

**Weierstraß-Institut  
für Angewandte Analysis und Stochastik  
Leibniz-Institut im Forschungsverbund Berlin e. V.**

Preprint

ISSN 2198-5855

**Quantum circuit simulation with a local time-dependent  
variational principle**

Aaron Sander<sup>1</sup>, Maximilian Fröhlich<sup>2</sup>, Martin Eigel<sup>2</sup>, Jens Eisert<sup>3,4</sup>, Mazen Ali<sup>5</sup>,  
Michael Hintermüller<sup>2,6</sup>, Richard M. Milbradt<sup>1</sup>, Robert Wille<sup>1,7,8</sup>, Christian B. Mendl<sup>1</sup>

submitted: August 26, 2025

<sup>1</sup> Technical University of Munich  
Arcisstr. 21  
80333 Munich, Germany  
E-Mail: aaron.sander@tum.de  
richard.milbradt@tum.de  
christian.mendl@tum.de  
robert.wille@tum.de

<sup>2</sup> Weierstrass Institute  
Mohrenstr. 39  
10117 Berlin, Germany  
E-Mail: maximilian.froehlich@wias-berlin.de  
martin.eigel@wias-berlin.de  
michael.hintermueller@wias-berlin.de

<sup>3</sup> Freie Universität Berlin  
Arnimallee 14  
14195 Berlin, Germany  
E-Mail: jense@zedat.fu-berlin.de

<sup>4</sup> Helmholtz-Zentrum Berlin  
für Materialien und Energie  
Hahn-Meitner-Platz 1  
14109 Berlin, Germany

<sup>5</sup> Fraunhofer ITWM  
Fraunhofer-Platz 1  
67663 Kaiserslautern, Germany  
E-Mail: mazen.ali90@gmail.com

<sup>6</sup> Humboldt-Universität zu Berlin  
Unter den Linden 6  
10099 Berlin, Germany  
E-Mail: hint@math.hu-berlin.de

<sup>7</sup> Software Competence Center  
Hagenberg GmbH (SCCH)  
Softwarepark 32a  
4232 Hagenberg im Mühlkreis, Austria

<sup>8</sup> Quantum Software  
Company GmbH  
Münchener Str. 34  
85748 Garching bei München, Germany

No. 3213  
Berlin 2025



---

2020 *Mathematics Subject Classification.* 65N75, 65C30, 60H35, 65J10.

*Key words and phrases.* Quantum computing, quantum circuit simulation, tensor networks, matrix product states, time-dependent variational principle.

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689, the Cluster of Excellence ML4Q (EXC 2004/1, project ID: 390534769).

Edited by  
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)  
Leibniz-Institut im Forschungsverbund Berlin e. V.  
Mohrenstraße 39  
10117 Berlin  
Germany

Fax: +49 30 20372-303  
E-Mail: [preprint@wias-berlin.de](mailto:preprint@wias-berlin.de)  
World Wide Web: <http://www.wias-berlin.de/>

# Quantum circuit simulation with a local time-dependent variational principle

Aaron Sander, Maximilian Fröhlich, Martin Eigel, Jens Eisert, Mazen Ali, Michael Hintermüller,  
Richard M. Milbradt, Robert Wille, Christian B. Mendl

## Abstract

Classical simulations of quantum circuits are vital for assessing potential quantum advantage and benchmarking devices, yet they require sophisticated methods to avoid the exponential growth of resources. Tensor network approaches, in particular *matrix product states* (MPS) combined with the *time-evolving block decimation* (TEBD) algorithm, currently dominate large-scale circuit simulations. These methods scale efficiently when entanglement is limited but suffer rapid bond dimension growth with increasing entanglement and handle long-range gates via costly SWAP insertions. Motivated by the success of the *time-dependent variational principle* (TDVP) in many-body physics, we reinterpret quantum circuits as series of discrete time evolutions, using gate generators to construct an MPS-based circuit simulation via a local TDVP formulation. This addresses TEBD's key limitations by (1) naturally accommodating long-range gates and (2) optimally representing states on the MPS manifold. By diffusing entanglement more globally, the method suppresses local bond growth and reduces memory and runtime costs. We benchmark the approach on five 49-qubit circuits: three Hamiltonian circuits (1D open and periodic Heisenberg, 2D  $7 \times 7$  Ising) and two algorithmic ones (*quantum approximate optimization*, *hardware-efficient ansatz*). Across all cases, our method yields substantial resource reductions over standard tools, establishing a new state-of-the-art for circuit simulation and enabling advances across quantum computing, condensed matter, and beyond.

## 1 Introduction

Classical simulation of quantum circuits plays a central role in the development of quantum technologies, providing a critical tool to validate algorithms, benchmark hardware, and probe quantum dynamics beyond current experimental reach. However, simulating large quantum systems remains notoriously difficult due to the growth of entanglement and exponential computational complexity with system size [1]. This state of affairs can be exploited in paradigmatic quantum computational tasks: Indeed, recent demonstrations which presumed quantum advantages over classical supercomputers have explicitly been constructed in order to exceed the computational capabilities of classical simulation [2, 3, 4, 5]. This, in principle, shows the computational capabilities even of near-term quantum devices. At the same time, it can be seen as an invitation to further develop scalable, structure-exploiting classical simulation methods that challenge such quantum advantages and can be used for their benchmarking.

Motivated by this line of thought, a variety of simulation strategies has been developed to tackle this challenge, ranging from tensor network families such as *projected entangled pair states* (PEPS) [6, 7] and the *multiscale entanglement renormalization ansatz* (MERA) [8], to efficient methods for special cases such as stabilizer circuits [9], as well as methods using well-established techniques from classical computing such as decision diagrams [10, 11]. More recently, contraction-optimized general

tensor networks [12, 13] have pushed classical simulation boundaries to match IBM's largest experiments [14]. Still, *matrix product states* (MPS) have remained the dominant practical tool, especially for one-dimensional systems, owing to their favorable scaling in low-entanglement regimes, algorithmic flexibility, and compatibility with gate-based quantum circuit simulation.

Among MPS-based approaches, the *time-evolving block decimation* (TEBD) algorithm [15, 16, 17, 18, 19, 1] has become a workhorse for simulating quantum circuits and many-body systems. By applying local gates sequentially and truncating intermediate states, TEBD can efficiently simulate circuits with limited entanglement. However, TEBD suffers from two fundamental limitations: (1) it introduces local truncation error after every gate, which compounds rapidly as bond dimensions grow; and (2) it cannot natively handle long-range gates, instead requiring costly SWAP networks that significantly increase overhead [20].

The *time-dependent variational principle* (TDVP) [21, 22, 23, 24, 25] provides an elegant alternative by evolving states entirely within the MPS manifold, making use of tangent spaces. Rather than truncating after gate application, TDVP projects the full time-evolution onto the tangent space of the MPS, producing a variationally optimal trajectory. TDVP can naturally accommodate long-range Hamiltonians [26] and better conserves physical quantities under limited bond dimension. These features have made TDVP a leading method in many-body physics [27, 28], though it has not yet been adapted to the discrete, gate-based setting of quantum circuit simulation, likely due to the incompatibility between TDVP's continuous-time formalism and the inherently discrete nature of quantum gates. Unlike Hamiltonian evolution, quantum circuits involve non-infinitesimal, localized unitary operations that do not readily fit within standard TDVP frameworks.

In this work, we close this gap by introducing the first TDVP-based algorithm specifically tailored for quantum circuits. The central idea is to interpret each quantum gate as a generator of local time evolution and then to apply a local TDVP update restricted to only a small region of the MPS. This reinterpretation enables to simulate entire quantum circuits with both local and long-range operations, while avoiding the stepwise truncation errors inherent to TEBD. It also leads to several key innovations that push the current limits of quantum circuit simulation. In particular:

- It enables direct simulation of long-range gates *without* SWAP insertion or scheduling overhead.
- It variationally adjusts bond dimensions using global state information, reducing local truncation errors.
- It suppresses bond dimension growth by distributing entanglement more evenly, improving scalability.

Unlike standard TDVP, which projects onto the full MPS tangent space, our method proves that a restricted projector acting only on a gate-local window yields *exactly* the same result. This theoretical insight allows us to implement the method with the same asymptotic cost as TEBD while preserving the advantages of variational evolution.

