Rev. Lett. **96**, 100503 (2006). URL https://link.aps.org/doi/10.1103/PhysRevLett.96.100503. 26

[136] S. Östlund and S. Rommer. *Thermodynamic limit of density matrix renormalization.* Phys. Rev. Lett. **75**, 3537 (1995). URL https://link.aps.org/doi/10.1103/PhysRevLett.75.3537. 28, 115

[137] S. R. White. *Density matrix formulation for quantum renormalization groups.* Phys. Rev. Lett. **69**, 2863 (1992). URL https://link.aps.org/doi/10.1103/PhysRevLett.69.2863. 28, 115

[138] M. Lubasch, J. I. Cirac, and M.-C. Bañuls. *Algorithms for finite projected entangled pair states.* Physical Review B **90**(6) (2014). URL http://dx.doi.org/10.1103/PhysRevB.90.064425. 29, 138

[139] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza. *Contracting projected entangled pair states is average-case hard.* Phys. Rev. Research **2**, 013010 (2020). URL https://link.aps.org/doi/10.1103/PhysRevResearch.2.013010. 29

[140] B. Swingle. *Constructing holographic spacetimes using entanglement renormalization* (2012). arXiv:1209.3304[hep-th]. 29

[141] A. Mollabashi, M. Naozaki, S. Ryu, and T. Takayanagi. *Holographic geometry of cmera for quantum quenches and finite temperature.* Journal of High Energy Physics **2014**(3), 98 (2014). URL https://doi.org/10.1007/JHEP03(2014)098.

[142] P. Hayden, S. Nezami, X.-L. Qi, N. Thomas, M. Walter, and Z. Yang. *Holographic duality from random tensor networks.* Journal of High Energy Physics **2016**(11) (2016). URL http://dx.doi.org/10.1007/JHEP11(2016)009. 29

[143] L. Kohn, P. Silvi, M. Gerster, M. Keck, R. Fazio, G. E. Santoro, and S. Montangero. *Superfluid-to-mott transition in a bose-hubbard ring: Persistent currents and defect formation.* Phys. Rev. A **101**, 023617 (2020). URL https://link.aps.org/doi/10.1103/PhysRevA.101.023617. 30, 33, 167

[144] G. Magnifico, T. Felser, P. Silvi, and S. Montangero. *Lattice quantum electrodynamics in (3+1)-dimensions at finite density with tensor networks* (2020). 2011.10658. iv, 30, 88, 171

[145] A. J. Ferris. *Area law and real-space renormalization.* Phys. Rev. B **87**, 125139 (2013). URL https://link.aps.org/doi/10.1103/PhysRevB.87.125139. 30

[146] N. Nakatani and G. K.-L. Chan. *Efficient tree tensor network states (ttns) for quantum chemistry: Generalizations of the density matrix renormalization group algorithm.* The Journal of Chemical Physics **138**(13), 134113 (2013). https://doi.org/10.1063/1.4798639, URL https://doi.org/10.1063/1.4798639. 30

[147] G. M. Crosswhite, A. C. Doherty, and G. Vidal. *Applying matrix product operators to model systems with long-range interactions.* Phys. Rev. B **78**, 035116 (2008). URL https://link.aps.org/doi/10.1103/PhysRevB.78.035116. 30, 56

[148] I. P. McCulloch. *From density-matrix renormalization group to matrix product states.* Journal of Statistical Mechanics: Theory and Experiment **2007**(10), P10014 (2007). URL https://doi.org/10.1088%2F1742-5468%2F2007%2F10%2Fp10014. 30, 56

[149] P. Czarnik and J. Dziarmaga. *Variational approach to projected entangled pair states at finite temperature.* Phys. Rev. B **92**, 035152 (2015). URL https://link.aps.org/doi/10.1103/PhysRevB.92.035152. 30, 134

[150] P. Czarnik, M. M. Rams, and J. Dziarmaga. *Variational tensor network renormalization in imaginary time: Benchmark results in the hubbard model at finite temperature.* Phys. Rev. B **94**, 235142 (2016). URL https://link.aps.org/doi/10.1103/PhysRevB.94.235142.

