

Adiabatic ground state preparation in an expanding lattice

Christopher T. Olund¹,[✉] Maxwell Block,¹ Snir Gazit,² John McGreevy³,[✉] and Norman Y. Yao^{1,4}

¹*Department of Physics, University of California at Berkeley, Berkeley, California 94720, USA*

²*Racah Institute of Physics and The Fritz Haber Research Center for Molecular Dynamics, The Hebrew University of Jerusalem, Jerusalem 91904, Israel*

³*Department of Physics, University of California at San Diego, La Jolla, California 92093, USA*

⁴*Materials Science Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*



(Received 2 March 2020; revised manuscript received 8 April 2020; accepted 9 April 2020; published 30 April 2020)

We implement and characterize a numerical algorithm inspired by the s source framework [B. Swingle and J. McGreevy, *Phys. Rev. B* **93**, 045127 (2016)] for building a quantum many-body ground state wave function on a lattice of size $2L$ by applying adiabatic evolution to the corresponding ground state at size L , along with L interleaved ancillae. The procedure can in principle be iterated to repeatedly double the size of the system. We implement the algorithm for several one-dimensional (1D) spin model Hamiltonians, and find that the construction works particularly well when the gap is large and, interestingly, at scale-invariant critical points. We explain this feature as a natural consequence of the lattice expansion procedure. This behavior holds for both the integrable transverse-field Ising model and nonintegrable variations. We also develop an analytic perturbative understanding of the errors deep in either phase of the transverse-field Ising model, and suggest how the circuit could be modified to parametrically reduce errors. In addition to sharpening our perspective on entanglement renormalization in 1D, the algorithm could also potentially be used to build states experimentally, enabling the realization of certain long-range correlated states with low-depth quantum circuits.

DOI: [10.1103/PhysRevB.101.155152](https://doi.org/10.1103/PhysRevB.101.155152)

I. INTRODUCTION

A deep lesson of late 20th-century physics is the renormalization group (RG) philosophy: many-body physics is organized scale by scale. The fruits of this lesson have been assimilated well into our understanding of classical statistical physics and of perturbative quantum field theory [1,2]. In strongly correlated quantum systems, however, we still have a great deal to learn, in particular about eigenstates and even ground states of local model Hamiltonians.

Most of many-body Hilbert space is fictional, at least in the sense that it cannot be reached from a product state by time evolution with local Hamiltonians in a time polynomial in system size [3]. Ground states of local Hamiltonians are even more special: generically (with few exceptions arising from an overabundance of gapless excitations) the entanglement entropy of large-enough subregions satisfies an area law [4]. This statement is supported by a great deal of evidence, and has been rigorously proved for gapped systems in one dimension (1D) [5].

Importantly, much of the area-law corner of Hilbert space can be efficiently parametrized using tensor networks. This has been done with several different tensor network geometries, such as matrix product states (MPS) [6–9] in 1D, and projected entangled pair states (PEPS) [10–13] and isometric tensor networks [14,15] in two dimensions (2D). These parametrizations have proven to be very effective variational

Ansätze in a wide range of circumstances [16,17].¹ In particular, the density matrix renormalization group (DMRG) algorithm can be understood as a variational optimization on the MPS manifold [22,23]. Despite their successes, variational algorithms based on area-law tensor network *Ansätze* face some limitations. Specifically, in gapless phases, or at critical points, entanglement entropy can diverge with subsystem size making these area-law tensor networks suboptimal variational manifolds. It is also known that there exist even area-law states that do not have an efficient MPS representation [24]. Finally, many tensor networks are difficult to efficiently optimize in $D > 1$ [17,25].

Developing numerical methods for gapless phases and critical points requires understanding a richer entanglement structure than area-law states exhibit; we must account for the amount of entanglement at each length scale. The process of organizing our understanding of the entanglement in a quantum state scale by scale is sometimes called *entanglement renormalization* [26,27]. So far, the best-developed implementation of this idea is the multiscale entanglement renormalization *Ansatz* (MERA), which is a state-of-the-art

¹We note that there are also *nonvariational* algorithms for finding ground states that make use of tensor networks, some of which are provably efficient in some circumstances [18–21].

variational *Ansatz* for the study of 1D quantum critical points [28–32]. MERA has also inspired several variants such as deep MERA (DMERA) [33] and an analytic construction continuous MERA (cMERA) [34,35].

Despite the successes of MERA, developing a deeper understanding of entanglement renormalization remains a key challenge in condensed matter physics. More generally, existing tensor network methods leave room for improvement in several ways. First, the numerical values of the optimal tensors found in this way are difficult to interpret or directly relate to analytic results; the procedure is essentially a black box (though we note that in certain situations analytical constructions of MERA have been found using wavelets [36,37]). Second, and more practically, the variational minimization of the expectation value of the Hamiltonian requires sweeping across the lattice many times, an often costly procedure which has many opportunities to get stuck in locally optimal configurations.

In this paper, we introduce and benchmark a numerical algorithm for entanglement renormalization that takes small steps toward alleviating some of these issues. In particular, we provide a numerical implementation of the so-called s source framework, originally introduced in Ref. [38]. We note that the purpose of this work is to implement s source and characterize its accuracy; we leave a rigorous resource analysis to future studies.

We now briefly describe the s source formalism; a more thorough explanation is provided in Ref. [38]. Let H_L be a Hamiltonian defined on a d -dimensional lattice of size L^d and $|\psi^L\rangle$ be the associated ground state. The Hamiltonian family $\{H_L\}$ belongs to an s source fixed point if $|\psi^{2L}\rangle$ can be constructed by applying a quasilocal unitary U to s copies of $|\psi^L\rangle$ and some unentangled ancilla degrees of freedom. In many cases, we expect the adiabatic theorem to provide a construction of such a quasilocal unitary: if there is a gapped path from H_L to H_{2L} then adiabatic evolution along this path will suffice. There is evidence that many known states are s source fixed points, including trivial insulators ($s = 0$), chiral insulators ($s = 1$), and various field theories [38]. Examples with $s > 1$ are known as well, including fracton models [39–41]. Belonging to an s source fixed point constrains the growth of entanglement with system size, and in particular when $s < 2^{d-1}$ implies an area law for the entanglement entropy of subregions [38]. While the construction in Ref. [38] is more general, we will specialize our numerical exploration to one-dimensional spin chains.

The key advantage of s source is its ability to generate long-range entangled states using a constant depth circuit for $s \geq 1$. We illustrate this by comparing it to some more intuitive state preparation schemes. It is well known that building highly entangled states from a product state with local gates requires extensively deep quantum circuits [42]. Even with $|\psi^L\rangle$ as a resource, not all renormalization schemes generate long-range entanglement. In 1D, for example, one might consider concatenating two copies of a ground state end to end, and then acting with a local unitary to “glue” the states together. Unfortunately, constructing a long-range entangled state in this manner is not possible since the local unitary cannot strongly entangle distant spins in the two halves. In the s source framework in 1D, we circumvent this issue by

intercalating L ancilla spins between the spins that make up $|\psi^L\rangle$, thus expanding the underlying lattice. Crucially, this implies that a quasilocal unitary only needs to *locally* redistribute the rescaled entanglement structure. We will refer to the state formed by interleaving ancillae and $|\psi^L\rangle$ as the “ $s = 1$ input state” or just the “ s source input state” where $s = 1$ is to be understood. In contrast, we would call a product state of $2L$ spins an “ $s = 0$ input state.”

To adapt the s source construction to a numerical setting, one must decide how to implement the quasilocal unitary. One possibility would be to perform quasiadiabatic evolution via an algorithm like time-evolving block decimation (TEBD) [43,44]. In this work, we take an alternative route and fix a tensor network structure informed by the analytic Trotterization of the adiabatic evolution, which we refer to as the s source tensor network. The tensor network we obtain closely resembles a single layer of the MERA tensor network, and repeated application indeed results in a flavor of MERA. In principle, the Trotterized adiabatic unitary provides an exact expression for the s source tensor network, but explicitly calculating it is generically computationally intractable. Instead, we seek to *find* the corresponding tensor network through other means. One approach, which is possible in certain limiting cases, is to make use of the Trotterized structure to determine analytic expressions for the constituent tensors. More generally, we can determine the s source network variationally by minimizing the expectation value of H_{2L} with $|\psi^L\rangle$ as input [Fig. 1(a)]. At its core, since we are still using a *variational* approach to identify the s source tensor network, our prior concerns of becoming stuck in a locally optimal configuration still apply.

Although we still resort to variational optimization, the existence of the analytic expression defining s source allows us to physically interpret the resulting network and encourages us to optimize it in novel ways. For example, when considering multilayer s source [red boxes in Fig. 1(b)], we think of each layer as an independent adiabatic expansion and optimize it *separately* by minimizing the energy of the appropriate Hamiltonian at that scale, given the input generated by the preceding layers. We emphasize that this is slightly different from the usual optimization of MERA, where one sweeps over the entire network multiple times. In general, our numerical procedure incurs larger errors than optimizing across all layers, and hence is suboptimal when compared to standard global optimization approaches. Even so, it is less computationally expensive and, as we will explore below, can still perform well in certain cases of interest. In addition, by utilizing a greedy numerical algorithm, our implementation allows us to numerically test the validity of the adiabatic construction at the heart of the s source approach.

