Nearly optimal state preparation for quantum simulations of lattice gauge theories

Christopher F. Kane [1], Niladri Gomes [2,*] and Michael Kreshchuk [3]
[1]*Department of Physics, University of Arizona, Tucson, Arizona 85719, USA*
[2]*Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*
[3]*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

(Received 28 January 2024; accepted 11 June 2024; published 23 July 2024)

We present several improvements to the recently developed ground-state preparation algorithm based on the quantum eigenvalue transformation for unitary matrices (QETU), apply this algorithm to a lattice formulation of U(1) gauge theory in $(2+1)$ dimensions, as well as propose an alternative application of QETU, a highly efficient preparation of Gaussian distributions. The QETU technique was originally proposed as an algorithm for nearly optimal ground-state preparation and ground-state energy estimation on early fault-tolerant devices. It uses the time-evolution input model, which can potentially overcome the large overall prefactor in the asymptotic gate cost arising in similar algorithms based on the Hamiltonian input model. We present modifications to the original QETU algorithm that significantly reduce the cost for the cases of both exact and Trotterized implementation of the time evolution circuit. We use QETU to prepare the ground state of a U(1) lattice gauge theory in two spatial dimensions, explore the dependence of computational resources on the desired precision and system parameters, and discuss the applicability of our results to general lattice gauge theories. We also demonstrate how the QETU technique can be utilized for preparing Gaussian distributions and wave packets in a way which outperforms existing algorithms for as little as $n_q \gtrsim 2$–5 qubits.

DOI: 10.1103/PhysRevA.110.012455

I. INTRODUCTION

Simulating many-particle quantum systems has always been seen as one of the most exciting potential applications of quantum computers [1]. While simulation of nonrelativistic many-body physics is already by itself complex enough to be considered as a candidate for quantum advantage [2–6], quantum field theory (QFT) adds to the picture such ingredients as particle number nonconservation, a multitude of particle types and their interactions, gauge symmetry, etc. Quantum simulation of realistic QFTs such as quantum chromodynamics (QCD) will likely require the access to fault-tolerant quantum computers. To extract as much physics as possible from such devices, it is imperative to investigate all theoretical aspects of the quantum simulation. The preliminary steps include constructing a model suitable for quantum computation by discretizing the original continuous theory [7,8], mapping the degrees of freedom in the discretized model onto logical qubits [9–45], as well as mapping the logical qubits onto the physical ones within an error-correcting procedure [46–54].[1] Those are followed by the quantum simulation itself, which typically involves preparing on a quantum computer states of specific form [56–69] (eigenstates, wave packets, thermal states) and evolving those in time. Lastly, once the desired final state is prepared, it is used for measuring the observables of interest.

Most early approaches to simulating time evolution and ground-state preparation were based on product formulas [2,4,70]. These methods do not require ancillary qubits and, due to the overall simplicity of quantum circuit design, could be readily applied to simple models. However, they are likely to not take advantage of neither noisy near-term devices nor fault-tolerant ones. The former direction has been revolutionized by the invention of variational [71–77] and subspace [78–82] methods, while the exploration of the latter one has led to the development of quantum simulation algorithms optimal and nearly optimal in problem parameters [83–91]. The asymptotic cost of these nearly optimal algorithms is typically given in terms of the number of calls to subroutines that provide information about the physical Hamiltonian. Most such algorithms rely on the usage of *block encoding* subroutines whose construction comes at a great expense and leads to large prefactors in the final gate cost [92,93]. A number of algorithms, however, use instead the Hamiltonian time evolution input model [94–97], in which case it is an approximate time evolution circuit that serves as an elementary block for the circuit construction. The development of this idea has ultimately lead to the quantum eigenvalue transformation for unitary matrices (QETU) algorithm [98], which enables one to assemble circuits applying a large class of polynomial transformations of a unitary operator to a given state, including those that implement nonunitary dynamics [99].

In this work, we present the first study using QETU for preparing the ground state of a lattice gauge theory. The test theory we consider is a particular formulation of a U(1) lattice gauge theory in two spatial dimensions [41,100,101]. Our

*Present address: The University of Chicago, Chicago, Illinois 60637, USA.

[1]Further optimization is likely to be achieved via cross-layer design, e.g., mapping physical degrees of freedom directly onto physical qubits [55].

study includes an analysis of how the parameters of the QETU algorithm, as well as the Trotter error from approximating the time evolution circuit, scale with the system size, number of qubits per site, and gauge coupling. Armed with the intuition from this numerical study, we provide a general discussion about how the cost of QETU for ground-state preparation of a QCD-like lattice gauge theory will scale with the system parameters.

We also present an alternative application of QETU for preparing Gaussian states in quantum mechanical systems. We show that, while a naïve application of QETU to this problem results in the error of the approximation decreasing polynomially with the number of calls to the time-evolution circuit, this scaling can be made exponential with simple modifications to the QETU procedure. Using our improved methods, we show that the cost of preparing Gaussian states with QETU outperforms existing state preparation methods for states represented using $n_q > 2$–5 qubits.

In addition to these more specialized studies, we developed modifications to the original QETU algorithm that can significantly reduce the cost for arbitrary Hamiltonians. These modifications can be applied to both the scenario when the time evolution circuit is prepared exactly and when prepared approximately using Trotter methods.

The rest of this work is organized as follows: In Sec. II A we review the general QETU algorithm. From there, in Sec. II B we review how QETU can be used for ground-state preparation. Next, in Sec. II B 2 we discuss our modifications to the original QETU algorithm and provide numerical demonstrations of the achievable cost reduction. In Sec. II C we discuss how to use QETU to prepare Gaussian states. Section III reviews the details of the test theory we consider, which is a particular formulation of a U(1) lattice gauge theory in two spatial dimensions. Numerical results for ground-state preparation of this U(1) lattice gauge theory using QETU are presented in Sec. IV A. From there, we present the results of our study using QETU to prepare Gaussian states in Sec. IV B. In Sec. V, we provide a general discussion of the scaling expected when extending QETU for ground-state preparation of a general lattice gauge theory with the same qualitative properties as QCD. Our conclusions, as well as a discussion of future applications, are presented in Sec. VI. The main notations used throughout the paper are listed in Table I.

TABLE I. Main notations used in the paper.

Quantum eigenvalue transformation for unitary matrices	
$\mathsf{f}(x)$	Function to be approximated by QETU
$\mathsf{F}(x)$	$\mathsf{f}(2\arccos(x))$
$F(x)$	Polynomial approximation to $\mathsf{F}(x)$
$f(x)$	$F(\cos(x/2))$
$H^{\rm phys} = \sum_n E_n^{\rm phys} \lvert\psi_n\rangle\langle\psi_n\rvert$	Physical Hamiltonian
$H = \sum_n E_n \lvert\psi_n\rangle\langle\psi_n\rvert$	Rescaled Hamiltonian
$c_{1,2}$	Constants in operator rescaling, Eqs. (3) and (23)
$P_{<\mu} = \lvert\psi_0\rangle\langle\psi_0\rvert$	Projector onto ground state of H
$\lvert\psi_{\rm init}\rangle$	Initial state
$\sigma_\pm$, $\sigma_{\rm min}$, $\sigma_{\rm max}$	Parameters of shifted sign function, Eq. (7)
$\Delta^{\rm phys}$	Initial (approximate) knowledge of $(E_1^{\rm phys} - E_0^{\rm phys})/2$
Δ	Initial (approximate) knowledge of $(E_1 - E_0)/2$
μ	Initial (approximate) knowledge of $(E_1 + E_0)/2$
τ	Evolution or scale parameter of $e^{-iH\tau}$
$N_{\rm steps}$	Number of steps in Trotter implementation of U
$\delta\tau$	$\tau/N_{\rm steps}$, Trotter step size
$\tau_{\rm max}$	Largest value of τ guaranteeing isolation of the ground state
η	Parameter defining H spectrum bounds, Eq. (4)
$\eta_{P_{<\mu}}$	Parameter of the shifted sign function, Eqs. (10) and (11)
$\gamma = \lvert\langle\psi_{\rm init}\rvert\psi_0\rangle\rvert \geqslant 0$	Overlap between initial guess and true ground state
d	Degree of (even) $F(x)$ polynomial
	Total number of calls to both U and $U^\dagger$
$n_{\rm Ch}$	Number of Chebyshev polynomials used to represent (even) $F(x)$
$N_{\rm steps}$	Number of Trotter steps in single call to U
$N_{\rm tot}$	$d \times N_{\rm steps}$, total number of Trotter steps in QETU circuit
$\mathcal{F}$	Discrete Fourier transformation matrix
U(1) gauge theory	
g	Coupling constant
a	Lattice spacing
V	Physical volume
N_s	Number of sites in each dimension of a lattice
N_p	Number of independent plaquettes
$b_{{\rm max},p}$, δb_p	Maximum value of magnetic field operator and discretization step for pth plaquette, Eq. (33)

II. QUANTUM EIGENVALUE TRANSFORMATION FOR UNITARY MATRICES

We begin this section by reviewing the QETU algorithm in its most general form, as well as its application to ground-state preparation. Readers familiar with QETU and the original state preparation procedure in Ref. [98] may proceed directly to Sec. II B 2. Next, we present modifications to the original algorithm that reduce the cost of ground-state preparation. Lastly, we consider an application of QETU to the preparation of Gaussian distributions and wave packets.

A. Algorithm review

Similarly to the quantum eigenvalue transformation [85,102] algorithm, the QETU circuit realizes a polynomial

TABLE I. (*Continued.*)

Quantum eigenvalue transformation for unitary matrices	
$r_{\max,p}$, δr_p	Maximum value of rotor operator and discretization step for pth plaquette, Eq. (34)
n_q	Number of qubits per lattice site
Wave-packet preparation	
x_0, p_0	Expectation values of position and momentum operators
σ_x	Wave-packet width
$\hat{x}_{\text{sh}}$	Shifted position operator
x_0^{QETU}	$c_1 x_0 + c_2$
σ_{QETU}	$c_1 \sigma_x$
$\mathsf{F}_+(x)$, $\mathsf{F}_-(x)$	Even and odd parts of $\mathsf{F}(x)$

transformation of $e^{-i\tau H}$, which, in turn, can be used for implementing a wide class of functions of H. While for a given Hermitian operator H constructing the exact circuit for $e^{-i\tau H}$ is, generally, a nontrivial problem—even for short times τ—QETU can potentially render useful results for approximate implementations of $e^{-i\tau H}$. The impact of such approximations on the method's performance will be explored in future sections.

Preparing the ground state of a Hamiltonian with the aid of the QETU algorithm is based on the concept of *filtering* [97,98,103–106], which implies constructing a circuit approximately implementing the action of a projector $P_{<\mu} = |\psi_0\rangle\langle\psi_0|$ onto the ground state of $H = \sum_n E_n |\psi_n\rangle\langle\psi_n|$ and applying this circuit to a state $|\psi_{\text{init}}\rangle$ having significant overlap $|\langle\psi_0|\psi_{\text{init}}\rangle| = \gamma \geqslant 0$ with the ground state:

$$P_{<\mu}|\psi_{\text{init}}\rangle \propto |\psi_0\rangle. \tag{1}$$

In Ref. [104] it was observed that $\cos^M H$ for large values of M is approximately proportional to the projector onto the ground state of H for a properly normalized H. The circuit for $\cos^M H$ in this approach is expressed as $\cos^M H = (e^{iH} + e^{-iH})^M / 2^M$, where the exponents are implemented using any available time evolution algorithm, while the linear combination of terms is obtained via the linear combination of unitaries (LCU) algorithm [85,86,102,107–111]. While Ref. [104] combined the ideas of the Hamiltonian time evolution input model [94–97] (i.e., utilizing e^{-iH} as a building block for the algorithm), this approach did not have optimal performance. In Ref. [105] algorithms based on the filtering concept have been developed based on the Hamiltonian oracle access model and QET technique [84]. Despite the asymptotically nearly optimal performance of algorithms in Ref. [105], they suffer from a high cost of the oracle subroutines.

Finally, nearly optimal algorithms for ground-state energy estimation [97] and ground-state preparation [98] based on the Hamiltonian time evolution input model have been proposed. While efficient ground-state energy estimation (without ground-state preparation) could be achieved with as little as a single implementation of $e^{-i\tau H}$ for several values of τ [97], preparing the ground state in QETU amounts to using a circuit similar to the circuit used in quantum eigenvalue transformation (QET) [85,98,102], see Fig. 1.

The QETU theorem [98] assumes the access to a circuit implementing e^{-iH} for an n-qubit Hermitian operator H. It states that, for any even real polynomial $F(x)$ of degree d satisfying $|F(x)| \leqslant 1 \;\; \forall\, x \in [-1, 1]$, one can find a sequence of symmetric phase factors $(\varphi_0, \varphi_1, \ldots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$ such that the circuit in Fig. 1 denoted by $\mathcal{U}$ satisfies $(\langle 0| \otimes \mathbb{1}_n)\mathcal{U}(|0\rangle \otimes \mathbb{1}_n) = F(\cos(H/2))$ [98]. In practice, the QETU circuit is used for approximately implementing $\mathsf{f}(H)$ by realizing a transformation $F(\cos(H/2))$, where $F(x)$ is a polynomial approximation to $\mathsf{F}(x) \equiv \mathsf{f}(2\arccos(x))$. Note that there also exists a control-free version of the QETU circuit that avoids having to implement controlled calls to the e^{-iH} circuit [98], which we explain in more detail in Appendix C.

The original algorithm in Ref. [98] proposed the use of e^{-iH} in the QETU circuit. It is possible, however, to use $e^{-i\tau H}$ instead. Doing so leads to a modification of the Lemma above, where the mapping is now $F(\cos(\tau H/2))$. Because $\cos(\tau x/2)$ is periodic, one must be careful when choosing τ to ensure construction of the desired function $\mathsf{f}(x)$. In the following section, we discuss how using $\tau \neq 1$ can reduce the cost of state preparation. In Sec. IV B we discuss how $\tau \neq 1$ can reduce the cost of using QETU for constructing Gaussian states.

B. Ground-state preparation via quantum eigenvalue transformation for unitary matrices

In this section we review the algorithm proposed in Ref. [98] for ground-state preparation using QETU. We first state the necessary assumptions and scaling of the algorithm. After reviewing the original algorithm in more detail, we discuss how using instead $e^{-i\tau H}$ as a building block in the QETU circuit can lead to significant cost reductions, both when implementing $e^{-i\tau H}$ exactly as well as using product formulas. We conclude by discussing how QETU can be used to prepare n_q-qubit Gaussian states with a cost that is linear in n_q.