[151] A. Kshetrimayum, H. Weimer, and R. Orús. *A simple tensor network algorithm for two-dimensional steady states.* Nature Communications **8**(1) (2017). URL http://dx.doi.org/10.1038/s41467-017-01511-6. 30, 134

[152] R. Orús, T.-C. Wei, O. Buerschaper, and A. García-Saez. *Topological transitions from multipartite entanglement with tensor networks: A procedure for sharper and faster characterization.* Phys. Rev. Lett. **113**, 257202 (2014). URL https://link.aps.org/doi/10.1103/PhysRevLett.113.257202. 32

[153] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete. *Time-dependent variational principle for quantum lattices.* Phys. Rev. Lett. **107**, 070601 (2011). URL https://link.aps.org/doi/10.1103/PhysRevLett.107.070601. 33

[154] P. Barmettler, M. Punk, V. Gritsev, E. Demler, and E. Altman. *Relaxation of antiferromagnetic order in spin-1/2 chains following a quantum quench.* Phys. Rev. Lett. **102**, 130603 (2009). URL https://link.aps.org/doi/10.1103/PhysRevLett.102.130603. 33

[155] C. Kollath, A. M. Läuchli, and E. Altman. *Quench dynamics and nonequilibrium phase diagram of the bose-hubbard model.* Phys. Rev. Lett. **98**, 180601 (2007). URL https://link.aps.org/doi/10.1103/PhysRevLett.98.180601.

[156] F. Pellegrini, S. Montangero, G. E. Santoro, and R. Fazio. *Adiabatic quenches through an extended quantum critical region.* Phys. Rev. B **77**, 140404 (2008). URL https://link.aps.org/doi/10.1103/PhysRevB.77.140404.

[157] B. Gardas, J. Dziarmaga, and W. H. Zurek. *Dynamics of the quantum phase transition in the one-dimensional bose-hubbard model: Excitations and correlations induced by a quench.* Phys. Rev. B **95**, 104306 (2017). URL https://link.aps.org/doi/10.1103/PhysRevB.95.104306. 33

[158] P. Doria, T. Calarco, and S. Montangero. *Optimal control technique for many-body quantum dynamics.* Phys. Rev. Lett. **106**, 190501 (2011). URL https://link.aps.org/doi/10.1103/PhysRevLett.106.190501. 33

[159] L. Arceci, P. Silvi, and S. Montangero. *Entanglement of formation of mixed many-body quantum states via tree tensor operators* (2020). arXiv:2011.01247[quant-ph]. 33, 167

[160] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero. *Positive tensor network approach for simulating open quantum many-body systems.* Phys. Rev. Lett. **116**, 237201 (2016). URL https://link.aps.org/doi/10.1103/PhysRevLett.116.237201. 33

[161] A. J. Daley. *Quantum trajectories and open many-body quantum systems.* Advances in Physics **63**(2), 77 (2014). https://doi.org/10.1080/00018732.2014.933502, URL https://doi.org/10.1080/00018732.2014.933502. 33

[162] J. Lowe. *Quantum chemistry* (Academic Press, 2005). 34

[163] P. R. Bunker. *Molecular Symmetry and Spectroscopy* (Academic Press, 1979).

