Holographic quantum algorithms for simulating correlated spin systems

Michael Foss-Feig [●],[1,*] David Hayes,[1] Joan M. Dreiling,[1] Caroline Figgatt,[1] John P. Gaebler [●],[1] Steven A. Moses,[1] Juan M. Pino,[1] and Andrew C. Potter[1,2,†]

[1]*Honeywell Quantum Solutions, 303 S. Technology Ct., Broomfield, Colorado 80021, USA*
[2]*Department of Physics, University of Texas at Austin, Austin, Texas 78712, USA*

We present a suite of "holographic" quantum algorithms for efficient ground-state preparation and dynamical evolution of correlated spin systems, which require far fewer qubits than the number of spins being simulated. The algorithms exploit the equivalence between matrix-product states (MPS) and quantum channels, along with partial measurement and qubit reuse, in order to simulate a D-dimensional spin system using only a $(D-1)$-dimensional subset of qubits along with an ancillary qubit register whose size scales logarithmically in the amount of entanglement present in the simulated state. Ground states can either be directly prepared from a known MPS representation or obtained via a holographic variational quantum eigensolver (holoVQE). Dynamics of MPS under local Hamiltonians for time t can also be simulated with an additional (multiplicative) poly(t) overhead in qubit resources. These techniques open the door to efficient quantum simulation of MPS with exponentially large bond dimension, including ground states of two- and three-dimensional systems, or thermalizing dynamics with rapid entanglement growth. As a demonstration of the potential resource savings, we implement a holoVQE simulation of the antiferromagnetic Heisenberg chain on a trapped-ion quantum computer, achieving within 10(3)% of the exact ground-state energy of an infinite chain using only a pair of qubits.

I. INTRODUCTION

One of the most promising near-term applications of quantum computers is the simulation of correlated quantum systems in which entanglement plays a crucial role, for which accurate classical simulations are often intractable. Examples include predicting low-temperature properties of correlated materials [1], calculating reaction rates or photoabsorption spectra of large molecules [2], and simulating lattice gauge theories of particle physics [3]. While these applications have generated considerable excitement, it is far from clear how large and how accurate a quantum computer will need to be in order to address classically hard questions of practical scientific and technological relevance. Many problems of interest require extracting information about systems in the thermodynamic limit, which often requires finite-size scaling to be performed on simulation results obtained from systems with hundreds (if not many thousands) of spins. At present, there are no circuit-model quantum computers that can directly simulate spin systems of these sizes.

However, it is well known that system size alone does not determine the classical hardness of simulating a quantum system. While the system size N determines the Hilbert space dimension ($\mathscr{D} \sim e^N$), which in turn sets the classical complexity of simulating the system's wave function by brute force, the Hilbert space actually explored by physical systems is highly structured, enabling efficient parametrizations of physical wave functions. The study of tensor networks over the past few decades has brought this point into sharp focus: Tensor network simulations generally require resources that scale no worse than algebraically with the system size and only suffer an exponential scaling with respect to the *amount of entanglement*, quantified as the bipartite entanglement entropy. This realization, and the scaling laws connecting entanglement entropy to equilibrium and nonequilibrium phases of matter, has made it possible to judge by inspection of general properties of a model—e.g., whether it is in equilibrium (and if so if it is at zero or finite temperature), its geometry, its spacial dimension, or its topological properties—whether it can be simulated efficiently on a classical computer or truly requires quantum resources to simulate. Situations in which the latter case is realized are compelling examples of hard (and practically relevant) problems for which quantum computing could provide a significant near-term benefit.

In general, the existence of a simple tensor-network representation for a state does not guarantee that properties of that state can be calculated efficiently, because the network may be difficult to contract. Important examples include when the size (bond dimension) of the tensors needs to be extremely large to achieve a good approximation, as happens generically for matrix-product state (MPS) simulations of higher dimensional systems or long-time evolution, or because the tensor network topology does not permit an efficient contraction [4]. In the past few years, several proposals have pointed out that near-term quantum computers may be capable of carrying out tensor-network calculations that are beyond the reach of

*michael.feig@honeywell.com
†acpotter@utexas.edu

classical computers [5–9], with a key insight that the size of that quantum computer can be far smaller than the physical system described by the tensor network. Very recently, these ideas have been exploited to provide variational energy estimates for the 2D Heisenberg model by simulating small quantum circuits [10].

In this paper, we present a toolbox for constructing and time-evolving high-bond dimension MPS states on a quantum computer. We refer to these techniques as "holographic" [6,7] because they enable simulation of a D-dimensional system using only a $(D-1)$-dimensional cross section's worth of qubits by simulating the transfer matrix for the MPS as a quantum channel [11] (this concept should not be confused with the common usage of "holography" in the context of the anti-de Sitter/conformal field theory correspondence). The channel effectively moves along the MPS by one unit of distance per channel iteration. Operationally, a purified version of this channel is implemented via unitary operations between a cross section of spins (physical qubits) and an ancillary quantum memory (bond qubits), followed by partial measurement of the physical qubits. Each iteration of the quantum channel moves one step along the stacking direction of the cross sections, with physical qubits reset between iterations and reused without duplication, thereby trading spatial resources (qubit number) for time resources (circuit depth). We present a detailed description of these techniques that applies to systems in any number of spatial dimensions and then benchmark them via classical simulations of algorithm performance on solvable one-dimensional (1D) spin chains.

Next, we show that this representation can actually be made to work by implementing it on a Honeywell trapped-ion quantum computer and using it to estimate the ground-state energy of the antiferromagnetic Heisenberg chain. A crucial technical ingredient required to perform holographic simulations is the ability to selectively measure and initialize a subset of qubits midcircuit, without affecting the remaining qubits. The quantum charge-coupled device (QCCD) architecture [12], in which individual qubits can be dynamically positioned far from other qubits during the execution of a circuit, enables individual addressing (including gates, measurement, and state preparation) with extremely low cross talk [13,14] and is therefore very well suited for these types of algorithms.

Finally, we extend these ideas to the simulation of quench dynamics starting from a holographically generated MPS. Naively, one would expect the simulation of N initially correlated qubits evolving for a time t under a local Hamiltonian to require a circuit of width $w = N$ and depth $d \sim \text{poly}(t)$. A constructive algorithm achieving such scaling for k-local Hamiltonians was found more than 20 years ago [15], and in the years since the dependence of the circuit depth d on t, N, and the error tolerance have all improved (see Refs. [16,17] and references therein). It might seem that the circuit width requirement $w \sim N$ is fundamental, or even tautological. While $w = N$ is indeed required in certain worst-case scenarios, in this paper we explore how far fewer than N qubits suffice in many cases of practical interest. Consider, for example, a one-dimensional (1D) system initially in a state with maximum bipartite entanglement entropy S: Time evolution of this state by a local Hamiltonian can be implemented by a circuit of width $\text{poly}(t) + S$ and depth $N \times \text{poly}(t)$. In other

words, the required number of qubits is determined by the *evolution time* (with a modest constant offset to accommodate the initial entanglement of the state), while the physical size of the system being simulated can be accommodated by increasing the depth of the circuit. Moreover, if the initial state has a finite correlation length ξ, the $N \to \infty$ limit can be well approximated by a circuit of depth $\sim \xi \times \text{poly}(t)$. This reshuffling of resources from circuit width to circuit depth is what we mean by the term "holographic."

II. HOLOGRAPHIC SIMULATION OF MATRIX PRODUCT STATES

The basis for our quantum simulation algorithms will be the MPS representation of quantum states, which provides an efficient compressed approximation of quantum states with less than maximal entanglement. We will construct methods for simulating ground-state properties and quench dynamics of local lattice models, defined on a Hilbert space $\mathcal{H}$ that decomposes into a tensor product of "sites" as $\mathcal{H} = \otimes_{i=1}^{\ell} \mathcal{H}_i$. Here the sites are arranged along a 1D line with open boundary conditions and have finite local dimension $|\mathcal{H}_i| = \mathcal{Q}$, and by local we mean that interactions act on at most k sequential sites. Note that this class of systems includes not only 1D spin chains, where each site simply represents a single spin, but also D-dimensional models that can be sliced into a 1D stack of $(D-1)$-dimensional cross sections (e.g., for a D-dimensional cube $\ell = N^{1/D}$).