1. Original algorithm

Suppose one has access to a Hamiltonian H via e^{-iH}, and that the spectrum of H is in the range $[\eta, \pi - \eta]$ for some $\eta > 0$. Furthermore, assume one has knowledge of parameters μ and Δ such that

$$E_0 \leqslant \mu - \Delta/2 \leqslant \mu + \Delta/2 \leqslant E_1, \tag{2}$$

where E_i is the ith excited state of H. Here μ represents the knowledge of the precise values of the energies, and Δ is a lower bound of the excited-state energy gap $E_1 - E_0$. Figure 2 shows an example of the exact projector onto the ground state, given by the step function $1 - \theta(x - \mu)$, which isolates the ground state even with only partial knowledge of E_0 and E_1. Lastly, assume one has an initial guess $|\psi_{\text{init}}\rangle$ for the ground state with overlap satisfying $|\langle\psi_0|\psi_{\text{init}}\rangle| \geqslant \gamma$. Under these assumptions, one can prepare a state $|\widetilde{\psi}_0\rangle$ such that $|\langle\psi_0|\widetilde{\psi}_0\rangle| \geqslant 1 - \epsilon$ using $\widetilde{O}(\gamma^{-2}\Delta^{-1}\log(1/\epsilon))$ controlled calls to the time evolution circuit for e^{-iH}. We spend the rest of this section describing the details of the algorithm, including how the scaling with γ, Δ, and ϵ arise.

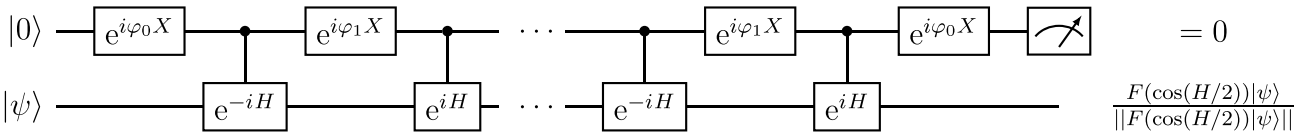

FIG. 1. QETU circuit diagram. The top qubit is the control qubit, and the bottom register is the state that the matrix function is applied to. Upon measuring the ancillary qubit to be in the zero state, one prepares the normalized quantum state $F(\cos(H/2))|\psi\rangle/||F(\cos(H/2))|\psi\rangle||$ for some polynomial $F(x)$. For symmetric phase factors $(\varphi_0, \varphi_1, \ldots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$, then $F(x)$ is a real even polynomial of degree d. The probability of measuring the control qubit in the zero state is $p = ||F(\cos(H/2))|\psi\rangle||^2$.

We begin by addressing the fact that the spectrum of our target Hamiltonian must be in the range $[\eta, \pi - \eta]$ for some $\eta > 0$. Because QETU returns a function with a periodic argument, i.e., $F(\cos(H/2))$, any projector constructed with QETU will repeat itself for large enough energies, as illustrated in Fig. 3. From this we see that unless the spectrum of H is in a limited range, we cannot guarantee that higher excited states will be filtered out. To avoid this problem, in Ref. [98] the Hamiltonian was first scaled so that its spectrum was in the range $[\eta, \pi - \eta]$ for some $\eta > 0$. We find that this constraint can be somewhat relaxed, which is described in detail in Sec. II B 2. If the physical, unshifted Hamiltonian is given by H^{phys} with energies E_i^{phys}, the Hamiltonian with the shifted spectrum is given by

$$H = c_1 H^{\text{phys}} + c_2 \mathbb{1}, \tag{3}$$

where

$$c_1 = \frac{\pi - 2\eta}{E_{\text{max}}^{\text{phys}} - E_0^{\text{phys}}}, \quad c_2 = \eta - c_1 E_0^{\text{phys}}. \tag{4}$$

Alternatively, one could replace $E_{\text{phys}}^{\text{max}}$ with an upper bound on the maximum eigenvalue. One important consequence of shifting the spectrum is that the energy gap of the shifted Hamiltonian shrinks as well. If Δ^{phys} is the energy gap of the physical Hamiltonian H^{phys}, the shifted gap is given by $\Delta = c_1 \Delta^{\text{phys}}$. This tells us that $\Delta \approx \Delta^{\text{phys}}/E_{\text{max}}$. Because the maximum eigenvalue of a Hamiltonian generally grows with the number of terms, the gap Δ used in the QETU algorithm will generally shrink with the number of lattice sites. Therefore, one generally expects $\Delta \sim 1/N_{\text{sites}}$, where N_{sites} denotes the number of lattice sites use in a simulation of some lattice gauge theory. We discuss this scaling in more detail in Secs. IV A and V.

Given some Hamiltonian H and the associated parameters μ and Δ, one can construct an approximate projector onto the ground state by constructing an approximation $f(x)$ to $\mathsf{f}(x) = 1 - \theta(x - \mu)$ satisfying

$$\begin{aligned} |f(x) - c| &\leqslant \epsilon \ \forall\, x \in [\eta, \mu - \Delta/2], \\ |f(x)| &\leqslant \epsilon \ \forall\, x \in [\mu + \Delta/2, \pi - \eta]. \end{aligned} \tag{5}$$

Furthermore, because a unitary matrix must have entries with magnitude less than or equal to one, we require $|f(x)| \leqslant 1\ \forall\, x$. The parameter c is chosen to be slightly smaller than one to avoid numerical overshooting when finding a Chebyshev polynomial satisfying the above constraints and is discussed in more detail later in the section.

To approximately implement the action of $\mathsf{f}(H)$ with QETU, we need the polynomial $F(x)$ to approximate the function $\mathsf{f}(2\arccos(x))$. The presence of the $\arccos(x)$ implies that

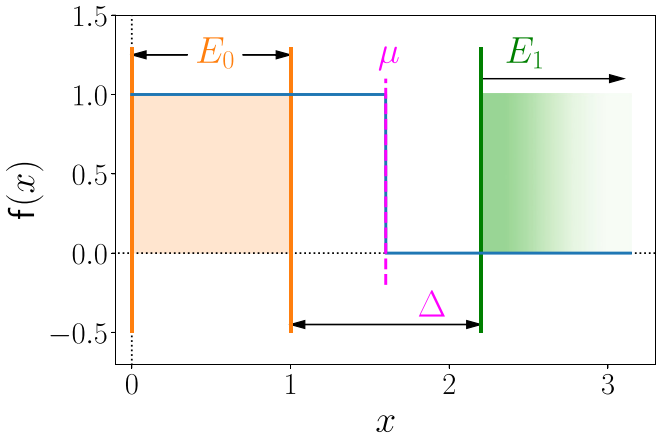

FIG. 2. Plot of the step function $\mathsf{f}(x) = 1 - \theta(x - \mu)$ encoding partial information known about E_0 and E_1. the x-axis variable x corresponds to the energy E. The ground-state energy E_0 is known to be in the shaded orange region between the orange vertical lines. The first-excited state E_1 is known to be larger than the value indicated by the green vertical line. The gap Δ used is the difference between the lower bound of E_1 and the upper bound of E_0. The central value μ is chosen to be halfway between the upper bound of E_0 and the lower bound of E_1. Isolating the ground state is possible without exact knowledge of E_0 and E_1.

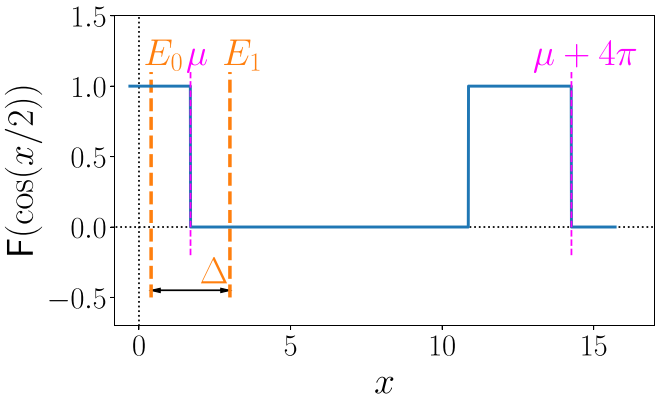

FIG. 3. To construct the approximate projector $P_{<\mu}$ onto the ground state, QETU implements an approximation of the function $\mathsf{F}(\cos(H/2))$, where $\mathsf{F}(x) = \mathsf{f}(2\arccos(x))$ and $\mathsf{f}(x)$ is a unit step function $\mathsf{f}(x) = 1 - \theta(x - \mu)$. The x-axis variable x corresponds to the energy E. Because $\cos(x/2)$ is periodic, the step function repeats itself with period 4π. Unless the spectrum of the Hamiltonian is in a limited range, one cannot guarantee that all the excited states will be filtered out.

$F(x)$ can only be defined for $x \in [-1, 1]$. Taking the cosine transform into account, we wish to construct $F(x)$ such that

$$\begin{aligned} |F(x) - c| &\leqslant \epsilon \; \forall \, x \in [\sigma_+, \sigma_{\max}], \\ |F(x)| &\leqslant \epsilon \; \forall \, x \in [\sigma_{\min}, \sigma_-], \\ |F(x)| &\leqslant c \; \forall \, x \in [-1, 1], \end{aligned} \tag{6}$$

where

$$\begin{aligned} \sigma_\pm &= \cos \frac{\mu \mp \Delta/2}{2}, \\ \sigma_{\min} &= \cos \frac{\pi - \eta}{2}, \\ \sigma_{\max} &= \cos \frac{\eta}{2}. \end{aligned} \tag{7}$$

For a nonzero gap Δ and a nonzero error ϵ, one candidate function for $F(x)$ is the shifted error function, which has the important property that Chebyshev approximations of it converge exponentially with the degree of the polynomial [98]. In principle, one could solve for $F(x)$ by choosing a specific error function such that the Chebyshev approximation to it has a constant error for all x. While it was shown in Ref. [112] that this procedure prepares the ground state of a system with a gap Δ to a precision ϵ using a polynomial with degree scaling as $O(\Delta^{-1} \log \epsilon^{-1})$, this method can lead to numerical instabilities if not performed carefully [98]. Instead, we follow a different procedure, also described in Ref. [98], that avoids the need to first choose a shifted error function while still producing a near-optimal approximation. This procedure is discussed in detail later in this section.

For the sake of argument, if one did choose $\mathsf{F}(x)$ to be a shifted error function, there would exist a Chebyshev polynomial approximation to it with fidelity $1 - \epsilon$ where the degree of the polynomial is $O(\Delta^{-1} \log(\epsilon^{-1}))$ [98]. While we do not in practice choose $\mathsf{F}(x)$ to be a shifted error function manually, the convex-optimization method we use to determine the Chebyshev approximation of $\mathsf{F}(x)$ has this same scaling. The scaling with Δ can be understood by first recognizing that, for a smaller Δ, the shifted error function rises more quickly. A steeper rising error function requires a higher degree polynomial to approximate the function to the same precision than a more slowly rising one. By inverting $O(\Delta^{-1} \log(\epsilon^{-1}))$, the error ϵ scales as $\epsilon = O(e^{-b\Delta d})$, where b is a constant. Figure 4 shows an example $F(x)$ for parameter values $\eta = 0.3$, $c = 0.999$, $E_0 = 1$, $E_1 = 1.6$, $\mu = (E_0 + E_1)/2$, $\Delta = (E_1 - E_0)$, corresponding to $\sigma_{\min} = 0.15$, $\sigma_- = 0.70$, $\sigma_+ = 0.88$, and $\sigma_{\max} = 0.99$. The function $F(x)$ is represented using a $d = 22$ degree polynomial, resulting in error $\epsilon \approx 0.0096$.

The final piece of the scaling has to do with the overlap of the initial guess for the ground state. Because the probability to measure zero in the ancillary register of the QETU circuit is given by $||F(\cos(H/2))|\psi_{\text{init}}\rangle|| = \gamma^2$, the number of times the circuit must be prepared in order to measure zero increases for a poor initial guess. It can therefore be beneficial to dedicate substantial resources to prepare a high-quality initial guess $|\psi_{\text{init}}\rangle$.

We now review the procedure in Ref. [98] for finding a degree d even Chebyshev polynomial for $F(x)$. The goal is to

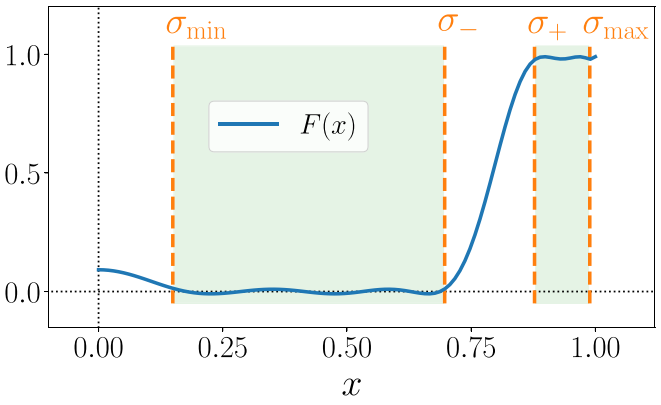

FIG. 4. The blue curve shows $F(x)$, where $F(x)$ is a Chebyshev approximation of the shifted error function. The orange vertical dashed lines indicate the values $\sigma_{\min}$, $\sigma_{\max}$, $\sigma_\pm$. The green shaded regions indicate x values that are included when solving for the Chebyshev expansion of the cosine transform shifted error function. For values of x outside the shaded green regions, the function $F(x)$ can take on any value as long as $|F(x)| \leqslant 1$.

solve for $n_{\text{Ch}} = d/2 + 1$ Chebyshev coefficients c_k such that

$$F(x) = \sum_{k=0}^{n_{\text{Ch}}-1} c_{2k} T_{2k}(x). \tag{8}$$

To do so, one first samples $F(x)$ at a set of M discrete points. It is known that polynomial interpolation sampling at equidistant points is exponentially ill conditioned [113]. To avoid this problem, $F(x)$ is instead sampled using the roots of Chebyshev polynomials $x_j = -\cos(\frac{j\pi}{M-1})$ for $j = 0, \ldots, M-1$ for some large value of M. Because the region where $F(x) = c$ shrinks with Δ, one has to be careful to choose a value of M large enough to resolve this region. Once we choose the values of x_j to sample $F(x)$, we define a coefficient matrix $A_{jk} = T_{2k}(x_j)$ such that $F(x_j) = \sum_k A_{jk} c_{2k}$. The coefficients are then determined by solving the following optimization problem:

$$\min_{\{c_k\}} \max \left\{ \max_{x_j \in [\sigma_+, \sigma_{\max}]} |F(x_j) - c|, \max_{x_j \in [\sigma_{\min}, \sigma_-]} |F(x_j)| \right\}. \tag{9}$$

This is a convex optimization problem and can be solved using, e.g., the Matlab code CVX [114] or the Python code CVXPY [115]. One chooses c smaller than one to allow the Chebyshev approximation to overshoot the step function, which, by the equioscillation theorem, is necessary to achieve an optimal Chebyshev approximations [113]. In practice, one can choose c to be sufficiently close to one such that the effect is negligible. Our numerical investigations showed that the cost of solving for the coefficients generally scales linearly or better in the degree of the polynomial.