We will benchmark our implementation’s performance on the standard transverse-field Ising model (TFIM)

$$H_{\text{TFIM}} = -J \sum_{\langle ij \rangle} \sigma_i^z \sigma_j^z - h \sum_i \sigma_i^x \quad (1)$$

as well as the TFIM with integrability-breaking perturbations [45,46]. The TFIM has several limits that provide intuition about how s source should behave in general. The first limit of interest is deep in the ferromagnetic phase ($h < J$), where the ground state at finite size is approximately a symmetric

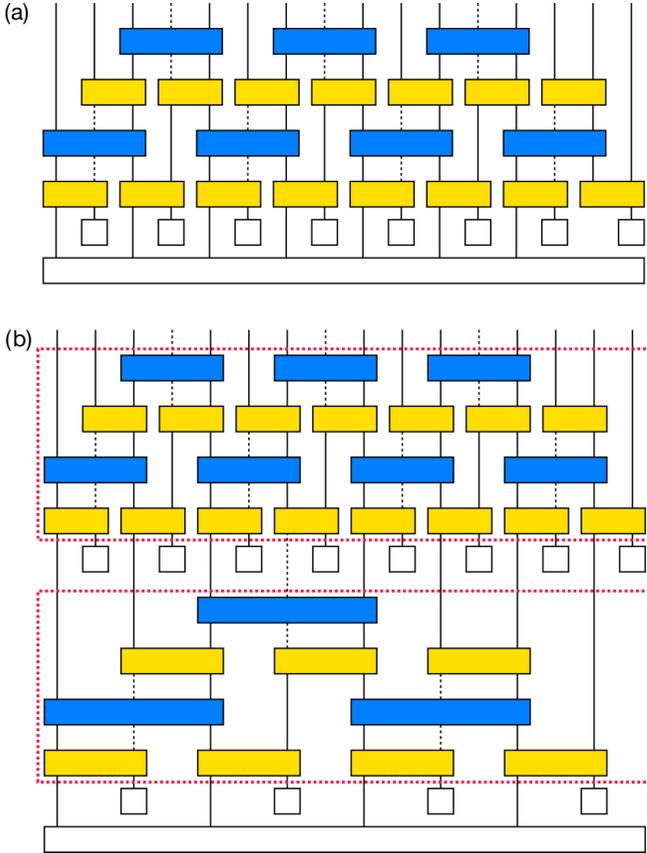


FIG. 1. Circuit diagrams for the s source renormalization procedure. (a) A single layer of the circuit which takes an 8-spin ground state (the large black box) and 8 ancillae (small black boxes) and after applying the turn “on” (yellow) and turn “off” (blue) unitaries produces an approximation to the 16-spin ground state. (b) A two-superlayer circuit which starts from the 4-spin ground state and produces an approximation to the 16-spin ground state; each block outlined by a red dashed line represents a superlayer made up of two on and two off layers.

Greenberger-Horne-Zeilinger (GHZ) state, i.e., $(|\uparrow \dots \uparrow\rangle + |\downarrow \dots \downarrow\rangle)/\sqrt{2}$ (as opposed to the symmetry-broken state such as $|\uparrow \dots \uparrow\rangle$) that one would typically consider in the thermodynamic limit). Building a GHZ state from a product state with local gates requires a circuit of extensive depth, but with a size L GHZ input the size $2L$ GHZ state can be prepared using a single layer of nearest-neighbor controlled-NOT (CNOT) gates. Similarly, the exact finite-size ferromagnetic ground state at finite magnetic field is another simple example of a state with long-range entanglement that cannot be built from a product state with a finite depth circuit. One would expect the same to be true of a gapped ground state with nontrivial topological order: while we would have no hope of building such a state with a low-depth circuit from a product input, an s source input could allow one to construct a good approximation.

In contrast, the paramagnetic phase ($h > J$) is easier to approximate with a product state input. Deep in the phase the ground state is almost a product state, and the correlations that do exist are short range. Using an s source input naturally doubles the length scale of those correlations and, thus, to

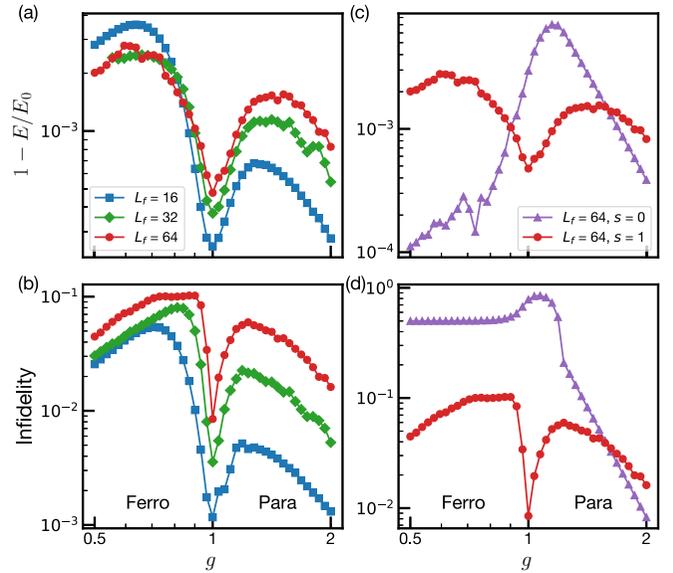


FIG. 2. (a) Relative energy error and (b) infidelity of the TFIM s source ground state as a function of g for several system sizes. Both go to 0 deep in either phase as one would expect, but there is also a local minimum at the critical point due to the scale invariance of the system. (c) Relative energy error and (d) infidelity for $L_f = 64$ with either the normal s source input ($s = 1$) or product state input ($s = 0$). Using a product input state gives better energies deep in both phases, but has an infidelity over 0.5 in the ferromagnetic phase as long-range entanglement cannot be generated. By either metric, $s = 1$ input gives substantially smaller errors near the critical point.

build the ground state one has to first *remove* those unwanted correlations before building the desired ones back in; for a product state input, we would only have to do the latter. Even if the s source constructed state has low error, our effort is wasted; we could have done even better with less work by starting with a product state.

The most interesting case is at the critical point. Here, we generically observe a local minimum in the error as a function of the transverse-field strength (Fig. 2). In some ways this is quite surprising; the existence of the analytic s source construction relies on the adiabatic theorem, which in turn requires a gap. Of course, there will always be a gap due to finite system size; however, one naively expects that such a small finite-size gap would force one to use a longer adiabatic evolution time, thus incurring larger Trotter errors when approximating the adiabatic unitary with a local circuit. However, the scale invariance of the TFIM critical point makes it particularly amenable to approximation by s source. At the critical point, the correlation length of the ground state scales with system size, so when we insert ancillae and hence trivially double the length scale of correlations, we actually achieve the proper long-range entanglement structure. We then correct the short-range details with the local circuit. We note that the location of the error minimum remains at the critical point even when one adds generic perturbations to the TFIM, consistent with the expectation that this behavior should generalize to other continuous phase transitions.

Our paper is organized as follows. In Sec. II, we give a precise description of both the s source algorithm and our

numerical implementation. In Sec. III, we benchmark our numerical implementation by applying it to several standard 1D spin chain models: first, the (integrable) transverse-field Ising chain (TFIM), next, the TFIM with a longitudinal field which is nonintegrable and has no symmetries, and finally the TFIM with a symmetry-preserving but integrability-breaking term. In Sec. IV, we develop some analytic understanding of the circuit in the large-gap limit. Finally, in Sec. V, we summarize our results and discuss potential future directions of study.

II. s SOURCE ALGORITHM AND NUMERICAL IMPLEMENTATION

In the s source framework, we regard the entanglement present in the ground state at linear system size L as a resource for constructing the ground state at system size $2L$. Rather than attempting to directly prepare the macroscopic ground state of a model Hamiltonian, we suppose we are given s copies of the ground state at system size L , and design a circuit which doubles the system size. That is, we seek a unitary map which produces the ground state at size $2L$ from s copies of the ground state at size L times a collection of factorized ancillary qubits. Iterating this doubling procedure yields a circuit which produces the ground state in the thermodynamic limit from s copies of the (easily determined) ground state of a small cluster of sites. We note, as previously discussed, that such a size-doubling map can exist even when the state represents a nontrivial phase and cannot be constructed from a product state via a low-depth local unitary circuit.