MPS describe quantum states on such 1D stacks using far fewer parameters—$\sim \mathcal{O}(\ell \mathcal{Q} \chi^2)$, with χ being the bond dimension—than the worst-case $\mathcal{Q}^{\ell}$ required to specify a generic state, and describe a very special class of states with entanglement entropy across any cut bounded above by $\log_2 \chi$. For ground states of 1D or quasi-1D systems, MPS states can be implemented using classical resources that scale at worst polynomially in the desired accuracy and system size and provide efficient classical algorithms for simulating low-dimensional ground states and short-time dynamics. However, classical MPS methods fail for 2D and 3D systems, and for longer time dynamics where substantial amounts of entanglement have been generated. Nevertheless, many of these systems have far less than the maximal amount of entanglement obtained by random states, and the MPS description provides a dramatic compression. Our goal is to devise an efficient method to prepare and time evolve an MPS representable state on a quantum computer with an economical use of qubit resources, which gives access to MPS with classically inaccessible bond dimension, while still leveraging the economical MPS representation for states with less than maximal entanglement.

To set the stage for these algorithms, we first briefly review a few key properties of MPS that are essential to understanding how they can be represented holographically on a small quantum computer.

A. Brief review of MPS formalism

An MPS with open boundary conditions can be written

$$|\Psi\rangle = \sum_{\sigma_1, \ldots, \sigma_\ell} L^T V_{\sigma_1}^{[1]} V_{\sigma_2}^{[2]} \cdots V_{\sigma_\ell}^{[\ell]} R |\sigma_1, \ldots, \sigma_\ell\rangle. \quad (1)$$
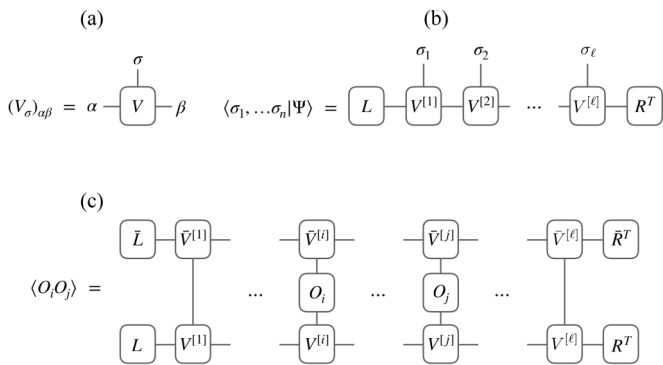
FIG. 1. Graphical description of an MPS. (a) An individual tensor, (b) the MPS wave function as a contraction over such tensors, and (c) the contraction of a tensor network to compute a correlation function.

For each site j, $V^{[j]}_{\sigma_j}$ is a set of $\mathcal{Q}$ square matrices [18], or equivalently a rank-3 tensor with "physical" index $\sigma_j = 1 \dots \mathcal{Q}$ and "bond" indices $\alpha, \beta = 1 \dots \chi$, where χ is referred to as the bond dimension. The χ-dimensional vectors L and R specify the left and right boundary conditions, respectively. The standard graphical representation of an MPS is shown in Fig. 1. An individual tensor is drawn as a box with a leg for each index, as in Fig. 1(a). Joined legs imply contraction of the associated tensor indices, such that Fig. 1(b) gives the wave function components of $|\Psi\rangle$ [the contractions are implied as matrix-matrix or matrix-vector multiplication in Eq. (1)].

One can imagine creating an ℓ-site MPS as a physical state of a quantum computer by letting an ancilla register (containing "bond qubits") interact unitarily and sequentially with ℓ physical registers [11], each representing one site of the MPS, as in Figs. 2(a)–2(c). In these circuit diagrams and elsewhere, open circles denote initialization of a qubit (or register of qubits) to the $|0\rangle$ state. To implement this construction with unitary circuit elements, one must exploit the gauge redundancy of the MPS description [19] to place the MPS in *right canonical form* (RCF) such that

$$\sum_{\sigma} V^{[j]}_{\sigma} V^{[j]\dagger}_{\sigma} = \mathbb{1} \quad \forall j. \tag{2}$$
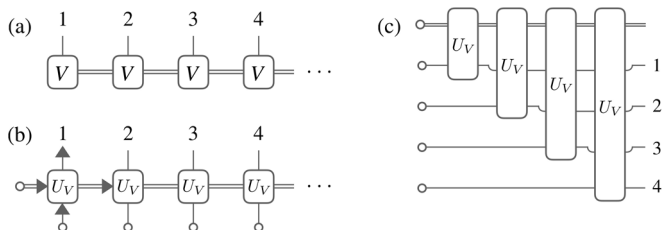

FIG. 2. MPS as a quantum circuit. (a) An MPS, assumed to be in right canonical form. The right canonical condition [Eq. (2)] guarantees that V is an isometry from $\mathbb{C}^{\chi} \to \mathbb{C}^{\chi} \otimes \mathbb{C}^{\mathcal{Q}}$, and it can therefore be embedded in a unitary U_V acting on $\mathbb{C}^{\chi} \otimes \mathbb{C}^{\mathcal{Q}}$ but restricted to a fixed input of the physical qubit (denoted with an open circle), as in panel (b). The unitary evolution in panel (b) is equivalent (up to simply rearranging the lines) to the circuit diagram in panel (c).

The relationship between RCF and unitary embedding [20] will be discussed further in the next section, but for now we simply want to emphasize that the aforementioned MPS gauge redundancy ensures that we can, without loss of generality, assume that Eq. (2) holds for the state in Eq. (1). Note that this definition of RCF is slightly nonstandard: The boundary tensors are typically also canonical, but we choose this version of RCF because it simplifies much of what follows. While this choice imposes some limitations on how faithfully the right-boundary condition can be imposed in a quantum circuit, bulk physics will be unaffected. Note also that for a normalized state $|\Psi\rangle$, the imposition of RCF implies that in general both L and R cannot be simultaneously normalized. Without loss of generality, we take L to be normalized but not necessarily R.

Correlation functions of local operators, $\langle \hat{O}_i \hat{O}_j \rangle$, can be computed by contracting the tensor network depicted in Fig. 1(c). An efficient method to contract such a network for a long chain is to first contract the physical legs of the tensors on each site to form transfer matrices

$$\mathcal{E}^{[k]}_{\alpha\gamma,\beta\delta} = \sum_{\sigma_k} \left(\bar{V}^{[k]}_{\sigma_k} \right)_{\gamma\delta} \left(V^{[k]}_{\sigma_k} \right)_{\alpha\beta} \tag{3}$$

on each site $k \neq i, j$ (with $\bar{V}$ being the complex conjugate of V) and begin multiplying these transfer matrices from left to right. When site i is encountered, we apply the modified transfer matrix

$$\mathcal{O}^{[i]}_{\alpha\gamma,\beta\delta} = \sum_{\sigma_i,\tau_i} \left(\bar{V}^{[i]}_{\tau_i} \right)_{\gamma\delta} \langle \tau_i | \hat{O}_i | \sigma_i \rangle \left(V^{[i]}_{\sigma_i} \right)_{\alpha\beta}, \tag{4}$$

and similarly for site j. If we interpret the bond vector space as the Hilbert space of some quantum system, then the transfer matrix is a linear superoperator acting on bond-space density matrices $\rho_{\beta\delta}$ ($\beta, \delta \in \{1, 2, \dots \chi\}$) via the mapping

$$\mathcal{E}: \rho \to \sum_{\sigma} V^{\dagger}_{\sigma} \rho V_{\sigma}. \tag{5}$$

Together with the RCF conditions in Eq. (2), Eq. (5) establishes $\mathcal{E}$ as a quantum channel (trace-preserving completely positive map) on the bond space [19], with the MPS matrices V_{σ} as the Kraus operators of the channel. In this language, the contraction depicted in Fig. 1(c) can be expressed as the overlap of a "time-evolved" initial bond-space density matrix $\rho_i = |L\rangle\rangle\langle\langle L|$ (with ket $| \dots \rangle\rangle$ indicating a state in the bond Hilbert space) with the final un-normalized state $|R\rangle\rangle$,

$$\langle \hat{O}_i \hat{O}_j \rangle = \langle\langle R | \mathcal{E}^{[\ell]} \circ \cdots \mathcal{O}^{[j]} \circ \cdots \mathcal{O}^{[i]} \circ \cdots \mathcal{E}^{[1]}(\rho_i) | R \rangle\rangle. \tag{6}$$

Indeed, the earliest comprehensive treatment of MPS in the literature (in which they were called *finitely correlated states*) defined them in terms of quantum channels [21].