We benchmark the method on quantum circuits simulating both Hamiltonian and non-Hamiltonian dynamics, including the 49-site Heisenberg model (with open and periodic boundaries), the  $7 \times 7$  two-dimensional Ising model, and circuits with randomly parameterized two-qubit gates such as QAOA and hardware-efficient ansätze. These examples span a wide range of entanglement structures and gate layouts, from strictly local to highly non-local. Across all cases, we observe substantial reductions in memory footprint and runtime compared to TEBD, while maintaining comparable accuracy.

Our results establish the local TDVP as a new state-of-the-art for classical simulation of quantum circuits. The method bridges discrete quantum gate models and continuous-time variational principles, offering a general-purpose simulation framework with broad implications for quantum algorithm development, hardware verification, and the study of strongly-correlated quantum matter. Additionally, all code is openly accessible as part of the *Yet Another Quantum Simulator* (YAQS) [29] within the *Munich Quantum Toolkit* (MQT) framework [30].

## 2 Lessons from Continuous-Time TDVP

The core of the approach developed in this work comes from important lessons that can be drawn from techniques used in simulating time evolution of quantum many-body systems. For illustration, a state vector of an  $N$ -site quantum system with local dimension  $d$  resides in a Hilbert space  $\mathcal{H} = (\mathbb{C}^d)^{\otimes N}$ . A general quantum state vector  $|\Psi\rangle \in \mathcal{H}$  is thus a complex valued vector in a space of dimension  $d^N$

$$|\Psi\rangle = \sum_{q_1, \dots, q_N=0}^{d-1} \Psi_{q_1 \dots q_N} |q_1, \dots, q_N\rangle, \quad (1)$$

which becomes prohibitively large for classical simulation as  $N$  grows. To address this, it is common to employ a *matrix product state* (MPS) representation

$$|\Psi\rangle = \sum_{q_1, \dots, q_N=0}^{d-1} M_1^{q_1} \dots M_N^{q_N} |q_1, \dots, q_N\rangle, \quad (2)$$

which expresses  $|\Psi\rangle$  as a chain – sometimes referred to as “tensor train” of  $N$  tensors, each with a physical leg  $q_n$  of dimension  $d$  and virtual (bond) legs  $a_{n-1}, a_n$  of dimension  $\chi_{n-1}, \chi_n$ , respectively, that connect neighboring tensors [15]. This structure encodes entanglement locally and reduces the representation cost from exponential to  $\mathcal{O}(Nd\chi^2)$ , where  $\chi$  is the maximum bond dimension. This structure allows efficient storage and localized operations, making it the state-of-the-art for the simulation of large-scale systems. Each physical dimension  $q_n$  with  $1 \leq i \leq N$  and dimensions  $0 \leq q_n \leq d-1$  can be viewed as a qudit (or qubit for  $d=2$ ).

Beyond static representations, we are typically interested in the time evolution of quantum states under a possibly time-dependent Hamiltonian  $t \mapsto H(t)$ . This evolution is governed by the Schrödinger equation,

$$\frac{d}{dt} |\Psi(t)\rangle = -\frac{i}{\hbar} H(t) |\Psi(t)\rangle, \quad (3)$$

with formal solution

$$|\Psi(t)\rangle = U(t) |\Psi(0)\rangle, \quad (4)$$

where  $t \mapsto U(t)$  is a family of unitary time-evolution operators defined by

$$U(t) = e^{-iHt}. \quad (5)$$

As is common practice,  $\hbar$  has been set to 1. In many applications, this evolution is approximated by a sequence of small, discrete time steps

$$U(n\delta t) \approx \prod_{j=1}^n U_j, \quad \text{with } U_j = e^{-iH_j\delta t}. \quad (6)$$

A widely used method for simulating time evolution within the MPS framework is the *time-evolving block decimation* (TEBD) algorithm [18, 19]. TEBD approximates the unitary evolution operators  $U_i(\delta t)$  by Trotter decomposing the Hamiltonian into a sequence of local two-site gates, which are applied sequentially across the chain. After each gate application, the resulting MPS is truncated by discarding small Schmidt values, ensuring the bond dimension remains bounded [31]. This approach is particularly effective for short-range Hamiltonians and systems with low entanglement growth, where the Trotter error and truncation error can both be controlled. However, TEBD has well-known limitations: systems with long-range interactions or rapidly growing entanglement can require prohibitively large bond dimensions. From this description, one may note that quantum circuits are an analogous simulation structure such that TEBD and MPS-based quantum circuit simulation are two sides of the same coin.

An alternative approach for unitary time-evolution is provided by the *time-dependent variational principle* (TDVP) [21, 22, 23, 24, 25], which evolves quantum states entirely within a variational manifold  $\mathcal{M}_\chi$  defined by its bond dimensions  $\chi$ . Rather than applying gates and truncating, TDVP projects the full time-evolution  $U_i(\delta t)$  onto the tangent space of the MPS at each timestep using a Hamiltonian encoded as a *matrix product operator* (MPO). This process finds the provably optimal path of the time-evolution along the manifold which can be understood as the most efficient set of bond dimensions and parameters of an MPS that represents the state at each timestep. TDVP can thus be expressed as the optimization problem

$$\frac{d}{dt} |\tilde{\Psi}(t)\rangle = \operatorname{argmin}_{|\tilde{\Psi}(t)\rangle \in \mathcal{M}_\chi} \left\| \frac{d}{dt} |\Psi(t)\rangle - \frac{d}{dt} |\tilde{\Psi}(t)\rangle \right\|_{\mathcal{H}}, \quad (7)$$

which is equivalent to solving the differential equation on the manifold

$$\frac{d}{dt} |\tilde{\Psi}(t)\rangle = -i P_{T_{|\tilde{\Psi}(t)\rangle} \mathcal{M}_\chi} H |\tilde{\Psi}(t)\rangle, \quad (8)$$

where  $P_{T_{|\tilde{\Psi}(t)\rangle} \mathcal{M}_\chi}$  is the  $\mathcal{H}$ -orthogonal projection onto the tangent space of  $\mathcal{M}_\chi$  at point  $|\tilde{\Psi}(t)\rangle$  [23].

The present work makes a momentous contribution by combining these ideas in a novel fashion, resulting in substantial algorithmical improvements. Concretely, in this work, we utilize the two-site TDVP for which this projector is defined by

$$\begin{aligned} P_{T,2|\tilde{\Psi}(t)\rangle \mathcal{M}_\chi} = & \sum_{n=1}^N P_L^{[1:n-1]} \otimes I_n \otimes I_{n+1} \otimes P_R^{[n+2:N]} \\ & - \sum_{n=1}^{N-1} P_L^{[1:n]} \otimes I_{n+1} \otimes P_R^{[n+2:N]}, \end{aligned} \quad (9)$$

where  $P_L^{[1:n-1]}$  and  $P_R^{[n+2:N-1]}$  are projection operators built with left- and right-environments around sites  $(n, n+1)$ . The first summation generates a forward-evolution of the Schrödinger equation and the second summation a backward-evolution. We call these the forward projectors and backwards projectors for brevity.

This yields a set of coupled differential equations at each pair of nearest-neighbor sites in the MPS to solve the time-dependent Schrödinger equation. Practically, this concept results in a method that preserves all symmetries and conservation laws. It ensures unitary, entanglement-respecting evolution within the chosen bond dimension and can be solved in a DMRG-like sweep of the MPS [21, 32, 24]. Unlike TEBD, TDVP is inherently norm-preserving and seems promising in realizing long-range interactions in physical systems [26]. In comparison to TEBD, TDVP has never been adapted or applied to quantum circuit simulation despite its recent successes and dominance in other areas of quantum science [27, 28].

### 3 Application of TDVP to quantum circuits

While TEBD naturally operates in terms of quantum gates, TDVP is traditionally formulated for continuous time evolution under a Hamiltonian. To apply TDVP in the context of quantum circuit simulation, we reinterpret quantum gates as *discrete time-evolution operators*, following a Trotterized prescription. It is important to note that the approach taken is versatile enough to capture both Hamiltonian dynamics encoded as quantum circuits, as well as more conventional algorithmic quantum circuits which will be explored later in the manuscript.

A quantum circuit  $G$  consists of a sequence of  $|G|$  unitary gates, which we write as

$$G = \prod_{j=1}^{|G|} g_j, \quad (10)$$

mirroring the structure of discrete time evolution operators in Eq. (6). Here, each gate  $g_j$  represents a localized operation that contributes to the overall evolution.

