


Quantum-classical eigensolver using multiscale entanglement renormalization

Qiang Miao  and Thomas Barthel

*Department of Physics, Duke University, Durham, North Carolina 27708, USA
and Duke Quantum Center, Duke University, Durham, North Carolina 27701, USA*



(Received 31 August 2021; accepted 22 May 2023; published 30 August 2023)

We propose a variational quantum eigensolver (VQE) for the simulation of strongly correlated quantum matter based on a multiscale entanglement renormalization ansatz (MERA) and gradient-based optimization. This MERA quantum eigensolver can have substantially lower computation costs than corresponding classical algorithms. Due to its narrow causal cone, the algorithm can be implemented on noisy intermediate-scale quantum (NISQ) devices and still describe large systems. It is particularly attractive for ion-trap devices with ion-shuttling capabilities. The number of required qubits is system-size independent and increases only to a logarithmic scaling when using quantum amplitude estimation to speed up gradient evaluations. Translation invariance can be used to make computation costs square-logarithmic in the system size and describe the thermodynamic limit. We demonstrate the approach numerically for a MERA with Trotterized disentanglers and isometries. With a few Trotter steps, one recovers the accuracy of the full MERA.

DOI: [10.1103/PhysRevResearch.5.033141](https://doi.org/10.1103/PhysRevResearch.5.033141)

I. INTRODUCTION

The complexity of quantum many-body systems makes it a formidable challenge to understand the properties of quantum matter, in particular in strongly correlated regimes where perturbative approaches fail. Hence, powerful classical simulation techniques like quantum Monte Carlo [1–4] and tensor networks states (TNSs) [5–10] have been developed. A strength of TNS techniques is that they are also applicable for frustrated quantum magnets and fermionic systems [11–15], where quantum Monte Carlo is hampered by the negative-sign problem [16,17]. These classes of systems include candidate spin liquid materials [18–22], fractional quantum Hall physics [23,24], and high-temperature superconductors [25,26].

Consider a lattice system with N sites, each associated with a site Hilbert space of dimension d such that the total Hilbert space has dimension d^N . The idea of TNSs is to approximate the many-body state by a network of partially contracted tensors. The tensors may carry physical indices that label site basis states and additional bond indices of dimension χ which are contracted with corresponding indices of other tensors. The structure of the network and the required bond dimension χ are adapted to the entanglement structure in the system. Typically, the more entangled a system is, the larger χ needs to be in order to achieve a desired approximation accuracy. To approximate the ground state of a given model \hat{H} by a TNS $|\Psi\rangle$, one minimizes the energy $\langle\Psi|\hat{H}|\Psi\rangle/\|\Psi\|^2$ with respect to the tensor elements. The beauty of the approach is that

computation costs for optimization steps are reduced from exponential in N to polynomial in N . In particular, they are linear in N for matrix product states (MPSs) [5,6,27–29], projected entangled pair states (PEPS) [7,8,30–32], and the multiscale entanglement renormalization ansatz (MERA) [9,33]. For homogeneous MERAs, one can reduce the cost to $O(\log N)$ and even access the thermodynamic limit $N \rightarrow \infty$. However, the classical computation time may scale with a high power of the bond dimension χ . While it is only $O(\chi^3)$ for MPSs in one-dimensional (1D) systems [6,28], it is $O(\chi^{7\cdots 9})$ for 1D MERAs [34], $O(\chi^{10\cdots 12})$ for 2D PEPS [35,36], and $O(\chi^{16\cdots 28})$ for 2D MERAs [37,38]. Hence, practicable χ are usually rather small, which limits the approximation accuracy.

In this work, we propose and analyze a hybrid quantum-classical variational eigensolver [39] to overcome these limitations, where many-body ground states are approximated by adapted MERA states and (small) quantum computers are employed to efficiently execute tensor contractions. In this context, MERAs have four advantages over other TNSs: (i) MERAs can be applied for systems with any number of spatial dimensions, (ii) all tensors are unitary or isometric, which allows for a rather direct implementation on quantum computers, (iii) MERA expectation values for local operators depend only on narrow causal cones such that they can be evaluated exactly and large systems can be simulated on noisy intermediate-scale quantum (NISQ) devices, and (iv) sets of MERAs are closed which implies that optimizers always exist [40]. Also, MERA optimizations are not hampered by barren plateaus [41,42].

Quantum algorithms using MPSs for 1D systems were recently suggested in Refs. [43–47]. Two prominent quantum-computing platforms are superconducting qubits [48,49] and ions in electromagnetic traps [50,51]. Ion-trap systems with qubit-shuttling capabilities [52–54] are particularly interesting for the quantum MERA scheme.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

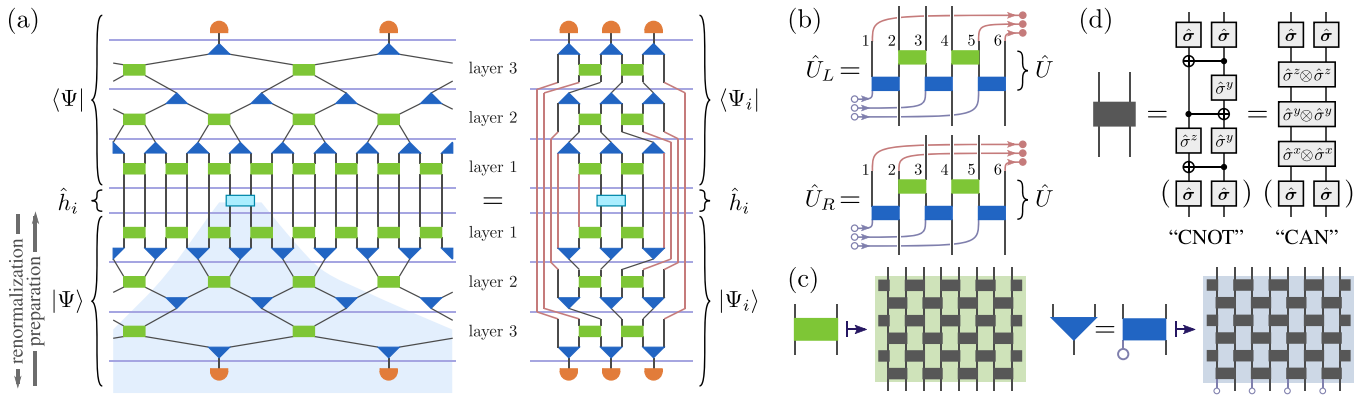


FIG. 1. TMERA structure and implementation. (a) Expectation value $\langle \Psi | \hat{h}_i | \Psi \rangle$ for a homogeneous binary 1D MERA $|\Psi\rangle$ with $T = 3$ layers. The MERA consists of disentglers (boxes) and isometries (triangles). Contraction lines between the tensors correspond to renormalized site vector spaces with dimension χ . The shaded region indicates the causal cone for a two-site operator \hat{h}_i . The cone has width three, i.e., contains at most three renormalized sites in each layer. Only causal-cone states like $|\Psi_i\rangle$, associated with \hat{h}_i , need to be generated on the quantum computer. (b) Isometries can be realized as unitaries where some input qubits are initialized in a reference state like $|0\rangle$ (open circles). Such $|0\rangle$ qubits can be moved in and others (filled circles) can be moved out of the quantum register or reset after applying the layer-transition maps. (c) A Trotter structure with t steps is imposed on each MERA tensor, making it a circuit of two-qubit gates. Here $t = 3$ and $\chi = 16$, i.e., $q = 4$ qubits per contraction line. (d) Each Trotter gate can be implemented using CNOTs and single-qubit rotations or, equivalently, single and two-qubit rotations. The specified Pauli operators generate the corresponding rotations (1), and $\hat{\sigma}$ refers to a general single-qubit rotation.

II. MERA ADAPTED FOR QUANTUM COMPUTERS

A MERA [9,33] is a hierarchical TNS motivated by the real-space renormalization group [55–57]: In each renormalization step $\tau = 1, \dots, T$, unitaries with small spatial support are applied to disentangle the system to some extent, before isometries are applied in order to map a block of b sites into a new renormalized site. In the process, states that are not important for the representation of the ground state are discarded. With a branching ratio of b , this process ends with one or a few renormalized sites after $T \sim \log_b N$ steps, and the resulting few-site problem can be solved exactly. While the physical site Hilbert spaces have dimension d , Hilbert spaces of renormalized sites have dimension χ . Seen in reverse, the renormalization group scheme defines a many-body state $|\Psi\rangle$. This state is a MERA with bond dimension χ . It consists of T layers, each comprising the unitary disentglers and isometries of a renormalization step. Optimizing the tensor elements to minimize the energy expectation value $E = \langle \Psi | \hat{H} | \Psi \rangle$, one obtains a ground-state approximation.

In principle, it is straightforward to prepare a MERA $|\Psi\rangle$ on quantum computers. Assume $\chi = 2^q$ such that every renormalized site corresponds to q qubits. A disentangler that acts on n (renormalized) sites is a $\chi^n \times \chi^n$ unitary acting on nq qubits. It can be decomposed into a circuit of $O(4^{nq})$ single-qubit and CNOT gates [58–60]. An isometry that maps n sites into $m > n$ can be implemented as a unitary acting on nq qubits and $(m - n)q$ additional ones initialized in state $|0\rangle$. It requires $O(2^{(n+m)q})$ single-qubit and CNOT gates [61].

For simplicity, we assume a Hamiltonian $\hat{H} = \sum \hat{h}_i$ with finite-range interaction terms \hat{h}_i . As exemplified in Fig. 1(a), many tensors cancel in expectation values $\langle \Psi | \hat{h}_i | \Psi \rangle$ due to their isometric property. The causal cone of \hat{h}_i comprises all tensors that can influence the expectation value, and we

define $|\Psi_i\rangle$ as the corresponding causal-cone TNS such that $\langle \Psi | \hat{h}_i | \Psi \rangle = \langle \Psi_i | \hat{h}_i | \Psi_i \rangle$. For a binary 1D MERA, disentglers act on $n = 2$ sites, and the cost to evaluate $\langle \Psi_i | \hat{h}_i | \Psi_i \rangle$ would scale in q as $O(4^{2q})$. Hence, the cost for the evaluation of an energy gradient would scale as $O(4^{4q} = \chi^8)$. This is only a modest improvement over the scaling $O(\chi^9)$ of the classical computation time. As discussed in Appendix A, the differences are generally more pronounced in higher dimensions. As an example, the quantum and classical gradient evaluation costs for the 2D $2 \times 2 \mapsto 1$ MERA of Ref. [37] scale as $O(4^{8q} = \chi^{16})$ and $O(\chi^{28})$, respectively. In any event, one also needs to account for the required number of measurement samples in the quantum case, and the quantum computational complexity can be reduced drastically by imposing further structure on the MERA.

III. TROTTERIZED TENSORS

There are many options for substructures. Here we choose to impose a Trotter structure on the MERA tensors. In particular, they shall consist of t Trotter steps, each comprising local unitary gates that act on, say, two nearest-neighbor qubits; see Fig. 1(c). For tensors that act on n (renormalized) sites, each Trotter step consists of $O(nq)$ local gates. For a Trotterized MERA (TMERA) with T layers, the measurement of a local expectation value $\langle \Psi_i | \hat{h}_i | \Psi_i \rangle$ then requires $O(Tt)$ time on the quantum computer. We will see that, using translation invariance, the measurement of the energy gradient requires $O(T^2 t^2 nq)$ time. As we will also see in benchmark simulations, the local unitary gates approach identities when increasing the number t of Trotter steps. This establishes a connection to Trotterization as used in time evolution problems [62–65].

Recently Refs. [66] and [67] appeared. The first studies the expressiveness of TMERAs numerically; the second is an experimental demonstration, measuring critical correlations in preoptimized TMERAs with bond dimension $\chi = 4$ using an ion-trap system.

IV. HYBRID OPTIMIZATION ALGORITHM

In classical computations, MERA states are optimized by evaluating the so-called environment for each tensor and updating tensors one by one [68]. On a quantum computer, we can measure only observables, and the tensor environment is not accessible. The hybrid algorithm works as follows. The Trotter gates can be written as small circuits, parametrized through the angles $\theta = (\theta_1, \theta_2, \dots)$ of rotations

$$\hat{R}_{\hat{\sigma}}(\theta) := e^{-i\theta\hat{\sigma}/2} = \mathbb{1} \cos \frac{\theta}{2} - i\hat{\sigma} \sin \frac{\theta}{2} \quad (1)$$

with respect to Hermitian unitary operators $\hat{\sigma}$ like the Pauli matrices $\{\mathbb{1}, \hat{\sigma}^x, \hat{\sigma}^y, \hat{\sigma}^z\}$ or tensor products thereof. A standard choice is depicted in Fig. 1(d). It comprises three CNOT gates, one $\hat{\sigma}^z$ and two $\hat{\sigma}^y$ single-qubit rotations, as well as four general single-qubit gates [69,70]. The number of angles per Trotter gate agrees with $\dim \text{SU}(4) = 15$ and reduces to nine angles when exploiting the unitary gauge freedoms in the TMERA. The energy gradient $\partial_{\theta} E$ can be evaluated by measuring

$$\partial_{\theta_j} E = \frac{1}{2}[E(\theta_j + \pi/2) - E(\theta_j - \pi/2)], \quad (2)$$

where all angles except for θ_j are kept fixed [71–73]. A derivation is given in Appendix C. The energy can now be minimized by a gradient-based algorithm like L-BFGS [74,75].