Our numerical implementation will focus on s source with $s = 1$. When the Hamiltonian is gapped, one can immediately write an expression for the s source unitary using the adiabatic theorem. Let \tilde{H}_L be the operator that acts as H_L on the odd lattice sites only. Now, consider a time-dependent Hamiltonian $H(t)$ which interpolates between

$$H(0) = \tilde{H}_L - \sum_{i \text{ even}} X_i \quad (2)$$

and

$$H(T) = H_{2L}. \quad (3)$$

Here, X_i are operators which put the ancillary qubits into a product ground state. The unitary operator which generates this time evolution is then

$$U = \mathcal{T} e^{-i \int_0^T H(t) dt}. \quad (4)$$

Of course, if we could generally compute the full adiabatic unitary explicitly, we could also solve the *easier* problem of just finding the exact ground state at system size $2L$! We can, however, imagine Trotterizing this unitary to get an approximation built out of *local* unitaries that is tractable enough to make further progress. When the gap is large, we can find these component local unitaries analytically, as we explore in Sec. IV A.

For the moment, however, we observe that even without actually doing the time-ordered integral, one can see upon which spins the local unitaries act; the terms in the leading order of the Trotter expansion will act on the same spins as do terms in either H_L or H_{2L} . Since we will work with nearest-neighbor

Hamiltonians, one can think of the terms coming from H_{2L} as turning on interactions between the spins of our original L site system and the ancillae (which are now nearest neighbors after the interleaving step), and we can interpret the terms coming from H_L as turning off interactions between the original spins (which are no longer nearest neighbors). Keeping these leading-order terms, we get an approximate tensor network for U as shown schematically in Fig. 1(a). Although we justify the circuit structure perturbatively, we will see from our numerics that it is still capable of generating approximate ground states even when a perturbative expansion would not converge. The order of the layers is in principle arbitrary, although some choices are more computationally efficient than others. We also note that one could choose to Trotterize into larger blocks and that doing so would improve the approximation in exchange for circuit optimization becoming much more computationally expensive. Later, in Sec. IV B, we will see exactly how introducing longer-range blocks reduces errors deep in either phase of the transverse-field Ising model.

In our numerical implementation, we treat the tensor network as a variational *Ansatz* built out of arbitrary unitaries. We minimize $\langle H_{2L} \rangle$ over those component unitaries to get an approximation for $|\psi^{2L}\rangle$. As the reader may have noted, the circuit that we end up obtaining is, in fact, a MERA, albeit one with a particular circuit structure and where we have cut off some number of layers at the smallest scale. However, we are thinking of this MERA as being “upside down”; rather than starting with a large state and repeatedly coarse graining, we start with a small state and scale up.

There is a fundamental tension between making the adiabatic evolution time T larger to reduce adiabaticity errors and making T smaller to reduce Trotterization errors for a fixed depth circuit. This tension disappears in the extreme limit of a large gap wherein we can determine U analytically, as we will describe in Sec. IV A.

We now describe the actual circuit *Ansatz* used, and explain how we numerically optimize it to find an approximate ground state. Suppose we have a solution for the ground state of H for an L spin system $|\psi^L\rangle$ in matrix product state (MPS) form. We construct the $2L$ spin input state $|\phi^{2L}\rangle$ by identifying spin i ($1 \leq i \leq L$) of the L particle system with spin $2i - 1$ of the $2L$ spin system, and then placing ancillary spins on the remaining sites. We note that the orientation of these ancillae does not matter as any single spin rotation can be absorbed into the circuit. Next, we construct a quantum circuit described by a total unitary U_T . We build this circuit in four layers: (i) applying two spin unitaries U_A^i to each pair of spins $(2i - 1, 2i)$ for $1 \leq i \leq L$, (ii) applying unitaries U_B^i to pairs of spins $(4i - 3, 4i - 1)$ for $1 \leq i \leq L/2$, (iii) applying unitaries U_C^i to pairs $(2i + 1, 2i)$ for $1 \leq i \leq L - 1$, and finally (iv) applying U_D^i to pairs $(4i + 1, 4i - 1)$ for $1 \leq i \leq L/2 - 1$. The unitaries U_A^i and U_C^i correspond to turning on the new couplings between the original L spins and the ancillae, and the unitaries U_B^i and U_D^i correspond to turning off the couplings between the original spins. A schematic of this setup for $L = 8$ can be seen in Fig. 1(a). We can also repeat this procedure multiple times, successively inserting ancillae and then applying four layers of the circuit (which we will henceforth call a “superlayer”) to repeatedly double the size

of the input state. As an example, a two-superlayer circuit is illustrated in Fig. 1(b).

In order to numerically optimize the circuit, we minimize the energy $E = \langle \tilde{\psi}^{2L} | H_{2L} | \tilde{\psi}^{2L} \rangle$ where $|\tilde{\psi}^{2L}\rangle = U_T |\phi^{2L}\rangle$. In particular, we begin with an initial circuit (which could either be a random circuit or an educated guess) and then consider E to be a function of each of the individual local unitaries comprising U_T . We then sweep over all of these component unitaries multiple times using the conventional MERA update procedure described in [30]. For the interested reader, we provide some additional details about the numerical optimization in Appendix A.

As mentioned above, this optimization procedure generally only finds a local minimum of the energy; if one wants to reliably find the global minimum, it is necessary to do this variational search many times with different initial conditions. To optimize a multilayer circuit, we pursue a greedy algorithm: for each superlayer, we minimize the expectation value of H_{2L} over the unitaries in the L to $2L$ layer with all preceding layers held fixed. The intuition behind this approach is that the adiabatic construction should in principle guarantee the existence of a multilayer circuit such that its first k superlayers generate the ground state at size $2^k L$. To be more precise, if we consider the analytic construction where we have the exact quasilocal adiabatic unitaries at our disposal, we know that one can construct the state $|\psi^{4L}\rangle$ from $|\psi^L\rangle$ by applying the adiabatic unitary $U^{L \rightarrow 2L}$ to get $|\psi^{2L}\rangle$ and then $U^{2L \rightarrow 4L}$ to get $|\psi^{4L}\rangle$. This suggests that a greedy approximation of each layer could in principle be effective. Of course, reoptimizing all superlayers at each scale is at least as accurate and in some cases may yield much lower errors. However, full circuit optimization comes at a significant computational cost, and we find that the greedy approach performs surprisingly well. Before presenting our numerical benchmarking results, we wish to emphasize that our approach, following the s source philosophy, attempts only to find an optimal *adiabatic* trajectory, namely, one which utilizes information from previous layers. In particular, we do *not* attempt a global energy minimization, as with standard MERA optimization schemes.

III. NUMERICAL BENCHMARKING RESULTS

To benchmark our numerical implementation we consider three 1D models: the transverse-field Ising model (TFIM), a mixed-coupling Ising model (MCIM), and a mixed-field Ising model (MFIM), with Hamiltonians

$$\begin{aligned} H_{\text{TFIM}} &= -J \sum_{(ij)} \sigma_i^z \sigma_j^z - h \sum_i \sigma_i^x, \\ H_{\text{MCIM}} &= - \sum_{(ij)} (J_x \sigma_i^x \sigma_j^x + J_z \sigma_i^z \sigma_j^z) - h \sum_i \sigma_i^x, \\ H_{\text{MFIM}} &= -J \sum_{(ij)} \sigma_i^z \sigma_j^z - \sum_i (h_x \sigma_i^x + h_z \sigma_i^z). \end{aligned} \quad (5)$$

The TFIM sets our baseline understanding for how $s = 1$ s source performs in three limits: a short-range correlated unique ground state (the paramagnetic phase), an almost-degenerate long-range correlated ground state (the ferromagnetic phase), and at a critical point.

We quantify our implementation's performance using both the relative error in energy (which we minimize) and the many-body infidelity, i.e., the overlap mismatch between the s source state obtained at size $2L$ and the "exact" DMRG wave function at the same size $1 - |\langle \tilde{\psi}^{2L} | \psi^{2L} \rangle|^2$. As a point of reference, we compare this performance to that of $s = 0$ s source, optimizing the same circuit structure with a product state input. We also study the consequences of truncating our approximation of the quasilocal unitary to include only nearest-neighbor gates. Finally, in order to understand the propagation of errors in our numerical s source algorithm, we analyze the performance of multilayer circuits.

Benchmarking via the TFIM model. Since the TFIM is integrable, in this case, we calculate energy errors relative to the exact values. For the MCIM and MFIM models, we benchmark against energies obtained via DMRG. In addition, we use DMRG to generate our initial input MPS states for s source for all three models (restricting to a specific \mathbb{Z}_2 parity sector when appropriate). In Fig. 2(a), we plot the relative error in energy for a single layer of s source for the TFIM as a function of $g = h/J$ for several values of the input system size L_0 . To be specific, this means that we start with the ground state at L_0 and perform a single layer of our s source algorithm to obtain an approximate ground state at $L_f = 2L_0$, whose energy we then compare with the exact value. Similarly, Fig. 2(b) depicts the many-body infidelity, which exhibits the same qualitative behavior. In all of our numerics, we ensure that the input state in the ferromagnetic ($g < 1$) phase is the non-symmetry-broken ground state.

