

<https://doi.org/10.1038/s42005-025-02125-x>

Tensor networks for lattice gauge theories beyond one dimension

Check for updates

Giuseppe Magnifico ^{1,2,3}, Giovanni Cataldi ^{3,4,5,6,7} ✉, Marco Rigobello ^{3,4,5}, Peter Majcen ^{3,4,5}, Daniel Jaschke ^{3,4,5,8}, Pietro Silvi ^{3,4,5} & Simone Montangero ^{3,4,5}

Tensor network methods are numerical tools and algorithms to study many-body quantum systems in and out of equilibrium, based on tailored variational wave functions. They have found significant applications in simulating lattice gauge theories that approach relevant problems in high-energy physics. Compared to Monte Carlo methods, they do not suffer from the sign problem, allowing them to explore challenging regimes such as finite chemical potentials and real-time dynamics. Further development is required to tackle fundamental challenges, such as accessing continuum limits or computations of large-scale quantum chromodynamics. This work reviews the state-of-the-art tensor network methods and discusses a possible roadmap for algorithmic development and strategies to enhance their capabilities and extend their applicability to open high-energy problems. We provide tailored estimates of the theoretical and computational resource scaling for attacking large-scale lattice gauge theories.

Gauge theories play a role of paramount importance in our understanding and description of the fundamental constituents of matter, their spectrum, and their interactions. At low energies, gauge theories characterize a large variety of collective phases of matter and phenomena, such as ferromagnetic superconductivity, spin liquids, topological order, and the fractional quantum Hall effect^{1–3}. At high energies, as elegantly summarized in the standard model of particle physics, they are at the heart of the microscopical description of the building blocks of our universe, i.e., quarks, leptons, and their interactions mediated by gauge bosons^{4,5}.

A powerful approach to studying and simulating gauge theories in nonperturbative regimes lies in lattice gauge theories (LGTs), where matter and gauge degrees of freedom are discretized and regularized on a finite lattice. LGTs were originally introduced by Wilson to encode quantum chromodynamics (QCD) on a lattice, as a model for quark confinement beyond the perturbative regime^{6,7}. In LGTs, matter and antimatter fermionic fields are defined on lattice sites, whereas the gauge fields live on the links connecting nearest-neighbor sites. This approach has opened the doors to the application of powerful numerical methods, such as Monte Carlo (MC), to the simulations of LGTs on classical computers⁸. In the last decades, MC methods have provided a wide variety of significant results on LGTs in the context of high-energy physics, such as phases diagrams at equilibrium, characterization of the quark-gluon plasma, precise determination of the masses of quarks, mesons, and baryons, hadronic and nucleon

form factors, hadronic spectra, and predictions for dark-matter models⁹. Furthermore, MC simulations of LGTs currently represent a powerful numerical tool to predict and interpret data from multiple large-scale high-energy experiments, such as the ones performed at the large hadron collider¹⁰.

Despite their impressive success, MC sampling methods are limited in some regimes of parameters of LGTs, such as in the presence of finite baryon chemical potentials, topological θ -terms, or for simulating out-of-equilibrium dynamics in real-time. In these cases, the notorious sign problems make the MC numerical approach ineffective and inaccurate¹¹.

Ranging from these problems, sign-problem circumventing methods based on tensor networks (TNs) have found significant applications in the simulations of LGTs in recent years^{12–14}. TNs were originally introduced as a class of variational wave functions in the field of quantum many-body (QMB) physics^{15–18}. They provide a compressed representation of physical states based on their entanglement content, capable of efficiently reproducing equilibrium properties, such as phase diagrams, and real-time dynamics of interacting quantum systems^{19–22}. TNs can also be employed to compute partition functions via transfer-matrix methods and tensor network renormalization algorithms^{23–25}. This approach can further be generalized to evaluate path integrals of gauge theories through field discretization²⁶. Here, we focus on the application of TNs to study ground-state properties and out-of-equilibrium dynamics in the Hamiltonian

¹Dipartimento di Fisica, Università di Bari, Bari, Italy. ²Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Bari, Bari, Italy. ³Dipartimento di Fisica e Astronomia “G. Galilei”, Università di Padova, Padova, Italy. ⁴Padua Quantum Technologies Research Center, Università degli Studi di Padova, Padova, Italy. ⁵Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Padova, Padova, Italy. ⁶Max Planck Institute of Quantum Optics, Garching, Germany. ⁷Munich Center for Quantum Science and Technology (MCQST), Munich, Germany. ⁸Institute for Complex Quantum Systems, Ulm University, Ulm, Germany. ✉e-mail: giovanni.cataldi@unipd.it

picture, where they represent one of the state-of-the-art numerical tools for simulating QMB systems, even with strong correlations.

In the context of high-energy physics, TN methods have proven noteworthy achievements in simulating LGTs in $(1+1)$ dimensions for both Abelian and non-Abelian gauge groups^{27–55}. They have recently found applications to Abelian LGTs up to $(3+1)$ dimensions^{56–64}, and also in simulating $(2+1)$ -dimensional non-Abelian $SU(2)$ models⁶⁵. Despite their effectiveness in these first important applications, further and intensive developments are still required to tackle, with TN methods, high-energy physics problems at the center of current research efforts, such as large-scale non-Abelian LGTs and their continuum limits.

In this work, we present a general overview of TN methods for LGTs, and we discuss a possible roadmap in terms of algorithmic development and strategies to improve TN capabilities toward the ambitious long-term goal of applying TNs to $(3+1)$ -dimensional QCD.

Importantly, TN methods share a common language with quantum computers and simulators. Thus, these developments could also be relevant for encoding, validating, and benchmarking the current and future quantum computations and simulations of LGTs on experimental quantum hardware^{66–76}.

TNs overview

In this section, we present a general overview of TN methods with a particular focus on tree tensor networks (TTNs), the latter being beneficial for LGTs in higher dimensions^{56,58,65}. However, this work's main concepts and ideas can be easily generalized and applied to other TN structures. Consider a QMB system defined on a lattice of N sites. The generic site j is described by a d -dimensional local Hilbert space \mathcal{H}_j , spanned by a local basis of vectors $|i\rangle_{1 \leq i \leq d}$. The quantum states of the whole system live on the total Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \dots \otimes \mathcal{H}_N$, that is the tensor product of the local Hilbert space of the lattice sites, with dimension d^N . Then, any pure state of the system $|\psi\rangle$ can be exactly expanded in terms of a complete basis set of \mathcal{H} :

$$|\psi\rangle = \sum_{i_1, i_2, \dots, i_N=1}^d c_{i_1, i_2, \dots, i_N} |i_1, i_2, \dots, i_N\rangle, \quad (1)$$

where $|i_1, i_2, \dots, i_N\rangle$ represents the tensor product of the local basis vectors, i.e. $|i_1\rangle \otimes |i_2\rangle \otimes \dots \otimes |i_N\rangle$. The coefficients of the linear combination c_{i_1, i_2, \dots, i_N} are, in general, complex scalars; their number scales as d^N , i.e., they scale exponentially with the system size N .

This is a fundamental limitation when solving a QMB problem on a classical computer since an exponential scaling with the number of degrees of freedom implies that the exact representation described in (1) is completely unfeasible from a computational and numerical point of view.

Nevertheless, a great variety of natural QMB systems happens to be described by ground and thermal equilibrium states that have little to moderate entanglement content. The physical states of these systems, instead of exploring the exponentially large dimension of the Hilbert space,

live in a small corner of it, which can be efficiently targeted and parameterized. This property is formally described by the entanglement area law, fulfilled by low-energy states of local Hamiltonians⁷⁷: the entanglement between a contiguous subsystem and the rest is proportional to the area of the boundary between them, instead of its volume, as happens for the majority of states in the Hilbert space. Thus, obeying area law implies that the state contains many fewer quantum correlations than expected for a generic (or random) QMB state. Small corrections to the area law exist, for instance, close to a transition point for one-dimensional quantum systems (logarithmic corrections). However, the entanglement remains overall moderate⁷⁸. From a theoretical point of view, the entanglement area law has been rigorously proven for (i) one-dimensional gapped local Hamiltonians, where the locality means that a lattice site interacts only with neighboring sites, without two-body all-to-all interactions^{79–81}; (ii) for a one-dimensional local Hamiltonian, if the Renyi entropy S_α , $\alpha < 1$ is bounded from above by $c \log(N)$ for any bipartition, where c is a constant and N the system size⁸²; (iii) for quantum states at thermal equilibrium, independently from the dimensionality of the system⁸³; (iv) if the to-be-approximated state has exponential decay of correlation functions for every pair of observables (even with unbounded support)⁸⁴. Even though rigorous proof for QMB systems in higher dimensions is lacking, several numerical and phenomenological results suggest that area law still holds in the presence of local interactions^{77,85–88}.

The area law has important implications on the TN simulation of quantum lattice models: indeed, it is possible to obtain an approximate but efficient representation capable of describing the main properties of these states if the entanglement content is low-to-moderate⁷⁷; the condition applies, for example, to ground-states and first excited states of local Hamiltonians. TNs give a natural language for this representation, by replacing the complete tensor of rank- N c_{i_1, i_2, \dots, i_N} of (1) with a chain of smaller tensors interconnected using auxiliary bond indices. The network keeps a number N of physical indices of dimension d (one for each lattice site), whereas the dimension of the bond indices (called bond dimension) is a control parameter χ that can be tuned in the numerical simulations and is related to the Schmidt decomposition^{22,30}.

The key advantage of passing from the exact representation of (1) to a TN representation is that the number of parameters in the TN is of the order $O(\text{poly}(d)\text{poly}(N)\text{poly}(\chi))$, e.g., $O(N\chi \max(\chi, d)^2)$ for the TTN. The scaling with the system size is now polynomial and not exponential. In this way, we obtain an efficient representation of the quantum state in terms of computational complexity. It is worth noting that the bond dimension χ determines the degree of entanglement and quantum correlations encoded in the TN, e.g., for $\chi = 1$ the TN describes a product state (no entanglement), whereas one recovers the exact but inefficient representation in the limit $\chi \lesssim d^N$. Tuning χ properly allows interpolation between these two extreme regimes, efficiently reproducing the entanglement of the quantum state.

The most widely used TN architectures are represented in Fig. 1. Matrix product states (MPS) are an established ansatz for one-dimensional systems, in which the structure of the lattice is reproduced with a network of tensors, one for each lattice site⁸⁹. As shown in Fig. 1a, each tensor in the bulk of the network has three indices: one physical leg of dimension d




Fig. 1 | Examples of tensor network structures. **a** Matrix product states (MPS); **b** Projected entangled pair states (PEPS); **c** Tree tensor networks (TTN) for an underlying one-dimensional system; **d** TTN for an underlying two-dimensional

square lattice. The physical links with local dimension d and the virtual links with bond dimension χ are highlighted in all the panels.

representing the local degrees of freedom and two virtual legs of dimension χ connected to the neighboring sites. In open boundary conditions, the tensors at the boundaries have one trivial leg together with one physical and one virtual leg. MPS intrinsically satisfies the area law and allows for efficient computation of scalar products between two states and physical observables. Currently, the MPS-based density matrix renormalization group (DMRG) stands as one of the most consolidated and accurate techniques for numerical simulations of one-dimensional QMB systems as well as quasi-one-dimensional systems such as ladder structures⁹⁰. The computational complexity of this ground state searching algorithm for MPS is of the order $O(Nd\chi^3 + Nd^2\chi^2)$ or $O(Nd^2\chi^3)$, depending if the algorithm optimizes a single MPS tensor at a time, or two-tensors simultaneously^{90,91}. Two-site optimization is usually important to avoid getting stuck in local minima or metastable configurations during the energy variational minimization^{92–94}. The subspace expansion is an intermediate approach with the benefits of the two-tensor update and a tunable computational cost between the two approaches³⁰.

The generalization of the MPS ansatz to two- or higher-dimensional lattices is represented by the projected entangled pair state (PEPS)^{22,95}. In this case, each tensor in the bulk has a physical leg of dimension d , and a number of χ -dimensional virtual legs depending on the coordination number of the considered lattice. For example, this coordination number is four in the case of a two-dimensional square lattice, as shown in Fig. 1b. PEPS directly encode in their structure the area law of entanglement; however, their exact contraction is an exponentially hard problem, meaning that PEPS can not be efficiently contracted for numerical computing, e.g., scalar products of states or physical observables⁹⁶. To circumvent this problem, approximate contraction methods have been developed during the last years and are still at the center of current research efforts²². But even with exploiting these approximate techniques, the computational complexity for ground state optimization remains quite high, e.g., of the order of $O(Nd^2\chi^8)$ ^{97,98}, limiting the maximum reachable bond dimensions (typical values are of the order $\chi \approx 10$).

Another important family of TN ansätze is represented by TTN, in which the wave function is decomposed into a hierarchical network of tensors that do not contain internal loops^{30,99}. This way, the network can be efficiently contracted and manipulated in polynomial time. A particular class of TTN is represented by binary TTNs, reported in Fig. 1c–d for one- and two-dimensional lattices. In these structures, tensors in the lowest layer have two physical legs of dimension d (representing two lattice sites) and a virtual leg of dimension χ , whereas, in the upper layers, they have three virtual legs of dimensions up to χ . The network intrinsically encodes a renormalization procedure, in which, at each layer, two sites are mapped into a single effective one. In models with finite-range interactions, ground search algorithms for binary TTN architectures display a numerical complexity of the order $O(Nd^2\chi^2 + N\chi^4)$, see³⁰. This is a much more favorable scaling concerning equivalent algorithms for other TN structures, such as PEPS, that allows reaching quite large values of bond dimensions ($\chi \approx 500$)¹⁰⁰. The drawback of loopless structures, such as binary TTN, is that the area law may not be explicitly reproduced in dimensions higher than one¹⁰¹, which becomes a limiting factor when large systems are addressed. The convergence and the precision of the numerical results obtained via a variational optimization of TTN can be analyzed and kept under control, e.g., by exploiting the wide range of available bond dimensions. Furthermore, it is possible to explicitly encode the area law of high-dimensional systems in the TTN ansatz by introducing an additional layer of independent disentanglers acting on different couples of lattice sites and connected to the corresponding physical legs. This process augments the expressive power of TTN, and the resulting ansatz is known as augmented Tree Tensor Network (aTTN)¹⁰². The computational complexity of variationally optimizing an aTTN structure, which means optimizing both the tensors and the disentanglers, is of the order $O(N\chi^4d^4 + N\chi d^7)$. We point out that the scaling of the computational costs with the local dimension d is particularly severe in the case of aTTN due to the presence of the disentanglers layer.