[164] D. A. McQuarrie and J. D. Simon. *Physical Chemistry: A Molecular Approach* (University Science Books: Sausalito, 1997). 34

[165] N. Wedell. *Female receptivity in butterflies and moths.* Journal of Experimental Biology **208**(18), 3433 (2005). https://jeb.biologists.org/content/208/18/3433.full.pdf, URL https://jeb.biologists.org/content/208/18/3433. 34

[166] C. Koshio, M. Muraji, H. Tatsuta, and S.-i. Kudo. *Sexual selection in a moth: effect of symmetry on male mating success in the wild.* Behavioral Ecology **18**(3), 571 (2007). https://academic.oup.com/beheco/article-pdf/18/3/571/17276366/arm017.pdf, URL https://doi.org/10.1093/beheco/arm017. 34

[167] J. Wagemans. *Characteristics and models of human symmetry detection.* Trends in Cognitive Sciences **1**, 346 (1997). 34

[168] E. Noether. *Invariante variationsprobleme.* Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen pp. 236–257 (1918). 34

[169] S. Singh. *Tensor network states and algorithms in the presence of abelian and non-abelian symmetries* (2012). arXiv:1203.2222[quant-ph]. 35, 38, 42, 77, 138

[170] E. Lieb. *The hubbard model: Some rigorous results and open problems.* arXiv:cond-mat/9311033 (1993). 38, 39

[171] M. V. Rakov and M. Weyrauch. *Bilinear-biquadratic spin-1 rings: an su(2)-symmetric mps algorithm for periodic boundary conditions.* Journal of Physics Communications **1**(1), 015007 (2017). URL http://dx.doi.org/10.1088/2399-6528/aa7470. 39

[172] P. Silvi. *Tensor networks: a quantum-information perspective on numerical renormalization groups* (2012). arXiv:1205.4198[quant-ph]. v, 42, 47

[173] E. R. Davidson. *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices.* Journal of Computational Physics **17**(1), 87 (1975). URL https://www.sciencedirect.com/science/article/pii/0021999175900650. 58

[174] W. E. Arnoldi. *The principle of minimized iterations in the solution of the matrix eigenvalue problem.* Quarterly Applied Mathematics **9**(1), 17 (1951). URL https://www.ams.org/journals/qam/1951-09-01/S0033-569X-1951-42792-9/. 58

[175] D. Sorensen, R. Lehoucq, C. Yang, K. Maschhoff, S. Ledru, and A. Cornet. *arpack-ng* (2017). URL https://github.com/opencollab/arpack-ng. 58

[176] J. Cervený. *gilbert* (2018). URL https://github.com/jakubcerveny/gilbert. 63

[177] J. Zhang, S.-i. Kamata, and Y. Ueshige. *A pseudo-hilbert scan algorithm for arbitrarily-sized rectangle region.* In N. Zheng, X. Jiang, and X. Lan, eds., *Advances in Machine Vision, Image Processing, and Pattern Analysis*, pp. 290–299 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006). 63

[178] N. Nakatani and G. K.-L. Chan. *Efficient tree tensor network states (ttns) for quantum chemistry: Generalizations of the density matrix renormalization group algorithm.* The Journal of Chemical Physics **138**(13), 134113 (2013). URL http://dx.doi.org/10.1063/1.4798639. 75, 167

[179] C. S. Applications and Innovation. *Cineca marconi a2 - supercomputing hardware* (2016). URL https://www.hpc.cineca.it/hardware/marconi. 86

[180] E. M. Stoudenmire and S. R. White. *Real-space parallel density matrix renormalization group.* Physical Review B **87**(15) (2013). URL http://dx.doi.org/10.1103/PhysRevB.87.155137. 86, 135, 169

[181] G. Baskaran and P. W. Anderson. *Gauge theory of high-temperature superconductors and strongly correlated fermi systems.* Phys. Rev. B **37**, 580 (1988). URL https://link.aps.org/doi/10.1103/PhysRevB.37.580. 87

[182] A. Mezzacapo, E. Rico, C. Sabín, I. L. Egusquiza, L. Lamata, and E. Solano. *Non-abelian su(2) lattice gauge theories in superconducting circuits.* Phys. Rev. Lett. **115**, 240502 (2015). URL https://link.aps.org/doi/10.1103/PhysRevLett.115.240502.