As one expects, the error decreases deep in either the ferromagnetic or paramagnetic phase. Indeed, because the gap is large in these regions, there must exist a suitable s source adiabatic unitary that minimizes both nonadiabatic and Trotter errors. Less expected, from this adiabatic perspective, is the existence of a local minimum in the error at the TFIM's critical point $g = 1$, despite the fact that the gap vanishes at this point. Naively, one might have expected that this would lead to an error maximum instead. In fact, this is exactly what does happen if we start with a product state input ($s = 0$) instead of the $s = 1$ s source input state, as can be seen in Fig. 2(c).

To understand this $s = 1$ local minimum, we note that the correlation length diverges at the critical point and it is impossible to capture these correlations starting from a product state and using a low-depth local quantum circuit. However, if we start with the size L ground state (as we do in $s = 1$ s source), then correlations of length L become correlations of length $2L$ upon ancillae insertion. In principle, at the critical point, this is exactly what we desire from the size $2L$ ground state; we emphasize once again that this is precisely the same intuition which underlies MERA and that our circuit is in fact a type of MERA with a "cutoff" at small scales.

To further check this intuition, we can define a single-site energy error for the TFIM as

$$\epsilon(i) = -\frac{J}{2} (\sigma_{i-1}^z \sigma_i^z + \sigma_i^z \sigma_{i+1}^z) - h \sigma_i^x, \quad (6)$$

and then take a Fourier transform to define a momentum-resolved energy error $\epsilon(k)$. Doing this, we found that the

momentum-resolved error was only significant for momenta of $k = 0, \pi/4$, and $\pi/2$ (k in units of inverse lattice spacing). The $k = 0$ component is just the total energy error, whereas the $k = \pi/4$ and $\pi/2$ components correspond to errors of characteristic length scale 2 and 1 lattice spacings, respectively. These are, of course, exactly the length scales at which the nearest-neighbor and next-nearest-neighbor gates comprising our circuit act. There is no corresponding dip in $\epsilon(k = \pi/4)$ or $\epsilon(k = \pi/2)$ at the critical point, consistent with our understanding that the dip in the overall energy error really is due to ancilla insertion and not the local action of the circuit.

While this built-in doubling of input correlations is beneficial at criticality, it can be detrimental in other regimes. This can be seen by comparing the performance of $s = 1$ s source with $s = 0$ s source deep in the paramagnetic phase [$g > 1$ in both Figs. 2(c) and 2(d)]. While both errors are scaling toward zero as g increases, the scaling is worse for the $s = 1$ input. Here, the true ground state is short-range correlated, approaching a product state for large g . Thus, constructing the size $2L$ ground state with an $s = 1$ input actually involves first getting *rid* of all the doubled correlations.

Looking only at energy errors [Fig. 2(c)], the above statement would also appear to apply deep in the ferromagnetic phase ($g < 1$). However, the many-body infidelity tells a different story. In particular, although the energy errors for $s = 0$ s source scale better than $s = 1$ s source, the fidelity does not [Fig. 2(d)]. To understand this behavior, we note that at $g = 0$, the ground state manifold of the TFIM is twofold degenerate, consisting of the symmetric and antisymmetric cat states. For finite but small g , these states will be split in energy by an exponentially small gap $\sim g^L$. Until g is nearly one, any linear combination of these two states will have approximately the same energy, and an “all up” like combination can be constructed from a product state input to give a low energy error. However, it is impossible to construct a cat state from a product state using a circuit with subextensive depth.

As a result, the infidelity of the $s = 0$ state is always greater than 0.5 throughout the entire ferromagnetic phase, as the zeroth-order piece of the true ground state cannot be constructed. In contrast, with an $s = 1$ input state, a size L cat state can be used to create a size $2L$ cat state by using controlled-NOT gates between each pair of original and ancilla spins. We expect that this behavior should generalize to certain classes of topological states. In particular, because one cannot change a topological invariant by acting with local unitaries, it is impossible to build such states from a product state input. On the other hand, using an $s = 1$ input preserves the topological character of the state.

Next, we turn to studying the effect of changing the range of the quasilocal unitary approximation by restricting our circuit to include *only* nearest-neighbor gates. A comparison of the resultant energy errors and infidelities is shown in Fig. 3. The nearest-neighbor circuit still exhibits a local error minimum at the critical point, and in fact nearly achieves the accuracy of the longer-ranged s source circuit there. This implies that the long-range correlations built in via ancillae insertion and the ability to perform nearest-neighbor corrective gates are the most important features for accurately constructing a state at the critical point. Moving away from

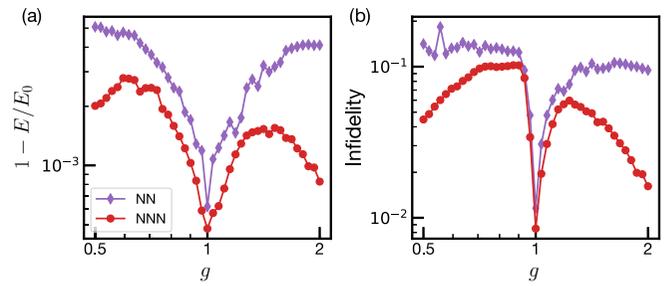


FIG. 3. (a) Relative energy error and (b) infidelity of the TFIM for both the “standard” s source circuit with both nearest-neighbor (NN) and next-nearest-neighbor (NNN) unitaries and a simplified circuit with only NN unitaries, both for $L_f = 64$. The general shape of the error curves, notably including the local minimum at the critical point, is similar for both circuits. The NN circuit does almost as well as the NNN circuit at the critical point, but the errors fall off more slowly than for the NNN circuit deep in either phase.

the critical point, one sees the advantage of the next-nearest-neighbor circuit geometry; we will show in Sec. IV that the range of individual gates in the circuit determines the scaling of the error with g deep in either phase.

Multilayer s source. Our preceding discussion focuses on single-layer s source, where one starts with a size L_0 input state and ends with a size $L_f = 2L_0$ final state. In multilayer s source, we start from a size L_0 state and perform the s source construction n times to get an approximate size $L_f = 2^n L_0$ state; we use the approximate state from one superlayer as the input state for the next. As aforementioned, in our numerics, we take a “greedy” approach where we optimize each superlayer in isolation rather than sweeping back and forth. In principle, simultaneously optimizing superlayers should improve accuracy, but it would come at a substantial computational cost.

Figure 4 depicts the energy errors for multilayer s source for the TFIM as a function of both L_0 and L_f . Although it is difficult to make sharp statements, it appears that errors are not accumulating, per layer, superlinearly. We hypothesize that the multilayer error obeys the following bound: $\mathcal{E}_{L \rightarrow 4L} \leq \mathcal{E}_{L \rightarrow 2L} + \mathcal{E}_{2L \rightarrow 4L}$, where $\mathcal{E}_{L_0 \rightarrow L_f}$ is the relative energy error for the optimal L_0 to L_f s source state. The analogous statement for accumulated infidelities holds trivially (if one were to optimize the s source circuit by minimizing infidelity instead of energy). In Fig. 4(c), one can see that this proposed bound appears to hold for a two-superlayer circuit.

Nonintegrable models. Finally, we now turn to applying the single-layer s source algorithm to the nonintegrable MCIM and MFIM models. In Fig. 5, we begin by showing the energy errors as a function of input size for the MCIM model. The qualitative features of the error curve are analogous to what we have already discussed in the TFIM case; the error decreases deep in either phase, and there is still a local minimum in the error at the phase transition. This minimum is consistent with our prior expectations since the critical point of the MCIM is still described by a conformal field theory. In Fig. 6, we plot the energy errors for a number of different longitudinal field strengths for the MFIM model. As h_z is

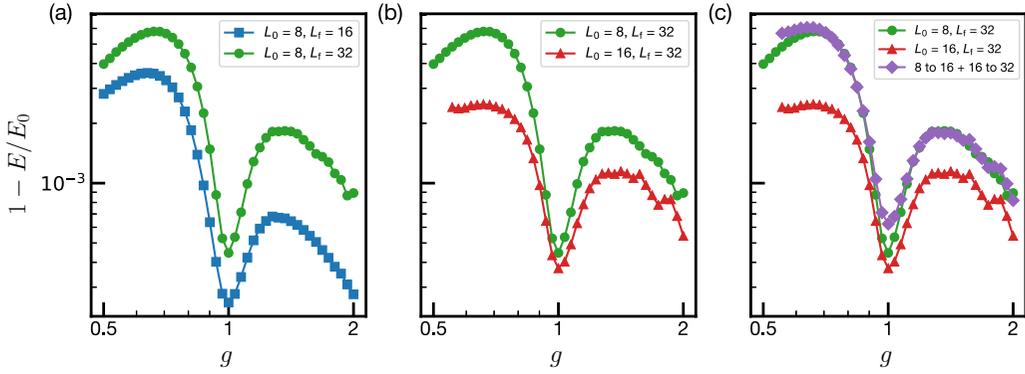


FIG. 4. Relative energy error in the TFIM for multilayer circuits. (a) Error comparison, starting with $L_0 = 8$ for one- and two-layer circuits. (b) Error comparison ending with $L_f = 32$ for one- and two-layer circuits. Error compounds reasonably with successive layers, and in particular is not much worse than single-layer optimization at the critical point. (c) We hypothesize that the multilayer error is subadditive, as illustrated here. Except where we have failed to find the global minimum, the $L_0 = 8$ to $L_f = 32$ error is bounded by the $L_0 = 8$ to $L_f = 16$ error plus the $L_0 = 16$ to $L_f = 32$ error, with the multilayer circuit substantially outperforming this bound at the critical point.

increased, the local minimum flattens out and then vanishes, consistent with the lack of a scale-invariant point.