Besides variational optimization for ground state searching, the previous TN families can also be exploited to simulate the real-time dynamics of local Hamiltonians¹⁰³; in the following, we describe two popular methods. The time evolved block decimation (TEBD) algorithm, is based on a Suzuki-Trotter decomposition of the time evolution exponential¹⁰⁴. The total evolution time is discretized in small time steps. The corresponding evolution operator is computed as products of local terms, such as two-body operators, and repeatedly applied to the TN wave function to generate the time-evolved state. Each application can determine an increase in the bond dimension of the network, so an optimized truncation is needed to maintain an efficient and manageable description of the quantum state. This truncation reduces the bond dimension back to χ and is performed through a singular value decomposition that minimizes the distance between the evolved and the truncated state. In general, the TEBD method allows the simulation of the real-time dynamics for nearest-neighbor or finite-range Hamiltonians; one time-step with an MPS for a one-dimensional system with local interactions comes with a computational cost that is below a two-tensor sweep for the ground-state search algorithm.

Another method for simulating the evolution of quantum states via TN is the time-dependent variational principle (TDVP), which does not rely on the Suzuki-Trotter decomposition^{105,106}. In general, TDVP constrains the time evolution to the specific TN manifold considered, such as MPS or TTN, of a given initial bond dimension¹⁰⁷. This is obtained by projecting the action of the Hamiltonian into the tangent space of the TN manifold and then solving the time-dependent Schrödinger equation within this manifold. This approach automatically preserves the energy and the norm of the quantum states during the time evolution. Both the TDVP algorithm and the variational ground state search rely on a set of Krylov vectors and, hence, have the same computational scaling for one time step compared to one sweep¹⁰⁸.

These algorithms represent important and efficient tools for simulating with TN the real-time dynamics of QMB systems. While equilibrium states satisfy the aforementioned area law, out-of-equilibrium time evolution can generate a linear growth of entanglement. In this case, the time-evolved state requires an exponential growth of the bond dimension as a function of the total time⁸². For this reason, TN methods are currently limited to studying the dynamics for low-to-moderate times or close-to-equilibrium phenomena¹⁰⁹. In this framework, further developments are extremely important to avoid or at least mitigate this barrier, by devising new algorithms or optimizing existing strategies^{110,111}.

Hamiltonian LGTs

In the traditional MC approach to LGT, the action of a continuum gauge theory is regularized by working on a finite and discrete Euclidean spacetime —i.e., both space and (imaginary) time are discretized⁸. Instead, TN (and quantum) simulations typically rely on the Hamiltonian formalism, where time remains a real, continuous variable while a cubic lattice Λ replaces D -dimensional space. Simultaneously, Hamiltonian LGTs present features that distinguish them from other lattice models commonly simulated via TN⁷. We now discuss these properties in more detail and outline the main steps that have to be taken to exploit TN algorithms in LGT. We focus on matter-coupled LGTs of the Yang-Mills type, such as those routinely employed in high-energy physics to describe nature's fundamental interactions; the paradigmatic example is lattice QCD, the SU(3) LGT describing quarks, gluons, and their strong interactions.

LGT building blocks

As depicted in Fig. 2, LGTs involve two types of degrees of freedom, matter, and gauge, hosted respectively on lattice sites $\mathbf{x} \in \Lambda$ and links $(\mathbf{x}, \boldsymbol{\mu})$, where $\boldsymbol{\mu}$ denotes one of the lattice basis vectors. Since fundamental interactions of nature involve gauge theories with compact Lie groups, the following discussion focuses on such cases. A continuous symmetry implies the existence of a color charge, which is carried by matter and gauge fields and labeled here using Greek letters.




Fig. 2 | Graphical representation of 2D LGT degrees of freedom. fermionic matter fields $\hat{\psi}_{x,\alpha}$, defined on lattice sites, and gauge fields (the parallel transporter, $\hat{U}_{x,\mu}^{\alpha\beta}$, and the chromo-electric fields, $\hat{L}_{x,\mu}^\nu$ and $\hat{R}_{x,\mu}^\nu$), living on lattice links. A local gauge transformation at \mathbf{x} acts on a matter site and all its attached links.

Matter fields. such as the electron or quark fields. We assume that matter consists of Dirac fermions. In general, discretizing the Dirac operator introduces additional unphysical solutions in the spectrum, known as doublers. Various approaches have been developed to tackle this problem, such as Wilson fermions^{46,112,113}, which, at the price of breaking the chiral symmetry, eliminate doublers in any spatial dimension. In this manuscript, we employ staggered fermions¹¹⁴, which partially resolve the doubling problem (some doublers remain in $D > 1$) but have a lower computational cost. Then, a color-multiplet $\hat{\psi}_{x,\alpha}$ of staggered fermion fields satisfies $\{\hat{\psi}_{x,\alpha}, \hat{\psi}_{y,\beta}^\dagger\} = \delta_{\alpha\beta}\delta_{xy}$. Assuming a single matter flavor in the fundamental representation of the gauge group \mathcal{G} , gauge transformations are generated by $\hat{Q}_x^\nu = \sum_{\alpha,\beta} \hat{\psi}_{x,\alpha}^\dagger \lambda_{\alpha\beta}^\nu \hat{\psi}_{x,\beta}$, where $\{\lambda_{\alpha\beta}^\nu\}_{\nu=1}^{\dim \mathcal{G}}$ are the hermitian generators of \mathcal{G} —e.g., the electric charge for U(1) or the three spin matrices for SU(2). The discussion can be generalized to multiple matter flavors and other types of lattice fermions.

Gauge fields. such as the photon or gluon fields. They are represented by bosonic operators, namely the parallel transporter $\hat{U}_{x,\mu}^{\alpha\beta}$ and the generators of its left and right gauge transformations, $\hat{L}_{x,\mu}^\nu$ and $\hat{R}_{x,\mu}^\nu$,

$$\hat{L}_{x,\mu}^\nu, \hat{U}_{y,\mu'}^{\alpha\beta} = -\delta_{xy}\delta_{\mu\mu'} \sum_\gamma \lambda_{\alpha\gamma}^\nu \hat{U}_{x,\mu}^{\gamma\beta}, \quad (2a)$$

$$\hat{R}_{x,\mu}^\nu, \hat{U}_{y,\mu'}^{\alpha\beta} = +\delta_{xy}\delta_{\mu\mu'} \sum_\gamma \hat{U}_{x,\mu}^{\alpha\gamma} \lambda_{\gamma\beta}^\nu, \quad (2b)$$

In the Abelian U(1) case there is only one generator: the electric field, $\hat{E}_{x,\mu} = \hat{L}_{x,\mu} = \hat{R}_{x,\mu}$, for which parallel transporters act as raising operators, $[\hat{E}_{x,\mu}, \hat{U}_{x,\mu}] = \hat{U}_{x,\mu}$. In terms of the above ingredients, a possible discretization of a matter-coupled Yang-Mills theory is defined by the *Kogut-Susskind Hamiltonian*¹¹⁵

$$\begin{aligned} \hat{H}_{\text{LGT}} = & -\frac{c\hbar}{2a} \sum_{x,\mu} \sum_{\alpha,\beta} \left(s_{x,\mu} \hat{\psi}_{x,\alpha}^\dagger \hat{U}_{x,\mu}^{\alpha\beta} \hat{\psi}_{x+\mu,\beta} + \text{h.c.} \right) \\ & + mc^2 \sum_x \sum_\alpha s_x \hat{\psi}_{x,\alpha}^\dagger \hat{\psi}_{x,\alpha} \\ & + \frac{c\hbar g^2}{2a^{D-2}} \sum_{x,\mu} \hat{E}_{x,\mu}^2 \\ & - \frac{c\hbar}{2g^2 a^{4-D}} \sum_{\square} \text{Tr}(\hat{U}_{\square} + \hat{U}_{\square}^\dagger), \end{aligned} \quad (3)$$

where c is the speed of light, \hbar is the Planck constant, a is the lattice spacing, g is the gauge coupling, and m is the mass parameter. The first \hat{H}_{LGT} term describes matter hopping between neighboring lattice sites, mediated by the

gauge field; the second term is the mass-energy; $s_{x,\mu}$ and s_x are phases that arise due to the use of staggered fermions¹¹⁴. Namely, we have:

$$s_{x,\mu_x} = 1 \quad s_{x,\mu_y} = (-1)^{x+y} = s_{x,\mu_z} \quad s_x = (-1)^{x+y+z}. \quad (4)$$

The last two terms of (3) represent the (chromo)electric and (chromo)magnetic energy of the gauge field, respectively. The Casimir operator gives the electric energy density:

$$\hat{E}_{x,\mu}^2 = \sum_\nu (\hat{L}_{x,\mu}^\nu)^2 = \sum_\nu (\hat{R}_{x,\mu}^\nu)^2. \quad (5)$$

The magnetic energy density is a plaquette term that, on a cubic lattice, corresponds to a four-body interaction:

$$\hat{U}_{\square} = \sum_{\alpha,\beta,\gamma,\delta} \hat{U}_{x,\mu}^{\alpha\beta} \hat{U}_{x+\mu,\mu'}^{\beta\gamma} \hat{U}_{x+\mu+\mu',\mu''}^{\gamma\delta} \hat{U}_{x,\mu''}^{\delta\alpha} \quad (6)$$

where μ and μ' span the plaquette's plane. Plaquette terms only exist in $D > 1$, contributing to the increased complexity of quantum and TN simulations of LGTs in higher dimensions⁷¹.

Gauge field truncation

The link Hilbert space is the space of square-integrable functions on the gauge group, $L^2(\mathcal{G})$, which is infinite-dimensional for a continuous \mathcal{G} ¹¹⁶.

In one space dimension, gauge degrees of freedom are unphysical (absence of transverse polarizations) and can thus be integrated out, albeit at the price of introducing non-local interactions¹¹⁷.

Beyond one dimension, the removal is much more delicate because it requires first decoupling the gauge field's longitudinal component^{118,119}. When impossible or inconvenient to remove, gauge degrees of freedom might have to be truncated to perform TN or quantum simulation. Among known truncation recipes are quantum link Models (QLM)^{120–123}, which have already been adopted for quantum simulation of LGTs^{67,71,73,75,124–128}, finite subgroups^{38,42,129}, digitization of gauge fields¹³⁰, and fusion-algebra deformation¹³¹.

Another adopted solution is truncating the spectrum of the electric energy density operator $\|\hat{E}_{x,\mu}^2\| \leq \Theta$ on each link^{65,132}. The cutoff is conveniently imposed in the irreducible representation (irrep) basis $\{|jmn\rangle\}$ ¹¹⁶ of $L^2(\mathcal{G})$, where $\hat{E}_{x,\mu}^2$ is diagonal:

$$\hat{E}_{x,\mu}^2 |jmn\rangle = C_2(j) |jmn\rangle. \quad (7)$$

Here, m and n are indices in the j -irrep of \mathcal{G} and $C_2(j)$ is the quadratic Casimir of j ¹¹⁶. In the strong coupling limit, where the electric energy term dominates \hat{H}_{LGT} , this truncation is equivalent to an energy cutoff¹³².

Gauss law and the dressed site

The most distinctive feature of gauge theories is arguably the presence of local constraints, analogous to the Gauss law of classical electrodynamics, relating the configuration of the gauge field to the spatial distribution of charges¹³³. At the quantum level, Gauss law is the statement that only gauge-invariant states are physical, namely, $\hat{G}_x^\nu |\Psi_{\text{phys}}\rangle = 0 \forall x, \nu$, where \hat{G}_x^ν are the generators of local gauge transformations at \mathbf{x} :

$$\hat{G}_x^\nu = \hat{Q}_x^\nu + q_x^\nu + \sum_\mu [\hat{L}_{x,\mu}^\nu + \hat{R}_{x-\mu,\mu}^\nu], \quad (8)$$

with q_x^ν representing eventual static background charges (typically vanishing). On the lattice, Gauss law provides a set of *vertex constraints*, each involving a lattice site and its 2D neighboring links. Due to Gauss law, the physical Hilbert space of LGTs is much smaller than the tensor product of all local sites and link Hilbert spaces. Properly exploiting gauge symmetries can thus significantly speed up numerical simulations²⁸. Strategies that solve Gauss law by eliminating (partially or entirely) either the gauge fields or the matter fields have been developed. Nonetheless, such approaches come with




Fig. 3 | Pictorial representation of the dressed site formalism adopted for TN simulations of LGTs. **a** Starting from the original formulation of matter and gauge fields, **b** each truncated gauge link is split into two representations, one per half-link; each is equipped with a proper fermionic rishon mode ζ . **c** All the half-links are absorbed into the attached matter site, forming a gauge-invariant dressed site **d** whose Hilbert space spans all the possible gauge singlets.

specific limitations: the range of interaction has to be extended, moreover, integrating-out gauge fields become problematic in $D > 1$ ¹¹⁸, while the recipe for removing matter is a model (matter content) dependent¹³⁴.

Another possibility is to enforce Gauss law using a dressed site construction, sketched in Fig. 3 and outlined below. Dressed sites have local dimensions typically larger than those resulting from the aforementioned approaches, but they are obtained from a model-independent prescription which has the advantage of preserving the locality of the interactions^{28,132}.

