

Quantum simulation of the qubit-regularized $O(3)$ sigma model

Alexander J. Buser

Caltech, Institute for Quantum Information and Matter, Pasadena, CA 91125

Tanmoy Bhattacharya, Lukasz Cincio, and Rajan Gupta

Los Alamos National Laboratory, Theoretical Division, Los Alamos, NM 87545

(Dated: September 7, 2020)

Recently, Singh and Chandrasekharan [1] showed that fixed points of the non-linear $O(3)$ sigma model can be reproduced near a quantum phase transition of a spin model with just two qubits per lattice site. In this paper, we demonstrate how to prepare the ground state of this model and measure a dynamical quantity of interest, the $O(3)$ Noether charge, on a quantum computer. In particular, we apply Trotter [2] methods to obtain results for the complexity of adiabatic ground state preparation in both the weak-coupling and quantum-critical regimes and use shadow tomography [3] to measure the dynamics of local observables. We then present and analyze a quantum algorithm based on non-unitary randomized simulation methods that may yield an approach suitable for intermediate-term noisy quantum devices.

I. INTRODUCTION

Considering the rapid pace of experimental progress in quantum computing, now is the time to think seriously about how quantum simulations of high energy physics will look in the near future. Quantum field theories generally involve infinitely many degrees of freedom per unit physical volume. A crucial step in simulating field theories, then, is choosing a method of truncating the Hilbert space and representing the states of the truncated model using qubits. According to the standard procedure advocated by Wilson [4], one regulates the theory via a hard cutoff, restricting it to a discrete finite lattice. Bosonic theories, however, involve infinitely many quantum states per lattice site even after regularization. Realizing the theory as a quantum simulation on a finite computer requires us to truncate the dimension of the Hilbert space, being careful to preserve the interesting physics of the theory.

One approach is to impose a hard cut-off on the occupation number in the Fourier space of the relevant symmetry group, and directly map the truncated Hamiltonian to one acting on a system of qubits [5]. This approach is provably efficient for Kogut-Susskind type lattice gauge theories [6], and trivially converges to the correct behavior as the local dimension is scaled back to infinity. In the near-term, however, one will be limited to small local dimension with only few qubits per lattice site. Such truncated models are not, in general, good approximations to the original field theory.

We therefore advocate working directly with the standard definition of the continuum field theory as describing the long distance physics near the critical points of a discretized theory. This is particularly interesting for asymptotically-free theories that arise near Gaussian fixed points of the renormalization group evolution. To this end, we construct a lattice Hamiltonian possessing the same symmetries as the desired continuum QFT, but

acting on a Hilbert space with small local dimension. This is the approach taken in Ref. 1 and extended in Appendix C for the case of the $O(3)$ non-linear sigma model. The present paper is concerned with the second step of this process—having found a lattice model in the right universality class, how does one go about simulating it on a quantum computer? We describe the implementation of two interesting and non-trivial tasks, namely ground state preparation and the measurement of dynamic quantities, in particular the $O(3)$ Noether current, and provide estimates of the complexity of each task. While the main focus of the paper concerns the capabilities of fault-tolerant quantum devices, we also consider approximate algorithms that could find applications in intermediate-term noisy quantum computers. Of particular interest to the quantum computing community is an algorithm developed here combining quantum circuit compiling and randomized simulation methods for adiabatic state preparation.

This paper is organized as follows. In Section II we describe the model and discuss the preparation of the ground state. In Section III the measurement of the $O(3)$ Noether charge is discussed. In Section IV we estimate the resources required for ground state preparation in both the weak-coupling and quantum-critical regimes. In Section V we find short-depth circuits approximating the adiabatic state-preparation algorithm and study them numerically. We end with a discussion of our main results in Section VI. Technical details are provided in appendices.

II. EXACT ALGORITHM

In this section, we describe the model studied in Ref. 1 using a slightly more convenient convention. We then describe a quantum circuit for preparing the ground state of this system using the Trotter approximation [2] on a fault-tolerant quantum computer and numerically study

the fidelity of the process. We then discuss the resource requirements for this algorithm.

A. The model

Singh and Chandrasekharan [1] recently demonstrated that, using just two qubits per lattice site, a direct truncation scheme reproduces the Wilson-Fisher fixed point of the $O(3)$ sigma model in two spatial dimensions and the Gaussian fixed point in three spatial dimensions [25]. The model considered there resides on a regular square or cubic lattice of length L in d spatial dimensions. We describe the system in the basis of the total angular momentum J of two spin- $\frac{1}{2}$ degrees of freedom representing the two qubits at each site. For the site at position x , let $|s, x\rangle$ be the $J = 0$ singlet state, and $|m, x\rangle$, $m = -1, 0, +1$ be the $J = 1$ triplet states with angular momentum component along the z -axis having the value m . The Hamiltonian consists of on-site and nearest-neighbor terms.

$$\begin{aligned}
 H &= H_1 + H_2, \\
 H_1 &= \sum_{x,m} (J + \mu m) |m, x\rangle \langle m, x|, \\
 H_2 &= J_r (H_p + H_h), \\
 H_p &= - \sum_{\langle x,x'\rangle, m} (-1)^m |m, x; -m, x'\rangle \langle s, x; s, x'| + \text{h.c.}, \\
 H_h &= \sum_{\langle x,x'\rangle, m} |s, x; m, x'\rangle \langle m, x; s, x'| + \text{h.c.}
 \end{aligned} \tag{1}$$

The choice of the sign of the couplings here is slightly different from that in Ref. 1, owing to a different choice of phases for the basis states.

The model contains three independent coupling constants: the on-site coupling J is the extra energy of the triplet states, μ is the splitting between the triplet states, and J_r is the nearest-neighbor coupling. When J_r and μ are small and J is positive, the vacuum is close to the all-singlet state, and one can think of the $|m\rangle$ as particle excitations with mass J , μ a ‘chemical potential’ corresponding to this particle, H_h the kinetic term, and H_p the pair creation/annihilation interaction.

This model is a qubit-representation of the non-linear $O(3)$ sigma model with a hard cut-off in angular momentum at $l = 1$ (see details in Appendix A). In this work, we study only the $\mu = 0$ case, when the theory has a global $O(3)$ symmetry (see details in Appendix A), and so conserves total angular momentum J and its z -component M . The Hamiltonian also has two more symmetries: the number of sites in state $|s\rangle$ modulo 2, and the parity symmetry. We normalize the Hamiltonian by setting $J = 1$, so that the only free parameter is J_r , hereafter referred to as the coupling constant. Since H_p and H_h separately have all the symmetries of the full model, we need not have the same coefficient for both, so this is actually a choice we are making.

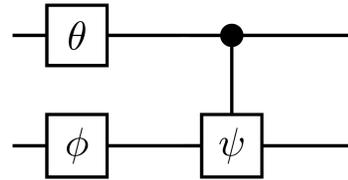


FIG. 1: Circuit performing arbitrary diagonal unitary on two qubits. Operators are z -rotations by the corresponding angle. One degree of freedom is set to zero by the global phase, and the other three uniquely determine θ , ϕ , and ψ .

B. Adiabatic state preparation

The first task of quantum simulation is to prepare initial states for a given Hamiltonian. The Hilbert space of a quantum field theory consists of the closure of the polynomials of particle creation and annihilation operators acting on the vacuum, i.e., the ground state—so preparing ground states is of fundamental importance. In the presence of an energy gap at all values of the couplings considered, we may prepare the ground state via the adiabatic algorithm [7]. Although many algorithms exist for simulating time-evolution with better asymptotic scaling [8–10], we choose to work with standard Trotter methods. This is because the symmetries of the present model make this kind of approach particularly simple to analyze. Furthermore, Trotter formulae are well-suited to the randomized algorithm presented in Section V.

The adiabatic algorithm works by choosing a Hamiltonian whose ground state is easy to prepare, and then slowly tuning a set of coupling constants to reach a target Hamiltonian without exciting the system away from the ground state appreciably throughout the whole computation. This requires the evolution to be slow on the time scale of the mass gap of the system.