Once the Chebyshev expansion for $F(x)$ is known, the final step is to calculate the phases $\{\varphi_j\}$ used in the QETU circuit. Note that multiple conventions for defining the phases $\{\varphi_j\}$ exist, and throughout this work we follow Ref. [98] and use the so-called *W-convention*. It has been found in Ref. [116] that using a quasi-Newton optimization method one can robustly find the symmetric phase factors for values of $d \approx 10\,000$. An example code to solve for the Chebyshev coefficients c_k and

the associated phase factors has been implemented in QSPPACK [117]. Note that the phases appearing in the QETU circuit $\{\varphi_j\}$ are related to the symmetric phases calculated using QSPPACK $\{\phi_j\}$ by $\varphi_j = \phi_j + (2-\delta_{j0})\pi/4$ (see Appendix B in Ref. [98]). Through numerical studies, we find that the cost of finding the phase factors scales roughly quadratically with the number of phases. More details regarding the calculation of the phases can be found in Appendix A.

To summarize, using QETU to prepare the ground state of a target Hamiltonian H^{phys} can be done according to the following steps:

(1) Construct $H = c_1 H^{\text{phys}} + c_2$ such that the spectrum of H is in $[\eta, \pi - \eta]$ for some $\eta > 0$.
(2) Determine μ and Δ.
(3) Solve for the Chebyshev approximation $F(x)$.
(4) Solve for the phase factors $\{\phi_j\}$.
(5) Implement the circuit in Fig. 1.

In the next section, we describe how the constraint that the spectrum of H should be in the range $[\eta, \pi - \eta]$ can be relaxed.

2. *Modified ground state preparation*

In the previous section, we reviewed the original QETU algorithm for ground-state preparation in which the Hamiltonian spectrum was assumed to be in the range $[\eta, \pi - \eta]$. This assumption was necessary because the function $\cos(x)$ is monotonic in the range $x \in [0, \pi]$. While this is true, QETU actually returns $F(\cos(x/2))$,[2] with $\cos(x/2)$ being monotonic in the range $x \in [0, 2\pi]$. This observation can be leveraged to increase the allowed range of the spectrum of H. This is useful because a larger range leads to a larger energy gap used in the QETU algorithm, which reduces the overall cost of the simulation. For the modified algorithm, we introduce the variable parameter τ to characterize the increased spectrum range. This adjustment to the original QETU algorithm can be viewed from two perspectives: one can either continue to use e^{-iH} as a building block and change the spectrum of H to be in the range $[\eta, \tau(\pi - \eta)]$, or, equivalently, consider $e^{-i\tau H}$ as a building block with the spectrum of H in the original range $[\eta/\tau, \pi - \eta]$. We choose the latter perspective, and in what follows derive the largest value of τ one can use while still guaranteeing isolation of the ground state. Depending on the context, it will be useful to think of τ as either a scale parameter of the Hamiltonian spectrum or as an evolution time; if necessary, we state explicitly which perspective is being used.

First, while the initial algorithm in Ref. [98] proposed using the same value of η when constructing the shifted sign function and when shifting the spectrum of the Hamiltonian, it is in principle possible to use different values. We continue to use η to denote the value used for shifting the spectrum of H to be in the range $[\eta, \pi - \eta]$, and introduce $\eta_{P_{<\mu}}$ to denote the value used when constructing the shifted error function. Note that, in order to avoid the scenario where the ground state is filtered out and excited states are not, one must ensure $\eta_{P_{<\mu}} \leqslant \eta$. The parameters of the shifted sign function for general τ are then given by

$$\sigma_{\pm}(\tau) = \cos\left(\tau \frac{\mu \mp \Delta/2}{2}\right), \tag{10}$$

$$\sigma_{\min}(\tau) = \cos\left(\tau \frac{\pi - \eta_{P_{<\mu}}}{2}\right), \tag{11}$$

$$\sigma_{\max}(\tau) = \cos\left(\tau \frac{\eta_{P_{<\mu}}}{2}\right). \tag{12}$$

Note that while there are no theoretical issues with using $\tau < 1$, doing so shrinks the energy gap and is not beneficial in the context of exact implementations of $e^{-i\tau H}$. We are now in a position to discuss the two possible pitfalls that can occur when using $\tau \neq 1$ and how to overcome them.

Recall from Eq. (6) that the approximate (even) step function $F(x)$ can be larger than ϵ in the region $-\sigma_{\min} \leqslant x \leqslant \sigma_{\min}$ (see Fig. 4 for an example of this behavior). The first caveat is that, for $\tau > 1$, it is possible that an excited-state energy falls into this region and will be suppressed by a factor larger than ϵ. This can be avoided by implementing the step function with $\eta_{P_{<\mu}} = 0$, i.e., set $\sigma_{\min} = 0$. Our numerical studies indicate the quality of the approximation is largely independent of $\eta_{P_{<\mu}}$, with the choice $\eta_{P_{<\mu}} = 0$ resulting in close to the smallest error in the range $0 \leqslant \eta_{P_{<\mu}} \leqslant \eta$.

The second caveat is that one must ensure that higher excited states fall in regions where they are filtered out by the approximate step function $F(x)$. Consider first using $\tau = 1$. Taking the cosine transform into account, the locations of the shifted energies E_i are $\cos(E_i/2)$, where $E_i \in [\eta, \pi - \eta]$. Because $\cos(x)$ is monotonically decreasing on $x \in [0, \pi]$, the cosine transformed energies get successively closer to zero for larger energies, with $\cos(E_{\max}/2) = \cos((\pi - \eta)/2)$ which is close to zero for small η. If one instead uses $\tau > 1$, the transformed energies are $\cos(\tau E_i/2)$. Taking advantage of the fact that the step function is even, we see that as long as $\cos(\tau E_i/2) > -\sigma_-(\tau)$, one still guarantees isolation of the ground state. This leads to a maximum value

$$\tau_{\max} = \frac{2\pi}{\pi - \eta + (\mu + \Delta/2)}. \tag{13}$$

Figure 5 shows the fidelity of the prepared ground state of a simple harmonic oscillator as a function of τ. The Hamiltonian is given by

$$\widehat{H} = \frac{g^2}{2}\hat{p}^2 + \frac{1}{2g^2}\hat{x}^2 \tag{14}$$

and is represented using n_q qubits. Working in the digitized eigenbasis of the position operator, we choose to sample its eigenvalues as

$$x_j = -x_{\max} + (\delta x) j, \quad j = 0, 1, \ldots, 2^{n_q} - 1, \tag{15}$$

where $\delta x = 2x_{\max}/(2^{n_q} - 1)$. Using fact that $[\hat{x}, \hat{p}] = i$, the momentum operator is implemented as

$$p_j^{(\hat{p})} = -p_{\max} + (\delta p) j, \quad j = 0, 1, \ldots, 2^{n_q} - 1, \tag{16}$$

$$\hat{p}^{(\hat{x})} = \mathcal{F}^{\dagger} \hat{p}^{(\hat{p})} \mathcal{F}, \tag{17}$$

[2]The control-free version of QETU, considered below in Sec. IV A and in Appendix C, implements instead $F(\cos(x))$. The analysis of this section can be adapted to that case upon replacing $F(\cos(x/2)) \to F(\cos(x))$.

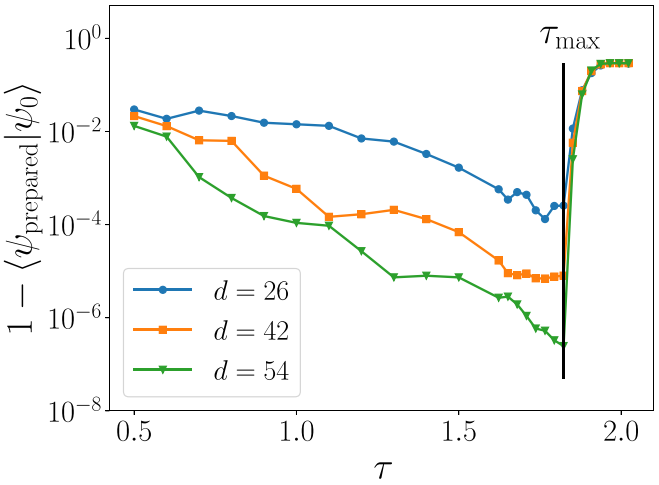

FIG. 5. Error of the state prepared using QETU as a function of τ for a simple harmonic-oscillator system. The parameters used are $g = 1$, $\eta_{P_{<\mu}} = 0$, $\eta = 0.05$, $\mu = 0.233$, and $\Delta = 0.244$. Different colored and shaped points indicated different degree polynomial approximations to the shifted error function. The horizontal black line indicates the maximum value of τ that still guarantees isolation of the ground state, which, for this choice of parameters, is $\tau_{\max} = 1.823$. Using $\tau = \tau_{\max}$ leads to a significant reduction in the error compared with using $\tau = 1$, while using values of $\tau > \tau_{\max}$ leads to significant excited-state contamination in the prepared state.

where $p_{\max} = \pi/(\delta x)$, $\delta p = 2\pi/(2^{n_q}\delta x)$, and $\mathcal{F}$ is the discrete Fourier transformation matrix. The superscripts $(\hat{p})$ and $(\hat{x})$ indicate that the momentum operator is written in the momentum and position basis, respectively. It will be useful for the circuit construction discussion to review the cost of exponentiation of $\hat{x}$ and $\hat{p}$. Because $\hat{x}$ and $\hat{p}^{(\hat{p})}$ are diagonal matrices with evenly spaced eigenvalues, $e^{ip_0\hat{x}}$ and $e^{ix_0\hat{p}^{(\hat{p})}}$ can be implemented using zero CNOT gates and n_q rotation gates [118]. One can then use the efficient quantum Fourier transform circuit to construct $e^{ix_0\hat{p}^{(\hat{x})}} = \mathcal{F}^\dagger e^{ix_0\hat{p}^{(\hat{p})}}\mathcal{F}$, with each Fourier transform requiring $O(n_q^2)$ gates [119].

The results in Fig. 5 are for a three qubit system with $g = 1$, $\eta_{P_{<\mu}} = 0$, $\eta = 0.05$, $\mu = 0.233$, and $\Delta = 0.244$. The maximum value of τ for these parameters is $\tau_{\max} = 1.823$. For various degree polynomials, we see that the fidelity in general improves for increasing τ up to $\tau_{\max}$, where increasing further leads to an increase in error. Relative to using $\tau = 1$, using $\tau = \tau_{\max}$ leads to a general improvement in precision by a factor of $O(\exp[d\Delta(\tau_{\max} - 1)])$.

We now discuss how using $\tau \neq 1$ can improve simulations when the approximating $e^{-i\tau H}$ using product formulas. For this discussion, it will be convenient to view τ as an evolution time. We denote by N_{steps} the number of Trotter steps per time evolution circuit, so that the building block of the QETU circuit is $(e^{-i\tau H/N_{\text{steps}}})^{N_{\text{steps}}}$. Each step is approximated using a first-order Trotter formula. If $H_x = \hat{x}^2/(2g^2)$ and $H_p = g^2\mathcal{F}^\dagger(\hat{p}^{(\hat{p})})^2\mathcal{F}/2$, then

$$\begin{aligned} e^{-i\delta\tau H} &\approx e^{-i\delta\tau H_x}e^{-i\delta\tau H_p} \\ &= e^{-i\delta\tau H_x}\mathcal{F}^\dagger e^{-i\delta\tau H_p^{(\hat{p})}}\mathcal{F}, \end{aligned} \tag{18}$$

where $\delta\tau = \tau/N_{\text{steps}}$. Recall that d denotes the total number of calls to the time evolution circuit (see Fig. 1), which corresponds to the polynomial of degree d. The total number of calls to the elementary time evolution circuit is then $N_{\text{tot}} = N_{\text{steps}}d$.

The error from approximating the step function and from a finite Trotter step size are denoted as ϵ_{QETU} and $\epsilon_{\text{Trotter}}$, respectively; the parameters are defined as before, such that the prepared ground state $|\widetilde{\psi}_0\rangle$ has overlap with the exact ground state given by $|\langle\widetilde{\psi}_0|\psi_0\rangle| = 1 - \epsilon_{\text{QETU}} - \epsilon_{\text{Trotter}}$. For concreteness, we assume the errors take the forms

$$\epsilon_{\text{QETU}} = ae^{-b(\tau\Delta)d}, \quad \epsilon_{\text{Trotter}} = c(\tau/N_{\text{steps}})^p, \tag{19}$$

where the first line is again obtained by inverting $d = \widetilde{O}(\Delta^{-1}\log(1/\epsilon_{\text{QETU}}))$, and the second line is the standard form for the error when using a pth-order Trotter formula. The parameters a, b, and c are constants. These expressions are expected to be correct to leading order. Replacing $d = N_{\text{tot}}/N_{\text{steps}}$, our total error is

$$\begin{aligned} \epsilon &= ae^{-b(\tau\Delta)N_{\text{tot}}/N_{\text{steps}}} + c(\tau/N_{\text{steps}})^p, \\ &= ae^{-b\Delta N_{\text{tot}}\delta\tau} + c(\delta\tau)^p. \end{aligned} \tag{20}$$

Notice that the total error ϵ only depends on N_{tot} and $\delta\tau$. Solving for N_{tot} gives

$$N_{\text{tot}} = \frac{1}{b\Delta\delta\tau}\log\left(\frac{a}{\epsilon - c(\delta\tau)^p}\right). \tag{21}$$

With this, we can now ask the question of what value of $\delta\tau$ minimizes the cost N_{tot} for some fixed error threshold ϵ. Before doing so, we emphasize the fact that N_{tot} depends only on $\delta\tau$, and not the individual values for N_{steps} and τ. One might expect that reducing τ would decrease the overall cost by reducing the number of Trotter steps per time evolution circuit N_{steps} for the same $\delta\tau$. However, reducing τ also shrinks the gap, and this increase in cost exactly cancels out the savings from decreasing N_{steps}.

