


Preparing Bethe Ansatz Eigenstates on a Quantum Computer

John S. Van Dyke^{1,*} George S. Barron¹ Nicholas J. Mayhall,² Edwin Barnes,¹ and Sophia E. Economou¹

¹*Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*

²*Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

 (Received 29 April 2021; revised 22 August 2021; accepted 4 October 2021; published 9 November 2021)

Several quantum many-body models in one dimension possess exact solutions via the Bethe ansatz method, which has been highly successful for understanding their behavior. Nevertheless, there remain physical properties of such models for which analytic results are unavailable and which are also not well described by approximate numerical methods. Preparing Bethe ansatz eigenstates directly on a quantum computer would allow straightforward extraction of these quantities via measurement. We present a quantum algorithm for preparing Bethe ansatz eigenstates of the spin-1/2 XXZ spin chain that correspond to real-valued solutions of the Bethe equations. The algorithm is polynomial in the number of T gates and the circuit depth, with modest constant prefactors. Although the algorithm is probabilistic, with a success rate that decreases with increasing eigenstate energy, we employ amplitude amplification to boost the success probability. The resource requirements for our approach are lower than for other state-of-the-art quantum simulation algorithms for small error-corrected devices and thus may offer an alternative and computationally less demanding demonstration of quantum advantage for physically relevant problems.

DOI: [10.1103/PRXQuantum.2.040329](https://doi.org/10.1103/PRXQuantum.2.040329)

I. INTRODUCTION

Quantum computers hold the promise of transformative applications in a variety of fields, including cryptanalysis [1], quantum chemistry [2,3], materials science [4,5], and, potentially, combinatorial optimization [6,7]. To realize the full potential of quantum computing, large-scale fault-tolerant devices will ultimately be necessary. As these do not yet exist, much recent work has studied possible near-term applications in the present era of noisy intermediate-scale quantum computers (NISQs) [8,9]. In this context, a key question concerns the demonstration of “quantum advantage”—that is, the ability to perform computations that cannot be done efficiently with classical methods. Recently, quantum advantage has been shown for a superconducting processor sampling random quantum circuits [10] and photonic-based Gaussian boson sampling [11]. Although these are important achievements, the specific tasks performed have not been closely related to the practical applications mentioned above but have essentially been designed for the purpose of demonstrating advantage.

Thus, the realization of quantum advantage for a problem of practical interest remains open.

A question of increasing interest is what applications become feasible with small-scale error-corrected devices, i.e., in the intermediate era between NISQs and fault-tolerant quantum computers with many logical qubits and a high clock speed for non-Clifford gates. Recent estimates suggest that algorithms that provide only a quadratic speed-up over classical methods may have difficulty achieving quantum advantage on problem sizes accessible with small devices [12–14]. The simulation of quantum systems, on the other hand, can yield exponential improvement over conventional approaches. These still require formidable resources, despite recent algorithmic advances [15–18]. This motivates the search for physically interesting problems and algorithms that can lead to quantum advantage with fewer resources in the near future.

We propose the study of Bethe ansatz (BA) states on a quantum computer as a computationally less demanding route to the demonstration of quantum advantage for problems relevant to physics, including quantum magnetism [19,20], ultracold atoms [21], and unconventional superconductivity [22]. BA methods allow for deep insight into the static and dynamic properties of these many-body systems and are able to explore not only ground states but also interactions between complex collective excitations, such as magnons, and the response to quantum quench experiments [23].

*jsvandyke86@gmail.com

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

More specifically, the BA technique yields exact solutions to a class of one-dimensional quantum many-body models, including the spin-1/2 Heisenberg and Hubbard models, among others [24–27]. The resulting wave functions depend on algebraic equations that can be efficiently solved classically. The exponential growth of the Hilbert space with the system size L has historically limited the direct computational studies of the eigenstates to small systems. Instead, various mathematical techniques have been extensively developed to access physical quantities in the thermodynamic limit $L \rightarrow \infty$, bypassing the calculation of the wave function itself. While many different quantities can be determined, the difficulty of their calculation varies widely. In particular, arbitrary-range and higher-order correlation functions have been very challenging to access and remain an active area of research [28–30]. Quantum computers, however, can compute such correlation functions [5,31] straightforwardly, thus suggesting the possibility of quantum advantage for this task. The importance of higher-order correlation functions for strongly correlated systems has recently been emphasized [32]. The calculation of such observables using a quantum computer in turn hinges on the possibility of efficiently preparing the Bethe ansatz states.

To this end, we demonstrate an efficient quantum algorithm that can prepare a subset of the Bethe ansatz eigenstates of the one-dimensional XXZ chain, a model that is fundamental to the study of quantum magnetism. Our algorithm uses the so-called coordinate Bethe ansatz method and has polynomial scaling in the circuit depth and T -gate complexity, along with low constant prefactors. As we show with explicit gate counts for the corresponding circuits, the approach scales to large enough systems for the calculation of classically inaccessible quantities in near-term error-corrected devices. While the algorithm we present is probabilistic, we also show that amplitude amplification can be used to increase the success rate [33].

Algorithms have been previously given (and also implemented) that prepare exact eigenstates of quantum many-body models [34–37]. However, these have been largely limited to cases that are equivalent to noninteracting fermions, for instance, under the Jordan-Wigner mapping. In contrast, the XXZ model corresponds to an interacting fermionic problem, which is computationally much more difficult. We note that the possibility of constructing circuits to diagonalize Bethe ansatz-solvable models and measure challenging correlation functions has previously been suggested, though without an indication of how this could be done [34].

The importance of Bethe ansatz-solvable models for benchmarking NISQ devices has been previously recognized [37–39], as they provide exact values for quantities (such as the energy) to compare against the results of noisy quantum computations. On the other hand, the direct preparation of Bethe ansatz states has been relatively

unexplored. This question has recently been studied in Ref. [40], which considered treating Bethe ansatz states variationally (using the algebraic Bethe ansatz) and concluded that the approach was not scalable. Furthermore, that work did not make a connection to the possibility of quantum advantage. Apart from direct preparation of Bethe ansatz eigenstates, other works have considered variational approaches using generic ansatzes [38,39,41–43]. The comparison of the computational complexity of these methods to that of the direct construction is an interesting question for future studies, as are probabilistic algorithms for preparing other strongly correlated states [44].

The paper is organized as follows. Section II introduces the XXZ model and the elements of the Bethe ansatz solution needed for the construction of the algorithm. Section III describes the Bethe ansatz state-preparation algorithm. Section IV presents numerical results that validate the method and studies its success probability. This section also includes resource estimates for classically intractable problem sizes. Section V describes the amplitude amplification procedure for our algorithm and presents numerical calculations that confirm its success. Section VI compares our algorithm with conceptually simpler but less efficient approaches to the same task, explicitly verifying the enormous speed-up of our method. Section VII argues that quantum advantage can be achieved with Bethe-state preparation by comparison with classical computational methods, and presents additional applications of the algorithm. Finally, we conclude in Sec. VIII.

II. MODEL AND SOLUTION

We consider the one-dimensional spin-1/2 XXZ chain on L sites with periodic boundary conditions, the Hamiltonian of which is given by

$$H = \sum_{i=0}^{L-1} J_{xy} (S_i^x S_{i+1}^x + S_i^y S_{i+1}^y) + J_z S_i^z S_{i+1}^z, \quad (1)$$

with $S_L^\alpha \equiv S_0^\alpha$ ($\alpha = x, y, z$). Here, the S_j^α are the spin operators with eigenvalues $\pm 1/2$. The exact solution of this model via the Bethe ansatz method has been presented in Ref. [45] and many introductions to the problem (in both the coordinate and algebraic formulations) exist [46–50]. We follow the account of the Bethe ansatz method given in Ref. [48]. The eigenstates of the above Hamiltonian are given by

$$\psi(x_1, \dots, x_M) = \sum_P A_P \exp \left[i \sum_{j=1}^M k_{Pj} x_j \right], \quad (2)$$

where x_1, \dots, x_M label the positions of the M down spins in the chain (the Hamiltonian conserves the z component of the total spin, $S_{\text{tot}}^z = \sum_i S_i^z$) and the momenta k_i label

the different states. The wave function of Eq. (2) gives the amplitude for the M down spins to occur on the sites $x_j = 0, \dots, L - 1$. The summation here is over the $M!$ permutations of the down-spin sites. These permutations arise from the fact that, within Bethe ansatz models, scattering processes exchange momenta between particles but do not alter their magnitudes. The coefficients A_P are related by

$$\frac{A_P}{A_{P'}} = -\frac{1 + e^{i(k_{Pl} + k_{P'l})} - 2\Delta e^{ik_{Pl}}}{1 + e^{i(k_{Pl} + k_{P'l})} - 2\Delta e^{ik_{P'l}}} \equiv -e^{-i\Theta(k_{Pl}, k_{P'l})}. \quad (3)$$

To fix the coefficients, we take $A_I = 1$, where I is the identity permutation. Here, $\Delta = J_z/J_{xy}$ is the anisotropy in the interactions and P and P' are permutations that differ by a single transposition between adjacent elements, $P(l+1) = P'(l)$ and $P(l) = P'(l+1)$. The momenta k_i are constrained by the quantization conditions,

$$Lk_i = 2\pi I_i + \sum_j \Theta(k_i, k_j), \quad (4)$$

where I_i is an integer (half-integer) for M odd (even). Physically, these constraints on k_i arise from imposing periodic boundary conditions on the model. Equations (3) and (4) are a set of algebraic equations (the Bethe equations) for the quantum numbers $\{k_i\}$. In general, these equations admit complex solutions but for the special case when all $\{k_i\}$ are real, $\Theta(k_i, k_j)$ is also real and is given by

$$\Theta(k_i, k_j) = 2 \arctan \left[\frac{\Delta \sin(k_{ij}^-/2)}{\Delta \cos(k_{ij}^-/2) - \cos(k_{ij}^+/2)} \right]. \quad (5)$$

where $k_{ij}^\pm = k_i \pm k_j$.