Specializing to the model in Eq. (1), when $J_r = 0$ the ground state is the trivial all-singlet state. We define our computational basis in such a way that, for a pair of qubits at a single site, $|00\rangle = |s\rangle$, $|01\rangle = |-1\rangle$, $|10\rangle = |0\rangle$, $|11\rangle = |1\rangle$. Note that, for instance, what we will refer to as a singlet state does not correspond to the singlet state of two physical qubits, but as long as this choice is made consistently throughout our algorithm, we never need to implement this unitary operation explicitly. Starting from the zero-coupling ground state, which, being a product state, is easy to prepare, we reach the ground state at finite coupling, $|\Omega(J_r)\rangle$, by a series of discrete time-steps;

$$|\Omega(J_r)\rangle \approx \prod_{i=1}^N e^{-i\Delta t_i (H_1 + J_{r,i} (H_p + H_h))} |\Omega(0)\rangle. \tag{2}$$

Here we have allowed the length of the time-step Δt_i and the coupling $J_{r,i}$ to change with each iteration, i , and have let N be the number of time-steps. We approximate

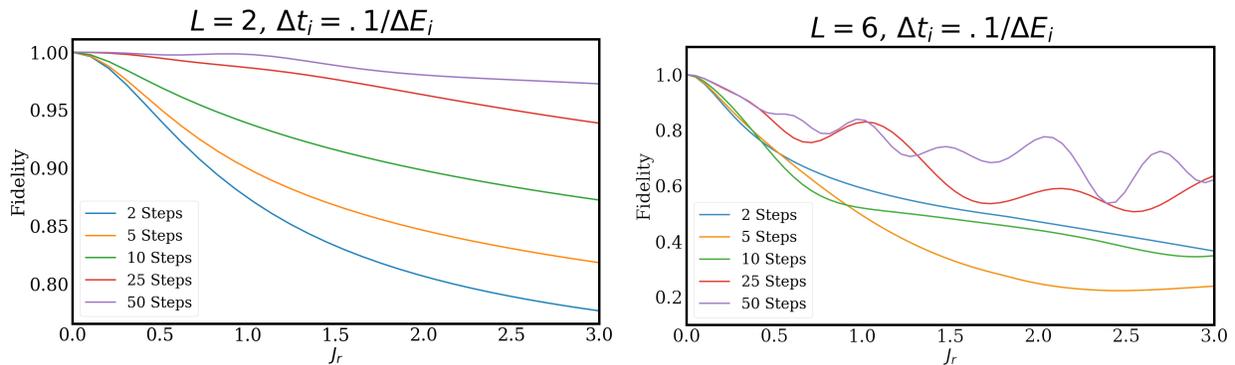


FIG. 2: Accuracy of adiabatic state preparation on two- (left) and six- (right) site system as a function of maximum coupling strength J_r for various numbers of Trotter steps. Here, the time step Δt_i is chosen to be one-tenth of the inverse of the energy gap, ΔE_i , at coupling $J_{r,i}$. Fidelity between the adiabatically-prepared state $|\psi\rangle$ and the true ground state $|\Omega\rangle$ is defined as $\sqrt{\langle\psi|\Omega\rangle\langle\Omega|\psi\rangle}$, where both $|\Omega\rangle$ and $|\psi\rangle$ are properly normalized. Fidelity can be further improved by simultaneously increasing the number of Trotter steps and reducing the time step to minimize non-adiabaticity and Trotterization error.

each time-step by separating out the single-site piece

$$e^{-i\Delta t_i(H_1+J_{r,i}(H_p+H_h))} \approx e^{-i\Delta t_i H_1} e^{-i\Delta t_i J_{r,i}(H_p+H_h)} + \mathcal{O}(L^d \Delta t^2). \quad (3)$$

Since H_1 is diagonal in our computational basis this term is easy to implement. As shown in Fig. 1, for a single site x , $\exp\{-i\Delta t_i(J+\mu m)|m,x\rangle\langle m,x|\}$ may be implemented using just two single-qubit rotations and one controlled rotation even for nonzero μ . Because the adiabatic evolution maintains the translational invariance of the ground state, any possible global phase may be removed consistently from each qubit so that the on-site term is implemented exactly with $3L^d$ gates.

Next, we decompose the nearest-neighbor term into a product of $2d$ non-commuting operators [11]. In one dimension, this involves splitting the links of the lattice into two disjoint sets of even and odd links;

$$e^{-i\Delta t_i J_{r,i}(H_p+H_h)} \approx e^{-i\Delta t_i J_{r,i} H_{even}} e^{-i\Delta t_i J_{r,i} H_{odd}} + \mathcal{O}(L^d \Delta t^2). \quad (4)$$

Here H_{even} and H_{odd} denote H_2 from Eq. (1) restricted to the sites connected by even and odd links, respectively. This is easily generalized to a cubic lattice in arbitrary dimension. Assuming L is even, each term within H_{even} commutes with all the rest (likewise for H_{odd}), since the links are disconnected from each other and H_h and H_p operate on different symmetry sectors (see Appendix C). Now it suffices to simulate the hopping and pair-creation terms on a single pair of adjacent sites. To do this exactly, we can follow an approach similar to Ref. 5 and introduce a unitary operator U_{CG} that implements the Clebsch-Gordan transform on the computational basis. Specifically, this takes a state $|j_1, m_1; j_2, m_2\rangle$ in the local angular momentum basis and gives a state $|J, M, p\rangle$ in the total angular momentum basis [26]. H_p acts nontrivially only on the 2-dimensional $J=0$ sector. Similarly,

the H_h term acts only on the 9-dimensional $J=1$ sector, decomposing into three 2×2 blocks and one 3×3 diagonal block (See Appendix C). Thus, evaluating the nearest-neighbor time-step reduces to implementing four two-qubit unitaries.

For a one-dimensional lattice of small size, the above procedure can be simulated exactly on a classical computer. The accuracy of the output is a complicated function of the adiabatic scheduling, that is, the number of time-steps and the values of J_r and Δt at each step. In Fig. 2 we plot the fidelity as a function of the target coupling J_r for five different numbers of time steps to simulate the adiabatic ground state preparation.

For each value of the maximum coupling, we prepare the ground state of our model at $J_r=0$, and apply the unitary operation in Eq. (2) under the approximations of Eqs. (3) and (4). The couplings are linearly interpolated between 0 and the maximum coupling. We require $\Delta t \Delta E \ll 1$ for adiabaticity, where ΔE is the energy gap of the model, so we set $\Delta t = 0.1/\Delta E$. As expected, the procedure works best for small values of the target coupling, and improves markedly as the number of time steps is increased. For a fixed J_r and number of time steps, however, the fidelity decreases significantly as the lattice size is increased from 2 to 6 as in Fig. 2. This reflects the introduction of Trotter error which increases linearly in the lattice volume, owing to the even-odd decomposition described in Eq. (4).

C. Resource Requirements

The number of quantum gates demanded by the algorithm described in the previous subsection makes it impractical to implement on noisy intermediate-scale quantum computers. Even neglecting the cost of the Clebsch-Gordan transformation, the circuit in Fig. 3 requires six single qubit gates and four controlled unitaries with three qubit control. In terms of arbitrary single-qubit

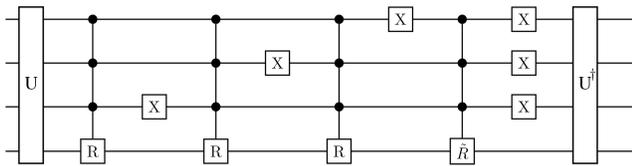


FIG. 3: Circuit implementing time evolution of the single-link nearest neighbor term. R performs the two-level unitary corresponding to the hopping term, while \bar{R} handles the pair creation. The U gate represents the Clebsch-Gordan transformation on four qubits, schematically shown in Fig. 4, ordering the resulting states such that the controlled two-level gates act on the right states. Using the Sleater-Weinfurter construction [12], the three-qubit control gates can be rewritten in terms of single-qubit controls.

and single-controlled-unitaries, this requires a total of 58 gates per time step [12, 13] (see Fig. 3).

In addition, one must implement the Clebsch-Gordan transformation and its inverse. A schematic description of a circuit implementing this transformation was provided by Bacon, Chuang, and Harrow [14] (see Fig. 4), however it requires ancilla qubits and many-qubit control gates. Although the number of gates required to implement a single time-step this way scales linearly with system size, implementing it exactly seems out of reach on current experimental platforms. In Section V we propose a resolution of this problem by finding short-depth circuits that approximate the desired unitary dynamics via numerical optimization.

III. NOETHER CURRENT

One promising application of quantum computing to high energy physics is the potential to measure dynamic quantities in real-time. This requires both the excitation of interesting initial states out of the vacuum (i.e., ground state) and measurement of the quantities of interest. In this study, we focus on the latter problem alone.

As a prototypical example of this, we consider the problem of measuring an interesting local observable which is simple to write down, namely the $O(3)$ Noether current. The z -component of the Noether charge $Q_{z,x}$ simply counts the angular momentum in the z -direction at the site x . Because of the $O(3)$ symmetry of our model, $Q_z = \sum_x Q_{z,x}$ is a conserved quantity. However, there may be interesting physics contained in the local fluctuations of $Q_{z,x}$. To find the current, $J_{z,x}$, of this charge we use the following relation

$$[Q_{z,x}, H] = -i \left(\sum_{x'} J_{z,\langle x,x' \rangle} \right) \quad (5)$$

where $\langle x, x' \rangle$ represent a pair of neighboring sites. Evaluating this commutator explicitly, we obtain the following

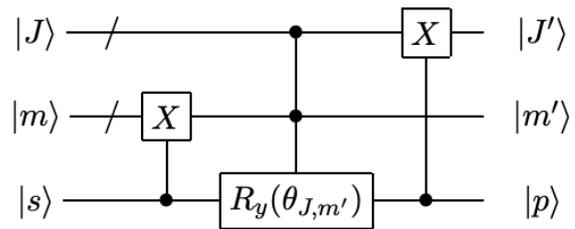


FIG. 4: Schematic description of the Clebsch-Gordan transformation by Bacon, Chuang, and Harrow [14]. Here s represents a single qubit which is added to a system with angular momentum J and m . On the right, J' and m' label the resulting angular momenta, and p keeps track of the pathway through which those values are obtained. The slash denotes a wire containing a register of qubits, the control-X gate adds angular momentum appropriately, and the controlled rotations (one for each J, m') produce the correct amplitudes.

two-site operator for $J_{z,\langle x,x' \rangle}$:

$$J_{z,\langle x,x' \rangle} = iJ_r \sum_m m (|sm\rangle \langle ms| - |ms\rangle \langle sm| + |ss\rangle \langle m-m| - |-mm\rangle \langle ss|), \quad (6)$$

where for convenience we have suppressed the position labels and used a compressed notation: $|pq\rangle \equiv |p,x\rangle |q,x'\rangle$. On a large scale fault-tolerant device, one could efficiently compute the expectation value of this observable using the standard technique of phase estimation [15]. This would allow one to determine the Noether current at a single site to precision ε using $\mathcal{O}(\log(1/\varepsilon))$ ancilla qubits and $\mathcal{O}(1/\varepsilon)$ controlled-U gates, where $U = \exp(iJ_{z,\langle x,x' \rangle})$, so that the Noether current everywhere can be computed in time $\mathcal{O}(\mathcal{V}/\varepsilon)$, where the number of links is $\mathcal{V} = dL^d$.