B. Holographic MPS generation

Equation (6) demonstrates that correlation functions of an MPS can be encoded in the dynamics of a quantum system with size independent of ℓ; the *spatial* structure of the physical Hilbert space has been converted into a (discrete) *time* direction of the bond Hilbert space (see also Ref. [20], where this mapping is exploited to apply classical open-system simulation techniques to the simulation of infinite-width circuits). Thus, one can simulate the state of a D-dimensional system
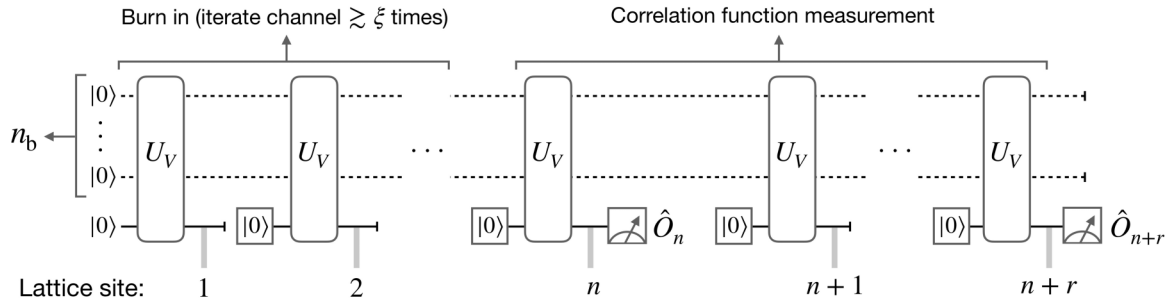
FIG. 3. A circuit implementing holographic state preparation and correlation function measurement of an MPS. In this simple example, we use one system qubit (enabling simulation of an infinite 1D spin-$\frac{1}{2}$ chain) and n_b bond qubits for an MPS with bond dimension $\chi = 2^{n_b}$. Measurements are implied to be in the eigenbasis of the operators $\hat{O}_n$ and $\hat{O}_{n+r}$. Averaging over the outcome of this circuit (with the two measurement results multiplied together) produces the correlation function $\langle \hat{O}_n \hat{O}_{n+r} \rangle$.

using a system of dimension $D - 1$, inspiring the moniker "holographic" [5]. However, it is important to keep in mind that this dynamics is not unitary. The holographic algorithms described here can be viewed as explicit purifications of this nonunitary dynamics in the form of quantum circuits. Alternatively, one can understand these algorithms starting from the known representation of MPS as quantum circuits, previously described in Sec. II A and illustrated in Fig. 2. From this perspective, the dimensional reduction can be understood by looking at the causal structure of that circuit and recognizing that the physical register corresponding to site j can be measured and reset before the bond qubit register interacts with the physical register of site $j + 1$, implying that only a single site worth of physical qubits is required to implement the entire circuit [8]. Despite the constant erasure of information in the physical qubits, long-range spatial correlations in the system are retained as memory in the bond qubits.

Generic versions of such holographic algorithms for 2D systems were previously outlined in Ref. [5], though without explicit discussion of connections to the MPS formalism, and the known representation of MPS as quantum channels. Subsequent work [6,7] revealed an intriguing element of noise resilience in holographic simulation techniques. Due to the repeated partial measurement and reset in the holographic technique, errors do not propagate indefinitely as for purely unitary circuits. Consequently, a finite density of errors produces a finite imprecision on the measured correlation function, in contrast to a purely unitary circuit, for which a single error can spread and contaminate all outputs.

In what follows, we unify these perspectives with the framework of MPS and develop concrete variational ground-state preparation and quantum dynamics simulation techniques using this framework. We begin with a detailed description of the holographic MPS preparation and measurement protocol, summarized in Fig. 3. This protocol utilizes a register of $n_b = \log_2 \chi$ "bond" qubits (representing the χ-dimensional bond Hilbert space) initialized in state $|L\rangle\!\rangle$, and a register of $n_p = \log_2 \mathcal{Q}$ "physical" qubits (representing the $\mathcal{Q}$-dimensional physical Hilbert space of a single lattice site) prepared in a fixed reference state $|0\rangle$. The channel $\mathcal{E}$ is realized by applying unitary gates between the physical qubits and bond qubits and then tracing out (i.e., discarding) the physical qubit. Such a unitary purification of the MPS channel can always be constructed via the Stinespring dilation.

Specifically, one can embed each MPS tensor $(V_\sigma)_{\alpha,\beta}$ as the columns of a unitary matrix $U_{\alpha,\sigma;\beta,\sigma'}$ with fixed index $\sigma' = 0$, i.e., $(V_\sigma)_{\alpha,\beta} = \langle\sigma|\langle\!\langle\alpha|U|0\rangle|\beta\rangle\!\rangle$. Since, in right canonical form, V forms an isometry from $\mathbb{C}^\chi \rightarrow \mathbb{C}^\chi \otimes \mathbb{C}^\mathcal{Q}$, the columns with $\sigma' = 0$ form an orthonormal set, which can always be completed into a full orthonormal basis for $\mathbb{C}^\chi \otimes \mathbb{C}^\mathcal{Q}$ to obtain U.

Any correlation function, $C = \langle\Psi|(\otimes_i \hat{O}_i)|\Psi\rangle$, of the corresponding MPS can be sampled by iterating this quantum channel to step through the sites of the chain from 1 to ℓ in the following sequence of steps:

(1) State prep: Start with $C = 1$. Prepare the bond-qubit register in a given state $|L\rangle\!\rangle$, which sets the left boundary condition for the MPS. Then, starting with site $i = 1$:

(2) While $i \leqslant \ell$: Iteratively apply the quantum channel for the MPS matrix to step along the chain from site 1 toward site ℓ by the following steps:

(a) Prepare the physical qubit register in a reference state, $|0\rangle$.

(b) Act on the physical qubits and bond qubits with a unitary circuit U, which is a purification of the MPS on-site tensor for site i.

(c) Measure the physical qubits in the eigenbasis of operator $\hat{O}_i$ (which may be the identity operator on most sites, for which the measurement is unnecessary). Denote the eigenvalue of $\hat{O}_i$ corresponding to the measurement outcome by λ_i, and multiply $C \leftarrow \lambda_i C_i$.

(d) Increment: $i \leftarrow i + 1$.

(3) Measure the bond qubits in a basis containing $|R\rangle\!\rangle$.

If we postselect on outcomes for which the bond qubits are found to be in state $|R\rangle\!\rangle$, then the above algorithm samples the correlator C in the state $|\Psi\rangle$, and the expectation value can be estimated by averaging over sufficiently many repetitions to achieve a desired statistical precision [22]. Note that fixing the right boundary condition through postselection incurs a multiplicative overhead $\sim\chi$, which will be very large in cases where quantum computation is required. However, since we are primarily interested in bulk properties and since correlations decay exponentially in distance from the boundary, it is generally sufficient to skip this postselection, which provides a weighted average of the correlation function over the right boundary condition. If we are interested in boundary effects, for example, when examining impurity models or boundary conformal field theories, one can study the left boundary

(where the boundary condition is set by the state preparation for the bond register).

C. Holographic entanglement measurements

Measures of entanglement provide detailed insights into quantum many-body systems beyond what can be drawn from local correlation functions, revealing nonlocal correlations and topological and symmetry-protected topological orders [23–25], and diagnosing thermalization, scrambling, and many-body localization [26]. Methods to measure Rényi entropies

$$S_A^{(n)} = \frac{1}{1-n} \log_2 \operatorname{tr} \rho_A^n \tag{7}$$

for a subsystem A, with integer index n, have been developed based on creating replica copies of the system and measuring operators that cyclically permute the quantum states of various copies [27], or by examining cross correlations in randomized measurements to virtually implement the desired replicas [28]. These methods can be directly adapted to holographically represented states by performing the desired measurements on the physical qubit registers, as described in the previous section. Such measurements would enable, for example, the estimation of free energy from holographically generated thermofield double states [29,30].

If one is interested in the entanglement entropy of a bipartition of the chain, it can be directly obtained from measurements of the bond-qubit register. Namely, since the holographic simulation method recreates an MPS in right canonical form, the entanglement spectrum for the physical qubits bipartitioned by cutting between sites j and $j-1$ is precisely equal to the spectrum of the density matrix for the bond qubits after j iterations of the holographic simulation algorithm. In this case, the replica-SWAP or randomized measurement techniques can be applied directly to the bond qubits, with a number of measurements that grows with χ, but not the interval size, offering a potential savings in measurement complexity. For replica-SWAP-based measurements of entanglement entropy, this holographic method provides a potentially huge savings in qubits required, as one needs to replicate only the bond qubits, without replicating an extensive number of physical qubits for every site in the chain.

