

Preparing valence-bond-solid states on noisy intermediate-scale quantum computers

Bruno Murta ^{1,2,*} Pedro M. Q. Cruz ^{1,3,4} and J. Fernández-Rossier ^{1,†}

¹*International Iberian Nanotechnology Laboratory (INL), 4715-330 Braga, Portugal*

²*Departamento de Física, Universidade do Minho, 4710-057 Braga, Portugal*

³*ICFO - Institut de Ciències Fotòniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels, Spain*

⁴*Faculdade de Ciências, Universidade do Porto, 4169-007 Porto, Portugal*



(Received 21 July 2022; accepted 10 February 2023; published 20 March 2023)

Quantum state preparation is a key step in all digital quantum simulation algorithms. Here we propose methods to initialize on a gate-based quantum computer a general class of quantum spin wave functions, the so-called valence-bond-solid (VBS) states, that are important for two reasons. First, VBS states are the exact ground states of a class of interacting quantum spin models introduced by Affleck, Kennedy, Lieb, and Tasaki (AKLT). Second, the two-dimensional VBS states are universal resource states for measurement-based quantum computing. We find that schemes to prepare VBS states based on their tensor-network representations yield quantum circuits that are too deep to be within reach of noisy intermediate-scale quantum (NISQ) computers. We then apply the general nondeterministic method herein proposed to the preparation of the spin-1 and spin-3/2 VBS states, the ground states of the AKLT models defined in one dimension and in the honeycomb lattice, respectively. Shallow quantum circuits of depth independent of the lattice size are explicitly derived for both cases, making use of optimization schemes that outperform standard basis gate decomposition methods. The probabilistic nature of the proposed routine translates into an average number of repetitions to successfully prepare the VBS state that scales exponentially with the number of lattice sites N . However, two strategies to quadratically reduce this repetition overhead for any bipartite lattice are devised. Our approach should permit to use NISQ processors to explore the AKLT model and variants thereof, outperforming conventional numerical methods in the near future.

DOI: [10.1103/PhysRevResearch.5.013190](https://doi.org/10.1103/PhysRevResearch.5.013190)

I. INTRODUCTION

Quantum many-body phenomena are ubiquitous in nature, but their study has been hampered by the difficulty [1] in carrying out large-scale numerical simulations to probe their defining emergent features [2]. Digital quantum computers hold the promise of a more scalable simulation of quantum many-body phenomena than what is possible with conventional numerical methods by exploiting the principle of superposition and the natural encoding of entanglement [3]. Indeed, some of the key limitations faced by leading numerical methods, notably the sign problem [4] in quantum Monte Carlo methods, the exponential wall problem [5] faced by exact diagonalization, and the entanglement area laws [6] that hinder the viability of tensor networks beyond one dimension, are overcome by gate-based quantum computation [7,8].

However, digital quantum simulation is yet to become a standard method in the study of quantum many-body systems. This is due to the limitations [9] of the current noisy

intermediate-scale quantum (NISQ) computers [10], namely the accumulation of both coherent and incoherent errors due to faulty gate operations [11], and the limited coherence times of qubits that restrict the depth of circuits that can be faithfully executed. Although these issues will eventually be overcome through quantum error correction [12], meanwhile can the strengths of NISQ devices be exploited to solve practical problems beyond what is possible via conventional numerical methods?

Two complementary approaches have been followed towards this goal. The first involves the development of hybrid variational algorithms (cf. Ref. [13] and references therein) that trade circuit depth for the parallel execution of independent circuits and delegate part of the computational task to conventional processors. The second consists of identifying problems for which the prospect of a quantum speed-up [14] is plausible, either because conventional methods are not scalable or because the problem itself is especially suited for NISQ devices.

In the vein of the second approach, several proposals [15–22] involving the preparation of exact eigenstates of integrable quantum many-body models on digital quantum computers have arisen in the literature recently. Even though analytical methods allow to probe many features of these models, the computation of other properties remains challenging. For example, although the spectrum of the quantum XXZ model was determined exactly over sixty years ago [23], calculating some of its correlation functions remains an active

*bpmurta@gmail.com

†Also at Departamento de Física Aplicada, 03690 San Vicente del Raspeig, Universidad de Alicante, Spain.

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

area of research [24], which motivated the development of a routine to prepare the corresponding Bethe ansatz states on quantum hardware [20,21]. Likewise, the addition of a sufficiently strong magnetic field to the Kitaev model [25] disrupts the fractionalization upon which its integrability is based, in which case a hybrid variational scheme on a NISQ device [22] may constitute a competitive strategy.

The class of quantum many-body states that will be the focus of this article are the so-called *valence-bond-solid* (VBS) states, which are the exact ground states of Affleck-Kennedy-Lieb-Tasaki (AKLT) models [26,27]. The ground state of the bilinear-biquadratic spin-1 AKLT model is the landmark VBS state, because it provided the first strong piece of evidence to support Haldane’s conjecture [28] that the integer-spin Heisenberg models are gapped in one dimension. Nevertheless, the VBS construction can be generalized to arbitrary lattices in higher dimensions [26,27]. One that will deserve our special attention is the spin-3/2 VBS state on the honeycomb lattice, as this is a resource state for universal quantum computation [29,30] within the measurement-based paradigm of Raussendorf and Briegel [31–33].

Measurement-based quantum computation rests upon the preparation of an initial state that has neither too high nor too low a degree of entanglement [34,35]. Given such a resource state, single-qubit operations and local measurements suffice to drive the quantum computation, thus avoiding the need to implement two-qubit operations. Ideally, a resource state should be the nondegenerate ground state of some naturally occurring or experimentally realizable physical system, as its preparation would then boil down to cooling the system [36]. The spin-3/2 VBS state is one such resource state [29,30], and its experimental realization has been pursued in different solid-state platforms [30,37,38]. Numerical simulations performed on conventional hardware [37,39–41] that could support this quest for the realization of the spin-3/2 VBS state are not scalable. Using hybrid variational methods [13,42–46], near-term quantum computers can in principle offer an advantage in the simulation of such quantum spin-3/2 models. To this end, the preparation of an educated initial state is crucial. Such state is the spin-3/2 VBS, and its preparation on quantum hardware will be explored in the remainder of this paper. It should be stressed, though, that our methods apply generally to any VBS state, regardless of the local spin or the dimensionality of the underlying lattice.

Recently, the problem herein considered of preparing VBS states on quantum hardware has received great attention. In particular, we would like to draw attention to three manuscripts [47–49] made available online after the submission for publication of this paper. Chen *et al.* [47] adopted a probabilistic strategy akin to ours, although the symmetrization operator was approximated through an ansatz that yields a greater CNOT count than our method, which is exact. Only the one-dimensional VBS state was considered, but proof-of-concept simulations were performed on an IBM quantum processor. Smith *et al.* [48] devised a constant-depth deterministic scheme to initialize the one-dimensional VBS state by exploiting its $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry to recover from unsuccessful measurement outcomes. The implementation of the method on IBM quantum hardware was considered, but its application to the preparation of two-dimensional VBS states was left

for future work. Wei, Malz, and Cirac [49] built upon earlier work [50] that proved injective PEPS defined on a N -site lattice can be prepared adiabatically in $\mathcal{O}(\text{polylog}(N))$ time, identifying a specific adiabatic path to prepare the spin-3/2 VBS state. Given the need to implement the time-evolution operator of local terms with support on six qubits, this method to prepare the spin-3/2 VBS state is likely to become the leading option only when (early stage) fault-tolerant quantum computers become available.

II. MOTIVATION AND OUTLINE

In this paper, we aim to convey two key points regarding the preparation of VBS states on a quantum computer. First, implementing generic digital quantum simulation methods or exploiting the tensor-network representation of VBS states generates quantum circuits that are too deep for near-term quantum hardware. Second, VBS states can be prepared with a shallow constant-depth circuit by considering a probabilistic implementation of the local symmetrization operator through an unconventional adaptation of the Hadamard test [51].

The VBS states are a canonical example of quantum many-body states that can be represented exactly in terms of tensor networks [52,53]. As a result, the natural approach to preparing the VBS states on quantum hardware corresponds to making use of a quantum routine that capitalizes on such exact tensor network representation. That is precisely our starting point in this paper: we prepare the one-dimensional spin-1 VBS state using the general method introduced by Schön *et al.* [54] to initialize a matrix product state (MPS) [52,55] on quantum hardware. It turns out that the resulting circuit involves sequential three-qubit operations, which may take up to $\mathcal{O}(20N)$ CNOTs of depth for N sites. Hence, even for a small chain, the MPS-based preparation of the spin-1 VBS state is challenging on NISQ hardware.

Furthermore, we have found that applying the method by Schön *et al.* [54] to the preparation of the spin-1 VBS state defined on a ring is challenging. One of the novel results presented in this paper is the adaptation of this method to the case of an MPS with periodic boundary conditions. All concepts and the outline of the technical steps are discussed in Sec. IV. To avoid an overly long digression in the main text, the full account of the technical aspects of the method, both for open and periodic boundary conditions, can be found in Appendix C.

More generally, in Sec. IV, we also argue that generic methods such as quantum phase estimation [3,56] and quantum cooling [57–59] are not scalable if the initial state is the conventional product state of valence bonds. Quantum schemes [50,60] based on the representation of two-dimensional VBS states in terms of projected entangled-pair states (PEPS) [61], typically based on the construction of an adiabatic path with gapped Hamiltonians related to the parent AKLT Hamiltonian, are also critically discussed, as far as their implementation on near-term quantum hardware is concerned.

Having attested the inadequacy of these quantum schemes for NISQ hardware, we proceed to the second key point of this paper. The main result, which is introduced in Sec. VI, is a nondeterministic method based on adapting the Hadamard test

to prepare a general spin- S VBS state. Not only is the resulting circuit depth independent of the lattice size but, through state-of-the-art basis gate decomposition techniques (Sec. VII), we achieve as few as 8 and 27 CNOT gates of depth for the spin-1 and spin-3/2 VBS states.

In the remaining sections, we address additional issues that are important to implement this probabilistic method on near-term quantum hardware: mitigation of repetition overhead due to nondeterministic nature of method (Secs. VIII and IX), adaptation of method to high local spin S to avoid basis gate decomposition of large $(2S + 1)$ -qubit operation (Sec. X), and consideration of qubit connectivity constraints (Sec. XI). A discussion of the most promising applications with a prospect of quantum advantage is also included in Sec. XI. For the sake of clarity, Sec. V provides a more detailed summary of the sections that follow it.

Since the topic of this paper lies at the intersection of quantum information and condensed matter physics, to accommodate the different backgrounds of its broad target readership, in the next section, we introduce the AKLT models and their ground states, the VBS states.

III. AKLT MODELS: VBS STATES

In one dimension, the spin-1 AKLT model [26,27] is a special case of the bilinear-biquadratic Hamiltonian (cf. Ref. [62] and references therein)

$$\mathcal{H}_{\text{BLBQ}}^{S=1}(\beta) = \sum_{n=1}^{N-1} \vec{S}_n \cdot \vec{S}_{n+1} + \beta (\vec{S}_n \cdot \vec{S}_{n+1})^2, \quad (1)$$

with $\beta = 1/3$ and $\vec{S}_n = (S_n^x, S_n^y, S_n^z)$, where S_n^i are the spin-1 matrices. The phase diagram of $\mathcal{H}_{\text{BLBQ}}^{S=1}(\beta)$ comprises a gapless phase for $\beta > 1$, a dimerized phase for $\beta < -1$ and the gapped Haldane phase for $\beta \in (-1, 1)$. The latter includes both the spin-1 Heisenberg model [28] ($\beta = 0$) and the spin-1 AKLT model [26,27] ($\beta = 1/3$).

The spin-1 AKLT model, $\mathcal{H}_{\text{AKLT}}^{S=1} \equiv \mathcal{H}_{\text{BLBQ}}^{S=1}(\beta = \frac{1}{3})$, can be expressed (up to a factor and an additive constant) in terms of a sum of the local projectors $P_{n,n+1}^{S=2}$ that map neighboring spins to the subspace of total spin $|\vec{S}_n + \vec{S}_{n+1}|^2 = 2(2+1)$ [26,27],

$$\mathcal{H}_{\text{AKLT}}^{S=1} = 2 \sum_{n=1}^{N-1} \left(P_{n,n+1}^{S=2}(\vec{S}_n, \vec{S}_{n+1}) - \frac{1}{3} \right), \quad (2)$$

where

$$P_{n,n'}^S = C \prod_{S' \neq S} ((\vec{S}_n + \vec{S}_{n'})^2 - S'(S' + 1)). \quad (3)$$

C is a constant that ensures that $P_{n,n'}^S$ acts trivially on states with $S_{\text{total}}^{n,n+1} = S$, where $\vec{S}_{\text{total}}^{n,n+1} \equiv \vec{S}_n + \vec{S}_{n+1}$. Ignoring the constants, all eigenenergies must be non-negative because each term is a projector. As a result, a state satisfying $P_{n,n+1}^{S=2} |\psi_0\rangle = 0$ for all pairs of neighboring spins must be a ground state of the AKLT model.

The VBS construction scheme ensures this ground state condition is satisfied as follows. Originally, the total spin of any pair of neighboring spins 1, $\vec{S}_{\text{total}}^{n,n+1} = \vec{S}_n + \vec{S}_{n+1}$, can take the values $1 \oplus 1 = 0, 1, 2$, where \oplus denotes angular momentum addition. To dispose of $S_{\text{total}}^{n,n+1} = 2$, each local spin 1

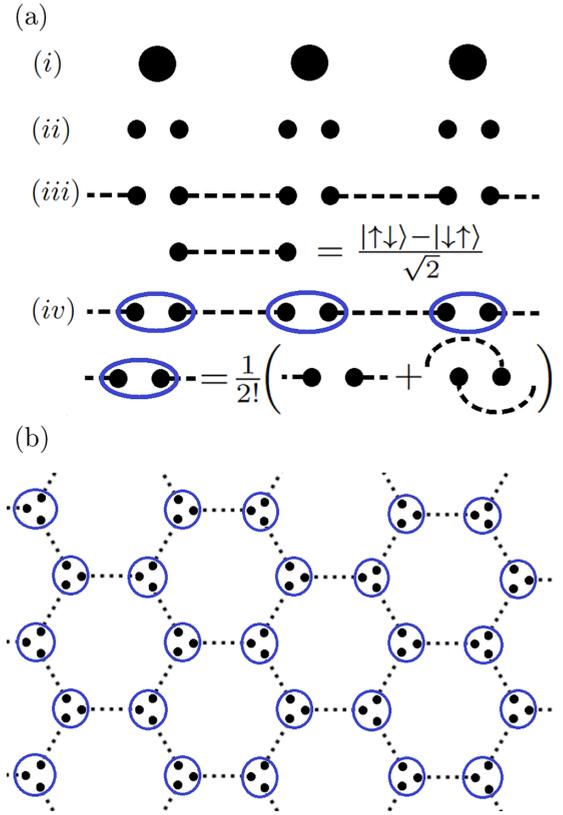


FIG. 1. (a) Illustrated description of construction of spin-1 valence-bond-solid (VBS) states in one-dimensional lattice. (i) Local degree of freedom at each lattice site is spin 1. (ii) Each spin 1 is decomposed into two spins 1/2. (iii) A valence bond $\frac{|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle}{\sqrt{2}}$ is prepared at every pair of nearest-neighbor sites, each site contributing a spin 1/2. This ensures the ground state condition for the AKLT Hamiltonian [cf. Eq. (2)] is satisfied. (iv) The local symmetrization operator $S^{(\mathcal{N}=2)}$ is applied at every pair of qubits that encode a single site to retrieve the original spin-1 local degree of freedom. (b) Extension of VBS construction to a honeycomb lattice. Since coordination number is 3, there are three spins 1/2 per lattice site, each involved in a single valence bond. Application of $S^{(\mathcal{N}=3)}$ at every site produces a spin 3/2. This spin-3/2 VBS state is the ground state of AKLT Hamiltonian stated in Eq. (7).

is decomposed into two spins 1/2, so that a valence bond $\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$ can be created between the two adjacent spins 1/2 for each pair of neighboring sites [cf. Fig. 1(a)]. Since the valence bond is a (spin-0) singlet, now $S_{\text{total}}^{n,n+1}$ can only take values $\frac{1}{2} \oplus 0 \oplus \frac{1}{2} = 0, 1$, as desired. Finally, to ensure that the physical degrees of freedom at each site are the expected spins 1, the respective pair of spins 1/2 must be symmetrized,

$$S^{(\mathcal{N}=2)}(|\sigma\rangle_L |\sigma'\rangle_R) = \frac{1}{2} (|\sigma\rangle_L |\sigma'\rangle_R + |\sigma'\rangle_L |\sigma\rangle_R), \quad (4)$$

with $\sigma, \sigma' \in \{\uparrow, \downarrow\}$. $S^{(\mathcal{N}=2)}$ refers to the symmetrization operator that acts on $\mathcal{N} = 2$ spins 1/2.

This VBS construction can be straightforwardly extended to arbitrary lattices in any dimensions [26,27,55]. Given some lattice (Ω, \mathcal{L}) , where Ω is the set of sites and \mathcal{L} the set of links, every site $n \in \Omega$ with coordination number \mathcal{N}_n is associated

with a local spin $\mathcal{N}_n/2$.¹ Each local spin is decomposed into \mathcal{N}_n spins $1/2$, one for each lattice link $(n, n') \in \mathcal{L}$ emanating from site n , with $n' \in \Omega$ a nearest-neighboring site of n . In the first step, a valence bond is created along each link, such that each spin $1/2$ is involved in one valence bond. Henceforth, this product state of valence bonds will be referred to as

$$|\psi_{\text{pre-VBS}}\rangle = \bigotimes_{(n,n') \in \mathcal{L}} \frac{|\uparrow\rangle_{(n,n')} |\downarrow\rangle_{(n',n)} - |\downarrow\rangle_{(n,n')} |\uparrow\rangle_{(n',n)}}{\sqrt{2}},$$

where $|\sigma\rangle_{(n,n')}$ denotes the state of the spin $1/2$ from site n that is associated with neighboring site n' . The second and final steps amount to encoding a spin $\mathcal{N}_n/2$ by applying the local symmetrization operator,

$$\mathcal{S}_n^{(\mathcal{N}_n)} \left(\bigotimes_{i=1}^{\mathcal{N}_n} |\sigma_i\rangle \right) = \frac{1}{\mathcal{N}_n!} \sum_{\mathcal{P}} \left(\bigotimes_{i=1}^{\mathcal{N}_n} |\mathcal{P}(\sigma_i)\rangle \right), \quad (5)$$

at every site n . \mathcal{P} is the generator of all $\mathcal{N}_n!$ permutations of the \mathcal{N}_n spins $1/2$ belonging to site n . The combination of these two steps yields the (unnormalized) valence-bond-solid (VBS) state:

$$|\psi_{\text{VBS}}\rangle = \bigotimes_{n \in \Omega} \mathcal{S}_n^{(\mathcal{N}_n)} |\psi_{\text{pre-VBS}}\rangle. \quad (6)$$

Figure 1(b) shows a schematic representation of the VBS state associated with the honeycomb lattice. Each dashed black line illustrates a valence bond, while the blue loops represent the local symmetrization. For both sublattices, every site has coordination number $\mathcal{N}_n = 3$, so the local physical degree of freedom corresponds to a spin-3/2, encoded by 3 spins $1/2$. As in one dimension, the valence bonds guarantee the total spin $\vec{S}_{\text{total}}^{n,n'} = \vec{S}_n + \vec{S}_{n'}$ resulting from the sum of the two neighboring spins $3/2$, \vec{S}_n and $\vec{S}_{n'}$, does not take its maximum value $S_{\text{total}}^{n,n'} = 3$, thus ensuring this VBS state is the ground state of the Hamiltonian [26,27]

$$\begin{aligned} \mathcal{H}_{\text{AKLT}}^{S=\frac{3}{2}} &= \sum_{(n,n') \in \mathcal{L}} P_{n,n'}^{S=3}(\vec{S}_n, \vec{S}_{n'}) \\ &= \sum_{(n,n') \in \mathcal{L}} \left(\vec{S}_n \cdot \vec{S}_{n'} + \frac{116}{243} (\vec{S}_n \cdot \vec{S}_{n'})^2 + \frac{16}{243} (\vec{S}_n \cdot \vec{S}_{n'})^3 \right). \end{aligned} \quad (7)$$

Despite the historic significance of the spin-1 VBS state in one dimension (1D), both for its support to Haldane's conjecture [28] and for its pioneering role in the development of matrix product states (MPS) [63,64], the spin-3/2 VBS state on the honeycomb lattice is a potentially more relevant trial state for digital quantum simulation. Indeed, while most properties of the 1D spin-1 AKLT model can be computed analytically or, if needed, numerically (e.g., via DMRG [52,65–67]), analytical results for the spin-3/2 AKLT model on the

honeycomb lattice are far scarcer and conventional numerical methods less scalable. In fact, in 1D, the existence of the spectral gap was proven already in the original AKLT papers [26,27] and even exact results for some excited states have been derived [68–70]. For the honeycomb lattice, in turn, the nonzero gap of the AKLT model was only recently established [40,41] and no exact excited states are known.

There are two noteworthy differences between the spin-1 and spin-3/2 VBS states. The first concerns their computational power within the context of measurement-based quantum computation [31–33]; while the spin-1 VBS state can only simulate restricted computations involving arbitrary single-qubit gates [71,72], the spin-3/2 VBS state has been shown to be a resource state for the simulation of universal quantum circuits [29,30]. The second difference lies in the relation to the corresponding Heisenberg model: In 1D the spin-1 AKLT and Heisenberg models are both in the Haldane phase [62], whereas in the honeycomb lattice the ground state of the spin-3/2 Heisenberg model is Néel-ordered [27], thus implying a phase transition separating it from the spin-3/2 AKLT model, which has a unique disordered ground state [27].

In the remainder of this paper, the methods will first be presented for the general case of a spin- S VBS state, with $S \in \{1, \frac{3}{2}, 2, \frac{5}{2}, 3, \dots\}$. In addition, the two important examples highlighted above, the spin-1 VBS state (the simplest and canonical example of a VBS state) and the spin-3/2 VBS state (the natural candidate VBS state to achieve quantum advantage), will be discussed in detail.

IV. VBS STATES ON QUANTUM HARDWARE: OUTLINE OF PROBLEM

A. Failure of standard quantum simulation

The preparation of the spin- S VBS state defined on a lattice with coordination number $\mathcal{N} = 2S$ involves the following two steps.

(1) Initializing $|\psi_{\text{pre-VBS}}\rangle$, the product state of valence bonds, one for each lattice link.

(2) Applying the local symmetrization operator $\mathcal{S}^{(\mathcal{N})}$ at each lattice site of $|\psi_{\text{pre-VBS}}\rangle$, yielding $|\psi_{\text{VBS}}\rangle$.

In the language of quantum information theory, a valence bond is nothing more than the Bell state $|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, where the mapping $\{|\uparrow\rangle \leftrightarrow |0\rangle, |\downarrow\rangle \leftrightarrow |1\rangle\}$ is implied. Hence, the first step can be straightforwardly implemented by applying the two-qubit subcircuit shown in Fig. 3(a) at each pair of qubits $|n; n'\rangle$ and $|n'; n\rangle$ representing the two spins $1/2$ along a lattice link $(n, n') \in \mathcal{L}$. All such subcircuits can be applied in parallel, thus resulting in a layer of depth 1 CNOT,² assuming merely connectivity between two such qubits.

The challenge lies in the second part of the VBS construction. This is due to the nonunitarity of the local symmetrization operator $\mathcal{S}^{(\mathcal{N})}$, which has two consequences.

¹Our notation includes the most general case of a graph with locally varying coordination number, in which case the local spin varies from site to site. However, the particular cases considered in this paper will correspond to regular lattices where every site is equivalent, and therefore all local spins are the same.

²Since the execution time and error rates of two-qubit gates are significantly greater than those of single-qubit gates (especially standard ones like Z , X or H), we ignore the latter in the calculation of the depth of this circuit.