To solve for $\delta\tau^*$ that minimizes N_{tot}, we first study the constraints on our parameters from the form of Eq. (21). The parameter N_{tot} is a positive real integer. Because the argument of the logarithm must be positive, we find $\delta\tau < (\epsilon/c)^{(1/p)}$. In words, one must choose a value of $\delta\tau$ such that the Trotter error is below the total target error ϵ. Furthermore, notice that the logarithm can return a negative result. A negative N_{tot} implies that one can achieve a precision of ϵ by setting $N_{\text{tot}} = 1$ and decreasing $\delta\tau$ until the target precision is achieved.

In Appendix B, a perturbative solution for $\delta\tau^*$ is presented. Setting $\frac{dN_{\text{tot}}}{d\tau}$ to zero and expanding the log to lowest order, we find $(\delta\tau^*)^p \approx \frac{\epsilon}{c}(1 - p/[\log(\frac{a}{\epsilon}) + p])$. Because $a > \epsilon$ in general, the lowest-order result for $\delta\tau^*$ is slightly below the maximum value of $(\epsilon/c)^{1/p}$. From this we learn that, for a given choice of N_{tot}, one should choose a time-step $\delta\tau$ such that most of the error comes from the Trotter error. This choice can be understood intuitively by comparing the rate of convergence of the two sources of error. Because the Trotter error converges as some power of $\delta\tau$, while the QETU error converges exponentially in N_{tot}, it is generally more cost effective for QETU to produce a smaller error than the Trotter error. Therefore, having the Trotter error

be the majority of the error results in the smallest value for N_{tot}.

Once the value of $\delta\tau^*$ has been chosen, one must choose values for N_{steps} and τ. One option would be to always set $N_{\text{steps}} = 1$ and $\tau = \delta\tau^*$. Another option would be to choose τ minimizing the degree of the Chebyshev polynomial achieving the target error ϵ, therefore reducing the classical cost of determining the angles $\{\varphi_j\}$. Since $N_{\text{tot}} = dN_{\text{steps}}$ is fixed, in order to decrease d, one needs to increase N_{steps} while ensuring that $\tau = N_{\text{steps}}\delta\tau^* \leqslant \tau_{\text{max}}$. This leads to the choice of τ being the largest natural number multiple of $\delta\tau^*$ that is less than τ_{max}. One possible caveat to keep in mind is that, because approximate time evolution using product formulas can be viewed as exact time evolution according to an effective Hamiltonian H_{eff} [120], one actually implements $F(\cos(\tau H_{\text{eff}}/2))$. Because the spectrum of the effective Hamiltonian will in general be different from the exact Hamiltonian by $O(\delta\tau^p)$, if one has shifted the spectrum of H to be in the range $[\eta/\tau, \pi - \eta]$, the spectrum of H_{eff} will be in the range $[\eta/\tau \pm O(\delta\tau^p), \pi - \eta \pm O(\delta\tau^p)]$. To avoid systematic errors from the possibility that the maximum energy of $H_{\text{eff}} > \pi - \eta$, one should decrease τ_{max} by $O(\delta\tau^p)$.

Figure 6 shows plots demonstrating the improvements one can achieve by choosing an optimal value of $\delta\tau$ from two perspectives, namely, fixed computational cost, and fixed target precision. The system studied is a compact U(1) lattice gauge theory in the weaved basis for a 2×2 lattice with $n_q = 2$ qubits per site using gauge coupling $g = 1$. The digitization scheme used is reviewed in Sec. III. The relevant parameters for state preparation are $\eta = 0.05$, $\eta_{P_{<\mu}} = 0$, $\mu = 0.1498$, and $\Delta = 0.1330$. The time evolution operator was implemented using a first-order Trotter method with $N_{\text{steps}} = 1$ Trotter steps, and therefore $\tau = \delta\tau$. The top plot in Fig. 6 shows ϵ as a function of $\delta\tau$ for different degree polynomial approximations. We see that for a fixed computational cost, significant precision improvements can be achieved by using an optimal choice for $\delta\tau$. The bottom plot in Fig. 6 studies the total computational cost needed to achieve a fixed target precision as a function of $\delta\tau$. The cost is given in terms of the total number of calls to the circuit that implements a single Trotter time-step. For all choices of fixed target precision, choosing an optimal value of $\delta\tau$ results in significant cost reductions.

We conclude this section with a discussion of additional improvements which can be gained with the aid of *zero error extrapolation* in cases when $e^{-i\tau H}$ is implemented approximately. For implementations based on product formulas, this would require running simulations at multiple values of ϵ and τ and extrapolating to $\epsilon \to 0$ and $\delta\tau \to 0$. Note that if the optimal value of $\delta\tau^*$ is chosen, the errors from QETU and from Trotter are comparable in magnitude. Therefore, extrapolating to $\epsilon \to 0$ and $\delta\tau \to 0$ would involve a fitting function of the form similar to that given in Eq. (20). While this is in principle possible, such a functional form is relatively complicated and undesirable. To simplify the extrapolation, one could instead work in a regime where the QETU error is negligible compared with the Trotter error. This would allow one to use a fitting function of simple form $c(\delta\tau)^p$. Because of QETU's fast convergence, it is possible that such slight increase in computational cost could be worth the trade-off of better control over systematic errors.

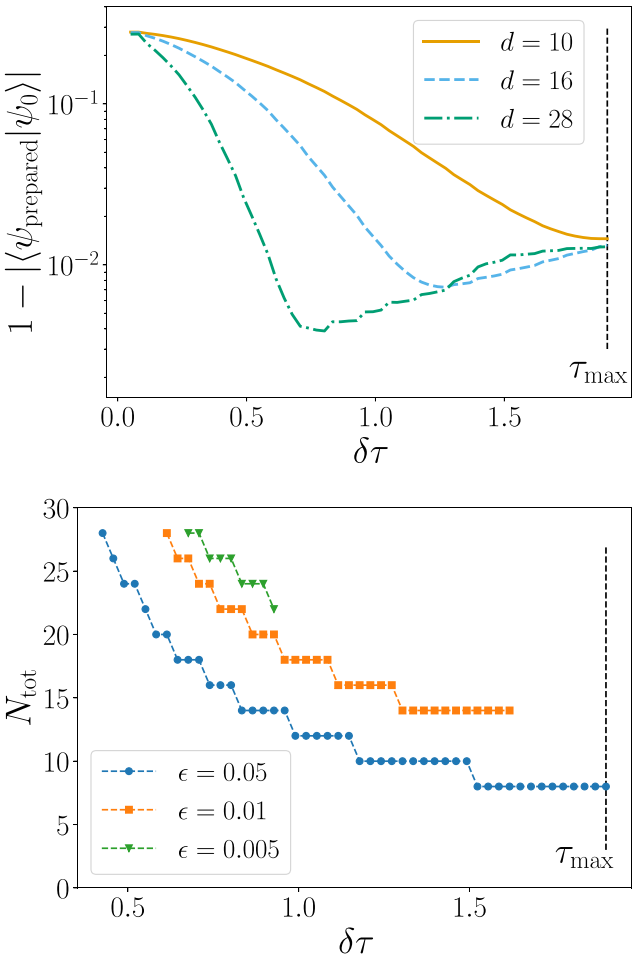

FIG. 6. Both plots show results for a compact U(1) lattice gauge theory in the waved basis for $N_p = 3$ and $n_q = 2$ with $g = 1$. The time evolution circuit was implemented using a single Trotter step with $\delta\tau = \tau$. The vertical black line indicates the value of τ_{max} for this system, with a value $\tau_{\text{max}} = 1.90$. (top) Fidelity of prepared ground state as a function of $\delta\tau$. Different colored lines indicate different degree polynomial approximations to the shifted error function. This figure demonstrates that choosing an optimal value of $\delta\tau$ results in significant precision improvements for a given computational cost. (bottom) Number of total calls to the circuit for a single first-order Trotter step as a function of $\delta\tau$. Each point is the smallest value of N_{tot} that achieves a given error ϵ, with different ϵ indicated by different colors and markers. The figure demonstrates that choosing an optimal value of $\delta\tau$ results in significant cost reductions for a fixed target precision.

C. Wave-packet construction via quantum eigenvalue transformation for unitary matrices

While the original application of the QETU algorithm in Ref. [98] was ground-state preparation, QETU is a flexible algorithm that can construct general matrix functions of any Hermitian operator. In this section, we provide a procedure for constructing Gaussian wave packets in the position basis of a quantum-mechanical system. We first describe the procedure at a high level. From there, we discuss the procedure for constructing an approximation to the Gaussian filter operator $e^{-\hat{x}^2/(2\sigma_x^2)}$ using QETU. We first discuss how the naïve

application of QETU to this problem leads to the error decreasing only polynomially with the number of Chebyshev polynomials. We then discuss a number of modifications to the QETU procedure that achieve an exponential scaling in the error for any desired value of the width. Using the improved methodology, we present a method that allows one to prepare the Gaussian state to high precision while avoiding the costly implementation of LCU, and find that this method outperforms existing methods for preparing Gaussian states for values of $n_q \gtrsim 2$–5. Throughout this section, we only discuss qualitative results; detailed scaling of the precision, as well as gate cost comparisons to direct state preparation, are presented in Sec. IV B. Note that the general method we employ is similar to that proposed in Ref. [90] for exactly constructing functions of diagonal operators.

The state we wish to prepare is given by

$$|\psi\rangle = \frac{1}{\sqrt{\mathcal{N}}} \sum_{j=0}^{2^{n_q}-1} e^{-\frac{1}{2}\left(\frac{x_j - x_0}{\sigma_x}\right)^2} e^{ip_0 x_j} |x_i\rangle, \tag{22}$$

where n_q is the number of qubits used to represent the state, $\hat{x}|x_i\rangle = x_i|x_i\rangle$, $\mathcal{N}$ is the normalization factor, x_0 is the central value of the Gaussian, p_0 is the expectation value of the momentum, and σ_x is the width of the wave packet in position space. To take advantage of the relative simplicity the $\hat{x}$ and $\hat{p}^{(\hat{x})}$ operators, we break the construction of the wave packet in Eq. (22) into three steps. The first, and the most costly step, is constructing a Gaussian with width σ_x centered at $x = 0$ using QETU. This will be done by applying an approximate implementation of the Gaussian filter operator $e^{-\frac{1}{2}\hat{x}^2/\sigma_x^2}$ to the state that is an equal superposition of all position eigenstates, i.e., $|\psi_{\text{init}}\rangle = \frac{1}{(2^{n_q})^{1/2}} \sum_{j=0}^{2^{n_q}-1} |x_i\rangle$. The wave packet is then shifted in position space by x_0 and in momentum space by p_0 with the aid of the operators $e^{-ix_0\hat{p}^{(\hat{x})}}$ and $e^{-ip_0\hat{x}}$, respectively.

To construct the approximate Gaussian filter operator, which is a matrix function of the $\hat{x}$ operator, we use QETU. The building block used in the QETU circuit is $e^{-i\tau\hat{x}_{\text{sh}}}$, where $\hat{x}_{\text{sh}}$ is the shifted and scaled position whose spectrum is in the range $[\eta, \pi - \eta]$. Because $\hat{x}_{\text{sh}}$ is also a diagonal matrix with evenly spaced eigenvalues, as previously discussed, $e^{-i\tau\hat{x}_{\text{sh}}}$ can be implemented exactly, using zero CNOT gates and n_q rotation gates. The controlled version can therefore be implemented using n_q CNOT and n_q rotation gates, leading to an asymptotic scaling linear in the number of qubits used to represent the operator, albeit with a large overall prefactor. In this context, we find it most natural to view τ as a parameter that scales the spectrum of $\hat{x}_{\text{sh}}$ rather than an evolution time.

The operator $\hat{x}_{\text{sh}}$ is given by

$$\hat{x}_{\text{sh}} = c_1\hat{x} + c_2, \tag{23}$$

where c_1 and c_2 are the same as in Eq. (4), upon replacing $E_0^{\text{phys}} \to \min(\hat{x})$ and $E_{\max}^{\text{phys}} \to \max(\hat{x})$. Note that because the spectrum of $\hat{x}$ is known, one does not have to use upper limits as is generally the case with state preparation where the spectrum of the Hamiltonian is not known *a priori* (this fact will also allow for other improvements, and will discussed in more detail later in this section).

Let $\mathsf{f}(\hat{x}) = c\, e^{-\frac{1}{2}\hat{x}^2/\sigma_x^2}$ denote the exact Gaussian filter operator we wish to approximate using QETU. Again, the parameter c is chosen to be slightly less than one to allow overshooting of the Chebyshev approximation. To produce a Gaussian centered at $x_0 = 0$ with width σ_x using $e^{-i\tau\hat{x}_{\text{sh}}}$ as a building block, one must approximate the function

$$\mathsf{F}(x) = c\, \exp\left[-\frac{1}{2\sigma_{\text{QETU}}^2}\left(\frac{2}{\tau}\arccos(x) - x_0^{\text{QETU}}\right)^2\right], \tag{24}$$

where $x_0^{\text{QETU}} = c_1x_0 + c_2$ and $\sigma_{\text{QETU}} = c_1\sigma_x$. Note that because the function in Eq. (24) does not have definite parity, one must use QETU to first prepare approximations to the even $[\mathsf{F}_+(x)]$ and odd $[\mathsf{F}_-(x)]$ parts separately, and then add them using, e.g., linear combinations of unitaries (LCUs) [121]. While performing an LCU does not change the overall scaling with n_q, it does lead to a larger constant factor in the asymptotic cost of preparing the Gaussian state.