D. Expressivity of holographic MPS

While a unitary circuit representing the (purified) quantum channel of any MPS is formally guaranteed to exist, the crux for practical use of holographic simulation techniques will be constructing effective methods for implementing channels for physically relevant systems using low-depth circuits. Namely, arbitrary unitary synthesis from a local gate set generically requires gate counts that scale exponentially with qubit number and is clearly not a viable technique for large bond dimension. However, physical systems with local Hamiltonians are far from "generic" and have considerable structure that could be exploited for efficient simulation. This observation poses the following basic question: What class of quantum states can be efficiently holographically represented on a quantum computer with exponentially large $\chi \sim 2^{n_b}$, but using low-depth [i.e., with poly(n_b) gates] quantum circuits? Though we can-

not conclusively answer this general question, in the following we develop holographic algorithms for simulating nonequilibrium dynamics starting from correlated ground states, and provide numerical evidence demonstrating that low-depth circuits may suffice in many situations of practical interest.

III. HOLOGRAPHIC VARIATIONAL QUANTUM EIGENSOLVER (holoVQE)

A central task for quantum materials simulation is to accurately approximate the ground-state correlations of local Hamiltonians $H = \sum_j h_j$, where each term h_j acts on sites within a distance at most k from site j. Hybrid classical-quantum variational algorithms, like the variational quantum eigensolver (VQE) [31], offer promising methodologies for attacking this problem on moderate-scale quantum computers. In VQE, one prepares a trial wave function $|\psi(\theta)\rangle$ on a quantum computer by evolving a fixed initial state with a quantum circuit composed of gates parameterized by rotation angles $\theta \in \mathbb{R}^p$ (p being the number of variational parameters). The expectation value of the energy, $E(\theta) = \sum_j \langle \psi(\theta)|h_j|\psi(\theta)\rangle$, is subsequently estimated by measuring each individual term in the sum to the desired precision. Then, a classical computer updates the parameters θ to lower the variational energy in order to find the best approximation of the true ground state within the family of states, $|\psi(\theta)\rangle$.

The holographic representation of MPS on a quantum computer naturally suggests a holographic extension of VQE (holoVQE), which uses the MPS representation described in the previous section with the unitary U_V represented by a parameterized circuit. The expectation value of energy can be computed by measuring each term in the Hamiltonian using the above-described procedure for measuring correlation functions. Then, the variational ansatz can be optimized by using a classical algorithm to minimize $E(\theta)$.

The implementation of holoVQE is simplified for crystalline materials with translation-invariant Hamiltonians, where there are only a finite number of terms in the Hamiltonian that must be independently measured in the thermodynamic limit. The holographic method produces an MPS with open boundary conditions which is not translation invariant. However, in an MPS the boundary's influence decays exponentially with the correlation length ξ, so one can simply measure the distinct terms in the Hamiltonian a distance $r \gtrsim \xi$ from the boundary. In the holographic correspondence, we move a large distance into the spatial bulk by iterating the MPS quantum channel for a long time to burn in its steady state, as in Fig. 3.

In the following sections, we demonstrate a simple application of the holoVQE technique for approximating the ground-state energy of an XXZ chain using only a single ancillary bond qubit, and then physically implement this technique on a trapped ion quantum computer to analyze the SU(2)-symmetric Heisenberg point. We show that symmetry principles can be incorporated into the variational circuit ansatz to reduce the number of variational parameters and simplify the optimization. Next, we consider a model with lower symmetry: the transverse-field Ising model (TFIM). We apply a more generic circuit ansatz and analyze the scaling of algorithm performance for ground-state preparation and

reconstruction of critical correlation functions with increasing number of bond qubits. For small circuit sizes, we are able to reproduce the optimal MPS ansatz at bond dimension $\chi = 2^{n_b}$, obtaining relative accuracy on the (infinite size) critical ground-state energy below 10^{-4} with only a few qubits.

A. Role of symmetries

The crux of effectively implementing a variational procedure is constructing a good variational ansatz. For present purposes, this entails identifying a parameterized unitary circuit family that effectively implements the purified transfer matrix of the desired MPS approximation to the ground state. Consideration of symmetries can guide the design of these circuits.

Symmetries of the model can be strictly enforced on the variational states by restricting attention to symmetry-preserving circuit families. Here, by symmetry-preserving circuit, we mean that we choose a particular linear representation of the symmetry action on the bond qubits (possibly including projective representations [32] when dealing with potential symmetry-protected topological states), and ensure that parametrizations of the variational circuit preserve the total symmetry quantum numbers of the physical and bond qubits together.

We note that it is not always desirable to explicitly enforce all symmetries of the model. For example, we may wish to assess whether the ground state spontaneously breaks symmetries, in which case one could build an ansatz around various possible symmetry-broken configurations. Moreover, it is often the case that, for a fixed bond dimension, the lowest energy state may not preserve the full set of symmetries possessed by the ground state. For example, when $\chi = 1$ (no bond qubits), an MPS simply corresponds to a mean-field (best product state) ansatz, for which energy minimization often yields symmetry-broken solutions. The examples below are indicative of these various possibilities.

B. *XXZ* chain

The *XXZ* spin chain is a canonical model of strongly correlated 1D systems, describing both one-dimensional quantum magnetism and superfluidity (by mapping the spins to hard-core bosons). For nearest neighbor interactions, the Hamiltonian is

$$H = J \sum_i (X_i X_{i+1} + Y_i Y_{i+1} + \Delta Z_i Z_{i+1}), \tag{8}$$

where (without loss of generality) we take $J > 0$. The model has a global U(1) symmetry, enlarged to a full SU(2) symmetry at the Heisenberg points $\Delta = \pm 1$. For $\Delta < -1$ [$\Delta > 1$], the spectrum is gapped and the symmetry-broken ground state is ferromagnetically [antiferromagnetically] ordered, while for $|\Delta| \leqslant 1$ the model is gapless and has no long-range spin order. For all values of Δ, the model is exactly solvable by Bethe ansatz [33].

In the antiferromagnetic phase ($\Delta \geqslant 1$), the mean-field solution is antiferromagnetically ordered, spontaneously breaking discrete-translational symmetry [and SU(2) symmetry for $\Delta = 1$]. Since this state is consistent with the known value of $S_z = 0$ for the true ground state, we can build a $\chi = 2$
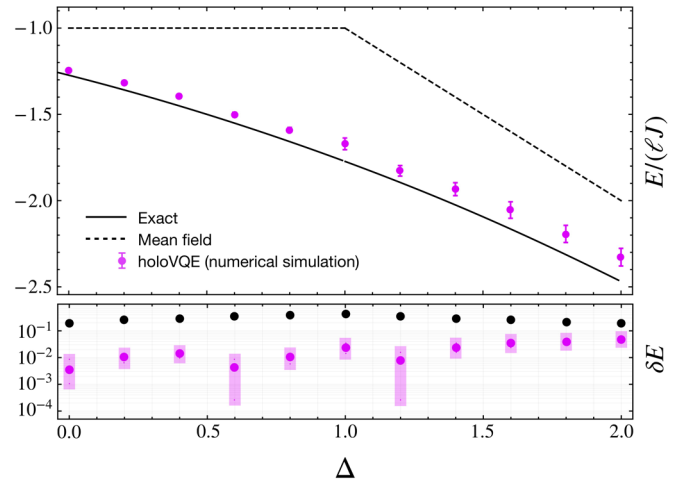

FIG. 4. Top: Energy per site of the *XXZ* chain in the thermodynamic limit. The dashed line is mean-field theory, the solid line is the exact energy obtained from Bethe ansatz, and the points are from simulated holoVQE using a single bond qubit (error bars represent 1σ sampling uncertainties). Bottom: Fractional energy error [$\delta E = |(E - E_{\text{exact}})/E_{\text{exact}}|$] of holoVQE (pink points with error bars) and mean-field theory (black points), showing that the addition of even a single bond qubit can drastically improve the ground-state energy estimate.

MPS by introducing a single-bond qubit and allowing it to interact with the system qubit via unitaries that conserve total S_z (note that by breaking discrete translational symmetry, this $\chi = 2$ MPS achieves the same energy as a $\chi = 4$ translationally invariant MPS). Choosing $U = \exp[-i\theta(X_p X_b + Y_p Y_b)]\exp[-i\phi(Z_p Z_b)]$, we can now use a holographic representation of the MPS to measure the energy for a given choice of parameters (θ, ϕ) and minimize using a classical feedback loop. Using gradient descent and simulating 256 shots per energy measurement, we obtain the results shown for $\Delta \geqslant 1$ in Fig. 4.