Thanks to their unitary nature, each gate  $g_j$  can be associated with a Hermitian generator  $H_j$  such that

$$g_j = e^{-iH_j} \quad (11)$$

with a unit timestep  $\delta t = 1$ . These generators correspond to the physical interactions that implement the gates and can be interpreted as *discrete Hamiltonians*, whose exponentials yield the observable gate dynamics. For reference, a list of generators of common quantum gates can be found in Appendix A.

TDVP requires a *matrix product operator* (MPO) representation of the Hamiltonian, which in our case corresponds to the generators of quantum gates. While methods exist to construct MPOs for arbitrary Hamiltonians [33, 34, 35], nearly all common two-qubit gates acting on sites  $k$  and  $k + q$  can be associated with generators of the form

$$H = H_k \otimes H_{k+q}, \quad (12)$$

where  $q \geq 1$  defines the range of the gate, with  $q = 1$  for nearest-neighbor interactions. This structure can be encoded as a length- $q$  factorization, i.e., an MPO with bond dimension one, where the generator is decomposed into local tensors and padded with identities [36, 7].

In standard 2TDVP, the projection operators defined in Eq. (9) act on the full Hilbert space, potentially undermining the gate-level locality of quantum circuits. However, we show that these global projectors can be restricted to a local window  $[k - 1, k + q + 1]$  around the gate generator without introducing additional error compared to the full TDVP. Specifically, the projected action of the Hamiltonian becomes

$$\begin{aligned} P_{\text{loc},2}^{[k,k+q]} H |\Psi\rangle &= \sum_{n=k-1}^{k+q} P_L^{[1:n-1]} \otimes I_n \otimes I_{n+1} \otimes P_R^{[n+2:N]} \\ &\quad - \sum_{n=k-1}^{k+q-1} P_L^{[1:n]} \otimes I_{n+1} \otimes P_R^{[n+2:N]}. \end{aligned} \quad (13)$$

We refer to this as the *local TDVP*. This reduction is crucial for scalability since the cost of applying TDVP now depends only on the gate range  $q$  rather than the total system size  $N$ . A high-level overview

of this method can be found in Sec. 7.2. For details on the implementation of these projectors in tensor networks as well as visualizations, we recommend several sources for a more detailed explanation [22, 24, 27].

This approach results in two key theoretical advantages that solve the main limitations experienced in TEBD-based quantum circuit simulation, observed purely in the structure of its projectors. The local TDVP does not have any inherent limitation in performing long-range gates. This eliminates the need for complex optimizations in SWAP networks or other optimizations required to simulate quantum circuits with arbitrary gates. Building on this, unlike TEBD, where each gate update modifies the bond dimension using only local information at that bond, TDVP incorporates neighboring environments into the update. This enables bond dimension adjustments at any location to reflect correlations across the entire state, resulting in a more globally informed and optimal evolution. Even in the case of nearest-neighbor gates where the number of required projectors is minimal, the method still operates within a small window around the gate, allowing entanglement (and thus bond dimension growth) to propagate into neighboring sites.

Combining these insights, we propose a novel quantum circuit simulation algorithm consisting of two primary components:

- 1 **Single-qubit quantum gates:** These gates are applied directly to the local tensors in the MPS representation without the use of generators as no practical advantage would be gained from exponentiating rather than using the unitary directly.
- 2 **Multi-qubit quantum gates:** These gates are applied using the local TDVP projection method described above. Each gate application consists of the following steps:
  - 2.1 **Construct generator MPO:** Form the MPO representation of the gate generator localized to the involved qubits in the window  $[k - 1, k + q + 1]$  by padding the form from Eq. (12) with identities.
  - 2.2 **Ensure MPS canonical form:** Set the MPS to mixed canonical form at the first site of the window  $k - 1$ .
  - 2.3 **Apply window to MPS:** Use only the MPS tensors within the window  $M_{k-1}^{q_{k-1}} \dots M_{k+q+1}^{q_{k+q+1}}$ .
  - 2.4 **Apply local TDVP projectors:** Apply projectors including forward- and backward-evolution in one sweep from  $k - 1 \rightarrow k + q + 1$ . This step includes creating effective generators, exponentiating them with the Lanczos method [37], and applying the result to the relevant tensors.

Repeating these steps for each gate in the quantum circuit provides a scalable, flexible method with provably optimal bond dimensions for simulating complex quantum circuits of arbitrary gate ranges and structures.

## 4 Justification for local TDVP

### 4.1 Equivalence with global TDVP

In this section, we discuss a conceptual insight on the equivalence of the local TDVP with a global TDVP scheme. Notably, the local scheme only uses a minimal subset of the TDVP projectors [23],



which ensures that opposite to a full TDVP no error is introduced in the computations while the computational effort of the time evolution is reduced drastically. In fact, the complexity of the local TDVP scheme does not depend on the system size anymore but only on the number of qubits which a given gate acts on. This effectively localizes the operation and exploits the benefits of the tensor network structure.

**Theorem 1** (Local 2TDVP). *Let  $H$  be a local generator as in (12) acting on the qubit range  $[k, k+q] \subset [1, N]$ . Then, for any  $|\Psi\rangle \in \mathcal{M}_\chi$ , we have  $P_{T,2|\Psi\rangle,\mathcal{M}_\chi} H |\Psi\rangle = P_{\text{loc},2}^{[k,k+q]} H |\Psi\rangle$  where  $P_{\text{loc},2}^{[k,k+q]} H |\Psi\rangle$  is defined in Eq. (13).*

*Proof.* The proof is a direct consequence of the projections in Eq. (9) acting on  $H |\Psi\rangle$ . By exploiting the gauge freedom in the MPS representation of  $|\Psi\rangle$ , we can first consider  $|\Psi\rangle$  with left-orthogonal and right-orthogonal tensors and the center of orthogonality at site  $n$ . Since  $H$  is local, it has an MPO representation with all sites to the left of  $k$  and to the right of  $k+q$  consisting of trivial identity tensors. Putting both together, for  $n < k-1$ , we observe that

$$\begin{aligned} & P_L^{[1:n-1]} \otimes I_n \otimes I_{n+1} \otimes P_R^{[n+2:N]} H |\Psi\rangle \\ &= I_1 \otimes \cdots \otimes I_{n+1} \otimes P_R^{[n+2:N]} H |\Psi\rangle \\ &= P_L^{[1:n]} \otimes I_{n+1} \otimes P_R^{[n+2:N]} H |\Psi\rangle. \end{aligned} \quad (14)$$

This holds similarly for  $n > k+q+1$ , where the orthogonality center is to the right of the non-trivial tensors. Contracting the product of the projections with the Hamiltonian and the quantum state vector  $|\Psi\rangle$  from the right, we observe that

$$\begin{aligned} & P_L^{[1:n-1]} \otimes I_n \otimes I_{n+1} \otimes P_R^{[n+2:N]} H |\Psi\rangle \\ &= P_L^{[1:n-1]} \otimes I_n \otimes \cdots \otimes I_N H |\Psi\rangle \\ &= P_L^{[1:n-1]} \otimes I_n \otimes P_R^{[n+1:N]} H |\Psi\rangle. \end{aligned} \quad (15)$$

Hence, all terms in (9) cancel out, except the ones corresponding to  $P_{\text{loc},2}^{[k,k+q]}$ . □

For completeness we note that this can also be directly transferred to the one-site TDVP (1TDVP) scheme, where we have  $P_{T|\Psi\rangle,\mathcal{M}_\chi} H |\Psi\rangle = P_{\text{loc}}^{[k,k+q]} H |\Psi\rangle$ , with

$$\begin{aligned} P_{\text{loc}}^{[k,k+q]} &= \sum_{n=k}^{k+q} P_L^{[1:n-1]} \otimes I_n \otimes P_R^{[n+1:N]} \\ &\quad - \sum_{n=k}^{k+q-1} P_L^{[1:n]} \otimes P_R^{[n+1:N]}. \end{aligned} \quad (16)$$

Note that the 1TDVP no longer needs to address sites neighboring the gate generator and hence only acts on the range  $n \in [k, k+q]$ .

## 4.2 Same computational complexity as TEBD

To analyze the efficiency of our local TDVP method for circuit simulations, we compare its computational complexity to that of the standard TEBD approach. We break down the cost into three categories:

single-qubit, nearest-neighbor, and long-range (multi-qubit) gates. Since asymptotic statements are primarily of interest, we make use of a Landau notation to qualitatively capture the computational effort.