We now describe the procedure to determine the Chebyshev approximation of $\mathsf{F}_+(x)$. The procedure for $\mathsf{F}_-(x)$ is analogous, except that one uses odd Chebyshev polynomials. Let $F_+(x) = \sum_{k=0}^{d/2} c_{2k}T_{2k}(x)$ denote the Chebyshev approximation to $\mathsf{F}_+(x)$. As before, we define the coefficient matrix $A_{jk} = T_{2k}(x_j)$, such that $F_+(x_j) = \sum_k c_{2k}A_{jk}$. The coefficients c_k are then determined by solving the following convex optimization problem:

$$\min_{\{c_k\}} \max_{x_j \in [\sigma_{\min}, \sigma_{\max}]} |F_+(x_j) - \mathsf{F}_+(x_j)|, \tag{25}$$

where the functions are sampled using the roots of Chebyshev polynomials $x_j = -\cos(\frac{j\pi}{M-1})$ for some large value of M. The main difference between this optimization problem and the one in Eq. (9) is that the error in Eq. (25) is minimized over the entire range of x values in $[\sigma_{\min}, \sigma_{\max}]$.

We now discuss how the error is expected to scale as the degree of the Chebyshev approximation is increased. Unlike the case of constructing the shifted error function for ground-state preparation, the functions $\mathsf{F}_\pm(x)$ are not infinitely differentiable on the interval $x \in [-1, 1]$ due to the presence of the $\arccos(x)$. It is well known that Chebyshev approximations of functions that are not infinitely differentiable converge polynomially to the true function with the typical rate $(1/n_{\text{Ch}})^m$, where m is the number of times the function is differentiable on $[-1, 1]$, and n_{Ch} is the number of Chebyshev polynomials used in the approximation [113]. We therefore expect the convergence of our Chebyshev approximations to be only polynomial.

This observation, however, leads to a natural method for avoiding the polynomial scaling. As one increases the value of η, the interval $x_j \in [\sigma_{\min}, \sigma_{\max}]$ will move farther from the nondifferentiable points at $x = \pm 1$ and improve the precision of approximation for a given degree polynomial. However, as η is increased, the functions $\mathsf{F}_\pm(x)$ become more sharply peaked and will eventually require more Chebyshev polynomials to achieve the same precision. We therefore expect that, for a given degree approximation, there will be a value of η that results in the smallest error. We find that this procedure of varying η results in the error decreasing more favorably with the degree of the polynomial. This scaling can made exponential, however, by also varying the τ parameter. By varying τ, the shape of the exact functions $\mathsf{F}_\pm$ change, and can improve the quality of the approximation. Note that, unlike the

case of constructing the filter operator, there are no theoretical issues with using a value of $\tau > \tau_{\max}$. This is due to the fact that we know exactly the values of the operator we sample. The optimization problem that determines the Chebyshev coefficients using this modification is

$$\min_{\{c_k,\eta,\tau\}} \max_{x_j\in[\sigma_{\min},\sigma_{\max}]} |F_+(x_j) - \mathsf{F}_+(x_j)|, \tag{26}$$

where the parameters $\sigma_{\min}$ and $\sigma_{\max}$ are functions of η and τ. Note that the optimization problem in this form is no longer a convex optimization problem. The parameters η and τ are found by passing the convex optimization problem in Eq. (25) to a numerical minimization procedure that solves for η and τ.

While varying η and τ already results in the error decreasing exponentially with the degree of the polynomial, we can further improve the rate of convergence, in particular for small values of n_q. Because we know the spectrum of the operator $\hat{x}_{\rm sh}$ exactly, we only need the Chebyshev expansion to approximate $\mathsf{F}(x)$ at those points, not the entire range $x \in [\sigma_{\min}, \sigma_{\max}]$. Taking into account the cosine transformation, the values of x we need to faithfully approximate $\mathsf{F}(x)$ are $\tilde{x}_j = \cos(\tau x_{{\rm sh},j}/2)$, where $x_{{\rm sh},j}$ is the jth eigenvalue of the $\hat{x}_{\rm sh}$ operator. If we denote the set of all $\tilde{x}_j$ values as $\tilde{\chi}$, then the Chebyshev coefficients are found by solving the modified optimization problem

$$\min_{\{c_k,\eta,\tau\}} \max_{\tilde{x}_j\in\tilde{\chi}} |F_+(x_j) - \mathsf{F}_+(x_j)|. \tag{27}$$

We find that, in general, this method has slightly better performance compared with sampling all x values, but for certain values of $\sigma_x/x_{\max}$ it can improve the rate of convergence significantly. Additionally, because we only include 2^{n_q} of $\tilde{x}_j$ values in the optimization, the error will be zero when the degree of the polynomial is equal to the number of points. Furthermore, by increasing τ, one can sample $\tilde{x}_j$ at negative values. This can be used to exploit the parity of the Chebyshev expansions, reducing the effective number of points we need to approximate $\mathsf{F}(\tilde{x}_j)$. One important consideration when using this method is that, as one varies η and τ, care must be taken to ensure that $|F_\pm(x)| \leqslant 1$ for all $x \in [-1, 1]$, and not just for $\tilde{x}_i$. If this condition is not satisfied, there are no possible values of phases $\{\varphi_j\}$ that implement the desired function $F(x)$.

Although we have developed a highly efficient procedure using QETU to construct Gaussian states, the cost of performing LCU to add the even and odd pieces introduces a large overall constant factor in the asymptotic cost. If one could modify the procedure to avoid using an LCU, the gate count reduction would be reduced by a factor of ten or more. We now discuss how to prepare Gaussian states using only the even component and eliminate the need for an LCU altogether. The main idea is to choose τ such that the function $\mathsf{F}(x)$ becomes a purely even function. To start, note that our choice of digitizing the $\hat{x}$ operator results in the parameter $c_2 = \pi/2$. This combined with the fact that $x_0 = 0$ leads to $x_0^{\rm QETU} = \pi/2$. With this, we see that setting $\tau = 2$ leads to $\mathsf{F}(x)$ becoming

$$\begin{aligned}\mathsf{F}(x) &= c\, e^{-[\arccos(x)-\pi/2]^2/(2\sigma^2_{\rm QETU})} \\ &= c\, e^{-[\arcsin(x)]^2/(2\sigma^2_{\rm QETU})},\end{aligned} \tag{28}$$

where we have used the relation $\arccos(x) - \pi/2 = -\arcsin(x)$. Because $\arcsin(x)$ has definite parity, $\mathsf{F}(x)$ is an even function for $\tau = 2$. After setting $\tau = 2$, the parameter η and the Chebyshev coefficients are found by solving an optimization problem similar to that in Eq. (27), except that τ is fixed to two. We find that this method offers the best precision for a given fixed gate count cost and outperforms existing methods for Gaussian state preparation for values of $n_q > 2$–3. More details of this comparison are given in Sec. IV B.

III. LATTICE FORMULATION OF U(1) GAUGE THEORY

In this section, we review a formulation of a compact U(1) gauge theory in two spatial dimensions with a Hilbert space that has been constrained to satisfy Gauss's law in the charge density $\rho(x) = 0$ sector. We also review the representation of the magnetic and electric operators introduced in Ref. [41], which can be used at all values of the gauge coupling. From there, we discuss the basis change introduced in Ref. [100], which is necessary to break the exponential volume scaling in the gate cost when performing time-evolution using Trotter methods. We conclude with a discussion of how the digitization of the theory in the weaved basis must be modified in order to maintain the efficiency of the representation [101]. Note that we provide only the details necessary to understand the numerical results presented in Sec. IV A. Further details can be found in the original references.

We chose the apply the QETU algorithm to the compact U(1) theory because it shares a number of important features with QCD, including the fact that, as the lattice spacing a goes to zero, the bare coupling $g(a)$ also goes to zero. Additionally, the gauge field in non-Abelian gauge theories is necessarily compact, proving another motivation for detailed studies of the compact U(1) theory.

The Hamiltonian considered in Ref. [41] is formulated in terms of electric rotor and magnetic plaquette operators, given by $\hat{R}(x)$ and $\hat{B}(x)$, respectively. These operators satisfy

$$[\hat{B}(x), \hat{R}(y)] = i\,\delta^3(x - y). \tag{29}$$

In the charge density $\rho(x) = 0$ sector, the rotors are defined such that $\vec{E}(x) = \vec{\nabla} \times R(x)$. In this way, electric Gauss's law is automatically satisfied.

The lattice version of the theory we consider introduces a periodic lattice of N_x and N_y evenly spaced lattice points in the $\hat{x}$ and $\hat{y}$ dimensions with a lattice spacing a. The lattice version of the continuum Hamiltonian is defined in terms of operators, $\hat{R}_p$ and $\hat{B}_p$, with the index p denoting a specific plaquette in the lattice volume. The Hamiltonian can be written in terms of an electric and magnetic component as

$$\hat{H} = \hat{H}_E + \hat{H}_B. \tag{30}$$

After solving the constraint from magnetic Gauss's law, the fully gauge fixed Hamiltonian contains $N_p \equiv N_x N_y - 1$ independent plaquette operators. The fully gauge-fixed electric and compact magnetic Hamiltonians are given by

$$\hat{H}_E = \frac{g^2}{2a} \sum_{p=1}^{N_p} (\vec{\nabla} \times \hat{R}_p)^2, \tag{31}$$

and

$$\hat{H}_B = \frac{N_p+1}{a\,g^2} - \frac{1}{a\,g^2}\left[\sum_{p=1}^{N_p}\cos\hat{B}_p + \cos\left(\sum_{p=1}^{N_p}\hat{B}_p\right)\right]. \quad (32)$$

The noncompact formulation of this theory can be found by making the replacements $\cos\hat{B}_p \to 1 - \frac{1}{2}\hat{B}_p^2$. The bases where the operators $\hat{H}_E$ and $\hat{H}_B$ are diagonal are referred to as the electric and magnetic basis. Furthermore, operators represented in the electric and magnetic basis will be denoted by superscripts (e) and (m), respectively.

We represent the operators $\hat{R}_p$ and $\hat{B}_p$ using the procedure given in Ref. [41], which involves carefully choosing the maximum value that the $\hat{B}_p$ operators are sampled, denoted by $b_{\max}$. The main idea for choosing $b_{\max}$ is that, in the magnetic basis, the low-energy eigenstates of the theory have a typical width proportional to the gauge coupling g. By choosing $b_{\max} \sim g$, one only samples the wave function where it has support, which results in an efficient representation for all values of g. We first describe how the operators are sampled, and then explain how to choose the precise value of $b_{\max}$.

Each plaquette is represented using n_q qubits. Working in magnetic basis, the magnetic operators $\hat{B}_p^{(m)}$ are diagonal and defined by

$$\hat{B}_p^{(m)}|b_{p,j}\rangle = (-b_{\max,p} + j\,\delta b_p)|b_{p,j}\rangle, \quad (33)$$

where $j = 0, 1, \ldots, 2^{n_q}-1$ and $\delta b_p = 2b_{\max,p}/2^{n_q}$. Because $\hat{B}_p$ and $\hat{R}_p$ are conjugate operators, the rotor operator in the magnetic basis can be written as

$$\hat{R}_p^{(m)} = \mathcal{F}^\dagger \hat{R}_p^{(e)} \mathcal{F}, \;\; \hat{R}_p^{(e)}|r_{p,j}\rangle = (-r_{\max,p} + j\,\delta r_p)|r_{p,j}\rangle, \quad (34)$$

where $\mathcal{F}$ denotes the usual quantum Fourier transform and

$$r_{\max,p} = \frac{\pi 2^{n_q}}{2b_{\max,p}}, \;\; \delta r_p = \frac{\pi}{b_{\max,p}}. \quad (35)$$

With these definitions, one can implement Trotter time evolution of the Hamiltonian in Eq. (30) in a similar way as in Eq. (18).

We now discuss the procedure for choosing $b_{\max,p}$, which can in general be different for different plaquettes p. In the compact formulation, it was shown that choosing $b_{\max,p}$ according to

$$b_{\max,p} = \min\left(g\frac{2^{n_q}}{2}\sqrt{\frac{\beta_{R,p}}{\beta_{B,p}}}\sqrt{\frac{2\pi}{2^{n_q}}}, \pi\right) \quad (36)$$

reproduces the low-lying spectrum of the theory to per-mille level precision while only sampling the operators a number of times that corresponds to $n_q = 3$. The variables $\beta_{R,p}$ and $\beta_{B,p}$ are found by matching the noncompact magnetic Hamiltonian to a Hamiltonian of the form

$$\widetilde{H} = \frac{g^2}{2}\beta_{R,p}^2\hat{R}_p^2 + \frac{1}{2g^2}\beta_{B,p}^2\hat{B}_p^2, \quad (37)$$

and ignoring the cross terms. Further details regarding the digitization of the $\hat{R}_p$ and $\hat{B}_p$ operators in this formulation can be found in Refs. [41,101].

While this formulation is efficient in terms of the number of qubits required per site n_q to achieve a high precision in the low-energy states, it was shown that performing time evolution using Trotter methods has a gate cost that scales exponentially with volume, i.e., the number N_p of plaquettes [100,101]. Specifically, the exponential scaling was shown to be caused by the $\cos\sum_p \hat{B}_p$ term in the magnetic Hamiltonian, which couples the entire lattice together. This exponential volume scaling can be broken, however, by performing a carefully chosen change of operator basis [100], which we now review.

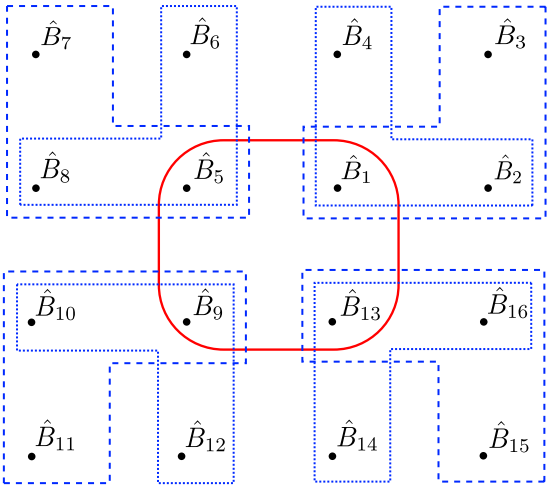

FIG. 7. Visual representation of the connectivity of the cosine terms in the weaved magnetic Hamiltonian for $N_p = 16$. Operators appearing inside the same box also appear as a sum inside a single cosine term in $\hat{H}_B$. Boxes with different line-styles or colors correspond to different cosine terms. The red solid square shows the reduced connectivity of the $\cos\sum_p \hat{B}_p$ term. The blue dashed and dotted rectangles show the increased connectivity of the previously local $\cos\hat{B}_p$ terms.