IV. ANALYTIC ANALYSIS OF ERRORS

A. Analytic tensors in the large field limit

Our effort to express the adiabatic s source unitary as a local, low-depth circuit faces two competing constraints. In order to have a good approximation, we need to be able to use both the adiabatic theorem (which requires that T^{-1} be small compared to the gap) and the Trotter decomposition (which requires that T^{-2} be large compared to commutators between different blocks of the Hamiltonian). In particular, in the large h/J limit of the TFIM, one can satisfy both of the above constraints. By moving into the interaction picture and expanding the time-ordered exponential to leading order, we find (see Appendix B for details) that the nearest-neighbor

(“on”) unitaries are (Fig. 1), to leading order in h/J , given by $U = e^{-iH_{\text{eff}}}$, where

$$H_{\text{eff}} = -\frac{J}{8h}(\sigma_z^1\sigma_y^2 + \sigma_y^1\sigma_z^2). \quad (7)$$

Similarly, the next-nearest-neighbor (“off”) unitaries (Fig. 1) are given by $U = e^{iH_{\text{eff}}}$ with the same H_{eff} .

One can perform the same calculation in the mixed field/coupling models. Defining $h = \sqrt{h_x^2 + h_z^2}$ and $\tan \eta = h_z/h_x$, we obtain an effective Hamiltonian

$$H_{\text{eff}} = \alpha(\sigma_x^1\sigma_y^2 + \sigma_y^1\sigma_x^2) + \beta(\sigma_y^1\sigma_z^2 + \sigma_z^1\sigma_y^2), \quad (8)$$

where

$$\alpha = \frac{J_x}{32h}(7 \sin \eta + 3 \sin 3\eta) - \frac{3J_z}{8h} \cos^2 \eta \sin \eta, \\ \beta = \frac{3J_x}{8h} \cos \eta \sin^2 \eta - \frac{J_z}{32h}(7 \cos \eta - 3 \cos 3\eta) \quad (9)$$

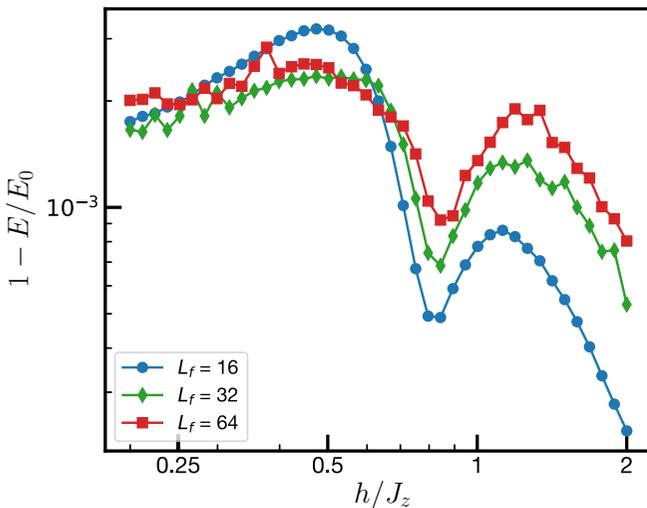


FIG. 5. Relative energy error for the mixed coupling Ising model with $J_x = 0.1J_z$. The features of the error curve are qualitatively similar to the transverse-field case, and the local minimum remains at the model’s critical point.

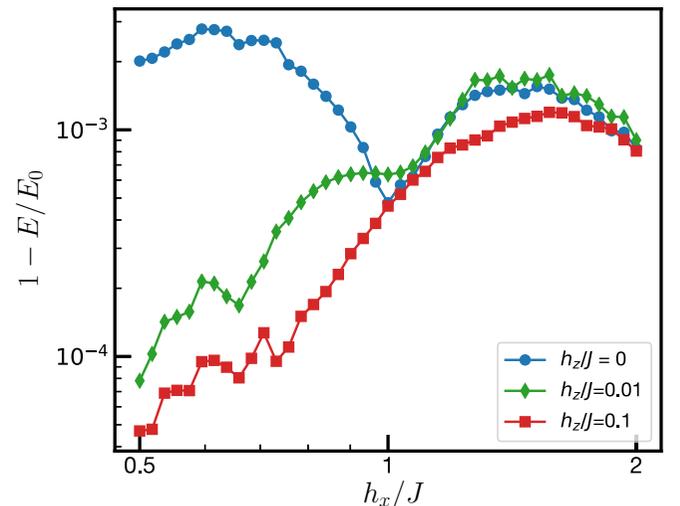


FIG. 6. Relative energy error for the mixed-field Ising model for $L_f = 64$. For sufficiently large h_z , there is no longer a local minimum in the error as there is no scale-invariant point.

and again the “on” and “off” unitaries are given by $U = e^{-iH_{\text{eff}}}$ and $e^{iH_{\text{eff}}}$, respectively.

B. Perturbative analysis

In this section, we will use a perturbative analysis to explore how the error of the optimal s source circuit varies with our system parameters. In our numerics, we are ultimately using a variational (in energy) method to solve for the circuit; thus, one cannot analytically calculate the error directly, but in a perturbative regime, we can compute how the leading-order correction to the s source wave function scales. As a bonus, this procedure naturally suggests additional tensors one could include in the circuit to further suppress errors. Although including such tensors would come with a computational cost for our numerical implementation, it is possible that they could be more natural for certain experimental geometries where long-range interactions are present [47–51].

Here, we present the results of this analysis for the TFIM. Our basic strategy is to first figure out the exact s source circuit at a fixed point (i.e., either $h/J = 0$ or $J/h = 0$), and then to perturb around this fixed point. In other words, once we factor out the fixed-point portion, we will parametrize each unitary in the circuit as $\exp(-iH) = \mathbb{I} - iH + \dots$ with $\|H\| = \mathcal{O}(h/J)$ or $\mathcal{O}(J/h)$. We then consider whether there exists such a circuit that will take our initial state to the target state correctly to a given order in perturbation theory.

For the $J/h = 0$ fixed point (with ancillae inserted in the direction of the field), the s source circuit is simply the identity and reproduces the size $2L$ ground state perfectly since both $|\psi^L\rangle$ and $|\psi^{2L}\rangle$ are product states. Deep in the paramagnetic phase, we can construct the correct size $2L$ ground state to linear order in J/h , and in fact, the circuit we use to do so is precisely based on the analytic unitaries found in the previous subsection. The leading-order errors that remain are pairs of spin flips at sites $4i$ and $4i + 4$ with coefficients of order $(J/h)^2$.

If we modify our circuit to contain two spin unitaries acting on pairs of spins located four sites away, we can construct the state *correctly* to order $(J/h)^2$. More generally, if one continues to add unitaries up to distance $2n$, one can faithfully construct the state to order $(J/h)^n$. The strategy is to first remove unwanted terms to bring the state back to the $J/h = 0$ state and then to build in the needed terms. Both distance n and $2n$ gates are required, the former to create the needed n th-order terms, and the latter to remove unwanted n th-order terms introduced when we add the ancillae. Furthermore, the effective Hamiltonians that parametrize these unitaries will be exponentially weak in distance, so the overall unitary will indeed be quasilocal as expected. Consistent with our numerical results, starting with a product state rather than the size L ground state allows us to do better; in particular, we will only need distance n gates in order to be correct to order $(J/h)^n$ since there are no unwanted terms to remove. We emphasize that this intuition is only true for sufficiently small J/h .

In the ferromagnetic phase things are a bit more subtle. We will denote the ground state which is “connected” to the symmetric cat state at $h = 0$ as $|\psi_0^L\rangle$ (even parity) and the analogous state which is connected to the antisymmetric

cat state as $|\psi_1^L\rangle$ (odd parity). The energy splitting between these states will scale like $(h/J)^L$. Were we to only care about energy errors, we might reasonably consider any linear combination of $|\psi_0^{2L}\rangle$ and $|\psi_1^{2L}\rangle$ to be our target state. However, we know that one cannot turn a product state into a cat state or vice versa with a circuit of subextensive depth. As previously discussed, at $h/J = 0$ one can go from a cat state input to a cat state output by inserting the ancillae in the $|\uparrow\rangle$ state and using CNOT gates for the bottom layer of the circuit. Thus, if our starting state is $|\psi_0^L\rangle$, then our target state will be $|\psi_0^{2L}\rangle$. We could also consider starting with either a product state or the symmetry-breaking linear combination $(|\psi_0^L\rangle + |\psi_1^L\rangle)/\sqrt{2}$ and building toward $(|\psi_0^{2L}\rangle + |\psi_1^{2L}\rangle)/\sqrt{2}$; in this case the circuit at $h/J = 0$ is the identity. If we insist that the target is $|\psi_0^{2L}\rangle$ and begin with a product state input, it is impossible to be correct to even zeroth order.