#### 4.2.1 General TDVP cost

We start by discussing the general TDVP cost. The full 2-site TDVP method has a known computational complexity of

$$c_{\text{TDVP}} = \mathcal{O} \left( N \left( \chi^2 d^3 D^2 + \chi^3 d^2 D + \chi^3 d^3 \right) \right), \quad (17)$$

where  $N$  is the number of sites,  $\chi$  is the maximum MPS bond dimension,  $d$  is the local Hilbert space dimension ( $d = 2$  for qubits), and  $D$  is the bond dimension of the local generator used in the TDVP update [38]. In the localized variant, where updates are applied to a block of qubits  $[k - 1, k + q + 1]$ , the cost reduces to

$$c_{\text{locTDVP}} = \mathcal{O} \left( q \left( \chi^2 D^2 + \chi^3 D + \chi^3 \right) \right) \approx \mathcal{O}(q\chi^3), \quad (18)$$

assuming  $D \ll \chi$ , which is valid for typical quantum gates.

#### 4.2.2 Single-qubit gates

For both TDVP and TEBD, single-qubit gates correspond to a local tensor contraction with complexity

$$c_{1\text{QB}} = \mathcal{O}(\chi^2), \quad (19)$$

which is negligible compared to multi-qubit operations.

#### 4.2.3 Nearest-neighbor gates

For nearest-neighbor quantum gates, when  $q = 1$ , the TDVP update acts on four consecutive qubits  $[k - 1, k + 2]$  with cost

$$c_{\text{NN}} = \mathcal{O}(\chi^3). \quad (20)$$

This is directly comparable to the TEBD cost for a two-site gate, which also scales as  $\mathcal{O}(\chi^3)$  and is dominated by the SVD step [19].

#### 4.2.4 Long-range gates

We now turn to discussing long-ranged quantum gates. To apply a gate on qubits  $[k, k + q]$ , we encounter the following approaches.

- **TEBD and SWAP networks:** This requires  $2(q - 1)$  SWAP gates to iteratively move the target qubits to adjacent positions, apply the gate, then undo the swaps. During this process, the SVD is the most expensive operation such that this leads to an overall complexity of

$$c_{\text{TEBD}} = \mathcal{O}(q\chi^3). \quad (21)$$

- **Local TDVP:** This avoids SWAP gates by directly evolving the state over the block  $[k - 1, k + q + 1]$ . This involves applying  $q + 1$  projectors and performing corresponding SVD decompositions, resulting in a cost of

$$c_{\text{locTDVP}} = \mathcal{O}(q\chi^3). \quad (22)$$

Since for large systems  $\chi$  always is the dominating factor, both complexities are asymptotically and practically equivalent at

#### 4.2.5 Summary

To summarize the discussion of the computational effort of the steps made use of, both TEBD and TDVP scale as  $\mathcal{O}(\chi^3)$  per gate application, with the main difference being how bond dimension grows during the simulation. The practical runtime is therefore determined primarily by the entanglement structure and resulting bond dimensions induced by each method rather than by asymptotic complexity constants.

## 5 Results

We now turn to presenting numerical results relating to large-scale quantum circuits, both derived from Trotterized Hamiltonian evolution as well as for families of natural algorithmic quantum circuits. We evaluate the performance and accuracy of our local TDVP method by comparing it to the established TEBD-based algorithm for simulating quantum circuits representing distinct quantum models as illustrated in Fig. 1.

We benchmark our method on five distinct 49-qubit scenarios, encompassing both physically motivated and algorithmic circuits. The first three are increasingly challenging Trotterized Hamiltonian circuits: an open-boundary one-dimensional Heisenberg model, a periodic-boundary one-dimensional Heisenberg model, and a  $7 \times 7$  two-dimensional Ising model. All are simulated at the critical point  $J = g = 1$  with a time step  $\delta t = 0.1$ .

These circuits represent highly entangling dynamics with progressively more long-range interactions: the open Heisenberg model includes only nearest-neighbor gates, the periodic variant adds a single long-range interaction, and the two-dimensional Ising model is dominated by long-range gates.

In addition to these structured models, we also simulate two general-purpose variational circuits: a *quantum approximate optimization* (QAOA) [39] circuit from the context of *variational quantum algorithms* [40] composed of alternating layers of  $R_x$  and  $R_{zz}$  gates, and a *hardware efficient ansatz* (HEA) circuit using  $U3$  and *Controlled-X* (CNOT) gates [41]. For both, each gate parameter is independently drawn from a uniform distribution over  $[-\pi, \pi]$  with no repeated values.

Each column in Fig. 1 corresponds to one of these five scenarios, where all subplots within a column are derived from a single, consistent dataset. To ensure a fair and robust comparison, all results are benchmarked against IBM Qiskit's MPS simulator [42], maintaining a consistent maximum bond dimension of  $\chi = 512$  and an SVD truncation threshold of  $s_{\max} = 10^{-9}$ . This threshold corresponds to a very mild truncation strategy, equivalent in both methods, allowing us to confidently attribute observed scalability differences solely to algorithmic distinctions rather than truncation artifacts. It is important to stress that our implementation is openly accessible as part of the *Yet Another Quantum Simulator* (YAQS) [29] within the *Munich Quantum Toolkit* (MQT) framework [30].

Fig. 1(a) compares the expectation values of the local correlator on the center bond  $\langle X_{24} X_{25} \rangle$  computed by TDVP (orange circles) and TEBD (blue triangles) to ensure accuracy. Both methods exhibit identical behavior within numerical precision until the point where TEBD simulations become infeasible due to excessive entanglement-induced bond dimension growth. Dashed vertical lines mark the maximum feasible simulation times for TEBD, highlighting TDVP's capability to accurately simulate circuits

well beyond TEBD's limits.

To understand the underlying reason for this performance disparity, we present the spatial-temporal bond dimension profiles in Fig. 1(b). While TEBD rapidly accumulates large bond dimensions at later times, TDVP maintains significantly lower bond dimensions throughout the simulation, reflecting more efficient management of entanglement. In particular, it spreads the bond dimension growth across a wider set of bonds, rather than just the bond which a gate acts on. This contrast is particularly pronounced in highly entangled scenarios, such as the two-dimensional Ising and periodic-boundary Heisenberg models, where TDVP shows clear advantages in controlling bond dimension growth. We see asymmetry in cases with long-range gates, likely stemming from the direction of the sweep. In the QAOA circuit, we see a decrease in the total bond dimension, however, similar simulability for a fixed maximum bond dimension. Finally, the HEA circuit causes the bond dimensions to grow uniformly exponentially in both methods. Notably, in all scenarios and at all Trotter steps, TDVP shows an overall reduction in total bond dimension of the MPS compared to TEBD, and in the worst case, an equivalent one.

Fig. 1(c) quantifies the computational efficiency of both methods in terms of runtime cost, approximated by  $\sum_j \chi_j^3$ , a metric directly related to computational complexity. In particular, this can be thought of as the complexity needed to perform consecutive SVDs across the MPS. TEBD exhibits steeply increasing computational costs, quickly reaching infeasible levels at relatively short simulation times. In contrast, TDVP's computational cost grows considerably more slowly, enabling simulations at larger circuit depths and longer evolution times.

The core insight is that by diffusing entanglement more globally, one encounters acceptable local bond dimensions. As a result, large-scale simulations of quantum circuits get within practical reach. In summary, our results demonstrate that the local TDVP provides a robust and computationally efficient approach for large-scale quantum circuit simulations, enabling exploration of regimes previously inaccessible by TEBD due to prohibitive entanglement growth.

## 6 Discussion

In this work, we have revisited – equipped with modern, powerful tools and recently developed insights – the question of classically simulating quantum circuits. Such endeavors are of key importance in efforts to benchmark large-scale quantum circuits and help assessing the performance of actual experimental quantum implementations of quantum circuits. Running classical against quantum simulations of quantum circuits has become a highly fruitful line of thought to improve both fields [43, 14, 12, 1, 6, 44]. In the light of this development, pushing the boundaries of efficient classical simulation and challenging quantum advantage claims constitutes an important milestone.