[183] D. Marcos, P. Widmer, E. Rico, M. Hafezi, P. Rabl, U.-J. Wiese, and P. Zoller. *Two-dimensional lattice gauge theories with superconducting quantum circuits.* Annals of Physics **351**, 634 (2014). URL https://www.sciencedirect.com/science/article/pii/S0003491614002711. 87

[184] C. Chamon, D. Green, and Z.-C. Yang. *Constructing quantum spin liquids using combinatorial gauge symmetry.* Phys. Rev. Lett. **125**, 067203 (2020). URL https://link.aps.org/doi/10.1103/PhysRevLett.125.067203. 87

[185] Y. Zhou, K. Kanoda, and T.-K. Ng. *Quantum spin liquid states.* Rev. Mod. Phys. **89**, 025003 (2017). URL https://link.aps.org/doi/10.1103/RevModPhys.89.025003. 87

[186] M. D. Schwartz. *Quantum Field Theory and the Standard Model* (Cambridge University Press, 2014). URL https://www.cambridge.org/us/academic/subjects/physics/theoretical-physics-and-mathematical-physics/quantum-field-theory-and-standard-model?format=HB&isbn=9781107034730. 87

[187] T. Cheng and L. Li. *Gauge theory of elementary particle physics* (Oxford University Press, 2006). URL https://global.oup.com/academic/product/gauge-theory-of-elementary-particle-physics-9780198519614. 87

[188] I. J. R. Aitchison and A. J. G. Hey. *Gauge theories in particle physics: A practical introduction. Vol. 1: From relativistic quantum mechanics to QED* (CRC Press, Bristol, UK, 2003). URL http://www-spires.fnal.gov/spires/find/books/www?cl=QC793.3.F5A34::2012. 87

[189] A. S. Kronfeld and C. Quigg. *Resource letter qcd-1: Quantum chromodynamics.* American Journal of Physics **78**(11), 1081 (2010). https://doi.org/10.1119/1.3454865, URL https://doi.org/10.1119/1.3454865.

[190] A. S. Kronfeld. *Twenty-first century lattice gauge theory: Results from the quantum chromodynamics lagrangian.* Annual Review of Nuclear and Particle Science **62**(1), 265–284 (2012). URL http://dx.doi.org/10.1146/annurev-nucl-102711-094942.

[191] X. Y. Xu, Y. Qi, L. Zhang, F. F. Assaad, C. Xu, and Z. Y. Meng. *Monte carlo study of lattice compact quantum electrodynamics with fermionic matter: The parent state of quantum phases.* Phys. Rev. X **9**, 021022 (2019). URL https://link.aps.org/doi/10.1103/PhysRevX.9.021022. 87

[192] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis. *Quantum supremacy using a programmable superconducting processor.* Nature **574**(7779), 505 (2019). URL https://doi.org/10.1038/s41586-019-1666-5. 115

[193] T. Barthel, M. Kliesch, and J. Eisert. *Real-space renormalization yields finite corre-lations.* Phys. Rev. Lett. **105**, 010502 (2010). URL https://link.aps.org/doi/10.1103/PhysRevLett.105.010502. 128

[194] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac. *Ground-state properties of quantum many-body systems: entangled-plaquette states and variational monte carlo.* New Journal of Physics **11**(8), 083026 (2009). URL http://dx.doi.org/10.1088/1367-2630/11/8/083026. 138

[195] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac. *Classical simulation of infinite-size quantum lattice systems in two spatial dimensions.* Physical Review Letters **101**(25) (2008). URL http://dx.doi.org/10.1103/PhysRevLett.101.250602. 138

[196] R. Orús and G. Vidal. *Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction.* Phys. Rev. B **80**, 094403 (2009). URL https://link.aps.org/doi/10.1103/PhysRevB.80.094403. 138

[197] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua. *Deep autoregressive models for the efficient variational simulation of many-body quantum systems.* Phys. Rev. Lett. **124**, 020503 (2020). URL https://link.aps.org/doi/10.1103/PhysRevLett.124.020503. 138