In experiments two-qubit gates are typically much more costly than single-qubit gates. For ion-trap and superconducting systems, Refs. [76–79] specify typical single-qubit gate times of $\sim 10 \mu\text{s}$ and $\sim 30 \text{ ns}$, respectively, whereas two-qubit gates require $\sim 100 \mu\text{s}$ and 200 ns , respectively. The CNOT parametrization [Fig. 1(d)] of the Trotter gates has the drawback that CNOT gates require two-qubit rotations with large angles. In the ion-trap and superconducting systems, CNOT is implemented using an effective Ising $\hat{\sigma}^{\alpha} \otimes \hat{\sigma}^{\alpha}$ interaction with rotation angle $\theta = \pi/2$ [79–82]. A better choice is then the canonical (CAN) parametrization in Fig. 1(d) that comprises three native $\hat{\sigma}^{\alpha} \otimes \hat{\sigma}^{\alpha}$ rotations ($\alpha = x, y, z$) and four general single-qubit rotations [83,84]. The benchmark simulations, discussed below, show that the occurring two-qubit angles for this parametrization are rather small. Furthermore, we find that the optimization actually works best in a parametrization-free fashion. Such a Riemannian quasi-Newton method on quantum circuits is described in Appendix C 2.

V. TRANSLATION INVARIANCE

For translation-invariant systems, the interaction terms \hat{h}_i are translates of the same operator \hat{h} . Correspondingly, we can reduce the number of variational parameters. A homogeneous MERA has translation-invariant layers, i.e., each layer consists of repeating identical groups of tensors. A binary

1D MERA, for example, is then characterized by a single disentangler and a single isometry for each layer. For a heterogeneous system, the derivatives (2) can be evaluated by measuring expectation values for all terms h_i that have the tensor of angle θ_j in their causal cone. In total, this requires $O(NT = N \log_b N)$ measurements. With translation invariance, this can be reduced to $O(T)$ by either using classical random bits or introducing auxiliary qubits: For the 1D case illustrated in Fig. 1, there are two unitary transition maps \hat{U}_L and \hat{U}_R . Either of them has to be applied to progress in the preparation of the causal-cone state $|\Psi_i\rangle$ from layer τ to $\tau - 1$. The specific sequence depends on the location i of the interaction term. In order to evaluate the energy density $e := \frac{1}{N} \langle \Psi | \hat{H} | \Psi \rangle = \frac{1}{N} \sum_i \langle \Psi_i | \hat{h}_i | \Psi_i \rangle$ for the entire system at once, we can replace $\hat{U}_{L,R}$ by their convex combination such that we obtain the state on layer $\tau - 1$ as

$$\hat{\rho}^{(\tau-1)} = \frac{1}{2}(\hat{U}_L \hat{\rho}^{(\tau)} \hat{U}_L^{\dagger} + \hat{U}_R \hat{\rho}^{(\tau)} \hat{U}_R^{\dagger}). \quad (3)$$

For experiments, such quantum channels can be implemented by randomly selecting \hat{U}_L or \hat{U}_R in each transition. Practically, this constant reprogramming of the hardware in gradient evaluations can be slow. So, alternatively, the channels can be lifted to fixed unitary evolutions on a larger Hilbert space [85,86]. For Eq. (3), adding a single auxiliary qubit per layer is sufficient. Initializing it in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ and applying \hat{U}_L or \hat{U}_R conditioned on the auxiliary qubit realizes the channel (3). Appendix B gives details on the efficient realization of layer-transition maps for various MERAs.

Homogeneous MERAs are necessarily defined with periodic boundary conditions. Let the linear system sizes L_x, L_y, \dots be large enough such that the causal cone of any local interaction term \hat{h}_i does not close upon itself along any of the spatial dimensions. Then, repeating the MERA tensor network in any spatial direction ($L_{\alpha} \mapsto nL_{\alpha}$) with accordingly adapted boundary conditions defines families of homogeneous MERAs, which all have the same energy density e . In particular, the results capture the thermodynamic limit $N \rightarrow \infty$.

VI. QUBIT RESETS AND ION SHUTTling

An attractive feature of the proposed quantum-classical TMERA algorithm is that, while we can simulate large systems, at any stage, only a system-size-independent number of qubits need to be acted upon with the unitary gates. When evaluating observables or gradients, as we progress from layer to layer, only the qubits inside the causal cone need to be in the quantum register. These are, e.g., $4q$ qubits for 1D binary and ternary MERAs, and $14q$ qubits for the 2D $2 \times 2 \mapsto 1$ MERA of Ref. [37]. In every layer transition, some contraction lines (groups of q qubits) leave the causal cone. The same number of (new) qubits, initialized in state $|0\rangle$, are needed to realize the isometries of the next MERA layer. When space efficiency is the highest priority, one can reset the qubits [89–93] that exit the causal cone to $|0\rangle$ for reuse. Auxiliary qubits can be reset as well. When one wants to minimize execution times, one can employ quantum amplitude estimation (QAE) [94,95] in the gradient evaluations. For a preparation $|\Psi_i\rangle = \hat{U}_i|0, \dots, 0\rangle$ of the causal-cone state, QAE requires application of powers




FIG. 2. Benchmark simulations. The plots show the convergence and accuracy of energy densities e for the transverse Ising model (4) using homogeneous modified binary MERAs with $T = 6$ layers and bond dimension $\chi = 8$ ($q = 3$). Panel (a) shows the convergence at $g = 1.25$ for TMERAs with $t = 2$, starting from the product state with state $e^{-i\frac{\pi}{4}\hat{\sigma}^z} e^{-i\frac{\pi}{4}\hat{\sigma}^y} |\uparrow\rangle$ on every site. In first optimization phases, disentangled are removed (set to $\mathbb{1}$), i.e., Trotterized tree tensor networks (TTNs) [87,88] are optimized. The resulting state is used to initialize the TMERA optimization. (b) Energy accuracy of optimized TMERAs and MERAs with full tensors (fMERAs) as a function of the field strength g . Local minima are avoided by scanning from $g = 1.25$ to $g = 0.75$ and back. Especially in the paramagnetic phase ($g > 1$) and at the critical point $g = 1$, the accuracy of the fMERA is recovered with only $t = 2$ Trotter steps per tensor.

$(\hat{U}_i)^m$. In this case, one cannot employ the (nonunitary) mid-circuit qubit resets. While the total number of required qubits is then logarithmic in the system size, still, only causal-cone qubits need to reside in the quantum register, and the others can be moved to a quantum memory. In ion-trap systems, this can be accomplished by shuttling as demonstrated in Refs. [52–54].

VII. BENCHMARK SIMULATIONS AND SCANNING

To demonstrate and benchmark TMERAs, we simulate the 1D transverse-field Ising model

$$\hat{H} = - \sum_i \hat{\sigma}_i^x \hat{\sigma}_{i+1}^x + g \sum_i \hat{\sigma}_i^z. \quad (4)$$

It has a critical point at $g = 1$ with the paramagnetic phase for $g > 1$ and the ferromagnetic phase for $g < 1$. Figure 2 shows results for homogeneous TMERAs with the modified binary network structure [34], using an L-BFGS optimization. The TMERA energy densities e are compared to the exact infinite-system value e_{gs}^∞ . The left panel shows the convergence for $g = 1.25$, which quickly reaches a high accuracy. The right panel shows TMERA accuracies for $0.75 \leq g \leq 1.25$. Local minima are avoided through *scanning*, i.e., starting at $g = 1.25$, g is lowered in steps, and the converged TMERA of the previous step is used to initialize the optimization of the next. Upon reaching $g = 0.75$, we start scanning back to $g = 1.25$. The numerical results confirm that a few Trotter steps t are sufficient to reach accuracies comparable to the full (non-Trotterized) MERA. In particular, $t = 2$ gives already excellent results for $\chi = 2^3$.

For the experimental implementation, the Trotter gates can be expressed in the CAN representation. Figure 3 shows distributions of the rotation angles in the converged TMERA at different g . They are peaked at small angles. This remains true even for the critical point $g = 1$. The fact that most angles

are small implies that these quantum gates can be executed quickly or at correspondingly higher fidelity.

VIII. COMPUTATION COST AND ACCURACY

Exploiting translation invariance, the $O(Ttq)$ components (2) of the energy gradient can be evaluated by preparing the corresponding causal-cone states, which costs $O(Tt)$ time, and then projectively measuring the local interaction term \hat{h} . With N_s samples per term, the statistical error of the gradient and, hence, the achievable energy accuracy scale as $\epsilon \propto 1/\sqrt{N_s}$. Thus, the quantum cost for each TMERA optimization step is $O(T^2 t^2 q / \epsilon^2)$. Using QAE [94,95], the cost reduces to $O(T^2 t^2 q \log(1/\epsilon)/\epsilon)$ while increasing the circuit depth by a factor $O(1/\epsilon)$. Our simulations show that the error ϵ decreases according to a power law $\epsilon \sim (t^2 q)^{-\alpha}$.




FIG. 3. Angle distributions in converged TMERAs. Trotter gates can be parametrized in the canonical (CAN) form based on four single-qubit rotations and three Ising-interaction gates (XX , YY , and ZZ). All angles are found to be peaked around 0, which is favorable for the experimental realization. The plots show angle distributions for converged TMERAs from Fig. 2 with $t = 2$ Trotter steps.

For fixed q , ϵ decreases until reaching the accuracy of the MERA with full tensors (fMERA) of bond dimension $\chi = 2^q$. Upon approaching the saturation, one should increase q . The fMERA computation cost also follows a power law $O(T\chi^r) = O(\epsilon^{-r/\beta})$, where exponent r is determined by the contraction cost and exponent β by the relation between χ and accuracy ϵ . For 1D MERAs, model-dependent exponents $\beta \approx 3.8, \dots, 6.8$ have been reported [34]. In simulations of the critical bilinear-biquadratic spin-1 chain [96–100] with modified binary MERAs ($r = 7$), we find that the QAE cost $O(\epsilon^{-1-1/\alpha})$ is already lower than the classical fMERA cost $O(\epsilon^{-r/\beta})$, providing a polynomial advantage; see Appendix A 5. As the exponent r for the classical fMERA cost is very large for higher-dimensional systems ($r \geq 16$), the quantum algorithm should substantially outperform the classical simulations for models in ≥ 2 spatial dimensions. It is numerically very expensive to determine the scaling exponents for higher-dimensional systems, and further investigations on this subject are needed.

IX. DISCUSSION

The presented TMERA quantum eigensolver allows for the approximation of many-body ground states with a system-size independent number of qubits, and it can substantially outperform classical MERA simulations. The Appendices provide details on the computational complexity for different MERA network structures (Appendix A), the realization of layer-transition maps for homogeneous TMERAs (Appendix B), optimization methods (Appendix C) including a Riemannian version of the L-BFGS algorithm, and the influence of different Trotter gate parametrizations (Appendix D).

Reference [101] discusses DMERAs, which are a special type of TMERAs, where the number of Trotter steps D in each layer is directly linked to the width $\sim 2D$ of the causal cone. The TMERAs that we consider here have more structure, which allows one to tune these quantities independently, and the imposed structure admits a more direct comparison with the typical MERAs used in classical simulations. The optimization based on the simultaneous perturbation stochastic approximation (SPSA), suggested in Ref. [101], is considerably less efficient because the energy derivative is evaluated only along random directions in the high-dimensional search space. Our gradient-based approach and the described utilization of translation invariance can be applied for any TMERA, including DMERAs. Conversely, the robustness to noise as analyzed in Ref. [101] also applies generally to TMERAs.

The decomposition of the MERA tensors into layers of nearest-neighbor Trotter gates is natural but not necessary. In future research, one could explore other network topologies to leverage, e.g., the all-to-all connectivity of ion-trap systems [102,103] and to increase the expressiveness of TMERAs at fixed cost. One could also consider other gate types, especially those that are naturally available in prominent quantum-computing architectures. An example are multiqubit Mølmer-Sørensen gates [104,105]. For an implementation on present-day devices, small two-qubit rotation angles are desirable. Hence, it will be interesting to explore how the angles and the TMERA accuracy are affected by adding large-angle penalty terms to the energy functional. See the follow-up paper Ref. [106] for further analysis of TMERAs.