Concretely, this work establishes the local TDVP as a scalable and accurate tool for quantum circuit simulation. Whereas TEBD only allows for control via the SVD threshold, the local TDVP approach introduces new algorithmic parameters – such as the choice of time integrator, exponentiation method, or number of sweeps – that can be tuned for further optimization.

Future research can further explore these exciting possibilities, combine the local TDVP with other tensor network structures, and adapt the method to simulate noisy quantum circuits. While the present approach makes use of the mathematically optimal representation of the quantum state on the MPS manifold, it would be interesting to put the results into the context of rigorous bounds for truncation errors based on entanglement measures [44]. We expect that ongoing algorithmic refinements and applications to emerging quantum hardware will reveal even broader potential for the approach taken

in this work. At a higher level, this work pushes the realm of classical simulation and challenges, in an interesting and novel fashion, the development of high accuracy quantum circuits, providing important benchmarks that may be used to challenge claims of quantum advantage.

## Acknowledgments

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689, the Cluster of Excellence ML4Q (EXC 2004/1, project ID: 390534769) and the CRC 183, the BMFTR (MuniQC-Atoms, QuSol, QPIC-1), the Munich Quantum Valley, Berlin Quantum, the Quantum Flagship (PasQuans2, Millenion), and the European Research Council (DebuQC). For MATH+, the Munich Quantum Valley, and Millenion, this is a joint project involving more than one project partner. It has also been funded by the European Union under the Horizon Europe Programme as part of the project NeQST (Grant Agreement 101080086) as well as by the European Research Council projects DA QC (Grant Agreement 101001318).

## 7 Methods

### 7.1 Matrix product states

Consider a quantum state composed of  $N \in \mathbb{N}$  qudits, where each site is associated with a local Hilbert space  $\mathcal{H}_d$  of dimension  $d \in \mathbb{N}$ . The Hilbert space of the entire system is constructed by taking the tensor product of the  $N$  local Hilbert spaces, denoted as  $\mathcal{H} = \bigotimes_{n=1}^N \mathcal{H}_d$ . State vectors  $|\Psi\rangle$  within this multi-qudit Hilbert space  $\mathcal{H}$  can be expressed as a *matrix product state* (MPS) [45, 19, 46], where

$$|\Psi\rangle = \sum_{q_1, \dots, q_N=0}^{d-1} M_1^{q_1} \dots M_N^{q_N} |q_1, \dots, q_N\rangle, \quad (23)$$

where  $q_n \in [0, \dots, d-1]$  represent the local qudit for each site  $n = 1, \dots, N$ . This MPS is composed of  $N$  degree-3 tensors

$$M := \{M_n \in \mathbb{C}^{d \times \chi_{n-1} \times \chi_n} \mid n = 1, \dots, N\}, \quad (24)$$

consisting of  $d$  matrices for each index  $q_n$

$$M_n := \{M_n^{q_n} \in \mathbb{C}^{\chi_{n-1} \times \chi_n} \mid q_n = 0, \dots, d-1\}. \quad (25)$$

The complexity of this structure is governed by its bond dimensions  $\chi_n \in \mathbb{N}$  (with  $\chi_0 = \chi_N = 1$ ), which grow with the entanglement entropy of the quantum state it represents.

The MPS representation is inherently non-unique and gauge-invariant, allowing a given quantum state to be expressed in multiple equivalent ways. This flexibility enables the individual tensors to be transformed into canonical forms, which in turn simplifies many computational operations.

These conditions are the left canonical form

$$\sum_{q_n=0}^{d-1} \sum_{a_{n-1}=1}^{\chi_{n-1}} \overline{M}_n^{q_n, a_{n-1}, a_n} M_n^{q_n, a_{n-1}, a_n} = I \in \mathbb{C}^{\chi_n \times \chi_n}, \quad (26)$$

and the right canonical form

$$\sum_{q_n=0}^{d-1} \sum_{a_n=1}^{\chi_n} \overline{M}_n^{q_n, a_{n-1}, a_n} M_n^{q_n, a_{n-1}, a_n} = I \in \mathbb{C}^{\chi_{n-1} \times \chi_{n-1}}, \quad (27)$$

where  $a_n \in [1, \dots, \chi_n]$  and  $\overline{M}$  is the conjugated tensor.

By combining these conditions, the MPS is fixed in a mixed canonical form with an orthogonality center at the site tensor  $j$

$$\begin{aligned} \sum_{q_n=0}^{d-1} \sum_{a_{n-1}=1}^{\chi_{n-1}} \overline{M}_n^{q_n, a_{n-1}, a_n} M_n^{q_n, a_{n-1}, a_n} &= I \text{ such that } n < j, \\ \sum_{q_n=0}^{d-1} \sum_{a_n=1}^{\chi_n} \overline{M}_n^{q_n, a_{n-1}, a_n} M_n^{q_n, a_{n-1}, a_n} &= I \text{ such that } n > j. \end{aligned} \quad (28)$$

Given the MPS  $|\Psi\rangle \in \mathcal{M}_\chi$ , for any site  $1 \leq n < N$ , there are orthonormal bases  $|\Phi_{L,\alpha}^{[1:n]}\rangle$  and  $|\Phi_{R,\beta}^{[n+1:N]}\rangle$  for the left and right blocks, respectively, such that

$$|\Psi\rangle = \sum_{\alpha, \beta=1}^{\chi_n} |\Phi_{L,\alpha}^{[1:n]}\rangle [M_n]_{\alpha, \beta} |\Phi_{R,\beta}^{[n+1:N]}\rangle,$$

which can be thought of as a partition around the tensor  $M_n$ .

## 7.2 Local time-dependent variational principle

This section elucidates how the projectors of the local TDVP, as defined in Eq. (13), can be implemented through a sequential sweep across an MPS. We begin by explicitly defining these projectors as

$$P_{\text{loc}}^{[k, k+q]} = \sum_{n=k-1}^{k+q} K_{n, n+1} - \sum_{n=k-1}^{k+q-1} G_{n, n+1}, \quad (29)$$

where  $K_{n, n+1}$  denote forward projectors acting on each site and  $G_{n, n+1}$  represent backward projectors acting on each bond. Substituting this definition into the projected TDVP equation (Eq. (8)) yields

$$\frac{d}{dt} |\Psi\rangle = -i \sum_{n=k-1}^{k+q} K_{n, n+1} H |\Psi\rangle + i \sum_{n=k}^{k+q} G_{n, n+1} H |\Psi\rangle, \quad (30)$$

which naturally separates into  $q + 1$  forward-evolving equations

$$\frac{d}{dt} |\Psi\rangle = -i K_{n, n+1} H |\Psi\rangle, \quad (31)$$

and  $q$  backward-evolving equations

$$\frac{d}{dt} |\Psi\rangle = +i G_{n, n+1} H |\Psi\rangle. \quad (32)$$

These terms are explicitly defined using the MPS partitions

$$\begin{aligned} K_{n,n+1} = & |\Phi_L^{[1:n-1]}\rangle \langle \Phi_L^{[1:n-1]}| \\ & \otimes I_n \otimes I_{n+1} \\ & \otimes |\Phi_R^{[n+2:k+q+1]}\rangle \langle \Phi_R^{[n+2:k+q+1]}|, \end{aligned} \quad (33)$$

and

$$\begin{aligned} G_{n,n+1} = & |\Phi_L^{[1:n-1]}\rangle \langle \Phi_L^{[1:n-1]}| \\ & \otimes I_n \\ & \otimes |\Phi_R^{[n+1:k+q+1]}\rangle \langle \Phi_R^{[n+1:k+q+1]}|. \end{aligned} \quad (34)$$

Fixing the MPS into a mixed canonical form at site  $n$  and applying the conjugate transpose of the single-site map  $|\Phi_{L,\alpha}^{[1:n-1]}\rangle \otimes |\Phi_{R,\beta}^{[n+2:N]}\rangle$  to each side of Eq. (31), we simplify to local ODEs

$$\frac{d}{dt} M_{n,n+1} = -i H_{n,n+1}^{\text{eff}} M_{n,n+1}, \quad n = k-1, \dots, k+q, \quad (35)$$

where  $H_{n,n+1}^{\text{eff}}$  is a local effective Hamiltonian governing the evolution of tensor pairs  $M_{n,n+1}$ .