To this end, with the standard circuit [Fig. 1(a)], one can prepare $|\psi_0^{2L}\rangle$ correctly only to $\mathcal{O}(1)$, whereas for the symmetry-broken or product inputs we can prepare $(|\psi_0^{2L}\rangle + |\psi_1^{2L}\rangle)/\sqrt{2}$ correctly to order h/J . The leading-order error in the former case comes in the form of double-spin flips on each of the two product states that make up the cat state. Using a three-site unitary with the third spin acting as a control would let us correct this error and prepare the true ground state correctly to order h/J . Similarly, blocks of $2n + 1$ sites will allow us to correctly prepare the state to order $(h/J)^n$ by eliminating correlated $2n$ spin-flip errors. In contrast with the paramagnetic phase, where the long-distance unitaries only needed to act on two sites, here the gates need to act on all of the sites within a block.

When targeting the symmetry-broken state, the standard circuit will give us the correct answer to order h/J for both the product state and symmetry-broken inputs. However, to correct higher-order errors using $(|\psi_0^L\rangle + |\psi_1^L\rangle)/\sqrt{2}$ as an input requires $(2n + 1)$ -site unitaries to obtain the correct result at order $(h/J)^{n+1}$; the analogous situation with a product state input requires n -site unitaries to be correct to order $(h/J)^n$. In all cases, the strengths of the required unitaries fall off exponentially with the diameter of the block, and the overall circuit is once again quasilocal.

V. DISCUSSION AND CONCLUSION

In this work, we have developed a numerical implementation of the s source algorithm for finding approximate ground states of local Hamiltonians [38]. We approximate the lattice-doubling unitary of the s source algorithm as an efficiently contractable tensor network, which we in turn variationally optimize to minimize the energy of the doubled-lattice ground state. To ensure our tensor network is efficiently contractable, we construct it from local rather than quasilocal components, although this decreases the accuracy of the approximation. We benchmark the resulting numerical algorithm on several 1D spin chain models, and find that the s source construction works particularly well at scale-invariant critical points. We ascribe this to the fact that ancillae insertion doubles the length scale of all correlations in the input state, much in the same spirit as MERA. In addition, to gain some analytic intuition, we computed the scaling of the wave-function errors deep in each phase of the TFIM, and determined how the

s source circuit could be modified to correct these errors. These corrections are consistent with the expectation that performing s source with a quasilocal unitary should permit the exact construction of the doubled-lattice ground state.

Our work suggests several interesting directions for future study. First, one could use multilayer s source as a numerical method to extract information about renormalization group flow. When creating a multilayer s source circuit, one obtains a sequence of unitaries U_1, U_2, \dots, U_n that each double the size of the system. By parametrizing how this sequence of unitaries changes, it should be possible to follow the renormalization group flow and to extract quantities such as the operator dimension. In a similar vein, since the Hamiltonian parameters also flow, one might expect that the ideal input state to build $|\psi^{2L}(g)\rangle$ would not be $|\psi^L(g)\rangle$ but rather $|\psi^L(g')\rangle$, where the Hamiltonian parameter g' at length L flows to g at length $2L$. Allowing for this may significantly improve the performance of s source away from criticality.

While we do not foresee s source outperforming established methods like DMRG in determining 1D ground states, it may be useful for constructing ground states in 2D where existing methods leave more room for improvement. Furthermore, the algorithm could be naturally adapted as an experimental method for preparing ground states.

In fact, the ability to interleave ancillae has recently become possible in Rydberg optical tweezer arrays [50–52]. An experimental implementation of s source would be particularly useful for generating states with long correlation lengths. Indeed, as we have previously discussed, strongly correlated many-body states often require deep quantum circuits in order to be built from product states. Absent error correction, deep circuits result in low fidelities due to compounding gate errors. Thus, the ability of s source to create certain classes of strongly correlated states with low-depth circuits could provide a significant advantage in the noisy intermediate scale quantum (NISQ) era [53]. Finally, although the numerical implementation we explore here is variational, the s source formalism provides a compelling connection to a nonvariational ground state construction which merits future exploration.

ACKNOWLEDGMENTS

We acknowledge the insights of and discussions with C. Laumann, T. Schuster, F. Machado, K. Akkaravarawong, J. Motruk, W.-W. Ho, G. Evenbly, and I. Cirac. We are particularly grateful to M. Zaletel for suggesting the current version of our multilayer optimization approach and R. Mong for suggesting the idea of using s source as a method for numerically computing renormalization group flows. J.M. is grateful to B. Swingle for collaboration on the s source program. DMRG calculations were performed using the ITensor Library [54]. This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program and the U.S. Department of Energy under cooperative research agreement DE-SC0009919. S.G. acknowledges support from the Israel Science Foundation, Grant No. 1686/18. C.O. was supported by the Department of

Defense through the National Defense Science & Engineering Graduate Fellowship program.

APPENDIX A: CIRCUIT OPTIMIZATION

Finding a quantum circuit that prepares a minimum energy eigenstate is a challenging problem [55,56]. In the generic case, we cannot deterministically find the optimal circuit or even verify that a given circuit does minimize the energy. We can, however, perform a variational search by starting with some circuit and then iteratively updating component tensors in order to continually lower the energy, as is common practice for MERA.

Within a superlayer of the s source circuit, we optimize each local unitary u while holding all others fixed. We then iterate this procedure for all u in the superlayer until the energy has converged. To update a single unitary, we utilize the optimization strategy outlined in [30]. To summarize, we pretend that the function we are trying to maximize, $f(u) = -\langle H \rangle$, is a linear function of u . Of course, it is actually quadratic, as u^\dagger is present in the dual circuit. If we treat u^\dagger as constant, however, we can write $f(u) = \text{Tr}(W^\dagger u)$, where the environment W^\dagger of u is found by contracting the tensor network formed by removing u from $-\langle H \rangle$'s tensor network. If W has singular value decomposition $W = XYZ^\dagger$, it follows that a linear function $f(u)$ achieves its maximum at $u = XZ^\dagger$.

For a nonlinear function of u , one should in principle update u multiple times until the energy converges. In practice, we only update each tensor u once during a full sweep of the superlayer; we have empirically found that this leads to a lower energy for a fixed total number of updates. We also tested an alternative optimization strategy referred to as “Linearization II” in [57], but this required using 10 updates per unitary per superlayer sweep for numerical stability, as well as tuning additional hyperparameters. We did not see any benefit of this approach for fixed computational cost.

As one might expect, this update procedure generally only finds a local energy minimum, not a global one. In order to get the global minimum, we perform this optimization many times over circuits initialized with Haar-random unitaries. The cost of simulation is linear in the number of samples, which can be large, so we note a few tricks that will improve either the speed or performance of optimization (although sampling over initializations is embarrassingly parallel).

First, we note that only a subset of the terms in the Hamiltonian will be within the light cone of a given u , so we only need to minimize the partial energy containing those terms when we update u . Crucially, the number of terms that contribute is *constant* as a function of system size, whereas the number of terms in the full Hamiltonian scales as L . We also note that each optimization step does not, generally, decrease the energy. For a quadratic function of a given u this procedure will actually maximize the absolute value of the function. In order to avoid this complication, we alter the spectrum of the (partial) energy we are minimizing to be negative definite by shifting the partial Hamiltonian by an appropriate multiple of the identity [29]: $\hat{h} \rightarrow \hat{h} - \alpha I$, where α is the maximum eigenvalue of the partial Hamiltonian \hat{h} . In practice, we find that we only need to do this for a few sweeps before all partial energies are negative, at which point we turn the shift off as it

seems to slow down convergence (this suggests the possibility that shifting the spectrum up could speed up convergence as long as we are careful to keep things negative definite).

We now describe some heuristics for efficient contraction of the next-nearest-neighbor $s = 1$ *s* source circuit. Suppose we want to evaluate the expectation value of a term in the Hamiltonian

$$\langle \hat{O} \rangle = \langle \psi^L | U_1^\dagger \dots U_n^\dagger \hat{O} U_n \dots U_1 | \psi^L \rangle, \quad (\text{A1})$$

where U_i is the i th superlayer of a multilayer s source circuit. We start with the operator \hat{O} (which is defined on the $2^n L$ site lattice) and then, in the language of MERA, act upon it with the ascending superoperator [30] (in other words, we conjugate by the innermost superlayer of the circuit). If \hat{O} was supported on at most six adjacent sites, the ascended version of \hat{O} will be supported on either four or six adjacent sites on the $2^{n-1} L$ spin lattice. This is because the causal cone extends by at most six sites, and contracting with the ancillae halves the support of the operator at the end.