For a noisy device with a limited number of qubits, however, the need for a large number of ancilla qubits and a circuit depth linear in system size makes applying full-scale phase estimation somewhat impractical. A recent algorithm by Huang et al. applies a technique known as shadow tomography to predict expectation values of M local observables using measurements on $\log(M)$ copies of a quantum state [3, 16]. By measuring a quantum state in a sequence of random bases, one can construct an efficient classical representation of the state, which suffices for predicting expectation values of a set of local observables. For a local observable like the Noether current, obtaining a precision δ uses only $\mathcal{O}(\log \mathcal{V}/\delta^2)$ independent copies of the system. This protocol, free from the need for ancilla qubits and long coherence times, appears to be more suitable for near-term experiments than phase estimation.

Concretely, we measure the time-dependence of local observables like the Noether current through several steps.

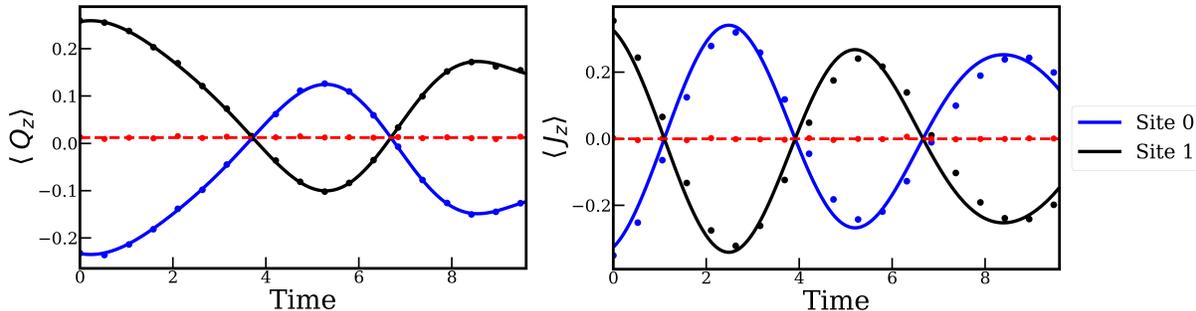


FIG. 5: Time-dependence of the expectation value of the z -component of the $O(3)$ Noether charge (left) and current (right) predicted using shadow tomography for a two-site system, prepared in a random initial state, with $J_r = 0.1$. Solid lines indicate the exact charge or current on each site while points show the values predicted using shadow tomography. Red dashed line and dots show the exact and predicted values, respectively, of the total charge or current across both sites. This is expected to be a constant for the charge and zero for the current. Since Q_z has range $k = 1$, the results converge more rapidly than suggested by Eq. (8). Here we have used $N = 1,000,000$ and $5,000,000$ random Pauli measurements for the charge and current respectively with $n = 10$ per time-step ($\delta \approx 0.027$). This implies $\varepsilon \approx 3.7 \times 10^{-3}$ and ≈ 0.030 for the charge and current respectively. Note that, even though we have used more measurements for the current, the prediction is less accurate than that for the charge owing to the four-fold increase in the range of the observable.

1. Prepare the desired initial state.
2. Simulate time-evolution for time t .
3. Perform a random unitary operation (this could be a random Clifford [27] circuit as in Ref. 3), then measure the state in the computational basis.
4. Repeat steps 1–3 N times. Use the results to predict, through a median-of-means estimator described below, the measurement outcome of the observable at time t to precision ε . It is known that one requires $N = \mathcal{O}(\log(\mathcal{V}/\varepsilon^2))$.
5. Repeat steps 1–4, sweeping t over a range of values from t_{initial} to t_{final} .

For the present case we choose a random unitary ensemble composed of single-qubit Clifford circuits. This choice is equivalent to measuring the state in a random Pauli basis at each step. Then the results of Ref. 3 tell us that to estimate the expectation value of a local observable \mathcal{O} acting non-trivially on only k qubits (at all sites on a lattice of volume \mathcal{V}) to within ε of its true value with probability $1 - \delta$, it suffices to repeat steps 1–3 N times, where

$$N = (2\log(2\mathcal{V}/\delta)) \frac{34}{\varepsilon^2} 4^k \|\mathcal{O}\|_\infty^2. \quad (7)$$

For $\mathcal{O} = J_{z,x}$ in one spatial dimension, we have

$$N = (2\log(2\mathcal{V}/\delta)) \frac{8704}{\varepsilon^2} |J_r|^2. \quad (8)$$

For instance, estimating the current on a two-site system to two decimal places with success probability 90% ($\delta = 0.1$, $\varepsilon = 0.01$, $k = 4$) is guaranteed provided

$N \gtrsim (5 \times 10^7) |J_r|^2$. Since each repetition can be performed in parallel, such large values of N pose no technical difficulties.

The precision parameters ε and δ are set independently via a simple median-of-means protocol. That is, we construct N classical representations of the initial state using the procedure outlined above, split them into n groups of N/n shadows each, and average over each group. This gives a set of n classical states $\{\rho_1, \dots, \rho_n\}$. We return the median expectation value, $\text{Median}(\{\text{Tr}(\mathcal{O}\rho_1), \dots, \text{Tr}(\mathcal{O}\rho_n)\})$, of the desired observable over this set, yielding rigorous guarantees on δ and ε as in Ref. 3;

$$\begin{aligned} n &= 2\log(2\mathcal{V}/\delta) \\ N &= \frac{34n}{\varepsilon^2} 4^k \|\mathcal{O}\|_\infty^2. \end{aligned} \quad (9)$$

As a demonstration of principle, we compute the time-evolution of the charge Q_z and its current on a two-site system prepared in a random initial pure state, that is, one whose complex-valued entries are identically and independently distributed according to a normal distribution. The results are shown in Fig. 5. In both cases, the results appear significantly better than the rigorous performance guarantees suggested by Eq. (9). The errors in the current are larger than those in the charge owing to the fact that the charge is a sum of single qubit operators, whereas the current acts on four qubits at a time. The closeness of the predicted values to the exact results suggest that this method may be an effective way to measure dynamic quantities in lattice models of quantum field theories on a quantum computer.

IV. ERROR ANALYSIS

Assuming perfect gate fidelity, there are two sources of errors we need to consider—Trotterization and violations of adiabaticity. Trotter error, which is $\mathcal{O}(\Delta t^2)$ in a first-order approximation, can be reduced at the cost of larger circuit depth by using a Trotter-Suzuki approximation of sufficiently high order, however we choose to work with a first-order approximation for several reasons. First, the algorithm of Section V was designed only to second order in the Trotter step size Δt , so any improvement via higher-order formulae would be lost. Second, the field theory of interest emerges at a quantum critical point and depends primarily on the symmetries preserved by the Hamiltonian, rather than its exact form. Under reasonable assumptions, the errors induced by Trotterization reside in the algebra of operators invariant under the same symmetries as the original Hamiltonian. Thus, treating Trotterized time-evolution as exact evolution under an effective Hamiltonian by resumming the Baker-Campbell-Hausdorff formula, the effective Hamiltonian will likely reside in the same universality class as the original [17]. In that case, close to a critical point, the effective Hamiltonian and the target field theory differ only in a set of irrelevant UV operators which should not affect the physical content of the lattice theory.

We mention, in passing, that throughout this work we use the fidelity of the prepared quantum state as a measure of accuracy. This is, however, known to be too strong. High fidelity implies that all observables are close to their desired values, but the continuum field theory cares only about products of local operators at separations scaling as the divergent correlation length. In other words, the relevant measure of fidelity should be calculated after tracing out the ultraviolet degrees of freedom. We ignore this subtlety in this work.

A. Weak-Coupling Limit

In this section we quantify errors due to Trotterization and violation of the adiabatic condition within the weak coupling approximation $J_r \ll 1$.