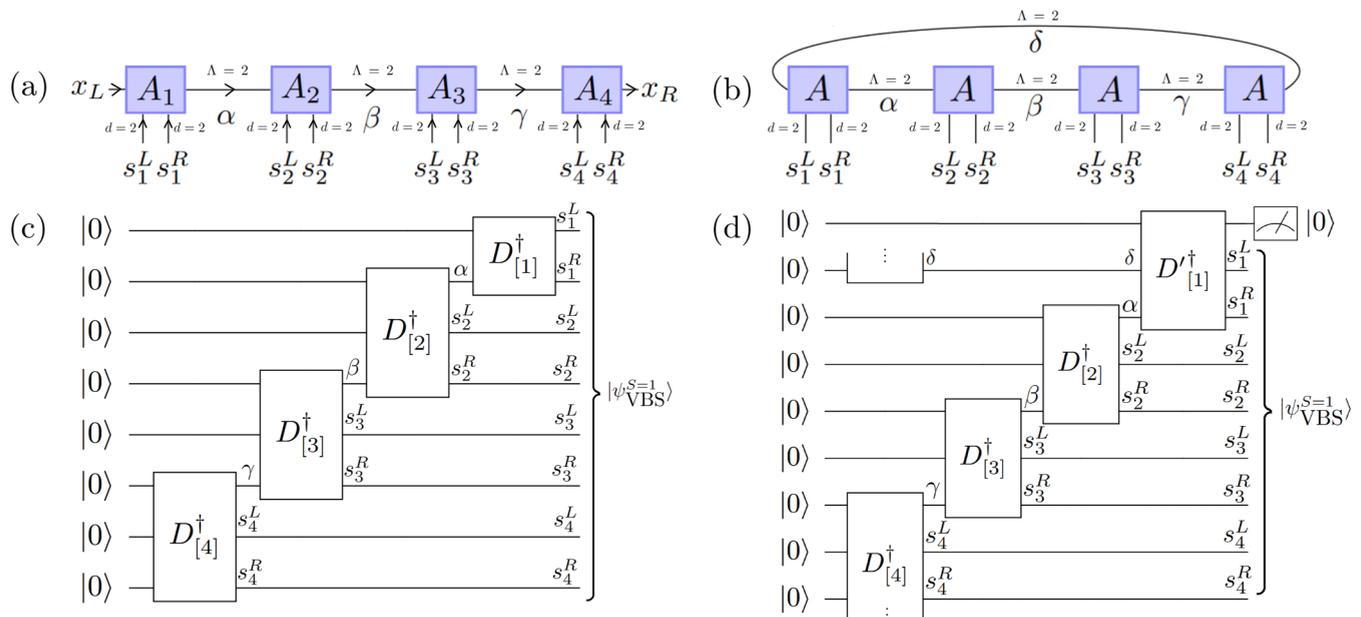


FIG. 2. [(a) and (b)] Matrix product state (MPS) representation of a one-dimensional spin-1 VBS state with open (a) and periodic (b) boundary conditions for $N = 4$ sites. Due to the translational invariance, all tensors in (b) are given by Eq. (11), where two physical indices, each of dimension $d = 2$, are used instead of the conventional single index with $d = 3$. The MPS with open boundary conditions in (a) can be obtained by opening the bond corresponding to the sum over δ in (b) and setting the two resulting free indices x_L and x_R to one of the two possible values, thus producing one of the four degenerate VBS states for a spin-1 AKLT chain. The arrows in (a) allude to the fact that the MPS is in left-canonical form, in which case the local tensors vary from site to site, hence the site subscript used. [(c) and (d)] High-level scheme of quantum circuits that initialize spin-1 VBS states with open (c) and periodic (d) boundary conditions for $N = 4$ sites by exploiting their MPS representation. $D_{[i]}^\dagger$ is the inverse of the unitary matrix product operator [73,74] $D_{[i]}$ that disentangles site i from the remainder of the MPS. The explanation of how to obtain $D_{[i]}$ from the corresponding local tensor A_i in the MPS can be found in Appendix C. For N sites and open boundary conditions, as in (c), the circuit involves the sequential application of $N - 1$ three-qubit operations $D_{[i]}^\dagger$, $i = N, N - 1, \dots, 2$, followed by the execution of the two-qubit operation $D_{[1]}^\dagger$, so the depth is $\mathcal{O}(N)$. The case of periodic boundary conditions introduces two differences. First, the final two-qubit operation, $D_{[1]}^\dagger$, is nonunitary, so it has to be embedded in a three-qubit unitary operation $D'_{[1]}^\dagger$ to be executed on quantum hardware. The additional ancilla must be initialized in $|0\rangle$ and measured in $|0\rangle$ (with probability of 50%, regardless of N) for the method to be successful. Second, the first operation, $D_{[N]}^\dagger$, acts on four qubits instead of three.

First, its implementation on quantum hardware is nontrivial, since quantum circuits are inherently unitary, as they ultimately amount to time-evolution operators of closed quantum systems (at least, within the gate-based paradigm) [9]. Second, applying $\mathcal{S}^{(N)}$ amounts to projecting $|\psi_{\text{pre-VBS}}\rangle$ onto the local subspace of exchange-symmetric states. As a result, ignoring finite-size effects, the overlap between the easy-to-prepare $|\psi_{\text{pre-VBS}}\rangle$ and the final (normalized, hence the tilde) $|\tilde{\psi}_{\text{VBS}}\rangle$ decreases exponentially with the number of lattice sites N ,

$$\langle \tilde{\psi}_{\text{VBS}} | \psi_{\text{pre-VBS}} \rangle = \left(\frac{2S+1}{2^{2S}} \right)^{N/2} \equiv p^{N/2}. \quad (8)$$

p corresponds to the fraction of symmetric spin states selected by $\mathcal{S}^{(N=2S)}$. In particular, $p = 3/4$ for the spin-1 VBS state and $p = 1/2$ for the spin-3/2 VBS state. Appendix B presents an explicit derivation of these two results and explains how the general approach is applied.

This exponentially vanishing overlap implies that attempting to prepare $|\psi_{\text{VBS}}\rangle$ starting from the easy-to-prepare $|\psi_{\text{pre-VBS}}\rangle$ via a standard digital quantum simulation strategy (e.g., quantum phase estimation [3,56], quantum cooling

methods [57–59]) becomes impractical for sufficiently large systems. Instead, Wang [59] started from a computational basis state, a good enough input state for the proof-of-concept simulation of spin-1 AKLT chain with just $N = 3$ sites but, given the exponential scaling of the size of the Hilbert space, certainly not for large enough systems to probe the thermodynamic limit.

B. Matrix product states in 1D

An alternative approach to the preparation of VBS states on quantum hardware involves exploiting their exact representation in terms of tensor networks of bond dimension $\Lambda = 2$. In particular, in one dimension, the spin-1 VBS state with periodic boundary conditions can be expressed as an MPS [52,55] [cf. Fig. 2(b)]

$$|\psi_{\text{VBS}}^{S=1}\rangle = \sum_{\vec{\sigma}} \text{Tr}(A^{\sigma_1} A^{\sigma_2} \dots A^{\sigma_N}) |\vec{\sigma}\rangle, \quad (9)$$

where $|\vec{\sigma}\rangle \equiv \bigotimes_{n=1}^N |\sigma_n\rangle$ represents the N physical spin-1 degrees of freedom [i.e., $|\sigma_i\rangle = (\sigma_i^{+1}, \sigma_i^0, \sigma_i^{-1})^T$] and the (left- and right-normalized) local rank-3 tensors A^σ are of the

form

$$A^{+1} = \begin{pmatrix} 0 & \sqrt{\frac{2}{3}} \\ 0 & 0 \end{pmatrix}, \quad A^0 = \begin{pmatrix} -\frac{1}{\sqrt{3}} & 0 \\ 0 & \frac{1}{\sqrt{3}} \end{pmatrix},$$

$$A^{-1} = \begin{pmatrix} 0 & 0 \\ -\sqrt{\frac{2}{3}} & 0 \end{pmatrix}. \quad (10)$$

On a digital quantum computer, these spins-1, $|\sigma_i\rangle$, have to be encoded (with redundancy) in terms of two qubits each, $(|s_i^L\rangle, |s_i^R\rangle)$. Hence, for our purposes, it is more useful to express the local tensor in the form

$$A^{\uparrow\uparrow} = \begin{pmatrix} 0 & \sqrt{\frac{2}{3}} \\ 0 & 0 \end{pmatrix}, \quad A^{\uparrow\downarrow} = \begin{pmatrix} -\frac{1}{\sqrt{6}} & 0 \\ 0 & \frac{1}{\sqrt{6}} \end{pmatrix},$$

$$A^{\downarrow\uparrow} = \begin{pmatrix} -\frac{1}{\sqrt{6}} & 0 \\ 0 & \frac{1}{\sqrt{6}} \end{pmatrix}, \quad A^{\downarrow\downarrow} = \begin{pmatrix} 0 & 0 \\ -\sqrt{\frac{2}{3}} & 0 \end{pmatrix}, \quad (11)$$

where the left- and right-normalization conditions are still satisfied. To obtain the MPS for the spin-1 VBS state with open boundary conditions, one can simply remove the virtual index corresponding to the bond linking the two ends of the chain (i.e., the bond between A^{σ_1} and A^{σ_N}) and set each of the resulting free indices x_L and x_R to one of two possible values [cf. Fig. 2(a)], thus yielding the expected four-fold degeneracy of the spin-1 AKLT chain. At this point, the two tensors at the ends of the chain, A^{σ_1} and A^{σ_N} , are only right- and left-normalized, respectively, so the MPS has to be brought into left-canonical form to apply the scheme discussed in the following paragraphs and explained in detail in Appendix C. In the left-canonical MPS, the tensors A^{σ_i} at different sites are no longer equal, which justifies the addition of a site subscript, as shown in Fig. 2(a).

A MPS with open boundary conditions, virtual index dimension $\Lambda = 2$ and physical index dimension $d = 2$ (e.g., for spins 1/2 or spinless fermions) can be initialized exactly on quantum hardware following the single-layer scheme discussed by Ran [75], which translated an earlier method by Schön *et al.* [54] to the language of gate-based quantum computing. As explained in Appendix C, it is straightforward to extend such method to $d = 4$ (i.e., two qubits at each site, so that a spin-1 can be encoded locally), which results in the quantum circuit schematically shown in Fig. 2(c) for the particular case of $N = 4$ sites. Apart from the last step, which amounts to a two-qubit operation $D_{[1]}^\dagger$, for a chain with N sites, this method corresponds to the sequential application of $N - 1$ three-qubit operations $D_{[i]}^\dagger$, $i = 2, 3, \dots, N$, thus allowing to prepare a 1D spin-1 VBS state with $\mathcal{O}(N)$ depth. $D_{[i]}^\dagger$ is the inverse of the unitary matrix product operator [73,74] $D_{[i]}$ that disentangles site i from the remainder of the MPS. Appendix C shows how to obtain $D_{[i]}$ from the corresponding local tensor A_i in the MPS.

The prefactor of such linear scaling can be prohibitively large for NISQ hardware, though. Using the state-of-the-art quantum Shannon decomposition [76], which takes at most 20 CNOT gates to decompose a three-qubit operation, the circuit depth in CNOT gates of this MPS-based method should scale at most as $\mathcal{O}(20N)$ for N lattice sites. For

concreteness, as discussed in Appendix C, using the Cirq `three_qubit_matrix_to_operations` method [77], the quantum Shannon decomposition was applied explicitly for the case of periodic boundary conditions described below [cf. Fig. 2(d)] to the three-qubit operation $D_{[i]}^\dagger$, $i = 2, 3, \dots, N - 1$ ³. This decomposition resulted in the maximum count of 20 CNOT gates.

It is also possible to adapt this scheme to the preparation of a 1D spin-1 VBS state with periodic boundary conditions [cf. Fig. 2(d)]. There are two main differences relative to the previous case. The first difference lies in the fact that the first operation, $D_{[N]}^\dagger$, acts on four qubits instead of three. However, since it is the first operation of the circuit, it is not necessary to find a circuit that acts on an arbitrary initial state, but only one that produces the desired effect on the fiducial state $|0\rangle^{\otimes 4}$. In other words, the basis gate decomposition of a 16×16 unitary matrix is replaced by the determination of the circuit that initializes a four-qubit state, which can be accomplished straightforwardly by exploiting its Schmidt decomposition, as discussed in Sec. IX.

The second difference appears in the final operation $D_{[1]}^\dagger$, which is nonunitary in this case. As discussed in Appendix C, this two-qubit nonunitary operation can be embedded in a three-qubit unitary one via the generic method presented in Appendix C of Ref. [78]. This involves the addition of an ancillary qubit initialized in state $|0\rangle$, which is measured in the computational basis at the end of the scheme. If the ancilla is found in state $|0\rangle$, which occurs with a probability of 50% (irrespective of N), the VBS state is prepared in the remaining $2N$ qubits.

C. Projected entangled-pair states in 2D

Beyond one dimension, VBS states can no longer be expressed as a MPS with constant bond dimension Λ , but they can nonetheless be cast exactly in the form of a PEPS [61] (a high-dimensional generalization of MPS) with bond dimension $\Lambda = 2$ as well. However, to the best of our knowledge, there is no method in the literature to initialize PEPSs on quantum hardware by deriving a quantum circuit from the respective local tensors, as described above for MPSs. In fact, the preparation of a considerable set of 2D PEPS is likely to require a circuit depth that scales exponentially with the number of lattice sites N [79,80].

Nevertheless, Schwarz *et al.* [60] developed a quantum routine to prepare injective PEPS on digital quantum computers with polynomial resources, provided that the parent Hamiltonian of which the injective PEPS is the unique ground state is gapped [81] and sequences of partial sums of local terms of the parent Hamiltonian can be defined such that each partial sum has a unique ground state of its own. The key idea [60] is to perform quantum phase estimation (QPE) [3,56] with such a partial sum of terms of the parent Hamiltonian, enlarging the number of lattice sites acted upon by one at a time. This therefore gives rise to a sequential application of at least N QPE executions—possibly more than N since the

³Thanks to the translational invariance, the unitary disentangler $D_{[i]}$ happens to be the same for $i = 2, 3, \dots, N - 1$.

measurement outcome of QPE may not be the ground state energy, in which case the measurement must be undone via the Marriott-Watrous trick [82–84], so that QPE can be repeated until achieving the successful outcome.

For the case of two-dimensional VBS states such as the spin-3/2 VBS state on the honeycomb lattice and possibly the spin-2 VBS state on the square lattice (given the tensor network renormalization group (TNRG) results suggesting the existence of a nonzero gap [85]), both these conditions are satisfied: The parent Hamiltonian (i.e., the respective AKLT model) is gapped, and the sequences of partial sums, which include only the projectors $P_{n,n'}^{S=S_{\max}}(\vec{S}_n, \vec{S}_{n'})$ acting on the sites considered up to that point, also have a unique ground state. Hence, starting from the product state of valence bonds, $|\psi_{\text{pre-VBS}}\rangle$, the VBS state can be grown one site at a time via QPE using the partial sums of terms of the AKLT Hamiltonian, thus effectively resulting in a $\mathcal{O}(N)$ circuit depth.

Alternatively, injective PEPS can be simply prepared via adiabatic evolution [86], thus forgoing the projection onto the ground-state manifold via QPE that is adopted in the method devised by Schwarz *et al.* [60]. Importantly, Ge *et al.* [50] have shown it is possible to prepare injective PEPS in $\mathcal{O}(\text{polylog}(N))$ time provided that the spectral gap along the whole adiabatic path is above a certain threshold, which can be computed in advance [87]. Further numerical studies are required to assess the feasibility of the adiabatic state preparation of the spin-3/2 VBS state for the intermediate system sizes attainable in NISQ devices.

D. Implementation in NISQ hardware

In summary, the use of standard digital quantum simulation algorithms such as QPE [3,56] or quantum cooling methods [57–59] to prepare VBS states on quantum hardware using their parent AKLT Hamiltonians is hampered by the exponentially vanishing overlap with the easy-to-prepare initial state $|\psi_{\text{pre-VBS}}\rangle$. As an alternative, VBS states can be constructed one site at a time by exploiting their MPS representation in one dimension or the fact that the AKLT Hamiltonian is a sum of local projectors in any number of dimensions. In any case, despite the $\mathcal{O}(N)$ depth of these sitewise approaches, even for a very small number of lattice sites N , the circuit depth is bound to be prohibitively large for NISQ hardware as the MPS construction [54,75] involves $N - 1$ sequential three-qubit operations and the injective PEPS preparation method by Schwarz *et al.* [60] requires at least N sequential QPE executions. Adiabatic state preparation could be a better alternative in light of the $\mathcal{O}(\text{polylog}(N))$ asymptotic scaling proven by Ge *et al.* [50], but the implementation of the time-evolution operators in NISQ hardware may be daunting, as the two-body operators acting on neighboring spin-3/2 that appear in the AKLT Hamiltonian become six-qubit operators upon encoding each spin-3/2 in terms of 3 qubits. In the next sections, we take the opposite approach by devising a method that implements the symmetrization at all sites in parallel, resulting in a constant-depth circuit at the cost of requiring the post-selection of the measurement outcomes of ancillary qubits. In a sense, circuit depth is traded for the

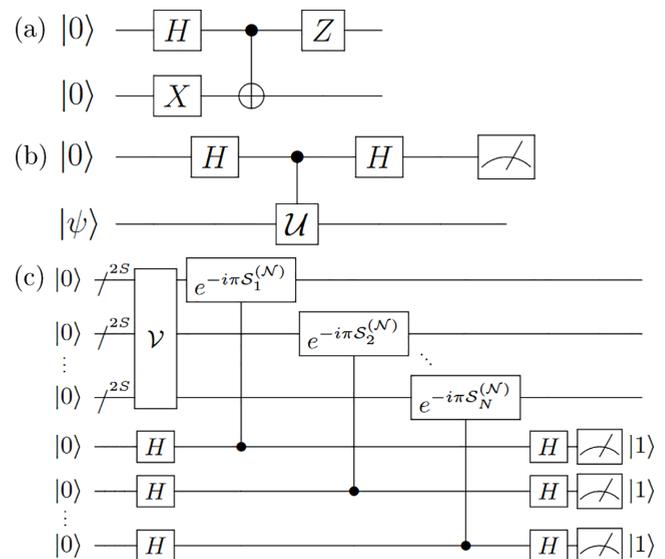


FIG. 3. (a) Two-qubit subcircuit that prepares a single valence bond $\frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) \equiv \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$. Full circuit \mathcal{V} that prepares $|\psi_{\text{pre-VBS}}\rangle = \mathcal{V}|0\rangle^{\otimes 2NS}$ amounts to repetition of this subcircuit across all pairs of qubits representing lattice links, with N the number of lattice sites and S the local spin. (b) Hadamard test [90,91] for generic unitary operator \mathcal{U} . Setting $\mathcal{U} = e^{-i\pi S^{(N)}}$ turns input $|0\rangle \otimes |\psi\rangle$ into $|0\rangle \otimes (\mathbb{1} - S^{(N)})|\psi\rangle + |1\rangle \otimes S^{(N)}|\psi\rangle$, so measuring ancilla in $|1\rangle$ ensures symmetrization operator $\mathcal{S}^{(N)}$ acting on $\mathcal{N} = 2S$ spins 1/2 (i.e., qubits) is applied at given site in main register. (c) Circuit that prepares spin- S VBS state, $|\psi_{\text{VBS}}\rangle = \bigotimes_{n=1}^N \mathcal{S}_n^{(N)} |\psi_{\text{pre-VBS}}\rangle$, in N -site lattice if all N ancillas are measured in $|1\rangle$. The basis gate decomposition of controlled $e^{-i\pi S}$ for $S = 1$ and $3/2$ can be found in Sec. VII. All N controlled- $e^{-i\pi S_n^{(N)}}$ are executed in parallel, so the circuit depth is constant, i.e., independent of the system size N .

repetition of the same shallow circuit, in the spirit of hybrid variational algorithms [13]. Thus the resulting method, though probabilistic, is suitable for NISQ devices.

V. PROBABILISTIC PREPARATION OF VBS STATES: SUMMARY OF METHODS

In light of the unsuitability of generic digital quantum simulation methods and quantum routines that exploit tensor-network representations to prepare VBS states on noisy intermediate-scale quantum processors, in the remainder of this paper we will introduce a probabilistic method that produces circuits with depth independent of the size of the lattice in which the VBS state is defined. For clarity, this section summarizes the results presented in the next sections.

The default version of this probabilistic scheme to prepare VBS states is introduced in the next section. It comprises the following two steps, for which the respective depth in CNOT gates for the spin-1 and spin-3/2 VBS states is given below.

(i) Prepare a valence bond at each pair of qubits encoding a lattice link through the two-qubit quantum circuit shown in Fig. 3(a).

Circuit depth: $S = 1$: 1 CNOT; $S = \frac{3}{2}$: 1 CNOT.

(ii) Apply the symmetrization operator $\mathcal{S}^{(N)}$ via the adapted Hadamard test shown in Fig. 3(c) at every set of $\mathcal{N} =$

$2S$ qubits encoding one lattice site. The successful application of $\mathcal{S}^{(\mathcal{N})}$ at all sites requires measuring all N ancillary qubits in state $|1\rangle$.

Circuit depth: $S = 1$: 7 CNOTs; and $S = \frac{3}{2}$: 26 CNOTs.

The circuit depth in CNOT gates for the second step, particularly for the spin-3/2 case, results from the optimization of the basis gate decomposition of the $(2S + 1)$ -qubit operation controlled- $e^{-i\pi\mathcal{S}^{(\mathcal{N})}}$, which is discussed in Sec. VII. In total, the circuit depth required to prepare the spin-1 and spin-3/2 VBS states is eight CNOT gates and 27 CNOT gates, respectively, ignoring qubit connectivity constraints.

Because the successful initialization of the VBS state demands measuring all ancillas in state $|1\rangle$, the average number of repetitions required to achieve success grows exponentially with the number of lattice sites N . To address this issue, particularly for the intermediate values of N for which quantum advantage is attainable with NISQ hardware, we have devised two strategies to quadratically reduce this repetition overhead for bipartite lattices (cf. Secs. VIII and IX). The repetition-overhead-mitigated preparation of the spin-1 and spin-3/2 VBS states takes the following two steps.

(i) For every site of one sublattice, using the general quantum state preparation method from Fig. 5, initialize the $4S$ -qubit states corresponding to the preparation of the $2S$ valence bonds emanating from the given site followed by the application of the symmetrization operator $\mathcal{S}^{(\mathcal{N}=2S)}$ at such site.

Circuit depth: $S = 1$: 4 CNOTs; $S = \frac{3}{2}$: 19 CNOTs.

(ii) For every site of the other sublattice, apply the symmetrization operator $\mathcal{S}^{(\mathcal{N})}$ via the adapted Hadamard test shown in Fig. 3(c) at every set of $\mathcal{N} = 2S$ qubits encoding one such site.

Circuit depth: $S = 1$: 7 CNOTs; and $S = \frac{3}{2}$: 26 CNOTs.

Figure 4 helps to visualize the sets of $4S$ qubits on which the first step is applied in parallel. Since half of the sites are deterministically symmetrized in the first step, the probabilistic approach in the second step takes quadratically fewer measurements on average to complete the preparation of the VBS state.

The probabilistic method herein introduced is valid for any VBS state, regardless of the local spin- S , which is determined by the coordination number of the underlying lattice. However, the basis gate decomposition of the $(2S + 1)$ -qubit operation controlled- $e^{-i\pi\mathcal{S}^{(\mathcal{N})}}$ may lead to a very deep circuit for a sufficiently large S . To address this issue, in Sec. X, we discuss how to implement the local symmetrization operator through the linear combination of unitaries (LCU) method [88], thus skipping the basis gate decomposition of a large unitary. The first repetition overhead mitigation strategy (Sec. VIII) is also more suitable for high-spin VBS states, as it avoids preparing a $4S$ -qubit state in the first step.