Rishons. As a first step, every link is decomposed into a pair of left (*L*) and right (*R*) *rishon* degrees of freedom, each associated with a Hilbert space spanned by the basis states $|jm\rangle$, identifying $|jmn\rangle \mapsto |jm\rangle_L \otimes |jn\rangle_R$, and writing parallel transporters as rishon bilinear⁶⁵:

$$\hat{U}_{x,\mu}^{\alpha\beta} \rightarrow \sum_k \hat{\zeta}_{x,\mu}^{L(k)\alpha} \hat{\zeta}_{x+\mu,-\mu}^{R(k)\beta\dagger} \quad (9)$$

Abelian link symmetries. Physical gauge link configurations, i.e. those with the left and right rishons in the same irrep, are selected introducing a local link symmetry at the TN simulation level^{28,132}. Notice this is always Abelian, regardless of the Abelian or non-Abelian nature of the gauge group.

Gauge invariant computational basis. Crucially, the gauge generators \hat{G}_x^μ now involve only the matter site at x and its 2D neighboring rishons. Fusing these degrees of freedom in a composite site, Gauss law becomes an internal constraint that singles out the dressed site Hilbert space as their gauge invariant subspace:

$$\mathcal{H}_{\text{dress}} = \ker G \subset \mathcal{H}_{\text{matt}} \otimes (\mathcal{H}_{\text{rish}})^{\otimes 2D}. \quad (10)$$

The expansion of the gauge singlet basis states of $\mathcal{H}_{\text{dress}}$ in terms of the matter and rishon bases is computed via Clebsch-Gordan decomposition¹³².

Defermionization. It is possible to effectively eliminate fermionic matter and treat dressed sites as large spins for any gauge theory where the gauge field has a well-defined parity. Specifically, a local parity operator $\hat{P}_{x,\mu} = \hat{P}_{x,\mu}^\dagger$ such that $\hat{P}_{x,\mu}^2 = \hat{1}$ must satisfy $\{\hat{U}_{x,\mu}^{\alpha\beta}, \hat{P}_{x,\mu}\} = 0$, as it happens for \mathbb{Z}_{2N} , $U(N)$, and $SU(2N)$ ^{65,134,135}. In this case, it is possible to take fermionic rishons, and as a result, all physical (gauge invariant) dressed site operators are genuinely local, i.e. they mutually commute at a nonzero distance (as spins or bosons)¹³⁶. This applies to the Hamiltonian H_{LGT} , making gauge-defermionization particularly convenient for higher-dimensional systems, where Jordan-Wigner strings result in long-range interactions¹³⁷.

Table 1 | Dressed site Hilbert space dimensions

ℓ	(2 + 1)-dimensions			(3 + 1)-dimensions		
	U(1)	SU(2)	SU(3)	U(1)	SU(2)	SU(3)
1	35	30	164	267	178	3096
2	165	168	752	3437	3670	52,476
3	455	600	3738	18,487	35,280	813,438
4	969	1650	19,878	64,953	214,958	17,490,134
5	1771	3822	43,698	177,155	967,466	69,232,482
6	2925	7840	82,128	408,421	3,509,062	228,461,186
7	4495	14,688	212,496	835,311	10,828,494	1,245,755,754
8	6545	25,650	333,538	1,561,841	29,473,038	2,782,999,996

For increasing number ℓ of allowed electric energy density levels, the corresponding dimensions d in some 2- and 3-dimensional paradigmatic LGTs with dynamical matter and gauge groups $U(1)$, $SU(2)$, and $SU(3)$ are shown.

Scaling of the local basis dimension

Table 1 lists the dressed site dimension $d = \dim \mathcal{H}_{\text{dress}}$ associated with the first few nontrivial gauge truncations of three representative LGT models in $D = 2, 3$ space dimensions. All the considered LGTs include dynamical matter, represented by one fermionic field multiplet in the fundamental representation of the gauge group. The ℓ -th row of Table 1 is obtained keeping only the first ℓ nonzero electric energy levels (i.e., using the ℓ -th smallest quadratic Casimir eigenvalue as cutoff Θ , see subsection “Gauge field truncation”). As Table 1 shows, the local dimension increases rapidly with ℓ . For 3-dimensional non-Abelian LGTs, $d \sim O(10^4)$ is reached already within the first two truncations, making the study of the untruncated limit prohibitive.

Differently from the models typically encountered in condensed matter physics, the local dimension of LGTs can thus be a limiting factor for TN simulation | especially when d becomes comparable to commonly used TN bond dimensions ($100 \lesssim \chi \lesssim 500$ for TTN). In these cases, strategies aimed at further compressing the local computational basis are needed (see subsection “Local basis truncation”). As just discussed, such truncation strategies are particularly crucial for high-dimensional LGTs.

Several numerical analyses suggest that, in some cases, a small-to-moderate truncation of the gauge group is enough for accurately approximating the continuum limits, at least for low-energy states^{36,38,44,138–143}. However, the optimal gauge truncation depends on the Hamiltonian parameters m and g .

As an example, we consider the (2+1)D $U(1)$ LGT including dynamical matter (QED), whose Hamiltonian can be obtained from (3):

$$\begin{aligned} \hat{H}_{\text{QED}} = & \frac{\hbar c}{2a} \sum_x \left[-i \hat{\psi}_x^\dagger \hat{U}_{x,\mu_x}^\dagger \hat{\psi}_{x+\mu_x} \right. \\ & \left. - (-1)^{x+y} \hat{\psi}_x^\dagger \hat{U}_{x,\mu_y}^\dagger \hat{\psi}_{x+\mu_y} + \text{H.c.} \right] \\ & + mc^2 \sum_x (-1)^{x+y} \hat{\psi}_x^\dagger \hat{\psi}_x \\ & + \frac{g^2 \hbar c}{2a} \sum_{x,\mu} \hat{E}_{x,\mu}^2 - \frac{\hbar c}{2g^2 a} \sum_{\square} \text{Tr}(\hat{U}_{\square} + \hat{U}_{\square}^\dagger), \end{aligned} \quad (11)$$

Since the Abelian $U(1)$ group has only one generator, the gauge fields \hat{E}_x and $\hat{U}_{x,\mu}$ can be represented as

$$\hat{E}_{x,\mu} = \hat{\sigma}_{x,\mu}^z(\ell) \quad \hat{U}_{x,\mu} = \hat{\zeta}_{x,\mu}^\dagger(\ell) \cdot \hat{\zeta}_{x+\mu,-\mu}(\ell) \quad (12)$$

where $\hat{\sigma}^z(\ell)$ is the spin z -operator in the ℓ $SU(2)$ irrep, while $\hat{\zeta}(\ell)$ is a $\ell \times \ell$ ladder operator with Fermi statistics.

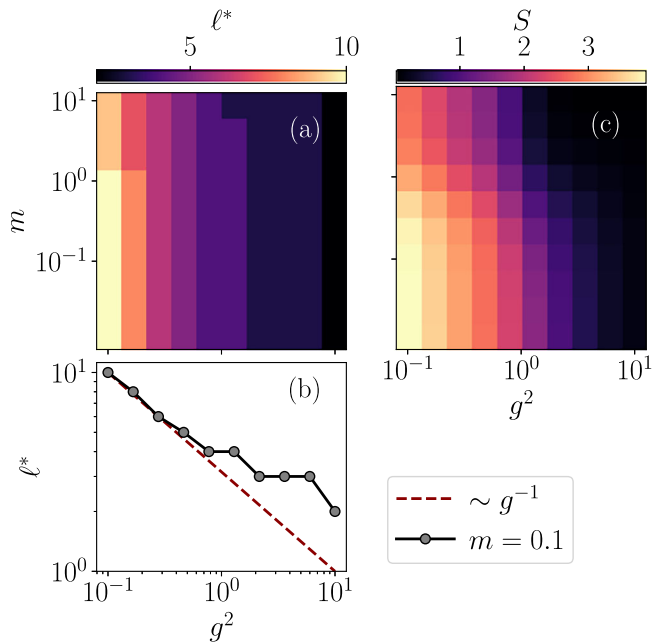


Fig. 4 | Exact diagonalization of a QED plaquette. a For a grid of masses and couplings $m \in [10^{-2}, 10^1]$ and $g^2 \in [10^{-1}, 10^1]$, minimal gauge truncation ℓ^* required to reach a precision $\epsilon_{\text{trunc}} = 10^{-5}$ in the magnetic energy ($\text{Re } \hat{U}_{\square}$). **b** Scaling of ℓ^* with g^2 for $m = 0.1$. **c** Corresponding entanglement entropy S associated with a symmetric bipartition of the plaquette (see also Supplementary Note I and Fig. S1).

We focus on a single plaquette in open boundary conditions, as it provides the minimal setting allowing for both electric and magnetic effects. Then, to characterize the convergence in the gauge truncation ℓ , we consider a candidate observable \hat{O} and compute its ground state expectation value $\langle \hat{O} \rangle_{\ell}$ for increasing ℓ . We iterate until the relative deviation between consecutive truncations drops below some threshold $\epsilon_{\text{trunc}}: |\langle \hat{O} \rangle_{\ell} - \langle \hat{O} \rangle_{\ell-1}| < \epsilon_{\text{trunc}} |\langle \hat{O} \rangle_{\ell}|$ for some ℓ^* . Figure 4a shows the minimal truncation ℓ^* at which the magnetic energy operator $\hat{O} = \text{Re } \hat{U}_{\square}$ is converged to $\epsilon_{\text{trunc}} = 10^{-5}$. We explore a grid of model parameters, whose extent has been chosen according to standard MC literature^{8,144–155}. $\text{Re } \hat{U}_{\square}$ is used as a benchmark due to its relevance in the weak coupling regime, where the continuum limit of $D < 3$ lattice QED is located⁸. As Fig. 4b shows, ℓ^* depends heavily on the coupling, growing asymptotically like $\ell^* \sim g^{-1}$ as g is decreased, while m plays almost no role. An analogous inverse dependence of the minimal gauge truncation on the coupling is expected for non-Abelian LGT in arbitrary dimensions. Moreover, the continuum limit of non-Abelian LGTs in $D \leq 3$ is also expected at $g \rightarrow 0^8$, further substantiating the need to compress the local dimension in TN simulations of LGT whenever extrapolation to the continuum is in order.

Apart from the growth of the local dimension, extrapolation to the continuum is further complicated by the fact that the continuum limit of a lattice quantum field theory corresponds to a critical point of the underlying lattice model¹⁵⁶. Close to criticality, quantum correlations are boosted and violations of the entanglement area law are expected⁹⁷. The higher entanglement entropy in the proximity of the continuum ($g, m \ll 1$ regime) is already captured by the single-plaquette analysis of Fig. 4c. Continuum limits of LGTs are thus an area of potential advantage for quantum computation over classical methods, as the former is not limited by entanglement. Nonetheless, quantum computation is also affected by the need to relax gauge truncations when $g \rightarrow 0$, either by increasing the number of qubits used to encode a dressed site, which is at least $\lceil \log_2(d) \rceil$, or by using hybrid devices with both qubits and bosonic modes¹⁵⁷.

Roadmap for advanced simulations

As detailed in the previous subsections, LGT models present some peculiar features that make TN simulations particularly challenging, especially for large system sizes and for studying the continuum limits in terms of gauge field truncation, lattice spacing, and volume.

State-of-the-art techniques, such as TTN algorithms, have been applied for simulating ground state properties of QED in $(2+1)$ - and $(3+1)$ -dimensions for small-to-intermediate sizes^{56,58}. Very recently, they have also been applied to the $\text{SU}(2)$ Yang-Mills model in $(2+1)$ -dimensions⁶⁵. In all these simulations, small non-trivial representations of the gauge groups have been exploited, e.g., three electric field levels for QED and the first two IRREP of $\text{SU}(2)$ for the Yang-Mills model.

Nowadays, lattice computations with MC-based techniques are performed on large lattices, e.g., of the order $64^3 \times 128$ sites (space and time discretization), and with no truncation of the gauge fields. These large sizes are required to control finite-volume effects and to perform extrapolations toward the continuum limits¹⁵⁸. In the last decades, the impressive progress in algorithmic development, high-performance optimizations, and the availability of increasingly powerful supercomputer facilities have played a major role in the advancement of MC-based LGT computations. Indeed, this progress has opened the doors to large-scale simulations, that currently represent the standard approach for studying non-perturbative phenomena in quantum field theory.

However, MC techniques are generally based on computations of path integrals, in which observables are calculated as weighted sums. Many physically relevant scenarios, such as finite baryon density regimes or real-time dynamics of quarks, give rise to a change in the sign of the weights and highly oscillating behaviors. Thus, numerical evaluations suffer from the near cancellation of the opposite-sign contributions to the integrals. This is the essence of the infamous sign problem, a long-standing issue of LGT simulations with MC methods¹¹.

Hence, the quest of conceiving, developing, and optimizing alternative approaches that enable simulating these regimes, being the latter at the heart of many open problems related to our understanding of high-energy physics.

As described in the previous subsections, TNs represent a promising complementary method, which found the first applications in simulating non-trivial instances of high-dimensional LGTs on small-to-intermediate lattice sizes. TNs are intrinsically sign-problem-free, enabling the simulation of both static properties at equilibrium, such as low-energy states, and real-time dynamics, even in the presence of finite chemical potentials or non-trivial topological terms. It is worth noting that, in addition to local observables and correlation functions, TNs allow the numerical computations of entanglement properties, such as entanglement entropies and central charges, that could potentially shine new light on LGT phenomena⁴⁴.

Nevertheless, TN simulations of high-dimensional LGTs still represent a challenging problem, especially for large lattice sizes or higher representations of gauge groups needed for analyzing continuum limits. In this framework, further and intensive developments are required to tackle TNs' current problems related to LGTs, such as QCD's non-perturbative effects on lattices of sizes comparable with MC simulations. In this regard, sign-problem-free TN ansätze can also be used in combination with variational MC methods to tackle high-dimensional LGTs with arbitrary gauge groups¹⁵⁹.