ACKNOWLEDGMENTS

We gratefully acknowledge helpful discussions with Marko Cetina, Kenneth R. Brown, Christopher R. Monroe, Jungsang Kim, Iman Marvian, Sarah Brandsen, and Yikang Zhang, and support through U.S. Department of Energy Grant No. DE-SC0019449.

APPENDIX A: COMPUTATIONAL COMPLEXITY FOR 1D AND 2D MERA

Let us discuss the quantum computational complexity for different TMERAs in one and two dimensions and compare to the corresponding classical simulation costs. Figure 4 shows the considered MERA networks: the 1D binary MERA, a modified 1D binary MERA [34], the 1D ternary MERA, a 2D $2 \times 2 \mapsto 1$ MERA [37], and a 2D $3 \times 3 \mapsto 1$ MERA [68]. We use the following labels:

- (1) b denotes the MERA branching ratio.
- (2) T denotes the number of layers.
- (3) $\tau = 1, \dots, T$ labels layers with layer $\tau = 1$ acting on the physical sites, i.e., τ increases in the renormalization direction.
- (4) $\chi = 2^q$ denotes the bond dimension with q being the corresponding number of qubits per renormalized site.
- (5) A denotes the cross section of the causal cone for local operators, defined as the maximum number of renormalized sites inside the causal cone at any layer interface ($\tau \rightarrow \tau \pm 1$).
- (6) t denotes the number of Trotter steps for each tensor in TMERA (or an upper bound).

For heterogeneous MERA, the total number of sites is denoted by N and assumed to be $\sim b^T$.

1. Classical time complexity

The costs for optimizing MERAs on classical computers are determined by the cost of computing the so-called environment of a tensor. This is in turn determined by the cost of applying a renormalization step ($\tau \mapsto \tau + 1$) to a local interaction term or, equivalently, for propagating a reduced density matrix inside the causal cone in the preparation direction ($\tau \mapsto \tau - 1$). In each step, one needs to contract the tensors of disentanglers and isometries and trace out sites that leave the causal cone. The costs for these operations, which one obtains by optimizing the contraction sequence, are given in Table I. Generally speaking, it is favorable to have a narrow causal cone, which explains why the classical costs for the 1D modified binary and ternary MERAs are smaller than those of the plain binary MERA. On the other hand, for a given bond dimension χ , the binary MERA can encode more entanglement than the other two 1D MERA types and generally achieves higher accuracy.

The costs shown in Table I refer to one evaluation of the global energy gradient on a classical computer, or equivalently, a single update of all tensors in the Evenbly-Vidal algorithm [68]. For a homogeneous MERA, this cost, like the number of different tensors, is linear in T . For heterogeneous MERAs, the total number of tensors and the cost are proportional to $\sum_{\tau=1}^T b^\tau \sim N$. The table shows the classical costs for MERAs with full tensors (fMERAs) because, on classical computers, there is not much to gain by exploiting the Trotter

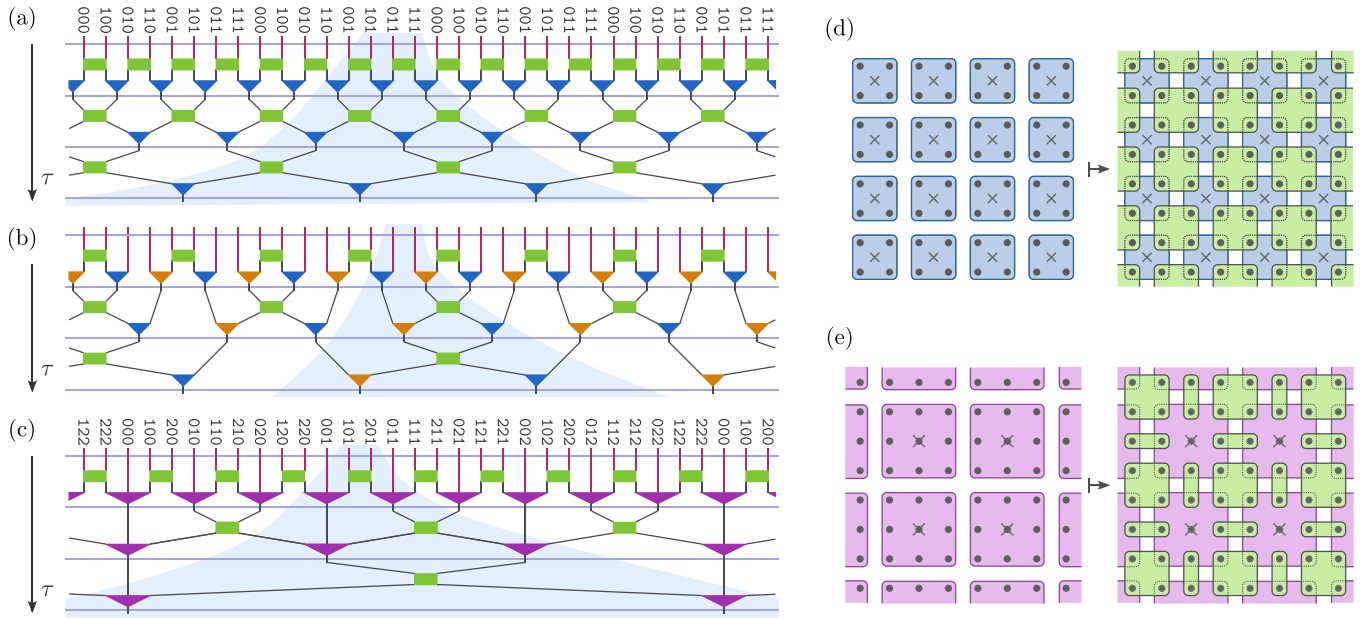


FIG. 4. 1D and 2D MERA tensor networks. Panels (a)–(c) each show three layers of a 1D MERA, where the preparation direction (decreasing τ) is upwards and causal cones for local operators are indicated: (a) the 1D binary MERA, (b) a modified 1D binary MERA [34], and (c) the 1D ternary MERA. Panels (d) and (e) each show one layer of a 2D MERA, where we progress in the preparation direction (decreasing τ) from left to right. (d) For the shown 2D $2 \times 2 \mapsto 1$ MERA [37], isometries map renormalized sites (crosses) into blocks of 2×2 sites (dots) before disentanglers are applied to shifted 2×2 site blocks. (e) For the shown 2D $3 \times 3 \mapsto 1$ MERA [68], isometries map a renormalized sites (crosses) into blocks of 3×3 sites (dots) before one applies 4-site and two 2-site disentanglers.

structure of TMERA tensors unless one is willing to introduce approximations.

2. Quantum computation time complexity

For the hybrid quantum-classical TMERA algorithm, the width of the causal cone is not as decisive for the computation costs. It determines primarily the number of qubits that need to be simultaneously in the interaction zone of the computer. Also, the specific network structure inside the cone, which influences the classical contraction costs, does not affect the scaling of the quantum computation costs. As discussed in

Sec. III, the quantum costs for evaluating the TMERA expectation value of a local interaction term \hat{h}_i is proportional to tT . For translation-invariant systems and homogeneous TMERA, the expectation value of the entire Hamiltonian $\hat{H} = \sum_i \hat{h}_i$ can in fact be measured in one go, i.e., with time $O(tT)$. As there are $\mathcal{N} = O(qtT)$ different Trotter gates, one needs $O(q(tT)^2)$ time to measure the energy gradient.

For heterogeneous TMERAs, there are $\mathcal{N} = O(qtN)$ different Trotter gates with $O(qtN/b^\tau)$ located in layer τ . To measure the gradient with respect to one Trotter gate in layer τ , we need to measure the $O(b^\tau)$ expectation values for

TABLE I. Computational complexity. For the five 1D and 2D MERA network structures shown in Fig. 4, this table states the number of qubits needed for the variational quantum eigensolver and compares the time complexity for the quantum measurement of energy gradients to the evaluation in the classical algorithms. Columns 2 and 3 show the branching ratio b and the cross section A of the causal cone for local operators like the considered Hamiltonian interaction terms. Columns 4 and 5 show the number of qubits needed in the quantum processor register and the number of additional auxiliary qubits (per layer) needed to avoid circuit reprogramming in the gradient evaluations for homogeneous MERAs. Columns 6–8 concern homogeneous MERAs. Column 6 shows the classical computation times for fMERAs, which agree with those for TMERAs as, on a classical computer, the Trotter structure of tensors does not admit substantial gains. Columns 7 and 8 show the quantum computation times for TMERAs and fMERAs. Columns 9 and 10 show the classical and the quantum computation times for heterogeneous TMERAs.

MERA type	Properties		Number of qubits		Times, homogeneous			Times, heterog.	
	b	A	Register	Auxiliary	Classical	Quantum	F-quantum	Classical	Quantum
1D binary	2	3	$4q$	1	$O(2^{9q}T)$	$O(q(tT)^2)$	$O(2^{8q}T^2)$	$O(2^{9q}N)$	$O(q(tT)^2N)$
1D mod. binary	2	2	$3q$	2	$O(2^{7q}T)$	$O(q(tT)^2)$	$O(2^{8q}T^2)$	$O(2^{7q}N)$	$O(q(tT)^2N)$
1D ternary	3	2	$4q$	2	$O(2^{8q}T)$	$O(q(tT)^2)$	$O(2^{8q}T^2)$	$O(2^{8q}N)$	$O(q(tT)^2N)$
2D $2 \times 2 \mapsto 1$	4	3×3	$14q$	2	$O(2^{28q}T)$	$O(q(tT)^2)$	$O(2^{16q}T^2)$	$O(2^{28q}N)$	$O(q(tT)^2N)$
2D $3 \times 3 \mapsto 1$	9	2×2	$15q$	4	$O(2^{16q}T)$	$O(q(tT)^2)$	$O(2^{20q}T^2)$	$O(2^{16q}N)$	$O(q(tT)^2N)$

all local interactions terms that it affects. Hence, one needs $\sum_{\tau=1}^T O(tT b^\tau qtN/b^\tau) = O(q(tT)^2N)$ time for the gradient measurement of heterogeneous TMERAs.

The comparison to the classical computation costs (Appendix A 1) is not trivial. In particular, one needs to take into account how many measurement samples are needed to reach a certain accuracy ϵ . As described in Sec. VIII, the comparison can be done by expressing the bond dimension χ as well as qt^2 in terms of ϵ . For critical models, they are related by power laws. The comparison for a specific model is discussed in Appendix A 5.

The classical component of the hybrid TMERA eigensolver controls the gradient evaluation and steers the gradient-based energy minimization. Its time complexity is always subleading to the quantum time complexity: For a TMERA with \mathcal{N} different Trotter gates, the classical component operates on a vector space of dimension $O(\mathcal{N})$. The classical time complexity for one iteration of gradient descent or L-BFGS is just $O(\mathcal{N})$.

Although it is much less efficient than TMERAs, one can also optimize fMERAs using a quantum computer. A disentangler that acts on n renormalized sites can be decomposed exactly into a circuit of $O(2^{2nq})$ single-qubit and CNOT gates [58–60]. An exact representation of an isometry that maps n sites into $m > n$ requires $O(2^{(n+m)q})$ single-qubit and CNOT gates [61]. The time cost for the energy gradient measurement of a homogeneous fMERA is then obtained by squaring the largest number of gates per tensor and multiplying by T^2 .

3. Quantum space complexity with resets

Concerning the quantum space complexity, we have to distinguish two settings. In the first, qubits corresponding to renormalized sites that leave the causal cone are reset to the reference state $|0\rangle$ in order to reuse them for the implementation of isometries in the next layer transition $\tau \mapsto \tau - 1$. The distribution of measurement results in the evaluation of energy gradients and local observables like $\langle \Psi_i | \hat{h}_i | \Psi_i \rangle$ is the same with and without such midcircuit resets. Qubit rests can be implemented through a projective measurement followed by a subsequent π rotation conditioned on the measurement result or through driven-dissipative reset schemes [89–93].

The number of qubits needed in the register is primarily determined by the cross section A of the causal cone. More precisely, from the sequence of contractions inside each layer of the causal cone, one can determine how many renormalized sites (groups of q qubits) are needed at any point in time. One can reduce this number by not insisting on parallel execution of gates, but shifting the MERA tensors of a layer-transition map temporally so that some qubits can already be reset before the application of further gates. The results for the considered MERA networks are shown in Table I.