VI. LOCAL SYMMETRIZATION VIA HADAMARD TEST

In this section, we will adapt the Hadamard test [cf. Fig. 3(b)], a canonical quantum routine to compute expectation values of unitary operators [51], to apply the symmetrization operator in the main register upon measuring the ancillary qubit in the state $|1\rangle$. This strategy is reminiscent

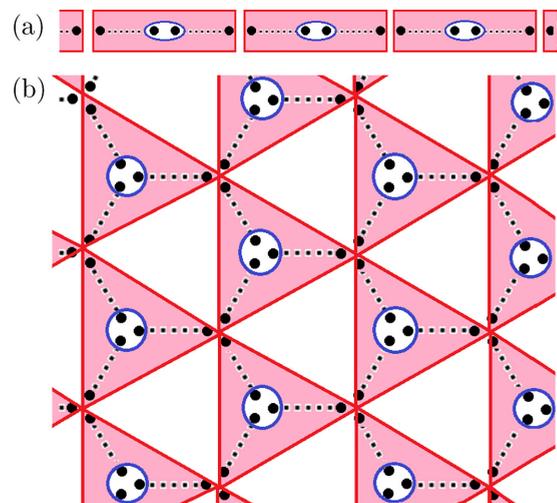


FIG. 4. Scheme of product state of valence bonds, $|\psi_{\text{pre-VBS}}\rangle$, in one-dimensional lattice (a) and honeycomb lattice (b) to support explanation of repetition overhead mitigation scheme. Dashed black lines represent valence bonds and black dots denote qubits. Four-qubit islands within filled rectangles (a) and six-qubit islands within filled triangles (b) are not entangled with one another. Hence, application of Hadamard test at sites marked with blue circles, corresponding to a single sublattice, can be repeated (with reset of qubits between consecutive trials) as many times as required to achieve success.

of the approach followed by Lacroix [89] to restore symmetries via quantum phase estimation [3,56].

The fact that the local symmetrization operator $\mathcal{S}^{(\mathcal{N})}$ is nonunitary means it cannot be expressed as a product of unitary operations (i.e., quantum gates). In any case, $\mathcal{S}^{(\mathcal{N})}$ is Hermitian, so $e^{-i\theta\mathcal{S}^{(\mathcal{N})}}$ is unitary for any $\theta \in \mathbb{R}$ and we can, in principle, find its quantum circuit.

Let us now consider the application of the Hadamard test [cf. Fig. 3(b)] with $\mathcal{U} = e^{-i\theta\mathcal{S}^{(\mathcal{N})}}$. Labelling the input state on the main register as $|\psi\rangle$, the output of this circuit before the measurement of the ancilla is

$$|0\rangle \otimes \left[\frac{\mathbb{1} + e^{-i\theta\mathcal{S}^{(\mathcal{N})}}}{2} \right] |\psi\rangle + |1\rangle \otimes \left[\frac{\mathbb{1} - e^{-i\theta\mathcal{S}^{(\mathcal{N})}}}{2} \right] |\psi\rangle. \quad (12)$$

Since $\mathcal{S}^{(\mathcal{N})}$ is idempotent ($(\mathcal{S}^{(\mathcal{N})})^2 = \mathcal{S}^{(\mathcal{N})}$), we have $e^{-i\theta\mathcal{S}^{(\mathcal{N})}} = \mathbb{1} - (1 - e^{-i\theta})\mathcal{S}^{(\mathcal{N})}$. Replacing in (12) gives

$$|0\rangle \otimes \left[\mathbb{1} - \frac{1 - e^{-i\theta}}{2} \mathcal{S} \right] |\psi\rangle + |1\rangle \otimes \left[\frac{1 - e^{-i\theta}}{2} \mathcal{S} \right] |\psi\rangle. \quad (13)$$

Setting $\theta = \pi$, we obtain

$$|0\rangle \otimes (\mathbb{1} - \mathcal{S}^{(\mathcal{N})}) |\psi\rangle + |1\rangle \otimes \mathcal{S}^{(\mathcal{N})} |\psi\rangle. \quad (14)$$

Hence, applying the Hadamard test with $\mathcal{U} = e^{-i\pi\mathcal{S}^{(\mathcal{N})}}$ to an input state $|\psi\rangle$ and retaining only the measurements of the ancilla qubit that yield $|1\rangle$ results in the application of $\mathcal{S}^{(\mathcal{N})}$ to $|\psi\rangle$ at the site in question.

This scheme can be generalized to the symmetrization of the input state $|\psi\rangle$ at all sites $n = 1, 2, \dots, N$ by considering

N Hadamard tests in parallel, with one ancilla qubit per site. The outcome is of the form

$$|\Phi\rangle \equiv \sum_{m=0}^{2^N-1} |m\rangle \otimes \left[\bigotimes_{n=1}^N (\mathbb{1} - \mathcal{S}_n^{(\mathcal{N})})^{1-m_n} (\mathcal{S}_n^{(\mathcal{N})})^{m_n} |\psi\rangle \right], \quad (15)$$

where $|m\rangle$ is a computational basis state and $m = m_N m_{N-1} \dots m_2 m_1$ is the binary representation of m . In words, whenever the n th ancilla qubit is measured in $|1\rangle$ (i.e., $m_n = 1$), the n th site of the input state $|\psi\rangle$ is locally symmetrized. As a result, if the N -qubit ancilla register is read out in state $|11\dots 11\rangle \equiv |2^N - 1\rangle$, the main register is found in $\mathcal{S}_1^{(\mathcal{N})} \mathcal{S}_2^{(\mathcal{N})} \dots \mathcal{S}_N^{(\mathcal{N})} |\psi\rangle \equiv \bigotimes_{n=1}^N \mathcal{S}_n^{(\mathcal{N})} |\psi\rangle$.

Figure 3(c) shows the quantum circuit that prepares the spin- S VBS state. It consists of three parts. First, the preparation of the product state of valence bonds, $|\psi_{\text{pre-VBS}}\rangle = \mathcal{V}|0\rangle$, then the application of the layer of local Hadamard tests, and finally the measurement of the site ancillas in the computational basis, retaining only the trials that yield all site ancillas in state $|1\rangle$. Each local spin- S is encoded by $2S$ qubits (e.g., three qubits for a spin-3/2), so the main register has a total of $2NS$ qubits for a N -site lattice. For each lattice site, an additional ancilla is required to perform the postprocessing measurement.

The fact that $|\psi_{\text{VBS}}\rangle$ is only prepared if all N ancillas are measured in $|1\rangle$ after the local Hadamard tests means only a few trials are successful. The average number of trials required to prepare $|\psi_{\text{VBS}}\rangle$ corresponds to the inverse of the probability of measuring all N site ancillas in state $|1\rangle$. Using Eq. (15) and the orthogonality of the computational basis states, the probability of measuring all N site ancillas in state $|1\rangle$ is

$$\begin{aligned} P_{11\dots 1} &= \text{Tr}(|11\dots 1\rangle\langle 11\dots 1| \otimes \mathbb{1}_{2^{2NS} \times 2^{2NS}} |\Phi\rangle\langle \Phi|) \\ &= |\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle|^2 = p^N, \end{aligned} \quad (16)$$

where in the last step the result derived in Appendix B was used. The average number of repetitions is thus exponential in the system size N , which may become prohibitively large for the intermediate values of N at which quantum advantage may be attainable.

VII. BASIS GATE DECOMPOSITION OF $e^{-i\pi \mathcal{S}^{(\mathcal{N})}}$ FOR LOCAL SPINS $S = 1$ and $3/2$

We now turn to the problem of compiling the derived algorithm into explicit quantum circuits for the important cases of $S = 1$ and $S = 3/2$ VBS states in order to better understand the practical cost of implementing the proposed methods. Concretely, we decompose the abstract controlled- $e^{-i\pi \mathcal{S}^{(\mathcal{N})}}$ gates, for local spins $S = 1$ and $3/2$, in terms of an elementary gate set $\{\text{CNOT}, U(\theta, \phi, \lambda)\}$, where $U(\theta, \phi, \lambda)$ is the general single-qubit operation

$$U(\theta, \phi, \lambda) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i(\phi+\lambda)} \cos(\theta/2) \end{pmatrix}. \quad (17)$$

To make the circuits as efficient as possible, we need to find their shallowest decompositions, as these enable faster

execution times and, more importantly, less error accumulation in near-term noisy platforms. Decompositions with fewer CNOT gates are prioritized because two-qubit operations are currently one to two orders of magnitude more error-prone than single-qubit ones [92,93].

We start with the decomposition of the controlled- $e^{-i\pi \mathcal{S}^{(\mathcal{N}=2)}}$ for local spin $S \equiv \mathcal{N}/2 = 1$. Taking the exponential of the matrix representation of $\mathcal{S}^{(\mathcal{N}=2)}$ stated in Eq. (A3) in Appendix A, one finds that

$$e^{-i\pi \mathcal{S}^{(\mathcal{N}=2)}} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} = e^{i\pi \text{SWAP}}. \quad (18)$$

Hence, the controlled- $e^{-i\pi \mathcal{S}^{(\mathcal{N}=2)}}$ is just a Fredkin gate (i.e., a controlled-SWAP, or CSWAP, for short) preceded by a Z gate acting on the control-qubit to account for the global phase factor $e^{i\pi}$ in Eq. (18) above, which becomes a relative phase factor upon controlling it. In practice, in the local Hadamard test, this Z gate can be skipped altogether; the only difference is that the desired outcome in the measurement of the site ancilla is $|0\rangle$ instead of $|1\rangle$.

Optimized basis gate decompositions of the Fredkin gate with the above gate set are given in Ref. [94] for three different coupling and qubit assignment setups. Building upon the result assuming all-to-all qubit connectivity, with seven CNOTs and a depth of ten operations, the circuit for the local Hadamard test can be straightforwardly compiled with the same depth and CNOT count, since the Hadamard and Z gates can be absorbed into the outer single-qubit operations of this Fredkin gate decomposition. Some processors, however, are restricted to nearest-neighbor couplings only. In that case, as will be described in Sec. XI, the natural placement of the ancilla qubit is between the two qubits that encode the spin-1 state at that site. The Hadamard test circuit can then be implemented with only nine CNOT gates and a depth of 19 operations by using the decomposition for the Fredkin gate in Ref. [94] where linear connectivity is assumed and the control qubit is placed at the center, dropping the unnecessary last three CNOT gates that perform a SWAP of the control back to the central position.

Moving on to the case of the controlled $e^{-i\pi \mathcal{S}^{(\mathcal{N}=3)}}$ for the scenario of local spin $S \equiv \mathcal{N}/2 = 3/2$ found in the honeycomb lattice, we first exponentiate the matrix representation of $\mathcal{S}^{(\mathcal{N}=3)}$, stated in Eq. (A5), to obtain

$$e^{-i\pi \mathcal{S}^{(\mathcal{N}=3)}} = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & -\frac{2}{3} & 0 & -\frac{2}{3} & 0 & 0 & 0 \\ 0 & -\frac{2}{3} & \frac{1}{3} & 0 & -\frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{3} & 0 & -\frac{2}{3} & -\frac{2}{3} & 0 \\ 0 & -\frac{2}{3} & -\frac{2}{3} & 0 & \frac{1}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{2}{3} & 0 & \frac{1}{3} & -\frac{2}{3} & 0 \\ 0 & 0 & 0 & -\frac{2}{3} & 0 & -\frac{2}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}. \quad (19)$$

Decomposing the controlled version of Eq. (19) using the state-of-the-art optimized quantum Shannon decomposition [76,95] takes 63 CNOT gates [96]. Starting from this method and then proceeding with a heuristic-based optimization approach, we have found a circuit that implements the controlled- $e^{-i\pi S^{(\mathcal{N}=3)}}$ with just 26 CNOT gates.

First, the 63-CNOT circuit yielded by the optimized quantum Shannon decomposition, assuming full connectivity, was recompiled with QISKIT's `transpile` function [97], which can perform some basic simplifications. The CNOT count was reduced to 54. Then, subsequent simplification steps were carried out in the ZX-calculus language; after some effort, the number of CNOT gates was reduced by more than half, as described below.

To accomplish that, the 54-CNOT circuit was converted into a ZX diagram with the aid of the PYZX software library [98], which contains a number of methods to simplify ZX diagrams and convert them back into quantum circuits [99]. We first tested the full bundle of ZX-diagram simplification rules available, including spider-fusion, identity removal, pivoting, local complementation and removal of interior nodes, followed by the gadgetization technique and its simplification. All these procedures are conveniently combined into the `full_reduce` function. The next step was extracting a new quantum circuit from the reduced ZX graph [100]. PYZX offers a few further methods to optimize the obtained circuit, which we employed before passing the result back through the QISKIT transpiler. However, in the end, the resulting quantum circuit was not simplified; rather, the CNOT count increased to 64.

Very different circuits are often extracted from equivalent ZX graphs. Hence, one way to try to avoid generating a deeper circuit than the initial 54-qubit one is to search for extracted circuits that minimize the gate count in the local space of ZX diagrams that are equivalent to the previously simplified version. PYZX implements two optimization techniques to perform this local search: simulated annealing and genetic algorithms. The next leg of our circuit simplification pipeline built upon these methods with the intensive search and optimization procedure described in Ref. [94], which often succeeds in escaping from local minima, thus optimizing decompositions further.

The final four-qubit circuit we have obtained for the controlled- $e^{-i\pi S^{(\mathcal{N}=3)}}$ operation (up to a global phase factor) in terms of the aforementioned gateset and assuming full qubit connectivity comprises 26 CNOT gates and 42 single-qubit gates, with depth 26 CNOT gates (or depth 45, if single-qubit layers are included). A similar optimization procedure was then repeated for the case of linear qubit connectivity, yielding a circuit with 39 CNOT gates and 53 single-qubit gates, with depth 39 CNOT gates (or depth 62, including single-qubit layers). These circuits can be extracted from the QASM files 5 and 6 introduced in Appendix H.

Finally, while the exact complexity of circuit extraction from a ZX diagram is not yet known, a recent preprint shows it is at least #P-hard [101]. Nevertheless, we were still able to dramatically reduce the number of operations necessary to build these circuits using a heuristic search.

VIII. MITIGATING REPETITION OVERHEAD VIA LOCAL HADAMARD TEST

The repetition overhead faced in the initialization of VBS states is common to other quantum many-body states, the preparation of which involves some nonunitary operation. This includes the Gutzwiller wave function [102,103], the ground state of the Kitaev honeycomb model [22], or the Bethe ansatz for the XXZ model [20,21]. In the latter case, for example, the number of repetitions was reduced via quantum amplitude amplification [104] (a generalization of Grover's algorithm [105]), which we may also apply to the preparation of VBS states. In fact, the layer of local Hadamard tests introduced in the previous section allows for the construction of an appropriate oracle that can be implemented with, at most, $\mathcal{O}(N)$ circuit depth (cf. Appendix D). However, the resulting quantum amplitude amplification scheme leads to a circuit with exponential depth in N , since the angle swept in the relevant two-dimensional subspace at each iteration is set by the exponentially vanishing overlap between the initial state $|\psi_{\text{pre-VBS}}\rangle$ and the final state $|\psi_{\text{VBS}}\rangle$. This would defeat our purpose of developing a low-depth preparation scheme for NISQ processors.

Returning to the probabilistic method introduced in the previous section, it is reasonable to expect to explore the locality of the Hadamard tests to reduce the repetition overhead. Indeed, measuring a single ancilla in state $|1\rangle$ is straightforward: it only takes $1/p$ repetitions on average. For example, for the spin-1 VBS state, 3 out of every 4 repetitions yield a site ancilla in state $|1\rangle$. Similarly, for the spin-3/2 VBS state, 1 out of 2 trials produces the desired outcome. The difficulty in preparing $|\psi_{\text{VBS}}\rangle$ lies in having to measure all N site ancillas in state $|1\rangle$ simultaneously. Considering a coin toss analogy, having a large number N of tossed coins land with “heads” facing up is an exponentially suppressed event, but obtaining “heads” from tossing a single coin is not.

If the input state happened to be a tensor product of sitewise states, one could repeat the local Hadamard test at each site as many times as required to obtain the successful outcome, resetting the respective qubits after an unsuccessful trial and before attempting a further one. Returning to the coin toss analogy, instead of having to obtain N “heads” at once, one would toss each coin separately until yielding “heads”, thus repeating only the toss of those coins that kept on yielding “tails.” In principle, different sites would attain success after a different number of attempts, so the first ones to succeed would have to be held until all remaining sites were symmetrized. The crucial point is that, because of the absence of entanglement between different sites, resetting the qubits at one site would not affect the wave function at any other site.

However, the initial state $|\psi_{\text{pre-VBS}}\rangle$ involves valence bonds between neighboring sites, so resetting the qubits of one site after an unsuccessful local Hadamard test affects the remainder of the sites. In fact, it does not simply affect the nearest-neighboring sites. The effective action of the Hadamard test on the state $|\psi\rangle$ in the main $2S$ -qubit register (i.e., $S^{(\mathcal{N}=2S)}|\psi\rangle$ if the ancilla is measured in $|1\rangle$, and $(1 - S^{(\mathcal{N}=2S)})|\psi\rangle$ otherwise) produces entanglement between the $2S$ qubits at that site. These qubits were already entangled

to those from neighboring sites, and within each such neighboring site all $2S$ qubits become entangled to one another due to the respective Hadamard test. Hence, the entanglement spreads across the whole system, as expected from the hidden order captured by the string order parameter in VBS states [106].

If, however, one attempts to symmetrize the wave function *only at the $N/2$ sites of one sublattice*—assuming, of course, the lattice is bipartite—the entanglement resulting from the action, whether successful or unsuccessful, of each local Hadamard test is no longer spread across the whole system, being confined to a $4S$ -qubit island (cf. Fig. 4) that only encompasses qubits from the respective site and its nearest neighbors (at which the local Hadamard test is *not* applied). The trial-and-error approach outlined above is now viable because next-nearest-neighboring sites are not entangled to one another in the input state, so all qubits within each island can be reset without affecting the remaining qubits.

In brief, the repetition overhead mitigation strategy consists of applying the Hadamard test with $\mathcal{U} = e^{-i\pi \mathcal{S}^{(N)}}$, as in the bare probabilistic method, but only to the $N/2$ sites of one sublattice, provided that the lattice in question is bipartite. Some of these local Hadamard tests will result in a successful symmetrization of the corresponding site (those for which the ancilla is measured in $|1\rangle$) while others will fail (those for which the ancilla is measured in $|0\rangle$). In the latter, one simply resets the qubits, reprepares the valence bonds that emanate from the respective site, and tries the Hadamard test again, repeating this as many times as required to obtain a successful outcome. For clarity, Fig. 4 details the independent $4S$ -qubit islands for the $S = 1, 3/2$ VBS states.

The cumulative probability of symmetrizing the $N/2$ sublattice sites after r local Hadamard test rounds is

$$P_r = (pR_r)^{N/2}, \quad (20)$$

where R_r follows the recursive relation:

$$R_r = 1 + (1 - p)R_{r-1}, \quad R_1 = 1. \quad (21)$$

As expected, for the first round, $P_1 = p^{N/2}$, which corresponds to the probability of measuring all $N/2$ ancillas in state $|1\rangle$ at once, as in the probabilistic method. As $r \rightarrow \infty$, R_r converges to $1/p$, so that the cumulative probability P_r approaches 1 in that limit.

The average number of rounds $\langle r \rangle$ required to symmetrize all $N/2$ sites of one sublattice can be computed numerically, yielding a logarithmic scaling with respect to the system size N . For the spin-1 VBS state ($p = 3/4$),

$$\langle r \rangle = 0.71 \ln N + 0.49 \quad (2 \text{ s.f.}), \quad (22)$$

and, for the spin-3/2 VBS state ($p = 1/2$),

$$\langle r \rangle = 1.4 \ln N + 0.49 \quad (2 \text{ s.f.}). \quad (23)$$

The logarithmic scaling with the system size N could be anticipated from the fact that, on average, at each round, $\mathcal{S}^{(N)}$ is successfully applied at a fraction p of the hitherto unsymmetrized sites. Since the circuit depth of a single layer of Hadamard tests [cf. Fig. 3(c)] is independent of N , this strategy to reduce the repetition overhead translates into an additional circuit layer of $\mathcal{O}(\ln N)$ depth.

Once this step is complete, one is left with the symmetrization of only half of the sites (those belonging to the other sublattice), which requires, on average, only the square root of the original repetitions using the probabilistic method introduced in the previous section. For example, for the spin-3/2 VBS state on a honeycomb lattice with $N = 20$ sites, the unmitigated probabilistic method requires an average of $2^{20} = 1\,048\,576$ repetitions, while the repetition overhead mitigation strategy reduces this overhead to $2^{10} = 1024$ trials.

Two final notes about the integration of this repetition overhead mitigation scheme in the probabilistic method are in order. First, the repetition overhead mitigation strategy must be applied *before* the probabilistic method, because the former exploits the entanglement structure of the initial state $|\psi_{\text{pre-VBS}}\rangle$, while the latter is agnostic to it, in that it applies the local symmetrization operator $\mathcal{S}^{(N)}$ at every site where the respective ancilla is measured in $|1\rangle$ regardless of the input state. Second, the $N/2$ ancillas used in the repetition overhead mitigation scheme can be reused in the probabilistic method, thus reducing the number of ancillas by half.

IX. DETERMINISTIC MITIGATION OF REPETITION OVERHEAD

Despite the low circuit depth added by the repetition overhead mitigation layer, its implementation in the currently available quantum computers may be challenging due to the need to process the outcome of mid-circuit measurements in real time. Indeed, although the mid-circuit measurement and reset of qubits are already available in some state-of-the-art platforms [107,108], the delay resulting from performing the measurements in the quantum processor, deciding the next step on the conventional processor, and finally conveying this decision to the quantum processor may be great enough to the point of extending the execution time of the circuit beyond the limits set by the decoherence of the qubits.

However, if the $4S$ -qubit islands are sufficiently small, for each island one can simply prepare the respective state deterministically using generic quantum state preparation methods. Specifically, to prepare the four-qubit and six-qubit (normalized) states corresponding to each island in the spin-1 and spin-3/2 VBS states, we will make use of a scheme [90,109] based on the Schmidt decomposition [3] of a general $2n$ -qubit state.

Let us consider the preparation of a $2n$ -qubit state

$$|\phi\rangle = \sum_{\substack{i_1, \dots, i_n=0 \\ j_1, \dots, j_n=0}}^1 M_{j_1, \dots, j_n}^{i_1, \dots, i_n} |i_1, \dots, i_n; j_1, \dots, j_n\rangle, \quad (24)$$

where the separation of indices $\{i_k\}_{k=1}^n$ and $\{j_k\}_{k=1}^n$ reflects the symmetric bipartition we will adopt to perform the Schmidt decomposition of $|\phi\rangle$. Merging the indices $\{i_k\}_{k=1}^n$ and $\{j_k\}_{k=1}^n$ into a single index \mathbf{i} and \mathbf{j} , respectively, we have

$$|\phi\rangle = \sum_{\mathbf{i}=0}^{2^n-1} \sum_{\mathbf{j}=0}^{2^n-1} M_{\mathbf{i}\mathbf{j}} |\mathbf{i}\rangle \otimes |\mathbf{j}\rangle. \quad (25)$$

The amplitudes of $|\phi\rangle$ are now cast in the form of a $2^n \times 2^n$ matrix \mathbf{M} . The singular value decomposition of \mathbf{M} gives $\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^\dagger$, in which case we can express $|\phi\rangle$ in terms of its

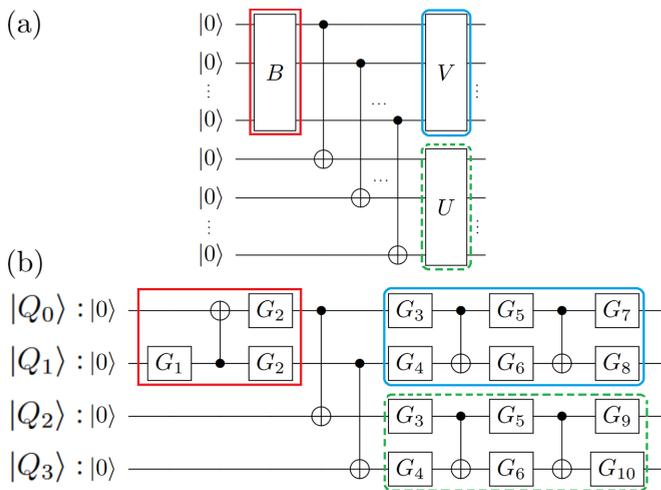


FIG. 5. (a) High-level scheme of general method [90,109] to prepare $2n$ -qubit state $|\phi\rangle = \sum_{i=0}^{2^n-1} \sum_{j=0}^{2^n-1} M_{ij} |\mathbf{i}; \mathbf{j}\rangle$ by exploiting its Schmidt decomposition. U and V are found through the singular value decomposition of \mathbf{M} , i.e., $\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^\dagger$. B prepares state $\sum_{k=0}^{2^n-1} s_k |k\rangle$, where $\mathbf{S} = \text{diag}(s_0, \dots, s_{2^n-1})$. (b) Application of this general method to preparation of four-qubit islands in spin-1 VBS state. Valence bonds are associated with qubit pairs (Q_0, Q_1) and (Q_2, Q_3) . Symmetrization operator \mathcal{S} is applied at qubits (Q_1, Q_2) , which encode a single site. A total of seven CNOT gates and a depth of four CNOT gates are required to prepare such state. Single-qubit gates G_i are presented in compact form for the sake of clarity; details can be found in the QASM file 3 introduced in Appendix H.