In the following, we present a possible roadmap in terms of algorithmic development and optimization strategies that we foresee as crucial for making the TN approach competitive as a complementary method to MC techniques. Therein, the first two subsections, i.e., “Local basis truncation” and “Tailored initial states”, can be approached with existing TN algorithms and a good intuition on LGT problems; then, subsections “Leverage HPC techniques for local optimization” and “Sweeps and HPC parallelization” outline optimization for existing algorithms to leverage HPC systems; finally, we discuss new classes of ansätze to tackle finite temperature problems in subsection “Finite temperature regime.” In some parts, we focus on




Fig. 5 | Graphical representation of methods for handling large local bases. a The pseudo site DMRG (PS-DMRG) approach, in which a single site having large local dimension d is replaced by n_{ps} pseudo sites with a smaller local dimension, e.g., two; **b** the local basis optimization embedded in the MPS ansatz (LBO-MPS), so that an additional layer of tensors connected to the physical legs of the MPS performs the reduction from a large local basis with dimension d into an effective basis with smaller dimension d' ; χ is the bond dimension of the MPS.

TTNs, but the vast majority of the presented concepts and techniques can be straightforwardly applied to other TN ansätze.

Local basis truncation

One of the main issues in simulating LGTs with TN methods lies in the very large local basis dimension d that one needs to handle to represent matter and gauge-field degrees of freedom properly. To some extent, this situation is very similar to some condensed matter models, e.g., the Holstein model, involving lattice fermions coupled with phonons¹⁶⁰, or bosonic systems coupled to optical cavities^{161,162}. Also, in these cases, the local Hilbert space is, in principle, infinite-dimensional, and a truncation to a fixed cutoff is needed for performing TN simulations. For one-dimensional systems of this type, some efficient algorithms based on MPS have been developed in the last years^{163,164}, e.g., the pseudosite DMRG (PS-DMRG) method¹⁶⁵, and the DMRG with local basis optimization (DMRG-LBO)^{166–168}.

The key idea of the PS-DMRG is to replace a single site having large local dimension d with $n_{ps} \approx \log_2(d)$ pseudo sites of local dimension 2, as shown Fig. 5a. Since a large class of optimization algorithms for TNs scale at least quadratically with the local dimension but linearly with the total number of sites, as detailed in subsection “TNs Overview”, one obtains a more efficient and manageable representation according to this procedure.

The price to pay lies in the range of the interactions: short-range interactions in the Hamiltonian are transformed into long-range operators due to the pseudo-site encoding of the local degrees of freedom. As a consequence, PS-DMRG should require a larger bond dimension and a larger number of variational steps to converge, and the benefits offered by the pseudo-sites could progressively fade for increasing values of n_{ps} at fixed bond dimension. PS-DMRG has been applied to one- and two-dimensional QMB systems with d up to $O(100)$ ^{163,165}.

The core of the DMRG-LBO algorithm, instead, lies in a LBO protocol that enables a controlled and efficient truncation of the local Hilbert space¹⁶⁹. For each site \mathbf{x}' of the lattice, the optimized local basis is computed by

starting from its reduced density matrix, i.e.

$$\rho_{\mathbf{x}'} = \text{Tr}_{\mathbf{x} \neq \mathbf{x}'} |\psi\rangle\langle\psi|, \quad (13)$$

where $|\psi\rangle$ is a general state of the whole system and the trace is over all the degrees of freedom which do not involve the site \mathbf{x}' . By performing the diagonalization procedure, the eigenvalues λ_α and the eigenvectors v_α of $\rho_{\mathbf{x}'}$ can be easily determined. The set of values λ_α represents the probabilities associated with the states v_α . If λ_α is small, e.g., below a certain numerical threshold, the related eigenvector v_α can be discarded from the local basis of the site \mathbf{x}' , with a controllable loss of accuracy for the state $|\psi\rangle$. Thus, for reducing the local basis dimension of \mathbf{x}' from d to a smaller value d' , an optimal choice is keeping the d' eigenvectors of $\rho_{\mathbf{x}'}$ with the largest probabilities. Then, the original state $|\psi\rangle$ can be projected on the new basis without losing the relevant physical information. Let us notice that the site \mathbf{x}' can also be a generic unit cell of the system, composed of a certain number of lattice sites.

A crucial point of the LBO procedure is the knowledge of the original state $|\psi\rangle$, generally the ground state of the system, that is not known prior. This issue can be overcome in several ways: by performing exact diagonalization of the system’s Hamiltonian with local dimension d for small lattice sizes, to determine $|\psi\rangle$, and then truncating the basis from d to d' for increasing values of the cutoff d' . From this procedure, we can obtain an optimized local basis, ensuring a controlled approximation of the original ground state. This optimized basis can then be exploited in optimization algorithms for simulating larger lattice sizes, such as DMRG. Another strategy directly incorporates the LBO procedure in the tensor network ansatz, as shown Fig. 5b: by inserting an additional tensor on each physical leg of the MPS, the large local basis with dimension d is transformed into an effective basis with a smaller dimension d' . In this way, the MPS tensors only see the effective basis in the optimization procedures, with a significant reduction of the computational costs described in subsection “TNs Overview”. This method has been used for both static DMRG and time evolution algorithms, such as TEBD, for the study of one-dimensional quantum impurity models and correlated electron-phonon systems^{166,167,170}.

Benchmarks and performance analysis concerning PS-DMRG and DMRG-LBO for QMB systems have been extensively discussed in¹⁶³. Let us mention that some of the main features of these two methods can be combined into more sophisticated algorithms, such as the projected purified DMRG¹⁷¹.

In the context of LGTs, the dimension of the local basis d can easily go beyond values of the order of 10^6 , especially for higher-dimensional non-Abelian models and large representations of the gauge groups, as shown in subsection “Hamiltonian LGTs.” In this scenario, numerical simulations with TNs are practically infeasible without an optimized scheme for truncating the local degrees of freedom. Techniques like PS-DMRG might offer some benefits for small system sizes, such as unveiling the most relevant degrees of freedom in the low-energy states, but it is not easy to scale them up for large sizes due to the long-range interactions induced between the pseudo-sites. Also, the local constraints imposed by Gauss law would become highly non-local when splitting a single site into multiple pseudo-sites representing the matter and link fields.

In principle, LBO-based procedures could instead represent a well-grounded route for addressing the problem. Their use in condensed matter, e.g., bosonic systems, is a rather consolidated approach, whereas their application to TN simulations of LGTs currently is an uncharted but promising territory. The main steps that we foresee as needed and important in this direction are the following:

- (i) Employing exact diagonalization, testing the convergence of LBO procedures for one-dimensional systems, such as the ϕ^4 -theory or the Schwinger model, for which analytical solutions, at least in some regimes of the phase diagram, and numerical results are widely available, also in the limit of no gauge field truncations. In this way, we can obtain valuable information about the scaling of the basis cutoffs

- concerning the final accuracy of the state representation for small system sizes.
- (ii) Performing the same analysis on $(2 + 1)$ -dimensional LGTs, such as QED or $SU(N)$ models, for one unit cell, like a single lattice-plaquette, to systematically study the effects of the magnetic interactions on the local degrees of freedom. Indeed, the QLM representation of LGTs, detailed in subsection “Hamiltonian LGT,” generally exploits the “electric field” basis, in which the electric field terms of the Hamiltonian and Gauss law are diagonal. In this scheme, the magnetic interactions correspond to non-diagonal operators which can increase the number of local electric states to include for an accurate description of the system, especially for small values of the coupling constant g of (3), as highlighted in the numerical analysis in subsection “Scaling of the local basis dimension”.
 - (iii) Exploiting the optimized local bases obtained from exact diagonalization as input of TN simulations for larger sizes, testing the effects on the global ground state accuracy and in computing physically relevant quantities, such as the mass gap¹⁵⁵.
 - (iv) In the same spirit of LBO-MPS, implementing LBO protocols directly in TN ansätze that are more suitable for simulating high-dimensional LGT models, such as TTN. This step could be of great benefit in particular for aTTNs, which encode the area law for the entanglement but are severely limited by large local bases (see subsection “TNs Overview”). By following and combining all these steps, we expect to reduce the effective local basis of LGT models, potentially enabling TN simulations for large representations of Abelian and non-Abelian gauge groups.

It is worth noting that constructing optimal bases for numerical and quantum computation of LGTs is an active area of research. Several approaches that have been recently proposed involve performing canonical transformations of the gauge degrees of freedom before truncation^{129,172–176}. By exploiting a resource-efficient protocol of this type, ref. 129 has shown that the number of local states required to reach a 1% accuracy level when computing the expectation value of the plaquette operators in two-dimensional $(2 + 1)$ -dimensional QED can be reduced by more than 94% compared to the unoptimized truncation. Integrating these approaches into TN algorithms could greatly benefit LGT simulations.

Tailored initial states

In TN algorithms for ground state searching, the optimization procedure generally starts from a random TN initial state $|\psi_{\text{init}}\rangle$, i.e., the tensors in the network are filled with random coefficients at the beginning. This strategy usually guarantees that the probability of overlapping with the true ground state $|\langle\psi_{\text{init}}|\psi_{\text{gs}}\rangle|^2$ does not vanish. To reach a small error in the final energy, this procedure typically requires from 10 to 50 optimization sweeps for LGTs simulations, depending on the specific models, the Hamiltonian parameters, and the lattice size. Since the time for completing a sweep can be very long, especially for large bond dimensions, strategies for reducing the number of needed sweeps could be beneficial for scaling up system sizes. From this perspective, constructing appropriate states to be used as initial guesses can speed up the convergence, similar to the choice of the trial wave function for variational MC simulations. We consider the following options:

- (i) *Physical insight*: Initial states can be constructed by following physical intuition, at least in those regimes in which analytical or partial numerical results are available. For instance, initial guesses can be constructed by starting from the TN ground states numerically obtained for a lower representation to simulate large spin representations of the gauge fields.
- (ii) *Machine learning*: Machine learning-assisted protocols can improve the construction of tailored initial states in the different regimes of the model parameters. For instance, feed-forward neural networks have been proposed as trial wave functions for quantum MC simulations¹⁷⁷, and machine learning techniques have been used to feed TN simulations¹⁷⁸. Also, neural network quantum states (NNQS) have demonstrated their ability to represent the quantum states of many-

body systems with great accuracy¹⁷⁹. Using this approach, neural networks could be trained as variational ansatzes to approximate low-energy states of LGTs. These states might then be used as initial guesses for TN simulations, potentially reducing the computational overhead associated with optimization sweeps. Furthermore, a recent study¹⁸⁰ has explored combining NNQS with TN methods to hybridize their strengths, offering a promising avenue for large-scale simulations.

- (iii) *TN results*: Following the idea of the physical insight, it is also possible to feed neighboring ground states as initial guesses into a ground state search. This option exists especially when scanning a phase diagram and varying parameters in a small increment, such that the overlap between neighboring wave functions is sufficient; this overlap decreases for two points on the opposite sides of a quantum critical point. The same idea can be implemented by preparing an initial guess quenching from an easily accessible ground state to the target parameters. Such an approach is well-known in the case of an adiabatic state preparation¹⁸¹, especially for challenging problems with strong localization. The quench does not necessarily have to be adiabatic or free of numerical errors, but it must only have sufficient overlap with the true ground state. The advantage here is that one quench can generate multiple initial guesses along the quench for different parameters.

Leverage HPC techniques for local optimization

We dedicate the two following subsections to the numerical optimization of the TN algorithms. In this subsection, we give an overview of the topic and discuss possible strategies to improve the optimization. The more technical steps are discussed in the following subsection, “Sweeps and HPC parallelization.”

To scale up TN simulations of LGTs, in particular regarding lattice sizes, another important factor is the number of optimization steps to be carried out. The number of optimization steps scales linearly with the number of sweeps as well as with the system size for MPS, PEPS, and TTN. The choice of the number of sweeps is set such that the algorithm reaches convergence when computing ground or low-energy states. Let us briefly describe the general procedure for ground state searches: we focus on the TTN optimization, however, the main points described here can be applied to other TN structures, such as MPS or PEPS. For a complete technical description of the algorithms and implementation details, see ref. 30.

Consider a generic QMB Hamiltonian H and a generic normalized state $|\psi\rangle$, defined on the same Hilbert space. To numerically determine the ground state of H , the following global minimization problem has to be solved:

$$\min_{|\psi\rangle} \{E(|\psi\rangle)\} = \min_{|\psi\rangle} \langle\psi|H|\psi\rangle. \quad (14)$$

If $|\psi\rangle$ is written in terms of the TTN ansatz introduced in subsection “TNs Overview”, the variational parameters are the coefficients in the TTN. In this case, the global optimization problem of (14) is broken down into a sequence of smaller optimizations, each involving only a minimal subset of tensors in the TTN. The algorithm solves the optimization via an eigenproblem searching for the eigenvector with the lowest eigenvalue. Without any loss of generality, one single tensor at a time is optimized in the simplest case, as shown in Fig. 6a. In detail, the energy is computed by contracting the Hamiltonian between the TTN and its complex conjugate. Then, we start the optimization procedure from a target tensor T , by computing its environment, i.e., the network without the tensors T and T^* , which represents the effective Hamiltonian H_{eff} for the local problem. At this stage, an eigenproblem of H_{eff} is solved, and the tensor T is updated with the newly found ground state. The whole procedure is sequentially iterated for all the tensors in the network, performing an optimization sweep.