As discussed in Sec. V, a major computation time reduction for the homogeneous MERAs is achieved by taking appropriate convex combinations of the layer-transition maps like \hat{U}_L and \hat{U}_R such that the energy density for the full (infinite) system is obtained in one go. This can be done either with classical randomness or, if one wants to avoid the corresponding reprogramming of pulse sequences in the experimental evaluation of gradients, by introducing auxiliary

qubits. In the latter approach, the number of required auxiliary qubits per MERA layer is proportional to \log_2 of the number of transition maps. The auxiliary qubits are also amenable to qubit rests.

4. Quantum space complexity without resets

The experimental evaluation of observables and energy gradients can be made more time efficient using quantum amplitude estimation (QAE) [94,95]. For the preparation $|\Psi_i\rangle = \hat{U}_i|0, \dots, 0\rangle$ of a causal-cone state, QAE requires application of powers $(\hat{U}_i)^m$. In this approach, one cannot employ the resetting of qubits that leave the causal cone, because subsequent factors \hat{U}_i will again act on them. Without the resets, every layer transition $\tau \mapsto \tau - 1$, requires $(b - 1)Aq$ new qubits, initialized in the reference state $|0\rangle$ in order to realize the isometries. For a TMERA with T layers, one then needs a total of

$$\sim (b - 1)AqT \tag{A1}$$

qubits, i.e., a number that grows logarithmically in the total system size N . In ion-trap systems, one can use shuttling [52–54] to move currently used qubits in and out of the quantum register.

5. Comparison of time complexities for critical spin-1 chains

To directly compare the time complexities of the classical fMERA and the quantum-classical TMERA algorithms, one has to take into account the number of measurement samples needed in each optimization step of the TMERA algorithm. As discussed in Sec. VIII, the required number of samples scales with the energy accuracy ϵ . Employing QAE [94,95], the total time complexity per iteration is $O(t^2q \log(1/\epsilon)/\epsilon)$. The time complexity per iteration for a classical fMERA simulation is $O(\chi^r) = O(2^{rq})$, where the exponent r depends on the MERA type as shown in Table I.

Figure 5 provides numerical results for the bilinear-biquadratic spin-1 chain [96–100]

$$\hat{H} = \sum_i [\cos \vartheta (\hat{S}_i \cdot \hat{S}_{i+1}) + \sin \vartheta (\hat{S}_i \cdot \hat{S}_{i+1})^2] \tag{A2}$$

at the critical Uimin-Lai-Sutherland point $\vartheta = \pi/4$. We choose this model because it corresponds to a conformal field theory with central charge $c = 2$ [107] and, hence, features significant entanglement. The figure shows the energy accuracy ϵ as a function of the classical fMERA and quantum-classical TMERA computation cost per iteration. In both cases, the energy accuracy follows a power law with a moderate polynomial advantage for TMERAs. For the TMERA curves with fixed q , ϵ decreases until reaching the accuracy of the fMERA with bond dimension $\chi = 2^q$. In practical simulations, one should hence increase q upon approaching the saturation, in order to follow the power-law decay as indicated by the dashed lines. The fMERA computation cost can be written as $O(\chi^r) = O(\epsilon^{-r/\beta})$, where the factor $1/\beta$ in the exponent captures the relation between bond dimension χ and accuracy ϵ . As the exponent r is very large ($r \geq 16$) for 2D MERAs, the quantum-classical TMERA algorithm will substantially outperform the classical fMERA simulations for models in ≥ 2 spatial dimensions.

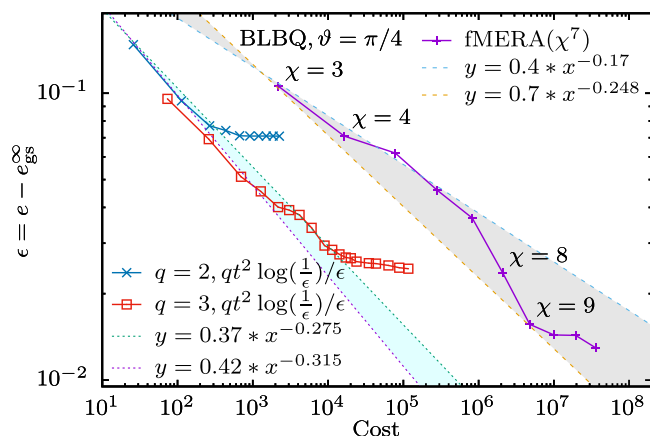


FIG. 5. Comparison of classical and VQE time complexities. Double-logarithmic plot of the energy accuracy ϵ as a function of the total time complexity per iteration for the quantum-classical TMERA algorithm (lower-left legend) and the corresponding classical fMERA simulation (upper-right legend) applied to the critical bilinear-biquadratic spin-1 model (A2). We use homogeneous modified binary MERAs with $T = 6$ layers. For the TMERA curves, the number of Trotter steps t is gradually increased at fixed $q = 2$ and 3 , respectively. For the fMERA curve, the bond dimension $\chi = 2^q$ increases by one from point to point. In both cases, ϵ and the (optimal) cost are related by a power law as indicated by dashed lines. The data suggests a polynomial advantage of the TMERA algorithm.

APPENDIX B: LAYER-TRANSITION MAPS FOR HOMOGENEOUS MERA

Let us discuss in more detail, how classical sampling or auxiliary qubits can be employed to realize convex combinations of layer-transition maps in the preparation. In this way, the spatially averaged A -site density matrices of homogeneous

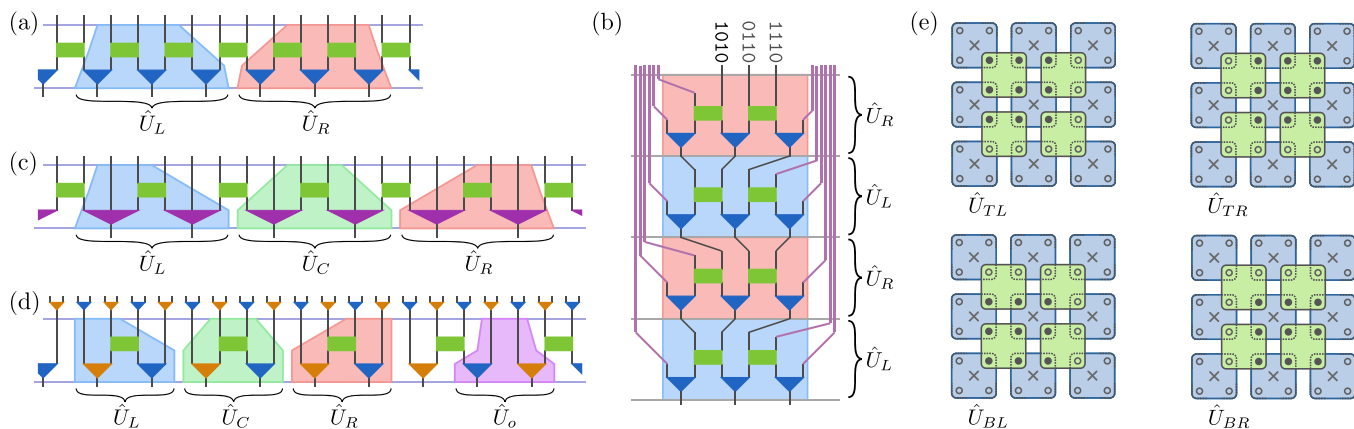


FIG. 6. Layer transition maps. (a) The 1D binary MERA has causal-cone width $A = 3$ and two layer-transition maps \hat{U}_L and \hat{U}_R . Panel (b) gives an example for the sequence of transition maps in a particular causal cone, here in correspondence with Fig. 4(a). For the binary MERA, the sequence can be determined from the binary representation of the index (four least significant digits shown) with 0 corresponding to \hat{U}_L and 1 to \hat{U}_R . (c) The 1D ternary MERA has $A = 2$ and three transition maps \hat{U}_L , \hat{U}_C , and \hat{U}_R . (d) The 1D modified binary MERA has $A = 2$, and one needs to distinguish even and odd bonds. Transition maps \hat{U}_L and \hat{U}_R map from even to odd bonds, \hat{U}_C from an even to an even bond, and \hat{U}_o from an odd to an even bond. (e) The 2D $2 \times 2 \mapsto 1$ MERA [37] has $A = 3 \times 3$ and four layer-transition maps \hat{U}_{TL} , \hat{U}_{TR} , \hat{U}_{BL} , and \hat{U}_{BR} . The renormalized causal-cone sites $\mathcal{A}_i^{(\tau)}$ that the maps act on are indicated by crosses, the sites that leave the causal cone during or after the map are indicated by empty circles, and the sites $\mathcal{A}_i^{(\tau-1)}$ that remain in the causal cone after the preparation step are indicated by filled circles. The layer-transition maps differ in the positions of the latter.

TMERAs can be prepared on the quantum computer in $O(tT)$ time. Hence, the energy density or a component of the energy gradient can be obtained in $O(tT)$ time.

1. 1D binary MERA

This case was already discussed shortly in Sec. V. For the 1D binary MERAs, local operators have causal cones of width $A = 3$. The sequence of layer-transition maps \hat{U}_L and \hat{U}_R for the causal cone of sites

$$\mathcal{A}_i := \{i, \dots, i + A - 1\} \quad (\text{B1})$$

on the physical lattice can be deduced from the binary representation $i_1 i_2 i_3 \dots$ of i with $i_\tau \in \{0, 1\}$, where i_1 is the least significant bit. For the choice shown in Figs. 4(a) and 6(a), $i_\tau = 0$ or 1 means that we progress from layer τ to layer $\tau - 1$ by applying transition map $\hat{U}_L \equiv \hat{U}_0$ or $\hat{U}_R \equiv \hat{U}_1$, respectively. Specifically, the causal-cone state for sites \mathcal{A}_i reads

$$|\Psi_i\rangle = \hat{U}_{i_1} \dots \hat{U}_{i_\tau} (|T\rangle^{\otimes A}), \quad (\text{B2})$$

where state $|T\rangle$ for a single renormalized site defines the so-called top tensor of the MERAs. Here and in the following, we do not explicitly denote the τ dependence of the layer-transition maps and use the convention that operators act on the A active sites in the causal cone and the remaining inactive sites are left untouched as indicated in Fig. 6(b).

For quantities like the energy density, we wish to evaluate the spatial average $\hat{\rho} = \frac{1}{N} \sum_i \hat{\rho}_i$ of the reduced density matrices $\hat{\rho}_i := \text{Tr}_{\mathcal{A}_i^c} |\Psi_i\rangle \langle \Psi_i|$ for the blocks of sites \mathcal{A}_i . The partial trace is over the inactive sites (outside the causal cone) and will not be denoted explicitly in the following. Starting from $\hat{\rho}^{(T)} := (|T\rangle \langle T|)^{\otimes A}$, we can recursively define

$$\hat{\rho}^{(\tau-1)} = \frac{1}{2} \sum_{i_\tau} \hat{U}_{i_\tau} \hat{\rho}^{(\tau)} \hat{U}_{i_\tau}^\dagger \quad (\text{B3a})$$

such that

$$\hat{\rho}^{(0)} = \frac{1}{2^T} \sum_{i_1, \dots, i_T} \hat{U}_{i_1} \dots \hat{U}_{i_T} \hat{\rho}^{(T)} \hat{U}_{i_T}^\dagger \dots \hat{U}_{i_1}^\dagger = \hat{\rho}. \quad (\text{B3b})$$

\hat{U}_L and \hat{U}_R are related to each other by a site permutation.

Hence, the quantum channel (B3a) can be implemented through a Stinespring dilation [85,86] with a single auxiliary qubit per layer. The auxiliary qubit is initialized in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ and used for a controlled site permutation such that, after tracing out the auxiliary qubit, one obtains the channel (B3a).

Alternatively, one can use classical random numbers. Selecting in each layer transition \hat{U}_L or \hat{U}_R with equal probability also implements the channel (B3a). Mathematically, this scheme is completely equivalent to the approach using auxiliary qubits and results in the same distribution of measurement results in the evaluation of observables and energy gradients: Whether one prepares states $|\Psi_1\rangle$ and $|\Psi_2\rangle$ with 50:50 probability or prepares $(|\Psi_1\rangle \otimes |0\rangle + |\Psi_2\rangle \otimes |1\rangle)/\sqrt{2}$ with an auxiliary qubit, the probability to observe outcome i associated with projector \hat{P}_i is in both cases $(\langle \Psi_1 | \hat{P}_i | \Psi_1 \rangle + \langle \Psi_2 | \hat{P}_i | \Psi_2 \rangle)/2$. Experimentally, the measurements need to be repeated until reaching a required accuracy ϵ . Practically, the approach using classical sampling involves a reprogramming of the experimental pulse generator hardware for each new sequence of transition maps. If this reprogramming is slow, the scheme with auxiliary qubits may be preferable as it employs the same quantum circuit in every measurement iteration.