Schmidt decomposition:

$$|\phi\rangle = \sum_{k=0}^{2^n-1} s_k (\mathbf{U}|k\rangle) \otimes (\mathbf{V}|k\rangle) \equiv \sum_{k=0}^{2^n-1} s_k |u_k\rangle \otimes |v_k\rangle. \quad (26)$$

$\{s_k\}_{k=0}^{2^n-1}$ are the singular values of \mathbf{M} , which are the real, non-negative entries of the diagonal matrix \mathbf{S} . $\{u_k\}_{k=0}^{2^n-1}$ and $\{v_k\}_{k=0}^{2^n-1}$ are the left and right singular vectors of \mathbf{M} , corresponding to the columns of the unitary matrices \mathbf{U} and \mathbf{V} . Importantly, since \mathbf{M} is itself unitary, $\sum_{k=0}^{2^n-1} s_k^2 = 1$, so $(s_0, s_1, \dots, s_{2^n-1})^T = \sum_{k=0}^{2^n-1} s_k |k\rangle$ defines a properly normalized n -qubit state.

Figure 5(a) shows a high-level scheme of the circuit that exploits the Schmidt decomposition to prepare $|\phi\rangle$ on a $2n$ -qubit register, which is split into two n -qubit subregisters. It comprises three parts.

(1) Initialize the state $\sum_{k=0}^{2^n-1} s_k |k\rangle$ on one of the two subregisters via the n -qubit subcircuit B .

(2) Apply a network of n CNOT gates, the i th CNOT having as control the i th most significant qubit of the subregister previously acted on by B and as target the i th most significant qubit of the other subregister. This prepares $\sum_{k=0}^{2^n-1} s_k |k\rangle \otimes |k\rangle$.

(3) Apply the n -qubit subcircuits U and V in parallel to one subregister each. This produces Eq. (26).

In short, this method allows to turn the preparation of a $2n$ -qubit state into the basis gate decomposition of n -qubit operations, thus simplifying the process considerably. For example, for $n = 1$, one can prepare an arbitrary entangled

two-qubit state $|\phi\rangle$ with a single CNOT gate, since B , U , and V are all single-qubit operations that can be decomposed into elementary operations trivially (e.g., via the ZYZ decomposition [3]). Of course, if $|\phi\rangle$ is separable, then $s_0 = 1$, $s_1 = 0$, in which case $|\phi\rangle = |u_0\rangle \otimes |v_0\rangle$ and the CNOT gate is redundant, as we only need to initialize two single-qubit states in parallel.

Let us now consider the application of this method to the preparation of the four-qubit state $|\psi_{\text{island}}^{S=1}\rangle$ corresponding to an island of a spin-1 VBS state [cf. Fig. 4(a)]. Concretely, $|\psi_{\text{island}}^{S=1}\rangle$ results from preparing a product state of two valence bonds, applying the symmetrization operator $\mathcal{S}^{(\mathcal{N}=2)}$ to the central pair of qubits, which encode a single site, and normalizing the resulting state.

$$\begin{aligned} |\psi_{\text{island}}^{S=1}\rangle &= \text{---} \bullet \text{---} \text{---} \bullet \text{---} \text{---} \bullet \text{---} \\ &= \sqrt{\frac{4}{3}} (\mathbb{1} \otimes \mathcal{S}^{(\mathcal{N}=2)} \otimes \mathbb{1}) \left[\begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \\ 0 \end{pmatrix} \right] \\ &= \frac{1}{\sqrt{3}} \left(0, 0, 0, \frac{1}{2}, 0, \frac{1}{2}, -1, 0, 0, 0, -1, 0, \frac{1}{2}, 0, \frac{1}{2}, 0, 0, 0 \right)^T. \end{aligned}$$

In general, preparing such a four-qubit state using the method from Fig. 5(a) requires the preparation of a two-qubit state (via subcircuit B), which takes at most one CNOT gate, and the application of two two-qubit subcircuits U and V , each taking at most three CNOT gates [110]. Hence, a maximum of nine CNOT gates and a circuit depth of at most five CNOT gates, ignoring connectivity constraints, are required to prepare an arbitrary four-qubit state. For the specific case of $|\psi_{\text{island}}^{S=1}\rangle$, its initialization involves a total of seven CNOT gates and a depth of four CNOT gates, since U and V can both be decomposed into a circuit with only two CNOT gates each using the QISKIT `two_qubit_cnot_decompose` method [111], which implements the KAK decomposition [112]. Figure 5(b) presents a scheme of the corresponding quantum circuit.

The application of the Schmidt-decomposition-based initialization method to a six-qubit state such as the spin-3/2 VBS state islands involves the preparation of a three-qubit state through subcircuit B and the application of two three-qubit operations U and V in parallel. The latter take at most 20 CNOT gates [76]. As for the former, even though the generic preparation method discussed in this section is only strictly applicable to states encoded in an even number of qubits, it is possible to consider an asymmetric bipartition, as discussed in Appendix E for the three-qubit case, which yields a maximum of four CNOT gates. Hence, ignoring connectivity constraints, a general six-qubit state can be prepared with at most 47 CNOT gates and a circuit depth of 25 CNOT gates. However, for the particular case of the six-qubit islands in the spin-3/2 VBS, the implementation of B takes three CNOT gates instead of the maximum 4, and the CNOT count of the decomposition of U and V was reduced from the 20 and 19 CNOT gates yielded by the Cirq `three_qubit_matrix_to_operations` method to 14 and 15 CNOT gates, respectively, using the same ZX-calculus-based procedure described in Sec. VII. To sum up, each spin-3/2 island can be prepared with a total of 35 CNOT

gates and a depth of 19 CNOT gates. The respective QASM file is introduced in Appendix H.

The use of this Schmidt-decomposition-based method to initialize the $4S$ -qubit islands of the spin- S VBS states as a strategy to reduce the repetition overhead of the probabilistic method is justified by the significant savings it generates relative to the default option of applying a standard state preparation method, such as that implemented via the QISKIT `initialize` function [113]. For the $S = 1$ case, QISKIT `initialize` prepares the four-qubit island state with a total of 22 CNOT gates, which compares with the seven CNOT gates yielded by the Schmidt-decomposition-based method. As for $S = 3/2$, the number of CNOT gates is reduced from 114 to 35.

Although this constant-depth strategy is clearly the preferred option for the $S = 1$ and $3/2$ cases, for larger local spins (i.e., larger coordination numbers of the lattice) it may be preferable to employ the logarithmic-depth scheme discussed in the previous section, as it only requires the basis gate decomposition of the controlled- $e^{-i\pi S^{(\mathcal{N}=2S)}}$ operation, for which a general implementation is discussed in the next section.

X. LOCAL SYMMETRIZATION VIA LINEAR COMBINATION OF UNITARIES

We have proposed a NISQ-friendly probabilistic method to prepare any spin- S VBS state on quantum hardware, with $S = 1, \frac{3}{2}, 2, \frac{5}{2}, 3, \dots$. The implementation of this method (and of the logarithmic-depth repetition overhead mitigation strategy from Sec. VIII) merely relies upon the determination of a $(2S + 1)$ -qubit subcircuit that implements the controlled $e^{-i\pi S^{(\mathcal{N}=2S)}}$. For the paradigmatic cases $S = 1$ and $3/2$, we have performed an efficient basis gate decomposition of the respective three-qubit and four-qubit operations, as discussed in Sec. VII.

However, for a larger local spin S , this process may become a bottleneck. Even for $S = 2$, corresponding to the local spin of a VBS state defined on a square lattice, the controlled- $e^{-i\pi S^{(\mathcal{N}=2S)}}$ is a five-qubit operation, for which it is already particularly challenging to find a shallow decomposition via standard methods.

This section presents a systematic method to find a circuit that, upon the measurement of ancillary qubits in the fiducial state, leads to the application of the local symmetrization operator $\mathcal{S}^{(\mathcal{N})}$ in the main register. The key idea is that, even though $\mathcal{S}^{(\mathcal{N})}$ is nonunitary, it can be expressed as a linear combination of unitaries, each of which is easy to decompose into basis gates. This approach therefore allows to bypass the basis gate decomposition of a large subcircuit, particularly for high local spin S . The potential savings generated by this method for the spin-2 VBS state, a resource state for universal quantum computation [114], will be discussed.

As a starting point, let us consider the simplest nontrivial case of the symmetrization operator $\mathcal{S}^{(\mathcal{N}=2)}$ acting on two spins $1/2$, which is relevant for the preparation of the one-dimensional spin-1 VBS. Noting that the symmetrization operator is defined as the uniform linear combination of the elements of the symmetric group S_n [cf. Eq. (4) for the

particular case of two spins $1/2$ and Eq. (5) for the general case], we have

$$\mathcal{S}^{(\mathcal{N}=2)} = \frac{1}{2!} \left(\begin{pmatrix} 1 & 2 \\ 1 & 2 \end{pmatrix} + \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix} \right),$$

where Cauchy's two-line notation is used to denote the permutations. In terms of quantum circuits,

$$\mathcal{S}^{(\mathcal{N}=2)} = \frac{1}{2!} \left(\begin{array}{c} \text{---} \times \text{---} \\ \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \\ \text{---} \times \text{---} \end{array} \right) = \frac{1}{2!} \left(\boxed{\mathcal{P}_1^{(\mathcal{N}=2)}} + \boxed{\mathcal{P}_2^{(\mathcal{N}=2)}} \right).$$

The notation used after the second equality will be relevant in the derivation of a systematic method to find the circuits for all permutations $\{\mathcal{P}_i^{(\mathcal{N})}\}_{i=1}^{\mathcal{N}!}$ of an arbitrary number \mathcal{N} of spins $1/2$.

Let us now consider the next case, $\mathcal{N} = 3$. We have

$$\mathcal{S}^{(\mathcal{N}=3)} = \frac{1}{3!} \left(\begin{pmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \end{pmatrix} + \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix} + \begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{pmatrix} \right. \\ \left. + \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix} + \begin{pmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{pmatrix} + \begin{pmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \end{pmatrix} \right),$$

or, in terms of quantum circuits,

$$\mathcal{S}^{(\mathcal{N}=3)} = \frac{1}{3!} \left(\begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} \right)$$

As in the $\mathcal{N} = 2$ case, not only is each term in the linear combination unitary but the corresponding quantum circuit is simple. More importantly, a systematic, inductive way of finding these circuits for arbitrary \mathcal{N} can be derived. We begin with the trivial case $\mathcal{N} = 1$, for which the symmetric group S_1 is the one-element group:

$$\mathcal{S}^{(\mathcal{N}=1)} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \text{---} = \boxed{\mathcal{P}_1^{(\mathcal{N}=1)}}$$

The expression for $\mathcal{S}^{(\mathcal{N}=2)}$ can be derived as follows:

$$\mathcal{S}^{(\mathcal{N}=2)} = \frac{1}{2!} \left(\begin{array}{c} \text{---} \times \text{---} \\ \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \\ \text{---} \times \text{---} \end{array} \right)$$

Similarly, we can obtain $\mathcal{S}^{(\mathcal{N}=3)}$ as

$$\mathcal{S}^{(\mathcal{N}=3)} = \frac{1}{3!} \left(\begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} + \begin{array}{c} \text{---} \times \text{---} \times \text{---} \\ \text{---} \times \text{---} \times \text{---} \end{array} \right)$$

The general induction relation to derive the linear expansion of $\mathcal{S}^{(\mathcal{N})}$ in terms of the $\mathcal{N}!$ circuits for all permutations

for details). As for the select oracle \mathcal{O}_S , there are $4! - 1 = 23$ terms (ignoring the identity), 6 with just 1 CSWAP gate, 11 with 2 CSWAP gates and the remaining 6 with 3 CSWAP gates. Hence, \mathcal{O}_S takes $6 \times 1 + 11 \times 2 + 6 \times 3 = 46$ CSWAP gates, which amounts to $46 \times 7 = 322$ CNOT gates. Noting that the Prepare oracle must be reversed, the total number of CNOT gates is thus $2 \times 46 + 322 = 414$. This number is below the 469 CNOT gates yielded by QISKIT transpile [97] for the decomposition of the five-qubit controlled $e^{-i\pi S^{(4)}}$ and the 414 CNOT gates required by the LCU method are also below the 444 CNOT gates [95] that the optimized quantum Shannon decomposition [76] requires, in general, to decompose a five-qubit gate.

The main limitation of the application of the LCU method to the implementation of $S^{(\mathcal{N})}$ is arguably the potentially large number $\mathcal{N}!$ of ancillary qubits per site (e.g., $4! = 24$ for $S = 2$), which compares with a single ancilla for the local Hadamard test. Nevertheless, the LCU method introduced above corresponds to a *sparse* version, for which $\mathcal{O}_P|0\rangle^{\otimes m} \sim \sum_{i=0}^{m-1} \sqrt{|\alpha_i|} |2^i\rangle$ and $\mathcal{O}_S = \sum_{i=0}^{m-1} |2^i\rangle\langle 2^i| \otimes e^{i\theta_i} V_i$. It is, however, possible to consider a *dense* version that uses only $a = \lceil \log_2(m) \rceil$ ancillas, with $\mathcal{O}_P|0\rangle^{\otimes a} \sim \sum_{i=0}^{m-1} \sqrt{|\alpha_i|} |i\rangle$ and $\mathcal{O}_S = \sum_{i=0}^{m-1} |i\rangle\langle i| \otimes e^{i\theta_i} V_i$. This leads to a more complex implementation of the Select and Prepare oracles, though. For the particular case of the symmetrization operator, each SWAP gate in the select oracle will be controlled by all a ancillas instead of just one. Moreover, the initialization of $\sum_{i=0}^{m-1} \frac{1}{\sqrt{m}} |i\rangle$ is generally more difficult than that of $|W_m\rangle$. The one relevant exception corresponds to the case where $m = \mathcal{N}!$ is a power of 2, in which case \mathcal{O}_P is just the Walsh-Hadamard transform $H^{\otimes \log_2 m}$. This happens to be case for $\mathcal{N} = 2$ (i.e., spin-1), for which the implementation of $S^{(2)}$ via this dense LCU method is entirely equivalent to the local Hadamard test. In general, however, the dense LCU method gives rise to a deeper circuit than the sparse version, which is just a manifestation of the ubiquitous width-depth trade-off (cf. Ref. [120] and references therein).

Finally, we note that East *et al.* [121] devised an implementation in the language of the ZXH calculus of the symmetrization operator acting on \mathcal{N} spins $1/2$ to generate a single spin $\mathcal{N}/2$ that is identical to the LCU-based method introduced in this section. East *et al.* [121] considered ZXH diagrams instead of quantum circuits to represent the VBS states and compute their properties in a fully diagrammatic way.

XI. DISCUSSION

In the previous sections, we have tackled the generally challenging problem [76,122,123] of preparing quantum many-body states on quantum hardware, namely, for the physically relevant class of VBS states of local spin $S = 1, \frac{3}{2}, 2, \frac{5}{2}, 3, \dots$ (cf. Sec. III). In particular, we have developed a probabilistic quantum scheme (cf. Sec. VI), inspired by the construction of the parent AKLT Hamiltonians [26,27], that gives rise to a circuit with depth independent of the lattice size N . In Sec. VII the detailed circuits to prepare the spin-1 and spin-3/2 VBS states were derived, achieving a depth of just 8 and 27 CNOT gates, respectively.

Such low depths, valid for arbitrarily large lattices, render this method especially suitable for NISQ hardware. This contrasts with previous proposals that exploited the tensor-network representation of the VBS states. In one dimension, as discussed in Appendix D, the three-qubit operations $D_{[i]}^T$ shown in Fig. 2(d) can take as many as 20 CNOT gates, which means that our probabilistic method yields a lower depth than the sequential MPS-based one for as few as $N = 3$ sites, with the difference in circuit depth increasing as $\mathcal{O}(20N)$ with further sites. In two dimensions, quantum phase estimation, upon which the method by Schwarz *et al.* [60] is based, and adiabatic state preparation [50,87] yield too deep circuits for NISQ hardware, especially taking into account the need to implement the time-evolution operator of two-spin-3/2 terms, which have support on 6 qubits.

However, thus far the issue of qubit connectivity constraints has been ignored in the determination of the circuit depth for the preparation of the VBS states. Although all-to-all connectivity has been achieved in trapped-ion quantum computers with as many as 20 qubits [124], it remains unclear if such degree of connectivity can be maintained as the number of qubits increases [92], despite ongoing efforts towards this end [125,126]. In superconducting-circuit-based quantum computers, restrictions in the connections between qubits are inevitable, with entangling gates between widely separated qubits being possible only through networks of SWAP gates [127]. Determining how the depth associated with preparing the spin-1 and spin-3/2 VBSs is affected by these qubit connectivity constraints is thus relevant to confirm its feasibility in near-term quantum hardware.

To this end, we will consider the heavy-hex lattice [128] adopted by IBM Quantum in their current and forthcoming devices [129–131]. This architecture would be perfectly suited for the preparation of the spin-3/2 VBS state, were it not for the fact that there is only one qubit per lattice link instead of the two required to create a valence bond. A conceptually simple, yet technically challenging solution to this problem would be to replace every qubit along a link by a ququart (i.e., a four-dimensional qudit), in the spirit of previous proposals [132–134] that exploit higher-dimensional local Hilbert spaces to simplify quantum circuits. A more feasible approach is to make use of only a fraction of the qubits available in the heavy-hex lattice, distributing them spatially in a way that suits the initialization of the spin-3/2 VBS. The resulting layout of the qubits across the heavy-hex lattice is illustrated in Figs. 7(a) and 7(b) for the cases where the bare probabilistic method and the combination of the preparation of the six-qubit islands and the probabilistic method at only one sublattice are employed, respectively. The circuit depth in CNOT gates for both cases is presented in Table I; the details of the calculations can be found in Appendix G. Even though, as expected, the depth does increase when qubit connectivity constraints are accounted for, the predicted depths of 51 CNOT gates (for the bare probabilistic method) and 105 CNOT gates (with mitigation of the repetition overhead) are still very likely to be within reach of noisy intermediate-scale quantum computers in the next few years.

For the spin-1 VBS state in one dimension, the restrictions related to the couplings between qubits are even less impactful: the circuit depth for the probabilistic method with

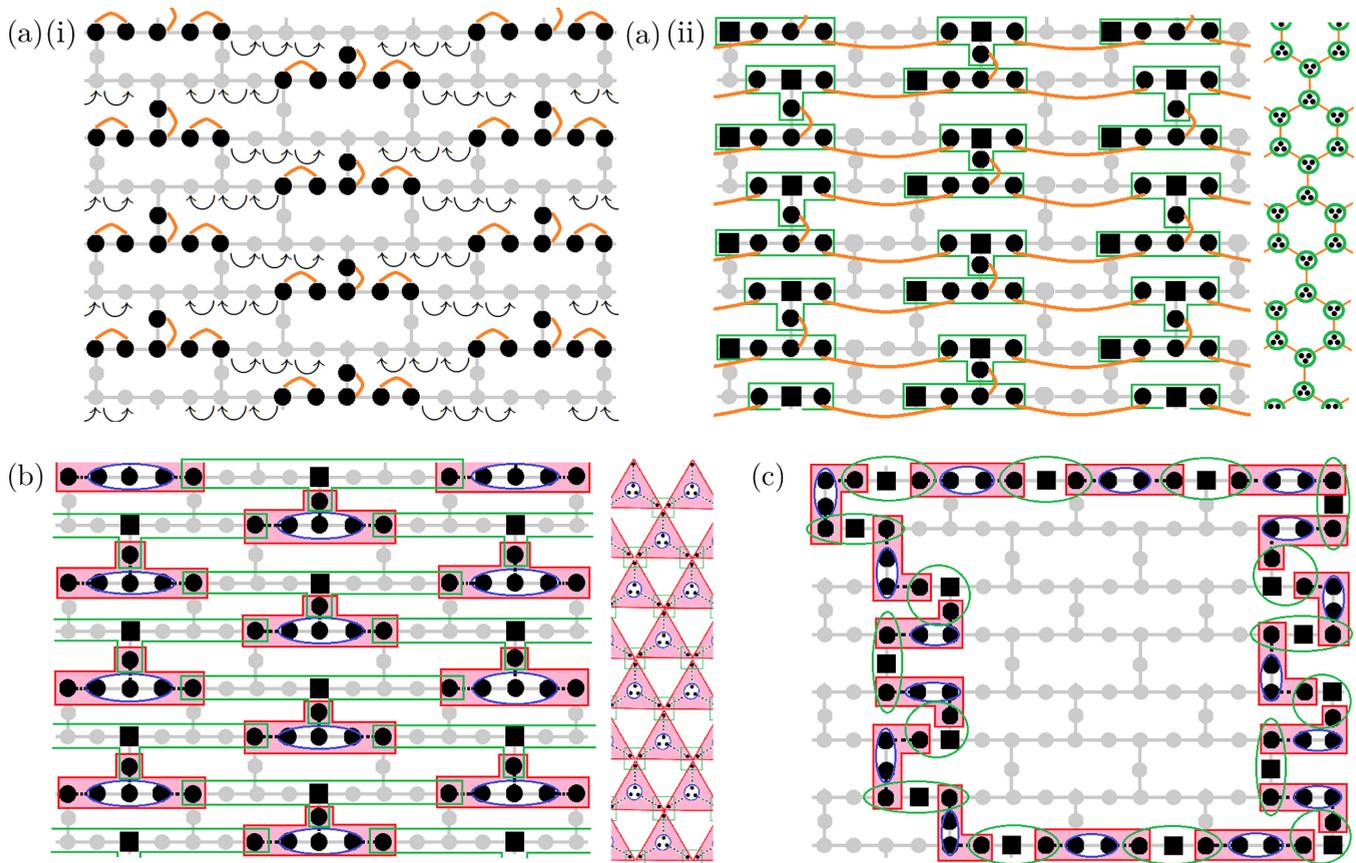


FIG. 7. Schemes of implementation of preparation of VBS states on IBM Quantum heavy-hex lattice [130]. (a) Two-step description of preparation of spin-3/2 VBS state via bare probabilistic method. (i) In each set of six adjacent qubits shown as black dots, valence bonds are prepared at pairs of qubits linked by orange curved lines by applying the two-qubit circuit from Fig. 3(a). Then, from each such six-qubit set, two qubits are moved past three redundant qubits (shown as gray dots), one towards the right and the other towards the left, as depicted by the curly arrows. This ensures that all three qubits that encode each site are nearby at the start of the second step. (ii) The local Hadamard test with $\mathcal{U} = e^{-i\pi S^{(N=3)}}$ [cf. Fig. 3(b)] is applied at the four-qubit sets enclosed by green boxes. Site ancillas, which were redundant qubits in (i), are now represented by black squares. T-shaped green boxes are associated with sites of one sublattice, and linear green boxes with those of the other. For the sake of clarity, a scheme of the segment of a honeycomb lattice corresponding to this circuit is shown on the right-hand side. (b) A single scheme describing the preparation of the spin-3/2 VBS state via the probabilistic method combined with the deterministic preparation of the six-qubit islands (cf. Sec. IX) to mitigate the repetition overhead. The main difference relative to (a) is the fact that the first step already symmetrizes the sites of one sublattice (as depicted by the blue loops), in addition to preparing all valence bonds (here represented as dashed black lines, as in Fig. 4). Hence, in second step only T-shaped green boxes require the application of the local Hadamard test. (c) Scheme of preparation of spin-1 VBS state for 32-site ring via the probabilistic method preceded by the initialization of the four-qubit islands. As in (b), islands are illustrated with the color scheme adopted in Fig. 4.

and without $\mathcal{O}(1)$ -depth mitigation of the repetition overhead increases from 10 to 17 CNOT gates and from 8 to 11 CNOT gates, respectively. For concreteness, Fig. 7(c) shows a scheme of the layout adopted in the IBM Quantum heavy-hex lattice [130] to prepare the spin-1 VBS state for a 32-site ring.