The rotor and magnetic operators in this so-called weaved basis are given by

$$\hat{B}_p \to \mathcal{W}_{pp'}\hat{B}_{p'}, \;\; \hat{R}_p \to \mathcal{W}_{pp'}\hat{R}_{p'}, \quad (38)$$

where $\mathcal{W}$ is an orthogonal matrix of dimension $N_p \times N_p$. For any value of N_p, there exists an efficient classical algorithm for choosing $\mathcal{W}$ that reduces the gate count scaling from exponential to polynomial in N_p [100]. Using this change of basis, the dominant contribution to the gate cost of a single Trotter step was shown to scale as $O(N_p^{n_q} + N_p(N_p/\log_2 N_p)^{n_q})$, which is polynomial in the volume, with the power of the polynomial determined by n_q. This scaling arises because the number of terms appearing in a single cosine in the weaved basis scales as $O(\log_2 N_p)$. An example demonstrating the connectivity of the magnetic Hamiltonian operators for $N_p = 16$ is shown in Fig. 7.

One important assumption that went into choosing the large coupling limit of $b_{\max,p}$ to be π in the original operator basis was that the coefficient of a magnetic-field operator $\hat{B}_p$ is equal to one anywhere it appears in the compact magnetic Hamiltonian. Because this is generally not true when working in the weaved basis, maintaining an efficient representation requires modifying the prescription for choosing $b_{\max,p}$ in the large-g limit [101]. To understand this, first notice that, in the weaved basis, some of the $\hat{B}_p$ operators will have coefficients smaller than one. Consequently, even if $b_{\max,p} = \pi$, an

operator $\hat{B}_p$ inside a given cosine will not get sampled between the full range of $[-\pi, \pi]$. It was shown that this problem could be fixed by scaling the upper limit for each $b_{\max,p}$ by the smallest coefficient of the operator $\hat{B}_p$ anywhere it appears in the Hamiltonian. To demonstrate this procedure, we review the example for the $N_p = 3$ case given in Ref. [101]. The rotation matrix used is given by

$$\mathcal{W} = \frac{1}{\sqrt{6}} \begin{pmatrix} \sqrt{2} & -2 & 0 \\ \sqrt{2} & 1 & -\sqrt{3} \\ \sqrt{2} & 1 & \sqrt{3} \end{pmatrix}, \tag{39}$$

which leads to the following weaved magnetic Hamiltonian

$$\begin{aligned} \hat{H}_B^{\mathrm{w}} = \frac{N_p + 1}{a\, g^2} - \frac{1}{a\, g^2} \Bigg(& \cos[\sqrt{3}\hat{B}_1] + \cos\left[\frac{\hat{B}_1 - \sqrt{2}\hat{B}_2}{\sqrt{3}}\right] \\ & + \cos\left[\frac{\sqrt{2}\hat{B}_1 + \hat{B}_2 - \sqrt{3}\hat{B}_3}{\sqrt{6}}\right] \\ & + \cos\left[\frac{\sqrt{2}\hat{B}_1 + \hat{B}_2 + \sqrt{3}\hat{B}_3}{\sqrt{6}}\right] \Bigg). \end{aligned} \tag{40}$$

To ensure that each $\hat{B}_p$ operator gets sampled between $[-\pi, \pi]$ in each of the cosine terms, one must scale the upper limit for each $b_{\max,p}$; in this case, the upper limits for $b_{\max,1}$, $b_{\max,2}$, and $b_{\max,3}$, are chosen as $\sqrt{2}\pi$, $\sqrt{6}\pi$, and $\sqrt{3}\pi$, respectively. It was shown in Ref. [101] that this procedure for scaling $b_{\max,p}$ results in a precision of the low-lying spectrum similar to that of the original basis.

We conclude this section by discussing the expected behavior of Δ with g for the compact U(1) gauge theory. As explained in Ref. [122], because the untruncated U(1) electric Hamiltonian is unbounded, as one approaches the continuum limit, the physical energy gap $E_1^{\mathrm{phys}} - E_0^{\mathrm{phys}}$ diverges; the value $a(E_1^{\mathrm{phys}} - E_0^{\mathrm{phys}})$ approaches a constant as $a \to 0$. This, combined with the fact that the physical energy difference $E_{\max}^{\mathrm{phys}} - E_0^{\mathrm{phys}}$ scales as $\approx 1/a$, implies that Δ will approach a constant as $a \to 0$. This scaling is qualitatively different than that of gauge theory with a finite physical energy gap (such as QCD), which is discussed in detail in Sec. V.

IV. NUMERICAL RESULTS

A. U(1) gauge theory

In this section, we prepare the ground state of the previously described formulation of a compact U(1) lattice gauge theory using QETU. We study the algorithm cost by determining how its parameters Δ and γ depend on the parameters of the Hamiltonian: the number of plaquettes N_p, number of qubits per plaquette n_q, and gauge coupling g (or lattice spacing a). We also study how the fidelity of the final state scales with the number of calls to $e^{-i\tau H}$, where the time evolution operator is implemented both exactly as well as using Trotter methods. The scaling results from our analysis are summarized in Table II. A general discussion regarding the physical reasoning for the observed scaling of these parameters is given in Sec. V.

TABLE II. Scaling of parameters in the cost of state preparation using QETU in terms of N_p, n_q, and g. The γ parameter defines the number of measurements needed to measure the ancillary qubit in the zero state, which scales as $O(1/\gamma^2)$. This scaling can be can be improved to $O(1/\gamma)$ using amplitude amplification at the cost of increasing the circuit depth by a factor of γ^{-1}. *Note that the dependence of γ on N_p and n_q is highly sensitive to the value of coupling constant g: for $g = 1.2$, γ scales with neither N_p nor with n_q, see Eq. (9). The Δ parameter defines the query depth of the time evolution circuit $e^{-i\tau H}$, which scales as $O(\Delta^{-1})$. Note that while the gauge coupling g does not appear in the asymptotic scaling of the γ and Δ parameters, their values still depend on on g. The $\tau/\delta\tau$ parameter defines the number of Trotter steps used when approximating $e^{-i\tau H}$, which, surprisingly, does not scale with N_p or n_q. The Gates$(e^{-i\delta\tau H})$ parameter defines the number of gates required to implement a single Trotter step for the particular formulation of U(1) gauge theory we consider [100,101]. This unusual scaling is due to the fact that the theory we consider is highly nonlocal.

Parameter	Scaling
γ	$O(\exp(-N_p)\exp(-n_q))^*$
Δ	$O\left(N_p^{-1} 2^{-2n_q}\right)$
$N_{\mathrm{steps}} = \tau/\delta\tau$	$O(1)$
Gates$(e^{-i\delta\tau H})$	$O\left(N_p^{n_q}\right)$

Preparing the ground state requires knowledge of aE_0^{phys}, aE_1^{phys}, and $aE_{\max}^{\mathrm{phys}}$. In a realistic quantum simulation, one would likely be able to estimate aE_0^{phys} beforehand. Calculating aE_1^{phys} and $aE_{\max}^{\mathrm{phys}}$ is generally more difficult; still, one can often use physical arguments to bound both $a(E_1^{\mathrm{phys}} - E_0^{\mathrm{phys}})$ and $E_{\max}^{\mathrm{phys}}$. For the purposes the current study, we calculate aE_0^{phys} and aE_1^{phys} using exact diagonalization. Regarding $aE_{\max}^{\mathrm{phys}}$, we provide arguments for placing upper bounds on its value, and compare our bounds to the exact result. One important consideration is that, because Δ is a ratio of energies, the explicit dependence on the lattice spacing a cancels. The value of Δ therefore only depends on the lattice spacing a through discretization effects.

We begin by placing an upper bound on $aE_{\max}^{\mathrm{phys}}$. In the digitization scheme we use, one can write the U(1) Hamiltonian as $H^{\mathrm{phys},(b)} = \mathcal{F}^\dagger H_E^{(e)} \mathcal{F} + H_B^{(b)}$, where the superscript e (b) indicates that the matrix is represented in the electric (magnetic) basis, where it is implied the Fourier transform is performed locally at each lattice site. In this section, it is understood that the symbols H_E and H_B denote the unscaled Hamiltonians. Using the fact that the Fourier transform is unitary and therefore does not change the eigenvalues of $H_E^{(e)}$, we see

$$\begin{aligned} \max(H^{\mathrm{phys}}) &\leqslant \max\left(\mathcal{F}^\dagger H_E^{(e)} \mathcal{F}\right) + \max\left(H_B^{(b)}\right) \\ &= \max\left(H_E^{(e)}\right) + \max\left(H_B^{(b)}\right), \end{aligned} \tag{41}$$

where $\max(A)$ denotes the maximum eigenvalue of A. Because $H_E^{(e)}$ and $H_B^{(b)}$ are diagonal matrices, the maximum eigenvalue of each matrix is simply the maximum diagonal entry. To proceed, we must look in more detail at the forms of the two Hamiltonians.

Because the magnetic Hamiltonian is a constant term $(N_p+1)/(ag^2)$ minus a sum of N_p+1 cosine terms, the maximum value any single diagonal entry can take is $2(N_p+1)/(ag^2)$. This upper bound, however, overestimates the actual value, especially at small values of g. To understand this, recall that the magnetic operators are sampled from $-b_{\max}$ to $b_{\max}$, where $b_{\max} \sim g$. As $g \to 0$, the compact magnetic Hamiltonian approaches the noncompact version, with each term of the form B^2/g^2. If $B \sim g$, then $B^2/g^2 \approx 1$ is roughly independent of g, while our upper bound scaled as $1/g^2$. This issue can be avoided by taking advantage of the structure of the weaved magnetic Hamiltonian. It was shown in Ref. [100] that, after changing to the weaved basis, each cosine term in the magnetic Hamiltonian contains a sum of no more than $O(\log_2 N_p)$ magnetic-field operators. Thus, the spectrum of each individual cosine term can be found exactly using classical resources that scale only polynomially with N_p. Once the maximum entry of each individual term is known, we can then place an upper bound on the maximum energy of H_B through

$$\max(H_B^{(b)}) \leqslant \frac{1}{a\,g^2}\left[(N_p+1)+\sum_{j=0}^{N_p-1}\max(-\cos\widetilde{B}_j) + \max\left(-\cos\sum_{j=0}^{N_p-1}\widetilde{B}_j\right)\right], \tag{42}$$

where $\widetilde{B}_j = \mathcal{W}_{jj'}B_{j'}$ is the jth magnetic operator in the weaved basis.

In a similar way, an upper bound of the maximum value of H_E can also be found. In both the original and weaved basis, we can write the Hamiltonian generally as $g^2/2\sum_{i,j=0}^{N_p-1} c_{ij}R_iR_j$ (note that many of the c_{ij} are zero). Because each R_iR_j term in this sum is a $2^{2n_q}\times 2^{2n_q}$ diagonal matrix, we can explicitly evaluate the spectrum of each term classically at a cost quadratic in N_p. The upper bound placed on the spectral norm of H_E is given by

$$\max\left(H_E^{(e)}\right) \leqslant \frac{g^2}{2a}\sum_{i,j=0}^{N_p-1}\max(c_{ij}R_iR_j). \tag{43}$$

Our final upper bound on the maximum energy of the full Hamiltonian is then found using Eq. (41). Using this upper bound, combined with the exact values for aE_1 and aE_0, we can place a lower bound on Δ. This lower bound is compared with the exact value in Fig. 8 as a function of N_p, n_q, and g. We discuss each plot individually.

The top plot shows the exact value and upper bound of Δ as a function of N_p for $n_q = 2$ and $g = 1.4$. Before discussing the results, we point out that $N_p = 3, 5, 7$ correspond to lattices with sites $N_x \times N_y$ of 2×2, 2×3, and 2×4, respectively. Due to the inherent limitations of classical simulation, we only increase the number of sites in a single dimension, and so we expect the finite-volume errors to remain roughly the same size for all values of N_p. From the plot we see that the upper bound is larger than the exact value, with the difference generally growing with N_p. The overall scaling of Δ is roughly as $1/[N_p\log^2(N_p)]$. To understand this behavior, we can study how the number of terms in the Hamiltonian grows with N_p. For H_B, the number of terms grows linearly with N_p, implying

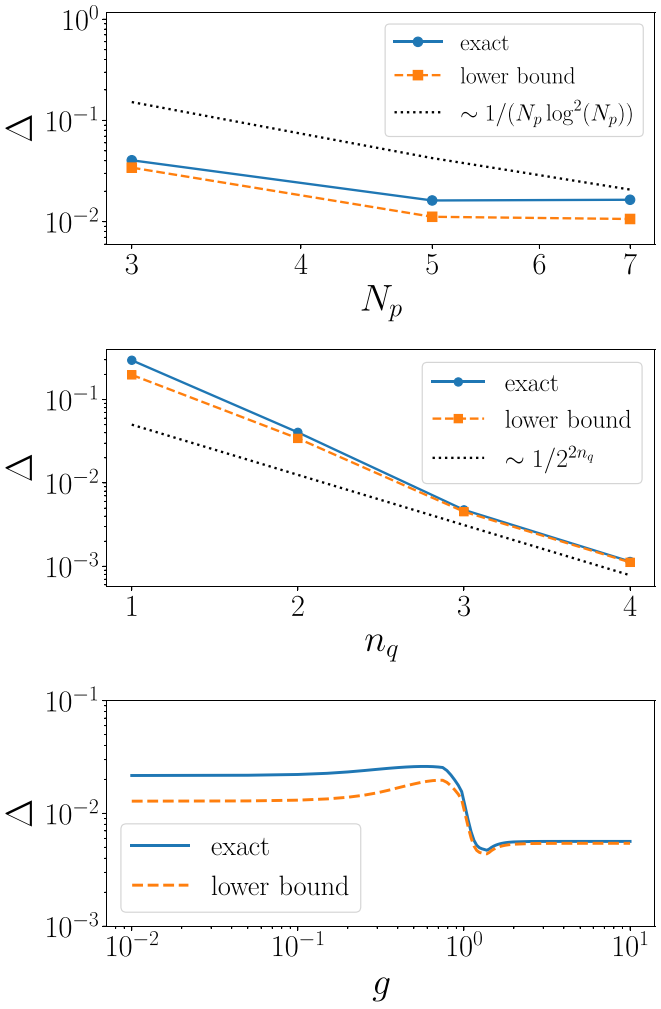

FIG. 8. Comparison of the exact value of Δ to the upper bound calculated using the procedure in Sec. IV A. (top) Δ as a function of N_p using $n_q = 2$ and $g = 1.4$. The gap scales asymptotically as $1/[N_p\log^2(N_p)]$. (middle) Δ as a function of n_q using $N_p = 3$ and $g = 1.4$. The gap generally decreases as $1/2^{2n_q}$. (bottom) Δ as a function of g using $N_p = 3$ and $n_q = 3$. The gap is generally independent of g, with a dip near $g \approx 1$, which is due to the gap $a(E_1 - E_0)$ becoming small. This behavior near $g = 1$ is an artifact of using the weaved basis.

that the maximum entry in H_B scales linearly with N_p as well. While the number of terms in H_E depends on the specific weaved matrix used, general statements can be made about the scaling. As demonstrated in Ref. [100], in order to break the exponential volume scaling in the gate count, a single rotor operator in the original basis is generally expressed as $O(\log N_p)$ operators in the weaved basis. Because H_E is a sum of N_p terms that are squares of differences of rotors, this leads to the number of terms scaling as $O(N_p\log_2^2 N_p)$. For these reasons, we expect gap to scale roughly as $N_p\log_2^2 N_p$. However, for smaller values of g, we expect that H_B and H_E become of similar magnitude, and cancellations between H_B and H_E could lead to a milder scaling with N_p.