First, we will consider errors due to non-adiabaticity. Let $\tilde{H}(t)$ be a time-varying Hamiltonian where $0 \leq t \leq T$, $H(\tau) = \tilde{H}(t/T)$, and the energy gap of $H(\tau)$ is $\Delta(\tau)$. A theorem due to Teufel [18] says that as long as the total time of the adiabatic evolution T satisfies

$$T \geq \frac{4}{\varepsilon} \left[\frac{\|\dot{H}(0)\|}{\Delta(0)^2} + \frac{\|\dot{H}(1)\|}{\Delta(1)^2} + \int_0^1 d\tau \left(10 \frac{\|\dot{H}(\tau)\|^2}{\Delta(\tau)^3} + \frac{\|\ddot{H}(\tau)\|^2}{\Delta(\tau)^3} \right) \right], \quad (10)$$

then the final state is within ε of the true ground state. We can schedule the adiabatic evolution by linearly in-

terpolating the coupling from its initial to its final value, using a series of discrete time steps Δt_i where Δt_i is set by the inverse spectral gap. According to our perturbative estimate of the spectral gap (Appendix B),

$$\Delta t_i := 1/\Delta_i = 1/(1 - 2dJ_{r,i} + 3dJ_{r,i}^2). \quad (11)$$

The linear interpolation ensures that $\ddot{H}(\tau) = 0$. Since the on-site term is time-independent, it does not factor into the adiabatic condition. Then, defining $H_{h,\langle x,x' \rangle}$ as the hopping term H_h from Eq. (1) restricted to the lattice sites x, x' (and similarly for H_p), we have

$$\dot{H}(\tau) = J_{r,\max} \sum_{\langle x,x' \rangle} (H_{p,\langle x,x' \rangle} + H_{h,\langle x,x' \rangle}), \quad (12)$$

where $J_{r,\max}$ is the final desired value of J_r and

$$\begin{aligned} \|\dot{H}(\tau)\| &= \|J_{r,\max} \sum_{\langle x,x' \rangle} (H_{p,\langle x,x' \rangle} + H_{h,\langle x,x' \rangle})\| \\ &\leq J_{r,\max} \sum_{\langle x,x' \rangle} \|H_{p,\langle x,x' \rangle} + H_{h,\langle x,x' \rangle}\| \\ &\propto \mathcal{V} J_{r,\max}. \end{aligned} \quad (13)$$

In this analysis, we ignored the fact that we change the Hamiltonian along a staircase approximation to the linear function. Assuming that $\Delta(\tau)$ decreases monotonically with τ , a standard result from real analysis lets us bound the error induced by replacing the integral in Eq. (10) with a sum as

$$\begin{aligned} 10\Delta t_{\max} \left(\frac{\|\dot{H}(0)\|^2}{\Delta(0)^3} - \frac{\|\dot{H}(1)\|^2}{\Delta(1)^3} \right) \\ = \mathcal{O}(\mathcal{V}^2 J_{r,\max}^3 \Delta t_{\max}) \end{aligned} \quad (14)$$

where we have defined $\Delta t_{\max} = \max_i \{\Delta t_i\}$. To second order in $J_{r,\max}$, we can obtain the mass gap perturbatively and choose

$$\Delta t_i = \left(1 - 2d \left(\frac{iJ_{r,\max}}{N} \right) + 3d \left(\frac{iJ_{r,\max}}{N} \right)^2 \right)^{-1}. \quad (15)$$

Since for small $J_{r,\max}$ the spectrum is completely gapped, Δt_i is bounded by a constant. Therefore, the leading discretization error is of order $\mathcal{V}^2 J_{r,\max}^3$.

To summarize, to second order in J_{\max} , the adiabatic condition becomes

$$T = \mathcal{O} \left(\frac{J_{r,\max}^2 \mathcal{V}^2}{\varepsilon} + \mathcal{V}^2 J_{r,\max}^3 \right) \quad (16)$$

$$= \mathcal{O} \left(\frac{J_{r,\max}^2 \mathcal{V}^2}{\varepsilon} \right). \quad (17)$$

This result holds so long as the energy gap is bounded from below by a constant. For fault-tolerant devices,

the limiting factor is then the adiabatic evolution near a phase transition where the gap shrinks to zero. In that case, knowledge of the critical exponent can determine the optimal scheduling as in Ref. 19.

Having obtained an upper bound on the error due to non-adiabaticity, we now include that arising from Trotterization. Suppose there are N time steps and the coupling is linearly interpolated between 0 and $J_{r,\max}$. Since each time-evolution operator has leading error corrections of order $\mathcal{V}\Delta t_i^2$ in a first-order Trotter approximation, the full evolution is valid to order $\mathcal{V}\sum_i \Delta t_i^2$. Naturally, this sum is bounded by $\mathcal{V}T^2$, so demanding that the Trotterization error is of order ε constrains the total time simulated as

$$T = \mathcal{O}\left(\sqrt{\frac{\varepsilon}{\mathcal{V}}}\right). \quad (18)$$

Naïvely, we assign as the total error the sum of those induced by Trotterization and non-adiabaticity, so that the total error is of order ε provided that

$$T = \mathcal{O}\left(\frac{(J_{r,\max}\mathcal{V})^2}{\varepsilon} + \sqrt{\frac{\varepsilon}{\mathcal{V}}}\right). \quad (19)$$

The total run-time is

$$T = \sum \Delta t_i = \mathcal{O}(N), \quad (20)$$

so equivalently we can write

$$N = \mathcal{O}\left(\frac{(J_{r,\max}\mathcal{V})^2}{\varepsilon} + \sqrt{\frac{\varepsilon}{\mathcal{V}}}\right). \quad (21)$$

Previously we showed that a single time-step can be implemented exactly with a constant number of gates per lattice site, so, multiplying by the total number of links, the total time-complexity of our algorithm is

$$\mathcal{O}\left(\frac{J_{r,\max}^2\mathcal{V}^3}{\varepsilon} + \sqrt{\varepsilon\mathcal{V}}\right) \quad (22)$$

in arbitrary spatial dimension. We will continue to display the subdominant term in this expression to keep track of the Trotterization error in comparison to the adiabaticity error.

B. Quantum-Critical Regime

The estimates in the previous section describe the resource requirements for preparing the ground-state of the discretized theory in the weak-coupling limit. Our goal, however, is to probe the physics of a quantum field theory, the $O(3)$ sigma model, which emerges only in the long-distance limit as the correlation length in the theory diverges. In two and three spatial dimensions the model in Eq. (1) resides in the universality class

of the $O(3)$ non-linear sigma model, meaning that, in infinite volumes, it undergoes a quantum phase transition at some critical value of the coupling $J_{r,c}$. Accordingly, the mass gap $\Delta(J_r)$ scales as $|J_r - J_{r,c}|^\nu$, where $J_{r,c}$ and ν are the critical coupling and exponent, respectively. For our model, we have efficient classical procedures for determining the critical parameters [1]: $\nu = 0.693(15)$ in (2+1)-dimensions and $\nu = 0.5050(96)$ in (3+1)-dimensions. The critical values of the coupling are $J_{r,c} = 4.81695(37)$ in (2+1)-dimensions and $J_{r,c} = 10.09817(55)$ in (3+1)-dimensions.

To estimate the time required to prepare the ground state at coupling J_r close to $J_{r,c}$, we apply the same methods as in Section IV A, starting from Eq. (10). We find that the error due to non-adiabaticity is less than ε provided that the total simulated time obeys

$$T \geq \frac{4\mathcal{V}\|J_r(H_p + H_h)\|}{\varepsilon} (1 + \alpha|J_r - J_{r,c}|^{-2\nu} + 10\mathcal{V}\|J_r(H_p + H_h)\| \int_0^1 ds \Delta(s)^{-3}), \quad (23)$$

where $\Delta(\tau) \sim \alpha|J_r - J_{r,c}|^\nu$ for $\tau \approx 1$, the linearly interpolated Hamiltonian is

$$H(s) = \sum_x H_{1,x} + J_r s \sum_{\langle x,x' \rangle} (H_{p,\langle x,x' \rangle} + H_{h,\langle x,x' \rangle}) \quad (24)$$

and s ranges from 0 to 1.

Because the gap approaches zero as J_r approaches $J_{r,c}$, we assume that $\Delta(s)$ is bounded from below by $\Delta(1)$. Since the norm of the nearest-neighbor Hamiltonian $\|(H_{p,\langle x,x' \rangle} + H_{h,\langle x,x' \rangle})\|$ on a single pair of sites is of unit order, we can lower bound the total required simulated time T in our algorithm as

$$T \geq \frac{4\mathcal{V}|J_r|}{\varepsilon} (1 + \alpha|J_r - J_{r,c}|^{-2\nu} + 10\mathcal{V}|J_r|\|J_r - J_{r,c}\|^{-3\nu}). \quad (25)$$

For J_r close to the critical value, then, the dominant contribution is

$$T = \mathcal{O}\left(\frac{|J_r\mathcal{V}|^2}{\varepsilon|J_r - J_{r,c}|^{3\nu}}\right). \quad (26)$$

Trotterization again requires that $T = \mathcal{O}(\sqrt{\varepsilon/\mathcal{V}})$, so the total time-complexity of adiabatic state preparation in the quantum-critical regime in arbitrary spatial dimension is

$$\mathcal{O}\left(\frac{|J_r|^2\mathcal{V}^3}{\varepsilon|J_r - J_{r,c}|^{3\nu}} + \sqrt{\varepsilon\mathcal{V}}\right). \quad (27)$$

In particular, this result implies that we can prepare the ground state of our model at coupling J_r close to the critical coupling $J_{r,c}$ in time polynomial in the distance to the critical point, lattice volume, and precision ε . The scaling of the time-complexity with the critical exponent ν

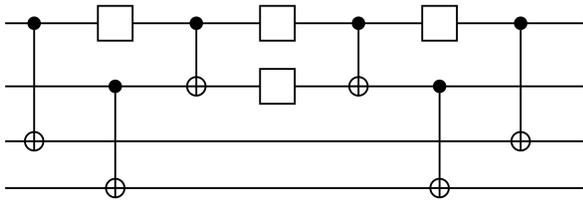


FIG. 6: Circuit for implementing nearest-neighbor time evolution step using ten gates. The single qubit operations are optimized over all possible unitary gates.

follows directly from the adiabatic condition in Eq. (10), and is independent of the spatial dimension. For infinite lattice size and J_r sufficiently close to $J_{r,c}$, physical properties of this ground state correspond to those of the vacuum state of the non-linear $O(3)$ sigma model.