In the spirit of hybrid quantum-classical algorithms, we solve the Bethe equations classically to obtain the momenta $\{k_i\}$ and phases $\Theta(k_i, k_j)$. These values are then used as input to our quantum algorithm for generating the corresponding eigenstate. Our algorithm allows for the preparation of Bethe ansatz eigenstates for which the $\{k_i\}$ are real, such that A_P and $e^{ik_{Pj}x_j}$ amount to complex phases applied to the second-quantized basis states of the system.

III. BETHE ANSATZ STATE-PREPARATION ALGORITHM

The quantum algorithm for preparing a Bethe ansatz state consists of several steps and is summarized in Algorithm 1. The general structure of our approach is based on the linear-combination-of-unitaries (LCU) method (which also finds application in the Taylor-series approach to Hamiltonian simulation) [51,52]. However, a key difference is that our algorithm aims to generate specific quantum states starting from a particular initial

-
- 1: Prepare the Dicke state $|D_{L,M}\rangle$ on the system qubits
 - 2: Create permutation labels while applying pieces of A_P
 - 3: Apply $e^{ik_{Pj}x_j}$ using the “faucet” method
 - 4: Reverse permutation label (without phases)
 - 5: Measure permutation label, with success on $|00\dots 0\rangle$
-

Algorithm 1. Bethe-state preparation.

state, rather than compiling a generic unitary evolution operator. In addition to the L qubits representing the system, a register of M^2 ancilla qubits are used to label the different permutation terms in Eq. (2). By preparing a superposition of the allowed label values on these ancillas, using these to apply controlled operations on the system, and finally disentangling the label and system registers, we perform the summation over all permutations present in Eq. (2). This process is facilitated by introducing a second ancillary register of M qubits that we call the “faucet register,” along with one additional ancilla work qubit. Thus, the algorithm requires a total of $M^2 + M + 1$ ancilla qubits.

The algorithm begins by preparing the Dicke state on L sites with M down spins, $|D_{L,M}\rangle$. Relabeling $|\uparrow\rangle \equiv |0\rangle$, $|\downarrow\rangle \equiv |1\rangle$, $|D_{L,M}\rangle$ is the equal superposition (that is, without relative phases) of all basis states on L qubits with Hamming weight M . This state forms the underlying “canvas” on which the phases in Eq. (2) are applied. Dicke-state preparation can be accomplished using the recent deterministic algorithm of Ref. [53], for which the gate count has been improved in Ref. [54]. This algorithm uses an inductive method to prepare smaller Dicke states that are subsequently combined to yield the desired one. We use this algorithm in our explicit circuit constructions, though any other deterministic method of preparing $|D_{L,M}\rangle$ would also work.

As discussed above, the amplitudes that must be applied to $|D_{L,M}\rangle$ to generate a Bethe ansatz state depend on the permutations $\{P\}$ of M objects. We use the permutation-label register to create the different permutations and their associated phases A_P . Naively, one could use an integer encoded in a binary representation to label each of the permutations. The difficulty with this approach is that the number of permutations is $M!$, so that imprinting the phases A_P and $e^{ik_{Pj}x_j}$ onto $|D_{L,M}\rangle$ would require combinatorially many operations. This leads to circuit depths and complexities that are superexponential in M , quickly becoming unfeasible as M grows (we explore a concrete realization of this approach in Sec. VI). To overcome this fundamental limitation of this method and design an efficient algorithm, we introduce a conceptually distinct approach for labeling the permutations. Rather than assigning an arbitrary number to a given permutation, we implement its explicit action on the string consisting of the numbers $1, \dots, M$. As described below, this allows for an

efficient generation of the permutation labels, while also generating the distinct A_P simultaneously.

The permutation-label register consists of M subregisters, each of which can store an integer value $k \in \{1, \dots, M\}$. To represent a valid permutation, the subregisters must contain distinct values (for instance, $|213\rangle$ is valid whereas $|233\rangle$ is not). We use a one-hot encoding such that each subregister consists of M qubits and the number k is represented by a 1 on the k th qubit and 0s on the rest. Thus, for $M = 3$, the allowed states on each subregister are $|1\rangle \equiv |001\rangle$, $|2\rangle \equiv |010\rangle$, and $|3\rangle \equiv |100\rangle$. This one-hot encoding requires M^2 qubits to represent the complete label. The use of the one-hot encoding facilitates a trade-off between time and space resources [15,55,56] by reducing the number of controls required to implement the necessary phase gates.

The goal of step 2 of Algorithm 1 is to create the state $(1/\sqrt{M!}) \sum_P A_P |P\rangle$ on the permutation-label register. The phases A_P are kicked back onto the system qubits, while the $|P\rangle$ are used to apply the conditional gates needed in step 3, as explained below. For clarity, we first describe the construction of the equal superposition of all permutation labels. We then show how to slightly modify this procedure to simultaneously generate the phases A_P for all $M!$ permutations. We use an iterative method to construct the permutation-label state, starting from the vacuum state $|00\dots 0\rangle$ on M^2 qubits. The complete label superposition state is built up sequentially from the first (rightmost) subregister to the last (leftmost) using a series of exchange-type gates. We describe the method inductively as follows. The zeroth sublabel is prepared by setting the zeroth qubit of the zeroth subregister to 1 (in the following, the index k is enumerated starting from 0). Assume that the k th sublabel (i.e., an equal superposition of permutations of integers 1 through $k+1$) is constructed on the $k+1$ rightmost subregisters. Set the $(k+1)$ th qubit of the $(k+1)$ th subregister to 1, thus introducing the next integer value to be included in the permutation-label state. Perform the exchange-type ASWAP gate [57,58],

$$A(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & e^{i\phi} \sin(\theta) & 0 \\ 0 & e^{-i\phi} \sin(\theta) & -\cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (6)$$

between the $(k+1)$ th qubits of subregisters $k+1$ and k , with $\theta = \arccos(1/\sqrt{k+2})$, $\phi = 0$. This generates a superposition state consisting of two sets of terms, those in which the 1 remains in the $(k+1)$ th subregister and those in which it is transferred to the k th. In the latter case, the $(k+1)$ th subregister now contains all zeros, while the k th has two qubits with 1, which is not valid. This is fixed by applying controlled-SWAP gates between all qubits $l < k+1$ in subregisters k and $k+1$, controlled on the state of qubit $k+1$ in subregister k . Taken together, these

operations produce a partial SWAP between subregisters $k+1$ and k ,

$$|k+1\rangle_{k+1} |m\rangle_k \rightarrow \frac{1}{\sqrt{k+2}} |k+1\rangle_{k+1} |m\rangle_k + \sqrt{\frac{k+1}{k+2}} |m\rangle_{k+1} |k+1\rangle_k, \quad (7)$$

where $|i\rangle_j$ is the one-hot encoded state for i on subregister j , and $m < k$ by construction. One repeats this partial swapping process, now between subregisters k and $k-1$, then between $k-1$ and $k-2$, and so on, until the last register is swapped. By implementing the inductive process up to the $(M-1)$ th subregister, the complete equally weighted superposition of permutation labels is formed. As an example, for $M = 3$, the above algorithm generates the following sequence of state transformations:

$$\begin{aligned} &|000\rangle |000\rangle |000\rangle \xrightarrow{X} |000\rangle |000\rangle |001\rangle \\ &\xrightarrow{X} |000\rangle |010\rangle |001\rangle \\ &\xrightarrow{\text{ASWAP}} \frac{1}{\sqrt{2}} |000\rangle (|010\rangle |001\rangle + |000\rangle |011\rangle) \\ &\xrightarrow{\text{CSWAPs}} \frac{1}{\sqrt{2}} |000\rangle (|010\rangle |001\rangle + |001\rangle |010\rangle) \\ &\xrightarrow{X} \frac{1}{\sqrt{2}} |100\rangle (|010\rangle |001\rangle + |001\rangle |010\rangle) \\ &\xrightarrow{\text{ASWAP}} \frac{1}{\sqrt{6}} |100\rangle |010\rangle |001\rangle + \frac{1}{\sqrt{3}} |000\rangle |110\rangle |001\rangle \\ &\quad + \frac{1}{\sqrt{6}} |100\rangle |001\rangle |010\rangle + \frac{1}{\sqrt{3}} |000\rangle |101\rangle |010\rangle \\ &\xrightarrow{\text{CSWAPs}} \frac{1}{\sqrt{6}} |100\rangle |010\rangle |001\rangle + \frac{1}{\sqrt{3}} |010\rangle |100\rangle |001\rangle \\ &\quad + \frac{1}{\sqrt{6}} |100\rangle |001\rangle |010\rangle + \frac{1}{\sqrt{3}} |001\rangle |100\rangle |010\rangle \\ &\xrightarrow{\text{ASWAP}} \frac{1}{\sqrt{6}} |100\rangle |010\rangle |001\rangle + \frac{1}{\sqrt{6}} |010\rangle |100\rangle |001\rangle \\ &\quad + \frac{1}{\sqrt{6}} |010\rangle |000\rangle |101\rangle + \frac{1}{\sqrt{6}} |100\rangle |001\rangle |010\rangle \\ &\quad + \frac{1}{\sqrt{6}} |001\rangle |100\rangle |010\rangle + \frac{1}{\sqrt{6}} |001\rangle |000\rangle |110\rangle \\ &\xrightarrow{\text{CSWAPs}} \frac{1}{\sqrt{6}} |100\rangle |010\rangle |001\rangle + \frac{1}{\sqrt{6}} |010\rangle |100\rangle |001\rangle \\ &\quad + \frac{1}{\sqrt{6}} |010\rangle |001\rangle |100\rangle + \frac{1}{\sqrt{6}} |100\rangle |001\rangle |010\rangle \\ &\quad + \frac{1}{\sqrt{6}} |001\rangle |100\rangle |010\rangle + \frac{1}{\sqrt{6}} |001\rangle |010\rangle |100\rangle. \end{aligned} \quad (8)$$

The explicit circuit for $M = 3$ is shown in Fig. 1, where the gates inside the red dashed rectangles are excluded at this point. Furthermore, gates acting on distinct qubits are pushed to the left, thereby reducing the circuit depth. This produces a different sequence of intermediate states than above, but the final state is the same.