For each operation, efficient algorithms from linear algebra are typically used, e.g., the Arnoldi algorithm implemented in the ARPACK library^{182,183}. We recall that the numerical complexity of this procedure for a

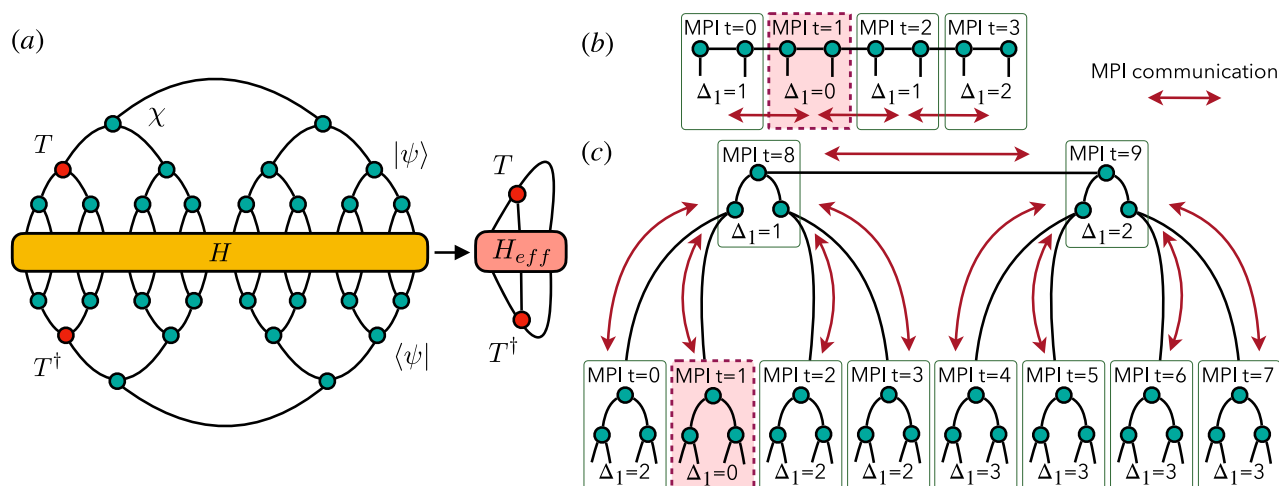


Fig. 6 | Effective operators and parallel tensor networks. **a** Procedure for optimizing a TTN to find the ground state of a QMB system: the energy is computed by contracting the Hamiltonian H (yellow tensor) with the TTN, representing the state $|\psi\rangle$, and its hermitian conjugate, representing $\langle\psi|$. The variational optimization starts from a target tensor T (red tensor), by computing its effective Hamiltonian H_{eff} and then solving the local eigenvalue problem for the latter. The tensor T is then updated with the newly found ground state, and the procedure iterates over all the tensors in the network (sweep). **b** The workload itself consists of optimizing each tensor held by the MPI thread t , which requires effective operators calculated by

other MPI threads. We dub delays_i the number of optimization cycles needed to obtain the information of tensors in the i -th MPI thread in another MPI thread via MPI communication. MPS naturally split into sub-chains, which communicate with one or two neighboring MPI threads to obtain updated effective operators. Delays for updates scale with the distance between two MPI threads along the chain. Each MPI thread can use threading or openMP, e.g., in a hybrid openMP-MPI approach. **c** Similarly, TTNs can be split into sub-trees for each MPI thread, allowing for optimizing the sub-tree without communication with other MPI threads. Delays due to updating scale logarithmically as any distance in a TTN.

single TTN-tensor is $O(d^2\chi^2 + \chi^4)$. Therefore, a single optimization can be time-consuming when χ is very large, e.g., $\chi \approx 1000$, as for simulating high-dimensional LGTs. We point out the established and promising future parallelization schemes for the single tensor optimization:

- (i) *openMP*: an efficient openMP implementation of the contraction between the effective operators with the tensor can speed up simulations. Moreover, the Arnoldi algorithm of ARPACK is optimized for large-scale linear algebra operations and supports intra-node multi-core parallelization based on openMP¹⁸⁴; thus, ARPACK does not become a bottleneck in the openMP implementation. Nonetheless, many simulations remain expensive in computation time even with 64 or more cores available in HPC facilities; therefore, we consider more approaches beyond the well-established openMP path.
- (ii) *Accelerators*: graphics processing units (GPU) and tensor processing units (TPU) offer a path to accelerate linear algebra routines, where both have demonstrated their usefulness¹⁸⁵: GPUs have reported speedups of up to a factor of 10 due to the efficient tensor manipulations^{186–189}; TPU have shown great potential in large-scale simulations of several quantum systems, e.g., drastically reducing the computational time of DMRG calculations with very large bond dimensions from months to hours^{190–193}. TPU are application-specific integrated circuits originally introduced for machine learning; we consider the integration and tuning of TPUs therefore as a step after the successful integration of GPUs. While single GPUs can solve TTN-problems up to a bond dimension of $\chi < 1000$, multi-GPU support is available for libraries; HPC systems typically provide hardware with four GPUs per node.
- (iii) *Multi-node approaches to local tensor optimizations*: both CPU and GPU algorithms can be further scaled by using multiple nodes. The underlying linear algebra routines of the local eigenvalue problem are parallelizable via libraries such as ScaLAPACK or MAGMA. Both libraries provide routines for distributed memory machines¹⁹⁴; moreover, MAGMA supports CPU and GPUs. In this way, the workload of the eigenproblem procedure can be split into several computation nodes. Then, it is important to analyze the performance as a function of the bond dimension χ to test the effectiveness of this approach against the latency of the inter-node communications.

- (iv) *Tuning of parameters and algorithms*: accelerators developed for machine learning applications have excellent support for lower and real precision. Tuning parameters over the different sweeps is beneficial, e.g., increasing the precision towards the end of the sweep. This approach profits from faster single-precision implementations during the first sweeps. Selecting algorithms like random singular value decompositions can also bring benefits¹⁹⁵.

Sweeps and HPC parallelization

So far, we have parallelized single tensor optimizations within a sweep, but the sweep itself was sequential, i.e., serial. Recent works formulated parallel versions of MPS algorithms for ground state search and time evolution, e.g., via the message passing interface (MPI)^{196–199}. The main difference between the serial and parallel algorithms is the effective operators used in the optimization. In the serial version, the effective operators contain the information of the most recent version of all other tensors. This update is delayed in the parallel version, i.e., the tensor that entered the effective operator is not necessarily the one of the current sweep, but can be the version of an earlier sweep.

If the delay becomes an obstacle to convergence, there is the option to modify parameters during the sweeps. Typically, ten to fifty sweeps are necessary to converge on a solution. As the initial state is usually random, MPI can be used especially at the beginning. To ensure convergence, one can consider serial steps at the end; even a gradual reduction of the MPI processes as the sweeps proceed is possible and gradually reduces the delays.

Considering the MPS scenario of a chain in Fig. 6b, we split the chain into equal parts of $(N/n_{\text{MPI-threads}})$ sites. Each part of the chain communicates with its two neighbors apart from the two boundaries. The effective operators take into account the tensors of the same MPI process with zero delay as in the serial case. The tensors of the i -th neighboring MPI process have a delay of i . The worst-case scenario of the delay scales linearly with the number of MPI processes. The delay can be avoided by communicating the effective operators after each update through the chain, which is a quasi-serial step with no more than two MPI processes active at the same time.

The problem becomes more complicated for the TTNs suggested for LGT, but we expect a benefit for the parallelization of a TTN versus an MPS for higher-dimensional systems. Figure 6c shows an example of how each

Table 2 | Extrapolated simulation time

System size	χ	Factor	Extrapolated walltime
64 × 64	450	T_{base}	4.16 days
64 × 64	900	$16 \cdot T_{\text{base}}$	66.6 days
256 × 256	450	$28 \cdot T_{\text{base}}$	116.5 days
256 × 256	900	$448 \cdot T_{\text{base}}$	5.1 years
16 × 16 × 16	450	$4 \cdot T_{\text{base}}$	16.6 days
16 × 16 × 16	900	$64 \cdot T_{\text{base}}$	266 days
64 × 64 × 64	450	$1984 \cdot T_{\text{base}}$	23 years
64 × 64 × 64	900	$31744 \cdot T_{\text{base}}$	362 years

We derive the baseline from a single-tensor optimization of a 64 × 64 quantum Ising model in (15) simulated with 7192 s on a single A100 GPU with 64 GB memory. Further, we assume a single-tensor update, one tensor and one GPU per MPI thread, and 50 sweeps for the baseline.

To extrapolate to larger systems, we assume a scaling with $\mathcal{O}(\chi^4 N^{D-1})$ as well as seven (31) tensors per MPI thread for 256 × 256 (64 × 64 × 64) systems. The empirical scaling is approximately a factor of 2.3 for doubling the system size of the two-dimensional system versus a theoretical value of 2; for doubling the bond dimension, the computation time grows by a factor of 13 versus a theoretical scaling of 16. The factors are obtained from smaller simulations with $\chi = 225$ and for 32 × 32 qubits. The predictions of times are based on a scaling where $d < \chi$.

MPI process gets assigned a sub-tree within the complete TTN. Unlike the MPS, the number of neighboring MPI processes for communication is at least three and increases with $(N/n_{\text{MPI-threads}})$ tensors per thread. For simplicity, we assume equally shaped subtrees for all MPI processes. The delay of the tensor update is now $2 \cdot \log_2(N)$ in the worst-case scenario.

One-dimensional systems with nearest-neighbor interactions thus exhibit a delay of 1 in the worst-case scenario in the MPS, while the delay is up to $2 \cdot \log_2(N)$ for the nearest neighbors in the center of the TTN. Rather, higher-dimensional systems change this aspect, e.g., for an $N \times N$ two-dimensional system mapped to 1D via a zig-zag mapping. The MPS has a worst-case delay of the tensor update of N for the slow index. In contrast, the TTN has the same log behavior and a maximum delay at the center of the TTN as $2 \cdot \log_2(N^2)$. Thus, the worst-case delay is equal for 16 × 16 systems; increasing N further, TTNs exhibit smaller maximum delays during parallel sweeps. Moreover, the TTN is unaffected by the type of mapping used; in contrast, the worst-case delay for the MPS grows to $2N$ for the snake mapping and to at least $N^2/2$ for the Hilbert curve²⁰⁰. Equal arguments hold for 3D systems and delays of N^2 (MPS, zig-zag) versus $2\log_2(N^3)$ (TTN, any mapping).

To get an intuition of what parallelized simulations can solve, we sketch out the specifications for a parallel simulation on the pre-exascale cluster *Leonardo* hosted by *Cineca*. As an example, we focus on a 2D quantum Ising model with a transverse field in the vicinity of the quantum critical point²⁰¹:

$$\hat{H} = \sum_{(i,j)} \hat{\sigma}_i^x \hat{\sigma}_j^x + h \sum_i \hat{\sigma}_i^z, \quad (15)$$

which, despite not resembling the previously described LGT scheme, has been demonstrated to provide the same Hilbert space dimension scaling of a gauge theory²⁸. Then, we choose an MPI approach together with the GPUs. *Leonardo* has 3456 nodes with four GPUs, totaling 13,824 GPUs available for the complete cluster. Bond dimensions on the order of $\chi = 450$ consume 54 GB of memory without effective operators (assuming double complex precision, 40 Lanczos vectors) and allow for solving the eigenproblem on the GPU without temporarily storing data on the CPU. We use the single-tensor per MPI-thread with $\chi = 450$ and initial tensor optimizations^{202,203} as a baseline where we extract a rough empirical estimate with *Leonardo*.

Due to the delay of the tensor in the effective operators, the minimum number of sweeps must be beyond 24. Then, we consider the scaling of the TTN previously introduced and generate Table 2 with an overview of different system sizes and bond dimensions. These results provide a coarse-grained estimate, since plaquette terms, different symmetries, entanglement

generation, and optimization time within later sweeps could further impact the computational time.

Our estimate predicts that a system of 256 × 256 sites takes about two months for bond dimension $\chi = 450$ on 1024 GPU nodes of *Cineca*'s *Leonardo*. Future improvements are likely to bring this simulation time down in comparison to this single-GPU approach, requiring significant data transfer, e.g., the next-generation GPUs beyond the A100, larger GPU memory, or further optimization in data movement. In contrast, a three-dimensional system with large entanglement and many sites requires three to four orders of magnitude improvement, where cluster size and other improvements have to come together to reach this challenge. Furthermore, Table 2 provides an estimate of the boundary for a potential quantum advantage in simulating LGTs with quantum computers or simulators.

Finite temperature regime

To date, TN simulations of high-dimensional LGTs, including dynamical matter, are exploring zero temperature regimes, which are important to understand the low-energy properties of the models. To explore finite temperature phenomena, particularly relevant for open research problems, such as the QCD phase diagram, technical and numerical challenges have to be tackled, e.g., devising and testing efficient TN algorithms for targeting quantum states at equilibrium. As suggested in the next paragraph, we foresee two possible paths toward finite temperature TN states.

Matrix product density operators and locally purified tensor networks (LPTNs) provide already today the option to tackle finite temperature regimes via an imaginary time evolution^{204–206}. Herein, the algorithm starts at the infinite temperature state and starts “cooling” the system via a specified number of time steps and a specified step size to reach a given temperature. In their original formulation, both approaches are one-dimensional chains. The matrix product density operator can be formulated as TTN but faces some challenges in terms of the question of positivity²⁰⁷ or integrating symmetries. In contrast, tree tensor operators (TTO) are the tree-equivalent of an LPTN; they are also a positive loopless representation of density matrices, recently introduced in ref. 208. However, TTOs cannot represent the infinite temperature state necessary for the imaginary time evolution approach. Instead, a possible use of the TTO employed in LGT simulations consists of a variational algorithm to target finite-temperature states or reconstruct open-system quantum dynamics, by efficiently compressing the relevant information. The TTO enables useful measures, e.g., computing the entanglement of formation as already shown for representative one-dimensional models at finite temperature²⁰⁸.