V. QUANTUM COMPILING AND RANDOMIZATION

Current quantum devices are severely limited by noisy gate implementations and short decoherence times. As a result, real near-term simulations demand circuits which can be implemented in time on the order of the qubit lifetime, strongly favoring short-depth circuits. Since, as in Section II, implementing the time-evolution operator exactly using the circuit we provided would require ancilla qubits and at least 58 gates per time-step, realizing the time-evolution operator of our model does not appear to be feasible for a near term device. Therefore, we should consider alternative approximate methods.

In order to find a short-depth circuit approximating the time-evolution operator for the nearest-neighbor Hamiltonian in Eq. (1), we use a classical simulation of quantum-assisted quantum compiling (QAQC), originally proposed by Khatri et al. [20]. Just as a compiler translates high-level source code to a low-level version which can be interpreted by a computer processor, quantum compiling seeks to translate a representation of a unitary, for instance, its matrix representation in a fixed basis, to a sequence of gates executable on a quantum device.

We fix an allowed set of gates specific to a device and set a desired circuit depth, then apply standard optimization techniques to obtain an approximation of the desired quantum circuit. Here, we allow only CNOT and arbitrary single-qubit gates, as in IBM’s QX architecture. The cost function is taken to be the square of the Frobenius norm of the difference between the exact operator and its approximation.

Performing this protocol for our model, we obtain 10, 20, and 30-gate circuits composed only of CNOTs and single-qubit gates (see Fig. 6) approximating the time evolution operator on a single pair of sites,

$$U(J_r \Delta t) = e^{iJ_r \Delta t (H_h + H_p)}, \quad (28)$$

for a fixed value of the coupling constant and time step ($J_r = 0.04$, $\Delta t = 0.2$). In each case, the optimizer was allowed to run for up to 48 CPU hours. One could repeat this procedure for each value of the coupling needed, however, for simulations involving many steps, this would be computationally demanding.

Instead, we run the procedure a single time and interpolate with a randomized approach [21]. By performing a probabilistically scheduled sequence of gates, the dynamics of the quantum state is represented as a non-unitary super-operator, also known as a quantum channel, mapping density matrices to density matrices. The aim is to produce a mixed state that still has a very high overlap with the desired ground-state, so that observables and expectation values are reproduced with little error.

The specific random sequence of gates is determined as follows. First we determine U_{approx} as the short-depth circuit best approximating $U(J_r \Delta t)$ for the largest $J_r \Delta t$ of interest. Then, for adiabatic state preparation with N steps, and coupling $J_{r,i}$ and time step $\Delta t(J_{r,i})$ at the i ’th step, let

$$p_i = \frac{J_{r,i} \Delta t(J_{r,i})}{\max_i (J_{r,i} \Delta t(J_{r,i}))}. \quad (29)$$

At each time step, we apply the identity circuit with probability $1 - p_i$ and U_{approx} with probability p_i . This has the advantage of optimality for fixed circuit depth at $p_i = 0, 1$. Since the randomized circuit is *not* a unitary operation, one needs to modify our measure of fidelity. We choose to use the channel fidelity, namely the diamond norm of the difference between the approximate and exact circuits. The diamond norm is a measure of how hard it is to distinguish two quantum channels with a single measurement. Specifically, if $\Phi, \Phi' \in \mathcal{L}(H \otimes H) \rightarrow \mathcal{L}(H \otimes H)$ are two unital quantum channels, their diamond distance is

$$\|\Phi - \Phi'\|_{\diamond} = \max_{\rho} \text{Tr}(\Phi \otimes I)(\rho) - (\Phi' \otimes I)(\rho), \quad (30)$$

where ρ is a density matrix on the k -qubit Hilbert space operated on by the circuit, with the addition of k possibly entangled ancilla qubits not acted on by the circuit. We use semidefinite programming to perform the maximization in Eq. (30) for each circuit. The results are shown in Fig. 7, Fig. 8 and Fig. 9.

A. Analysis of randomized algorithm

Consider an arbitrary Hamiltonian H which is a function of a set of coupling constants λ_i . Suppose we can implement $U(\lambda_i, \Delta t) = e^{-iH(\lambda_i)\Delta t}$ exactly for two sets of values for the couplings and time step; $\lambda_i(0), \Delta t(0)$ and $\lambda_i(1), \Delta t(1)$. In order to find the ground state of $H(\lambda_i(1))$ using adiabatic state preparation (or similar quantum quench), one could prepare the system in the ground state of $H(\lambda_i(0))$ and linearly interpolate the cou-

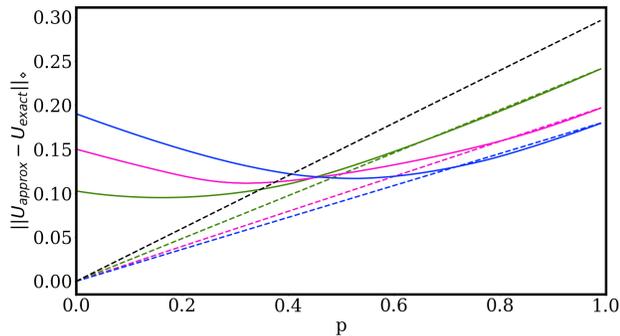


FIG. 7: Diamond distance, given by Eq. (30), between the exact time-evolution operator $e^{iJ_r\Delta t(J_r)(H_p+H_h)}$ and its approximation by short-depth unitaries, on four qubits as a function of $p = \frac{J_r\Delta t(J_r)}{\max(J_r\Delta t(J_r))}$, shown here for $\max(J_r\Delta t(J_r)) = 0.1$. Curves are shown for 10 (blue), 20 (magenta), and 30 (green) gate circuits. Solid lines indicate deterministically applying the approximate circuit at each time-step. Dashed lines correspond to applying the identity with probability $1-p$ and the approximate circuit with probability p . The black dashed line shows the identity circuit as a benchmark for performance. Randomized approach appears to work better than deterministic one for all circuit depths, since the randomized circuit is closer to the exact time-evolution operator in diamond distance.

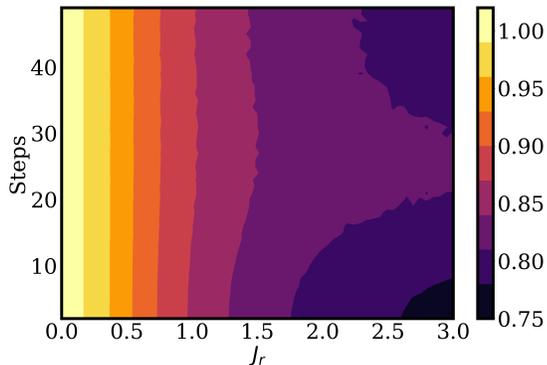


FIG. 8: Fidelity of adiabatic state preparation on four qubits for several numbers of Trotter steps using the short-depth probabilistic method with circuit depth 10. Fidelity is averaged over 10 repetitions. There appears to be only slight variation in the accuracy as the number of steps is increased. Note that the fidelity does not increase monotonically with the number of steps. This is because the error from quantum circuit compiling grows additively with the number of steps, eventually out-pacing whatever gains are made by reducing Trotter error through more steps.

plings between $\lambda_i(0)$ and $\lambda_i(1)$.

Now we form the quantum channel

$$\Lambda_t(\rho) = (1-p)U(0)\rho U^\dagger(0) + pU(1)\rho U^\dagger(1), \quad (31)$$

where $U(t)$ is shorthand for $U(\lambda_i(t), \Delta t(t))$. Letting $\tilde{\lambda}_i(t) = \lambda_i(t)\Delta t(t)$ and λ be any of the various couplings,

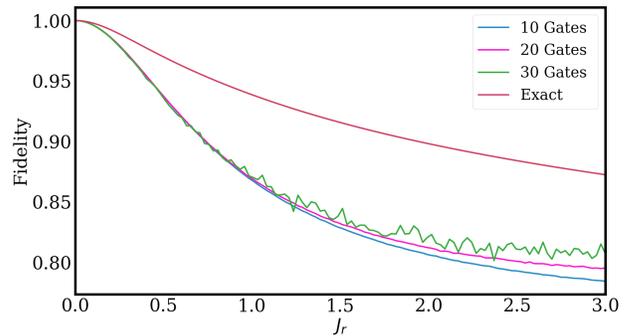


FIG. 9: Adiabatic state preparation on four qubits using probabilistic method with 10, 20, and 30-gate circuits, and exact nearest-neighbor time-evolution operator. For the curve labeled 'exact,' deviations from one are due entirely to Trotterization and non-adiabatic errors. Results are averaged over 100 repetitions, using 10 Trotter steps. Results for short-depth circuits are slightly worse than an exact Trotterized approach, with small improvements through larger depth. We expect the short-depth circuits to perform better in the presence of noise.

we set for the probability

$$p(t) = \frac{\tilde{\lambda}(t) - \tilde{\lambda}(0)}{\tilde{\lambda}(1) - \tilde{\lambda}(0)}. \quad (32)$$

Expanding Eq. (31) in powers of $\Delta t(t)$, we obtain

$$\Lambda_t(\rho) = U(t)\rho U^\dagger(t)(1 + \mathcal{O}(\Delta t^2)). \quad (33)$$

Thus, we see that the time-dependent dynamics is reproduced to the same order in Δt as in first-order Trotterization, implying that the randomized scheme does not change the asymptotic scaling of the time-complexity of our algorithm.