Additional care must be taken when computing the energy for $0 < \Delta < 1$. In this case, the mean-field ground state breaks the U(1) symmetry of the model by spontaneously aligning (antiferromagnetically) along some direction in the *XY* plane. Since this state does not live in the correct symmetry sector with respect to U(1), it is not sufficient to restrict our attention to circuits that conserve total S_z. In these cases, we find that the idea MPS at $\chi = 2$ can be obtained using the three-parameter ansatz $U = \exp[-i(\theta X_p X_b + \phi Y_p Y_b + \psi Z_p Z_b)]$.

C. Trapped ion implementation: Heisenberg chain

At the antiferromagnetic Heisenberg point, we implement the holoVQE procedure experimentally using Honeywell's QCCD trapped-ion quantum computer described in Ref. [12]. We utilize a subset of the five designated "gate zones" [orange and blue in Fig. 5(a)], which suffices to run two parallel instances of the holographic state preparation protocol with a single bond qubit for each [34]. At the Heisenberg point, it can be shown that the two-parameter ansatz described above for $\Delta \geqslant 1$ is actually unnecessarily flexible, and it suffices to
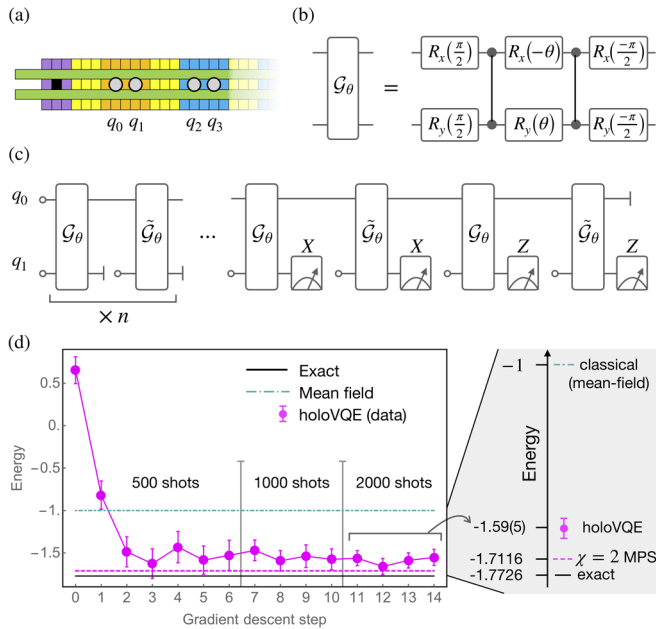
FIG. 5. Experimental implementation of holoVQE on a Honeywell trapped-ion quantum computer. (a) Schematic of the ion trap; for this experiment we used two neighboring gate zones, each loaded with two data qubits and two sympathetic cooling ions (not shown). (b) Decomposition of $\mathcal{G}_\theta$ into two cZ gates. (c) The circuits used for holoVQE involve one physical qubit and one bond qubit, though we utilize the two gate zones to parallelize the data taking. (d) Representative data from holoVQE.

restrict our attention to $\phi = 0$. Thus, the entangling unitary between physical and bond qubit is

$$\mathcal{G}_\theta = \exp[-i\theta(X_\mathrm{p}X_\mathrm{b} + Y_\mathrm{p}Y_\mathrm{b})/2]. \qquad (9)$$

The native two-qubit gate for our architecture (the Mølmer-Sørensen gate [35]) is local unitary equivalent to a controlled-Z (CZ) gate, at least two of which are necessary to synthesize $\mathcal{G}_\theta$ for arbitrary θ [a minimal decomposition is shown in Fig. 5(b)]. The holographic MPS circuit is then built by alternating applications of $\mathcal{G}_\theta$ and $\tilde{\mathcal{G}}_\theta = X_\mathrm{p}\mathcal{G}_\theta$ (this alternation corresponds to starting with a classical antiferromagnet, as discussed above), with reinitialization of the physical qubit in between each entangler [see Fig. 5(c)]. For the small bond dimension accessed in this example, we find that a "burn in" distance of four lattice sites is sufficient to approximate bulk expectation values to well within shot noise. Figure 5(d) shows the results of holoVQE. Starting with a randomly chosen parameter θ, we use gradient descent with derivatives estimated from finite differences of the measured energies $E(\theta)$, increasing the shot count for each energy measurement from 500 to 2000 as gradient descent proceeds. Error bars are 2σ confidence intervals obtained from a nonparametric bootstrap resampling of the data. Averaging over the final four data points (taken at the highest shot counts), we obtain an estimate of $E = -1.59(5)J$ for the per-site ground-state energy. For comparison, the mean-field ground state—which lowers the energy as much as possible without entanglement—achieves $E = -J$ [blue dot-dashed line in Fig. 5(d)]. The optimal MPS with bond dimension $\chi = 2^{n_\mathrm{b}} = 2$ achieves $E \approx -1.712J$,

which the experimental results would converge to if the circuits were perfectly implemented [purple dashed line in Fig. 5(d)], while the exact ground-state ($\chi = \infty$) has $E = J(1 - 4\log 2) \approx -1.773J$ [from Bethe ansatz, black solid line in Fig. 5(d)]. Considering the performance metrics for the computer used to obtain these results [12], we expect the discrepancy between the VQE estimate and the minimum energy of a $\chi = 2$ MPS is primarily due to errors on the two-qubit gates comprising $\mathcal{G}_\theta$.

Note that this measured value provides a proper variational upper bound for the *infinite chain*, despite being obtained from a small quantum circuit. To highlight the resource savings of holoVQE, we note that achieving comparable accuracies using brute-force simulation of an L-site chain would require $L = 6$ (rather than 2) qubits to sufficiently suppress finite-size effects. If the circuit infidelities were reduced by either improving the gate fidelities or using error mitigation techniques, the minimum achievable energy with two qubits for holoVQE, $E_\mathrm{min} \approx -1.712J$, would require 10 (perfect) qubits to achieve by brute force.

D. Increasing the bond dimension: Transverse field Ising model (TFIM)

In the previous examples, we used just a single bond qubit corresponding to MPS with bond dimension 2. These results already demonstrate the dramatic compression of resources enabled by the holographic simulation method in achieving reasonable accuracy on an infinite, critical spin chain using only a pair of qubits. However, turning holoVQE into a useful algorithm requires a method to systematically improve the accuracy of the simulations. We now explore the performance of holoVQE upon including additional bond qubits, focusing on the task of ground-state energy estimation for the 1D transverse-field Ising model (TFIM), with Hamiltonian

$$H_\mathrm{TFIM} = -\sum_j (JZ_jZ_{j+1} + hX_j). \qquad (10)$$

The TFIM exhibits a ground-state phase transition from an ordered ($h < J$) to a disordered ($h > J$) phase, both of which are gapped and can be well described by MPS of fixed (system-size-independent) bond dimension. These phases are separated by a self-dual critical point at $J = h$, described by a conformal field theory (CFT) with central charge $c = 1/2$, whose nonconstant entanglement scaling requires a bond dimension that grows with system size as $\chi \gtrsim L^{c/3} = L^{1/6}$ to achieve asymptotically accurate correlations. Nevertheless, it turns out that the ground-state energy and moderate-range spin-correlation functions of this model can be captured with fairly high accuracy using modest bond dimension MPS, even at the critical point [36].