We emphasize that we can ignore contraction with any gate in U_n outside of \hat{O} 's causal cone since it will contract with its inverse in U_n^\dagger to form an identity [30]. There is no computational advantage to starting with a block of fewer than four sites (or six sites if the block would be ascended to a six-site block), so if we want to evaluate the sum of expectation values of many operators, we should group them into blocks of operators living on either four or six adjacent sites. Doing so allows us to do a single contraction to find the sum of the expectation values of several adjacent local observables instead of multiple contractions to find the expectation value of each term separately. This reduces the time it takes to evaluate the expectation value of the Hamiltonian considerably. Repeatedly applying ascending superoperators, we eventually obtain an operator defined on the L -site lattice, which when contracted with the MPS and its dual will give the desired expectation value.

It is advantageous to cache various partial contractions of the MPS portion of the tensor network (i.e., $\langle \psi^L |$ and $| \psi^L \rangle$). In particular, at the final step of evaluating an expectation value, we will contract the ascended \hat{O} with four or six pairs of adjacent physical indices of the MPSs, with all other MPS indices already contracted. By storing all contractions of the MPSs with four or six pairs of adjacent dangling bonds, we can avoid repeating this costly computation. The evaluation of the W^\dagger 's needed to optimize the circuit is done in much the same way, simply omitting the contraction with the specific u that is to be updated.

Energy minimization appears to take on the order of 1000 sweeps for the models we tested, with that number growing slightly with system size. This also varies from run to run; sometimes it might take 1000 sweeps, and sometimes it may take 10 000. For most of the figures in this paper starting with an initial state of 32 spins and creating an s source state of 64 spins we ran ~ 1000 initializations with 1000 sweeps each and took the best energy among them. For smaller systems, e.g., 8 to 16, we performed < 100 initializations.

We note that for this work we were particularly interested in characterizing the error of the s source algorithm, and as such we needed to find the global minimum as reliably as possible. For some other applications, one may be perfectly

content to have, say, twice the minimum error, in which case it is not as necessary to run so many randomly initialized optimizations. In this case one can often do pretty well by starting with a good guess for the initial circuit, adding some noise, and optimizing just a few initial states. For the TFIM, a good guess may be the leading-order analytic solution that we discuss in Sec. IV A, where noise is added by multiplying each unitary by another random unitary close to the identity. Here, we make two notes. First, with fewer parallel optimizations it is more important to do more sweeps for each one (several thousand rather than one thousand, say). Second, it is important to make sure that the initial condition of the circuit is not entirely real, as updating a real valued circuit will keep the circuit in the real manifold.

We note that further improvements are likely possible. It seems, for example, that it should be possible to reduce the average number of required sweeps by monitoring for convergence. However, checking for convergence can be quite deceptive here; one typically sees plateaus where the energy appears to converge, and then sudden jumps down to new local minima. A more careful analysis may reveal an effective way to anticipate whether or not further sweeps will result in an improved energy. In our experience, the energy would sometimes continue to improve beyond 1000 sweeps, but it was more efficient to sample more initial conditions than execute more sweeps per sample.

Finally, the cost of contraction scales roughly exponentially in the width of the circuit's causal cone. In practice, this might motivate the use of the simplified s source circuit comprised of only the nearest-neighbor gates. We analyzed this circuit in Sec. III and found that the energy error was qualitatively similar to that of the circuit containing both nearest-neighbor and next-nearest-neighbor gates, and was quantitatively not much worse at the critical point.

APPENDIX B: DERIVATION OF ANALYTIC UNITARIES FOR LARGE MAGNETIC FIELDS

Here, we derive the analytic expressions for the s source unitaries of the TFIM in the limit $h \gg J$, previously described in Sec. IV A. In this regime, we can simultaneously make the adiabatic evolution time T long enough to be adiabatic, but short enough that we can do a Trotter expansion. The former condition requires, for $h \gg J$, $hT \gg 1$. In the interaction picture that we will consider shortly, the Trotter expansion requires $JT \ll 1$. We consider a single term in the Trotter expansion of Eq. (4) and thus reduce the problem to considering two spins that are initially in a field of strength h in the x direction and then turning on an interaction of strength J in the z direction. If we slowly turn on the interaction over a time T , then we have

$$H(t) = H_0 + \frac{t}{T} H_1 \quad (\text{B1})$$

with

$$H_0 = -h(\sigma_x^1 + \sigma_x^2), \quad (\text{B2})$$

$$H_1 = -J\sigma_z^1\sigma_z^2. \quad (\text{B3})$$

Then, the adiabatic unitary associated with moving from H_0 to $H_0 + H_1$ is

$$U = \mathcal{T} \exp -i \int_0^T H(t) dt. \quad (\text{B4})$$

It is helpful for us to move to the interaction picture before proceeding. Doing so gives us the interaction picture unitary

$$\begin{aligned} U_I &= \mathcal{T} \exp -i \int_0^T \frac{t}{T} e^{iH_0 t} H_1 e^{-iH_0 t} dt \\ &\approx 1 - i \int_0^T \frac{t}{T} e^{iH_0 t} H_1 e^{-iH_0 t} dt. \end{aligned} \quad (\text{B5})$$

Upon integrating and discarding higher-order terms, we get

$$\begin{aligned} U_I &= 1 + i \left(\frac{JT}{4} (\sigma_z^1 \sigma_z^2 + \sigma_y^1 \sigma_y^2) \right. \\ &\quad \left. + \frac{J}{8h} e^{-ih(\sigma_x^1 + \sigma_x^2)T} (\sigma_z^1 \sigma_y^2 + \sigma_y^1 \sigma_z^2) e^{ih(\sigma_x^1 + \sigma_x^2)T} \right). \end{aligned} \quad (\text{B6})$$

Moving back to the Schrödinger picture and continuing to work to leading order,

$$\begin{aligned} U &= e^{-iH_0 T} U_I = (e^{-iH_0 T} U_I e^{iH_0 T}) e^{-iH_0 T} \\ &= 1 + i \left(hT (\sigma_x^1 + \sigma_x^2) + \frac{JT}{4} (\sigma_z^1 \sigma_z^2 + \sigma_y^1 \sigma_y^2) \right. \\ &\quad \left. + \frac{J}{8h} (\sigma_z^1 \sigma_y^2 + \sigma_y^1 \sigma_z^2) \right) \\ &= \exp i \left(hT (\sigma_x^1 + \sigma_x^2) + \frac{JT}{4} (\sigma_z^1 \sigma_z^2 + \sigma_y^1 \sigma_y^2) \right. \\ &\quad \left. + \frac{J}{8h} (\sigma_z^1 \sigma_y^2 + \sigma_y^1 \sigma_z^2) \right) \\ &= \exp(-iH_{\text{eff}}) \end{aligned} \quad (\text{B7})$$

which corresponds to an effective Hamiltonian

$$\begin{aligned} H_{\text{eff}} &= -hT (\sigma_x^1 + \sigma_x^2) - \frac{JT}{4} (\sigma_z^1 \sigma_z^2 + \sigma_y^1 \sigma_y^2) \\ &\quad - \frac{J}{8h} (\sigma_z^1 \sigma_y^2 + \sigma_y^1 \sigma_z^2). \end{aligned} \quad (\text{B8})$$

We have not as of yet specified a value for T , so its presence in our effective Hamiltonian may appear, at first glance, to be troubling. However, we expect from the adiabatic theorem that, as long as the assumptions are met, there should be no strong T dependence. Indeed, one can explicitly verify that in the limit $T \gg J/h^2$ the effect of the T -dependent terms is a phase shift. Dropping them, we end up with a particularly simple form for H_{eff} :

$$H_{\text{eff}} = -\frac{J}{8h} (\sigma_z^1 \sigma_y^2 + \sigma_y^1 \sigma_z^2). \quad (\text{B9})$$

The “off” unitaries, on the other hand, are given by $\exp(iH_{\text{eff}})$. We can see this by considering running the process backward in time. This is, of course, just the “turn on” problem we just solved. There is one additional complication: these unitaries are acting not on the ground state, but on the first layer of the circuit. However, corrections due to the noncommutation of the layers will come in at a higher order, and since we are only working to first order anyway we can simply ignore them.