A slight modification of the above procedure allows one to simultaneously apply the phases A_P to the appropriate terms in the superposition. After each partial subregister swap, one applies a controlled-phase gate with angle $\Theta(k_i, k_j) + \pi$, where i, j are the integer values that are swapped and where the additional phase π implements the signature of the permutation (red dashed rectangles in Fig. 1). Suppose, as in Eq. (7), that the value $(k + 1)$ is swapped to the right. Then the $(k + 1)$ th qubit of the right subregister can serve as the target qubit in the required controlled phase. Since the left subregister can store any value $m < k + 1$, a separate controlled phase is used for each possibility, where the control bits are given by the values m . This leads to a total of $M^3/3 - M^2/2 + M/6$ controlled-phase gates for this part of the algorithm.

By the end of the construction, the phases A_P have been applied to the corresponding permutation-label state, having been successively built up from the elementary transpositions of which they are composed. The explicit circuit for $M = 3$ is displayed in Fig. 1, where now the gates in the red dashed rectangles are included, in order to produce the phases A_P . At this point in the construction, the total state of the physical system and the permutation label is

$$\left(\frac{1}{\sqrt{M!}} \sum_P A_P \otimes_{j=M}^1 |Pj\rangle_P \right) |D_{L,M}\rangle_s, \quad (9)$$

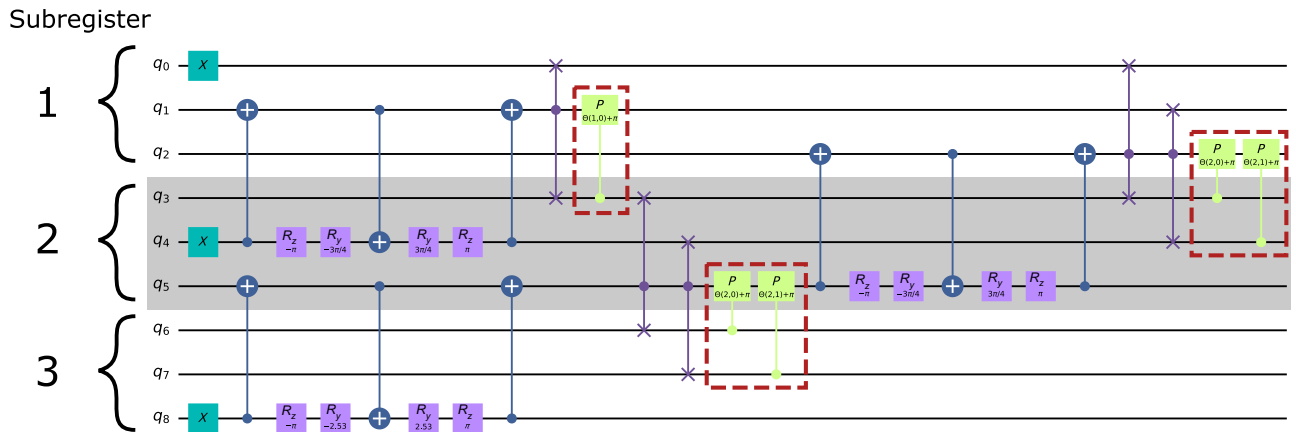


FIG. 1. The circuit to prepare the permutation-label ancilla state for $M = 3$. When the gates inside the red dashed rectangles are excluded, the circuit realizes the equally weighted superposition encoding the permutations of three objects. With these gates included, the circuit applies the additional permutation-dependent phases A_P . The three subregisters are indicated by shading and the initial state of each qubit is $|0\rangle$. The alternating CNOT gates interleaved with R_y and R_z rotations realize the ASWAP gates.

where $|\cdots\rangle_P$ is a state of the permutation-label qubits and $|\cdots\rangle_S$ is a state of the system qubits.

In the next step, one applies the position-dependent phase factors $e^{ik_{Pj}x_j}$ to the relevant basis states on the system qubits (step 3 of Algorithm 1). To do this, we introduce an efficient method that acts on all the $\binom{L}{M}$ basis states in $|D_{L,M}\rangle$ simultaneously, yielding an enormous speed-up over classical approaches. The technique, which we call the “faucet” method, is based on the observation that the positions x_j take integer values $x_j = 0, \dots, L - 1$, so that the total phase $e^{ik_{Pj}x_j}$ can be produced by x_j repetitions of the phase $e^{ik_{Pj}}$.

For this part of the algorithm, we use the M additional ancilla qubits comprising the faucet register. Each of these new qubits is initialized to $|1\rangle$. In the outer loop of the faucet subroutine, one traverses the register of the system qubits site by site from $x = 0$ to $x = L - 1$. At each site, if it is occupied by a down spin (i.e., the bit is 1), one turns off the next faucet ancilla qubit, $|1\rangle \rightarrow |0\rangle$. This is achieved through a sequence of multicontrolled X gates, which are controlled on the previous ancilla being in the state $|0\rangle$ and the next one being in the state $|1\rangle$, along with the additional control that the current system site is a $|1\rangle$. Since the meaning of the “next faucet ancilla” at a given site depends on the bit string, one must generically apply a multicontrolled X gate for each ancilla at every step. We note that one can decrease the number of gates for sites near the edges of the chain. For instance, at site 2 at most two faucets could have been turned off, so that the later ones do not need to be checked.

Next, the phases $e^{ik_{Pj}}$ ($j = 1, \dots, M$) are applied to the system qubits, each being controlled on the state of one of the faucet ancillas. For the j th ancilla, this gate is also controlled on the state of the permutation-label subregister j (since the value of k_{Pj} is permutation dependent). By the

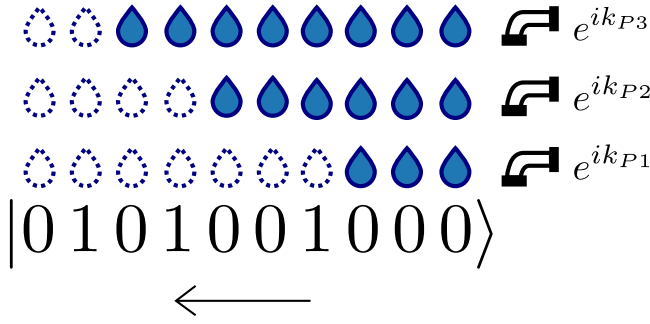


FIG. 2. A schematic diagram illustrating the idea of the faucet method. Phases $e^{ik_{Pj}}$ are applied for each system qubit while traversing the bit string. When a “1” is encountered, the next faucet in the list is turned off, so that no more phases with the given k_{Pj} value are applied.

end of the system bit string, all the faucet-register qubits are in state $|0\rangle$. Thus, the subroutine can be compared to a set of M running faucets (that correspond to applying the phases $e^{ik_{Pj}}$), which are turned off at the appropriate times (upon encountering a “1” in the traversal of the system register). This analogy is illustrated in Fig. 2. In total, this step requires M^2L doubly controlled-phase gates and a number of Toffoli gates that scales as ML (as mentioned above, some gates can be skipped near the edge of the system). In our numerical implementation of the method, we introduce additional work qubits to facilitate the construction of these gates using chains of Toffolis [59].

After the relevant phases have been applied, it remains to disentangle the system from the permutation label (the entanglement having been generated during the faucet method, since the phases there are permutation dependent). This is accomplished by applying the inverse of the circuit that generates the permutation-label superposition, without the additional controlled-phase gates that are used to produce the A_P phases. The fact that the phases A_P and $e^{ik_{Pj}x_j}$ have been applied to the initial Dicke state implies that the permutation-label reversal will not completely disentangle the system qubits from the permutation-label register. However, it turns out that the $|00\dots 0\rangle$ component of the permutation state corresponds precisely with the occurrence of the target Bethe ansatz state on the system qubits. That is, the full state vector takes the form

$$|\psi\rangle = \alpha|00\dots 0\rangle_p |\psi_B\rangle_s + \beta|\phi_j\rangle, \quad (10)$$

where $|\psi_B\rangle_s$ is the normalized target Bethe ansatz state on the system qubits $|\phi_j\rangle$ is a junk state with $(|00\dots 0\rangle\langle 00\dots 0|_p \otimes \mathbb{1}_s)|\phi_j\rangle = 0$. Thus, by measuring the permutation-label qubits, the target Bethe ansatz state is successfully prepared on the outcome $|00\dots 0\rangle$. This result is essentially that obtained from LCU methods [52],

though here we have the additional construction of A_P during the label-preparation step, which is not present in the standard LCU. To illustrate the full algorithm, the complete circuit to construct a BA eigenstate with $L = 4$, $M = 2$ is given in Appendix A. As discussed in our numerical simulations below, the success probability $|\alpha|^2$ depends on the system parameters and also varies between different eigenstates. Thus, in general one must repeat the procedure multiple times to obtain the correct state, which can then be used to calculate physical quantities or in other applications, as we discuss in Sec. VII. We also show in Sec. V that amplitude amplification can be used to boost the success rate, thereby reducing the overall resource requirements.