If, as in the present case, we replace the exact implementation of $U(1)$ with an approximate unitary $\bar{U}(1) = U(1)(1 + \mathcal{O}(\delta))$, the result becomes

$$\Lambda_t(\rho) = U(t)\rho U^\dagger(t)(1 + \mathcal{O}(\Delta t^2, \delta)). \quad (34)$$

In this way, it suffices to compile a constant number of circuits and interpolate between them using randomized simulation, rather than compiling a number of circuits growing linearly with the number of time steps.

VI. CONCLUSION

Quantum computing has enormous potential for future investigations in high energy physics. In this paper, we focused on the task of simulating a qubit-regularized version of the non-linear $O(3)$ sigma model, in particular, preparing its ground state and measuring dynamic quantities. We found that, in d -dimensions, for a lattice of size $\mathcal{V} = dL^d$, the ground state may be prepared near the quantum critical point to precision ε in

time $\mathcal{O}\left(\frac{\mathcal{V}^3|J|^2}{\varepsilon|J-J_c|^{3\nu}} + \sqrt{\varepsilon\mathcal{V}}\right)$ using ordinary Trotter methods and provided an explicit circuit representation. We described how to use shadow tomography to measure the time-dependent $O(3)$ Noether current efficiently on a near-term device to precision δ in time $\mathcal{O}(\log(\mathcal{V}/\delta^2))$. Lastly, we performed numerical experiments simulating a heuristic algorithm for obtaining short-depth circuits approximating adiabatic ground state preparation for applications on intermediate-term devices, as well as an improved version implementing techniques from randomized quantum simulation.

Acknowledgment

We would like to thank R. Somma, S. Chandrasekharan, H. Singh and J. Preskill for helpful discussions.

TB and RG were funded under Department of Energy (DOE) Office of Science High Energy Physics Contract #89233218CNA000001. LC was supported by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20190065DR. LC was also supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under the Quantum Computing Application Teams program. AJB also acknowledges support from the Los Alamos Quantum Computing Summer School (QCSS) and its organizers.

-
- [1] H. Singh and S. Chandrasekharan, Phys. Rev. D **100**, 054505 (2019), doi:[10.1103/PhysRevD.100.054505](https://doi.org/10.1103/PhysRevD.100.054505).
- [2] H. F. Trotter, Proc. Amer. Math. Soc. **10**, 545 (1959), doi:[10.1090/S0002-9939-1959-0108732-6](https://doi.org/10.1090/S0002-9939-1959-0108732-6).
- [3] H.-Y. Huang, R. Kueng, and J. Preskill, (2020), [arXiv:2002.08953](https://arxiv.org/abs/2002.08953) [quant-ph].
- [4] K. G. Wilson, p. 45 (1974), doi:[10.1103/PhysRevD.10.2445](https://doi.org/10.1103/PhysRevD.10.2445).
- [5] T. Byrnes and Y. Yamamoto, Phys. Rev. A **73**, 022328 (2006), doi:[10.1103/PhysRevA.73.022328](https://doi.org/10.1103/PhysRevA.73.022328).
- [6] J. B. Kogut and L. Susskind, Phys. Rev. D **11**, 395 (1975), doi:[10.1103/PhysRevD.11.395](https://doi.org/10.1103/PhysRevD.11.395).
- [7] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, (2000), [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [8] D. W. Berry, A. M. Childs, and R. Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 792–809, 2015.
- [9] A. M. Childs, A. Ostrander, and Y. Su, Quantum **3**, 182 (2019), doi:[10.22331/q-2019-09-02-182](https://doi.org/10.22331/q-2019-09-02-182).
- [10] J. Haah, M. Hastings, R. Kothari, and G. H. Low, Quantum algorithm for simulating real time evolution of lattice hamiltonians, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 350–360, 2018.
- [11] A. M. Childs and Y. Su, (2019), [arXiv:1901.00564](https://arxiv.org/abs/1901.00564) [quant-ph].
- [12] T. Sleator and H. Weinfurter, Phys. Rev. Lett. **74**, 4087 (1995), doi:[10.1103/PhysRevLett.74.4087](https://doi.org/10.1103/PhysRevLett.74.4087).
- [13] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, New York, NY, USA, 2011).
- [14] D. Bacon, I. L. Chuang, and A. W. Harrow, Phys. Rev. Lett. **97**, 170502 (2006), doi:[10.1103/PhysRevLett.97.170502](https://doi.org/10.1103/PhysRevLett.97.170502).
- [15] A. Yu. Kitaev, (1995), [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026) [quant-ph].
- [16] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, p. 325338, New York, NY, USA, 2018, Association for Computing Machinery, doi:[10.1145/3188745.3188802](https://doi.org/10.1145/3188745.3188802).
- [17] A. J. Buser, T. Bhattacharya, and S. Chandrasekharan, Trotterization and universality in quantum simulation of quantum field theories, 2020, In Preparation.
- [18] S. Teufel, Lecture Notes in Mathematics, Berlin Springer Verlag **1821** (2003), doi:[10.1007/b13355](https://doi.org/10.1007/b13355).
- [19] S. P. Jordan, K. S. M. Lee, and J. Preskill, Science **336**, 1130 (2012), doi:[10.1126/science.1217069](https://doi.org/10.1126/science.1217069).
- [20] S. Khatri *et al.*, Quantum **3**, 140 (2019), doi:[10.22331/q-2019-05-13-140](https://doi.org/10.22331/q-2019-05-13-140).
- [21] E. Campbell, Phys. Rev. Lett. **123**, 070503 (2019), doi:[10.1103/PhysRevLett.123.070503](https://doi.org/10.1103/PhysRevLett.123.070503).
- [22] H. Singh, (2019), [arXiv:1911.12353](https://arxiv.org/abs/1911.12353) [hep-lat].
- [23] NuQS Collaboration, A. Alexandru, P. F. Bedaque, H. Lamm, and S. Lawrence, Phys. Rev. Lett. **123**, 090501 (2019), doi:[10.1103/PhysRevLett.123.090501](https://doi.org/10.1103/PhysRevLett.123.090501).
- [24] E. Zohar and M. Burrello, Phys. Rev. D **91**, 054506 (2015), doi:[10.1103/PhysRevD.91.054506](https://doi.org/10.1103/PhysRevD.91.054506).
- [25] This method can be easily generalized to the $O(N)$ nonlinear sigma model, see Ref. 22.
- [26] In order to make this unitary we also need to keep track of the channel p that produced the $|J, M, p\rangle$ state; for example, $J = 1$ can arise out of the three channels $|m, x; m, x'\rangle$, $|m, x; s, x'\rangle$ and $|s, x; m, x'\rangle$.
- [27] The Clifford group is useful for many concepts in quantum information. The n -qubit Clifford group C_n is defined as the set of unitaries which normalize the n -qubit Pauli group P_n . That is, if $U \in U(2^n)$, then $U \in C_n$ if for every $V \in P_n$ there is a $W \in P_n$ such that $UVU^\dagger = W$.