These Hamiltonians are then matricized and exponentiated via the Lanczos method [37]. Post-exponentiation, the updated tensors are

$$M_{n,n+1} = e^{-i H_{n,n+1}^{\text{eff}}} M_{n,n+1}. \quad (36)$$

Following this, a *singular value decomposition* (SVD) with threshold  $s_{\max}$  is applied, updating bond dimensions  $\chi_n$  to control error and move the orthogonality center forward from  $n$  to  $n+1$ .

Subsequently, the backward evolution at site  $n+1$  involves simplified backward ODEs obtained by applying the conjugate transpose of the shifted single-site map  $|\Phi_{L,\alpha}^{[1:n]}\rangle \otimes |\Phi_{R,\beta}^{[n+2:N]}\rangle$  to each side of Eq. (32)

$$\frac{d}{dt} M_n = +i H_n^{\text{eff}} M_n, \quad n = k, \dots, k+q, \quad (37)$$

since the bond tensor  $C_{n,n+1}$  is effectively identical to a single-site tensor  $M_n$ .

Applying the Lanczos method again, this tensor evolves as

$$M_n = e^{+i H_n^{\text{eff}}} M_n. \quad (38)$$

Thus, the local TDVP method practically mirrors a DMRG-like algorithm [31], performing a spatial sweep from  $n = k$  to  $n = k+q$ , alternating forward evolution updates for tensor pairs with backward evolution updates for individual tensors.

## Supplementary material

### A Generators of quantum gates

We list the generators of quantum gates used in our experiments. Here,  $\theta$  denotes the rotation angle parameter. The generators are provided in the convention  $U(\theta) = \exp(-i\theta H)$ , where  $H$  is the

generator. Note that we only use the generators of two-qubit gates,

$$R_{zz}(\theta) = \exp\left(-i\frac{\theta}{2}Z \otimes Z\right), \quad R_{yy}(\theta) = \exp\left(-i\frac{\theta}{2}Y \otimes Y\right), \quad (\text{S1})$$

$$R_{xx}(\theta) = \exp\left(-i\frac{\theta}{2}X \otimes X\right), \quad \text{CNOT} = \exp\left(-i\frac{\pi}{4}(I - Z) \otimes (I - X)\right). \quad (\text{S2})$$

since single qubit gates are contracted directly into the corresponding site tensors without using TDVP.

## B Benchmark circuits

In this work, we benchmark local TDVP on four classes of quantum circuits that reflect a broad range of structures found in many-body simulation and variational quantum algorithms. These include: (1) the one-dimensional Heisenberg model with open and periodic boundary conditions, (2) the two-dimensional Ising model on a square lattice, (3) quantum approximate optimization algorithm (QAOA) circuits using  $R_x$  and  $R_{zz}$  gates, and (4) hardware-efficient ansatz (HEA) circuits with random single- and two-qubit gates. All circuits were constructed to exhibit nontrivial entanglement dynamics and include both local and non-local interactions. For the Hamiltonian circuits, standard first-order Trotterization techniques are used, while the variational circuits follow common layouts from quantum algorithm design.

### B.1 One-dimensional Heisenberg model

The simulated circuit corresponds to the Trotterized evolution under the one-dimensional XXX Heisenberg Hamiltonian (with open and periodic boundaries)

$$H_{\text{Heis}} = \sum_{\langle i,j \rangle} J(X_i X_j + Y_i Y_j + Z_i Z_j) + h \sum_i Z_i, \quad (\text{S3})$$

where the sum is over nearest neighbors (and includes a term connecting  $L - 1$  and  $0$  for periodic boundary conditions). These are performed at the critical point  $J = h = 1$  with a timestep  $\delta t = 0.1$ . Each Trotter step consists of the following gates, parameterized as

$$\begin{aligned} R_z &= \exp(-i\delta t h Z), & R_{xx} &= \exp(-i\delta t J X \otimes X), \\ R_{yy} &= \exp(-i\delta t J Y \otimes Y), & R_{zz} &= \exp(-i\delta t J Z \otimes Z). \end{aligned} \quad (\text{S4})$$

The gates are applied in an even-odd/odd-even decomposition across the chain.  $R_z$  rotations are applied to all qubits, followed by  $R_{zz}$ ,  $R_{xx}$ , and  $R_{yy}$  gates acting on all nearest-neighbor pairs in two sweeps (even-odd and odd-even); for periodic boundaries, gates are additionally applied between sites  $(L - 1, 0)$ . The non-periodic and periodic Heisenberg circuit for 5 sites is illustrated in Fig. S1(a) and Fig. S1(b).

### B.2 Two-dimensional Ising model

The simulated circuit corresponds to the Trotterized evolution under the two-dimensional Ising Hamiltonian on a rectangular  $N = n_{\text{rows}} \times n_{\text{cols}}$  lattice

$$H_{\text{Ising}} = J \sum_{\langle i,j \rangle} Z_i Z_j + g \sum_i X_i, \quad (\text{S5})$$



where  $\langle i, j \rangle$  denotes all nearest-neighbor pairs on the grid. This is simulated at the critical point  $J = g = 1$  with a timestep  $\delta t = 0.1$ . Each Trotter step consists of the following gates, parameterized as

$$R_x = \exp(-i\delta t g X), \quad R_{zz} = \exp(-i\delta t J Z \otimes Z). \quad (\text{S6})$$

A snaking (serpentine) mapping of the two-dimensional grid to a one-dimensional qubit ordering is used for efficient MPS simulation.  $R_x$  rotations are applied to all qubits, followed by  $R_{zz}$  gates acting on all horizontally and vertically adjacent pairs, using even-odd and odd-even decompositions within each row and column which is illustrated for 9 sites in Fig. S1(c).

### B.3 QAOA circuits

To benchmark performance on variational quantum algorithms, we simulate the Quantum Approximate Optimization Algorithm (QAOA) for a 1D Ising cost Hamiltonian. Each circuit consists of  $p$  layers of alternating unitaries derived from a cost Hamiltonian and a mixing Hamiltonian:

$$U(\vec{\gamma}, \vec{\beta}) = \prod_{l=1}^p \left[ \exp\left(-i\beta_l \sum_i X_i\right) \exp\left(-i\gamma_l \sum_{\langle i,j \rangle} Z_i Z_j\right) \right]. \quad (\text{S7})$$

Each layer is implemented with a sequence of single-qubit  $R_x(\beta_\ell) = \exp(-i\beta_\ell X)$  gates, followed by two-qubit  $R_{zz}(\gamma_\ell) = \exp(-i\gamma_\ell Z \otimes Z)$  gates applied to all nearest-neighbor pairs, using even-odd and odd-even sweeps as in the Trotterized models. The angles  $(\vec{\gamma}, \vec{\beta})$  are randomly sampled from a uniform distribution (such that each gate's rotation angle is uniquely defined) to create diverse and entangling circuits, rather than optimized for a specific problem instance. This setup allows us to probe TDVP's behavior on variational, non-Hamiltonian circuits with structured but non-integrable dynamics.

### B.4 Hardware-efficient ansatz (HEA) circuits

We also benchmark the method on *hardware-efficient ansatz* (HEA) circuits, which are widely used in variational quantum algorithms due to their shallow depth and compatibility with quantum hardware constraints. Each HEA layer consists of arbitrary single-qubit rotations followed by entangling two-qubit gates between neighboring qubits

$$U_{\text{HEA}} = \prod_{l=1}^d \left[ \bigotimes_i U_3^{(i)}(\theta_l) \right] \cdot \left[ \prod_{\langle i,j \rangle} \text{CZ}_{i,j} \right], \quad (\text{S8})$$

where  $U_3(\theta) = R_z(\theta_1)R_y(\theta_2)R_z(\theta_3)$  is a universal single-qubit gate, and CZ is the controlled- $Z$  gate. We use a brickwork entangling pattern, alternating between even-odd and odd-even layers. The single-qubit parameters are sampled randomly from a uniform distribution, generating circuits with high expressivity and low regularity. This circuit class serves as a stress test for TDVP under irregular, high-depth gate sequences and random entanglement growth.

## References

- [1] Yiqing Zhou, E. Miles Stoudenmire, and Xavier Waintal. What limits the simulation of quantum computers? *Phys. Rev. X*, 10:041038, November 2020.