IV. NUMERICAL SIMULATIONS

To calculate the momenta $\{k_i\}$ defining the Bethe eigenstates, we solve the Bethe equations iteratively using the approach presented in Ref. [48]. We then perform numerical simulations of Algorithm 1 using the IBM QISKIT library’s state-vector simulator [60]. These calculations verify the correctness of our algorithm and reveal its success probabilities for the sufficiently small systems that can be studied on a classical computer. However, we can also explicitly construct the circuits that would need to be run for much larger instances, far beyond what is classically tractable. The corresponding circuit depths and gate counts in these cases indicate that our algorithm is feasible for near-term error-corrected quantum computers. We stress that our analysis does not rely on asymptotic resource scaling arguments but, rather, provides exact gate counts, since the corresponding circuits are precisely known. Although we do not compile our algorithm down to an error-correcting code such as the surface code, we estimate the required number of T gates below.

In Fig. 3(a) we present the numerically calculated success probability of the algorithm for preparing selected eigenstates when $L = 2M$ and $M = 2, 3, 4$, with $J_{xy} = 1$, $J_z = -1/2$. The interaction strength in this case corresponds to the critical regime of the ferromagnetic model. We note the eigenstates included in Fig. 3 are not meant to comprise the complete set of real-valued solutions but, rather, are simply those for which we obtain numerical solutions of the Bethe equations by using the algorithm presented in Ref. [48]. Two general trends are apparent: a significant suppression of the success rate with increasing M and a more moderate suppression as a function of the eigenstate energy, within each set of system parameters L, M . The worst-case probabilities are roughly consistent with $1/M!$, although we have only limited values of M to support this (larger M being outside of our computational resources for classical simulation). This value is further supported by Fig. 3(b), which shows the success

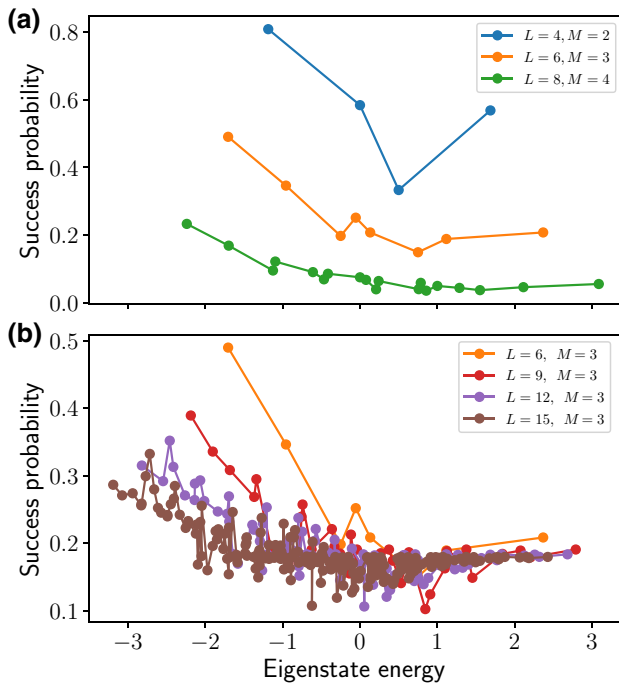


FIG. 3. (a) The success probabilities for the preparation of selected eigenstates as a function of their energies, when $L = 2M$ and $M = 2, 3, 4$. (b) The success probabilities of selected eigenstates for $M = 3$ and varying L . The Hamiltonian parameter values are $J_{xy} = 1$ and $J_z = -1/2$.

probabilities for various eigenstates when $M = 3$ and for different values of L , with the clear trend that increasing L tends to flatten the success probability across the spectrum. Although the low-energy states enjoy less of an advantage over the higher-energy ones in this case, the lowest probabilities are still around $1/M!$ on average. However, the behavior of the success rate changes significantly depending on the value of J_z . These effects are considered in Appendix B.

These results suggest that one can go to very large system sizes L while still preserving a reasonably large success probability, if M is sufficiently small. We note that the case of small M is particularly interesting for physics applications in the ferromagnetic regime of the model. In this case, the all-up state $|0\rangle^{\otimes L}$ is a ground state of the model, while small M states are low-lying excited states of interacting magnons (spin waves). The ability to study these states as a function of M , as enabled by our algorithm, would yield deeper insight into the development of strong correlations in these systems as the number of interacting quasiparticles grows. While the ferromagnetic regime is especially natural for our algorithm, we note that interesting physics in the paramagnetic and anti-ferromagnetic regimes can also be explored at small values of M . These correspond to highly excited eigenstates,

which are relevant, for instance, in the study of quantum thermalization [61,62].

For these applications (and others discussed below), it appears feasible to access values of L and M that would not be classically simulable (even by approximate methods), while maintaining relatively modest resource requirements for the algorithm. This is demonstrated in Fig. 4, which provides circuit depths [Fig. 4(a)], Toffoli-gate counts [Fig. 4(b)], controlled-phase gate counts [Fig. 4(c)], and the number of qubits required [Fig. 4(d)] for the Bethe-state-preparation algorithm. The linear scaling of all these metrics in L is immediately apparent. The slopes of these lines for the case $M = 5$ are (a) 39, (b) 11, (c) 25, and (d) 1, respectively. The results for the number of controlled-phases and qubits are in exact agreement with the analytical results of Sec. III.

Apart from the asymptotic scaling behavior, the absolute values of the circuit depths and gate counts are seen to be very low, on the order of 10^3 – 10^4 , even for large systems of $L \sim 100$ sites. Furthermore, the total number of qubits required (approximately 10^2) is also quite reasonable for small error-corrected devices. In Fig. 5, we show the total gate and measurement counts for the case $M = 5$ as L is varied. This indicates that the controlled-phase, R_y , and Toffoli gates are the most prevalent non-Clifford operations in the algorithm. To estimate the fault-tolerant resources needed, we therefore convert the counts for these gates into the corresponding numbers of T gates. Following Ref. [63], we assume a worst-case scenario for the number of T gates needed to realize an arbitrary z rotation to be given by $4 \log_2(1/\epsilon) + 11$, where ϵ is the rotation synthesis error [64]. Similarly, arbitrary y rotations can be performed by conjugation with Clifford operations. As in Ref. [16], we replace each Toffoli gate with two T gates [65]. For $L = 100$, $M = 5$, this leads to a T count of approximately 6.2×10^5 for a single run of the state-preparation algorithm, with $\epsilon = 10^{-10}$. Assuming a worst-case success probability of $1/M!$, approximately 120 attempts would need to be performed on average to correctly generate the target eigenstate. This yields approximately 7.4×10^7 T gates overall, which is comparable to the estimates for simulating the Hubbard model using the methods of Ref. [16]. We note that the estimates in that work involve optimizing an error budget between multiple sources (Trotterization, phase estimation, and rotation synthesis) and do not appear to include the cost of preparing a good initial state for the phase estimation routine. Furthermore, the above estimate for our algorithm assumes the seemingly worst-case scenario in the number of repetitions (approximately $M!$), whereas the results in Fig. 3 suggest that lower-energy states require fewer repetitions in general. To reduce the number of repetitions required for our algorithm, we implement amplitude amplification in the following section.

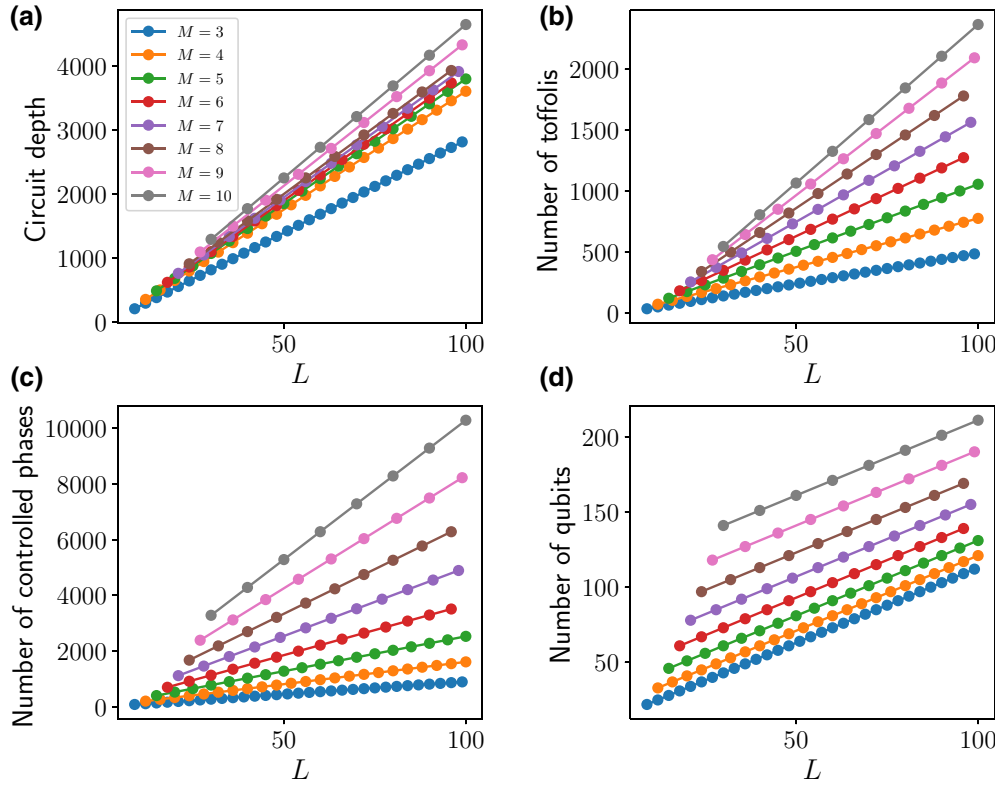


FIG. 4. The Bethe-state-preparation (a) circuit depth, (b) number of Toffoli gates, (c) number of controlled-phase gates, and (d) number of qubits versus the system size L , for different numbers of down spins M .