To explore the efficacy of holoVQE for this paradigmatic toy model, we numerically simulate the holoVQE procedure at the critical point ($h = J$) for a sequence of variationally parameterized circuits with a variable number of bond qubits. For n_b bond qubits, we construct a "star" circuit that involves only two-qubit gates that sequentially entangle the physical qubit and the jth bond qubit for $j = 1, \ldots, n_\mathrm{b}$, allowing each individual gate to be an arbitrary $\in SU(4)$ two-qubit unitary, as in Fig. 6(a). The primary challenge in these calculations is
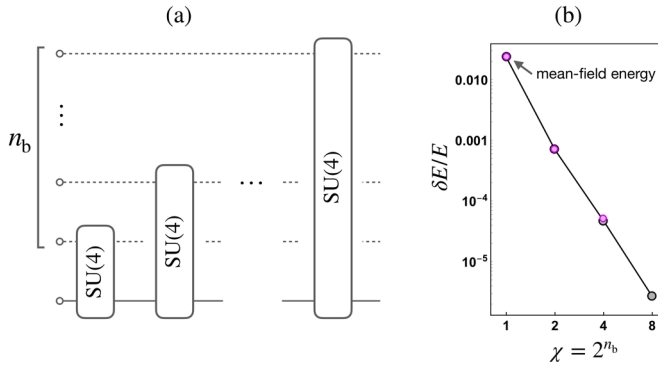
FIG. 6. (a) Circuits used for holographic MPS representation of the TFIM ground state. Each arbitrary SU(4) is decomposed into three native two-qubit gates and eight single-qubit gates, with a total of 15 real variational parameters (as in Ref. [37]). (b) Energy obtained using "star" circuits with an increasing number of bond qubits. The agreement with exact MPS energy minimization at bond dimension $\chi = 2^{n_b}$ is excellent, although we were not able to obtain reliable energies from this ansatz for $n_b \geqslant 3$.
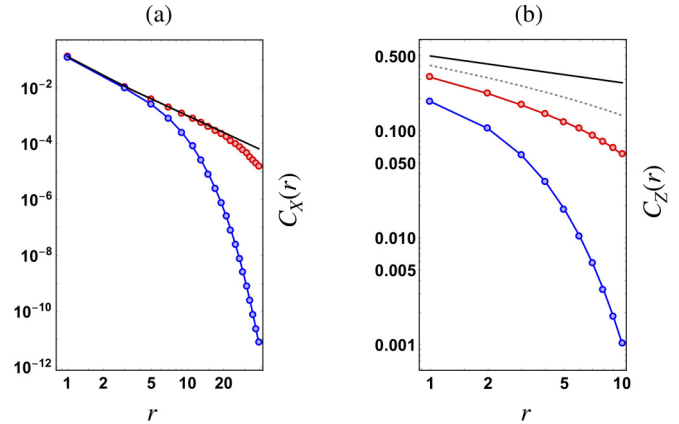

FIG. 7. Correlation functions of the transverse (a) and longitudinal (b) spin directions in the 1D TFIM. Black lines are exact from fermionization, while the points are from holoVQE (blue points for $n_b = 1$ and red points for $n_b = 2$). In panel (b), the $\chi = 8$ curve is shown as well (black dashed line), though this has been obtained from direct MPS optimization as the numerical minimization for an $n_b = 3$ star circuit was inconclusive.

to reliably find the global minimum of a constrained nonlinear optimization problem; we were only able to find reliable results for $n_b = 0, 1, 2$, for which simulated annealing worked well. We note that classical MPS calculations also suffer from this challenge, which is typically overcome by breaking translational invariance in order to make the problem linear (at the cost of greatly expanding the parameter space). It is clear that scaling holoVQE to large circuits (and therefore large effective bond dimension) will require significant further development along these lines.

The results obtained by brute-force global optimization are shown in Fig. 6(b), along with those obtained by an unconstrained MPS optimization using bond dimension $\chi = 2^{n_b}$. Surprisingly, this simple circuit design finds the best possible MPS even for $n_b = 2$, for which the parametrization is not exhaustive of all $n_b + 1 = 3$ qubit unitaries. Because global optimization strategies did not yield reliable improvements for $n_b \geqslant 3$, we do not know if this feature is generic or restricted to small circuits.

Part of the challenge in achieving further improvements for the TFIM is the extremely rapid convergence of variational energy with bond dimension to the exact ground-state energy. We note that, in an actual quantum computation, resolving very small energy differences will become impractical due to large statistical sampling overhead. This issue is especially pronounced in the TFIM, likely due to its integrability and small central charge, $c = \frac{1}{2}$ (the smallest of any minimal model). In particular, assuming a near-optimal variational ansatz, the effective correlation length scales with the number of bond qubits like $\xi \sim e^{n_b(3/c)\log 2}$ (at criticality), so that smaller c yields larger correlation length and more rapid convergence of finite-range correlations with the addition of bond qubits.

Since the Hamiltonian is composed of nearest neighbor interactions, the holoVQE procedure only requires measurement of nearest neighbor correlation functions. Once the circuit is optimized, one can freely use the holographic state preparation subroutine to extract any desired observables in

the variational solution. For example, in Fig. 7 we show the critical ($J = h$) transverse and longitudinal connected spin-spin correlation functions:

$$C_X(r) = \langle X_j X_{j+r} \rangle - \langle X_j \rangle^2,$$
$$C_Z(r) = \langle Z_j Z_{j+r} \rangle - \langle Z_j \rangle^2, \qquad (11)$$

where we have assumed translational invariance. The points are obtained from the holoVQE method, while the solid black lines are exact results from fermionization. The holoVQE results show clear signs of the universal scaling behavior for the Ising transition, $C_X(r) \sim r^{-1}$, and $C_Z(r) \sim r^{-1/4}$ over moderate length scales (as much as 20 sites for the transverse correlations), despite using no more than three qubits.

IV. DYNAMICS

Calculating the dynamical properties of interacting quantum systems is an essential challenge for practical applications such as predicting chemical kinetics, computing nonequilibrium electronic and optical properties of quantum materials and devices, and analyzing NMR spectra [38]. Quantum dynamics also underpins foundational scientific questions ranging from the nature of thermalization and quantum chaos, to properties of quark-gluon plasmas in heavy-ion collisions, to understanding cosmological scenarios for defect production. Despite its fundamental importance, simulating quantum dynamics remains among the most challenging tasks for classical computers, generically requiring exponential classical resources even in low-dimensional systems where ground states can be efficiently simulated. This area is therefore a promising candidate for achieving a practical quantum advantage on near-term quantum hardware.

In the following, we develop a quantum simulation algorithm that incorporates the holographic representation of MPS initial states described above to simulate time evolution of an

initial state under a quantum quench:

$$|\psi(t)\rangle = \mathcal{T}\{e^{-i\int_0^t H(s)ds}\}|\psi(0)\rangle. \tag{12}$$

Here, the initial state $|\psi(0)\rangle$ is represented by an MPS (potentially with interesting correlations and entanglement), and $H(t) = \sum_\alpha h_\alpha$ is any geometrically local time-dependent Hamiltonian, where each term h_α acts on at most k adjacent physical sites. We dub this technique holographic quantum dynamics simulation (holoQUADS). Together with well-established interferometric techniques [39,40], HoloQUADS enables a simulation of arbitrary unequal-time correlation functions using only $\sim$poly$(t)\log_2 \mathcal{Q} + \log_2(\chi)$ qubits, which is *independent* of ℓ [41]. By comparison, the classical resources required to simulate time evolution from MPS initial states generically scale exponentially with t due to rapid entanglement growth, even in 1D systems [42].

A. Holographic quantum dynamics simulation (holoQUADS)

If the dynamics we care about is naturally generated by some circuit of depth r, we can proceed immediately with a holographic circuit construction as detailed below. If we are concerned with continuous time evolution under a Hamiltonian $H(t)$, the first step is to approximate the resulting unitary by a circuit consisting of r layers. There are many ways this approximation can be accomplished (see Refs. [16,17] for a helpful review of some state-of-the-art Hamiltonian simulation techniques), with different techniques having different scalings of r with time, qubit number, and desired simulation accuracy. For our purposes, it suffices to note that algorithms exist for which, at fixed error, r scales no worse than $\sim\ell^\varepsilon t^{1+\varepsilon}$ for any $\varepsilon > 0$ [16,43]. We consider this scaling to be essentially linear in t and *independent* of ℓ. Taking a more pragmatic approach, we note that simple product formulas based on Trotterization generally require a number of layers that is far smaller than most rigorous bounds suggest, and—at least for local observables like two-body correlation functions—produce accurate results using a modest (and system-size independent) number of time steps.