The downside of this low circuit depth is the probabilistic nature of the method, which means that the circuit must be repeated multiple times to successfully prepare the VBS state. In a sense, this reduction of the circuit depth by increasing the number of repetitions follows the spirit of hybrid variational algorithms [13], which are naturally suited for NISQ hardware. Two strategies that achieve a quadratic reduction of this repetition overhead were devised by exploiting the entanglement structure of the initial state $|\psi_{\text{pre-VBS}}\rangle$. The latter

(cf. Sec. IX) produces a $\mathcal{O}(1)$ -depth overhead but requires the generic preparation of $4S$ -qubit states, thus being more suitable for low spin- S . The former (cf. Sec. VIII) results in a $\mathcal{O}(\ln N)$ -depth layer but only involves applying the controlled $e^{-i\pi S^{(N=2S)}}$ used in the probabilistic method, for which a systematic scheme that avoids the basis gate decomposition of the corresponding $(2S + 1)$ -qubit unitary operation was developed in Sec. X, thus making it a more suitable option for high spin S . It should be noted that the real-time processing of mid-circuit measurements, which is a key element of the former method, is an active field of research [135], particularly due to its importance within the context of quantum error correction [12]. In fact, a number of quantum routine proposals that make use of mid-circuit measurements have been put forth recently [136–139].

TABLE I. Circuit depth in CNOT gates for preparation of spin-1 and spin-3/2 VBS states assuming different cases of qubit connectivity constraints, with and without mitigation of repetition overhead. Cf. Appendix G for details of calculations.

Local spin, Qubit connectivity	Probabilistic method at all lattice sites	4S-Qubit Islands and probabilistic method at one sublattice
$S = 1$, all-to-all	8	11
$S = 1$, linear	10	17
$S = 3/2$, all-to-all	27	45
$S = 3/2$, heavy-hex	51	105

Table II presents the average number of repetitions required to prepare the spin-1 and spin-3/2 VBS states with and without employing a mitigation overhead repetition strategy. Although the scaling of the average number of repetitions is exponential in the number of lattice sites N in both cases, the mitigation of the repetition overhead can be decisive to make the probabilistic preparation of VBS states feasible for the intermediate values of N at which quantum advantage can be achieved with near-term quantum computers.

An important application with a plausible prospect of achievable quantum advantage to which the preparation of the spin-3/2 VBS state on quantum hardware makes a relevant contribution is the simulation of a naturally occurring gapped Hamiltonian such as the spin-3/2 AKLT model, or possibly a nearby but nonintegrable model, to aid the experimental realization of a resource state for measurement-based quantum computation (MBQC) [31–33]. This could be framed within the wider effort of using near-term quantum hardware to support the development of fault-tolerant quantum computers [140,141].

The idea of employing the ground state of a naturally occurring gapped Hamiltonian with an appropriate entanglement structure [34] as the resource state for MBQC is appealing for the flexible state preparation via cooling, on the one hand, and the stability against local perturbations, on the other [142]. Having prepared such a resource state, no entangling operations are required to perform the computations; single-qubit gates and local measurements suffice. The cluster state [143] is the canonical example of a resource state for MBQC, but it cannot occur as the exact ground state of any naturally occurring physical system [36]. Given the proof that the spin-

TABLE II. Average number of repetitions required to successfully prepare spin-1 and spin-3/2 VBS states via probabilistic method introduced in Sec. VI with and without employing the repetition overhead mitigation strategies proposed in Secs. VIII and IX. Scaling of average number of repetitions for spin-1 and spin-3/2 is $(\frac{4}{3})^N$ and 2^N in unmitigated case, and $(\frac{4}{3})^{N/2}$ and $2^{N/2}$ if mitigation is applied.

N	10	20	30	40	50
$S = 1$, unmitigated	18	315	5600	99 000	1.8×10^6
$S = 1$, mitigated	4	18	75	315	1,300
$S = 3/2$, unmitigated	1000	10^6	10^9	10^{12}	10^{15}
$S = 3/2$, mitigated	32	1000	33 000	10^6	3.3×10^7

3/2 VBS state is a resource state for MBQC [29,30], its experimental realization in solid-state platforms has been considered [30,37,38] to develop a measurement-based quantum computer.

Computing the expected properties of the ground state of such a naturally occurring parent Hamiltonian, calculating its spectral gap, and determining the nature of its low-lying excitations could provide valuable assistance in the experimental realization of the desired robust resource state. However, computing these properties using conventional numerical methods is a formidable challenge, as demonstrated by the recent confirmation of the nonzero spectral gap of the spin-3/2 AKLT model [40,41]. The semianalytical proof developed by Lemm *et al.* [41] involved a DMRG calculation on a 36-site cluster, although the convergence to the excited states in all spin sectors would not have been possible without making use of the AKLT construction [26,27], which is only strictly possible at the integrable point. The challenge is certainly even more daunting away from the integrable point, where experimental systems are likely to be found [38]. Ganesh *et al.* [39] investigated the phase diagram of a spin-3/2 model that interpolates between the Heisenberg and AKLT models using exact diagonalization on clusters with up to $N = 18$ spins 3/2.

Digital quantum simulation methods, in turn, may allow for a more viable calculation of these quantities. In the NISQ era, a hybrid variational algorithm such as the variational quantum eigensolver (VQE) [42] may be used to compute the ground state and low-lying excited states [43–46] of such nonintegrable models close to the spin-3/2 AKLT model. In this paper we have addressed one of the key challenges involved in this VQE simulation: The preparation of an initial state that overlaps significantly with the target ground state of the nearby nonintegrable model. Such educated guess is essential to simplify the optimization process by reducing the number of layers of the ansatz applied on top of it and by avoiding barren plateaus [144,145]. An additional challenge that remains unaddressed in the literature is how to adapt VQE to the simulation of spin-3/2 (or, more generally, higher-spin) models, as all applications of VQE to the study of quantum magnetism have been restricted to local spin 1/2 degrees of freedom [22,115,146–153]. Improvements in the quantum hardware are also necessary, although the consistent rise in the quantum volume achieved by the leading manufacturers [154,155] and a recent simulation involving a two-qubit gate depth of up to 159 on 20 qubits [156] suggest the required hardware may already be at our disposal in the next few years. It should be noted that the measurements of the ancillas in the bare probabilistic method or in the second part of the probabilistic method with repetition overhead mitigation can be deferred until the end of the circuit, thus avoiding mid-circuit measurements.

An alternative application where quantum advantage can be attained potentially sooner is the simulation of the quench dynamics [157] of the spin-1 VBS state in one dimension. Even though DMRG [52,65–67] and time-dependent variants thereof (e.g., tDMRG [158,159] and tangent-space methods [160,161]) have been extremely successful at simulating static properties and short-time dynamics of one-dimensional systems (cf. Refs. [52,53,67] and references therein), for a

sufficiently long duration of the time evolution the truncation error due to the bounded bond dimension of the underlying MPS becomes exceedingly large, signaling the “runaway” time at which the numerical simulation ceases to be reliable [162]. Digital quantum computation offers an exponential speed-up in the simulation of quantum dynamics [163], so it could become the leading method to validate experiments performed with cold atoms in optical lattices [164]. Schemes to realize the spin-1 VBS state with cold atoms have been put forth [165,166]. By making use of the probabilistic method herein proposed to prepare the spin-1 VBS state on quantum hardware instead of the deterministic MPS-based method discussed in Sec. IV, considerable savings of the circuit depth devoted to the state initialization can be achieved, leaving more circuit depth available to implement the time evolution.

Understanding the impact of noise in its multiple forms (e.g., gate errors, decoherence effects, measurement errors) in the preparation of VBS states via the probabilistic method we have proposed is an important line of future work. Although it is likely that our probabilistic method will handle decoherence effects far more effectively than the existing deterministic methods thanks to the significant reduction of circuit depth, it is interesting to understand the impact of cross-talk errors [167] in the implementation of multiple CNOT gates in parallel and whether the additional measurement errors arising from the implementation of the local symmetrization via ancillary qubits can be adequately addressed through scalable quantum error mitigation strategies [168]. Large-scale *in silico* simulations with realistic noise models and implementations on actual quantum processors will be essential to address these questions.

XII. CONCLUSIONS

In summary, we have proposed a method to prepare VBS states of arbitrary local spin S on quantum hardware. Inspired by the construction of the parent AKLT Hamiltonians, this method consists of initializing a product state of valence bonds, after which the local symmetrization operator is applied at all sites in parallel. As a result, the depth of the resulting quantum circuit is independent of the lattice size N , at the cost of requiring multiple independent repetitions to achieve success. Two schemes to reduce the average number of repetitions were developed by exploiting the entanglement structure of the initial state, one introducing a constant depth overhead but requiring the initialization of $4S$ -qubit states and another leading to a $\mathcal{O}(\ln N)$ -depth overhead but involving the same building block as the main probabilistic method. The former is therefore appropriate for low S , while the latter is the method of choice for larger S . An alternative approach to implement the local symmetrization operator was also devised, bypassing the basis gate decomposition of the building block of the probabilistic method, which can be too onerous for high S .

These general methods were then applied to the particular cases of $S = 1$ and $3/2$. Shallow circuits to implement the probabilistic method were devised using state-of-the-art basis gate decomposition methods, yielding circuits with depth 8 and 27 CNOT gates for the preparation of the spin-

1 and spin-3/2 VBS states, respectively. Two applications with prospective quantum advantage were identified, one for each case. The impact of qubit connectivity constraints and the size of the repetition overhead required to prepare these valence bond states at the intermediate lattice sizes N for which such quantum advantage is foreseeable were considered.

More broadly, in comparison with the tensor-network-based methods to prepare VBS states on quantum hardware, the quantum routine herein proposed trades circuit depth for the repetition of the same shallow circuit, which follows the spirit of hybrid variational algorithms. Extending this depth-repetitions trade-off to the preparation of other classes of quantum many-body states could be a prolific strategy.

ACKNOWLEDGMENTS

B.M. acknowledges financial support from the FCT PhD scholarship No. SFRH/BD/08444/2020. P.M.Q.C. acknowledges financial support from FCT Grant No. SFRH/BD/150708/2020. J.F.R. acknowledges financial support from the Ministry of Science and Innovation of Spain (grant No. PID2019-109539GB-41), from Generalitat Valenciana (Grants No. Prometeo2021/017 and MFA/2022/045), and from FCT (Grant No. PTDC/FIS-MAC/2045/2021).

APPENDIX A: MATRIX REPRESENTATION OF SYMMETRIZATION \mathcal{S}

Although the matrix representation of the local symmetrization operator \mathcal{S} can be obtained via its general definition [cf. Eq. (5)], a leaner approach can be adopted instead. Concretely, at site $n \in \Omega$,

$$\mathcal{S}_n^{(\mathcal{N})} = C' \prod_{S' \neq S} ((\vec{\mathcal{S}}_{\text{total}}^n)^2 - S'(S' + 1)), \quad (\text{A1})$$

where, in this case, $\vec{\mathcal{S}}_{\text{total}}^n = \sum_{i=1}^{\mathcal{N}_n} \vec{\mathcal{S}}_i$ is the total spin resulting from the sum of all \mathcal{N}_n spins $1/2$ and $S = \mathcal{N}_n/2$ is the spin that we wish to associate with the local degree of freedom. C' is a normalization constant that ensures that spin- S states are acted on trivially by \mathcal{S}_n .

1. Spin 1

For the spin-1 case, we have two spins $1/2$ per site, in which case the matrix representations of the three Cartesian components of $\vec{\mathcal{S}}_{\text{total}}$ are of the form $S_{\text{total}}^i = S^i \otimes \mathbb{1} + \mathbb{1} \otimes S^i$, with

$$S^x = \frac{1}{2} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad S^y = \frac{1}{2} \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad S^z = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The matrix representation of $(\vec{\mathcal{S}}_{\text{total}})^2 = (S_{\text{total}}^x)^2 + (S_{\text{total}}^y)^2 + (S_{\text{total}}^z)^2$ is then

$$(\vec{\mathcal{S}}_{\text{total}})^2 = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}. \quad (\text{A2})$$

The symmetrization operator removes $S_{\text{total}} = 0$ states, so $\mathcal{S}^{(\mathcal{N}=2)} = C'(\vec{S}_{\text{total}}^2 - 0(0+1)\mathbb{1})$. Setting $C' = 1/2$ gives

$$\mathcal{S}^{(\mathcal{N}=2)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (\text{A3})$$

2. Spin 3/2

For the spin-3/2 case, there are three spins 1/2 per site, so the matrix representations of the three Cartesian components of \vec{S}_{total} are of the form $S_{\text{total}}^i = S^i \otimes \mathbb{1} \otimes \mathbb{1} + \mathbb{1} \otimes S^i \otimes \mathbb{1} + \mathbb{1} \otimes \mathbb{1} \otimes S^i$, with

$$S^x = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad S^y = \frac{1}{\sqrt{2}i} \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix},$$

$$S^z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}.$$

The matrix representation of $(\vec{S}_{\text{total}})^2 = (S_{\text{total}}^x)^2 + (S_{\text{total}}^y)^2 + (S_{\text{total}}^z)^2$ is therefore

$$(\vec{S}_{\text{total}})^2 = \frac{1}{4} \begin{pmatrix} 15 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 7 & 4 & 0 & 4 & 0 & 0 & 0 \\ 0 & 4 & 7 & 0 & 4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 7 & 0 & 4 & 4 & 0 \\ 0 & 4 & 4 & 0 & 7 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 & 7 & 4 & 0 \\ 0 & 0 & 0 & 4 & 0 & 4 & 7 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 15 \end{pmatrix}. \quad (\text{A4})$$

Symmetrizing removes $S_{\text{total}} = 1/2$ states, so $\mathcal{S}^{(\mathcal{N}=3)} = C'(\vec{S}_{\text{total}}^2 - \frac{1}{2}(\frac{1}{2} + 1)\mathbb{1})$. Setting $C' = 1/3$ gives

$$\mathcal{S}^{(\mathcal{N}=3)} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (\text{A5})$$

APPENDIX B: NORM REDUCTION DUE TO LOCAL SYMMETRIZATION $\mathcal{S}^{(\mathcal{N})}$ AND OVERLAP $\langle \psi_{\text{VBS}} | \psi_{\text{pre-VBS}} \rangle$

First, we show that $\langle \psi_{\text{VBS}} | \psi_{\text{pre-VBS}} \rangle = \langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle$, where $|\psi_{\text{VBS}}\rangle$ is the unnormalized VBS state. Noting that the symmetrization operator $\mathcal{S}^{(\mathcal{N})}$ is Hermitian ($\mathcal{S}^{(\mathcal{N})} = (\mathcal{S}^{(\mathcal{N})})^\dagger$) and idempotent ($(\mathcal{S}^{(\mathcal{N})})^2 = \mathcal{S}^{(\mathcal{N})}$), it follows that

$$\begin{aligned} \langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle &= (\langle \psi_{\text{pre-VBS}} | \otimes_n (\mathcal{S}_n^{(\mathcal{N})})^\dagger) (\otimes_n \mathcal{S}_n^{(\mathcal{N})} | \psi_{\text{pre-VBS}} \rangle) \\ &= \langle \psi_{\text{pre-VBS}} | (\otimes_n (\mathcal{S}_n^{(\mathcal{N})})^2) | \psi_{\text{pre-VBS}} \rangle \end{aligned}$$

$$\begin{aligned} &= (\langle \psi_{\text{pre-VBS}} | \otimes_n (\mathcal{S}_n^{(\mathcal{N})})^\dagger) | \psi_{\text{pre-VBS}} \rangle \\ &= \langle \psi_{\text{VBS}} | \psi_{\text{pre-VBS}} \rangle. \end{aligned} \quad (\text{B1})$$

In terms of the normalized VBS state, $|\tilde{\psi}_{\text{VBS}}\rangle \equiv \frac{|\psi_{\text{VBS}}\rangle}{\sqrt{\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle}}$, Eq. (B1) can be expressed as

$$\langle \tilde{\psi}_{\text{VBS}} | \psi_{\text{pre-VBS}} \rangle = \sqrt{\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle}. \quad (\text{B2})$$

Now, to arrive at Eq. (8) in the main text, we need to show that, at each application of the symmetrization operator $\mathcal{S}^{(\mathcal{N})}$, the norm of $|\psi_{\text{pre-VBS}}\rangle$ is reduced by a constant factor p , ultimately leading to an exponentially vanishing overlap in the number of sites N . We will derive this result explicitly for the spin-1 and spin-3/2 VBS states first, and then we will present the general case.

1. Spin-1 valence-bond-solid State

The matrix representation of $\mathcal{S}^{(\mathcal{N}=2)}$ for the spin-1 case is given by (A3). Focusing only on the valence bonds involving site n , $|\psi_{\text{pre-VBS}}\rangle$ can be written as

$$\begin{aligned} |\psi_{\text{pre-VBS}}\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\rangle_{n-1,R} |\downarrow\rangle_{n,L} - |\downarrow\rangle_{n-1,R} |\uparrow\rangle_{n,L}) \\ &\otimes \frac{1}{\sqrt{2}} (|\uparrow\rangle_{n,R} |\downarrow\rangle_{n+1,L} - |\downarrow\rangle_{n,R} |\uparrow\rangle_{n+1,L}) \otimes \dots, \end{aligned} \quad (\text{B3})$$

where the ellipsis includes all remaining valence bonds, which do not involve any of the two spins 1/2 from site n . Expanding the tensor product of the two highlighted valence bonds in the eigenbasis of $\mathcal{S}^{(\mathcal{N}=2)} - \{ |s, m_s\rangle \}$, $s \in \{0, 1\}$, $m_s \in \{-s, -s+1, \dots, s-1, s\}$ —at site n gives

$$\begin{aligned} |\psi_{\text{pre-VBS}}\rangle &= \frac{1}{2} \left[-|1, 1\rangle_n \otimes (|\downarrow\rangle_{n-1,R} |\downarrow\rangle_{n+1,L}) \right. \\ &\quad - |1, -1\rangle_n \otimes (|\uparrow\rangle_{n-1,R} |\uparrow\rangle_{n+1,L}) \\ &\quad + |1, 0\rangle_n \otimes \frac{1}{\sqrt{2}} (|\uparrow\rangle_{n-1,R} |\downarrow\rangle_{n+1,L} \\ &\quad + |\downarrow\rangle_{n-1,R} |\uparrow\rangle_{n+1,L}) + |0, 0\rangle_n \\ &\quad \left. \otimes \frac{1}{\sqrt{2}} (|\uparrow\rangle_{n-1,R} |\downarrow\rangle_{n+1,L} + |\downarrow\rangle_{n-1,R} |\uparrow\rangle_{n+1,L}) \right] \\ &\otimes \dots \end{aligned} \quad (\text{B4})$$

Upon application of $\mathcal{S}^{(\mathcal{N}=2)}$ at site n , the first three terms on the right-hand side of Eq. (B4) remain unchanged (since $|1, m_s\rangle$, $m_s \in \{1, 0, -1\}$, are eigenstates of $\mathcal{S}^{(\mathcal{N}=2)}$ with eigenvalue 1), while the fourth term vanishes (since $|0, 0\rangle$ is an eigenstate of $\mathcal{S}^{(\mathcal{N}=2)}$ with eigenvalue 0). As a result, assuming $|\psi_{\text{pre-VBS}}\rangle$ is normalized, the norm is reduced from 1 to 3/4. In general, applying the local symmetrization operator $\mathcal{S}^{(\mathcal{N}=2)}$ at one site of a spin-1 $|\psi_{\text{pre-VBS}}\rangle$ state reduces the norm by a factor 3/4.

Starting from a normalized spin-1 $|\psi_{\text{pre-VBS}}\rangle$, the application of the N symmetrization operators in parallel leads to an exponential decrease of the norm of the resulting $|\psi_{\text{VBS}}\rangle = \otimes_{n=1}^N \mathcal{S}_n^{(\mathcal{N}=2)} |\psi_{\text{pre-VBS}}\rangle$. Asymptotically,

$$\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle = \left(\frac{3}{4}\right)^N \quad (\text{spin } 1). \quad (\text{B5})$$

However, for finite one-dimensional spin-1 AKLT models, there is a deviation from the asymptotic limit that depends on the boundary conditions.

(i) For periodic boundary conditions (i.e., N -site ring),

$$\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle = (3/4)^N + 3(-1/4)^N. \quad (\text{B6})$$

(ii) For open boundary conditions (i.e., a N -site chain where the outermost spins $1/2$ are not connected):

(a) If outermost spins are aligned (e.g., both $|\uparrow\rangle$),

$$\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle = (3/4)^N - (-1/4)^N. \quad (\text{B7})$$

(b) If the outermost spins are antialigned (e.g., one $|\uparrow\rangle$ and another $|\downarrow\rangle$),

$$\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle = (3/4)^N + (-1/4)^N. \quad (\text{B8})$$

(c) If the outermost spins are neither aligned nor antialigned, the finite-size correction takes values between Eqs. (B7) and (B8).

Naturally, as $N \rightarrow \infty$, in all cases the finite-size effects vanish and the boundary conditions become redundant.

2. Spin-3/2 valence-bond-solid State

The local symmetrization operator for the spin-3/2 case is given by Eq. (A5). Highlighting only the valence bonds involving site n , and denoting the three nearest neighbors by n'_s , $s \in \{1, 2, 3\}$, $|\psi_{\text{pre-VBS}}\rangle$ can be written as

$$\begin{aligned} |\psi_{\text{pre-VBS}}\rangle &= \frac{1}{\sqrt{2}}(|\uparrow\rangle_{(n;n'_1)}|\downarrow\rangle_{(n'_1;n)} - |\downarrow\rangle_{(n;n'_1)}|\uparrow\rangle_{(n'_1;n)}) \\ &\otimes \frac{1}{\sqrt{2}}(|\uparrow\rangle_{(n;n'_2)}|\downarrow\rangle_{(n'_2;n)} - |\downarrow\rangle_{(n;n'_2)}|\uparrow\rangle_{(n'_2;n)}) \\ &\otimes \frac{1}{\sqrt{2}}(|\uparrow\rangle_{(n;n'_3)}|\downarrow\rangle_{(n'_3;n)} - |\downarrow\rangle_{(n;n'_3)}|\uparrow\rangle_{(n'_3;n)}) \\ &\otimes \dots \end{aligned} \quad (\text{B9})$$

Expanding the tensor product of the three highlighted valence bonds in the eigenbasis of $\mathcal{S}^{(\mathcal{N}=3)}$ at site n gives

$$\begin{aligned} |\psi_{\text{pre-VBS}}\rangle &= \frac{1}{2\sqrt{2}} \left[\left| \frac{3}{2}, \frac{3}{2} \right\rangle_n \otimes |\downarrow\downarrow\downarrow\rangle_{n'} - \left| \frac{3}{2}, -\frac{3}{2} \right\rangle_n \otimes |\uparrow\uparrow\uparrow\rangle_{n'} - \left| \frac{3}{2}, \frac{1}{2} \right\rangle_n \otimes \frac{1}{\sqrt{3}}(|\uparrow\downarrow\downarrow\rangle_{n'} + |\downarrow\uparrow\downarrow\rangle_{n'} + |\downarrow\downarrow\uparrow\rangle_{n'}) \right. \\ &\quad - \left. \left| \frac{3}{2}, -\frac{1}{2} \right\rangle_n \otimes \frac{1}{\sqrt{3}}(|\uparrow\uparrow\downarrow\rangle_{n'} + |\uparrow\downarrow\uparrow\rangle_{n'} + |\downarrow\uparrow\uparrow\rangle_{n'}) + \left| \frac{1}{2}, \frac{1}{2} \right\rangle_n \otimes \frac{1}{\sqrt{2}}(|\uparrow\downarrow\downarrow\rangle_{n'} - |\downarrow\downarrow\uparrow\rangle_{n'}) \right. \\ &\quad - \left. \left| \frac{1}{2}, \frac{1}{2} \right\rangle_n' \otimes \left(\frac{1}{\sqrt{6}}|\downarrow\downarrow\uparrow\rangle_{n'} - \frac{\sqrt{2}}{\sqrt{3}}|\downarrow\uparrow\downarrow\rangle_{n'} + \frac{1}{\sqrt{6}}|\uparrow\downarrow\downarrow\rangle_{n'} \right) + \left| \frac{1}{2}, -\frac{1}{2} \right\rangle_n \otimes \frac{1}{\sqrt{2}}(|\downarrow\uparrow\uparrow\rangle_{n'} - |\uparrow\uparrow\downarrow\rangle_{n'}) + \left| \frac{1}{2}, -\frac{1}{2} \right\rangle_n' \right. \\ &\quad \left. \otimes \left(\frac{1}{\sqrt{6}}|\downarrow\uparrow\uparrow\rangle_{n'} - \frac{\sqrt{2}}{\sqrt{3}}|\uparrow\downarrow\uparrow\rangle_{n'} + \frac{1}{\sqrt{6}}|\uparrow\uparrow\downarrow\rangle_{n'} \right) \right] \otimes \dots \end{aligned} \quad (\text{B10})$$

Upon application of $\mathcal{S}^{(\mathcal{N}=3)}$ at site n , the last four terms on the right-hand side of Eq. (B10) vanish, which reduces the norm from 1 to $1/2$. Hence, applying the local symmetrization operator $\mathcal{S}^{(\mathcal{N}=3)}$ at one site of a spin- $\frac{3}{2}$ $|\psi_{\text{pre-VBS}}\rangle$ state decreases the norm by a factor $1/2$. In the asymptotic limit, regardless of the boundary conditions, the norm of the VBS state $|\psi_{\text{VBS}}\rangle = \bigotimes_{n=1}^N \mathcal{S}_n^{(\mathcal{N}=3)} |\psi_{\text{pre-VBS}}\rangle$ at a lattice with N sites is

$$\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle = \left(\frac{1}{2} \right)^N \quad (\text{spin } 3/2). \quad (\text{B11})$$

3. General valence-bond-solid state

In general, ignoring finite-size effects, the norm of the VBS, $|\psi_{\text{VBS}}\rangle = \bigotimes_n \mathcal{S}_n |\psi_{\text{pre-VBS}}\rangle$, is given by

$$\langle \psi_{\text{VBS}} | \psi_{\text{VBS}} \rangle = \left(\frac{2S+1}{2^{2S}} \right)^N \equiv p^N. \quad (\text{B12})$$

In words, $p \equiv \frac{2S+1}{2^{2S}}$ corresponds to the fraction of symmetric spin states selected by \mathcal{S} at each site or, equivalently, p is the ratio between the number of states corresponding to the desired local spin $\mathcal{N}_n/2$ and the total number of states resulting from adding the \mathcal{N}_n spins $1/2$.