V. AMPLITUDE AMPLIFICATION

Amplitude amplification, a generalization of the well-known Grover search algorithm, is a quantum subroutine that can boost the probability of a desired measurement outcome, leading in general to a square-root improvement in the number of repetitions required for the success of a probabilistic algorithm [33]. For our problem, we use \mathcal{B} to denote Algorithm 1 with the measurement step removed.

Amplitude amplification defines an operator

$$\mathcal{Q} = -\mathcal{B}S_0\mathcal{B}^{-1}S_B, \quad (11)$$

where S_B changes the relative sign of the “good” states in the Hilbert space, while S_0 changes the relative sign of the vacuum state $|00\dots 0\rangle$. In the present case, the good states are the components of the Bethe ansatz state, which correspond to $|00\dots 0\rangle_p$ on the permutation-label qubits. S_B can therefore be implemented using a OR circuit on the permutation label, followed by Z on the work qubit that stores the result, after which the OR is uncomputed. In our numerical calculations of the success probability, we use the ancilla-free implementation of OR in the standard QISKIT circuit library. The ancilla-free approach is used here to decrease the number of qubits needed for the simulation, allowing us to study larger system sizes. Below, we examine an ancilla-based method for which the gate counts at large sizes are reduced. To produce S_0 , we use the same approach, with the OR circuit extended to include the system qubits (we do not implement the -1 in Eq. (11), as it is an overall phase). We present numerical results for amplitude amplification in Fig. 6(a). This confirms the clear enhancement of the algorithm success probability using this method. Although only one round of amplification is applied here, the protocol can be repeated in the standard way to further increase the success rate. Applying this improvement to the resource estimate of the previous section, the $M = 5$ worst-case eigenstates should

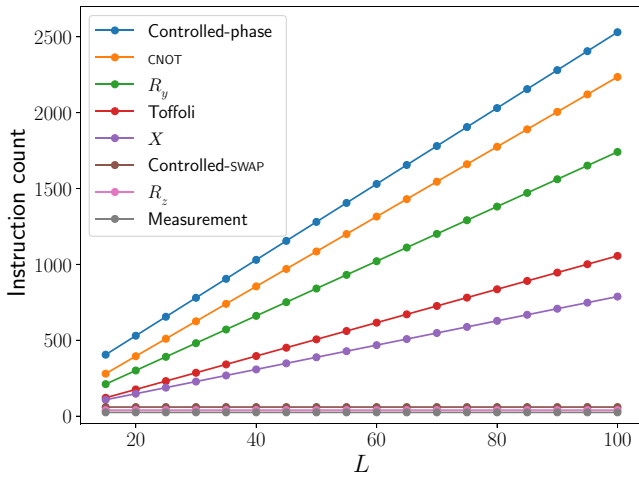


FIG. 5. The Bethe-state-preparation gate and measurement counts for $M = 5$ as a function of L .

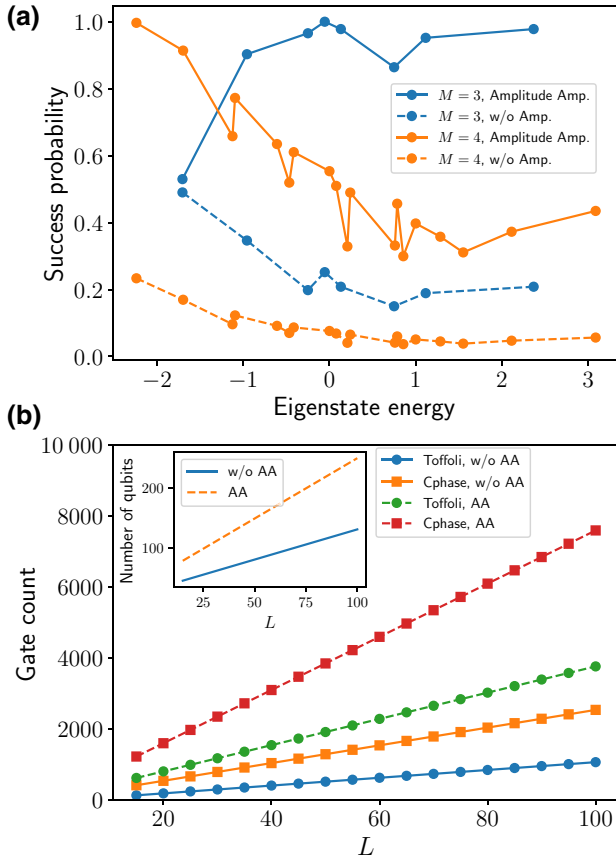


FIG. 6. (a) The success probability for Bethe-state preparation with (solid lines) and without (dashed lines) amplitude amplification (“AA”). In the former case, a single round of amplification is applied. The system parameters are $J_{xy} = 1$, $J_z = -1/2$, $L = 2M$. (b) The gate counts as a function of the system size L for Toffoli (circles) and controlled-phase (squares) gates (“Cphase”) without (solid lines) and with one round of amplitude amplification (dashed lines). The inset shows the total number of qubits without (solid) and with (dashed) amplitude amplification.

require on average $\sqrt{120} \approx 11$ repetitions of the algorithm \mathcal{B} to achieve success, leading to an overall T count of approximately 4.1×10^6 (neglecting the costs of S_B and S_0).

To estimate the resources needed for amplitude amplification more precisely, we implement the multicontrolled OR operation using elementary gates and ancillas [59]. The Toffoli count, controlled-phase count, and number of qubits are shown in Fig. 6(b) for a single round of amplification when $M = 5$. Since the S_0 reflection depends on the state of both the system and the permutation label, the implementation requires an additional $L + M^2 - M$ ancillas (we can reuse the M faucet ancillas for the amplification). The additional gates required are dominated by the cost of executing \mathcal{B} three times. Although the total number of qubits is approximately doubled in this approach,

we note that it may be possible to reduce this number by acting with the OR operation on only a subset of the qubits that are nominally necessary to identify the state. For example, although S_B is an OR circuit acting on the M^2 permutation-label qubits, in practice the computational basis states appearing in the junk state $|\phi_j\rangle$ can be distinguished from $|00 \dots 0\rangle_p$ by only acting on a reduced number of qubits. While the particular subset of qubits needed will depend on the state under consideration, it is possible to confirm the success of this approach by measuring the energy or other quantities that can be compared to exact analytic expressions.

We also implement a different version of amplitude amplification, which is a modified form of the oblivious amplitude amplification of Ref. [52]. Unfortunately, this method leads to a reduced fidelity of the actually prepared state with the exact target state, though in some cases the overlap remains quite high (> 0.99). We attribute this reduced fidelity to the nonunitarity of summing exponentials with unit modulus, since $|e^{ia} + e^{ib}| \neq 1$ in general. We note, however, that nearly deterministic success of the oblivious amplitude amplification procedure was obtained for the application of Ref. [52] (simulation of Hamiltonian dynamics with Taylor-series expansions). For this reason, it is less clear how the approach will fare for the Bethe-state-preparation problem at larger values of M . In addition, further modification to the method of Ref. [52] may ameliorate some of the difficulties with applying it to Bethe-state preparation in its present form.

VI. COMPARISON WITH ALTERNATIVE ALGORITHMS

To highlight the advantages of Algorithm 1, we compare it with conceptually simpler, but much less efficient, methods of preparing Bethe ansatz states on a quantum computer. First, one could imagine applying controlled-phase gates directly to each term in the Dicke-state superposition to generate the desired eigenstate. This approach still requires permutation-label ancillas to generate the linear combination of phases needed in Eq. (2). However, it has the seeming advantage of allowing one to combine the phases A_P and $e^{ikp_j x_j}$ into a single controlled-phase rotation, whereas the former term requires order M^3 and the latter one order M^2 controlled phases to implement using Algorithm 1. Nevertheless, it is clear that this benefit is vastly outweighed by the large number of terms in the superposition that must be separately addressed, $M! \binom{L}{M}$. For the case $L = 100$, $M = 5$ this amounts to approximately 9.0×10^9 controlled phases, compared to the 2530 of Algorithm 1.

A more promising approach is to use the “faucet” method of Algorithm 1 to handle the $e^{ikp_j x_j}$ phases while still applying the full A_P in a single multicontrolled

phase gate, rather than decomposing it into its elementary transpositions. This replaces the $M! \binom{L}{M}$ dependence above with $M!LM$. While the scaling is still inferior to that of Algorithm 1 for large M , it is conceivable that for small M this alternative method may be competitive. In particular, one can replace the complicated permutation label of Algorithm 1, which requires M^2 qubits to construct A_P in terms of individual transpositions, with a compressed label that simply assigns a number to each permutation. This approach uses significantly fewer qubits, at the expense of requiring more controls for the relevant phase gates. Since the permutation-label construction still needs to be reversed to disentangle the system from the ancillas, it is important that it can still be executed in a unitary fashion. This in turn requires an efficient method for generating an equal superposition of $M!$ states. We implement such states using the prime factorization $M! = 2^{n_2} 3^{n_3} \dots$. We then construct the permutation label as the tensor product of the binary representation of 2^{n_2} (using n_2 qubits) and W_n states for the odd prime factors. The equal superposition for the binary part of the label is easily generated by applying Hadamard gates to the relevant qubits, while various efficient algorithms exist for W_n state preparation [66–68]. We implement this algorithm numerically and verify that it successfully prepares Bethe ansatz eigenstates. Since the fundamental approach for creating the linear combination of phases is the same between this method and Algorithm 1, their success probabilities are equal. Unfortunately, explicit construction of the circuits for the alternative method indicates that the resource requirements are significantly higher, even for small M . This is shown in Fig. 7, which reveals that the number of controlled-phase and Toffoli gates required for the alternative method vastly exceeds those of Algorithm 1, even for $M = 5$.