Outlook

The field of TN methods for LGTs has shown great potential in the last decade, during which significant efforts have been devoted to developing numerical algorithms and strategies that can be seen as a complementary approach to MC simulations for high-energy physics. The validity of sign-problem-free TN algorithms has been proven for one-dimensional LGT models for both Abelian and non-Abelian scenarios. Recently, TN methods have also been applied to higher-dimensional LGTs with simple truncated gauge groups and small-to-intermediate lattice sizes. On the one hand, these results prove the effectiveness of TN methods for simulating LGTs even in regimes that are problematic for other numerical methods; on the other hand, they highlight the challenges that one needs to tackle to address state-of-the-art research problems, such as accessing the continuum limits or simulating high-dimensional QCD on large lattices.

In this work, after a general overview of TN methods and their use for LGT simulations, we have described these challenges, starting from the problem of the very large local basis required for complex gauge groups and then discussing several computational bottlenecks in terms of bond dimensions and system sizes.

To mitigate and possibly overcome these problems soon, we have presented a feasible roadmap, in terms of algorithmic developments and

numerical strategies that might have a concrete impact in extending the range of applicability of TN algorithms to current research problems in high-energy physics. We foresee that all the presented steps could potentially have a key role in making the TN approach competitive as a complementary method to MC techniques for simulating high-dimensional LGTs, such as large-scale QCD.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability

The results presented in this study have been obtained using ref. 209 for Exact Diagonalization and ref. 202 for the TNs simulations.

Received: 29 July 2024; Accepted: 5 May 2025;

Published online: 09 August 2025

References

- Kleinert, H. *Gauge Fields in Condensed Matter* (WORLD SCIENTIFIC, 1989).
- Fradkin, E. *Field Theories of Condensed Matter Physics* (Cambridge University Press, 2013).
- Ichinose, I. & Matsui, T. Lattice gauge theory for condensed matter physics: ferromagnetic superconductivity as its example. *Mod. Phys. Lett. B* **28**, 1430012 (2014).
- Peskin, M. E. & Schroeder, D. V. *An Introduction to Quantum Field Theory* (Westview, 1995).
- Schwartz, M. D. *Quantum Field Theory and the Standard Model* (Cambridge University Press, 2013).
- Wilson, K. G. Confinement of quarks. *Phys. Rev. D* **10**, 2445–2459 (1974).
- Kogut, J. B. An introduction to lattice gauge theory and spin systems. *Rev. Mod. Phys.* **51**, 659–713 (1979).
- Creutz, M., Jacobs, L. & Rebbi, C. Monte Carlo computations in lattice gauge theories. *Phys. Rep.* **95**, 201–282 (1983).
- Davoudi, Z. et al. Report of the snowmass 2021 topical group on lattice gauge theory. *arXiv e-prints* arXiv:2209.10758 <https://doi.org/10.48550/arXiv.2209.10758> (2022).
- Aoki, Y. et al. Flag review 2021. *Eur. Phys. J. C* **82**, 869 (2022).
- Nagata, K. Finite-density lattice QCD and sign problem: current status and open problems. *Prog. Part. Nucl. Phys.* **127**, 103991 (2022).
- Bañuls, M. C., Cichy, K., Cirac, J. I., Jansen, K. & Kühn, S. Tensor networks and their use for lattice gauge theories. In: *Proc. 36th Annual International Symposium on Lattice Field Theory — PoS(LATTICE2018)*, vol. 334, 022 (Sissa Medialab, 2019).
- Bañuls, M. C. & Cichy, K. Review on novel methods for lattice gauge theories. *Rep. Prog. Phys.* **83**, 024401 (2020).
- Bañuls, M. C. et al. Simulating lattice gauge theories within quantum technologies. *Eur. Phys. J. D* **74**, 165 (2020).
- Fannes, M., Nachtergaele, B. & Werner, R. F. Finitely correlated states on quantum spin chains. *Commun. Math. Phys.* **144**, 443–490 (1992).
- Östlund, S. & Rommer, S. Thermodynamic limit of density matrix renormalization. *Phys. Rev. Lett.* **75**, 3537–3540 (1995).
- Rommer, S. & Östlund, S. Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group. *Phys. Rev. B* **55**, 2164–2181 (1997).
- Schollwöck, U. The density-matrix renormalization group. *Rev. Mod. Phys.* **77**, 259–315 (2005).
- Montangero, S. *Introduction to Tensor Network Methods: Numerical Simulations of Low-Dimensional Many-Body Quantum Systems* (Springer, 2018).
- Orús, R. Tensor networks for complex quantum systems. *Nat. Rev. Phys.* **1**, 538–550 (2019).
- Bañuls, M. C. Tensor network algorithms: a route map. *Annu. Rev. Condens. Matter Phys.* **14**, null (2023).
- Cirac, J. I., Pérez-García, D., Schuch, N. & Verstraete, F. Matrix product states and projected entangled pair states: concepts, symmetries, theorems. *Rev. Mod. Phys.* **93**, 045003 (2021).
- Evenbly, G. & Vidal, G. Tensor network renormalization. *Phys. Rev. Lett.* **115**, 180405 (2015).
- Levin, M. & Nave, C. P. Tensor renormalization group approach to two-dimensional classical lattice models. *Phys. Rev. Lett.* **99**, 120601 (2007).
- Evenbly, G. Algorithms for tensor network renormalization. *Phys. Rev. B* **95**, 045117 (2017).
- Meurice, Y., Sakai, R. & Unmuth-Yockey, J. Tensor lattice field theory for renormalization and quantum computing. *Rev. Mod. Phys.* **94**, 025005 (2022).
- Byrnes, T. M. R., Sriganesh, P., Bursill, R. J. & Hamer, C. J. Density matrix renormalization group approach to the massive Schwinger model. *Phys. Rev. D* **66**, 013002 (2002).
- Silvi, P., Rico, E., Calarco, T. & Montangero, S. Lattice gauge tensor networks. *N. J. Phys.* **16**, 103015 (2014).
- Pichler, T., Dalmonte, M., Rico, E., Zoller, P. & Montangero, S. Real-time dynamics in U(1) lattice gauge theories with tensor networks. *Phys. Rev. X* **6**, 011023 (2016).
- Silvi, P. et al. The tensor networks anthology: simulation techniques for many-body quantum lattice systems. *SciPost Phys. Lecture Notes* 008 (2019).
- Funcke, L., Jansen, K. & Kühn, S. Topological vacuum structure of the Schwinger model with matrix product states. *Phys. Rev. D* **101**, 054507 (2020).
- Buyens, B., Haegeman, J., Van Acoleyen, K., Verschelde, H. & Verstraete, F. Matrix product states for gauge field theories. *Phys. Rev. Lett.* **113**, 091601 (2014).
- Rico, E., Pichler, T., Dalmonte, M., Zoller, P. & Montangero, S. Tensor networks for lattice gauge theories and atomic quantum simulation. *Phys. Rev. Lett.* **112**, 201601 (2014).
- Haegeman, J., Van Acoleyen, K., Schuch, N., Cirac, J. I. & Verstraete, F. Gauging quantum states: from global to local symmetries in many-body systems. *Phys. Rev. X* **5**, 011024 (2015).
- Silvi, P., Rico, E., Dalmonte, M., Tschirsich, F. & Montangero, S. Finite-density phase diagram of a $(1+1)-d$ non-abelian lattice gauge theory with tensor networks. *Quantum* **1**, 9 (2017).
- Buyens, B., Montangero, S., Haegeman, J., Verstraete, F. & Van Acoleyen, K. Finite-representation approximation of lattice gauge theories at the continuum limit with tensor networks. *Phys. Rev. D* **95**, 094509 (2017).
- Bañuls, M. C., Cichy, K., Cirac, J. I., Jansen, K. & Kühn, S. Density induced phase transitions in the schwinger model: a study with matrix product states. *Phys. Rev. Lett.* **118**, 071601 (2017).
- Ercolessi, E., Facchi, P., Magnifico, G., Pascazio, S. & Pepe, F. V. Phase transitions in Z_n gauge models: towards quantum simulations of the Schwinger-Weyl QED. *Phys. Rev. D* **98**, 074503 (2018).
- Magnifico, G. et al. Symmetry-protected topological phases in lattice gauge theories: topological qed₂. *Phys. Rev. D* **99**, 014503 (2019).
- Magnifico, G. et al. Z_N gauge theories coupled to topological fermions: QED2 with a quantum mechanical θ angle. *Phys. Rev. B* **100**, 115152 (2019).
- Sala, P. et al. Gaussian states for the variational study of $(1+1)$ -dimensional lattice gauge models. In *The 36th Annual International Symposium on Lattice Field Theory. 22–28 July*, 230 (Lattice, 2018).
- Magnifico, G. et al. Real time dynamics and confinement in the Z_n Schwinger-Weyl lattice model for $1+1$ QED. *Quantum* **4**, 281 (2020).

43. Bañuls, M. C., Cichy, K., Cirac, J. I., Jansen, K. & Kühn, S. Efficient basis formulation for $(1+1)$ -dimensional $su(2)$ lattice gauge theory: Spectral calculations with matrix product states. *Phys. Rev. X* **7**, 041046 (2017).
44. Rigobello, M., Notarnicola, S., Magnifico, G. & Montangero, S. Entanglement generation in $(1+1)$ D QED scattering processes. *Phys. Rev. D* **104**, 114501 (2021).
45. Funcke, L., Jansen, K. & Kühn, S. Exploring the CP-violating Dashen phase in the Schwinger model with tensor networks. *Phys. Rev. D* **108**, 014504 (2023).
46. Angelides, T., Funcke, L., Jansen, K. & Kühn, S. Computing the mass shift of Wilson and staggered fermions in the lattice Schwinger model with matrix product states. *Phys. Rev. D* **108**, 014516 (2023).
47. Chanda, T., Dalmonte, M., Lewenstein, M., Zakrzewski, J. & Tagliacozzo, L. Spectral properties of critical $1+1$ d abelian-higgs model. *Phys. Rev. B* **109**, 045103 (2023).
48. Schmoll, P., Naumann, J., Nietner, A., Eisert, J. & Sotiriadis, S. Hamiltonian truncation tensor networks for quantum field theories. <https://doi.org/10.48550/arXiv.2312.12506> (2023).
49. Hayata, T., Hidaka, Y. & Nishimura, K. Dense QCD₂ with matrix product states. *J. High Energy Phys.* **2024**, 106 (2023).
50. Florio, A., Weichselbaum, A., Valgushev, S. & Pisarski, R. D. Mass gaps of a \mathbb{Z}_3 gauge theory with three fermion flavors in $1+1$ dimensions. *Phys. Rev. D* **110**, 045013 (2023).
51. Osborne, J., McCulloch, I. P. & Halimeh, J. C. Probing confinement through dynamical quantum phase transitions: from quantum spin models to lattice gauge theories. <https://doi.org/10.48550/arXiv.2310.12210> (2023).
52. Kebrič, M., Halimeh, J. C., Schollwöck, U. & Grusdt, F. Confinement in $1+1$ d \mathbb{Z}_2 lattice gauge theories at finite temperature. <https://doi.org/10.48550/arXiv.2308.08592> (2024).
53. Belyansky, R. et al. High-energy collision of quarks and mesons in the Schwinger model: from tensor networks to circuit QED. *Phys. Rev. Lett.* **132**, 091903 (2024).
54. Papaefstathiou, I., Knolle, J. & Bañuls, M. C. Real-time scattering in the lattice Schwinger model. *Phys. Rev. D* **111**, 014504 (2024).
55. Calajò, G. et al. Quantum many-body scarring in a non-abelian lattice gauge theory. *Phys. Rev. Res.* **7**, 013322 (2024).
56. Felser, T., Silvi, P., Collura, M. & Montangero, S. Two-dimensional quantum-link lattice quantum electrodynamics at finite density. *Phys. Rev. X* **10**, 041040 (2020).
57. Emonts, P. et al. Finding the ground state of a lattice gauge theory with fermionic tensor networks: a $2+1$ d \mathbb{Z}_2 demonstration. *Phys. Rev. D* **107**, 014505 (2023).
58. Magnifico, G., Felser, T., Silvi, P. & Montangero, S. Lattice quantum electrodynamics in $(3+1)$ -dimensions at finite density with tensor networks. *Nat. Commun.* **12**, 3600 (2021).
59. Knaute, J., Feuerstein, M. & Zohar, E. Entanglement and confinement in lattice gauge theory tensor networks. *J. High Energy Phys.* **2024**, 174 (2024).
60. Pradhan, S., Maroncelli, A. & Ercolessi, E. Discrete abelian lattice gauge theories on a ladder and their dualities with quantum clock models. *Phys. Rev. B* **109**, 064410 (2024).
61. Su, G.-X., Osborne, J. & Halimeh, J. C. A cold-atom particle collider **5**, 040310 (2024).
62. Kuramashi, Y. & Yoshimura, Y. Three-dimensional finite temperature z_2 gauge theory with tensor network scheme. *J. High. Energy Phys.* **2019**, 23 (2019).
63. Robaina, D., Bañuls, M. C. & Cirac, J. I. Simulating $(2+1)$ d z_3 lattice gauge theory with an infinite projected entangled-pair state. *Phys. Rev. Lett.* **126**, 050401 (2021).
64. Yosprakob, A. & Okunishi, K. Tensor renormalization group study of the three-dimensional $su(2)$ and $su(3)$ gauge theories with the reduced tensor network formulation. *arXiv preprint arXiv:2406.16763* (2024).
65. Cataldi, G., Magnifico, G., Silvi, P. & Montangero, S. Simulating $(2+1)$ d $su(2)$ yang-mills lattice gauge theory at finite density with tensor networks. *Phys. Rev. Res.* **6**, 033057 (2024).
66. Villalonga, B. et al. A flexible high-performance simulator for verifying and benchmarking quantum circuits implemented on real hardware. *npj Quant. Inf.* **5**, 86 (2019).
67. Mathis, S. V., Mazzola, G. & Tavernelli, I. Toward scalable simulations of lattice gauge theories on quantum computers. *Phys. Rev. D* **102**, 094501 (2020).
68. Zhou, Y., Stoudenmire, E. M. & Waintal, X. What limits the simulation of quantum computers? *Phys. Rev. X* **10**, 041038 (2020).
69. Huang, C. et al. Efficient parallelization of tensor network contraction for simulating quantum computation. *Nat. Comput. Sci.* **1**, 578–587 (2021).
70. Haghshenas, R., Gray, J., Potter, A. C. & Chan, G. K.-L. Variational power of quantum circuit tensor networks. *Phys. Rev. X* **12**, 011047 (2022).
71. Zohar, E. Quantum simulation of lattice gauge theories in more than one space dimension—requirements, challenges and methods. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **380**, 20210069 (2021).
72. Catterall, S. et al. Report of the snowmass 2021 theory frontier topical group on quantum information science. Tech. Rep. FERMILAB-FN-1199-T; arXiv:2209.14839, Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States); Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA (United States) <https://doi.org/10.48550/arXiv.2209.14839> (2022).
73. Pomarico, D. et al. Dynamical quantum phase transitions of the Schwinger model: real-time dynamics on IBM quantum. *Entropy* **25**, 608 (2023).
74. Funcke, L., Hartung, T., Jansen, K. & Kühn, S. Review on Quantum Computing for Lattice Field Theory. In *Proc. 39th International Symposium on Lattice Field Theory — PoS(LATTICE2022)*, vol. 430, 228 (SISSA Medialab, 2023).
75. Mariani, A., Pradhan, S. & Ercolessi, E. Hamiltonians and gauge-invariant Hilbert space for lattice Yang-Mills-like theories with finite gauge group. *Phys. Rev. D* **107**, 114513 (2023).
76. Zache, T. V., González-Cuadra, D. & Zoller, P. Fermion-qudit quantum processors for simulating lattice gauge theories with matter. *Quantum* **7**, 1140 (2023).
77. Eisert, J., Cramer, M. & Plenio, M. B. Colloquium: area laws for the entanglement entropy. *Rev. Mod. Phys.* **82**, 277–306 (2010).
78. Calabrese, P. & Cardy, J. Entanglement entropy and quantum field theory. *J. Stat. Mech. Theory Exp.* **2004**, P06002 (2004).
79. Hastings, M. B. An area law for one-dimensional quantum systems. *J. Stat. Mech. Theory Exp.* **2007**, P08024 (2007).
80. Kuwahara, T. & Saito, K. Area law of noncritical ground states in 1d long-range interacting systems. *Nat. Commun.* **11**, 4478 (2020).
81. Cho, J. Realistic area-law bound on entanglement from exponentially decaying correlations. *Phys. Rev. X* **8**, 031009 (2018).
82. Schuch, N., Wolf, M. M., Verstraete, F. & Cirac, J. I. Entropy scaling and simulability by matrix product states. *Phys. Rev. Lett.* **100**, 030504 (2008).
83. Wolf, M. M., Verstraete, F., Hastings, M. B. & Cirac, J. I. Area laws in quantum systems: mutual information and correlations. *Phys. Rev. Lett.* **100**, 070502 (2008).
84. Brandão, F. G. S. L. & Horodecki, M. Exponential decay of correlations implies area law. *Commun. Math. Phys.* **333**, 761–798 (2015).
85. Masanes, L. Area law for the entropy of low-energy states. *Phys. Rev. A* **80**, 052104 (2009).
86. Hamza, E., Michalakakis, S., Nachtergaele, B. & Sims, R. Approximating the ground state of gapped quantum spin systems. *J. Math. Phys.* **50**, 095213 (2009).
87. Kastoryano, M. J., Lucia, A. & Perez-Garcia, D. Locality at the boundary implies gap in the bulk for 2D PEPS. *Commun. Math. Phys.* **366**, 895–926 (2019).