2. 1D ternary MERA

We can proceed along the same lines for the 1D ternary MERAs. The only modifications are due to, now, having the three transition maps $\hat{U}_L \equiv \hat{U}_0$, $\hat{U}_C \equiv \hat{U}_1$, and $\hat{U}_R \equiv \hat{U}_2$ shown in Fig. 6(c). Correspondingly, we now use the ternary representation $i_1 i_2 i_3 \dots$ of the site index i with $i_\tau \in \{0, 1, 2\}$ and replace the prefactors 1/2 in Eqs. (B3) by 1/3. For the Stinespring dilation, we now require two auxiliary qubits per layer, which are, for example, initialized in the state $(|0, 0\rangle + |1, 0\rangle + |0, 1\rangle)/\sqrt{3}$ to perform the controlled site permutations that relate the three unitaries \hat{U}_{i_τ} .

3. 2D 2 x 2 to 1 MERA

Also the 2D $2 \times 2 \mapsto 1$ MERAs can be treated similarly. Now we use binary representations $x_1 x_2 x_3 \dots$ and $y_1 y_2 y_3 \dots$ for the x and y coordinates of site i on the square lattice with $x_\tau, y_\tau \in \{0, 1\}$. There are four layer-transition maps $\hat{U}_{TL} \equiv \hat{U}_{0,0}$, $\hat{U}_{TR} \equiv \hat{U}_{1,0}$, $\hat{U}_{BL} \equiv \hat{U}_{0,1}$, $\hat{U}_{BR} \equiv \hat{U}_{1,1}$, where \hat{U}_{TL} corresponds to the causal cone continuing along the inner (top-left corner) 3×3 square indicated in Fig. 6(e). The layer-transition channel reads

$$\hat{\rho}^{(\tau-1)} = \frac{1}{4} \sum_{x_\tau, y_\tau} \hat{U}_{x_\tau, y_\tau} \hat{\rho}^{(\tau)} \hat{U}_{x_\tau, y_\tau}^\dagger. \quad (\text{B4})$$

Its Stinespring dilation can be implemented with two auxiliary qubits, each initialized in $(|0\rangle + |1\rangle)/\sqrt{2}$, to realize the controlled site permutations that relate the maps \hat{U}_{x_τ, y_τ} to each

other. The first auxiliary qubit can be used to permute rows of sites and the second to permute columns of sites on the inner 4×4 square in Fig. 6(e).

4. 1D modified binary MERA

The cases of the 1D modified binary MERAs and the 2D $3 \times 3 \mapsto 1$ MERAs are a bit more involved.

First, note that the layer-transition channel (B3a) for the (unmodified) 1D binary MERAs can also be explained as follows: Let $\mathcal{A}_i^{(\tau)}$ denote the block of renormalized sites $\{i, \dots, i + A - 1\}$ after τ renormalization steps, i.e., at the interface of layers τ and $\tau + 1$; $\mathcal{A}_i^{(0)}$ agrees with the block \mathcal{A}_i of physical sites in Eq. (B1). Let $\hat{\rho}_i^{(\tau)}$ denote the reduced density matrix for $\mathcal{A}_i^{(\tau)}$, constructed from all MERA tensors in the causal cone of $\mathcal{A}_i^{(\tau)}$, which ranges from layer $\tau + 1$ to the final layer T . Given all density matrices $\hat{\rho}_i^{(\tau)}$ for layer τ , we want to compute those for layer $\tau - 1$. The density matrices $\hat{\rho}_j^{(\tau-1)}$ and $\hat{\rho}_{j+1}^{(\tau-1)}$ for the two blocks $\mathcal{A}_j^{(\tau-1)}$ and $\mathcal{A}_{j+1}^{(\tau-1)}$ whose causal cones contain $\mathcal{A}_i^{(\tau)}$ are obtained by conjugating $\hat{\rho}_i^{(\tau)}$ with the appropriate layer- τ transition maps

$$\hat{\rho}_j^{(\tau-1)} = \hat{U}_L \hat{\rho}_i^{(\tau)} \hat{U}_L^\dagger \quad \text{and} \quad \hat{\rho}_{j+1}^{(\tau-1)} = \hat{U}_R \hat{\rho}_i^{(\tau)} \hat{U}_R^\dagger. \quad (\text{B5})$$

Now the spatially averaged density matrix $\hat{\rho}^{(\tau-1)} = \frac{b^{\tau-1}}{N} \sum_j \hat{\rho}_j^{(\tau-1)}$ is obtained by averaging over all sites i in layer τ and the two corresponding sites in layer $\tau - 1$. Thus, we recover Eq. (B3a).

The 1D modified binary MERA with $A = 2$ can be addressed similarly. In this case, not all bonds are equivalent. Depending on whether we are on an even or an odd bond, we apply either one of the three layer-transition maps \hat{U}_L , \hat{U}_C , \hat{U}_R or the map \hat{U}_o , respectively, as shown in Fig. 6(d). While \hat{U}_L and \hat{U}_R move the causal cone from an even to an odd bond, \hat{U}_C maps from an even to an even bond, and \hat{U}_o maps from an odd to an even bond. Let $\hat{\rho}_e^{(\tau)}$ denote the bond density matrix averaged over all even bonds for layer τ and $\hat{\rho}_o^{(\tau)}$ the average over all odd bonds. Then the same considerations as above lead to the layer transitions

$$\hat{\rho}_e^{(\tau-1)} = \frac{1}{2} (\hat{U}_C \hat{\rho}_e^{(\tau)} \hat{U}_C^\dagger + \hat{U}_o \hat{\rho}_o^{(\tau)} \hat{U}_o^\dagger), \quad (\text{B6a})$$

$$\hat{\rho}_o^{(\tau-1)} = \frac{1}{2} (\hat{U}_L \hat{\rho}_e^{(\tau)} \hat{U}_L^\dagger + \hat{U}_R \hat{\rho}_o^{(\tau)} \hat{U}_R^\dagger). \quad (\text{B6b})$$

For the evaluation of energy densities and gradients we want to prepare the spatially averaged density matrices $\hat{\rho}_e^{(0)}$ and $\hat{\rho}_o^{(0)}$ for even and odd bonds of the physical lattice.

To realize this in the approach using auxiliary qubits, we can first introduce an additional flag qubit in the register of the quantum computer which indicates whether we are on an even or odd bond, i.e., progressing in the preparation direction, we want to prepare the states

$$\hat{\rho}^{(\tau)} := \frac{1}{2} (|0\rangle\langle 0|_f \otimes \hat{\rho}_e^{(\tau)} + |1\rangle\langle 1|_f \otimes \hat{\rho}_o^{(\tau)}). \quad (\text{B7})$$

For these, the layer-transition channel reads

$$\hat{\rho}^{(\tau-1)} = \frac{1}{2} \sum_{k=1}^4 \hat{L}_k \hat{\rho}^{(\tau)} \hat{L}_k^\dagger \quad \text{with}$$

$$\begin{aligned} \hat{L}_1 &= |0\rangle\langle 0|_f \otimes \hat{U}_C, & \hat{L}_2 &= |0\rangle\langle 1|_f \otimes \hat{U}_o, \\ \hat{L}_3 &= |1\rangle\langle 0|_f \otimes \hat{U}_L, & \hat{L}_4 &= |1\rangle\langle 0|_f \otimes \hat{U}_R. \end{aligned} \quad (\text{B8})$$

It can be implemented by a Stinespring dilation that employs two auxiliary qubits per layer as specified in the second row of Table I.

APPENDIX C: GRADIENT EVALUATION AND RIEMANNIAN OPTIMIZATION

The goal of the variational quantum eigensolver is to minimize the energy expectation value $E = \langle \Psi | \hat{H} | \Psi \rangle$ over a TMERA variety $\{|\Psi\rangle\}$, where the TMERA is characterized by the network structure, bond dimensions, and the tensor Trotterization as previously discussed. This minimization can be carried out by evaluating energy gradients and employing them in gradient descent methods or, preferably, quasi-Newton methods like the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm [74,75].

1. Gradients in the CNOT and CAN parametrizations

In the CNOT and CAN parametrizations, the Trotter gates are expressed in terms of single- and two-qubit rotations $\hat{R}_\delta(\theta) = e^{-i\theta\hat{\sigma}}/2$ [Eq. (1)]. The rotation angles θ parametrize the TMERA variety. To compute the energy derivative for one of these angles, we can write the energy expectation value in the form

$$E(\theta) = \text{Tr}(\hat{A}[\hat{R}_\delta^\dagger(\theta) \otimes \mathbb{1}_\perp] \hat{B}[\hat{R}_\delta(\theta) \otimes \mathbb{1}_\perp]), \quad (\text{C1})$$

where the Hermitian operators \hat{A} and \hat{B} comprise the remaining tensors of $\langle \Psi |$, $|\Psi \rangle$, and the Hamiltonian. For brevity of notation, in the following we will drop the “ $\otimes \mathbb{1}_\perp$ ”, which indicates that $\hat{R}_\delta(\theta)$ acts only on a subspace corresponding to one or two qubits. The derivative is

$$\partial_\theta E(\theta) = \frac{i}{2} \text{Tr}(\hat{A}[\hat{R}_\delta^\dagger(\theta)[\hat{\sigma}, \hat{B}]\hat{R}_\delta(\theta)). \quad (\text{C2})$$

For the Hermitian and unitary operators $\hat{\sigma}$, $\hat{R}_\delta(\pm\frac{\pi}{2}) = (\mathbb{1} \mp i\hat{\sigma})/\sqrt{2}$ and, hence,

$$i[\hat{\sigma}, \hat{B}] = \hat{R}_\delta^\dagger\left(\frac{\pi}{2}\right)\hat{B}\hat{R}_\delta\left(\frac{\pi}{2}\right) - \hat{R}_\delta^\dagger\left(-\frac{\pi}{2}\right)\hat{B}\hat{R}_\delta\left(-\frac{\pi}{2}\right). \quad (\text{C3})$$

such that Eq. (2) follows. In this way, the gradient can be evaluated on the quantum computer by measuring energy expectation values [71–73].

For homogeneous TMERAs, the same rotation occurs multiple times as tensors are repeated in the translation-invariant MERA layers. Applying the product rule, this just means that the derivative $\partial_\theta E$ will contain one term for each occurrence of the rotation $\hat{R}(\theta)$. This sum can be evaluated efficiently as described in the main text and Appendix B.

With the gradient in hand, one can apply standard implementations of gradient descent or quasi-Newton methods like L-BFGS.

2. An alternative Riemannian version of the optimization

Instead of employing an explicit parametrization of the Trotter gates, one can formulate the problem as a minimization over the manifold

$$\mathcal{M} := \text{U}(4)^{\times \mathcal{N}} \quad (\text{C4})$$

formed by the product of the unitary groups for the \mathcal{N} Trotter gates of the TMERA. For homogeneous TMERAs, repeated gates are counted once. This Riemannian approach turns out to have somewhat improved convergence properties as discussed in Appendix D. Note that one can take unitary gauge freedoms on the inputs of the Trotter gates into account and consider quotient groups $\text{U}(4)/(\text{U}(2) \times \text{U}(2))$. For the simplicity of notation, we stick to the full $\text{U}(4)$ in the following.

For the optimization, we can regard \mathcal{M} as embedded in the Euclidean space

$$\mathcal{E} = \text{End}(\mathbb{C}^4)^{\times \mathcal{N}} \simeq \mathbb{R}^{32\mathcal{N}}, \quad (\text{C5})$$

which is the space of the Trotter gates without the unitarity constraint. Let $\mathbf{u} \in \mathcal{M} \subset \mathcal{E}$ denote the vector that contains the matrix elements of all gates and $E(\mathbf{u}) = \langle \Psi(\mathbf{u}) | \hat{H} | \Psi(\mathbf{u}) \rangle$ the energy functional. To apply gradient-based optimization algorithms in this setting, we need to compute the derivative $\partial_{\mathbf{u}} E(\mathbf{u})$, project it onto the tangent space $\mathcal{T}_{\mathbf{u}}$ of \mathcal{M} at \mathbf{u} to obtain the gradient direction, construct retractions for line search, and vector transport to be able to sum gradient vectors from different points on the manifold. This is the program of Riemannian optimization as discussed generally in Refs. [108,109] and recently demonstrated for MERAs in Refs. [110,111].