- [2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [3] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, and et al. Phase transitions in random circuit sampling. *Nature*, 634(8033):328–333, October 2024.
- [4] Reza Haghshenas, Eli Chertkov, Michael Mills, Wilhelm Kadow, Sheng-Hsuan Lin, Yi-Hsiang Chen, Chris Cade, Ido Niesen, Tomislav Begušić, Manuel S. Rudolph, Cristina Cirstoiu, Kevin Hemery, Conor Mc Keever, Michael Lubasch, Etienne Granet, Charles H. Baldwin, John P. Bartolotta, Matthew Bohn, Julia Cline, Matthew DeCross, Joan M. Dreiling, Cameron Foltz, David Francois, John P. Gaebler, Christopher N. Gilbreth, Johnnie Gray, Dan Gresh, Alex Hall, Aaron Hankin, Azure Hansen, Nathan Hewitt, Ross B. Hutson, Mohsin Iqbal, Nikhil Kotibhaskar, Elliot Lehman, Dominic Lucchetti, Ivaylo S. Madjarov, Karl Mayer, Alistair R. Milne, Steven A. Moses, Brian Neyenhuis, Gunhee Park, Boris Ponsioen, Michael Schechter, Peter E. Siegfried, David T. Stephen, Bruce G. Tiemann, Maxwell D. Urmei, James Walker, Andrew C. Potter, David Hayes, Garnet Kin-Lic Chan, Frank Pollmann, Michael Knap, Henrik Dreyer, and Michael Foss-Feig. Digital quantum magnetism at the frontier of classical simulations, 2025.
- [5] Dominik Hangleiter and Jens Eisert. Computational advantage of quantum random sampling. *Rev. Mod. Phys.*, 95:035001, 2023.
- [6] Siddhartha Patra, Saeed S. Jahromi, Sukhbinder Singh, and Román Orús. Efficient tensor network simulation of IBM’s largest quantum processors. *Phys. Rev. Res.*, 6:013326, Mar 2024.
- [7] Román Orús. Tensor networks for complex quantum systems. *Nature Rev. Phys.*, 1(9):538–550, August 2019.
- [8] A. V. Berezutskii, I. A. Luchnikov, and A. K. Fedorov. Simulating quantum circuits using the multi-scale entanglement renormalization ansatz. *Phy. Rev. Res.*, 7:013063, January 2025.
- [9] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70(5):052328, November 2004.
- [10] Alwin Zulehner and Robert Wille. Advanced simulation of quantum computations. *IEEE Trans. Comp.-Aided Design Int. Circ. Sys.*, 38(5):848–859, 2019.
- [11] Aaron Sander, Lukas Burgholzer, and Robert Wille. Towards hamiltonian simulation with decision diagrams. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 01, pages 283–294, 2023.
- [12] Joseph Tindall, Matthew Fishman, E. Miles Stoudenmire, and Dries Sels. Efficient tensor network simulation of IBM’s Eagle kicked ising experiment. *PRX Quantum*, 5(1):010308, 2024.
- [13] Johnnie Gray and Stefanos Kourtis. Hyper-optimized tensor network contraction. *Quantum*, 5:410, March 2021.
- [14] Youngseok Kim, Andrew Eddins, Sajant Anand, Ken Wei, Ewout Berg, Sami Rosenblatt, Hasan Nayfeh, Yantao Wu, Michael Zaletel, Kristan Temme, and Abhinav Kandala. Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618:500–505, 06 2023.

- [15] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sc. Comp.*, 33(5):2295–2317, 2011.
- [16] F. Verstraete, V. Murg, and J. I. Cirac. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.*, 57(2):143–224, 2008.
- [17] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac. Matrix product density operators: Simulation of finite-temperature and dissipative systems. *Phys. Rev. Lett.*, 93(20):207204, November 2004.
- [18] Guifré Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91:147902, October 2003.
- [19] Guifré Vidal. Efficient simulation of one-dimensional quantum many-body systems. *Phys. Rev. Lett.*, 93:040502, 2004.
- [20] E M Stoudenmire and Steven R White. Minimally entangled typical thermal state algorithms. *New J. Phys.*, 12(5):055026, may 2010.
- [21] Jutho Haegeman, J. Ignacio Cirac, Tobias J. Osborne, and Iztok Pižorn. Time-dependent variational principle for quantum lattices. *Phys. Rev. Lett.*, 107(7):070601, August 2011.
- [22] Jutho Haegeman, Tobias J. Osborne, and Frank Verstraete. Post-matrix product state methods: To tangent space and beyond. *Phys. Rev. B*, 88(7):075133, August 2013.
- [23] Jutho Haegeman, Christian Lubich, Ivan Oseledets, Bart Vandereycken, and Frank Verstraete. Unifying time evolution and optimization with matrix product states. *Phys. Rev. B*, 94:165116, October 2016.
- [24] Sebastian Paeckel, Thomas Köhler, Andreas Swoboda, Salvatore R. Manmana, Ulrich Schollwöck, and Claudius Hubig. Time-evolution methods for matrix-product states. *Ann. Phys.*, 411:167998, December 2019.
- [25] C. M. Dawson, J. Eisert, and T. J. Osborne. Unifying variational methods for simulating quantum many-body systems. *Phys. Rev. Lett.*, 100:130501, 2008.
- [26] Michael P. Zaletel, Roger S. K. Mong, Christoph Karrasch, Joel E. Moore, and Frank Pollmann. Time-evolving a matrix product state with long-ranged interactions. *Phys. Rev. B*, 91(16):165112, April 2015.
- [27] Aaron Sander, Maximilian Fröhlich, Martin Eigel, Jens Eisert, Patrick Gelß, Michael Hintermüller, Richard M. Milbradt, Robert Wille, and Christian B. Mendl. Large-scale stochastic simulation of open quantum systems, 2025.
- [28] M. Moroder, M. Grundner, F. Damanet, U. Schollwöck, S. Mardazad, S. Flannigan, T. Köhler, and S. Paeckel. Stable bipolarons in open quantum systems. *Phys. Rev. B*, 107(21):214310, June 2023.
- [29] Aaron Sander. YAQS: Yet another quantum simulator. <https://github.com/munich-quantum-toolkit/yaqs>, 2025.
- [30] R. Wille, L. Berent, T. Forster, J. Kunasaikaran, K. Mato, T. Peham, N. Quetschlich, D. Rovara, A. Sander, L. Schmid, D. Schonberger, Y. Stade, and L. Burgholzer. The MQT handbook : A summary of design automation tools and software for quantum computing. In *International Conference on Quantum Software*, pages 1–8. IEEE, 2024.

- [31] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.*, 326(1):96–192, January 2011.
- [32] Christian Lubich, Ivan V. Oseledets, and Bart Vandereycken. Time integration of tensor trains. *SIAM J. Num. Ana.*, 53(2):917–941, 2015.
- [33] Ian P. McCulloch. From density-matrix renormalization group to matrix product states. *J. Stat. Mech.*, 2007(10):P10014, 2007.
- [34] F. Fröwis, V. Nebendahl, and W. Dür. Tensor operators: Constructions and applications for long-range interaction systems. *Phys. Rev. A*, 81(6):062337, 2010.
- [35] Hazar Çakır, Richard M Milbradt, and Christian B Mendl. Optimal symbolic construction of matrix product operators and tree tensor network operators. *arXiv*, 2025.
- [36] Aaron Sander, Lukas Burgholzer, and Robert Wille. Equivalence checking of quantum circuits via intermediary matrix product operator. *Phys. Rev. Res.*, 7:023261, Jun 2025.
- [37] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat. Bur. Stand.*, 45(4):255, October 1950.
- [38] Angus J. Dunnett and Alex W. Chin. Dynamically evolving bond-dimensions within the one-site time-dependent-variational-principle method for matrix product states: Towards efficient simulation of non-equilibrium open quantum dynamics, 2020.
- [39] E. Farhi, J. Goldstone, and S. Gutmann. A quantum approximate optimization algorithm, 2014.
- [40] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles. Variational quantum algorithms. *Nature Rev. Phys.*, 3:625–644, 2021.
- [41] A. Kandala, A. Mezzcapo, K. Temme, M. Takita, M. Brink, J. W. Chow, and J. M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549:242, 2017.
- [42] Ali Javadi-Abhari, Matthew Treinish, Kevin Krsulich, Christopher J. Wood, Jake Lishman, Julien Gacon, Simon Martiel, Paul D. Nation, Lev S. Bishop, Andrew W. Cross, Blake R. Johnson, and Jay M. Gambetta. Quantum computing with Qiskit. *arXiv:2405.08810*, 2024.
- [43] S. Trotzky, Y.-A. Chen, A. Flesch, I. P. McCulloch, U. Schollwöck, J. Eisert, and I. Bloch. Probing the relaxation towards equilibrium in an isolated strongly correlated one-dimensional Bose gas. *Nature Phys.*, 8:325–330, 2012.
- [44] S. Cichy, P. K. Faehrmann, L. Bittel, J. Eisert, and H. Pashayan. Classical simulation of noisy quantum circuits via locally entanglement-optimal unravelings, 2025.
- [45] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863, November 1992.
- [46] D. Perez-Garcia, F. Verstraete, M.M. Wolf, and J.I. Cirac. Matrix product state representations. *Quant. Inf. Comp.*, 7:401, 2007.