It is also useful to compare Algorithm 1 with other more standard techniques, such as adiabatic state preparation [2,69,70]. As is well known, the evolution time to produce a large overlap with the desired eigenstate is expected to scale as the inverse square of the gap between the energy of the given state and that of the next closest one. This time can be short, for instance, in the ordered regime of the XXZ chain, where the gap between the ground and first excited state remains finite in the large system limit. However, in the critical regime of the model, the spectrum is gapless and so the evolution time is expected to diverge. Furthermore, even in the ordered regime, the spectrum at higher energies above the ground state generically has dense regions, preventing an efficient preparation of those eigenstates by the adiabatic approach. Our algorithm does not suffer from such complications, as it implements the exact analytic expression for the wave function, irrespective of the gap between the target eigenstate and the other ones.

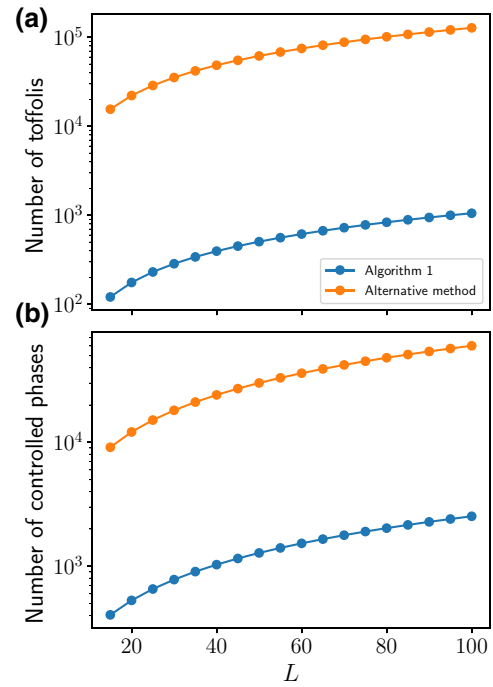


FIG. 7. A comparison between Algorithm 1 and the alternative method of the number of (a) Toffoli and (b) controlled-phase gates used, for $M = 5$.

VII. DISCUSSION

To achieve quantum advantage for a physically relevant problem, it should be the case that no classical method can deliver results of a comparable accuracy. Since the present quantum algorithm exactly prepares eigenstates of the XXZ chain, it is reasonable to compare it to the numerical exact diagonalization of finite-size systems (i.e., using classical computers). In a recent study, a matrix-free approach has been used to investigate Heisenberg spin chains up to length $L = 26$ [71]. In this work, the memory requirements have been vastly reduced compared to conventional methods, though the scaling remains exponential with the system size. Specifically the $S^z = 0$ subspace ($M = 13$) has been considered, for which the dimension is approximately 10^7 . In contrast, the $L = 100$, $M = 5$ subspace is roughly seven times larger (dimension approximately 7.5×10^7). Although state vectors of this size can still be stored in memory, we note that the computation time is also exponential in the system size, ultimately limiting the practicality of exact diagonalization.

In addition to numerically exact calculations, approximate tensor-network methods have been highly successful for studying one-dimensional quantum many-body systems with the matrix product state (MPS) ansatz [72,73]. However, these approaches are best suited for states with a relatively low amount of entanglement, such as gapped ground states obeying an area law for the entanglement entropy. This makes simulation of long-time dynamics

challenging, due to the growth of entanglement from, for instance, an initial product state. Our Bethe ansatz algorithm can prepare eigenstates throughout the spectrum at the same computational cost, including highly excited states the entanglement entropy of which grows more quickly than logarithmically, even in the small- M limit [74]. This yields an advantage over MPS methods for large systems when targeting these strongly entangled states. An explicit link between the Bethe ansatz and MPS has been developed in Refs. [75,76], which uses the algebraic Bethe ansatz to produce exact tensor-network representations for generic eigenstates. These networks have a PEPS-like structure (but with fewer physical indices), which underscores the computational intractability of such states for large systems.

In addition to computing arbitrary-range and higher-order correlation functions that are inaccessible with traditional Bethe ansatz methods, our algorithm has a number of other applications. Simulation of the real-time dynamics of many-body systems is widely recognized as a task allowing for quantum advantage. Such simulations often take the form of quench experiments, for which the system is initialized in an easy-to-prepare product state, then allowed to evolve under the influence of an interacting many-body Hamiltonian. Our algorithm would enable interesting variations on this approach, for instance by initializing the system in an eigenstate of a given value of the interaction strength, then subsequently evolving it with a different value. The evolution here can be performed using any of the known algorithms for quantum simulation, whether by Trotterization [16,77,78], Taylor expansions [52,79], or other approaches [51,80].

In a different direction, one could use our algorithm as a starting point to explore integrability-breaking perturbations. Thus, we consider a Hamiltonian of the form $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_p$, where \mathcal{H}_0 is solvable by the Bethe ansatz and \mathcal{H}_p includes perturbations that break the integrability of the total Hamiltonian \mathcal{H} (such as disorder in the coupling strengths). In this case, the Bethe-state-preparation algorithm is used to prepare an eigenstate of \mathcal{H}_0 , which then serves as an initial state for quantum annealing or phase estimation on \mathcal{H} . For sufficiently weak perturbations, the overlap of this state with the corresponding exact eigenstate of \mathcal{H} should be much greater than that of a mean-field or noninteracting trial state.

Although we focus on deploying our algorithm on small error-corrected quantum computers, one may also consider implementing it on present-day or near-term NISQ devices. Many of the controlled-phase gates in our algorithm involve repetitions of the same basic rotation angles, of which there are only $(M^2 + M)/2$ distinct values. This suggests replacing the exact values with variational parameters, similarly to Ref. [40]. The resulting variational form can then be optimized under the cost function $|E - E_B|$, where E is the energy calculated on the

quantum computer and E_B is the exact value, known analytically from the Bethe ansatz solution. We note that the present optimization problem should be significantly easier than that of a standard VQE, since the ideal values of the rotation angles can serve as a good initial guess. Updates to the parameters then serve to directly mitigate systematic errors due to over- or under-rotation in the controlled-phase gates.

VIII. CONCLUSIONS

We present a quantum algorithm for the efficient preparation of Bethe ansatz eigenstates of the XXZ model. To our knowledge, this is the first quantum algorithm for the direct preparation of eigenstates of an interacting many-body problem. The circuit depth and gate counts of the algorithm scale linearly in the system size, for a fixed number of down spins. Our algorithm is feasible to perform on small error-corrected devices of order 100 qubits, provided that the number of down spins is small. The usefulness of the approach can be extended through amplitude amplification. In particular, quantum advantage over classical computational methods appears to be achievable, with resource estimates that are comparable to the most efficient known quantum simulation algorithms. Our work suggests directions for future research, including the modification of the algorithm for the case of complex-valued $\{k_i\}$, and its generalization to other Bethe ansatz-solvable models, such as the one-dimensional Hubbard model.