[198] L. He, H. An, C. Yang, F. Wang, J. Chen, C. Wang, W. Liang, S. Dong, Q. Sun, W. Han, W. Liu, Y. Han, and W. Yao. *Peps++: Towards extreme-scale simulations of strongly correlated quantum many-particle models on sunway taihulight.* IEEE Transactions on Parallel and Distributed Systems **29**(12), 2838 (2018). 138

[199] W.-Y. Liu, Y.-Z. Huang, S.-S. Gong, and Z.-C. Gu. *Accurate simulation for finite projected entangled pair states in two dimensions* (2020). arXiv:1908.09359[cond-mat.str-el]. 138

[200] S. Singh and G. Vidal. *Global symmetries in tensor network states: Symmetric tensors versus minimal bond dimension.* Phys. Rev. B **88**, 115147 (2013). URL https://link.aps.org/doi/10.1103/PhysRevB.88.115147. 138

[201] M. A. Nielsen. *Neural networks and Deep Learning*, vol. 3 (Determination Press, 2015). 143

[202] S. S. Haykin. *Neural networks and learning machines*, vol. 3 (Pearson Upper Saddle River, NJ, USA, 2018).

[203] C. M. Bishop. *Neural Networks for Pattern Recognition* (Oxford University Press, 1996). 143

[204] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez. *A guide to constraining effective field theories with machine learning.* Physical Review D **98**(5) (2018). URL http://dx.doi.org/10.1103/PhysRevD.98.052004. 143

[205] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta. *Re-inforcement learning in different phases of quantum control.* Phys. Rev. X **8**, 031086 (2018). URL https://link.aps.org/doi/10.1103/PhysRevX.8.031086.

[206] J. Carrasquilla and R. G. Melko. *Machine learning phases of matter.* Nature Physics **13**(5), 431 (2017). URL https://doi.org/10.1038/nphys4035.

[207] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. *Automatic chemical design using a data-driven continuous representation of molecules.* ACS Central Science **4**(2), 268–276 (2018). URL http://dx.doi.org/10.1021/acscentsci.7b00572. 143

[208] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová. *Machine learning and the physical sciences.* Reviews of Modern Physics **91**(4) (2019). URL http://dx.doi.org/10.1103/RevModPhys.91.045002. 143

[209] S. Das Sarma, D.-L. Deng, and L.-M. Duan. *Machine learning meets quantum physics.* Physics Today **72**(3), 48–54 (2019). URL http://dx.doi.org/10.1063/PT.3.4164.

[210] M. Schuld, I. Sinayskiy, and F. Petruccione. *An introduction to quantum machine learning.* Contemporary Physics **56**(2), 172–185 (2014). URL http://dx.doi.org/10.1080/00107514.2014.964942.

[211] M. Schuld and F. Petruccione. *Supervised Learning with Quantum Computers* (Springer, Cham, 2018). 143

[212] Y. Levine, D. Yakira, N. Cohen, and A. Shashua. *Deep learning and quantum entanglement: Fundamental connections with implications to network design* (2017). arXiv:1704.01552[cs.LG]. 143

[213] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang. *Equivalence of restricted boltzmann machines and tensor network states.* Phys. Rev. B **97**, 085104 (2018). URL https://link.aps.org/doi/10.1103/PhysRevB.97.085104. 143, 167

[214] D.-L. Deng, X. Li, and S. Das Sarma. *Machine learning topological states.* Physical Review B **96**(19) (2017). URL http://dx.doi.org/10.1103/physrevb.96.195145. 143

[215] Y. Nomura, A. S. Darmawan, Y. Yamaji, and M. Imada. *Restricted boltzmann machine learning for solving strongly correlated quantum systems.* Physical Review B **96**(20) (2017). URL http://dx.doi.org/10.1103/physrevb.96.205152.