Appendix A: Representation Theory

The Hamiltonian for the lattice-regulated $O(3)$ nonlinear sigma model is essentially that of an $O(3)$ rotor

model;

$$\hat{H} = J_1 \sum_i \frac{\hat{L}_i^2}{2} - J_2 \sum_{\langle i,j \rangle} \hat{\phi}_i \cdot \hat{\phi}_j := \hat{H}_L + \hat{H}_\phi, \quad (\text{A1})$$

where i and j are nearest neighbor sites on a Euclidean spatial lattice in d -dimensions, $\vec{\phi}_i$ is a unit 3-vector associated to site i and \hat{L}_i is the angular momentum. In spherical coordinates the dot product of two unit Euclidean 3-vectors can be expressed as

$$\vec{\phi}_1 \cdot \vec{\phi}_2 = \sin(\theta_1) \sin(\theta_2) \cos(\phi_1 - \phi_2) + \cos(\theta_1) \cos(\theta_2). \quad (\text{A2})$$

According to the Peter-Weyl theorem, the space of square integrable functions on the unit sphere is isomorphic to the direct sum of all unitary irreducible representations of $SO(3)$. This is precisely the content of spherical harmonic analysis. In terms of spherical harmonics $Y_l^m(\theta, \phi)$,

$$\begin{aligned} \vec{\phi}_i \cdot \vec{\phi}_j &= \frac{4\pi}{3} (Y_1^0(\theta_i, \phi_i) Y_1^0(\theta_j, \phi_j) \\ &\quad - Y_1^1(\theta_i, \phi_i) Y_1^{-1}(\theta_j, \phi_j) \\ &\quad - Y_1^{-1}(\theta_i, \phi_i) Y_1^1(\theta_j, \phi_j)). \end{aligned} \quad (\text{A3})$$

These spherical harmonics may be recast as Hermitian operators acting on a Hilbert space where the states are labelled by the irreps of $SO(3)$, as in Ref. 5. We define

$$\hat{Y}_l^m |s\rangle = \sqrt{2l+1} |l, m\rangle, \quad (\text{A4})$$

where $|s\rangle$ is the $l = m = 0$ state, and the irrep states satisfy

$$\langle l, m | l', m' \rangle = \delta(l, l') \delta(m, m'). \quad (\text{A5})$$

The remaining matrix elements of the \hat{Y}_l^m are determined by Clebsch-Gordan decomposition

$$\hat{Y}_{l_1}^{m_1} \hat{Y}_{l_2}^{m_2} = \sum_{L=|l_1-l_2|}^{L=l_1+l_2} \hat{Y}_L^M \langle L, M | l_1, m_1; l_2, m_2 \rangle, \quad (\text{A6})$$

with $M = m + m'$.

We now imagine truncating the Hilbert space by imposing a hard cutoff on the irrep labels $L \leq L_{max}$. Setting $L_{max} = 1$ for the Hamiltonian of the non-linear $O(3)$ sigma model, we have 4 states per lattice site and a 16×16 Hamiltonian;

$$\hat{H}_\phi = \frac{4\pi J_2}{3} \sum_{\langle i, j \rangle, m} (-1)^m \hat{Y}_1^m(i) \hat{Y}_1^{-m}(j). \quad (\text{A7})$$

Allowing this to act on the singlet-singlet state, we can identify a pair-creation term

$$\hat{H}_p = \frac{4\pi J_2}{\sqrt{3}} \sum_{\langle i, j \rangle} (-1)^m |m, -m\rangle \langle ss |_{i, j} + \text{h.c.} \quad (\text{A8})$$

The Hamiltonian must either increase or decrease l_1 and l_2 by one, so the only other states on which \hat{H}_ϕ acts non-

trivially are those of the form $|sm\rangle$ or $|ms\rangle$. Subtracting \hat{H}_p from \hat{H}_ϕ yields a simple hopping term

$$\hat{H}_h = \frac{4\pi J_2}{\sqrt{3}} \sum_{\langle i, j \rangle} |sm\rangle \langle ms |_{i, j} + \text{h.c.}, \quad (\text{A9})$$

where

$$\hat{H}_\phi = \hat{H}_h + \hat{H}_p. \quad (\text{A10})$$

These results also follow by directly evaluating the integral

$$\begin{aligned} \int d\{\theta, \phi\} (\vec{\phi}_i \cdot \vec{\phi}_j) \bar{Y}_{l_1}^{m_1}(\theta_1, \phi_1) \bar{Y}_{l_2}^{m_2}(\theta_2, \phi_2) \\ \times Y_{l_3}^{m_3}(\theta_1, \phi_1) Y_{l_4}^{m_4}(\theta_2, \phi_2). \end{aligned} \quad (\text{A11})$$

Performing this calculation in Mathematica for all $\{l, m\}$, $l \leq 1$, we obtain precisely the terms above (up to normalization).

The kinetic term, \hat{H}_L , is diagonal in the $|l_1, m_1; l_2, m_2\rangle$ basis;

$$\hat{H}_L |l_1, m_1; l_2, m_2\rangle = J_1 \frac{l_1^2 + l_2^2}{2} |l_1, m_1; l_2, m_2\rangle, \quad (\text{A12})$$

yielding an onsite potential term in the truncated Hilbert space,

$$\hat{H}_L = \frac{J_1}{2} \sum_{m, i} |m\rangle \langle m |_{i,}. \quad (\text{A13})$$

In this way, the qubit model from Eq. (1) can be interpreted as a truncation of the lattice-regulated Hamiltonian for the $O(3)$ non-linear sigma model. This provides a connection between our work and that of other groups who consider direct truncation schemes for simulating quantum field theories, as in Ref. 5, 23, 24.

Appendix B: Perturbation Theory

We consider the model of Eq. (1) on a regular square or cubic lattice of length L in d dimensions. We wish to know the energy gap near the weak-coupling limit, since this will inform our adiabatic algorithm, and the fidelity between the strong-coupling ground state and that at finite coupling. Both questions may be handled by ordinary Rayleigh-Schrödinger perturbation theory.

The weak-coupling vacuum $|\Omega(0)\rangle$ has the singlet on all spatial sites, $|\Omega(0)\rangle = \otimes_r |s\rangle_r$. In the degenerate case $\mu = 0$ at weak coupling, the first excited state manifold is spanned by all states with a single site in a triplet state. Thus, the energy gap at weak coupling is 1 in units where $J = 1$. The first order correction to the ground state energy is zero.

Since the first excited state manifold for a lattice with finite volume is highly degenerate, we must apply de-

generate perturbation theory. Luckily, within this manifold the perturbation is easily diagonalized. Since the pair-creation term does not act within this subspace we may ignore it. The hopping term conserves M and is translation invariant, so its eigenstates are also momentum eigenstates. Thus the hopping term decomposes into three equal blocks, each of which is diagonalized by the Fourier transform within that subspace.

In one dimension the momentum eigenstates are represented by the L th roots of unity $e^{2\pi in/L}$. The corresponding eigenvalue of the hopping term is readily shown to be $2J_r \cos(2\pi n/L)$. In d -dimensions the first-order corrections to the energy eigenvalues are

$$E_{n_i} = 2dJ_r \cos(2\pi n_i/L). \quad (\text{B1})$$

For large L the first-order correction to the energy gap is therefore $-2J_r$;

$$\Delta_E = 1 - 2dJ_r + \mathcal{O}(J_r^2). \quad (\text{B2})$$

We now consider the first-order correction to the weak-coupling ground state. Since $|\Omega(0)\rangle$ has zero momentum it can only couple to a state which also has zero momentum. The hopping term is nonzero only in the subspace with an odd number of triplet states, so it does not contribute. The pair creation term couples the weak-coupling vacuum to the subspace of the second-excited state manifold with a $|p, m\rangle, |p, -m\rangle$ pair on adjacent sites and zero momentum in each spatial direction (call this manifold D).

The first-order correction is

$$|\Omega(0)^{(1)}\rangle = -J_r \sum_{k \in D} \frac{\langle k | H_p | \Omega(0) \rangle}{2} |k\rangle. \quad (\text{B3})$$

Note that $H_p |\Omega(0)\rangle$ is the sum of all states with adjacent $|p, m\rangle, |p, -m\rangle$ pairs. An orthonormal basis of D is the equal superposition of $|p, m\rangle, |p, -m\rangle$ pairs for each of the 3 choices of m . This gives an overlap of

$$\langle k | H_p | \Omega(0) \rangle = \sqrt{L^d},$$

so that

$$|\Omega(0)^{(1)}\rangle = -\frac{J_r \sqrt{L^d}}{2} \sum_{k \in D} |k\rangle = -\frac{J_r}{2} H_p |\Omega(0)\rangle. \quad (\text{B4})$$

We can now calculate the overlap between the ground state at finite J_r and that at $J_r = 0$. The normalized finite-coupling ground state is (let $-\frac{J_r}{2} H_p |\Omega(0)\rangle = |C\rangle$)

$$|\Omega(J_r)\rangle = \frac{|\Omega(0)\rangle + |C\rangle}{\sqrt{1 + 3J_r^2 L^d/4}}, \quad (\text{B5})$$

so that

$$\langle \Omega(0) | \Omega(J_r) \rangle = (1 + 3J_r^2 L^d/4)^{-1/2}. \quad (\text{B6})$$

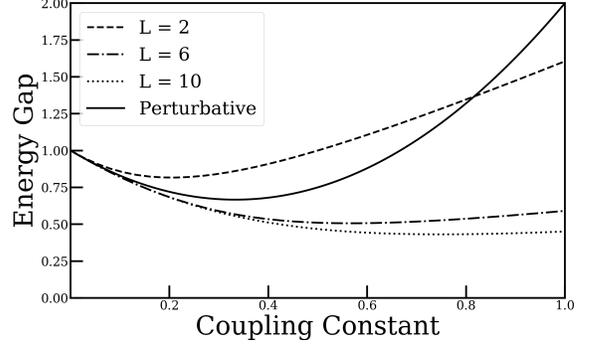


FIG. 10: Ground state energy gap as a function of coupling constant in one spatial dimension for various lattice sizes, compared to perturbative result. The energy gap is nearly linear in coupling for larger values of J_r (not shown). The spectrum appears to converge to something which is gapped at all values of J_r as L is increased.