For concreteness, let us consider the previous two examples plus the spin-2 VBS state on a square lattice.

(i) Spin-1 VBS: the addition of two spins $1/2$ gives $\frac{1}{2} \oplus \frac{1}{2} = 0, 1$, so we have $2 \times 1 + 1 = 3$ spin-1 states and $2 \times 0 + 1 = 1$ spin-0 state, yielding $p = 3/4$.

(ii) Spin-3/2 VBS: adding three spins $1/2$ gives $\frac{1}{2} \oplus \frac{1}{2} \oplus \frac{1}{2} = \frac{1}{2}, \frac{1}{2}, 3/2$, so there are $2 \times \frac{3}{2} + 1 = 4$ spin-3/2 states and $2 \times (2 \times \frac{1}{2} + 1) = 4$ spin $1/2$ states, thus $p = 4/8 = 1/2$.

(iii) Spin-2 VBS: adding four spins $1/2$ gives $\frac{1}{2} \oplus \frac{1}{2} \oplus \frac{1}{2} \oplus \frac{1}{2} = 0, 0, 1, 1, 2$, so there are $2 \times (2 \times 0 + 1) = 2$ spin-0 states, $3 \times (2 \times 1 + 1) = 9$ spin-1 states and $2 \times 2 + 1 = 5$ spin-2 states, in which case $p = 5/16$.

APPENDIX C: INITIALIZATION OF MATRIX PRODUCT STATES WITH PHYSICAL INDEX DIMENSION $D = 4$ AND VIRTUAL INDEX DIMENSION $D = 2$ ON QUANTUM HARDWARE

This section explains how the MPS form of the 1D spin-1 VBS state can be exploited to construct the circuits shown in Fig. 2 of the main text. In fact, this discussion applies to an arbitrary left-canonical MPS with virtual and physical index dimensions $\Lambda = 2$ and $d = 4$, of which the spin-1 VBS state is one example. First, the case of open boundary conditions will be considered, resulting in a deterministic

clearer. Naturally, the resulting sum over x is redundant. Prior to the application of $D_{[1]}$, there were two open legs s_1^L and s_1^R encoding the local degree of freedom at site 1. After applying $D_{[1]}$, one of those open legs, now termed Q_1 , is separated from the MPS, as desired. The other encodes the virtual index α that connects sites 1 and 2. This will be an input for the next MPD, $D_{[2]}$.

The only missing element is how to construct the enlarged tensor A'_1 from the original tensor A_1 . Merging all incoming/outgoing legs into a single incoming/outgoing leg, A_1 and A'_1 are, respectively, 4×2 and 4×4 matrices. Because A_1 and A'_1 share the same incoming legs, the columns of A_1 can be set as the first two columns of A'_1 :

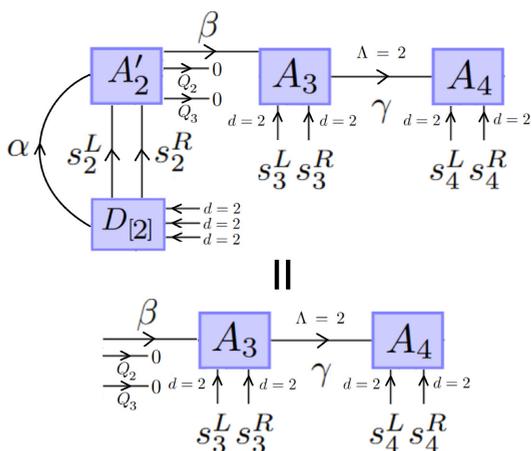
$$A'_1 = \begin{pmatrix} \uparrow & \uparrow & \uparrow & \uparrow \\ A_1^{s_1^L, s_1^R} & A_1^{s_1^L, s_1^R} & \vec{u} & \vec{v} \\ \downarrow & \downarrow & \downarrow & \downarrow \end{pmatrix}. \tag{C2}$$

The fact that A_1 is left-normalized ensures that $A_1^{s_1^L, s_1^R}$ and $A_1^{s_1^L, s_1^R}$ are orthogonal to each other and normalized. For A'_1 to be unitary, its columns (and rows) must form an orthonormal set. Hence, to complete the construction of A'_1 , we just need to find the two remaining columns \vec{u} and \vec{v} , which must be normalized four-dimensional vectors orthogonal to each other and to the first two columns. Concretely, (\vec{u}, \vec{v}) span the kernel of the 2×4 matrix

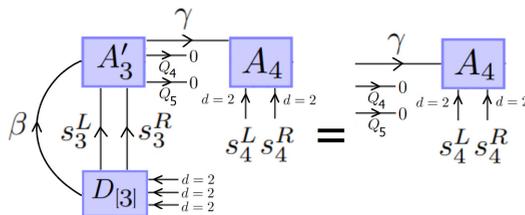
$$\begin{pmatrix} \leftarrow & (A_1^{s_1^L, s_1^R} 0)^\dagger & \rightarrow \\ \leftarrow & (A_1^{s_1^L, s_1^R} 1)^\dagger & \rightarrow \end{pmatrix}. \tag{C3}$$

The MPD $D_{[1]}$ is simply given by the adjoint of Eq. (C2).

For the remaining sites, the construction of the respective MPDs proceeds in a similar spirit to what was described for the first site. The main difference is that more dummy outgoing legs need to be introduced to ensure the enlarged tensor, when cast in matrix form, is unitary. The resulting MPDs are three-qubit operations. For the second site,

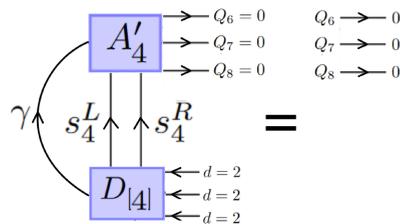


where $A_2^{\alpha, s_2^L, s_2^R} = A_2^{\prime \alpha, s_2^L, s_2^R} \beta_{\beta, 0, 0}$ and $D_{[2]} = A_2^{\prime \dagger}$. Similarly, for the third site, we have



where $A_3^{\beta, s_3^L, s_3^R} = A_3^{\prime \beta, s_3^L, s_3^R} \gamma_{\gamma, 0, 0}$ and $D_{[3]} = A_3^{\prime \dagger}$. Of the three incoming legs, two encode the local degree of freedom at site $i = 2, 3$, and the other is associated with the virtual index that links sites $i - 1$ and i . After the application of the MPD $D_{[i]}$, two of the legs become disentangled from the remainder of the MPS, while the third contains the virtual index that connects sites i and $i + 1$.

Finally, at the last site, since there are three incoming legs and no outgoing ones, three dummy legs must be added, with $A_4^{\gamma, s_4^L, s_4^R} = A_4^{\prime \gamma, s_4^L, s_4^R} 0_{0, 0, 0}$. After applying the respective MPD $D_{[4]} = A_4^{\prime \dagger}$, all three remaining legs are disentangled.

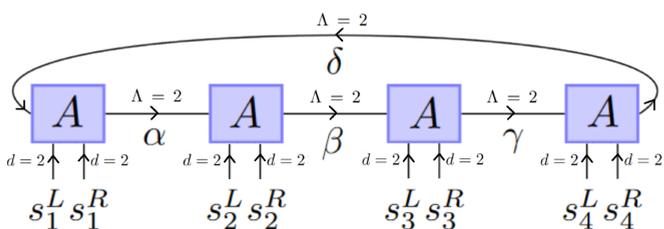


Putting all pieces together, the sequential application of the N MPDs $\{D_{[i]}\}_{i=1}^N$ reverses the black-box initialization of the MPS [cf. Fig. 8(b)]. Hence, the quantum circuit that prepares such MPS is given by the inverse of the circuit within the dashed-line box in Fig. 8(b), in agreement with Fig. 2(c) in the main text.

2. Periodic boundary conditions

The adaptation of the method discussed for a MPS with open boundary conditions to the case of periodic boundary conditions appears challenging at first glance. In the former case, starting from one of the two ends of the chain (the left end, by hypothesis, since we assumed the MPS is in left-canonical form) meant that only the physical indices s_1^L and s_1^R appeared as inputs for the first MPD $D_{[1]}$, which was just as well, since, given the black-box initialization of the MPS, we had access to the physical indices but not to the virtual ones. However, for a MPS with periodic boundary conditions, regardless of the site from which we start, there will always be an incoming virtual leg that we cannot access. For example, let us consider the following translationally invariant MPS for which the local rank-4 ($2 \times 2 \times 2 \times 2$) tensor A is assumed to be left-normalized (hence the arrows defining the directions of the legs in

the diagram).



If we start from site 1 (any site will do, since they are all equivalent), the first MPD, $D_{[1]}$, will take as inputs s_1^L, s_1^R , and δ , outputting the virtual index α . We do have access to the physical indices s_1^L, s_1^R in our retrosynthetic approach [cf. Fig. 8(b)] but not to the virtual index δ .

We therefore need a different approach to the first site. Specifically, we will cast the local tensor $A^{\delta, s_1^L, s_1^R, \alpha}$ at site 1 into a 4×4 matrix $A^{(s_1^L, s_1^R)}_{(\alpha, \delta)} \equiv \tilde{A}$ by fusing the two physical indices (s_1^L, s_1^R) into a single index labeling the rows and the two virtual indices (α, δ) into a single index labeling the columns. Notice that, in this case, the dimensions of the inputs and outputs coincide, so there seems to be no need to add dummy legs to construct the MPD $D_{[1]}$. The problem, however, is that \tilde{A} is, in general, nonunitary, because we are fusing one incoming index, δ , with an outgoing index, α . Hence, although the MPD defined as $D_{[1]} = \tilde{A}^\dagger$ does produce the desired effect of turning the two physical indices from site 1, s_1^L , and s_1^R , into the two virtual indices α and δ , implementing such 4×4 nonunitary matrix on quantum hardware is not immediately obvious.

To address this problem, we will embed this 4×4 nonunitary matrix \tilde{A} in an 8×8 unitary matrix $U_{\tilde{A}}$. \tilde{A} will correspond to the top-left quadrant of $U_{\tilde{A}}$, in which case the most significant qubit on which the three-qubit operation $U_{\tilde{A}}$ acts is an ancilla prepared in $|0\rangle$ and projected onto $|0\rangle$ by postselection, thus ensuring that the effective action of $U_{\tilde{A}}$ on the remaining two qubits is the desired \tilde{A} . Then, the three-qubit $D'_{[1]}$ that appears in Fig. 2(d) in the main text is simply given by $D'_{[1]} = U_{\tilde{A}}^\dagger$. The probability of success is, in general, $\mathcal{O}(1)$; for the spin-1 VBS state, 50% of trials are successful regardless of N .

A general scheme to accomplish this embedding of a N -qubit nonunitary operation in a $(N + 1)$ -qubit gate was presented by Lin *et al.* [78]. Applying it to the specific case considered here, the goal is to find the 8×8 unitary matrix

$$U_{\tilde{A}} = \begin{pmatrix} n\tilde{A} & \\ & B \end{pmatrix}, \quad (C4)$$

where C is a 4×4 matrix, B is a 8×4 matrix, and n is a scalar that ensures that no column of \tilde{A} has norm greater than 1, otherwise it is impossible to ensure that all columns of $U_{\tilde{A}}$ are orthonormal, which is a defining property of unitary matrices. C can be derived from \tilde{A} by performing its singular value decomposition, $\tilde{A} = USV^\dagger$, so that

$$C = U(\mathbb{1} - n^2S)^{\frac{1}{2}}V^\dagger, \quad (C5)$$

where n is any scalar that satisfies the requirement

$$0 < n < 1/\sqrt{\max(\{s_i^2\}_{i=1}^4)}, \quad (C6)$$

with $\{s_i\}_{i=1}^4$ the singular values of \tilde{A} [i.e., $S = \text{diag}(s_1, s_2, s_3, s_4)$]. Having derived n and C from \tilde{A} , we already have the first four columns of $U_{\tilde{A}}$, which are orthonormal. We just need to find the remaining four columns (i.e., the matrix B), which can be accomplished similarly to the way we completed the enlarged tensors for the MPS with open boundary conditions. That is, the four columns of B are the orthonormal basis that spans the kernel of the 4×8 matrix

$$\begin{pmatrix} n\tilde{A} \\ C \end{pmatrix}^\dagger. \quad (C7)$$

For sites $i = 2, 3, \dots, N - 1$, we can proceed exactly as for the bulk sites in the MPS with open boundary conditions. In this case, however, because of the translational invariance, all MPDs $D_{[2]}, D_{[3]}, \dots, D_{[N-1]}$ are equal. For the particular case of the spin-1 VBS state, using the definition of the local tensor A stated in Eq. (11) in the main text and following the steps described in the previous section of this Appendix, we obtain

$$D_{[i]}^\dagger = \begin{pmatrix} 0 & \sqrt{\frac{2}{3}} & 0 & 0 & 0 & 0 & 0 & -\frac{1}{\sqrt{3}} \\ -\frac{1}{\sqrt{6}} & 0 & 0 & a & a & c & 0 & 0 \\ -\frac{1}{\sqrt{6}} & 0 & 0 & -\frac{1}{\sqrt{12}} & -\frac{1}{\sqrt{12}} & d & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{6}} & 0 & \frac{1}{2} & -\frac{1}{2} & 0 & 0 & \frac{1}{\sqrt{3}} \\ 0 & \frac{1}{\sqrt{6}} & 0 & -\frac{1}{2} & \frac{1}{2} & 0 & 0 & \frac{1}{\sqrt{3}} \\ -\sqrt{\frac{2}{3}} & 0 & 0 & b & b & e & 0 & 0 \end{pmatrix}, \quad (C8)$$

with $a = \frac{2\sqrt{2}}{3} + \frac{\sqrt{3}}{30}$, $b = -\sqrt{\frac{5}{12} - a^2}$, $c = -(1 + f^2 + \frac{1}{4}(1 + f)^2)^{-\frac{1}{2}}$, where $f = \frac{a - \frac{1}{2}b}{\frac{1}{\sqrt{12}} + \frac{1}{2}b}$, $d = fc$, and $e = -\frac{c+d}{2}$. The circuit for this three-qubit operation can be obtained via the Cirq `three_qubit_matrix_to_operations` method [77], resulting in a count of 20 CNOT gates.

For the final site, the rank-4 tensor $A^{\gamma, s_{N=4}^L, s_{N=4}^R, \delta}$ is turned into a 16-dimensional vector by fusing all four indices into a single one. This corresponds to the first column of the enlarged $A'_{N=4}$, and the remaining columns have to be found as in the case of open boundary conditions. There is, however, one important difference: Because we fused three incoming legs ($\gamma, s_{N=4}^L, s_{N=4}^R$) with an outgoing leg (δ), the vectorized form of $A_{N=4}$ is not guaranteed to be normalized, so we have to normalize it explicitly. Of course, due to the extra bond linking sites 1 and N , $D_{[N=4]}$ is a four-qubit gate as opposed to the three-qubit gate found for a chain.

In any case, in the initialization of the MPS, $D_{[N=4]}^\dagger$ is the first operation, so only its action on the fiducial state $|0\rangle^{\otimes 4}$

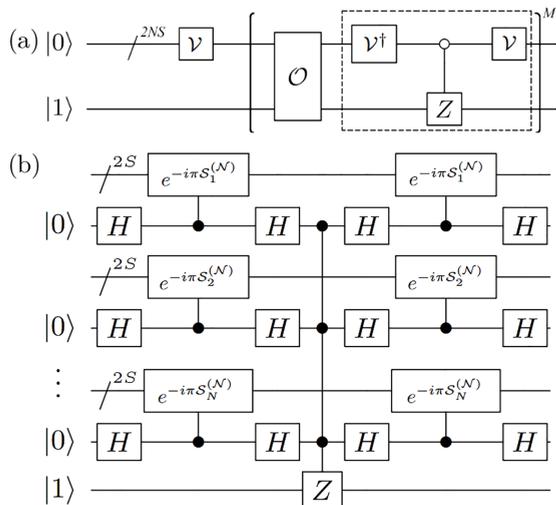


FIG. 9. (a) High-level circuit for quantum amplitude amplification algorithm. (b) Quantum circuit for oracle \mathcal{O} that is suitable for the application of quantum amplitude amplification to prepare spin- S VBS state $|\psi_{\text{VBS}}\rangle$ defined on N -site lattice starting from $|\psi_{\text{pre-VBS}}\rangle = \mathcal{V}|0\rangle^{\otimes 2NS}$, where \mathcal{V} is defined in Fig. 3(a). As in the Hadamard-test-based method described in Fig. 3(c), all N site ancillas are initialized in state $|0\rangle$. The extra ancilla required to apply the phase shift per the phase kickback effect is initialized in state $|1\rangle$. All remaining $2NS$ qubits correspond to the main register where $|\psi_{\text{VBS}}\rangle$ is ultimately prepared.

matters. As a result, instead of having to find the basis gate decomposition of a four-qubit gate $D_{[N=4]}^\dagger$, we only have to find the circuit that prepares the (normalized) vectorized form of $A_{N=4}$, which is a four-qubit state. For example, for the spin-1 VBS state, the following state must be initialized:

$$\frac{\sqrt{3}}{6}(0, 2, -1, 0, -1, 0, 0, 0, 0, 0, 0, 0, 1, 0, -2, 0)^T.$$

Using the Schmidt-decomposition-based method discussed in Sec. IX of the main text, this can be accomplished with a total of at most nine CNOT gates.

APPENDIX D: QUANTUM AMPLITUDE AMPLIFICATION APPLIED TO PREPARATION OF VBS STATES

Figure 9(a) shows a schematic representation of the circuit corresponding to the quantum amplitude amplification for the preparation of VBS states. The subcircuit \mathcal{V} prepares the initial state, $|\psi_{\text{pre-VBS}}\rangle$, and its structure can be found in Fig. 3 in the main text. The only missing element is the oracle that marks the target state, $|\psi_{\text{VBS}}\rangle = \bigotimes_n S_n^{(N=2S)}|\psi_{\text{pre-VBS}}\rangle$, with a phase shift $e^{i\pi}$. An implementation of such oracle that results in a circuit depth of $\mathcal{O}(N)$ is presented in Fig. 9(b). The first part corresponds to the layer of local Hadamard tests, of constant depth relative to N , used in the probabilistic method discussed in the main text. Then, the N ancillas are used as control-qubits of a controlled- Z gate acting on an additional ancilla initialized in $|1\rangle$. Such a gate can be implemented with $\mathcal{O}(N)$ depth (cf. Lemma 7.11 in Ref. [169]). If all ancillas are in $|1\rangle$, the main $2NS$ -qubit register is in the fully symmetrized state $|\psi_{\text{VBS}}\rangle$, and the controlled- Z gate is triggered, applying a phase shift $e^{i\pi}$ to the main register per the phase kickback

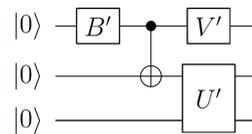


FIG. 10. Adaptation of Schmidt-decomposition-based method to preparation of arbitrary three-qubit state $|\phi\rangle = \sum_{i=0}^3 \sum_{j=0}^1 \mathbf{M}'_{i,j}|i\rangle \otimes |j\rangle$. U' and V' can be obtained from the singular value decomposition of $\mathbf{M}' = U'S'V'$. Single-qubit operation B' prepares state $B'|0\rangle = \sum_{l=0}^1 s'_l|l\rangle$, where $\{s'_0, s'_1\}$ are the singular values of \mathbf{M}' .

effect [170]. The final layer of local Hadamard tests resets the ancillas, so that they can be reused in the following iteration.

Despite the $\mathcal{O}(N)$ depth of the circuit that realizes the oracle \mathcal{O} , the number of iterations M corresponding to the repetition of the subcircuit shown inside brackets in Fig. 9(a) is exponential in N . Indeed, quantum amplitude amplification amounts to a rotation in a two-dimensional subspace spanned by $\{|\tilde{\psi}_{\text{VBS}}\rangle, |\psi_{\text{VBS}}^\perp\rangle\}$, where $|\psi_{\text{VBS}}^\perp\rangle \equiv |\psi_{\text{pre-VBS}}\rangle - |\tilde{\psi}_{\text{VBS}}\rangle$. Since $|\langle \tilde{\psi}_{\text{VBS}} | \psi_{\text{pre-VBS}} \rangle|^2 = p^N$ [cf. Eq. (8)], for $N \gg 1$, $p^N \ll 1$, so the angle covered in each iteration is $\sim 2p^{N/2}$ [170], in which case maximizing the overlap with the target state $|\psi_{\text{VBS}}\rangle$ takes $\mathcal{O}(p^{-N/2})$ iterations. Since all iterations have to be implemented consecutively, this results in a total circuit depth that is far beyond the capabilities of near-term quantum hardware even for small N .