Next, we seek a holographic description of the state resulting from applying this discretized evolution to the initial MPS. To do so, it is useful to adopt space-time-inspired terminology to discuss different geometrical regions of the circuit. Denote the layer of the circuit by a discrete time index τ and the position along the spin chain by x. Each wire in the circuit has an implied directionality, as shown by the arrows in Fig. 8. We define the past light cone of the point (x, τ) to be the set of all points (x', τ') from which one can arrive at (x, τ) by flowing along the circuit in a forward direction, exiting gates along any outgoing wire. Inspection of Fig. 8 shows that measurements on the first k ($k = 2$ in this example) sites of the chain depend on the past light cone of the kth qubit (gray-shaded region of Fig. 8). The circuit in this region can be implemented using only the physical qubits for the first $r + k - 1$ sites (with r scaling polynomially in t for fixed error) along with the $\log_2(\chi)$ bond qubits: First implement the unitary circuits to prepare the MPS state within the gray-shaded region at $\tau = 0$, and then apply the layers of gates from $\tau = 0$ to $\tau = r$ that fall within this region. The first k physical qubits at $\tau = r$ can be measured in any desired basis, and then reset and reused
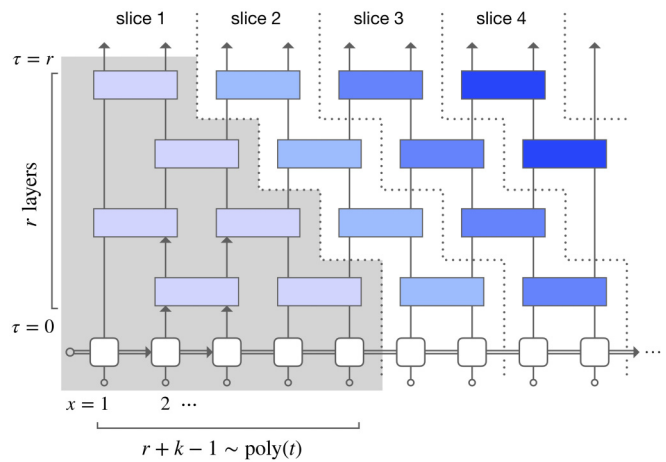

FIG. 8. Holographic time evolution of a matrix product state for nearest neighbor interactions ($k = 2$). The circuit can be evaluated by first executing all gates in the past causal cone (gray shaded region) of the qubits exiting the top left corner of the circuit. The remaining slices can be executed in order by (a) resetting the qubits exiting the top of the previous slice, (b) using the reset qubits to extend the MPS in space, and (c) applying all time-evolution gates in the current slice.

to represent the next k physical qubits at $\tau = 0$. These can be initialized into the correct state for the initial MPS by a horizontal (left-to-right) sequence of interactions with the bond qubits and then propagated diagonally to $\tau = r$ by acting with the remaining gates lying within the past causal cone of the $2k$th qubit at $\tau = r$, labeled as the diagonal "slice 2" in Fig. 8.

Repeating this process, one implements the full time-evolution circuit from left to right by sequentially implementing left-facing diagonal slices of the circuit. Effectively, the top of the circuit is being sheared and reattached to the bottom, such that the actual circuit fits naturally on the geometry of a cylinder as drawn in Fig. 9. By measuring the physical qubits at desired space-time points, one can reconstruct the time-ordered correlation functions of any local operators (relevant for dynamical response both near and far from equilibrium). Further, out-of-time ordered correlators that provide insight into thermalization, scrambling, and many-body quantum chaos can be simulated by including intervals of reverse-time evolution.

One can alternatively view this construction as a generalization of the holographic-MPS representation obtained by slicing the time-evolved-MPS circuit into diagonal slices whose boundaries are left-future-null-trajectories, and considering all qubit lines entering the slice from the left as "bond qubits" and those exiting the slice vertically as physical qubits. With this interpretation, one can measure entanglement Renyi entropies of the physical chain by measuring the corresponding quantities of the bond qubits as described above for general holographic MPS.

B. Comparison to classical methods

We now contrast this method with classical time-evolution techniques for simulating quantum dynamics. In 1D systems, the leading classical method for simulating time evolution of
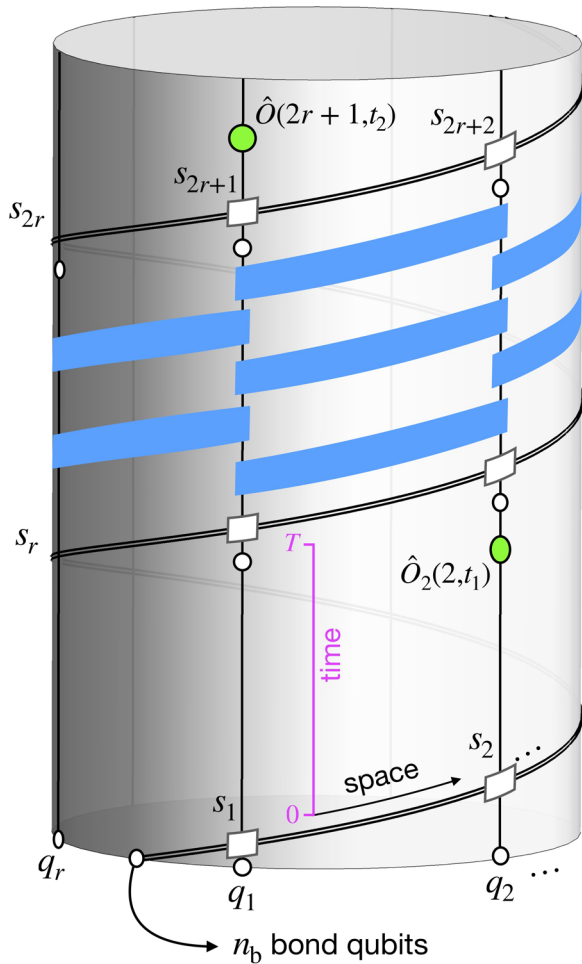
FIG. 9. Circuit for holographic time evolution of an MPS, obtained from Fig. 8 by attaching wires exiting the top of the circuit to those at the bottom wherever the respective exiting qubits are reset and reused at the bottom of the circuit. The space direction wraps diagonally around the cylinder indefinitely (the site of the physical system at position x is labeled s_x), and the circumference of the cylinder, which determines the required number of qubits, is determined by the evolution time. A time-ordered correlation function $\langle \hat{O}(x_1, t_1) \hat{O}(x_2, t_2) \rangle$ is obtained by measuring the corresponding operators at appropriate places in the circuit, as shown here (green circles representing measurements) for the example $x_1 = 2$, $x_2 = 2r + 1$.

an MPS is time-evolving block decimation (TEBD). TEBD works by converting an infinitesimal time evolution step $e^{-iH\Delta t} \approx (1 - iH\Delta t)$ into a matrix-product operator (MPO), applying the MPO to the initial MPS state, and reinterpreting the result as an MPS with bond space given by the tensor product of the MPO and MPS bond spaces.

At each stage, the tensors of the MPS are compressed (if possible) by discarding subleading singular values below the target accuracy. The compression step is effective if little entanglement is generated during the Trotter step compared to the maximum possible, but does not save classical resources in cases where significant entanglement is generated during each Trotter step, e.g., as in a strong quench with a thermalizing Hamiltonian. In fact, in extreme examples with maximally entangling dynamics, such as simulating stroboscopic dynam-

ics of random circuits [44,45], no compression whatsoever can be achieved. In contrast, holoQUADS exhibits polynomial scaling of the required qubit resources with evolution time regardless of the amount of entanglement generated per Trotter step, and will exhibit a maximal advantage in cases where low-rank classical compression is ineffective, or in higher dimensions where bond dimension can be prohibitively high for classical simulation from the outset.

As an aside, we note that applying an MPO to an MPS (which forms the basis of many classical methods) cannot be directly implemented holographically by unitary circuits plus measurement, since application of an MPO does not generally preserve the right canonical form of an MPS. We leave as an open question for future work whether holoQUADS can be generalized to incorporate different, non-MPO-based quantum analogs of such compression schemes that operate efficiently on exponentially large bond spaces [46–48] to further save on qubit resources.

V. DISCUSSION

These holographic methods will provide a quantum advantage for situations where classical MPS techniques are intractable due to prohibitively high bond dimension. Physically relevant examples include ground-state preparation of systems in dimensions $D > 1$, and time evolution under thermalizing Hamiltonians. State-of-the-art classical techniques run out of steam for 2D spin systems of widths of around 10 spins (less for gapless systems), or 1D time evolution with thermalizing dynamics over a few tens of interaction times. Holographic quantum algorithms might provide quantum advantage on these tasks with as few as 30–40 qubits, which are employed to directly tackle the difficult, highly entangled quantum aspects of these problems, rather than spending these resources to capture large sections of Hilbert space that are not accessed in physically relevant systems.

For higher dimensional simulations, the classical simulation cost for MPS techniques grows exponentially in the system width, even for area-law entangled states. One could employ a holoVQE method analogous to higher dimensional DMRG, obtained by slicing the system into a 1D stack of $(D - 1)$-dimensional cross-sectional slices, and specifying a circuit architecture to implement a quantum channel connecting one slice to another. The exponentially large bond dimension of the resulting MPS could be captured with polynomially many bond qubits. Importantly, for physically relevant low-energy states, MPS provide an exponential compression over a full wave-function description even in $D > 1$, so that the holographic representation affords substantial savings in qubit resources.