The middle plot shows the exact value and upper bound of Δ as a function of n_q for $N_p = 3$ and $g = 1.4$. These results show that the quality of the upper bound generally increases

with n_q. Furthermore, we see that the lower bound generally scales as $O(2^{-2n_q})$. To understand this scaling, we start with H_B, which is a sum of the cosine of sums of magnetic operators. Even though the number of Pauli-Z operators in each magnetic operator grows exponentially with n_q, the maximum value a given cosine can take is one, regardless of n_q. For H_E on the other hand, the rotor operators do not appear inside a cosine, and therefore the maximum eigenvalue grows with the number of terms. Additionally, the maximum eigenvalue of a single rotor operator scales as $O(2^{n_q})$, as seen from the relation $r_{\max,p} = \pi 2^{n_q}/(2b_{\max,p})$. Because each term in H_E is bilinear, the maximum eigenvalue of each term grows as $O(2^{2n_q})$, leading to the observed scaling. For smaller values of g, however, we expect that H_B and H_E become of similar magnitude, and the scaling with n_q will likely be more mild due to cancellations between H_B and H_E.

Lastly, the bottom plot shows Δ as a function of g (which is a function of the lattice spacing a) for $N_p = 3$ and $n_q = 3$. First, notice that the quality of the upper bound is higher for large g. Second, except for the region of $g \approx 1$, the value of Δ is roughly independent of g. The roughly constant behavior for small and large g can be understood by the weak- and strong-coupling limits of the Hamiltonian. For large g, the electric Hamiltonian dominates. This, combined with the fact that $b_{\max,p}$ approaches a constant for large g and $H_E \sim g^2$, leads to the spectrum of H_E increasing as g^2 for large g. Because Δ depends only on ratios of energy differences, the g dependence cancels, and we expect Δ to approach a constant for large g. Similarly, Δ approaches a constant for small g, which can be understood by recalling that as $g \to 0$, the compact theory approaches the noncompact version. The noncompact theory is a free theory, i.e., a theory of noninteracting harmonic oscillators, with the gauge coupling g playing the role of the mass m of the canonical quantum harmonic oscillator. The spectrum of the compact U(1) Hamiltonian in the small-g limit is therefore independent of g, leading to the observed behavior. The large dip near $g \approx 1$ is an artifact of using the weaved basis, and is not present in the original basis. It was found in Ref. [101] that controlling digitization errors near $g = 1$ required a tuning of the choices of $b_{\max,p}$, which we did not perform here. It is therefore possible that, after performing this tuning, the dip near $g = 1$ will disappear.

We conclude this discussion by pointing out that, even though one can argue how Δ scales with various parameters, significant savings can still be achieved by either improving the lower bound, or performing a dedicated calculation to determine the exact value. Our studies indicate that the cost reduction of such a study will be more significant as one increases N_p towards realistic values, and at small g.

Next, we study how the parameter γ scales with N_p, n_q, and g. As will be argued in Sec. V, using direct state preparation methods to implement the initial guess wave function $|\psi_{\text{init}}\rangle$ results in an overlap γ that is exponentially suppressed in the number of sites. In this study, we instead consider implementing $|\psi_{\text{init}}\rangle$ using adiabatic state preparation, with the objective of developing a preliminary understanding of how γ scales with N_p and n_q.

The general strategy is to start in the ground state of the large-coupling limit of our theory. At large g, the electric Hamiltonian dominates and the ground-state approaches a state with constant entries; this state can be prepared by applying a Hadamard gate on each qubit. Once the large coupling ground state is prepared, the ground state of the target theory at some smaller coupling is prepared by adiabatically evolving between the two Hamiltonians.

We follow closely the procedure and notation given in Ref. [123]. The initial strong-coupling Hamiltonian is denoted by H_1 and the target Hamiltonian by H_2, with gauge couplings g_1 and g_2, respectively. The adiabatic evolution is performed according to a time-dependent Hamiltonian $H[u(t)] = [1 - u(t)]H_1 + u(t)H_2$, where $u(t) \in [0, 1]$ is the ramping function satisfying $u(0) = 0$ and $u(T) = 1$. The parameter T is the total time the system is adiabatically evolved. The exact time-ordered time-evolution operator $U(T) = \mathcal{T} e^{-i\int_0^T H[u(t)]dt}$ is implemented using Trotter methods in two stages. First, the integral over t is split into M discrete steps of size $\delta t = T/M$, such that

$$U(T) \approx \prod_{k=0}^{M-1} U_1(k, \delta t) U_2(k, \delta t), \tag{44}$$

where $U_1(k, \delta t) = e^{-iH_1 \delta t_1}$ and $U_2(k, \delta t) = e^{-iH_2 \delta t_2}$. The values of δt_1 and δt_2 are given by

$$\delta t_1 = \int_{k\delta t}^{(k+1)\delta t} dt[1 - u(t)], \tag{45}$$

$$\delta t_2 = \int_{k\delta t}^{(k+1)\delta t} dt\, u(t). \tag{46}$$

Second, for $i = 1, 2$, each operator $U_i(k, \delta t)$ is approximated using a single step of a first-order Trotter formula, written explicitly as $U_i(k, \delta t) \approx e^{-i\delta t_i H_E} e^{-i\delta t_i H_B}$.

For our study, we choose the strong-coupling Hamiltonian with $g_1 = 10$ as H_1. The ground state is prepared by applying a Hadamard gate on each qubit. We choose the simple linear ramp with $u(t) = t/T$. Furthermore, we set $T = 1$ and perform the adiabatic evolution using $M = 2$ steps. Note that in order for the adiabatic theorem to be satisfied, the parameter T needs to scale as one over the square of the smallest energy gap along the adiabatic trajectory of the unscaled and unshifted Hamiltonian, i.e., $T = O((aE_1 - aE_0)^{-2})$, see, e.g., Ref. [124].

The top plot in Fig. 9 shows the value of γ as a function of N_p for multiple values of g. We observe that the dependence of γ on N_p is highly sensitive to the value of g: while for $g = 1.2$, the value of γ stays relatively constant, for $g = 0.2$, it decreases exponentially with N_p. The middle plot in Fig. 9 shows the value of γ as a function of n_q for multiple values of g. Similarly to the dependence N_p, γ depends weakly on n_q for $g = 1.2$, and decreases for $g = 0.2, 0.7$. Lastly, the bottom plot in Fig. 9 shows γ as a function of g for $N_p = 3$ and $n_q = 3$. For $g \gtrsim 2$ we find γ to be close to 1. This is expected as the large coupling ground state with $g_1 = 10$ and the ground state for $g \gtrsim 2$ have reasonable overlap. As one decreases g further, we observe a steep decrease in γ near $g = 1$. The value of γ approaches a constant in the small-g limit. This behavior is consistent with the decrease in overlap between the ground state of the strong-coupling Hamiltonian H_1 and the ground state of the target Hamiltonian H_2.

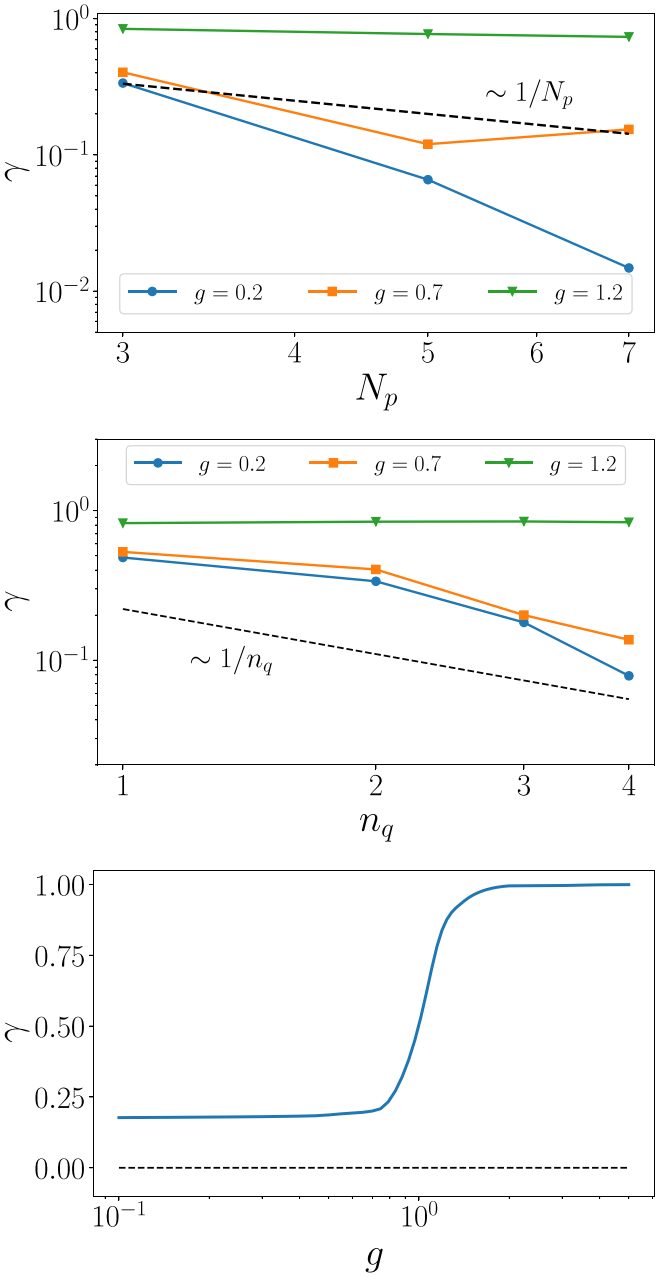

FIG. 9. This figure shows how γ scales with the parameters of our system when the initial guess is prepared using the adiabatic state preparation procedure described in the main text. (top) γ as a function of N_p for $n_q = 2$. The different colored lines indicate different values of g. The small-volume results indicate that γ is not exponentially decreasing as a function of N_p. The black dashed line shows a curve scaling as $\sim 1/N_p$. (middle) γ as a function of n_q for $N_p = 3$. The different colored lines indicate different values of g. The black dashed line shows a curve scaling as $\approx 1/n_q$. (bottom) γ as a function of the gauge coupling g for $N_p = 3$ and $n_q = 3$. The value is near one for large g and approaches a constant value for small g.

For the purpose of quoting how γ scales with the parameters of our system, we make the conservative choice $\gamma = O(\exp(-N_p)\exp(-n_q))$; this scaling is consistent with the worst-case observed scaling. In light of our simple procedure for preparing $|\psi_{\text{init}}\rangle$, this inefficient scaling is perhaps

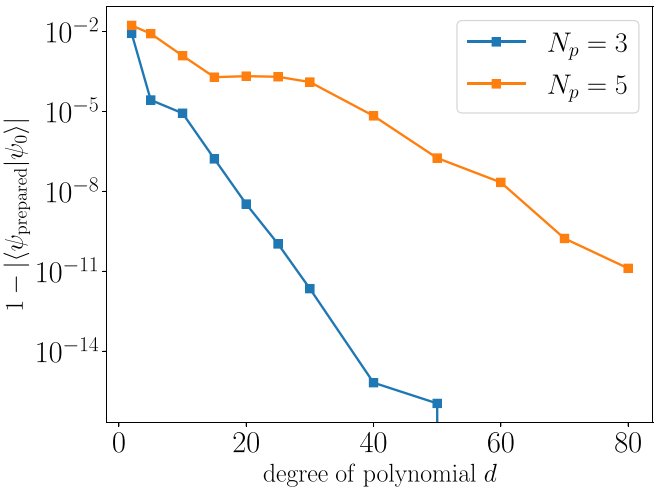

FIG. 10. Error of the ground state prepared using QETU as a function of the degree of the Chebyshev approximation d, where the time evolution operator was implemented *exactly*. Different colored points correspond to different number of plaquettes N_p. All results used $n_q = 1$ and $g = 1.4$. The error decreases exponentially with d, with the rate being slower for $N_p = 5$ due to a smaller gap Δ.

not unexpected. It is interesting, however, that for $g = 1.2$ we observe γ to be essentially independent of N_p and n_q for the values studied. Due to the simplicity of our adiabatic procedure combined with the small volumes that are accessible to classical simulation, predicting the scaling for realistic lattice sizes will require further dedicated studies.

We now study the fidelity of the prepared state as a function of the number of calls to the *exact* time evolution operator $e^{-i\tau H}$. The spectral norm and energy gap were determined exactly by diagonalizing the Hamiltonian. The value of $\Delta = (E_1 - E_0)/1.5$ was used, with $\tau = 1$. We denote the state prepared using QETU by $|\psi_{\text{prepared}}\rangle$, and the exact ground state by $|\psi_0\rangle$. The error is defined to be $1 - |\langle\psi_{\text{prepared}}|\psi_0\rangle|$. Figure 10 shows the error as a function of the degree of the polynomial d used to approximate the shifted sign function for $N_p = 3, 5$, $n_q = 1$, and $g = 1.4$. The error decreases exponentially as we increase d, with the rate of convergence being slower for $N_p = 5$ due to the smaller value of Δ.