[216] G. Torlai and R. G. Melko. *Learning thermodynamics with boltzmann machines.* Phys. Rev. B **94**, 165134 (2016). URL https://link.aps.org/doi/10.1103/PhysRevB.94.165134. 143

[217] M. Rizzi, S. Montangero, and G. Vidal. *Simulation of time evolution with multiscale entanglement renormalization ansatz.* Phys. Rev. A **77**, 052328 (2008). URL https://link.aps.org/doi/10.1103/PhysRevA.77.052328. 167

[218] T. Gao and V. Jojic. *Degrees of freedom in deep neural networks* (2016). arXiv:1603.09260[cs.LG]. 167

[219] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac. *Neural-network quantum states, string-bond states, and chiral topological states.* Phys. Rev. X **8**, 011006 (2018). URL https://link.aps.org/doi/10.1103/PhysRevX.8.011006. 167

[220] I. Glasser. *Tensor networks, conformal fields and machine learning: applications in the description of quantum many-body systems.* Dissertation, Technische Universität München, München (2018). 167

[221] M. P. Zaletel. *Exact tensor network ansatz for strongly interacting systems.* Dissertation, University of California, Berkeley (2014). URL https://escholarship.org/uc/item/8xh9n43q. 168

[222] M. Bañuls, K. Cichy, J. Cirac, and K. Jansen. *The mass spectrum of the schwinger model with matrix product states.* Journal of High Energy Physics **2013**(11), 158 (2013). URL https://doi.org/10.1007/JHEP11(2013)158. 168

[223] P. Silvi, Y. Sauer, F. Tschirsich, and S. Montangero. *Tensor network simulation of an su(3) lattice gauge theory in 1d.* Phys. Rev. D **100**, 074512 (2019). URL https://link.aps.org/doi/10.1103/PhysRevD.100.074512. 168

[224] S. V. Mathis, G. Mazzola, and I. Tavernelli. *Toward scalable simulations of lattice gauge theories on quantum computers.* Phys. Rev. D **102**, 094501 (2020). URL https://link.aps.org/doi/10.1103/PhysRevD.102.094501. 168

[225] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller, and R. Blatt. *Real-time dynamics of lattice gauge theories with a few-qubit quantum computer.* Nature **534**, 516 (2016). URL https://doi.org/10.1038/nature18318. 168

[226] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. J. Savage. *Quantum-classical computation of schwinger model dynamics using quantum computers.* Phys. Rev. A **98**, 032331 (2018). URL https://link.aps.org/doi/10.1103/PhysRevA.98.032331.

[227] A. Mil, T. V. Zache, A. Hegde, A. Xia, R. P. Bhatt, M. K. Oberthaler, P. Hauke, J. Berges, and F. Jendrzejewski. *A scalable realization of local u(1) gauge invariance in cold atomic mixtures.* Science **367**(6482), 1128 (2020). URL https://science.sciencemag.org/content/367/6482/1128.

[228] B. Yang, H. Sun, R. Ott, H.-Y. Wang, T. V. Zache, J. C. Halimeh, Z.-S. Yuan, P. Hauke, and J.-W. Pan. *Observation of gauge invariance in a 71-site bose–hubbard quantum simulator.* Nature **587**(7834), 392–396 (2020). URL http://dx.doi.org/10.1038/s41586-020-2910-8.

[229] C. Schweizer, F. Grusdt, M. Berngruber, L. Barbiero, E. Demler, N. Goldman, I. Bloch, and M. Aidelsburger. *A scalable realization of local u(1) gauge invariance in cold atomic mixtures.* Science **367**(6482), 1128 (2020). URL https://science.sciencemag.org/content/367/6482/1128. 168

[230] K. Wang, L. Xiao, W. Yi, S.-J. Ran, and P. Xue. *Quantum image classifier with single photons* (2020). arXiv:2003.08551[quant-ph]. 169

[231] S. Efthymiou, J. Hidary, and S. Leichenauer. *Tensornetwork for machine learning* (2019). arXiv:1906.06329[cs.LG]. 169