A more intrinsically higher dimensional generalization of these holographic methods would be to implement holographic simulation of isometric tensor networks [49]. These recently constructed higher dimensional generalizations of MPS capture a broad range of correlated states including (nonchiral) long-range entangled topological orders [50], and can be implemented straightforwardly as unitary circuits between physical qubits (now for each site in the lattice rather than each cross-sectional slice) and ancillary bond qubits. The isometric tensor networks invoke an explicit ordering

of operations and can be recast as a quantum channel that can be implemented holographically by resetting and reusing physical qubits that have already completed their participation in the circuit. The advantage of this technique over the usual boustrophedonic sweeping for 2D DMRG techniques is that it imposes a natural geometrically local 2D structure onto the physical and bond qubits.

It would also be desirable to extend these techniques to treat fermionic systems, using fermionic MPS representations [51], in order to simulate electronic materials. Other topics for future study include identifying effective heuristics for it-

eratively increasing bond dimension to improve the holoVQE accuracy, and developing a formal and systematic understanding to the class of states that can be efficiently represented holographically using large qubit numbers, but reasonable circuit depths, to implement each MPS tensor.

ACKNOWLEDGMENTS

[1] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer, Hybrid Quantum-Classical Approach to Correlated Materials, Phys. Rev. X **6**, 031045 (2016).

[2] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya *et al.*, Quantum chemistry in the age of quantum computing, Chem. Rev. **119**, 10856 (2019).

[3] M. C. Banuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero *et al.*, Simulating lattice gauge theories within quantum technologies, Eur. Phys. J. D **74**, 165 (2020).

[4] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, Phys. Rev. Lett. **98**, 140506 (2007).

[5] I. H. Kim, Holographic quantum simulation, arXiv:1702.02093.

[6] I. H. Kim, Noise-resilient preparation of quantum many-body ground states, arXiv:1703.00032.

[7] I. H. Kim and B. Swingle, Robust entanglement renormalization on a noisy quantum computer, arXiv:1711.07500.

[8] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, Towards quantum machine learning with tensor networks, Quantum Sci. Technol. **4**, 024001 (2019).

[9] F. Barratt, J. Dborin, M. Bal, V. Stojevic, F. Pollmann, and A. G. Green, Parallel quantum simulation of large systems on small quantum computers, npj Quantum Inf. **7**, 79 (2021).

[10] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, Variational quantum eigensolver with fewer qubits, Phys. Rev. Research **1**, 023025 (2019).

[11] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf, Sequential Generation of Entangled Multiqubit States, Phys. Rev. Lett. **95**, 110503 (2005).

[12] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, Demonstration of the trapped-ion quantum CCD computer architecture, Nature (London) **592**, 209 (2021).

[13] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland, Deterministic quantum teleportation of atomic qubits, Nature **429**, 737 (2004).

[14] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland, Realization of quantum error correction, Nature **432**, 602 (2004).

[15] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).

[16] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, Quantum algorithm for simulating real time evolution of lattice Hamiltonians, SIAM J. Comput., FOCS18-250 (2018), doi: 10.1137/18M1231511.

[17] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, A Theory of Trotter Error, Phys. Rev. X **11**, 011020 (2021).

[18] In general, these matrices can be rectangular, and the bond dimension can depend on j. In this case, we can embed every matrix in a larger square matrix with dimension $\max_j \chi_j$.

[19] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, Quantum Inf. Comput. **7**, 401 (2007).

[20] S. Gopalakrishnan and A. Lamacraft, Unitary circuits of finite depth and infinite width from quantum channels, Phys. Rev. B **100**, 064309 (2019).

[21] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, Commun. Math. Phys. **144**, 443 (1992).

[22] For local correlation functions where $\hat{O}_i = \mathbb{1}_i$ on all but a fixed number of sites, the variance of the observable should not scale with system size, and statistical accuracy can be achieved with a system-size independent number of repetitions. Nonlocal observables with weight on extensively many sites may have high variance and cannot in general be efficiently estimated.

[23] M. Levin and X.-G. Wen, Detecting Topological Order in a Ground State Wave Function, Phys. Rev. Lett. **96**, 110405 (2006).

[24] A. Kitaev and J. Preskill, Topological Entanglement Entropy, Phys. Rev. Lett. **96**, 110404 (2006).

[25] F. Pollmann, A. M. Turner, E. Berg, and M. Oshikawa, Entanglement spectrum of a topological phase in one dimension, Phys. Rev. B **81**, 064439 (2010).

[26] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, Rev. Mod. Phys. **91**, 021001 (2019).

[27] A. K. Ekert, C. M. Alves, D. K. L. Oi, M. Horodecki, P. Horodecki, and L. C. Kwek, Direct Estimations of Linear and Nonlinear Functionals of a Quantum State, Phys. Rev. Lett. **88**, 217901 (2002).

[28] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Rényi entanglement entropy via randomized measurements, Science **364**, 260 (2019).

[29] J. Wu and T. H. Hsieh, Variational Thermal Quantum Simulation Via Thermofield Double States, Phys. Rev. Lett. **123**, 220502 (2019).

[30] D. Zhu, S. Johri, N. M. Linke, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Y. Matsuura, T. H. Hsieh, and C. Monroe, Variational generation of thermofield double states and critical ground states with a quantum computer, Proc. Nat. Academy Sci. **117**, 25402 (2020).

[31] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 4213 (2014).

[32] F. Pollmann, Symmetry protected topological phases in one-dimensional systems, *Topological Aspects of Condensed Matter Physics: Lecture Notes of the Les Houches Summer School 2014* (Oxford University Press, Oxford, 2017), Vol. 103, p. 361.

[33] L. D. Faddeev, Algebraic aspects of the Bethe ansatz, Int. J. Mod. Phys. A **10**, 1845 (1995).

[34] Running two instances in parallel simply allows us to achieve smaller error bars for a given total run time.

[35] K. Mølmer and A. Sørensen, Multiparticle Entanglement of Hot Trapped Ions, Phys. Rev. Lett. **82**, 1835 (1999).

[36] We note that the anomalous exponent η controlling the decay of the longitudinal spin correlations is an exception to the above rule and is only qualitatively captured for the small bond dimensions explored here.

[37] D. Hanneke, J. P. Home, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland, Realization of a programmable two-qubit quantum processor, Nat. Phys. **6**, 13 (2010).

[38] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. Eriksson, K.-M. C. Fu *et al.*, Quantum simulators: Architectures and opportunities, PRX Quantum **2**, 017003 (2021).

[39] M. Knap, A. Kantian, T. Giamarchi, I. Bloch, M. D. Lukin, and E. Demler, Probing Real-Space and Time-Resolved Correlation Functions with Many-Body Ramsey Interferometry, Phys. Rev. Lett. **111**, 147205 (2013).

[40] B. Swingle, G. Bentsen, M. Schleier-Smith, and P. Hayden, Measuring the scrambling of quantum information, Phys. Rev. A **94**, 040302(R) (2016).

[41] The exact polynomial dependence on the evolution time t, as well as the dependence on k, depends on the chosen Hamiltonian simulation algorithm.

[42] A few interesting but nongeneric exceptions to this rule include disordered, glassy dynamics of many-body localized systems, or dynamics very close to the ground state such as in a quench under a local perturbation.

[43] A. M. Childs and Y. Su, Nearly Optimal Lattice Simulation by Product Formulas, Phys. Rev. Lett. **123**, 050503 (2019).

[44] A. Nahum, J. Ruhman, S. Vijay, and J. Haah, Quantum Entanglement Growth Under Random Unitary Dynamics, Phys. Rev. X **7**, 031016 (2017).

[45] C. W. Von Keyserlingk, T. Rakovszky, F. Pollmann, and S. L. Sondhi, Operator Hydrodynamics, Otocs, and Entanglement Growth in Systems Without Conservation Laws, Phys. Rev. X **8**, 021013 (2018).

[46] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, Nat. Phys. **10**, 631 (2014).

[47] P. Rebentrost, A. Steffens, I. Marvian, and S. Lloyd, Quantum singular-value decomposition of nonsparse low-rank matrices, Phys. Rev. A **97**, 012327 (2018).

[48] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, Quantum Sci. Technol. **2**, 045001 (2017).

[49] M. P. Zaletel and F. Pollmann, Isometric Tensor Network States in Two Dimensions, Phys. Rev. Lett. **124**, 037201 (2020).

[50] T. Soejima, K. Siva, N. Bultinck, S. Chatterjee, F. Pollmann, M. P. Zaletel *et al.*, Isometric tensor network representation of string-net liquids, Phys. Rev. B **101**, 085117 (2020).

[51] N. Bultinck, D. J. Williamson, J. Haegeman, and F. Verstraete, Fermionic matrix product states and one-dimensional topological phases, Phys. Rev. B **95**, 075108 (2017).