In the following, consider a single unitary $\hat{u} \in \text{U}(n)$, with $n = 4$ for the considered Trotter gates, and we employ the Euclidean metric (real part of the Hilbert-Schmidt inner product)

$$\langle \hat{u}, \hat{u}' \rangle := \text{Re Tr}(\hat{u}^\dagger \hat{u}'). \quad (\text{C6})$$

The extension to the product manifold \mathcal{M} is straightforward. As in Eq. (C1), let us write the energy expectation value in the form

$$\begin{aligned} E(\hat{u}) &= \langle \Psi(\hat{u}) | \hat{H} | \Psi(\hat{u}) \rangle \\ &= \text{Tr}(\hat{A}[\hat{u}^\dagger \otimes \mathbb{1}_\perp] \hat{B}[\hat{u} \otimes \mathbb{1}_\perp]), \end{aligned} \quad (\text{C7})$$

where “ $\otimes \mathbb{1}_\perp$ ” indicates that \hat{u} only acts on an n -dimensional subspace. Then the energy gradient in the embedding space $\text{End}(\mathbb{C}^n)$ is

$$\hat{d} = 2 \text{Tr}_\perp(\hat{B}[\hat{u} \otimes \mathbb{1}_\perp] \hat{A}), \quad (\text{C8})$$

where Tr_\perp is the partial trace over the subspace that \hat{u} does not act on. The gradient \hat{d} fulfills $\partial_\varepsilon E(\hat{u} + \varepsilon \hat{w})|_{\varepsilon=0} = \langle \hat{d}, \hat{w} \rangle$ for all \hat{w} . An element \hat{w} of the tangent space $\mathcal{T}_{\hat{u}}$ for $\text{U}(n)$ at \hat{u} needs to obey $(\hat{u} + \varepsilon \hat{w})^\dagger (\hat{u} + \varepsilon \hat{w}) = \mathbb{1} + O(\varepsilon^2)$, i.e., $\hat{u}^\dagger \hat{w} + \hat{w}^\dagger \hat{u} = 0$. So $\hat{u}^\dagger \hat{w}$ needs to be skew-Hermitian and, hence,

$$\mathcal{T}_{\hat{u}} = \{i\hat{u}\hat{\eta} \mid \hat{\eta} = \hat{\eta}^\dagger \in \text{End}(\mathbb{C}^n)\}. \quad (\text{C9})$$

The Riemannian energy gradient \hat{g} for the manifold $\text{U}(n)$ at \hat{u} is obtained by projecting \hat{d} onto the tangent space such that $\langle \hat{w}, \hat{g} \rangle = \langle \hat{w}, \hat{d} \rangle$ for all $\hat{w} \in \mathcal{T}_{\hat{u}}$. This gives

$$\hat{g} = (\hat{d} - \hat{u} \hat{d}^\dagger \hat{u})/2 \in \mathcal{T}_{\hat{u}}. \quad (\text{C10})$$

For a line search on the manifold, we need a retraction, i.e., a curve $\hat{r}_{\hat{u},\hat{p}}(\tau)$ on the manifold that starts from $\hat{u} = \hat{r}_{\hat{u},\hat{p}}(0)$ in direction $\hat{p} = \partial_\tau \hat{r}_{\hat{u},\hat{p}}(\tau)|_{\tau=0} \in \mathcal{T}_{\hat{u}}$. We use

$$\hat{r}_{\hat{u},\hat{p}}(\tau) := e^{\tau \hat{p} \hat{u}} \hat{u} \in \mathcal{U}(n). \quad (\text{C11})$$

For quasi-Newton methods, we also need to compute differences of Riemannian gradients from different points on the manifold, specifically, for two points on a retraction. This is accomplished by vector transport, i.e., a map between the two corresponding tangent spaces. We use

$$\hat{T}_{\hat{u},\hat{p}}(\tau) \hat{w} := e^{\tau \hat{p} \hat{u}} \hat{w} \quad \text{for } \hat{w} \in \mathcal{T}_{\hat{u}}, \tau \in \mathbb{R}. \quad (\text{C12})$$

This gives an element of the tangent space at $\hat{r}_{\hat{u},\hat{p}}(\tau)$ and $\hat{T}_{\hat{u},0} \hat{w} = \hat{w}$. The vector transport is isometric in the sense that $(\hat{T}_{\hat{u},\hat{p}}(\tau) \hat{w}, \hat{T}_{\hat{u},\hat{p}}(\tau) \hat{w}') = (\hat{w}, \hat{w}')$ for all $\hat{p}, \hat{w}, \hat{w}' \in \mathcal{T}_{\hat{u}}$ and $\tau \in \mathbb{R}$.

In analogy to Eq. (2), the Riemannian gradient (C10) can be obtained on the quantum computer by measuring energies of TMERAs where one Trotter gate is modified: The tangent space $\mathcal{T}_{\hat{u}}$ is a real n^2 -dimensional vector space, and we can choose a basis $\{\hat{u} \hat{\sigma}_j | j = 1, \dots, n^2\}$ with Hermitian unitaries $\hat{\sigma}_j$, i.e., $\hat{\sigma}_j = \hat{\sigma}_j^\dagger$ and $\hat{\sigma}_j^2 = \mathbb{1}$. With $(\hat{\sigma}_j, \hat{\sigma}_k) = \delta_{j,k}$, we can expand the energy gradient in the form $\hat{g} = i \sum_{j=1}^{n^2} \alpha_j \hat{u} \hat{\sigma}_j / n$ and evaluate the expansion coefficients α_j using the energy expectation values (C7),

$$\begin{aligned} \alpha_j &= (\hat{g}, i \hat{u} \hat{\sigma}_j) = (\hat{d}, i \hat{u} \hat{\sigma}_j) = \text{Re Tr}(i \hat{d}^\dagger \hat{u} \hat{\sigma}_j) \\ &= -i \text{Tr}(\hat{A}[\hat{\sigma}_j \otimes \mathbb{1}_\perp, (\hat{u}^\dagger \otimes \mathbb{1}_\perp) \hat{B}(\hat{u} \otimes \mathbb{1}_\perp)]) \\ &= E\left(\hat{u} \hat{R}_{\hat{\sigma}_j}\left(-\frac{\pi}{2}\right)\right) - E\left(\hat{u} \hat{R}_{\hat{\sigma}_j}\left(\frac{\pi}{2}\right)\right). \end{aligned} \quad (\text{C13})$$

Algorithm 1. Riemannian version of the L-BFGS algorithm to minimize (T)MERA energies, adapted from Ref. [108]. The optimization is started at a point $\mathbf{u}_0 \in \mathcal{M}$ and stopped when the gradient norm falls below ε . For the approximation \tilde{H}_k of the inverse Hessian, ℓ vector pairs $(\mathbf{s}_i, \mathbf{y}_i)$ with $i = k - \ell, \dots, k - 1$ are kept in memory. The constants c_1 and c_2 enter the Wolfe conditions.

```

1:  $\mathbf{u}_0 \in \mathcal{M}, \varepsilon > 0, \ell \in \mathbb{N}, 0 < c_1 < \frac{1}{2} < c_2 < 1, \gamma = 1, k = 0, m = 0$  ▷ Initialization.
2: while  $\|\mathbf{g}_k\| > \varepsilon$  do ▷ Riemannian gradient  $\mathbf{g}_k$  at  $\mathbf{u}_k$  determined according to Eq. (C10).
3:    $\mathbf{p}_k = -\mathbf{g}_k$  ▷ Lines 3-11 determine the search direction  $\mathbf{p}_k = -\tilde{H}_k \mathbf{g}_k$  with Eq. (C14).
4:   for  $i = k - 1, k - 2, \dots, m$  do ▷ Compute  $-V_m^\dagger V_{m+1}^\dagger \dots V_{k-1}^\dagger \mathbf{g}_k$  and  $\xi_i := -(\mathbf{s}_i, V_i^\dagger \dots V_{k-1}^\dagger \mathbf{g}_k)$ .
5:      $\xi_i \leftarrow \rho_i(\mathbf{s}_i, \mathbf{p}_k)$ 
6:      $\mathbf{p}_k \leftarrow \mathbf{p}_k - \xi_i \mathbf{y}_i$ 
7:   end for
8:    $\mathbf{p}_k \leftarrow \gamma \mathbf{p}_k$  ▷ Multiply with  $\tilde{H}_{m-1} := \gamma \mathbb{1}$ .
9:   for  $i = m, m + 1, \dots, k - 1$  do ▷ Multiply with  $V_{k-1} \dots V_m$  and insert  $\rho_i \mathbf{s}_i \mathbf{s}_i^\dagger$  terms.
10:     $\mathbf{p}_k \leftarrow \mathbf{p}_k - \rho_i \mathbf{s}_i (\mathbf{y}_i, \mathbf{p}_k) + \mathbf{s}_i \xi_i$ 
11:   end for ▷ Now,  $\mathbf{p}_k = -\tilde{H}_k \mathbf{g}_k$ .
12:   Do line search to find  $\tau_k \in \mathbb{R}$  that satisfies the Wolfe conditions. With  $\mathbf{u}_{k+1} := \mathbf{r}_{\mathbf{u}_k, \mathbf{p}_k}(\tau_k)$ :
13:      $E(\mathbf{r}_{\mathbf{u}_k, \mathbf{p}_k}(\tau_k)) \leq E(\mathbf{u}_k) + c_1 \tau_k (\mathbf{g}_k, \mathbf{p}_k),$ 
14:      $\partial_\tau E(\mathbf{r}_{\mathbf{u}_k, \mathbf{p}_k}(\tau))|_{\tau=\tau_k} \geq c_2 \partial_\tau E(\mathbf{r}_{\mathbf{u}_k, \mathbf{p}_k}(\tau))|_{\tau=0}.$ 
15:      $\mathbf{u}_{k+1} = \mathbf{r}_{\mathbf{u}_k, \mathbf{p}_k}(\tau_k)$  ▷ For the following,  $T_k := T_{\mathbf{u}_k, \mathbf{p}_k}(\tau_k)$  from Eq. (C12).
16:      $\mathbf{s}_k = T_k \tau_k \mathbf{p}_k, \mathbf{y}_k = \mathbf{g}_{k+1} - T_k \mathbf{g}_k, \rho_k = 1/(\mathbf{s}_k, \mathbf{y}_k), \gamma \leftarrow (\mathbf{s}_k, \mathbf{y}_k) / \|\mathbf{y}_k\|^2$ 
17:      $m \leftarrow \max(k - \ell, 0)$ 
18:     Discard  $(\rho_{m-1}, \mathbf{s}_{m-1}, \mathbf{y}_{m-1})$  from memory if  $m > 0$ .
19:     for  $i = m, m + 1, \dots, k - 1$  do ▷ Transport  $\{\mathbf{s}_i\}$  and  $\{\mathbf{y}_i\}$  from  $\mathcal{T}_{\mathbf{u}_k}$  to  $\mathcal{T}_{\mathbf{u}_{k+1}}$ .
20:        $\mathbf{s}_i \leftarrow T_k \mathbf{s}_i, \mathbf{y}_i \leftarrow T_k \mathbf{y}_i$ 
21:     end for
22:      $k \leftarrow k + 1$ 
23: end while

```

For the third line, we have used Eq. (C3) and $\hat{R}_{\hat{\sigma}}(\theta) = e^{-i\theta \hat{\sigma}/2}$.