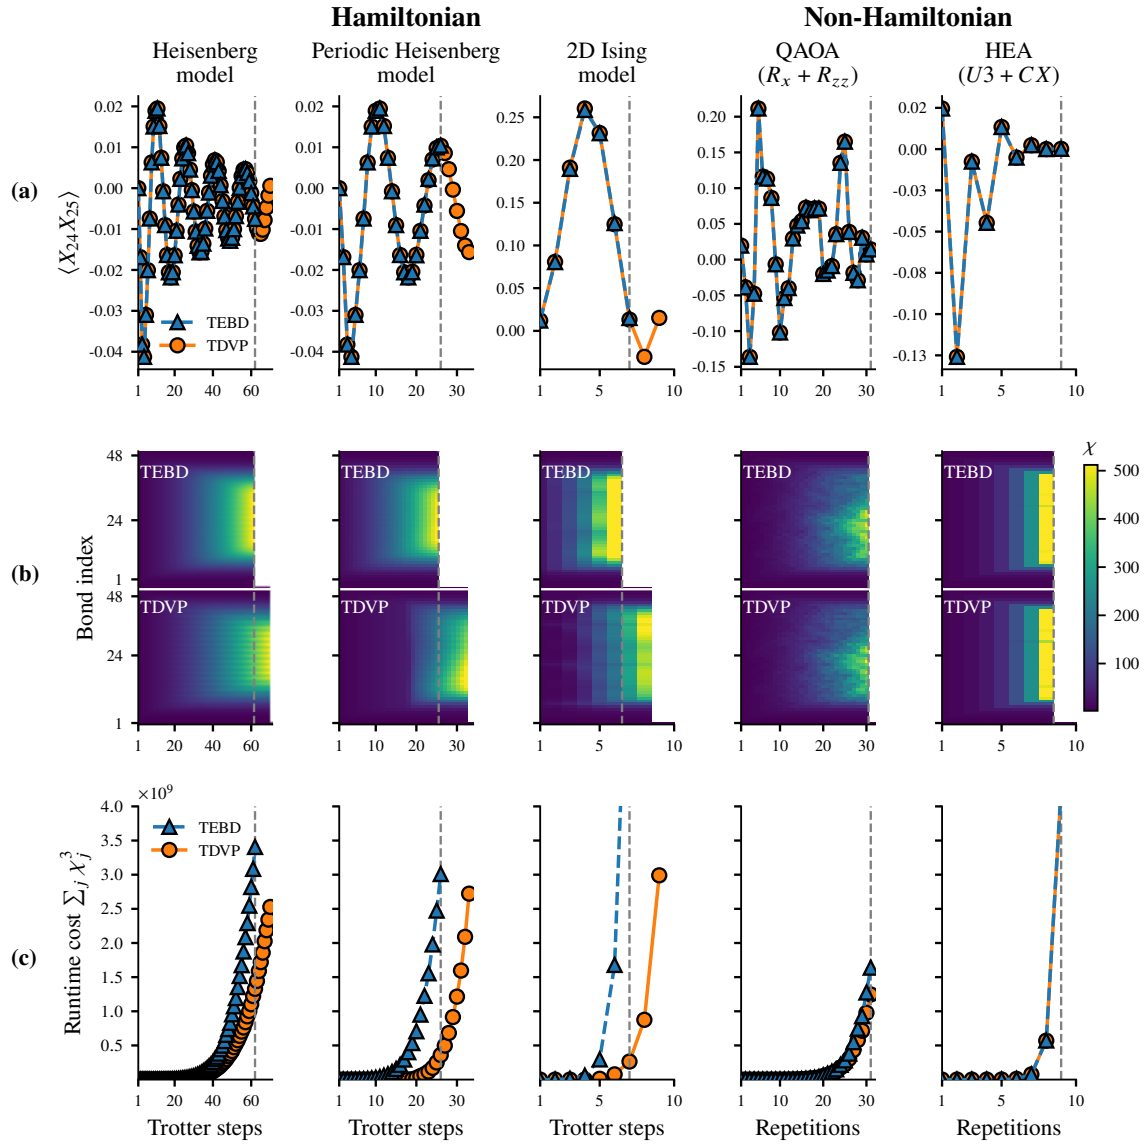


Figure 1: **Comparison of TDVP and TEBD for large-scale quantum circuit simulation.** This figure includes the results for simulating circuits of three Trotterized Hamiltonian circuits and two non-Hamiltonian circuits, all with 49 qubits. This includes a one-dimensional Heisenberg model (column 1), periodic Heisenberg model (column 2), two-dimensional Ising model (column 3), as well as a *quantum approximate algorithm* (QAOA) circuit (column 4) and a *hardware-efficient ansatz* (HEA) (column 5). (a) Expectation values of the correlator  $\langle X_{24}X_{25} \rangle$ . All circuits are simulated using TEBD (blue triangles) and localized TDVP (orange circles). Both methods match within numerical precision. (b) Bond dimension profiles for TEBD (top) and TDVP (bottom) up to a maximum  $\chi = 512$ . TDVP maintains significantly lower bond dimensions throughout the simulation by spreading bond dimension growth across neighboring bonds of the MPS. Dashed lines mark TEBD simulation limits. (c) Runtime cost  $\sum_j \chi_j^3$  versus Trotter steps. TDVP reduces computational costs through its decreased bond dimension, enabling the simulation of deeper circuits inaccessible to previous MPS-based quantum circuit simulation methods.

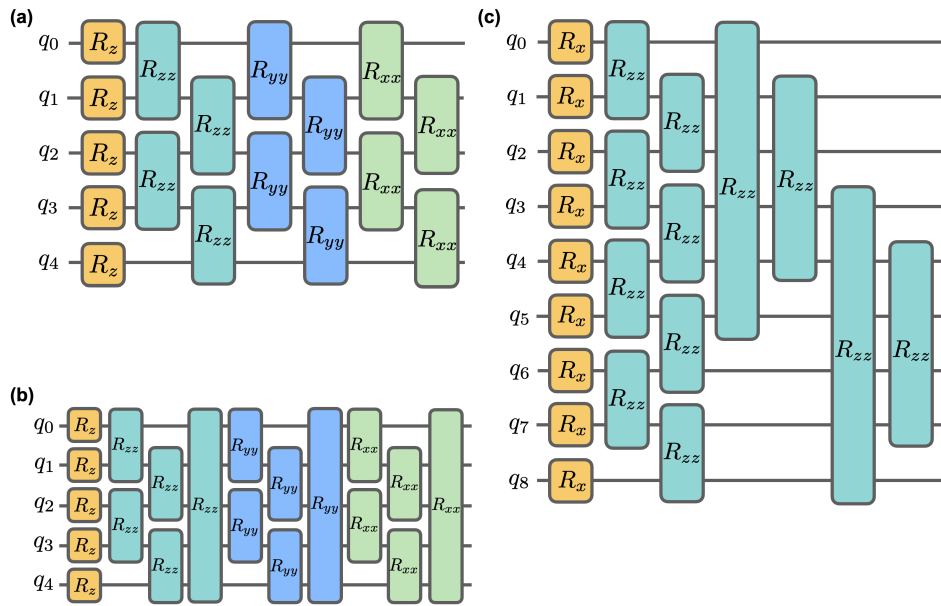


Figure S1: **Trotterized quantum circuits for the models benchmarked in this work.** (a) Circuit for the open boundary one-dimensional Heisenberg model. (b) Circuit for the periodic one-dimensional Heisenberg model. (c) Circuit for the  $3 \times 3$  two-dimensional Ising model, mapped to a one-dimensional snaking MPS layout.