Next, to calculate the energy gap to $\mathcal{O}(J_r^2)$, we need to obtain the first-order corrections to the first-excited states. The contributions here are from the third-excited states, which couple to the former via the pair annihilation operator. Again, momentum must be preserved. Now, however, there is a large degeneracy owing to where we decide to create the pair (for simplicity we imagine the state as having an excitation at one spatial site, create a pair somewhere else on the lattice, and extrapolate to the full state using the translation operator).

Let $|k, m\rangle_1$ be a first-excited state with momentum k and $M = m$. We can create a pair at any of $L^d - 2d$ sites (call the site r), and we have 3 options for the m -value of that pair (call this m'), so that $|k, m\rangle_1$ couples to exactly $(3L^d - 6d)$ third-excited states (call these $|k, m, r, m'\rangle$). This gives an overlap

$$\langle k, m | {}_1 H_p | k, m, r, m' \rangle = 1, \quad (\text{B7})$$

and

$$|k, m^{(1)}\rangle_1 = -J_r \sum_{r, m'} \frac{1}{2} |k, m, r, m'\rangle = -\frac{J_r}{2} H_p |k, m\rangle_1 \quad (\text{B8})$$

We find second-order corrections to the ground and first excited state energies of

$$\begin{aligned} E_0^{(2)} &= -3J_r^2 L^d/2, \\ E_1^{(2)} &= -3J_r^2 (L^d - 2d)/2, \end{aligned} \quad (\text{B9})$$

so that the energy gap to second order in the coupling is

$$\Delta E = 1 - 2dJ_r + 3dJ_r^2 + \mathcal{O}(J_r^3). \quad (\text{B10})$$

We find reasonable agreement between this result and that obtained via exact diagonalization on finite lattices (see Figures 10 and 11).

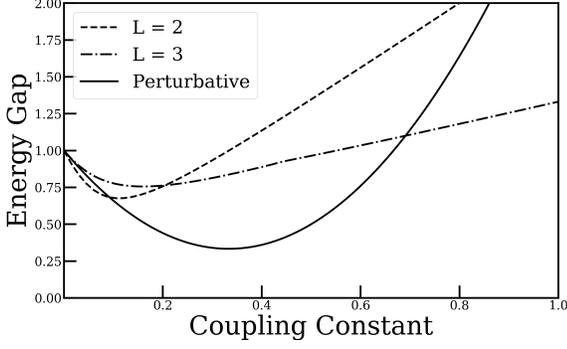


FIG. 11: Ground state energy gap as a function of coupling constant in two spatial dimensions for $L = 2, 3$, compared to perturbative result. The energy gap is nearly linear in coupling for larger values of J_r (not shown).

Appendix C: General Symmetric Model

In this section we provide a complete characterization of all translation invariant nearest-neighbor Hamiltonians with two qubits per lattice site that are invariant under $SO(3)$. This treatment serves to clarify some of the statements made in the main text. Let $g \in SO(3)$, and let U_g^ρ be the representation of g under the irrep ρ . The Hilbert space on each lattice site transforms as the direct sum $0 \oplus 1$. If H is a Hamiltonian acting on two sites, H transforms as the direct sum $0 \oplus 1 \oplus 2$, with multiplicities 2, 3, and 1, respectively.

In order to conjugate H by a group element g on each site, we apply $(U_g^0 \oplus U_g^1) \otimes (U_g^0 \oplus U_g^1)$. This is unitarily equivalent to the following 16×16 matrix;

$$U_g = \begin{bmatrix} U_g^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & U_g^1 & 0 & 0 & 0 & 0 \\ 0 & 0 & U_g^1 & 0 & 0 & 0 \\ 0 & 0 & 0 & U_g^1 & 0 & 0 \\ 0 & 0 & 0 & 0 & U_g^0 & 0 \\ 0 & 0 & 0 & 0 & 0 & U_g^0 \end{bmatrix}.$$

If H is $SO(3)$ -invariant, then $U_g H U_g^\dagger = H$ for all $g \in SO(3)$. Writing H in the same basis as that which block diagonalizes U_g , we can think of H in block-form, where each block defines a linear map between representations j and j' . According to Schur's lemma, if a block commutes with the group action it is either 0 or invertible. Thus the blocks coupling *different* representations are 0, and the blocks coupling equivalent representations are proportional to the identity. This implies that H

takes the following simple form

$$H = \begin{bmatrix} aI_5 & 0 & 0 & 0 & 0 & 0 \\ 0 & bI_3 & cI_3 & dI_3 & 0 & 0 \\ 0 & eI_3 & fI_3 & gI_3 & 0 & 0 \\ 0 & hI_3 & iI_3 & jI_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & k & l \\ 0 & 0 & 0 & 0 & m & n \end{bmatrix},$$

where I_n is the $n \times n$ identity matrix.

The most general $SO(3)$ invariant four-qubit Hamiltonian can then be written as

$$H = I_5 \oplus (I_3 \otimes H_{j=1}) \oplus (H_{j=0}), \quad (C1)$$

where $H_{j=1}$ and $H_{j=0}$ are 3×3 and 2×2 Hermitian matrices, respectively. Any additional symmetries of the Hamiltonian arise from symmetries of $H_{j=1}$ and $H_{j=0}$.

For convenience we work in the standard basis, in which the three $l = 1$ blocks are those of the ordinary symmetry channels, $1 \oplus 1$, $1 \oplus 0$, and $0 \oplus 1$, in that order. The $l = 0$ blocks are the $1 \oplus 1$ and the $0 \oplus 0$ channels, in that order.

1. Parity

In this context, by parity we mean ordinary spatial inversion; for example, in one-dimension site i is mapped to site $L - i$. Since our general symmetric Hamiltonian is translation-invariant and nearest-neighbor, under parity it is mapped to another translation-invariant nearest-neighbor Hamiltonian. Identifying the nearest-neighbor terms on a given pair of lattice sites before and after a parity transformation, in the basis considered above the parity operator is represented as

$$P = \begin{bmatrix} I_5 & 0 & 0 & 0 & 0 & 0 \\ 0 & -I_3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I_3 & 0 & 0 \\ 0 & 0 & I_3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (C2)$$

If H is $SO(3)$ invariant and $PHP^{-1} = H$, this implies that

$$H = \begin{bmatrix} aI_5 & 0 & 0 & 0 & 0 & 0 \\ 0 & bI_3 & cI_3 & dI_3 & 0 & 0 \\ 0 & c^*I_3 & eI_3 & fI_3 & 0 & 0 \\ 0 & d^*I_3 & f^*I_3 & gI_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & h & i \\ 0 & 0 & 0 & 0 & i^* & j \end{bmatrix} \\ = \begin{bmatrix} aI_5 & 0 & 0 & 0 & 0 & 0 \\ 0 & bI_3 & -dI_3 & -cI_3 & 0 & 0 \\ 0 & -d^*I_3 & gI_3 & f^*I_3 & 0 & 0 \\ 0 & -c^*I_3 & fI_3 & eI_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & h & i \\ 0 & 0 & 0 & 0 & i^* & j \end{bmatrix}. \quad (C3)$$

Without imposing parity there are 14 real degrees of freedom. Under parity, four of these are removed ($c = -d$, $e = g$, $f = f^*$).

The most general $SO(3)$ and parity-invariant Hamiltonian is then

$$H = \begin{bmatrix} aI_5 & 0 & 0 & 0 & 0 & 0 \\ 0 & bI_3 & cI_3 & -cI_3 & 0 & 0 \\ 0 & c^*I_3 & dI_3 & eI_3 & 0 & 0 \\ 0 & -c^*I_3 & eI_3 & dI_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & f & g \\ 0 & 0 & 0 & 0 & g^* & h \end{bmatrix}. \quad (C4)$$

Lastly, imposing time-reversal symmetry, we derive that $c = 0$, $g = g^*$. Using our freedom to subtract an overall constant, The full $SO(3)$, P, and T-symmetric Hamiltonian (with 6 real degrees of freedom) is

$$H = \begin{bmatrix} aI_5 & 0 & 0 & 0 & 0 & 0 \\ 0 & bI_3 & 0 & 0 & 0 & 0 \\ 0 & 0 & dI_3 & eI_3 & 0 & 0 \\ 0 & 0 & eI_3 & dI_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & f & g \\ 0 & 0 & 0 & 0 & g & 0 \end{bmatrix}. \quad (C5)$$

We would like to connect this general model to our original qubit-regularized Hamiltonian. From now on, we neglect the cumbersome I_n in each block. The terms in our original Hamiltonian take the form (setting $\mu = 0$)

$$H^h = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (C6)$$

$$H^p = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \quad (C7)$$

$$H_1 = \begin{bmatrix} 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (C8)$$

We get back to the most general model by adding three additional couplings

$$H^X = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (C9)$$

$$H^= = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (C10)$$

$$H^{sm} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (C11)$$

Including all the couplings, we obtain the Hamiltonian

$$H = \begin{bmatrix} 2J_t + J_X & 0 & 0 & 0 & 0 & 0 \\ 0 & J_t - J_X & 0 & 0 & 0 & 0 \\ 0 & 0 & J_t + J_{sm} & J_h & 0 & 0 \\ 0 & 0 & J_h & J_t + J_{sm} & 0 & 0 \\ 0 & 0 & 0 & 0 & J_+ + J_X J_p & 0 \\ 0 & 0 & 0 & 0 & J_p & 0 \end{bmatrix}. \quad (C12)$$