To minimize TMERA energies, we can employ a Riemannian version of the L-BFGS algorithm. For the following, we return to the global optimization problem on the product manifold (C4) with the embedding space (C5), vectors $\mathbf{u} \in \mathcal{M}$ comprising the matrix elements of all Trotter gates, and the Euclidean metric $(\mathbf{u}, \mathbf{u}') = \text{Re}(\mathbf{u}^\dagger \mathbf{u}')$. Similarly, gradients \hat{g} , retractions \hat{r} , and so on are now written in a vectorized form. In the Newton method, one generates a sequence of points $\mathbf{u}_1, \mathbf{u}_2, \dots \in \mathcal{M}$ that converges quadratically fast to a minimum of $E(\mathbf{u})$. In each step one obtains a second-order model of $E(\mathbf{u})$ using the Riemannian gradient \mathbf{g}_k and the inverse Hessian H_k at \mathbf{u}_k . The vector $\mathbf{p}_k := -H_k \mathbf{g}_k \in \mathcal{T}_{\mathbf{u}_k}$ that points from \mathbf{u}_k to the minimum of the quadratic model is used for an inexact line search, and one chooses the next point $\mathbf{u}_{k+1} = \mathbf{r}_{\mathbf{u}_k, \mathbf{p}_k}(\tau_k)$ on the corresponding line (retraction curve) with $\tau_k \in \mathbb{R}$ such that the Wolfe conditions are obeyed. The latter require that both the function value and the gradient norm decrease sufficiently, where, in the Riemannian version, one rather considers the energy derivatives in the search direction. The BFGS algorithm [74,108] modifies this procedure, avoiding the costly evaluation of the Hessian. Instead, one updates a positive definite approximation $\tilde{H}_k \in \text{End}(\mathcal{T}_{\mathbf{u}_k})$ of the inverse Hessian. \tilde{H}_{k+1} is determined by requiring that the gradient of the new quadratic model at \mathbf{u}_{k+1} , evaluated at \mathbf{u}_k , should agree with the actual \mathbf{g}_k . This is equivalent to the secant equation $\mathbf{s}_k = \tilde{H}_{k+1} \mathbf{y}_k$ with $\mathbf{s}_k := T_k \tau_k \mathbf{p}_k$ and the gradient change $\mathbf{y}_k := \mathbf{g}_{k+1} - T_k \mathbf{g}_k$, which are both elements of the tangent space $\mathcal{T}_{\mathbf{u}_{k+1}}$ as $T_k := T_{\mathbf{u}_k, \mathbf{p}_k}(\tau_k)$ denotes the vector transport. From the solution space of the secant equation, one chooses the matrix \tilde{H}_{k+1} that is closest to \tilde{H}_k in a suitable

metric. Specifically, the BFGS update reads

$$\tilde{H}_{k+1} = V_k \tilde{H}_k V_k^\dagger + \rho_k s_k s_k^\dagger \quad (\text{C14})$$

with $V_k := (\mathbb{1} - \rho_k s_k s_k^\dagger) T_k$ and $\rho_k := 1/(\mathbf{y}_k, s_k)$. Finally, the L-BFGS algorithm [74,108] avoids the increasing cost of operating with \tilde{H}_k by keeping the ℓ most recent triples $(\rho_k, s_k, \mathbf{y}_k)$ in memory and computing an approximation of the inverse Hessian from them in each iteration.

The Riemannian L-BFGS algorithm that we employ, adapted from Ref. [108], is shown in Algorithm 1. Note that, for numerical stability, after every retraction (C11), it may be necessary to project the resulting point onto the manifold in order to avoid the accumulation of small numerical errors. For example, this can be done by a singular value decomposition for every Trotter gate and setting all singular values to one. If this is done, the finite numerical precision for the vector transport (C12) should have negligible effects, i.e., need not be corrected.

The (classical) computation costs for each iteration of the Riemannian L-BFGS algorithm are linear in the total number \mathcal{N} of Trotter gates and linear in the number ℓ of retained terms (rank of \tilde{H}_k). The latter does not need to be scaled with the problem size, and we choose $\ell = 9$ in all computations. Retractions (C11), vector transport (C12), and so on can be applied separately for every Trotter gate, each corresponding to $2n^2$ entries of the full \mathbf{u} vector with $n = 4$. The cost is hence indeed $O(n^3 \ell \mathcal{N}) = O(\mathcal{N})$.

APPENDIX D: DIFFERENT TROTTER-GATE PARAMETRIZATIONS AND XX TMERAS

The CNOT and CAN parametrizations for the Trotter gates of the TMERA are equivalent to the parametrization-free representation as $U(4)$ unitaries, and, up to an irrelevant phase factor, one can transform between the parametrizations as described in Refs. [70,83]. This equivalence is tested and




FIG. 7. Optimization and scanning with different parametrizations. Starting from a product state at $g = 1.25$, we scan forth and back on the interval $0.75 \leq g \leq 1.25$, minimizing the MERA energy density e for the 1D transverse-field Ising model with the L-BFGS algorithm. This is done for homogeneous modified binary MERAs with $T = 6$ layers and bond dimension $\chi = 2^q = 8$. The TMERA tensors consist of $t = 2$ Trotter steps, and XX TMERAs are simulated once with two and once with six XX Trotter steps per tensors.

confirmed by optimizing TMERAs for the 1D transverse-field Ising model (4) while scanning forth and back on the parameter interval $g \in [0.75, 1.25]$. Starting at $g = 1.25$, the optimization was initialized by a product state with $|\phi\rangle = \hat{R}_{\hat{\sigma}^z}(\frac{\pi}{4}) \hat{R}_{\hat{\sigma}^y}(\frac{\pi}{4}) \hat{R}_{\hat{\sigma}^z}(\frac{\pi}{4}) |0\rangle$ on every site. Figure 7 shows the accuracies of the energy densities e during the scanning procedure for TMERAs in the CNOT and CAN parametrizations as well as the parametrization-free form (“Riemannian”). The results are compared to the corresponding fMERA optimization. The Euclidean L-BFGS algorithm [74,75] was employed for the TMERA in the CAN and CNOT parametrizations and the Riemannian L-BFGS algorithm, as discussed in Appendix C2, was employed for the parametrization-free TMERA and fMERA. In all cases, the L-BFGS parameters were chosen as $\varepsilon = 10^{-12}$, $\ell = 9$, $c_1 = 0.1$, and $c_2 = 0.9$.

Although different TMERA parametrizations show differing energies in the early stages of the scanning procedure, they quickly converge. Without scanning, the parametrization-free TMERA is somewhat favorable. The explicit parametrizations in terms of rotation angles are more prone to getting stuck in local minima, whereas the local minima and saddle points for the parametrization-free form are entirely due to the structure of the TMERA manifold and the Hamiltonian. For random initial states, the parametrization-free form shows better convergence than the CNOT and CAN forms. As a simple example, note that the product state $|0\rangle^{\otimes N}$ is a stationary point in the energy landscape of the transverse Ising model (4) in the CAN parametrization. The data in Fig. 2 were obtained by Riemannian optimization.

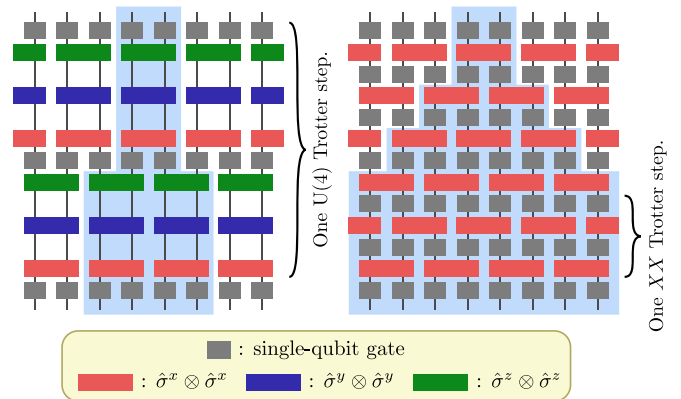


FIG. 8. XX Trotterization. Left: In the TMERAs considered so far, tensors are Trotterized into regular circuits of general two-qubit Trotter gates $\in U(4)$. For 1D systems, one Trotter step consists of such gates applied on all odd bonds and then on all even bonds (or vice versa). In the CAN parametrization, each $U(4)$ Trotter gate is realized by a sequence of two single-qubit gates, followed by three Ising rotations generated by $\hat{\sigma}^\alpha \otimes \hat{\sigma}^\alpha$, and two final single-qubit gates. In every Trotter step, causal cones widen by at most four sites. Right: In the XX TMERA, layers of single-qubit gates alternate with layers of $\hat{\sigma}^x \otimes \hat{\sigma}^x$ Ising gates on even and odd bonds, respectively. One XX Trotter step contains two layers of Ising gates. In terms of the number of Ising gates, one $U(4)$ Trotter step corresponds to three XX Trotter steps. For the latter, causal cones grow correspondingly faster.

For the 1D TMERAs, we chose the Trotter circuit of each tensor to consist of steps with U(4) gates on all odd qubit bonds and all even qubit bonds, alternatingly, as shown in Fig. 1(c) and in the left panel of Fig. 8. The goal of this choice is to admit the generation of entanglement between any pair of qubits in a few Trotter steps. However, in the CNOT and CAN parametrizations, each Trotter gate features three elementary two-qubit gates; CNOT and $\hat{\sigma}^\alpha \otimes \hat{\sigma}^\alpha$ Ising rotations with $\alpha = x, y, z$, respectively. We have explored a different Trotterization approach, where layers of generic single-qubit gates alternate with XX Ising rotations. One XX

Trotter step contains two layers of XX Ising rotations on odd and even bonds, respectively. For the same computational cost, characterized by the total number of two-qubit Ising rotations, causal cones in the XX Trotterization of the MERA tensors grow substantially faster than in the U(4) Trotterization. However, the benchmark simulations in Fig. 7 show no enhanced approximation accuracy. The energies for six XX Trotter steps per MERA tensor converge to approximately the same values as the energies for two U(4) Trotter steps per tensor. Improvements along these lines are a topic for future work.

-
- [1] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* **73**, 33 (2001).
- [2] M. Suzuki, S. Miyashita, and A. Kuroda, Monte Carlo simulation of quantum spin systems. I, *Prog. Theor. Phys.* **58**, 1377 (1977).
- [3] O. F. Syljuåsen and A. W. Sandvik, Quantum Monte Carlo with directed loops, *Phys. Rev. E* **66**, 046701 (2002).
- [4] N. V. Prokof'ev and B. V. Svistunov, Polaron Problem by Diagrammatic Quantum Monte Carlo, *Phys. Rev. Lett.* **81**, 2514 (1998).
- [5] R. J. Baxter, Dimers on a rectangular lattice, *J. Math. Phys.* **9**, 650 (1968).
- [6] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [7] H. Niggemann, A. Klümper, and J. Zittartz, Quantum phase transition in spin-3/2 systems on the hexagonal lattice-optimum ground state approach, *Z. Phys. B* **104**, 103 (1997).
- [8] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066) (2004).
- [9] G. Vidal, Entanglement Renormalization, *Phys. Rev. Lett.* **99**, 220405 (2007).
- [10] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys.* **349**, 117 (2014).
- [11] T. Barthel, C. Pineda, and J. Eisert, Contraction of fermionic operator circuits and the simulation of strongly correlated fermions, *Phys. Rev. A* **80**, 042333 (2009).
- [12] P. Corboz and G. Vidal, Fermionic multiscale entanglement renormalization ansatz, *Phys. Rev. B* **80**, 165129 (2009).
- [13] C. V. Kraus, N. Schuch, F. Verstraete, and J. I. Cirac, Fermionic projected entangled pair states, *Phys. Rev. A* **81**, 052338 (2010).
- [14] P. Corboz, G. Evenbly, F. Verstraete, and G. Vidal, Simulation of interacting fermions with entanglement renormalization, *Phys. Rev. A* **81**, 010303(R) (2010).
- [15] C. Pineda, T. Barthel, and J. Eisert, Unitary circuits for strongly correlated fermions, *Phys. Rev. A* **81**, 050303(R) (2010).
- [16] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Sign problem in the numerical simulation of many-electron systems, *Phys. Rev. B* **41**, 9301 (1990).
- [17] M. Troyer and U.-J. Wiese, Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
- [18] L. Balents, Spin liquids in frustrated magnets, *Nature (London)* **464**, 199 (2010).
- [19] Y. Zhou, K. Kanoda, and T.-K. Ng, Quantum spin liquid states, *Rev. Mod. Phys.* **89**, 025003 (2017).
- [20] Y. Shimizu, K. Miyagawa, K. Kanoda, M. Maesato, and G. Saito, Spin Liquid State in an Organic Mott Insulator with a Triangular Lattice, *Phys. Rev. Lett.* **91**, 107001 (2003).
- [21] F. L. Pratt, P. J. Baker, S. J. Blundell, T. Lancaster, S. Ohira-Kawamura, C. Baines, Y. Shimizu, K. Kanoda, I. Watanabe, and G. Saito, Magnetic and non-magnetic phases of a quantum spin liquid, *Nature (London)* **471**, 612 (2011).
- [22] A. Banerjee, C. A. Bridges, J.-Q. Yan, A. A. Aczel, L. Li, M. B. Stone, G. E. Granroth, M. D. Lumsden, Y. Yiu, J. Knolle *et al.*, Proximate Kitaev quantum spin liquid behaviour in a honeycomb magnet, *Nat. Mater.* **15**, 733 (2016).
- [23] H. L. Stormer, D. C. Tsui, and A. C. Gossard, The fractional quantum Hall effect, *Rev. Mod. Phys.* **71**, S298 (1999).
- [24] R. de-Picciotto, M. Reznikov, M. Heiblum, V. Umansky, G. Bunin, and D. Mahalu, Direct observation of a fractional charge, *Nature (London)* **389**, 162 (1997).
- [25] J. G. Bednorz and K. A. Müller, Possible high- T_c superconductivity in the Ba-La-Cu-O system, *Z. Phys. B* **64**, 189 (1986).
- [26] A. J. Leggett, What DO we know about high T_c ? *Nat. Phys.* **2**, 134 (2006).
- [27] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, *Commun. Math. Phys.* **144**, 443 (1992).
- [28] S. Rommer and S. Östlund, A class of ansatz wave functions for 1D spin systems and their relation to DMRG, *Phys. Rev. B* **55**, 2164 (1997).
- [29] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [30] T. Nishino, K. Okunishi, Y. Hieida, N. Maeshima, and Y. Akutsu, Self-consistent tensor product variational approximation for 3D classical models, *Nucl. Phys. B* **575**, 504 (2000).
- [31] M. A. Martín-Delgado, M. Roncaglia, and G. Sierra, Stripe ansätze from exactly solved models, *Phys. Rev. B* **64**, 075117 (2001).
- [32] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of