APPENDIX E: DETERMINISTIC PREPARATION OF GENERAL 3-QUBIT STATES VIA SCHMIDT DECOMPOSITION

As mentioned in Sec. IX of the main text, the application of the Schmidt-decomposition-based method to the initialization of the six-qubit islands $|\psi_{\text{island}}^{S=3/2}\rangle = \sum_{i=0}^7 \sum_{j=0}^7 \mathbf{M}_{i,j}|i\rangle \otimes |j\rangle$ of the spin-3/2 VBS state involves the preparation of a three-qubit state $\sum_{k=0}^7 s_k|k\rangle$, where $\{s_k\}_{k=0}^7$ are the singular values of \mathbf{M} . Even though the number of qubits is not even, the Schmidt-decomposition-based method can be adapted to this three-qubit case. Specifically, we can consider an asymmetric bipartition with two qubits on one side and one qubit on the other: $\sum_{k=0}^7 s_k|k\rangle = \sum_{i=0}^3 \sum_{j=0}^1 \mathbf{M}'_{i,j}|i\rangle \otimes |j\rangle$. Performing the singular value decomposition of the 4×2 matrix \mathbf{M}' whilst retaining the redundant all-zero rows of \mathbf{S} (thus ensuring \mathbf{U} is unitary and not just left-normalized) gives $\mathbf{M}' = \mathbf{U}'\mathbf{S}'\mathbf{V}'^\dagger$, with

$$\mathbf{U}' = \begin{pmatrix} \uparrow & \uparrow & \uparrow & \uparrow \\ |u'_0\rangle & |u'_1\rangle & |u'_2\rangle & |u'_3\rangle \\ \downarrow & \downarrow & \downarrow & \downarrow \end{pmatrix},$$

$$\mathbf{S}' = \begin{pmatrix} s'_0 & 0 \\ 0 & s'_1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathbf{V}' = \begin{pmatrix} \uparrow & \uparrow \\ |v'_0\rangle & |v'_1\rangle \\ \downarrow & \downarrow \end{pmatrix}. \quad (\text{E1})$$

Hence, $\sum_{k=0}^7 s_k|k\rangle = \sum_{l=0}^1 s'_l|u'_l\rangle \otimes |v'_l\rangle$, which can be prepared via the quantum circuit shown in Fig. 10 below. B' and V' are both single-qubit operations, so their decomposition involves no CNOT gates. U' is a two-qubit operation, so it can

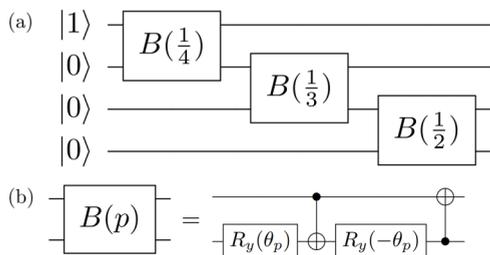


FIG. 11. (a) Scheme of quantum circuit that prepares $|W_4\rangle$ state. For the general case $|W_m\rangle$, a sequence of two-qubit blocks $B(\frac{1}{m}), B(\frac{1}{m-1}), \dots, B(\frac{1}{3}),$ and $B(\frac{1}{2})$ is applied in the staircaselike structure of (a), with the first qubit initialized in $|1\rangle$ and all remaining $m - 1$ qubits in $|0\rangle$. (b) Basis gate decomposition of two-qubit block $B(p)$, with $0 < p < 1$. $\theta_p = \arcsin(\cos(\arctan(\sqrt{\frac{1-p}{p}})))$.

be decomposed in terms of, at most, three CNOT gates [110], yielding a maximum of four CNOT gates to prepare an arbitrary three-qubit state via this method.

APPENDIX F: PREPARATION OF $|W_m\rangle$ STATES

The $|W_m\rangle$ state is a m -qubit state defined as

$$\begin{aligned} |W_m\rangle &:= \sum_{i=0}^{m-1} \frac{|2^i\rangle}{\sqrt{m}} \\ &= \left(\frac{|10\dots 0\rangle}{\sqrt{m}} + \frac{|01\dots 0\rangle}{\sqrt{m}} + \dots + \frac{|00\dots 1\rangle}{\sqrt{m}} \right). \end{aligned}$$

Figure 11(a) shows the circuit that prepares such W state for $m = 4$, as first presented in [116]. The key building block of this circuit is the parameterized two-qubit subcircuit $B(p)$, with $p \in (0, 1)$, shown in Fig. 11(b).

The generalization of this preparation scheme of $|W_m\rangle$ for arbitrary $m \in \mathbb{N}$ is self-evident. It consists of applying the sequence of two-qubit blocks $B(\frac{1}{m}), B(\frac{1}{m-1}), \dots, B(\frac{1}{3}),$ and $B(\frac{1}{2})$ in the staircase-like structure of Fig. 11(a), with the first qubit initialized in $|1\rangle$ and all remaining $m - 1$ qubits in $|0\rangle$. Hence, a total of $m - 1$ executions of the $B(p)$ block are required to prepare $|W_m\rangle$. However, as alluded to in Sec. X of the main text, the execution of such $m - 1$ $B(p)$ blocks need not be sequential; it is possible to exploit the freedom to choose the target qubit of the final CNOT in the $B(p)$ building block to execute multiple such blocks in parallel, thus yielding a $\mathcal{O}(\ln m)$ -depth circuit, as opposed to the $\mathcal{O}(m)$ depth of the sequential approach. In any case, the total number of blocks is the same for both cases: $m - 1$ $B(p)$.

Let us consider the case $m = 4! = 24$, which is relevant for the construction of the Prepare oracle \mathcal{O}_P for the LCU implementation of the symmetrization operator acting on $\mathcal{N} = 4$ spins- $\frac{1}{2}$, $\mathcal{S}^{(\mathcal{N}=4)}$. $24 - 1 = 23$ $B(p)$ blocks are involved in the preparation of $|W_{24}\rangle$. Each $B(p)$ takes two CNOT gates, so the implementation of the prepare oracle \mathcal{O}_P for $\mathcal{S}^{(\mathcal{N}=4)}$ takes $23 \times 2 = 46$ CNOT gates.

APPENDIX G: CIRCUIT DEPTH IN PREPARATION OF SPIN-1 AND SPIN-3/2 VALENCE-BOND-SOLID STATES

In this Appendix, the depth of the circuits that prepare the spin-1 and spin-3/2 VBS states is determined, taking into account the qubit connectivity constraints of the NISQ device. The $S = 1$ and $S = 3/2$ cases will be examined separately. All-to-all qubit connectivity will be assumed first, and then the qubit connectivity arrangement obtained by embedding these circuits on the IBM Quantum heavy-hex lattice [130], as shown in Fig. 7 in Sec. XI, will be considered. The results herein discussed are summarized in Table I on Sec. XI.

1. Spin-1 valence-bond-solid state

Regarding the spin-1 VBS state defined on a one-dimensional lattice, assuming all-to-all qubit connectivity for the moment, the probabilistic method without any mitigation of the repetition overhead results in a circuit of depth eight CNOT gates, 7 of which are associated with the local Hadamard test (cf. decomposition of Fredkin gate in Ref. [94]) and 1 with the initialization of the product state of valence bonds [cf. Fig. 3(a)]. The average number of repetitions in such case is $(\frac{4}{3})^N$ (ignoring the finite-size effects discussed in Appendix B).

A quadratic reduction of this repetition overhead can be achieved by preparing the four-qubit islands deterministically, as discussed in Sec. IX, before applying the probabilistic method to the remaining sublattice. This merely increases the circuit depth from 8 to 11 CNOT gates, as the repetition overhead mitigation layer takes four CNOT gates of depth, three more than the initialization of the product state of valence bonds.

Let us now assume the most restrictive case of linear qubit connectivity. If the probabilistic method alone is implemented, then all N ancillary qubits are required, one for each site. These ancillas are placed between the two qubits that encode the spin-1 of the respective site, so that the two qubits forming a valence bond are connected to each other, thus allowing all valence bonds to be initialized with depth one CNOT. Then, in the implementation of the local Hadamard test, given that the ancilla is between the two site qubits, we make use of the circuit for the CSWAP in [94] that assumes linear qubit connectivity and the placement of the control-qubit in the central position of the three-qubit register. This takes nine CNOT gates, ignoring the final network of three CNOT gates that swap the ancilla with one of the site qubits, which is immaterial, since the ancilla is measured immediately afterwards. Having discarded this SWAP gate, the qubit to be measured must be changed accordingly; this is already accounted for in the respective QASM file (cf. Appendix H). In total, the depth is of $1 + 9 = 10$ CNOT gates.

Alternatively, the $\mathcal{O}(1)$ -depth repetition overhead mitigation layer can be used to reduce the average number of repetitions. The QISKIT transpile function was applied to the original four-qubit island preparation circuit, imposing the linear connectivity constraint via the coupling_map input. After a few trials, a circuit with a depth of eight CNOT gates was generated. The layer of local Hadamard tests at the sites of the other sublattice must then be added, taking the same

nine CNOT gates stated in the previous paragraph. The only difference relative to the bare probabilistic method is that this scheme requires only $N/2$ ancillas, placed between consecutive four-qubit islands, as illustrated in Fig. 7(c) in Sec. XI. In total, with linear qubit connectivity, the probabilistic method preceded by the repetition overhead mitigation layer takes $8 + 9 = 17$ CNOT gates of depth.

2. Spin-3/2 valence-bond-solid state

Ignoring qubit connectivity constraints, applying the probabilistic method without mitigating the repetition overhead to prepare the spin-3/2 VBS state defined on a honeycomb lattice takes 27 CNOT gates of depth, 26 of which due to the local Hadamard test at every site (cf. Sec. VII) and the remaining CNOT due to the initialization of the product state of valence bonds $|\psi_{\text{pre-VBS}}\rangle$.

If, instead, one initializes the six-qubit islands to symmetrize all sites of one sublattice deterministically before applying the probabilistic method to the other sublattice, the average repetition overhead is decreased from 2^N to $2^{\frac{N}{2}}$ at the cost of increasing the circuit depth. Concretely, the initialization of the six-qubit islands takes 19 CNOT gates of depth, which adds to the 26 CNOT gates from the local-Hadamard-test layer to yield a total circuit depth of $19 + 26 = 45$ CNOT gates.

Let us now take realistic qubit connectivity constraints into account by considering the preparation of the spin-3/2 VBS state on the IBM Quantum heavy-hex lattice, as illustrated in Figs. 7(a) and 7(b). First, we will adopt the local Hadamard test at all lattice sites, which corresponds to the case depicted in Fig. 7(a). Then, we will assume the six-qubit islands are initialized deterministically before the probabilistic method is employed at half of the sites at which the wave function is yet to be locally symmetrized, which is the method contemplated in Fig. 7(b).

Regarding the bare probabilistic method, the first step amounts to the initialization of all valence bonds, which are represented by the orange curved lines in Fig. 7(a). This simply involves the execution of the two-qubit subcircuit shown in Fig. 3(a) for each valence bond. All such circuits can still be executed in parallel, so the depth of this first step is of 1 CNOT gate. At this point, we have the prepared the $|\psi_{\text{pre-VBS}}\rangle$ state, with sets of three adjacent valence bonds [the six qubits of which are represented as black dots in Fig. 7(a)(i)] separated from one another in the qubit heavy-hex lattice by redundant qubits (shown as gray dots). Before we can proceed to the local symmetrization at every site, we must ensure that all three qubits that encode a local spin 3/2, plus the respective site ancilla, are next to one another. To accomplish this, two qubits from each set of three adjacent valence bonds are displaced in the qubit lattice by swapping them with three consecutive redundant qubits, as depicted by the curly arrows included in Fig. 7(a)(i). One of these displaced qubits moves to the left, while the other moves to the right. As a result of these displacements, two out of every three valence bonds are now defined between pairs of separated qubits, as illustrated in Fig. 7(a)(ii) by the stretching of the orange curved lines.

We are now in a position to apply the local Hadamard test with $\mathcal{U} = e^{-i\pi S^{(N=3)}}$ [cf. Fig. 3(b)] at the sets of four qubits

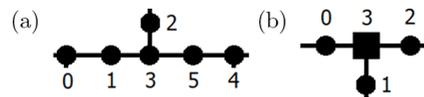


FIG. 12. Illustration of the qubit coupling maps considered in the implementation in the IBM Quantum heavy-hex lattice of the initialization of the six-qubit islands (a) and the local Hadamard test for the remaining sublattice (b) for the preparation of the spin-3/2 VBS state via the probabilistic method combined with the $\mathcal{O}(1)$ -depth repetition overhead mitigation scheme. Numbering of qubits in (a) ensures that qubit pairs (1,2), (3,4), and (5,6) form valence bonds, and site encoded by qubit trio (1,3,5) is symmetrized, in agreement with Fig. 7(b). Black square in (b) corresponds to site ancilla.

enclosed by the green boxes in Fig. 7(a)(ii). Notice that there are two types of arrangements of these sets of four qubits, one being T-shaped and the other having a linear shape. Each corresponds to a different sublattice. The ancillary qubits are identified as black squares. For the T-shaped arrangements, the ancillary qubit corresponds to the redundant qubit from Fig. 7(a)(i) between the two qubits that joined it from either side. For the linear arrangement, in turn, the site ancilla is the very first redundant qubit that is swapped with the left-moving qubit that departed from there. The qubit connectivity layout for the T-shaped set is shown in Fig. 12(b), while that of the linear set can be identified with the (0, 1, 3, 5) subset from Fig. 12(a), with the ancilla being qubit 0 in this case. Applying the QISKIT transpile [97] function to the controlled $e^{-i\pi S^{(N=3)}}$ and providing the coupling_map input according to the layout of the T-shaped and linear sets, the circuit depth increases from the 26 CNOT gates found without any qubit connectivity constraints to 39 and 41 CNOT gates, respectively. For the sake of reproducibility, the QASM files corresponding to these basis gate decompositions are provided (cf. Appendix H), since the QISKIT transpile function does not always produce the same outcome, as this is a hard optimization problem. The circuits provided are the outcomes with lowest CNOT depth and total count yielded by QISKIT transpile over many trials. To sum up, the total circuit depth corresponding to the preparation of the spin-3/2 VBS state on the IBM Quantum heavy-hex lattice via the bare probabilistic method is $1 + 9 + 41 = 51$ CNOT gates, where the first term is due to the initialization of the product state of valence bonds, the second to the 3 SWAP gates required to place all qubits in the right positions for the final step, and the third to the local Hadamard test.

Alternatively, the spin-3/2 VBS state can be prepared by first initializing the six-qubit islands deterministically and then applying the probabilistic method at only half of the sites, resulting in a quadratic reduction of the average number of repetitions. Figure 7(b) illustrates the execution of this method on the IBM Quantum heavy-hex lattice. The six-qubit islands can be prepared in the same sets of six qubits considered in the bare probabilistic method above for the initialization of the valence bonds. Considering the qubit coupling map shown in Fig. 12(a), the circuit depth is found to increase from 19 CNOT gates (assuming all-to-all connectivity) to 57 CNOT gates. At this point, not only are all valence bonds already initialized but all sites of one sublattice are symmetrized as

well. Before the probabilistic method is applied at the T-shaped sets of four qubits corresponding to the sites of the other sublattice, the same three SWAP gates with redundant qubits considered above need to be performed, even though they are not explicitly shown in Fig. 7(b). The total circuit depth of this method is therefore $57 + 9 + 39 = 105$ CNOT gates.

APPENDIX H: OpenQASM INSTRUCTION FILES

To foster the reproducibility of the methods herein presented, we have made available the following QASM files in Ref. [171] as .txt files.

- (1) S_1_Hadamard_test_all_to_all.qasm
- (2) S_1_Hadamard_test_linear.qasm
- (3) S_1_4Q_island_all_to_all.qasm
- (4) S_1_4Q_island_linear.qasm
- (5) S_3_2_Hadamard_test_all_to_all.qasm
- (6) S_3_2_Hadamard_test_heavy_hex_linear_box.qasm
- (7) S_3_2_Hadamard_test_heavy_hex_T_box.qasm
- (8) S_3_2_6Q_island_all_to_all.qasm
- (9) S_3_2_6Q_island_heavy_hex.qasm

We note that the .txt files can be resaved in .qasm format in any text editor. The use of these QASM files is described below.

1. $S = 1$ VBS state via bare probabilistic method

First, the product state of valence bonds, $|\psi_{\text{pre-VBS}}\rangle$, must be initialized by applying at every pair of qubits representing a lattice link the two-qubit subcircuit presented in Fig. 3(a), for which no QASM file is provided, given its simple structure.

Then, the local Hadamard test subcircuit must be applied at all sets of three qubits associated with one site (two encoding the local spin-1 and the remaining acting as the site ancilla). Depending on whether the qubit connectivity between these three qubits is all-to-all or linear, the three-qubit circuits corresponding to the QASM files 1 and 2 from the list above should be used, respectively.

In both cases, the whole structure of the local Hadamard test is already included, namely, the Hadamard gates acting on the ancilla before and after the controlled $e^{-i\pi S(N=2)}$, the Z gate on the ancilla before the Fredkin gate to account for the global phase factor in Eq. (18) (so that the desired outcome of the ancilla measurement is $|1\rangle$), and the actual measurement of the ancilla at the end. In QASM file 1, the ancilla is the most significant qubit, q[2]. In QASM file 2, the ancilla qubit is initially q[1], in agreement with the fact that the ancilla should be between the two site qubits to ease the initialization of the valence bonds [cf. Fig. 7(c)]. However, at the end of the circuit of QASM file 2, the ancilla becomes encoded in q[0] because a SWAP gate is skipped, as discussed in Appendix G. As a result, the measured qubit is q[0], not q[1].

2. $S = 1$ VBS state via preparation of four-qubit islands and probabilistic method at one sublattice

With respect to the bare probabilistic method described above, the initialization of the product state of valence bonds is replaced by the deterministic preparation of the four-qubit

islands [cf. Fig. 4(a)] using the circuits from the QASM files 3 and 4 for all-to-all and linear qubit connectivity, respectively.

In both cases the valence bonds are prepared at qubit pairs (q[0],q[1]) and (q[2],q[3]), and the symmetrized site is encoded by the qubit pair (q[1],q[2]), in agreement with the discussion from Sec. IX. A schematic diagram of the circuit corresponding to QASM file 3 is shown in Fig. 5(b), and the details of the single-qubit operations can be inferred from the QASM file.

Then, the local Hadamard test must be applied at the remaining half of the sites using the circuits from QASM files 1 or 2, as described in the previous case.

3. $S = 3/2$ VBS state via bare probabilistic method

First, all valence bonds are prepared in parallel by applying the two-qubit subcircuit in Fig. 3(a) at every qubit pair representing a lattice link, just like in the spin-1 case. Then, if the IBM Q heavy-hex lattice [130] is considered, the three consecutive SWAP gates represented by the curly arrows in Fig. 7(a)(i) must be applied to place all three qubits encoding each local spin-3/2 next to one another.

The local Hadamard test is then applied to these three qubits, plus the respective ancilla, represented by the black squares in Fig. 7(a)(ii). In the heavy-hex lattice, there are two different circuits for this four-qubit operation, one for each sublattice. QASM file 6 applies to the sublattice associated with the linear-shaped green boxes in Fig. 7(a)(ii), and QASM file 7 to the other sublattice corresponding to the T-shaped green boxes [cf. numbered coupling scheme in Fig. 12(b)].

If instead a quantum processor with all-to-all connectivity is used, the circuit from QASM file 5 should be used to implement the local Hadamard test. In QASM files 5, 6, and 7, the site ancilla is always the most significant qubit, q[3]. As in the spin-1 case, all elements of the Hadamard test are already included, including the measurement of the ancilla.

4. $S = 3/2$ VBS state via preparation of six-qubit islands and probabilistic method at one sublattice

The starting point corresponds to the deterministic initialization of the 6-qubit islands [cf. Fig. 4(b)]. For a quantum computer with all-to-all connectivity between the six qubits, the circuit to be used is provided in QASM file 8 from the list above. If, instead, the IBM Quantum heavy-hex [130] architecture is used, the circuit to prepare the 6-qubit islands is given by QASM file 9, which adopts the coupling map shown in Fig. 12(a). In both cases, the valence bonds are associated with the qubit pairs (q[0],q[1]), (q[2],q[3]), and (q[4],q[5]). The site at the center of this island that is already symmetrized in this step is encoded by (q[1],q[3],q[5]).

The remaining part corresponds to applying the local symmetrization operator at the sites of the other sublattice. For the case of all-to-all connectivity, the corresponding local Hadamard test is implemented via the 4-qubit subcircuit given in QASM file 5. In the case of the IBM Quantum heavy-hex architecture [130], first, the sequence of three SWAP gates

must be applied to place all three qubits associated with each site of this sublattice next to each other, and then the

local Hadamard test is executed through the circuit in QASM file 7.

-
- [1] R. B. Laughlin and D. Pines, The Theory of Everything, *Proc. Natl. Acad. Sci. USA* **97**, 28 (2000).
- [2] P. W. Anderson, More Is Different, *Science* **177**, 393 (1972).
- [3] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, UK, 2010)
- [4] M. Troyer and U.-J. Wiese, Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
- [5] W. Kohn, Nobel Lecture: Electronic structure of matter—wave functions and density functionals, *Rev. Mod. Phys.* **71**, 1253 (1999).
- [6] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Rev. Mod. Phys.* **82**, 277 (2010).
- [7] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Lafamme, Quantum algorithms for fermionic simulations, *Phys. Rev. A* **64**, 022319 (2001).
- [8] J. Preskill, Quantum computing and the entanglement frontier, [arXiv:1203.5813](https://arxiv.org/abs/1203.5813) (2012).
- [9] J. Stolze and D. Suter, *Quantum Computing: A Short Course from Theory to Experiment* (John Wiley & Sons, Hoboken, NJ, USA, 2004).
- [10] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [11] P. M. Q. Cruz and J. Fernández-Rossier, Testing complementarity on a transmon quantum processor, *Phys. Rev. A* **104**, 032223 (2021).
- [12] D. A. Lidar and T. A. Brun (Eds.), *Quantum Error Correction* (Cambridge University Press, Cambridge, UK, 2013).
- [13] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [14] R. Jozsa and N. Linden, On the role of entanglement in quantum-computational speed-up, *Proc. R. Soc. London A* **459**, 2011 (2003).
- [15] F. Verstraete, J. I. Cirac, and J. I. Latorre, Quantum circuits for strongly correlated quantum systems, *Phys. Rev. A* **79**, 032316 (2009).
- [16] P. Scholl and R. Orús, Kitaev honeycomb tensor networks: Exact unitary circuits and applications, *Phys. Rev. B* **95**, 045112 (2017).
- [17] A. Cervera-Lierta, Exact Ising model simulation on a quantum computer, *Quantum* **2**, 114 (2018).
- [18] X. Xiao, J. K. Freericks, and A. F. Kemper, Determining quantum phase diagrams of topological Kitaev-inspired models on NISQ quantum hardware, *Quantum* **5**, 553 (2021).
- [19] M. J. Cervia, A. B. Balantekin, S. N. Coppersmith, C. W. Johnson, P. J. Love, C. Poole, K. Robbins, and M. Saffman, Lipkin model on a quantum computer, *Phys. Rev. C* **104**, 024305 (2021).
- [20] J. S. Van Dyke, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, Preparing Bethe Ansatz Eigenstates on a Quantum Computer, *PRX Quantum* **2**, 040329 (2021).
- [21] J. S. Van Dyke, E. Barnes, S. E. Economou, and R. I. Nepomechie, Preparing exact eigenstates of the open XXZ chain on a quantum computer, *J. Phys. A: Math. Theor.* **55**, 055301 (2022).
- [22] T. A. Bespalova and O. Kyriienko, Quantum simulation and ground state preparation for the honeycomb Kitaev model, [arXiv:2109.13883](https://arxiv.org/abs/2109.13883) (2021).
- [23] R. Orbach, Linear antiferromagnetic chain with anisotropic coupling, *Phys. Rev.* **112**, 309 (1958).
- [24] C. Babenko, F. Göhmann, K. K. Kozłowski, J. Sirker, and J. Suzuki, Exact Real-Time Longitudinal Correlation Functions of the Massive XXZ Chain, *Phys. Rev. Lett.* **126**, 210602 (2021).
- [25] A. Kitaev, Anyons in an exactly solved model and beyond, *Ann. Phys.* **321**, 2 (2006).
- [26] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Rigorous Results on valence-Bond Ground States in Antiferromagnets, *Phys. Rev. Lett.* **59**, 799 (1987).
- [27] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Valence bond ground states in isotropic quantum antiferromagnets, *Commun. Math. Phys.* **115**, 477 (1988).
- [28] F. D. M. Haldane, Nonlinear Field Theory of Large-Spin Heisenberg Antiferromagnets: Semiclassically Quantized Solitons of the One-Dimensional Easy-Axis Néel State, *Phys. Rev. Lett.* **50**, 1153 (1983).
- [29] T.-C. Wei, I. Affleck, and R. Raussendorf, Affleck-Kennedy-Lieb-Tasaki State on a Honeycomb Lattice is a Universal Quantum Computational Resource, *Phys. Rev. Lett.* **106**, 070501 (2011).
- [30] A. Miyake, Quantum computational capability of a 2D valence bond solid phase, *Ann. Phys.* **326**, 1656 (2011).
- [31] R. Raussendorf and H. J. Briegel, A One-Way Quantum Computer, *Phys. Rev. Lett.* **86**, 5188 (2001).
- [32] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, *Phys. Rev. A* **68**, 022312 (2003).
- [33] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest, Measurement-based quantum computation, *Nat. Phys.* **5**, 19 (2009).
- [34] D. Gross, S. T. Flammia, and J. Eisert, Most Quantum States are Too Entangled to be Useful as Computational Resources, *Phys. Rev. Lett.* **102**, 190501 (2009).
- [35] M. J. Bremner, C. Mora, and A. Winter, Are Random Pure States Useful for Quantum Computation? *Phys. Rev. Lett.* **102**, 190502 (2009).
- [36] M. A. Nielsen, Cluster-state quantum computation, *Rep. Math. Phys.* **57**, 147 (2006).
- [37] M. Koch-Janusz, D. I. Khomskii, and E. Sela, Affleck-Kennedy-Lieb-Tasaki State on a Honeycomb Lattice from t_{2g} Orbitals, *Phys. Rev. Lett.* **114**, 247204 (2015).
- [38] S. Mishra, G. Catarina, F. Wu, R. Ortiz, D. Jacob, K. Eimre, J. Ma, C. A. Pignedoli, X. Feng, P. Ruffieux, J. Fernández-Rossier, and R. Fasel, Observation of fractional edge excitations in nanographene spin chains, *Nature* **598**, 287 (2021).
- [39] R. Ganesh, D. N. Sheng, Y.-J. Kim, and A. Paramekanti, Quantum paramagnetic ground states on the honeycomb lattice and field-induced Néel order, *Phys. Rev. B* **83**, 144414 (2011).