88. Cirac, J. I., Garre-Rubio, J. & Pérez-García, D. Mathematical open problems in projected entangled pair states. *Rev. Matemática Complut.* **32**, 579–599 (2019).
89. Perez-Garcia, D., Verstraete, F., Wolf, M. M. & Cirac, J. I. Matrix product state representations. *Quant. Info Comput.* **7**, 401–430 (2007).
90. Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **326**, 96–192 (2011).
91. Chan, G. K.-L., Keselman, A., Nakatani, N., Li, Z. & White, S. R. Matrix product operators, matrix product states, and ab initio density matrix renormalization group algorithms. *J. Chem. Phys.* **145**, 014102 (2016).
92. Gleis, A., Li, J.-W. & von Delft, J. Controlled bond expansion for density matrix renormalization group ground state search at single-site costs. *Phys. Rev. Lett.* **130**, 246402 (2023).
93. Hubig, C., McCulloch, I. P., Schollwöck, U. & Wolf, F. A. Strictly single-site DMRG algorithm with subspace expansion. *Phys. Rev. B* **91**, 155115 (2015).
94. White, S. R. Density matrix renormalization group algorithms with a single center site. *Phys. Rev. B* **72**, 180403 (2005).
95. Scheb, M. & Noack, R. M. Finite projected entangled pair states for the Hubbard model. *Phys. Rev. B* **107**, 165112 (2023).
96. Schuch, N., Wolf, M. M., Verstraete, F. & Cirac, J. I. Computational complexity of projected entangled pair states. *Phys. Rev. Lett.* **98**, 140506 (2007).
97. Eisert, J. Entanglement and tensor network states <https://doi.org/10.48550/arXiv.1308.3318> (2013).
98. Vanderstraeten, L. et al. Variational methods for contracting projected entangled-pair states. *Phys. Rev. B* **105**, 195140 (2022).
99. Shi, Y.-Y., Duan, L.-M. & Vidal, G. Classical simulation of quantum many-body systems with a tree tensor network. *Phys. Rev. A* **74**, 022320 (2006).
100. Qian, X. & Qin, M. From tree tensor network to multiscale entanglement renormalization ansatz. *Phys. Rev. B* **105**, 205102 (2022).
101. Ferris, A. J. Area law and real-space renormalization. *Phys. Rev. B* **87**, 125139 (2013).
102. Felser, T., Notarnicola, S. & Montangero, S. Efficient tensor network ansatz for high-dimensional quantum many-body problems. *Phys. Rev. Lett.* **126**, 170603 (2021).
103. Jaschke, D., Wall, M. L. & Carr, L. D. Open source matrix product states: opening ways to simulate entangled many-body quantum systems in one dimension. *Comput. Phys. Commun.* **225**, 59–91 (2018).
104. Vidal, G. Efficient simulation of one-dimensional quantum many-body systems. *Phys. Rev. Lett.* **93**, 040502 (2004).
105. Haegeman, J. et al. Time-dependent variational principle for quantum lattices. *Phys. Rev. Lett.* **107**, 070601 (2011).
106. Bauernfeind, D. & Aichhorn, M. Time dependent variational principle for tree tensor networks. *SciPost Phys.* **8**, 024 (2020).
107. Kohn, L. et al. Superfluid-to-mott transition in a bose-hubbard ring: persistent currents and defect formation. *Phys. Rev. A* **101**, 023617 (2020).
108. Yang, M. & White, S. R. Time-dependent variational principle with ancillary Krylov subspace. *Phys. Rev. B* **102**, 094315 (2020).
109. Paeckel, S. et al. Time-evolution methods for matrix-product states. *Ann. Phys.* **411**, 167998 (2019).
110. White, C. D., Zaletel, M., Mong, R. S. K. & Refael, G. Quantum dynamics of thermalizing systems. *Phys. Rev. B* **97**, 035127 (2018).
111. Surace, J., Piani, M. & Tagliacozzo, L. Simulating the out-of-equilibrium dynamics of local observables by trading entanglement for mixture. *Phys. Rev. B* **99**, 235115 (2019).
112. Akiyama, S., Maurice, Y. & Sakai, R. Tensor renormalization group for fermions. *J. Phys. Condens. Matter* **36**, 343002 (2024).
113. Zache, T. V. et al. Quantum simulation of lattice gauge theories using Wilson fermions. *Quant. Sci. Technol.* **3**, 034010 (2018).
114. Susskind, L. Lattice fermions. *Phys. Rev. D* **16**, 3031–3039 (1977).
115. Kogut, J. & Susskind, L. Hamiltonian formulation of Wilson's lattice gauge theories. *Phys. Rev. D* **11**, 395–408 (1975).
116. Zohar, E. & Burrello, M. Formulation of lattice gauge theories for quantum simulations. *Phys. Rev. D* **91**, 054506 (2015).
117. Sala, P. et al. Variational study of $u(1)$ and $su(2)$ lattice gauge theories with Gaussian states in $1 + 1$ dimensions. *Phys. Rev. D* **98**, 034505 (2018).
118. Bender, J. & Zohar, E. Gauge redundancy-free formulation of compact qed with dynamical matter for quantum and classical computations. *Phys. Rev. D* **102**, 114517 (2020).
119. Ligtnerink, N., Walet, N. & Bishop, R. Toward a many-body treatment of Hamiltonian lattice $su(n)$ gauge theory. *Ann. Phys.* **284**, 215–262 (2000).
120. Horn, D. Finite matrix models with continuous local gauge invariance. *Phys. Lett. B* **100**, 149–151 (1981).
121. Orland, P. & Rohrlich, D. Lattice gauge magnets: local isospin from spin. *Nucl. Phys. B* **338**, 647–672 (1990).
122. Chandrasekharan, S. & Wiese, U. J. Quantum link models: a discrete approach to gauge theories. *Nucl. Phys. B* **492**, 455–471 (1997).
123. Brower, R., Chandrasekharan, S. & Wiese, U.-J. QCD as a quantum link model. *Phys. Rev. D* **60**, 094502 (1999).
124. Byrnes, T. & Yamamoto, Y. Simulating lattice gauge theories on a quantum computer. *Phys. Rev. A* **73**, 022328 (2006).
125. Davoudi, Z. et al. Towards analog quantum simulations of lattice gauge theories with trapped ions. *Phys. Rev. Res.* **2**, 023015 (2020).
126. Mazzola, G., Mathis, S. V., Mazzola, G. & Tavernelli, I. Gauge-invariant quantum circuits for $U(1)$ and Yang-Mills lattice gauge theories. *Phys. Rev. Res.* **3**, 043209 (2021).
127. Kan, A. et al. Investigating a $(3 + 1)D$ topological θ -term in the Hamiltonian formulation of lattice gauge theories for quantum and classical simulations. *Phys. Rev. D* **104**, 034504 (2021).
128. Bauer, C. W., Davoudi, Z., Klcio, N. & Savage, M. J. Quantum simulation of fundamental particles and forces. *Nat. Rev. Phys.* **5**, 420–432 (2023).
129. Haase, J. F. et al. A resource efficient approach for quantum and classical simulations of gauge theories in particle physics. *Quantum* **5**, 393 (2021).
130. Hackett, D. C. et al. Digitizing gauge fields: lattice Monte Carlo results for future quantum computers. *Phys. Rev. A* **99**, 062341 (2019).
131. Zache, T. V., González-Cuadra, D. & Zoller, P. Quantum and classical spin network algorithms for q -deformed Kogut-Susskind gauge theories. *Phys. Rev. Lett.* **131**, 171902 (2023).
132. Rigobello, M., Magnifico, G., Silvi, P. & Montangero, S. Hadrons in $(1 + 1)D$ Hamiltonian hardcore lattice QCD. <https://doi.org/10.48550/arXiv.2308.04488> (2023).
133. Strocchi, F. *An Introduction to Non-Perturbative Foundations of Quantum Field Theory*. International series of monographs on Physics (Oxford University Press, 2013).
134. Zohar, E. & Cirac, J. I. Removing staggered fermionic matter in $U(N)$ and $SU(N)$ lattice gauge theories. *Phys. Rev. D* **99**, 114511 (2019).
135. Zohar, E. & Cirac, J. I. Eliminating fermionic matter fields in lattice gauge theories. *Phys. Rev. B* **98**, 075119 (2018).
136. Ballarin, M. et al. Digital quantum simulation of lattice fermion theories with local encoding. *Quantum* **8**, 1460 (2024).
137. Verstraete, F., Murg, V. & Cirac, J. I. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.* **57**, 143–224 (2008).
138. Ciavarella, A., Klcio, N. & Savage, M. J. Trailhead for quantum simulation of $su(3)$ Yang-Mills lattice gauge theory in the local multiplet basis. *Phys. Rev. D* **103**, 094501 (2021).