ACKNOWLEDGMENTS

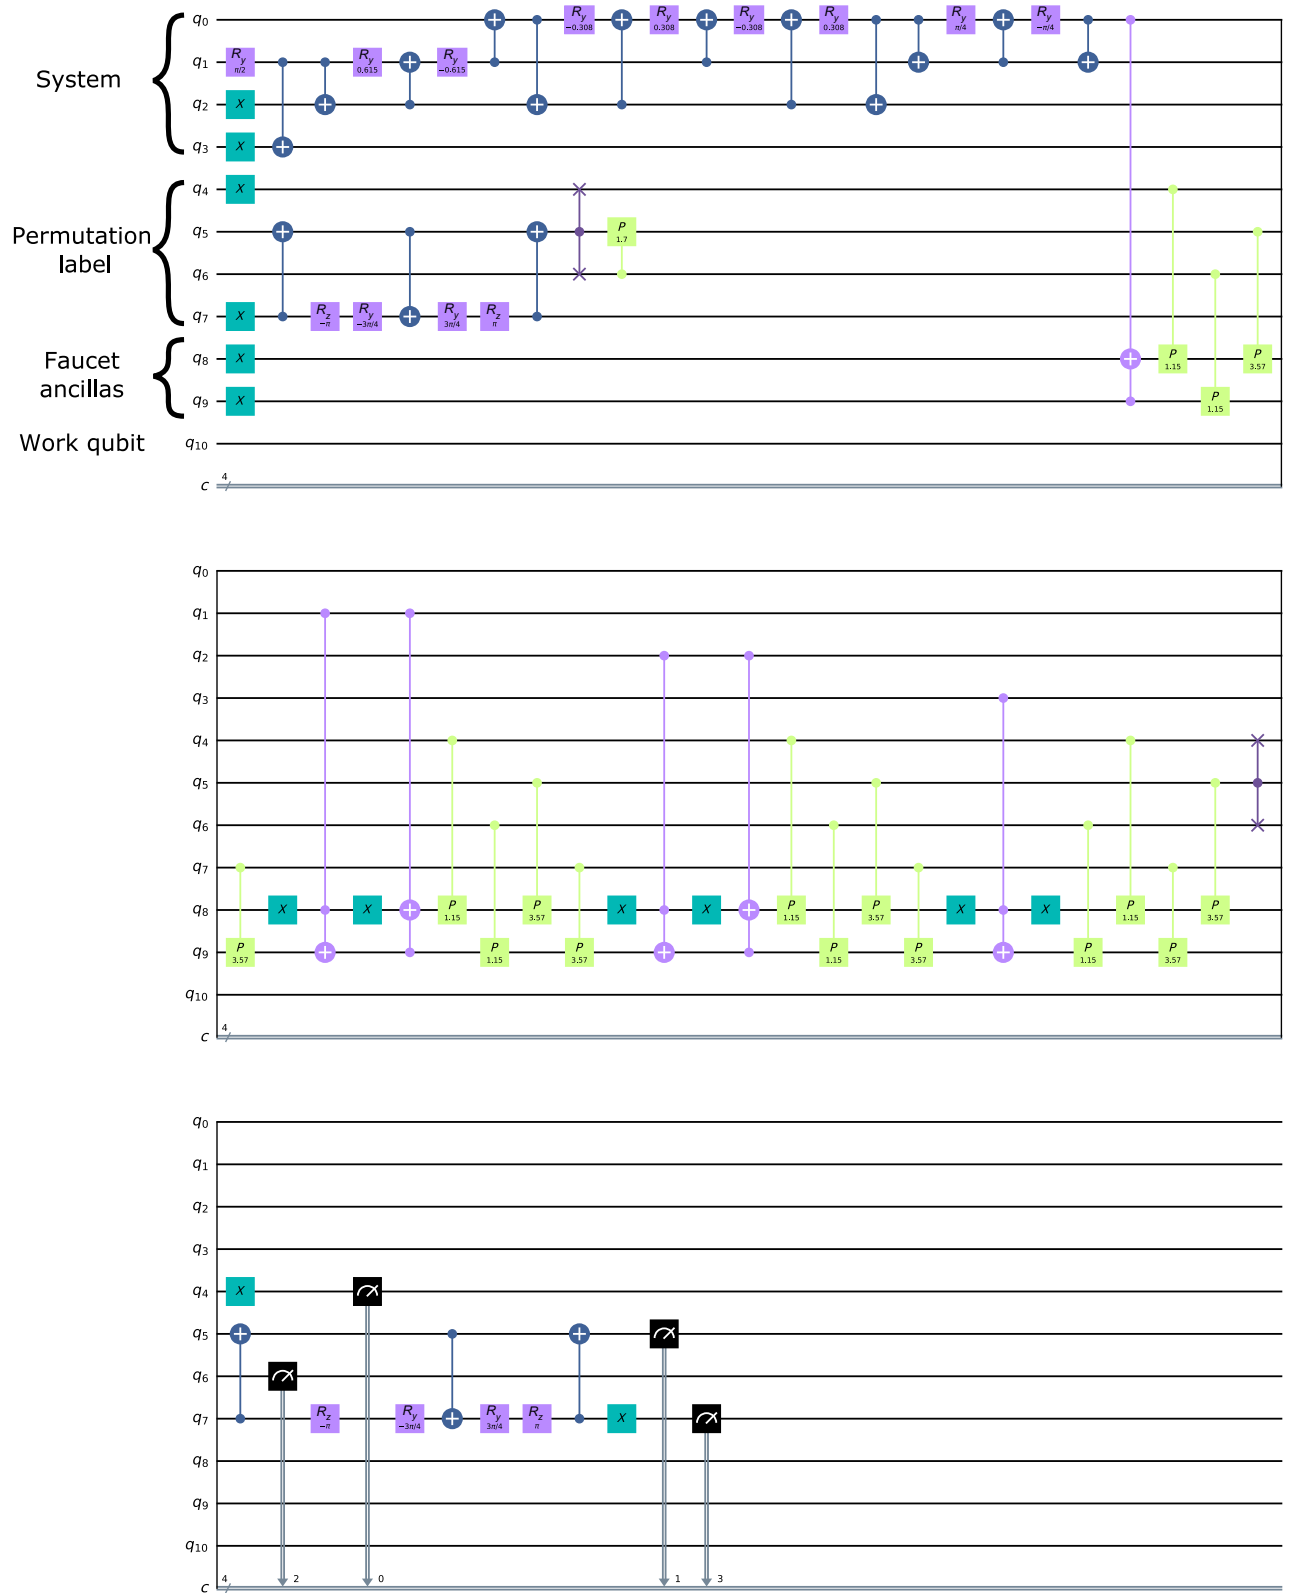
The numerical simulations were performed with IBM QISKIT and the QuSpin PYTHON library [81]. We would like to thank Chandra Sekhar Mukherjee, Jesko Sirker, and Rafael Nepomechie for helpful discussions. This work was supported by the U.S. Department of Energy (DOE). E.B. and N.J.M. acknowledge Award No. DE-SC0019199 and S.E.E. acknowledges the DOE Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA), Contract No. DE-SC0012704.

APPENDIX A: FULL CIRCUIT FOR $L = 4, M = 2$

In Fig. 8, we present the full quantum circuit to prepare a Bethe ansatz eigenstate with $L = 4$ and $M = 2$. We note that the work qubit is not needed in the $M = 2$ case but we include it here as a reminder that it is used to implement multicontrolled gates for $M > 2$.

APPENDIX B: SUCCESS PROBABILITY FOR DIFFERENT J_z

The success rate of the Bethe ansatz state preparation changes qualitatively with both the sign of J_z and whether

FIG. 8. The quantum circuit to prepare the Bethe ansatz eigenstate with $L = 4$, $M = 2$, $k_1 = 1.14676529$, and $k_2 = 3.56562369$.

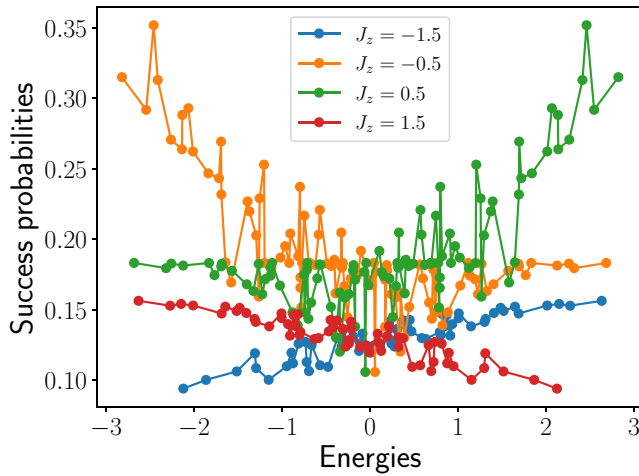


FIG. 9. The success probabilities for the preparation of selected eigenstates as a function of their energies, for $L = 12$, $M = 3$, $J_{xy} = 1$, with varying interaction strengths J_z .

it lies in the critical ($|J_z| < 1$) or noncritical ($|J_z| \geq 1$) regimes. These effects are illustrated in Fig. 9. Whereas the ferromagnetic model in the critical regime has higher success probabilities for lower-energy states, the antiferromagnetic case shows the opposite behavior. In fact, the success probabilities for different states are exactly mirrored across the $E = 0$ axis. This is especially interesting since the corresponding eigenstates at $E = -E_i$ for the FM case and $E = E_i$ for the AFM one are different in general. It is unclear at present why these distinct states should have the same success probability. When J_z is outside the critical regime, we find that the success probability can be lower than the $1/M!$ value (approximately 0.17 for $M = 3$) that appears to bound the critical case.

[1] P. Shor, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc. Press, 1994), p. 124.
[2] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
[3] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
[4] D. S. Abrams and S. Lloyd, Simulation of Many-Body Fermi Systems on a Universal Quantum Computer, *Phys. Rev. Lett.* **79**, 2586 (1997).
[5] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, Solving strongly correlated electron models on a quantum computer, *Phys. Rev. A* **92**, 062318 (2015).
[6] L. K. Grover, Quantum Mechanics Helps in Searching for a Needle in a Haystack, *Phys. Rev. Lett.* **79**, 325 (1997).
[7] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028* (2014).

[8] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
[9] I. H. Deutsch, Harnessing the power of the second quantum revolution, *PRX Quantum* **1**, 020101 (2020).
[10] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, and B. Burkett, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
[11] H.-S. Zhong, *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
[12] E. Campbell, A. Khurana, and A. Montanaro, Applying quantum algorithms to constraint satisfaction problems, *Quantum* **3**, 167 (2019).
[13] Y. R. Sanders, D. W. Berry, P. C. S. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, Compilation of fault-tolerant quantum heuristics for combinatorial optimization, *PRX Quantum* **1**, 020312 (2020).
[14] R. Babbush, J. McClean, C. Gidney, S. Boixo, and H. Neven, Focus beyond quadratic speedups for error-corrected quantum advantage, *arXiv:2011.04149* (2020).
[15] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity, *Phys. Rev. X* **8**, 041015 (2018).
[16] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, *Quantum* **4**, 296 (2020).
[17] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, *arXiv:2007.14460* (2020).
[18] J. Lee, D. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *arXiv:2011.03494* (2020).
[19] S. E. Nagler, D. A. Tennant, R. A. Cowley, T. G. Perring, and S. K. Satija, Spin dynamics in the quantum antiferromagnetic chain compound KCuF_3 , *Phys. Rev. B* **44**, 12361 (1991).
[20] Z. Wang, J. Wu, W. Yang, A. K. Bera, D. Kamenskyi, A. T. M. N. Islam, S. Xu, J. M. Law, B. Lake, C. Wu, and A. Loidl, Experimental observation of Bethe strings, *Nature* **554**, 219 (2018).
[21] T. Fukuhara, P. Schauf, M. Endres, S. Hild, M. Cheneau, I. Bloch, and C. Gross, Microscopic observation of magnon bound states and their dynamics, *Nature* **502**, 76 (2013).
[22] P. R. Pasnoori, N. Andrei, and P. Azaria, Edge modes in one-dimensional topological charge conserving spin-triplet superconductors: Exact results from Bethe ansatz, *Phys. Rev. B* **102**, 214511 (2020).
[23] P. Calabrese and J. Cardy, Time Dependence of Correlation Functions Following a Quantum Quench, *Phys. Rev. Lett.* **96**, 136801 (2006).
[24] H. Bethe, Zur Theorie der Metalle i. Eigenwerte und Eigenfunktionen der linearen Atomkette, *Z. Phys.* **71**, 205 (1931).
[25] E. H. Lieb and F. Y. Wu, Absence of Mott Transition in an Exact Solution of the Short-Range, One-Band Model in One Dimension, *Phys. Rev. Lett.* **20**, 1445 (1968).