- [40] N. Pomata and T.-C. Wei, Demonstrating the Affleck-Kennedy-Lieb-Tasaki Spectral Gap on 2D Degree-3 Lattices, *Phys. Rev. Lett.* **124**, 177203 (2020).
- [41] M. Lemm, A. W. Sandvik, and L. Wang, Existence of a Spectral Gap in the Affleck-Kennedy-Lieb-Tasaki Model on the Hexagonal Lattice, *Phys. Rev. Lett.* **124**, 177204 (2020).
- [42] 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).
- [43] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states, *Phys. Rev. A* **95**, 042308 (2017).
- [44] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, Variational quantum algorithms for discovering Hamiltonian spectra, *Phys. Rev. A* **99**, 062304 (2019).
- [45] O. Higgott, D. Wang, and S. Brierley, Variational quantum computation of excited states, *Quantum* **3**, 156 (2019).
- [46] K. M. Nakanishi, K. Mitarai, and K. Fujii, Subspace-search variational quantum eigensolver for excited states, *Phys. Rev. Res.* **1**, 033062 (2019).
- [47] T. Chen, R. Shen, C. H. Lee, and B. Yang, High-fidelity realization of the AKLT state on a NISQ-era quantum processor, [arXiv:2210.13840](https://arxiv.org/abs/2210.13840) v1(2022).
- [48] K. C. Smith, E. Crane, N. Wiebe, and S. M. Girvin, Deterministic constant-depth preparation of the AKLT state on a quantum processor using fusion measurements, [arXiv:2210.17548](https://arxiv.org/abs/2210.17548) v1 (2022).
- [49] Z.-Y. Wei, D. Malz, and J. I. Cirac, Efficient Adiabatic Preparation of Tensor Network States, [arXiv:2209.01230](https://arxiv.org/abs/2209.01230) v1 (2022).
- [50] Y. Ge, A. Molnár, and J. I. Cirac, Rapid Adiabatic Preparation of Injective Projected Entangled Pair States and Gibbs States, *Phys. Rev. Lett.* **116**, 080503 (2016).
- [51] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Quantum algorithms revisited, *Proc. R. Soc. A* **454**, 339 (1998).
- [52] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [53] R. Orús, Tensor networks for complex quantum systems, *Nat. Rev. Phys.* **1**, 538 (2019).
- [54] 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).
- [55] H. Tasaki, Affleck-Kennedy-Lieb-Tasaki Model, in *Physics and Mathematics of Quantum Many-Body Systems* (Springer, 2020), pp. 177-224.
- [56] A. Yu. Kitaev, Quantum measurements and the Abelian Stabilizer Problem, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026) (1995).
- [57] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [58] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, Variational ansatz-based quantum simulation of imaginary time evolution, *npj Quantum Inf.* **5**, 75 (2019).
- [59] H. Wang, Quantum algorithm for preparing the ground state of a system via resonance transition, *Sci. Rep.* **7**, 16342 (2017).
- [60] M. Schwarz, K. Temme, and F. Verstraete, Preparing Projected Entangled Pair States on a Quantum Computer, *Phys. Rev. Lett.* **108**, 110502 (2012).
- [61] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).
- [62] A. Läuchli, G. Schmid, and S. Trebst, Spin nematics correlations in bilinear-biquadratic $S = 1$ spin chains, *Phys. Rev. B* **74**, 144426 (2006).
- [63] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, *Commun. Math. Phys.* **144**, 443 (1992).
- [64] A. Klümper, A. Schadschneider, and J. Zittartz, Matrix product ground states for one-dimensional spin-1 quantum antiferromagnets, *Europhys. Lett.* **24**, 293 (1993).
- [65] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [66] S. Rommer and S. Östlund, Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group, *Phys. Rev. B* **55**, 2164 (1997).
- [67] U. Schollwöck, The density-matrix renormalization group, *Rev. Mod. Phys.* **77**, 259 (2005).
- [68] D. P. Arovas, Two exact excited states for the $S = 1$ AKLT chain, *Phys. Lett. A* **137**, 431 (1989).
- [69] S. Moudgalya, S. Rachel, B. A. Bernevig, and N. Regnault, Exact excited states of nonintegrable models, *Phys. Rev. B* **98**, 235155 (2018).
- [70] S. Moudgalya, N. Regnault, and B. A. Bernevig, Entanglement of exact excited states of Affleck-Kennedy-Lieb-Tasaki models: Exact results, many-body scars, and violation of the strong eigenstate thermalization hypothesis, *Phys. Rev. B* **98**, 235156 (2018).
- [71] D. Gross and J. Eisert, Novel Schemes for Measurement-Based Quantum Computation, *Phys. Rev. Lett.* **98**, 220503 (2007).
- [72] G. K. Brennen and A. Miyake, Measurement-Based Quantum Computer in the Gapped Ground State of a Two-Body Hamiltonian, *Phys. Rev. Lett.* **101**, 010502 (2008).
- [73] J. I. Cirac, D. Perez-Garcia, N. Schuch, and F. Verstraete, Matrix product unitaries: Structure, symmetries, and topological invariants, *J. Stat. Mech.* (2017) 083105.
- [74] M. B. Şahinoğlu, S. K. Shukla, F. Bi, and X. Chen, Matrix product representation of locality preserving unitaries, *Phys. Rev. B* **98**, 245122 (2018).
- [75] S.-J. Ran, Encoding of matrix product states into quantum circuits of one- and two-qubit gates, *Phys. Rev. A* **101**, 032310 (2020).
- [76] V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of quantum-logic circuits, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **25**, 1000 (2006).
- [77] Cirq Documentation: three_qubit_matrix_to_operations, (Google Quantum AI), last accessed on 22-02-2023, https://quantumai.google/reference/python/cirq/optimizers/three_qubit_matrix_to_operations.
- [78] S.-H. Lin, R. Dilip, A. G. Green, A. Smith, and F. Pollmann, Real- and imaginary-time evolution with compressed quantum circuits, *PRX Quantum* **2**, 010342 (2021).
- [79] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Criticality, the Area Law, and the Computational Power of

- Projected Entangled Pair States, *Phys. Rev. Lett.* **96**, 220601 (2006).
- [80] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational Complexity of Projected Entangled Pair States, *Phys. Rev. Lett.* **98**, 140506 (2007).
- [81] D. Perez-Garcia, F. Verstraete, J. I. Cirac, and M. M. Wolf, PEPS as unique ground states of local Hamiltonians, *Quantum Inf. Comput.* **8**, 650 (2008).
- [82] C. Jordan, Essai sur la géométrie à n dimensions, *Bull. Soc. Math. France* **3**, 103 (1875).
- [83] C. Marriott and J. Watrous, Quantum Arthur–Merlin games, *Comput. Complex.* **14**, 122 (2005).
- [84] D. Nagaj, P. Wocjan, and Y. Zhang, Fast amplification of QMA, *Quantum Info. Comput.* **9**, 1053 (2009).
- [85] A. Garcia-Saez, V. Murg, and T.-C. Wei, Spectral gaps of Affleck–Kennedy–Lieb–Tasaki Hamiltonians using tensor network methods, *Phys. Rev. B* **88**, 245118 (2013).
- [86] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [87] E. Cruz, F. Baccari, J. Tura, N. Schuch, and J. I. Cirac, Preparation and verification of tensor network states, *Phys. Rev. Res.* **4**, 023161 (2022).
- [88] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combination of unitary operations, *Quantum Inf. Comput.* **12**, 901 (2012).
- [89] D. Lacroix, Symmetry-Assisted Preparation of Entangled Many-Body States on a Quantum Computer, *Phys. Rev. Lett.* **125**, 230502 (2020).
- [90] J. Abhijith *et al.*, Quantum Algorithm Implementations for Beginners, *ACM Transactions on Quantum Computing*, **3**, 1 (2018).
- [91] K. Mitarai and K. Fujii, Methodology for replacing indirect measurements with direct measurements, *Phys. Rev. Res.* **1**, 013006 (2019).
- [92] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, *Appl. Phys. Rev.* **6**, 021314 (2019).
- [93] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [94] P. M. Q Cruz and Bruno Murta, In preparation.
- [95] A. M. Krol, A. Sarkar, I. Ashraf, Z. Al-Ars, and K. Bertels, Efficient decomposition of unitary matrices in quantum circuit compilers, *Appl. Sci.* **12**, 759 (2022).
- [96] R. Iten, O. Reardon-Smith, E. Malvetti, L. Mondada, G. Pauvert, E. Redmond, R. S. Kohli, and R. Colbeck, Introduction to UniversalQCompiler, [arXiv:1904.01072](https://arxiv.org/abs/1904.01072) (2019).
- [97] Qiskit Documentation: transpile, (Qiskit), last accessed on 22-02-2023.
- [98] A. Kissinger and J. van de Wetering, Pyzx: Large scale automated diagrammatic reasoning, in *Proceedings 16th International Conference on Quantum Physics and Logic, Chapman University, Orange, CA, USA., 10-14 June 2019*, edited by Bob Coecke and Matthew Leifer, Electron. Proc. Theor. Comput. Sci. EPTCS Vol. 318 (Open Publishing Association, 2020), pp. 229–241.
- [99] R. Duncan, A. Kissinger, S. Perdrix, and J. van de Wetering, Graph-theoretic Simplification of quantum circuits with the ZX-calculus, *Quantum* **4**, 279 (2020).
- [100] M. Backens, H. Miller-Bakewell, G. de Felice, L. Lobski, and J. van de Wetering, There and back again: A circuit extraction tale, *Quantum* **5**, 421 (2021).
- [101] N. de Beaudrap, A. Kissinger, and J. van de Wetering, Circuit Extraction for ZX-Diagrams Can Be #P-Hard, in *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, edited by M. Bojańczyk, E. Merelli, and D. P. Woodruff, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 229 (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2022), pp. 119:1–119:19.
- [102] B. Murta and J. Fernández-Rossier, Gutzwiller wave function on a digital quantum computer, *Phys. Rev. B* **103**, L241113 (2021).
- [103] K. Seki, Y. Otsuka, and S. Yunoki, Gutzwiller wave function on a quantum computer using a discrete Hubbard–Stratonovich transformation, *Phys. Rev. B* **105**, 155119 (2022).
- [104] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, [arXiv:quant-ph/0005055](https://arxiv.org/abs/quant-ph/0005055) (2000).
- [105] L. K. Grover, Quantum Mechanics Helps in Searching for a Needle in a Haystack, *Phys. Rev. Lett.* **79**, 325 (1997).
- [106] T.-C. Wei, R. Raussendorf, and I. Affleck, Some Aspects of Affleck–Kennedy–Lieb–Tasaki Models: Tensor Network, Physical Properties, Spectral Gap, Deformation, and Quantum Computation, in *Entanglement in Spin Chains: From Theory to Quantum Technology Applications*, edited by A. Bayat, S. Bose, and H. Johannesson (Springer International Publishing, Cham, 2022), pp. 89–125.
- [107] Honeywell System Model H1, (Honeywell), last accessed on 22-02-2023, <https://www.honeywell.com/us/en/company/quantum/quantum-computer>.
- [108] P. Nation and B. Johnson, How to measure and reset a qubit in the middle of a circuit execution. (IBM Research), last accessed on 22-02-2023, <https://www.ibm.com/blogs/research/2021/02/quantum-mid-circuit-measurement/>.
- [109] M. Plesch and Č. Brukner, Quantum-state preparation with universal gate decompositions, *Phys. Rev. A* **83**, 032302 (2011).
- [110] G. Vidal and C. M. Dawson, Universal quantum circuit for two-qubit transformations with three controlled-NOT gates, *Phys. Rev. A* **69**, 010301(R) (2004).
- [111] Qiskit Documentation: two_qubit_cnot_decompose, (Qiskit), last accessed on 22-02-2023, https://qiskit.org/documentation/stubs/qiskit.quantum_info.two_qubit_cnot_decompose.html.
- [112] R. R. Tucci, An introduction to Cartan’s KAK decomposition for QC programmers, [arXiv:quant-ph/0507171](https://arxiv.org/abs/quant-ph/0507171) v1 (2005).
- [113] Qiskit Documentation: initialize, (Qiskit), last accessed on 22-02-2023, <https://qiskit.org/documentation/stubs/qiskit.compiler.transpile.html>.
- [114] T.-C. Wei and R. Raussendorf, Universal measurement-based quantum computation with spin-2 Affleck–Kennedy–Lieb–Tasaki states, *Phys. Rev. A* **92**, 012310 (2015).
- [115] K. Seki, T. Shirakawa, and S. Yunoki, Symmetry-adapted variational quantum eigensolver, *Phys. Rev. A* **101**, 052340 (2020).
- [116] D. Cruz, R. Fournier, F. Gremion, A. Jeannerot, K. Komagata, T. Tosić, J. Thiesbrummel, C. L. Chan, N. Macris, M.-A. Dupertuis, and C. Javerzac-Galy, Efficient quantum algorithms for ghz and W states, and implementation on the

- IBM quantum computer, *Adv. Quantum Technol.* **2**, 1900015 (2019).
- [117] A. Bärttschi and S. Eidenbenz, Deterministic Preparation of Dicke States, in *Fundamentals of Computation Theory* (Springer, 2019), pp. 126–139.
- [118] C. S. Mukherjee, S. Maitra, V. Gaurav, and D. Roy, Preparing dicke states on a quantum computer, *IEEE Trans. Quantum Eng.* **1**, 1 (2020).
- [119] S. Aktar, A. Bärttschi, A.-H. A. Badawy, and S. Eidenbenz, A divide-and-conquer approach to dicke state preparation, *IEEE Trans. Quantum Eng.* **3**, 1 (2022).
- [120] X.-M. Zhang, T. Li, and X. Yuan, Quantum State Preparation with Optimal Circuit Depth: Implementations and Applications, *Phys. Rev. Lett.* **129**, 230504 (2022).
- [121] R. D. P. East, J. van de Wetering, N. Chancellor, and A. G. Grushin, AKLT-states as ZX-diagrams: Diagrammatic reasoning for quantum states, *PRX Quantum* **3**, 010302 (2022).
- [122] D. E. Deutsch, A. Barenco, and A. Ekert, Universality in quantum computation, *Proc. R. Soc. London A* **449**, 669 (1995).
- [123] E. Knill, Approximation by quantum circuits, [arXiv:quant-ph/9508006](https://arxiv.org/abs/quant-ph/9508006) (1995).
- [124] Quantinuum Completes Hardware Upgrade; Achieves 20 Fully Connected Qubits, (Quantinuum), last accessed on 22-02-2023, <https://www.quantinuum.com/pressrelease/quantinuum-completes-hardware-upgrade-achieves-20-fully-connected-qubits>.
- [125] Y. Wan, R. Jördens, S. D. Erickson, J. J. Wu, R. Bowler, T. R. Tan, P.-Y. Hou, D. J. Wineland, A. C. Wilson, and D. Leibfried, Ion transport and reordering in a 2D trap array, *Adv. Quantum Technol.* **3**, 2000028 (2020).
- [126] J. Ramette, J. Sinclair, Z. Vendeiro, A. Rudelis, M. Cetina, and V. Vuletić, Any-to-any connected cavity-mediated architecture for quantum computing with trapped ions or Rydberg arrays, *PRX Quantum* **3**, 010344 (2022).
- [127] B. O’Gorman, W. J. Huggins, E. G. Rieffel, and K. B. Whaley, Generalized swap networks for near-term quantum computing, [arXiv:1905.05118](https://arxiv.org/abs/1905.05118) (2019).
- [128] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits, *Phys. Rev. X* **10**, 011022 (2020).
- [129] Hardware-aware approach for fault-tolerant quantum computation, (IBM Quantum), last accessed on 22-02-2023, <https://www.ibm.com/blogs/research/2020/09/hardware-aware-quantum>.
- [130] The IBM Quantum heavy hex lattice, (IBM Quantum), last accessed on 22-02-2023, <https://research.ibm.com/blog/heavy-hex-lattice>.
- [131] Expanding the IBM Quantum roadmap to anticipate the future of quantum-centric supercomputing, (IBM Quantum), last accessed on 22-02-2023, <https://research.ibm.com/blog/ibm-quantum-roadmap-2025>.
- [132] T. C. Ralph, K. J. Resch, and A. Gilchrist, Efficient Toffoli gates using qudits, *Phys. Rev. A* **75**, 022313 (2007).
- [133] B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O’Brien, A. Gilchrist, and A. G. White, Simplifying quantum logic using higher-dimensional Hilbert spaces, *Nat. Phys.* **5**, 134 (2009).
- [134] T. Bækkegaard, L. B. Kristensen, N. J. S. Loft, C. K. Andersen, D. Petrosyan, and N. T. Zinner, Realization of efficient quantum gates with a superconducting qubit-qutrit circuit, *Sci. Rep.* **9**, 13389 (2019).
- [135] D. Ristè, L. C. G. Govia, B. Donovan, S. D. Fallek, W. D. Kalfus, M. Brink, N. T. Bronn, and T. A. Ohki, Real-time processing of stabilizer measurements in a bit-flip code, *npj Quantum Inf.* **6**, 71 (2020).
- [136] L. Botelho, A. Glos, A. Kundu, J. A. Miszczyk, Ö. Salehi, and Z. Zimborás, Error mitigation for variational quantum algorithms through mid-circuit measurements, *Phys. Rev. A* **105**, 022441 (2022).
- [137] R. Yalovetzky, P. Minssen, D. Herman, and M. Pistoia, NISQ-HHL: Portfolio optimization for near-term quantum hardware, [arXiv:2110.15958](https://arxiv.org/abs/2110.15958) (2021).
- [138] D. Zhu *et al.*, Interactive protocols for classically verifiable quantum advantage, [arXiv:2112.05156](https://arxiv.org/abs/2112.05156) (2021).
- [139] J. M. Koh, S.-N. Sun, M. Motta, and A. J. Minnich, Experimental realization of a measurement-induced entanglement phase transition on a superconducting quantum processor, [arXiv:2203.04338](https://arxiv.org/abs/2203.04338) (2022).
- [140] P. Iyer and D. Poulin, A small quantum computer is needed to optimize fault-tolerant protocols, *Quantum Sci. Technol.* **3**, 030504 (2018).
- [141] T. H. Kyaw, T. Menke, S. Sim, A. Anand, N. P. D. Sawaya, W. D. Oliver, G. G. Guerreschi, and A. Aspuru-Guzik, Quantum Computer-Aided Design: Digital Aspinum Simulation of Quantum Processors, *Phys. Rev. Appl.* **16**, 044042 (2021).
- [142] J. Cai, A. Miyake, W. Dür, and H. J. Briegel, Universal quantum computer from a quantum magnet, *Phys. Rev. A* **82**, 052309 (2010).
- [143] H. J. Briegel and R. Raussendorf, Persistent Entanglement in Arrays of Interacting Particles, *Phys. Rev. Lett.* **86**, 910 (2001).
- [144] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [145] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [146] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, Variational quantum eigensolver with fewer qubits, *Phys. Rev. Res.* **1**, 023025 (2019).
- [147] G. Mazzola, P. J. Ollitrault, P. Kl. Barkoutsos, and I. Tavernelli, Nonunitary Operations for Ground-State Calculations in Near-Term Quantum Computers, *Phys. Rev. Lett.* **123**, 130501 (2019).
- [148] A. C. Y. Li, M. S. Alam, T. Iadecola, A. Jahin, D. M. Kurkcuoglu, R. Li, P. P. Orth, A. B. Özgüler, G. N. Perdue, and N. M. Tubman, Benchmarking variational quantum eigensolvers for the square-octagon-lattice Kitaev model, [arXiv:2108.13375](https://arxiv.org/abs/2108.13375) (2021).
- [149] J.-B. You, D. E. Koh, J. F. Kong, W.-J. Ding, C. E. Png, and L. Wu, Exploring variational quantum eigensolver ansatzes for the long-range XY model, [arXiv:2109.00288](https://arxiv.org/abs/2109.00288) (2021).
- [150] D. Hueriga, Variational quantum simulation of valence-bond solids, *Quantum* **6**, 874 (2022).
- [151] M. S. Jattana, F. Jin, H. De Raedt, and K. Michielsen, Assessment of the variational quantum eigensolver: Application to the heisenberg model, *Front. Phys.* **10**, 907160(2022).
- [152] J. L. Bosse and A. Montanaro, Probing ground-state properties of the kagome antiferromagnetic Heisenberg model using the

- variational quantum eigensolver, *Phys. Rev. B* **105**, 094409 (2022).
- [153] V. Feulner and M. J. Hartmann, Variational quantum eigensolver ansatz for the J_1 – J_2 -model, *Phys. Rev. B* **106**, 144426 (2022).
- [154] Pushing quantum performance forward with our highest Quantum Volume yet, (IBM Quantum), last accessed on 22-02-2023, <https://research.ibm.com/blog/quantum-volume-256>.
- [155] Quantinuum Announces Quantum Volume 4096 Achievement, (Quantinuum), last accessed on 22-02-2023, <https://www.quantinuum.com/pressrelease/quantinuum-announces-quantum-volume-4096-achievement>.
- [156] P. Niroula, R. Shaydulin, R. Yalovetzky, P. Minssen, D. Herman, S. Hu, and M. Pistoia, Constrained quantum optimization for extractive summarization on a trapped-ion quantum computer, *Sci. Rep.* **12**, 17171 (2022).
- [157] A. Mitra, Quantum quench dynamics, *Annu. Rev. Condens. Matter Phys.* **9**, 245 (2018).
- [158] S. R. White and A. E. Feiguin, Real-Time Evolution Using the Density Matrix Renormalization Group, *Phys. Rev. Lett.* **93**, 076401 (2004).
- [159] A. J. Daley, C. Kollath, U. Schollwöck, and G. Vidal, Time-dependent density-matrix renormalization-group using adaptive effective Hilbert spaces, *J. Stat. Mech.* (2004) P04005.
- [160] J. Haegeman, T. J. Osborne, and F. Verstraete, Post-matrix product state methods: To tangent space and beyond, *Phys. Rev. B* **88**, 075133 (2013).
- [161] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [162] D. Gobert, C. Kollath, U. Schollwöck, and G. Schütz, Real-time dynamics in spin 1/2 chains with adaptive time-dependent density matrix renormalization group, *Phys. Rev. E* **71**, 036102 (2005).
- [163] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [164] M. Lewenstein, A. Sanpera, and V. Ahufinger, *Ultracold Atoms in Optical Lattices: Simulating Quantum Many-Body Systems* (Oxford University Press, Oxford, UK, 2012).
- [165] V. Sharma and E. J. Mueller, Driven-dissipative control of cold atoms in tilted optical lattices, *Phys. Rev. A* **103**, 043322 (2021).
- [166] M. P. Zaletel, A. Kaufman, D. M. Stamper-Kurn, and N. Y. Yao, Preparation of Low Entropy Correlated Many-Body States via Conformal Cooling Quenches, *Phys. Rev. Lett.* **126**, 103401 (2021).
- [167] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, Detecting crosstalk errors in quantum information processors, *Quantum* **4**, 321 (2020).
- [168] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, Scalable mitigation of measurement errors on quantum computers, *PRX Quantum* **2**, 040326 (2021).
- [169] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [170] J. A. Jones and D. Jaksch, *Quantum Information, Computation and Communication* (Cambridge University Press, Cambridge UK, 2012).
- [171] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.5.013190> includes the QASM files (in .txt file format) of the optimized versions of the different subcircuits required to prepare the spin-1 and spin- $\frac{3}{2}$ VBS states. The instructions for the use of these circuits can be found in Appendix H of this paper.