139. Davoudi, Z., Raychowdhury, I. & Shaw, A. Search for efficient formulations for Hamiltonian simulation of non-abelian lattice gauge theories. *Phys. Rev. D* **104**, 074505 (2021).
140. Tong, Y., Albert, V. V., McClean, J. R., Preskill, J. & Su, Y. Provably accurate simulation of gauge theories and bosonic systems. *Quantum* **6**, 816 (2022).
141. Kühn, S., Cirac, J. I. & Bañuls, M.-C. Quantum simulation of the Schwinger model: a study of feasibility. *Phys. Rev. A* **90**, 042305 (2014).
142. Bañuls, M. C., Cichy, K., Cirac, J. I., Jansen, K. & Kühn, S. Efficient basis formulation for $(1 + 1)$ -dimensional $su(2)$ lattice gauge theory: Spectral calculations with matrix product states. *Phys. Rev. X* **7**, 041046 (2017).
143. Zache, T. V., Van Damme, M., Halimeh, J. C., Hauke, P. & Banerjee, D. Toward the continuum limit of a $(1 + 1)$ D quantum link Schwinger model. *Phys. Rev. D* **106**, L091502 (2022).
144. Fiore, R. et al. QED₃ on a space-time lattice: A comparison between compact and noncompact formulation. In *Proceedings of XXIIIrd International Symposium on Lattice Field Theory — PoS(LAT2005)*, vol. 20, 243 (SISSA Medialab, 2005).
145. Raviv, O., Shamir, Y. & Svetitsky, B. Nonperturbative beta function in three-dimensional electrodynamics. *Phys. Rev. D* **90**, 014512 (2014).
146. Svetitsky, B., Raviv, O. & Shamir, Y. Beta function of three-dimensional QED. In *Proceedings of The 32nd International Symposium on Lattice Field Theory — PoS(LATTICE2014)*, vol. 214, 051 (SISSA Medialab, 2015).
147. Xu, X. Y. et al. Monte Carlo study of lattice compact quantum electrodynamics with fermionic matter: the parent state of quantum phases. *Phys. Rev. X* **9**, 021022 (2019).
148. Creutz, M. Lattice gauge theory and Monte Carlo methods. Tech. Rep. BNL-42086 (Brookhaven National Laboratory, 1988).
149. Creutz, M. Lattice gauge theories and Monte Carlo algorithms. *Nucl. Phys. B Proc. Suppl.* **10**, 1–22 (1989).
150. Loan, M., Brunner, M., Sloggett, C. & Hamer, C. Path integral Monte Carlo approach to the $U(1)$ lattice gauge theory in $2+1$ dimensions. *Phys. Rev. D* **68**, 034504 (2003).
151. Rothe, H. *Lattice Gauge Theories: An Introduction* (Fourth Edition) (World Scientific Publishing Company, 2012).
152. Funcke, L. et al. Hamiltonian limit of lattice QED in $2+1$ dimensions. In *Proceedings of The 39th International Symposium on Lattice Field Theory — PoS(LATTICE2022)*, vol. 430, 292 (SISSA Medialab, 2023).
153. Strouthos, C. & Kogut, J. B. The phases of non-compact QED(3) In *Proc. XXV International Symposium on Lattice Field Theory* (2008).
154. Bender, J. *Quantum and Classical Methods for Lattice Gauge Theories in Higher Dimensions*. Ph.D. thesis, (Technische Universität, 2023).
155. Clemente, G., Crippa, A. & Jansen, K. Strategies for the determination of the running coupling of $(2+1)$ -dimensional QED with quantum computing. *Phys. Rev. D* **106**, 114511 (2022).
156. Hernández, M. P. *Lattice field theory fundamentals*. (Oxford University Press, 2011).
157. Kang, C., Soley, M. B., Crane, E., Girvin, S. M. & Wiebe, N. Leveraging Hamiltonian simulation techniques to compile operations on bosonic devices. *J. Phys. A Math. Theor.* **58**, 175301 (2023).
158. Workman, R. L. et al. Review of particle physics. *PTEP* **2022**, 083C01 (2022).
159. Kelman, A. et al. Gauged Gaussian peps – a high dimensional tensor network formulation for lattice gauge theories. *Phys. Rev. D* **110**, 054511 (2024).
160. Holstein, T. Studies of polaron motion: Part I. The molecular-crystal model. *Ann. Phys.* **8**, 325–342 (1959).
161. Bloch, I. Ultracold quantum gases in optical lattices. *Nat. Phys.* **1**, 23–30 (2005).
162. Schlawin, F., Kennes, D. M. & Sentef, M. A. Cavity quantum materials. *Appl. Phys. Rev.* **9**, 011312 (2022).
163. Stolpp, J. et al. Comparative study of state-of-the-art matrix-product-state methods for lattice models with large local hilbert spaces without $u(1)$ symmetry. *Comput. Phys. Commun.* **269**, 108106 (2021).
164. Kamar, N. A. & Maghrebi, M. Splitting the local Hilbert space: MPS-based approach to large local dimensions. *Phys. Rev. B* **110**, 075114 (2023).
165. Jeckelmann, E. & White, S. R. Density-matrix renormalization-group study of the polaron problem in the holstein model. *Phys. Rev. B* **57**, 6376–6385 (1998).
166. Guo, C., Weichselbaum, A., von Delft, J. & Vojta, M. Critical and strong-coupling phases in one- and two-bath spin-boson models. *Phys. Rev. Lett.* **108**, 160401 (2012).
167. Brockt, C., Dorfner, F., Vidmar, L., Heidrich-Meisner, F. & Jeckelmann, E. Matrix-product-state method with a dynamical local basis optimization for bosonic systems out of equilibrium. *Phys. Rev. B* **92**, 241106 (2015).
168. Stolpp, J., Herbrych, J., Dorfner, F., Dagotto, E. & Heidrich-Meisner, F. Charge-density-wave melting in the one-dimensional Holstein model. *Phys. Rev. B* **101**, 035134 (2020).
169. Zhang, C., Jeckelmann, E. & White, S. R. Density matrix approach to local Hilbert space reduction. *Phys. Rev. Lett.* **80**, 2661–2664 (1998).
170. Schröder, F. A. Y. N. & Chin, A. W. Simulating open quantum dynamics with time-dependent variational matrix product states: towards microscopic correlation of environment dynamics and reduced system evolution. *Phys. Rev. B* **93**, 075105 (2016).
171. Köhler, T., Stolpp, J. & Paeckel, S. Efficient and flexible approach to simulate low-dimensional quantum lattice models with large local Hilbert spaces. *SciPost Phys.* **10**, 058 (2021).
172. Mathur, M. & Rathor, A. Exact duality and local dynamics in $su(n)$ lattice gauge theory. *Phys. Rev. D* **107**, 074504 (2023).
173. Bauer, C. W., D’Andrea, I., Freytsis, M. & Grabowska, D. M. A new basis for Hamiltonian $SU(2)$ simulations. *Phys. Rev. D* **109**, 074501 (2023).
174. Kaplan, D. B. & Stryker, J. R. Gauss’s law, duality, and the hamiltonian formulation of $u(1)$ lattice gauge theory. *Phys. Rev. D* **102**, 094515 (2020).
175. Raychowdhury, I. & Stryker, J. R. Loop, string, and hadron dynamics in $su(2)$ hamiltonian lattice gauge theories. *Phys. Rev. D* **101**, 114502 (2020).
176. Kadam, S. V., Raychowdhury, I. & Stryker, J. R. Loop-string-hadron formulation of an $su(3)$ gauge theory with dynamical quarks. *Phys. Rev. D* **107**, 094513 (2023).
177. Kessler, J., Calcavecchia, F. & Kühne, T. D. Artificial neural networks as trial wave functions for quantum Monte Carlo. *Adv. Theory Simul.* **4**, 2000269 (2021).
178. Schröder, F. A. Y. N., Turban, D. H. P., Musser, A. J., Hine, N. D. M. & Chin, A. W. Tensor network simulation of multi-environmental open quantum dynamics via machine learning and entanglement renormalisation. *Nat. Commun.* **10**, 1062 (2019).
179. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
180. Chen, Z., Newhouse, L., Chen, E., Luo, D. & Soljačić, M. Antr: bridging autoregressive neural networks and tensor networks for quantum many-body simulation. In *Proc. 37th International Conference on Neural Information Processing Systems*, vol. 22 450–476 (NIPS, 2024).
181. Ljubotina, M., Roos, B., Abanin, D. A. & Serbyn, M. Optimal steering of matrix product states and quantum many-body scars. *PRX Quant.* **3**, 030343 (2022).
182. Arnoldi, W. E. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Q Appl. Math.* **9**, 17–29 (1951).

183. Lehoucq, R. B., Sorensen, D. C. & Yang, C. *ARPACK: solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*. Available from netlib@ornl.gov (1997).
184. Dagum, L. & Menon, R. OpenMP: an industry standard api for shared-memory programming. *Comput. Sci. Eng.* **5**, 46–55 (1998).
185. Jaschke, D., Ballarin, M., Reinić, N., Pavešić, L. & Montangero, S. Benchmarking quantum red tea on CPUS, GPUS, and TPUS <https://arxiv.org/abs/2409.03818> (2024).
186. Shi, Y., Niranjan, U. N., Anandkumar, A. & Cecka, C. Tensor Contractions with Extended BLAS Kernels on CPU and GPU. In *2016 IEEE 23rd International Conference on High Performance Computing*, 193–202 (IEEE, 2016).
187. Abdelfattah, A. et al. High-performance tensor contractions for GPUs. *Procedia Comput. Sci.* **80**, 108–118 (2016).
188. Vincent, T. et al. Jet: fast quantum circuit simulations with parallel task-based tensor-network contraction. *Quantum* **6**, 709 (2022).
189. Pan, F. & Zhang, P. Simulation of quantum circuits using the big-batch tensor network method. *Phys. Rev. Lett.* **128**, 030501 (2022).
190. Jouppi, N. P. et al. In-datacenter performance analysis of a tensor processing unit In: *Proceedings of the 44th Annual International Symposium on Computer Architecture* 1–12 <https://doi.org/10.48550/arXiv.1704.04760> (ACM, 2017).
191. Hauru, M. et al. Simulation of quantum physics with tensor processing units: brute-force computation of ground states and time evolution. *PRX Quant.* **3**, 020331 (2021).
192. Morningstar, A. et al. Simulation of quantum many-body dynamics with tensor processing units: floquet prethermalization. *PRX Quant.* **3**, 020331 (2022).
193. Ganahl, M. et al. Density matrix renormalization group with tensor processing units. *PRX Quant.* **4**, 010317 (2023).
194. Blackford, L. S. et al. *ScaLAPACK Users' Guide*. Software, environments, and tools (Society for Industrial and Applied Mathematics, 1997).
195. Lu, Y., Ino, F. & Matsushita, Y. High-performance out-of-core block randomized singular value decomposition on GPU <https://doi.org/10.48550/arXiv.1706.07191> (2017).
196. Gabriel, E. et al. Open MPI: goals, concept, and design of a next generation MPI implementation. In *Proc. 11th European PVM/MPI Users' Group Meeting*, 97–104 (EuroPVM/MPI, 2004).
197. Stoudenmire, E. M. & White, S. R. Real-space parallel density matrix renormalization group. *Phys. Rev. B* **87**, 155137 (2013).
198. Secular, P. et al. Parallel time-dependent variational principle algorithm for matrix product states. *Phys. Rev. B* **101**, 235123 (2020).
199. Menczer, A. & Legeza, Ö. Massively parallel tensor network state algorithms on hybrid CPU-GPU based architectures. *J. Chem. Theor. Comput.* **21**, 1572–1587 (2023).
200. Cataldi, G. et al. Hilbert curve vs Hilbert space: exploiting fractal 2D covering to increase tensor network efficiency. *Quantum* **5**, 556 (2021).
201. Suzuki, S., Inoue, J.-i. & Chakrabarti, B. K. *Quantum Ising Phases and Transitions in Transverse Ising Models*, vol. 862 of *Lecture Notes in Physics* (Springer, 2013).
202. Baccari, F. et al. Quantum TEA: qtealeaves (1.5.8). Zenodo. <https://doi.org/10.5281/zenodo.14883066> (2025).
203. Textarossa Dibona: using node of the EuroHPC project with Nvidia A100 GPU with 40GB memory.
204. Verstraete, F., García-Ripoll, J. J. & Cirac, J. I. Matrix product density operators: simulation of finite-temperature and dissipative systems. *Phys. Rev. Lett.* **93**, 207204 (2004).
205. Zwolak, M. & Vidal, G. Mixed-state dynamics in one-dimensional quantum lattice systems: a time-dependent superoperator renormalization algorithm. *Phys. Rev. Lett.* **93**, 207205 (2004).
206. Werner, A. H. et al. Positive tensor network approach for simulating open quantum many-body systems. *Phys. Rev. Lett.* **116**, 237201 (2016).
207. Kliesch, M., Gross, D. & Eisert, J. Matrix-product operators and states: Np-hardness and undecidability. *Phys. Rev. Lett.* **113**, 160503 (2014).
208. Arceci, L., Silvi, P. & Montangero, S. Entanglement of formation of mixed many-body quantum states via tree tensor operators. *Phys. Rev. Lett.* **128**, 040501 (2022).
209. Cataldi, G. Ed-Igt. Exact diagonalization code for lattice gauge theories and quantum many-body Hamiltonians. Zenodo. <https://doi.org/10.5281/zenodo.11145318> (2024).

Acknowledgements

The authors acknowledge financial support from the following institutions: the European Union (EU) via the Horizon 2020 research and innovation program (Quantum Flagship) projects PASQuaS2 and Euryqa, via the NextGenerationEU project CN00000013 - Italian Research Center on HPC, Big Data and Quantum Computing (ICSC), and via the QuantERA projects T-NISQ and QuantHEP; the Italian Ministry of University and Research (MUR) via the Departments of Excellence grants 2023-2027 projects Quantum Frontiers and Quantum Sensing and Modelling for One-Health (QuaSiModO), and via PRIN2022 project TANQU; the Italian Istituto Nazionale di Fisica Nucleare (INFN) via specific initiatives IS-QUANTUM and IS-NPQCD; the German Federal Ministry of Education and Research (BMBF) via project QRYdDemo; the University of Bari via the 2023-UNBACLE-0244025 grant; the World-Class Research Infrastructure – Quantum Computing and Simulation Center (QCSC) of Padova University. We acknowledge computational resources by the Cloud Veneto, CINECA, the BwUniCluster, the University of Padova Strategic Research Infrastructure Grant 2017: “CAPRI: Calcolo ad Alte Prestazioni per la Ricerca e l’Innovazione”, and ReCaS Bari. We acknowledge also the financial support from the European Union’s Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant Agreement No. 101034267. G.C. acknowledges funding by the Max Planck Society, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2111 – 390814868, and the European Research Council (ERC) under the European Union’s Horizon Europe research and innovation program (Grant Agreement No. –101165667)—ERC Starting Grant QuSiGauge. This work is part of the Quantum Computing for High-Energy Physics (QC4HEP) working group.

Author contributions

G.M. designed the project. G.M., G.C., M.R., P.M., D.J., and P.S. conducted the theoretical study. G.C. and D.J. performed the numerical simulations and data analysis. S.M. conceived, supervised, and directed the project. All authors discussed the results, contributed to writing and refining the article, and approved it.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02125-x>.

Correspondence and requests for materials should be addressed to Giovanni Cataldi.

Peer review information *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025