- [26] D. C. Mattis, *The Many-Body Problem: An Encyclopedia of Exactly Solved Models in One Dimension (3rd Printing with Revisions and Corrections)* (World Scientific, Singapore, 1993).
- [27] F. H. L. Essler, H. Frahm, F. Göhmann, A. Klümper, and V. E. Korepin, *The One-Dimensional Hubbard Model* (Cambridge University Press, Cambridge, 2005).
- [28] F. Göhmann, M. Karbach, A. Klümper, K. K. Kozłowski, and J. Suzuki, Thermal form-factor approach to dynamical correlation functions of integrable lattice models, *J. Stat. Mech.: Theory Exp.* **2017**, 113106 (2017).
- [29] C. Babenko, F. Göhmann, K. K. Kozłowski, and J. Suzuki, A thermal form factor series for the longitudinal two-point function of the Heisenberg-Ising chain in the antiferromagnetic massive regime, *J. Math. Phys.* **62**, 041901 (2021).
- [30] C. Babenko, F. Göhmann, K. K. Kozłowski, J. Sirker, and J. Suzuki, Exact real-time longitudinal correlation functions of the massive XXZ chain, [arXiv:2012.07378](#) (2020).
- [31] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Simulating physical phenomena by quantum networks, *Phys. Rev. A* **65**, 042323 (2002).
- [32] A. Bohrdt, Y. Wang, J. Koepsell, M. Kánasz-Nagy, E. Demler, and F. Grusdt, Dominant Fifth-Order Correlations in Doped Quantum Antiferromagnets, *Phys. Rev. Lett.* **126**, 026401 (2021).
- [33] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, *Quantum Comput. Quantum Inf.* 305, 53 (2002). [ArXiv:aquant-ph/0005055](#)
- [34] F. Verstraete, J. I. Cirac, and J. I. Latorre, Quantum circuits for strongly correlated quantum systems, *Phys. Rev. A* **79**, 032316 (2009).
- [35] P. Scholl and R. Orús, Kitaev honeycomb tensor networks: Exact unitary circuits and applications, *Phys. Rev. B* **95**, 045112 (2017).
- [36] A. Cervera-Lierta, Exact Ising model simulation on a quantum computer, *Quantum* **2**, 114 (2018).
- [37] K. Robbins and P. J. Love, Benchmarking near-term quantum devices with the variational quantum eigensolver and the Lipkin-Meshkov-Glick model, *Phys. Rev. A* **104**, 022412 (2021).
- [38] P.-L. Dallaire-Demers, M. Stęchły, J. F. Gonthier, N. T. Bashige, J. Romero, and Y. Cao, An application benchmark for fermionic quantum simulations, [arXiv:2003.01862](#) (2020).
- [39] M. J. Cervia, A. B. Balantekin, S. N. Coppersmith, C. W. Johnson, P. J. Love, C. Poole, K. Robbins, and M. Saffman, Exactly solvable model as a testbed for quantum-enhanced dark matter detection, [arXiv:2011.04097](#) (2020).
- [40] R. I. Nepomechie, Bethe ansatz on a quantum computer? *Quantum Inf. Comput.* **32**, 255 (2021).
- [41] W. W. Ho and T. H. Hsieh, Efficient variational simulation of non-trivial quantum states, *SciPost Phys.* **6**, 029 (2019).
- [42] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [43] A. Cervera-Lierta, J. S. Kottmann, and A. Aspuru-Guzik, The meta-variational quantum eigensolver (Meta-VQE): Learning energy profiles of parameterized Hamiltonians for quantum simulation, [arXiv:2009.13545](#) (2020).
- [44] B. Murta and J. Fernández-Rossier, Gutzwiller wave function on a digital quantum computer, [arXiv:2103.15445](#) (2021).
- [45] R. Orbach, Linear antiferromagnetic chain with anisotropic coupling, *Phys. Rev.* **112**, 309 (1958).
- [46] V. E. Korepin, N. M. Bogoliubov, and A. G. Izergin, *Quantum Inverse Scattering Method and Correlation Functions* (Cambridge University Press, Cambridge, 1993).
- [47] M. Takahashi, *Thermodynamics of One-Dimensional Solvable Models* (Cambridge University Press, Cambridge, 1999), 1st ed.
- [48] T. Giamarchi, *Quantum Physics in One Dimension*, The International Series of Monographs on Physics No. 121 (Clarendon, Oxford University Press, Oxford, New York, 2004), oCLC: ocm52784724.
- [49] B. Sutherland, *Beautiful Models: 70 Years of Exactly Solved Quantum Many-Body Problems* (World Scientific, Singapore, 2004).
- [50] M. Gaudin and J.-S. Caux, *The Bethe Wavefunction* (Cambridge University Press, Cambridge, 2014).
- [51] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Inf. Comput.* **12**, 901 (2012).
- [52] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [53] A. Bärttschi and S. Eidenbenz, in *Fundamentals of Computation Theory* (Springer, Cham, 2019) p. 126.
- [54] C. S. Mukherjee, S. Maitra, V. Gaurav, and D. Roy, Preparing Dicke states on a quantum computer, *IEEE Trans. Quantum Eng.* **1**, 3102517 (2020).
- [55] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proc. Natl. Acad. Sci.* **115**, 9456 (2018).
- [56] K. Wan, Exponentially faster implementations of $\text{Select}(H)$ for fermionic Hamiltonians, *Quantum* **5**, 380 (2021).
- [57] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions, *Phys. Rev. A* **98**, 022322 (2018).
- [58] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm, *npj Quantum Inf.* **6**, 00 (2020).
- [59] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2019), oCLC: 1129119093.
- [60] H. Abraham, *et al.*, QISKIT: An open-source framework for quantum computing (2019).
- [61] F. H. L. Essler and M. Fagotti, Quench dynamics and relaxation in isolated integrable quantum spin chains, *J. Stat. Mech.: Theory Exp.* **2016**, 064002 (2016).

- [62] L. Vidmar and M. Rigol, Generalized Gibbs ensemble in integrable lattice models, *J. Stat. Mech.: Theory Exp.* **2016**, 064007 (2016).
- [63] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, *Proc. Natl. Acad. Sci.* **114**, 7555 (2017).
- [64] P. Selinger, Efficient Clifford+ T approximation of single-qubit operators, *Quantum Inf. Comput.* **15**, 159 (2015).
- [65] C. Gidney and A. G. Fowler, Efficient magic state factories with a catalyzed $|CCZ\rangle$ to $2|T\rangle$ transformation, *Quantum* **3**, 135 (2019).
- [66] C. Yesilyurt, S. Bugu, F. Ozaydin, A. A. Altintas, M. Tame, L. Yang, and Ş. K. Özdemir, Deterministic local doubling of W states, *JOSA B* **33**, 2313 (2016).
- [67] 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).
- [68] C. Gidney, General construction of W_n state, <https://quantumcomputing.stackexchange.com/questions/4350/general-construction-of-w-n-state>.
- [69] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv:quant-ph/0001106v1](https://arxiv.org/abs/quant-ph/0001106v1) (2000).
- [70] B. W. Reichardt, In *Proc. 36th ACM Symp. on Theory of Computing (STOC)*, (2004), p. 502.
- [71] R. Van Beeumen, G. D. Kahanamoku-Meyer, N. Y. Yao, and C. Yang, in *HPCAsia2020: Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region* (ACM, New York, 2020), p. 179.
- [72] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, *Phys. Rev. B* **73**, 094423 (2006).
- [73] F. Verstraete, V. Murg, and J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
- [74] J. Mölter, T. Barthel, U. Schollwöck, and V. Alba, Bound states and entanglement in the excited states of quantum spin chains, *J. Stat. Mech.: Theory Exp.* **2014**, P10029 (2014).
- [75] V. Murg, V. E. Korepin, and F. Verstraete, Algebraic Bethe ansatz and tensor networks, *Phys. Rev. B* **86**, 045125 (2012).
- [76] Y. Q. Chong, V. Murg, V. E. Korepin, and F. Verstraete, Nested algebraic Bethe ansatz for the supersymmetric $t-j$ model and tensor networks, *Phys. Rev. B* **91**, 195132 (2015).
- [77] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Simulation of electronic structure Hamiltonians using quantum computers, *Mol. Phys.* **109**, 735 (2011).
- [78] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, Improving quantum algorithms for quantum chemistry, *Quantum Inf. Comput.* **15**, 1 (2015).
- [79] R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, Exponentially more precise quantum simulation of fermions in second quantization, *New J. Phys.* **18**, 033032 (2016).
- [80] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [81] P. Weinberg and M. Bukov, QuSpin: A PYTHON package for dynamics and exact diagonalisation of quantum many body systems part I: Spin chains, *SciPost Phys.* **2**, 003 (2017).