We can repeat this analysis for a mixed coupling and field Ising model. Here we have Hamiltonians

$$H_0 = -h_x (\sigma_x^1 + \sigma_x^2) - h_z (\sigma_z^1 + \sigma_z^2) \quad (\text{B10})$$

and

$$H_1 = -J_x \sigma_x^1 \sigma_x^2 - J_z \sigma_z^1 \sigma_z^2. \quad (\text{B11})$$

If we define $h = \sqrt{h_x^2 + h_z^2}$, $\tan \eta = h_z/h_x$, and go through the same steps, we find that to order J/h

$$\begin{aligned} H_{\text{eff,XX}} &= -\frac{J_x T}{32} (9 + 4 \cos 2\eta + 3 \cos 4\eta) \sigma_x^1 \sigma_x^2 - \frac{J_x T}{4} \sin^2 \eta \sigma_y^1 \sigma_y^2 - \frac{3J_x T}{16} \sin^2 2\eta \sigma_z^1 \sigma_z^2 \\ &\quad - \frac{J_x T}{32} (2 \sin 2\eta + 3 \sin 4\eta) (\sigma_x^1 \sigma_z^2 + \sigma_z^1 \sigma_x^2) + \frac{J_x}{32h} (7 \sin \eta + 3 \sin 3\eta) (\sigma_x^1 \sigma_y^2 + \sigma_y^1 \sigma_x^2) \\ &\quad + \frac{3J_x}{8h} \cos \eta \sin^2 \eta (\sigma_y^1 \sigma_z^2 + \sigma_z^1 \sigma_y^2), \end{aligned} \quad (\text{B12})$$

$$\begin{aligned} H_{\text{eff,ZZ}} &= -\frac{3J_z T}{16} \sin^2 2\eta \sigma_x^1 \sigma_x^2 - \frac{J_z T}{4} \cos^2 \eta \sigma_y^1 \sigma_y^2 - \frac{J_z T}{32} (9 - 4 \cos 2\eta + 3 \cos 4\eta) \sigma_z^1 \sigma_z^2 \\ &\quad - \frac{J_z T}{32} (2 \sin 2\eta - 3 \sin 4\eta) (\sigma_x^1 \sigma_z^2 + \sigma_z^1 \sigma_x^2) - \frac{3J_z}{8h} \cos^2 \eta \sin \eta (\sigma_x^1 \sigma_y^2 + \sigma_y^1 \sigma_x^2) \\ &\quad - \frac{J_z}{32h} (7 \cos \eta - 3 \cos 3\eta) (\sigma_y^1 \sigma_z^2 + \sigma_z^1 \sigma_y^2), \end{aligned} \quad (\text{B13})$$

with an overall effective Hamiltonian $H_{\text{eff}} = H_0 T + H_{\text{eff,XX}} + H_{\text{eff,ZZ}}$. If we again drop the T -dependent terms, we get

$$\begin{aligned} H_{\text{eff}} &= \left(\frac{J_x}{32h} (7 \sin \eta + 3 \sin 3\eta) - \frac{3J_z}{8h} \cos^2 \eta \sin \eta \right) (\sigma_x^1 \sigma_y^2 + \sigma_y^1 \sigma_x^2) \\ &\quad + \left(\frac{3J_x}{8h} \cos \eta \sin^2 \eta - \frac{J_z}{32h} (7 \cos \eta - 3 \cos 3\eta) \right) (\sigma_y^1 \sigma_z^2 + \sigma_z^1 \sigma_y^2). \end{aligned} \quad (\text{B14})$$

- [1] H. E. Stanley, *Rev. Mod. Phys.* **71**, S358 (1999).
- [2] K. G. Wilson and J. Kogut, *Phys. Rep.* **12**, 75 (1974).
- [3] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, *Phys. Rev. Lett.* **106**, 170501 (2011).
- [4] M. M. Wolf, F. Verstraete, M. B. Hastings, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 070502 (2008).
- [5] M. B. Hastings, *J. Stat. Mech.: Theory Exp.* (2007) P08024.
- [6] S. Rommer and S. Östlund, *Phys. Rev. B* **55**, 2164 (1997).
- [7] I. P. McCulloch, *J. Stat. Mech.: Theory Exp.* (2007) P10014.
- [8] A. Szasz, J. Motruk, M. P. Zaletel, and J. E. Moore, [arXiv:1808.00463](https://arxiv.org/abs/1808.00463) [Phys. Rev. X (to be published)].
- [9] W. W. Ho, S. Choi, H. Pichler, and M. D. Lukin, *Phys. Rev. Lett.* **122**, 040603 (2019).
- [10] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [11] N. Schuch, D. Pérez-García, and I. Cirac, *Phys. Rev. B* **84**, 165139 (2011).
- [12] L. Vanderstraeten, J. Haegeman, and F. Verstraete, *Phys. Rev. B* **99**, 165121 (2019).
- [13] M. C. Bañuls, D. Pérez-García, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. A* **77**, 052306 (2008).
- [14] M. P. Zaletel and F. Pollmann, *Phys. Rev. Lett.* **124**, 037201 (2020).
- [15] T. Soejima, K. Siva, N. Bultinck, S. Chatterjee, F. Pollmann, and M. P. Zaletel, *Phys. Rev. B* **101**, 085117 (2020).
- [16] R. Orus, *Ann. Phys.* **349**, 117 (2014).
- [17] L. Vanderstraeten, J. Haegeman, P. Corboz, and F. Verstraete, *Phys. Rev. B* **94**, 155123 (2016).
- [18] Z. Landau, U. Vazirani, and T. Vidick, *Nat. Phys.* **11**, 566 EP (2015).
- [19] F. Verstraete and J. I. Cirac, *Phys. Rev. B* **73**, 094423 (2006).
- [20] I. Arad, Z. Landau, U. Vazirani, and T. Vidick, *Commun. Math. Phys.* **356**, 65 (2017).
- [21] B. Roberts, T. Vidick, and O. I. Motrunich, *Phys. Rev. B* **96**, 214203 (2017).
- [22] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [23] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
- [24] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 030504 (2008).
- [25] A. Anshu, I. Arad, and A. Jain, *Phys. Rev. B* **94**, 195143 (2016).
- [26] G. Vidal, *Phys. Rev. Lett.* **99**, 220405 (2007).
- [27] G. Vidal, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [28] R. N. C. Pfeifer, G. Evenbly, and G. Vidal, *Phys. Rev. A* **79**, 040301(R) (2009).
- [29] G. Evenbly and G. Vidal, Quantum criticality with the multi-scale entanglement renormalization ansatz, in *Strongly Correlated Systems: Numerical Methods*, edited by A. Avella and F. Mancini (Springer, Berlin, 2013), pp. 99–130.
- [30] G. Evenbly and G. Vidal, *Phys. Rev. B* **79**, 144108 (2009).
- [31] L. Cincio, J. Dziarmaga, and M. M. Rams, *Phys. Rev. Lett.* **100**, 240603 (2008).
- [32] Z. Li and R. S. K. Mong, *Phys. Rev. B* **99**, 241105(R) (2019).
- [33] I. H. Kim and B. Swingle, [arXiv:1711.07500](https://arxiv.org/abs/1711.07500).
- [34] J. Haegeman, T. J. Osborne, H. Verschelde, and F. Verstraete, *Phys. Rev. Lett.* **110**, 100402 (2013).
- [35] J. S. Cotler, M. Reza Mohammadi Mozaffar, A. Mollabashi, and A. Naseh, *Phys. Rev. D* **99**, 085005 (2019).
- [36] G. Evenbly and S. R. White, *Phys. Rev. Lett.* **116**, 140403 (2016).
- [37] J. Haegeman, B. Swingle, M. Walter, J. Cotler, G. Evenbly, and V. B. Scholz, *Phys. Rev. X* **8**, 011003 (2018).
- [38] B. Swingle and J. McGreevy, *Phys. Rev. B* **93**, 045127 (2016).
- [39] J. Haah, *Phys. Rev. B* **89**, 075119 (2014).
- [40] W. Shirley, K. Slagle, Z. Wang, and X. Chen, *Phys. Rev. X* **8**, 031051 (2018).
- [41] W. Shirley, K. Slagle, and X. Chen, *Phys. Rev. B* **99**, 115123 (2019).
- [42] A. Harrow and S. Mehraban, [arXiv:1809.06957](https://arxiv.org/abs/1809.06957).
- [43] G. Vidal, *Phys. Rev. Lett.* **91**, 147902 (2003).
- [44] G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [45] R. Cole, F. Pollmann, and J. J. Betouras, *Phys. Rev. B* **95**, 214410 (2017).
- [46] J. Li, R. Fan, H. Wang, B. Ye, B. Zeng, H. Zhai, X. Peng, and J. Du, *Phys. Rev. X* **7**, 031011 (2017).
- [47] B. Gadway and B. Yan, *J. Phys. B: At., Mol. Opt. Phys.* **49**, 152002 (2016).
- [48] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. Hollenberg, *Phys. Rep.* **528**, 1 (2013).
- [49] C. Monroe, W. C. Campbell, L. M. Duan, Z. X. Gong, A. V. Gorshkov, P. Hess, R. Islam, K. Kim, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, [arXiv:1912.07845](https://arxiv.org/abs/1912.07845).
- [50] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, *Science* **354**, 1021 (2016).
- [51] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, *Science* **354**, 1024 (2016).
- [52] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Phys. Rev. Lett.* **121**, 123603 (2018).
- [53] J. Preskill, *Quantum* **2**, 79 (2018).
- [54] ITensor Library, version 2.0.11, <http://itensor.org>.
- [55] I. Arad and Z. Landau, *SIAM J. Comput.* **39**, 3089 (2010).
- [56] I. L. Markov and Y. Shi, *SIAM J. Comput.* **38**, 963 (2008).
- [57] G. Vidal, [arXiv:0707.1454v1](https://arxiv.org/abs/0707.1454v1) (2009).