- Projected Entangled Pair States, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [33] G. Vidal, Class of Quantum Many-Body States That Can Be Efficiently Simulated, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [34] G. Evenbly and G. Vidal, Quantum criticality with the multiscale entanglement renormalization ansatz, [arXiv:1109.5334](https://arxiv.org/abs/1109.5334) (2011).
- [35] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [36] R. Orús and G. Vidal, Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction, *Phys. Rev. B* **80**, 094403 (2009).
- [37] L. Cincio, J. Dziarmaga, and M. M. Rams, Multiscale Entanglement Renormalization Ansatz in Two Dimensions: Quantum Ising Model, *Phys. Rev. Lett.* **100**, 240603 (2008).
- [38] G. Evenbly and G. Vidal, Entanglement Renormalization in Two Spatial Dimensions, *Phys. Rev. Lett.* **102**, 180406 (2009).
- [39] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [40] T. Barthel, J. Lu, and G. Friesecke, On the closedness and geometry of tensor network state sets, *Lett. Math. Phys.* **112**, 72 (2022).
- [41] T. Barthel and Q. Miao, Absence of barren plateaus and scaling of gradients in the energy optimization of isometric tensor network states, [arXiv:2304.00161](https://arxiv.org/abs/2304.00161) (2023).
- [42] Q. Miao and T. Barthel, Isometric tensor network optimization for extensive Hamiltonians is free of barren plateaus, [arXiv:2304.14320](https://arxiv.org/abs/2304.14320) (2023).
- [43] F. Barratt, J. Dborin, M. Bal, V. Stojevic, F. Pollmann, and A. G. Green, Parallel quantum simulation of large systems on small NISQ computers, *npj Quantum Inf.* **7**, 79 (2021).
- [44] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, Variational quantum eigensolver with fewer qubits, *Phys. Rev. Res.* **1**, 023025 (2019).
- [45] M. Foss-Feig, D. Hayes, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, J. M. Pino, and A. C. Potter, Holographic quantum algorithms for simulating correlated spin systems, *Phys. Rev. Res.* **3**, 033002 (2021).
- [46] A. Smith, B. Jobst, A. G. Green, and F. Pollmann, Crossing a topological phase transition with a quantum computer, *Phys. Rev. Res.* **4**, L022020 (2022).
- [47] E. Chertkov, J. Bohnet, D. Francois, J. Gaebler, D. Gresh, A. Hankin, K. Lee, D. Hayes, B. Neyenhuis, R. Stutz *et al.*, Holographic dynamics simulations with a trapped-ion quantum computer, *Nat. Phys.* **18**, 1074 (2022).
- [48] R. J. Schoelkopf and S. M. Girvin, Wiring up quantum systems, *Nature (London)* **451**, 664 (2008).
- [49] M. H. Devoret and R. J. Schoelkopf, Superconducting circuits for quantum information: An outlook, *Science* **339**, 1169 (2013).
- [50] J. I. Cirac and P. Zoller, Quantum Computations with Cold Trapped Ions, *Phys. Rev. Lett.* **74**, 4091 (1995).
- [51] R. Blatt and D. Wineland, Entangled states of trapped atomic ions, *Nature (London)* **453**, 1008 (2008).
- [52] M. A. Rowe, A. Ben-Kish, B. DeMarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. Itano, B. Jelenković *et al.*, Transport of quantum states and separation of ions in a dual RF ion trap, *Quantum Info. Comput.* **2**, 257 (2002).
- [53] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk, T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation, *Appl. Phys. Lett.* **88**, 034101 (2006).
- [54] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger, Controlling Fast Transport of Cold Trapped Ions, *Phys. Rev. Lett.* **109**, 080501 (2012).
- [55] L. P. Kadanoff, Scaling laws for Ising models near T_c , *Phys. Phys. Fiz.* **2**, 263 (1966).
- [56] R. Jullien, J. Fields, and S. Doniach, Kondo Lattice: Real-Space Renormalization-Group Approach, *Phys. Rev. Lett.* **38**, 1500 (1977).
- [57] S. D. Drell, M. Weinstein, and S. Yankielowicz, Quantum field theories on a lattice: Variational methods for arbitrary coupling strengths and the Ising model in a transverse magnetic field, *Phys. Rev. D* **16**, 1769 (1977).
- [58] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [59] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Quantum Circuits for General Multiqubit Gates, *Phys. Rev. Lett.* **93**, 130502 (2004).
- [60] V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of quantum-logic circuits, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **25**, 1000 (2006).
- [61] R. Iten, R. Colbeck, I. Kukuljan, J. Home, and M. Christandl, Quantum circuits for isometries, *Phys. Rev. A* **93**, 032318 (2016).
- [62] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [63] M. Suzuki, Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, *Commun. Math. Phys.* **51**, 183 (1976).
- [64] T. Barthel and Y. Zhang, Optimized Lie-Trotter-Suzuki decompositions for two and three non-commuting operators, *Ann. Phys.* **418**, 168165 (2020).
- [65] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter Error with Commutator Scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [66] R. Haghshenas, J. Gray, A. C. Potter, and G. K.-L. Chan, Variational Power of Quantum Circuit Tensor Networks, *Phys. Rev. X* **12**, 011047 (2022).
- [67] R. Haghshenas, E. Chertkov, M. DeCross, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, N. Hewitt, C. V. Horst, M. Matheny *et al.*, Probing critical states of matter on a digital quantum computer, [arXiv:2305.01650](https://arxiv.org/abs/2305.01650) (2023).
- [68] G. Evenbly and G. Vidal, Algorithms for entanglement renormalization, *Phys. Rev. B* **79**, 144108 (2009).
- [69] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, *Phys. Rev. A* **69**, 032315 (2004).
- [70] V. V. Shende, I. L. Markov, and S. S. Bullock, Minimal universal two-qubit controlled-NOT-based circuits, *Phys. Rev. A* **69**, 062321 (2004).
- [71] J. Li, X. Yang, X. Peng, and C.-P. Sun, Hybrid Quantum-Classical Approach to Quantum Optimal Control, *Phys. Rev. Lett.* **118**, 150503 (2017).

- [72] G. G. Guerreschi and M. Smelyanskiy, Practical optimization for hybrid quantum-classical algorithms, [arXiv:1701.01450](#) (2017).
- [73] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
- [74] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. (Springer, New York, 2006).
- [75] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, *Math. Program.* **45**, 503 (1989).
- [76] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, *Appl. Phys. Rev.* **6**, 021314 (2019).
- [77] S. Sheldon, L. S. Bishop, E. Magesan, S. Filipp, J. M. Chow, and J. M. Gambetta, Characterizing errors on qubit operations via iterative randomized benchmarking, *Phys. Rev. A* **93**, 012301 (2016).
- [78] D. C. McKay, S. Sheldon, J. A. Smolin, J. M. Chow, and J. M. Gambetta, Three-Qubit Randomized Benchmarking, *Phys. Rev. Lett.* **122**, 200502 (2019).
- [79] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [80] A. Sørensen and K. Mølmer, Quantum Computation with Ions in Thermal Motion, *Phys. Rev. Lett.* **82**, 1971 (1999).
- [81] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Demonstration of a small programmable quantum computer with atomic qubits, *Nature (London)* **536**, 63 (2016).
- [82] D. Maslov, Basic circuit compilation techniques for an ion-trap quantum machine, *New J. Phys.* **19**, 023035 (2017).
- [83] B. Kraus and J. I. Cirac, Optimal creation of entanglement using a two-qubit gate, *Phys. Rev. A* **63**, 062309 (2001).
- [84] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, Geometric theory of nonlocal two-qubit operations, *Phys. Rev. A* **67**, 042313 (2003).
- [85] W. F. Stinespring, Positive functions on C^* -algebras, *Proc. Am. Math. Soc.* **6**, 211 (1955).
- [86] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
- [87] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
- [88] V. Murg, F. Verstraete, O. Legeza, and R. M. Noack, Simulating strongly correlated quantum systems with tree tensor networks, *Phys. Rev. B* **82**, 205105 (2010).
- [89] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, Fast reset and suppressing spontaneous emission of a superconducting qubit, *Appl. Phys. Lett.* **96**, 203110 (2010).
- [90] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff, Fast and Unconditional All-Microwave Reset of a Superconducting Qubit, *Phys. Rev. Lett.* **121**, 060502 (2018).
- [91] D. Egger, M. Werninghaus, M. Ganzhorn, G. Salis, A. Fuhrer, P. Müller, and S. Filipp, Pulsed Reset Protocol for Fixed-Frequency Superconducting Qubits, *Phys. Rev. Appl.* **10**, 044030 (2018).
- [92] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt, Experimental repetitive quantum error correction, *Science* **332**, 1059 (2011).
- [93] J. P. Gaebler, C. H. Baldwin, S. A. Moses, J. M. Dreiling, C. Figgatt, M. Foss-Feig, D. Hayes, and J. M. Pino, Suppression of midcircuit measurement crosstalk errors with micromotion, *Phys. Rev. A* **104**, 062440 (2021).
- [94] E. Knill, G. Ortiz, and R. D. Somma, Optimal quantum measurements of expectation values of observables, *Phys. Rev. A* **75**, 012328 (2007).
- [95] D. Wang, O. Higgott, and S. Brierley, Accelerated Variational Quantum Eigensolver, *Phys. Rev. Lett.* **122**, 140504 (2019).
- [96] G. V. Uimin, One-dimensional problem for $S = 1$ with modified antiferromagnetic Hamiltonian, *JETP Lett.* **12**, 225 (1970).
- [97] C. K. Lai, Lattice gas with nearest-neighbor interaction in one dimension with arbitrary statistics, *J. Math. Phys.* **15**, 1675 (1974).
- [98] B. Sutherland, Model for a multicomponent quantum system, *Phys. Rev. B* **12**, 3795 (1975).
- [99] A. Läuchli, G. Schmid, and S. Trebst, Spin nematics correlations in bilinear-biquadratic $S = 1$ spin chains, *Phys. Rev. B* **74**, 144426 (2006).
- [100] M. Binder and T. Barthel, Low-energy physics of isotropic spin-1 chains in the critical and Haldane phases, *Phys. Rev. B* **102**, 014447 (2020).
- [101] I. H. Kim and B. Swingle, Robust entanglement renormalization on a noisy quantum computer, [arXiv:1711.07500](#) (2017).
- [102] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pienti, M. Chmielewski, C. Collins *et al.*, Benchmarking an 11-qubit quantum computer, *Nat. Commun.* **10**, 5464 (2019).
- [103] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Experimental comparison of two quantum computing architectures, *Proc. Natl. Acad. Sci. USA* **114**, 3305 (2017).
- [104] K. Mølmer and A. Sørensen, Multiparticle Entanglement of Hot Trapped Ions, *Phys. Rev. Lett.* **82**, 1835 (1999).
- [105] A. Sørensen and K. Mølmer, Entanglement and quantum computation with ions in thermal motion, *Phys. Rev. A* **62**, 022311 (2000).
- [106] Q. Miao and T. Barthel, Convergence and quantum advantage of Trotterized MERA for strongly-correlated systems, [arXiv:2303.08910](#) (2023).
- [107] C. Itoi and M.-H. Kato, Extended massless phase and the Haldane phase in a spin-1 isotropic antiferromagnetic chain, *Phys. Rev. B* **55**, 8295 (1997).
- [108] W. Huang, K. A. Gallivan, and P.-A. Absil, A Broyden class of quasi-Newton methods for Riemannian optimization, *SIAM J. Optim.* **25**, 1660 (2015).
- [109] S. T. Smith, in *Hamiltonian and Gradient Flows, Algorithms, and Control*, Vol. 3 of *Fields Institute Communications* (AMS, Providence, RI, 1994), p. 113.
- [110] M. Hauru, M. Van Damme, and J. Haegeman, Riemannian optimization of isometric tensor networks, *SciPost Phys.* **10**, 040 (2021).
- [111] I. A. Luchnikov, M. E. Krechetov, and S. N. Filippov, Riemannian geometry and automatic differentiation for optimization problems of quantum physics and quantum technologies, *New J. Phys.* **23**, 073006 (2021).