We now move to the studies of *approximate* implementations of the time-evolution operator with the aid of Trotter methods. The first study compares the overall Trotter error between the standard controlled version of QETU and the control-free version. As explained in Sec. II A, the control-free version of QETU algorithm is designed to avoid using controlled calls to the time-evolution circuit. This method requires a Hamiltonian dependent procedure; the details for the U(1) case are given in Appendix C. By repeatedly calling $e^{-i\tau H}$, the standard version returns $F(\cos(\tau H/2))$, while the control-free version of QETU returns $F(\cos(\tau H))$. This implies that one can use instead $e^{-i\tau H/2}$ as a building block for control-free QETU. Because one only has to time evolve by half the total time $\tau/2$, for the same number of Trotter steps, one can use a step size half as large, leading to a smaller overall error relative to the standard controlled version of QETU. Figure 11 shows the error as a function of d when using both the standard and control-free versions of QETU.

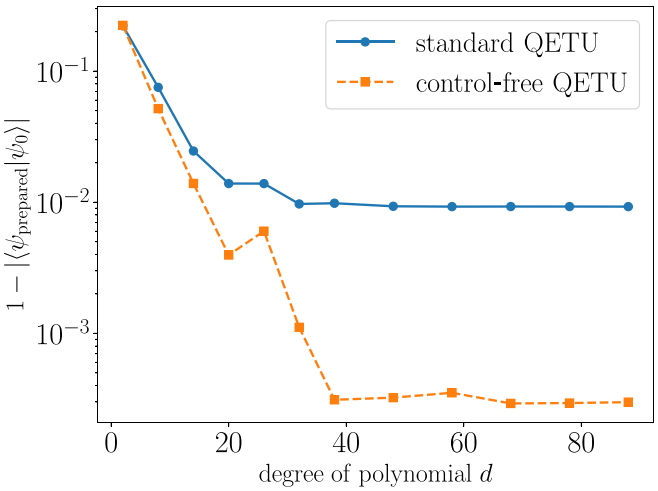

FIG. 11. Error as a function of d using both the standard and control-free versions of QETU. For both cases, a single Trotter step was used to approximate $e^{-i\tau H}$ with $\tau = 1.78$. The step sizes used were $\delta\tau = \tau$ and $\delta\tau = \tau/2$ for the standard and control-free versions, respectively. The Trotter error in the control-free version is smaller than the standard version by an order of magnitude.

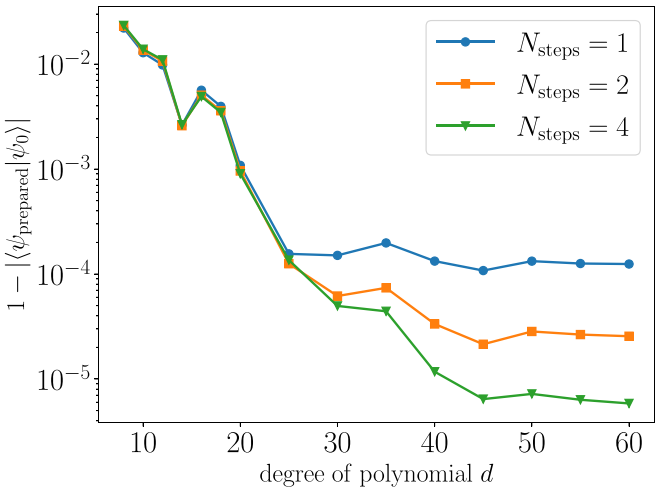

FIG. 12. Error of the state prepared using the control-free version of QETU as a function of d, where different colored and shaped data points correspond to different numbers of Trotter steps used to approximate the time-evolution circuit. The results are shown for $N_p = 3$, $n_q = 2$, $g = 0.6$, and $\tau = 1.5$. As the number of steps is increased, the error saturates at large d to smaller values due to the reduced Trotter errors.

The system parameters used are $N_p = 3$, $n_q = 2$, and $g = 0.6$. We approximated $e^{-i\tau H}$ with $\tau = 1.78$ using a single Trotter step of $\delta\tau = \tau$ and $\delta\tau = \tau/2$ for the standard and control-free versions of QETU, respectively. As d is increased, the error of both methods leveled out, with the error of the control-free version being an order of magnitude smaller due to the smaller time-step used. We see that, in addition to requiring less gates, the control-free version of QETU also results in a smaller Trotter error for the same number of calls to the time-evolution circuit. The remainder of the results in this section were obtained using the control-free version of QETU.

We now study how the maximum achievable precision depends on the Trotter step size $\delta\tau$. Figure 12 shows the error as a function of calls to the time evolution operator $e^{-i\tau H}$ for $\tau = 1.5$, approximated using a first-order Trotter formula with $N_{\text{steps}} = 1, 2, 4$. The parameters of the system studied are $N_p = 3$, $n_q = 2$, and $g = 0.6$. We see that the error decreases exponentially at first and then levels out for large number of calls. This occurs because the error is now dominated by the Trotter error, and improving the quality of the approximation of the projector is no longer beneficial. Even though we use the first order Trotter formula, the maximum precision achievable scales as $O(\delta\tau^2)$. This is due to the fact that the leading order error term of the form $\delta\tau[H_E, H_B]$ is zero for this system.

The final study we perform regarding the Trotter error is how, for some fixed total precision, the Trotter step size $\delta\tau$ must scale with the volume. Typically, as one increases the number of terms in a Hamiltonian, for fixed step size, the Trotter error increases due to the increased number of terms that do not commute. From this argument, as we increase N_p or n_q, we expect that $\delta\tau$ must be decreased accordingly if one wants to maintain a constant level of precision. However, scaling H_E and H_B by the parameter c_1 is equivalent to scaling the Trotter step size by c_1. Because c_1 generally increases with N_p and n_q, the effective Trotter step size decreases with N_p and n_q. If the decrease in error from the smaller effective step size is more significant than the increase in error from the extra noncommuting terms, then the Trotter error will *decrease* as we increase N_p or n_q. We observe this to be the case, and show an example of this counterintuitive behavior in Fig. 13. In this plot, we show the error as a function of d for three valus of $N_p = 3, 5, 7$ using $n_q = 1$ and $g = 1.4$. For each value of N_p, we use a single Trotter step with $\delta\tau = 1.5$. Notice that we can use a value of $\delta\tau > 1$ and still see convergence due to the fact that the effective Trotter step size is scaled by c_1.

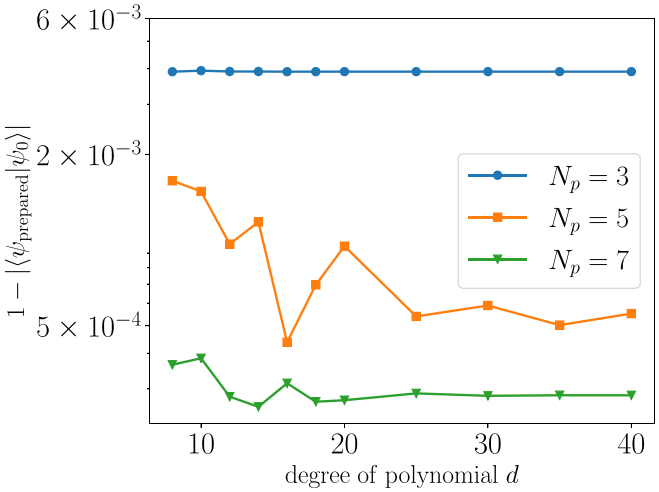

FIG. 13. Error of the state prepared with the control-free version of QETU as a function of d, where different colored and shaped data points correspond to different values of N_p. The results are shown for $n_q = 1$, $g = 1.4$, and $\tau = 1.5$. A single Trotter step was used to approximate the time-evolution circuit. As the number of plaquettes N_p increases, the Trotter error decreases. This behavior is due to the effective Trotter step size decreasing with N_p.

Looking at Fig. 13, we see that as we increase N_p, the maximum achievable precision increases. This behavior was also observed as n_q was increased. Because this behavior is expected to continue as one increases N_p and n_q towards realistic values, the Trotter error will eventually become negligible for any realistic precision requirements. Practically speaking, this implies that for a realistic calculation, one can approximate the time evolution operator using a single Trotter step of a first-order Trotter formula with $\delta\tau = \tau_{\max}$, independent of N_p or n_q. What at first seemed like a technical feature of QETU, turns out to offer a powerful protection against further N_p or n_q scaling. We conclude by stressing that, even though the effective step size is scaled by c_1, the un-scaled $\delta\tau$ must still satisfy $\delta\tau \leqslant \tau_{\max}$ in order to guarantee isolation of the ground state when applying the filter operator.

We conclude with the classical computational cost of calculating the angles $\{\varphi_j\}$ needed in the QETU circuit. We found numerically that the cost scales quadratically with the number of angles. Because the number of angles is proportional to Δ^{-1}, the associated classical cost scales as $(\Delta^{-1})^2 = O(N_p^2 2^{4n_q})$.

B. Wave-packet construction

In this section, we use QETU to prepare a Gaussian state, defined in Eq. (22), with $x_0 = 0$ and $\sigma_x = 0$. We first show an example of how naïvely applying QETU according to Eq. (25) leads to the error decreasing only polynomially with the number of Chebyshev polynomials. Next, we show how the modifications described in Sec. II C achieve an exponential scaling in the error for any desired value of the width, including a method that avoids the costly implementation of LCU to add the even and odd components of $\mathsf{F}(x)$. From there, we compare the gate count cost of our method to that of exact state preparation methods and find that our method requires less gates than exact state preparation methods for states represented by >2–5 qubits.

As explained in Sec. II C, due to the presence of the $\arccos(x)$ term, we expect the error of the approximation to decrease only polynomially with the number of Chebyshev polynomials. An example of the polynomial convergence is shown in Fig. 14, using parameters $n_q = 4$, $\sigma_x/x_{\max} = 0.4$, $\tau = 1$, and $\eta = 0$. Looking at the data labeled Method I, we observe that the error converges quadratically as the number of Chebyshev polynomials is increased. We now discuss modifications that can improve the scaling.

The first modification we study is to determine the parameter η and the Chebyshev coefficients according to the optimization problem in Eq. (26), except with τ fixed to $\tau = 1$. Looking at the data labeled Method II in Fig. 14, we see that while varying η improves the scaling to a point, the error eventually starts to decrease polynomially.

To improve upon the previous method, we now allow the parameter τ to be chosen, along with η and the Chebyshev coefficients, according to the optimization problem in Eq. (26). During our numerical tests, we found that standard minimization techniques were highly sensitive to the initial values of η and τ.

The next modification we study is now τ to determine the parameter η and the Chebyshev coefficients according to the optimization problem in Eq. (26), except with τ fixed to $\tau = 1$. Looking at the data labeled Method II in Fig. 14, we see that, while varying η improves the scaling to a point, the error eventually starts to decrease polynomially. To avoid this problem, we first determine good starting values for η and τ using a brute force approach and then use these values as inputs into a standard minimization procedure. While brute force approaches are technically inefficient, for our two-dimensional parameter space and cheap cost function, we find this method to be an appropriate choice. Note that negative values of η are allowed and often result in the smallest error. Our numerical studies routinely found that using a value of $\tau = 4$ and varying η leads to the smallest errors. One possible explanation for this is that, for $\tau = 4$, the parameter η can be varied to change the shape of the function $\mathsf{F}(\tilde{x}_i)$ to be approximately linear and quadratic, for the odd and even components, respectively; functions resembling linear and quadratic dependence are approximated well using only a few Chebyshev polynomials. Figure 14 shows the error using this procedure as a function of the number of Chebyshev polynomials, labeled as Method

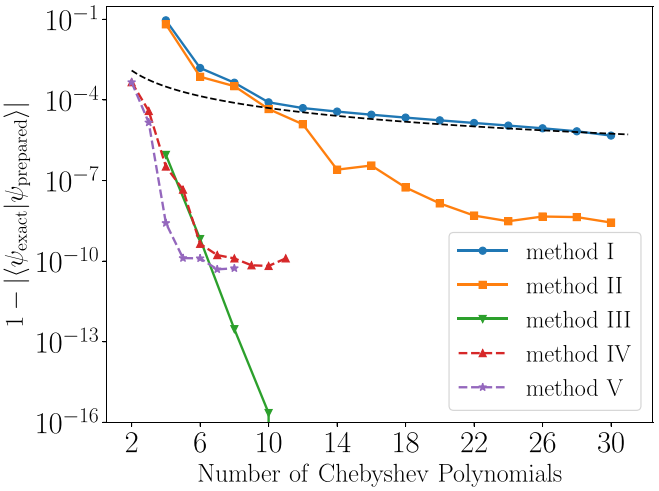

FIG. 14. Error of the prepared Gaussian state using QETU as a function of the number of Chebyshev polynomials included in both the even and odd components. Different colored points correspond to different methods of construction. Method I determined the Chebyshev expansion by solving the optimization problem in Eq. (25), using $\eta = 0$ and $\tau = 1$. Method II shows results solving the optimization problem in Eq. (26) for τ fixed to $\tau = 1$. Method III used values for η, τ, and the Chebyshev coefficients by solving the optimization problem in Eq. (26). Method IV first set $\tau = 2$ to make $\mathsf{F}(x)$ purely even and used values for η and the Chebyshev coefficients by solving the optimization problem in Eq. (26). Method V also set $\tau = 2$ and used values for η and the Chebyshev coefficients by solving the optimization problem in Eq. (27) where one only samples the function at points $\tilde{x}_j$. The black dashed line shows the curve $\approx 1/n^2$. The results are shown for $n_q = 4$ and $\sigma_x/x_{\max} = 0.4$. The error using Method I decreases quadratically with the number of Chebyshev polynomials. Using Method II, the error appears to at first decrease exponentially, then decreases only polynomially. The error using Method III decreases exponentially, and reaches floating point precision with only five Chebyshev polynomials. Using methods IV and V, the error first decreases exponentially, then levels out at $\approx 10^{-10}$, with method V generally outperforming method IV.

III. The error is five orders of magnitude smaller than using the previous methods, and decreases exponentially.

While we have demonstrated that, for fixed σ_x and n_q, QETU can be used to implement Gaussian states using $O(\log(1/\epsilon))$ controlled calls to $e^{-i\tau\hat{x}_{\rm sh}}$, the cost of performing LCU to add the even and odd pieces introduces a large overall coefficient. If one could modify the procedure to avoid using LCU, the gate count would be reduced by a factor of 10 or more. As discussed in Sec. II C, because of our specific digitization of the $\hat{x}$ operator, choosing a value of $\tau = 2$ results in the function $\mathsf{F}(x)$ being purely even. After setting $\tau = 2$, the parameter η and Chebyshev coefficients are determined by solving the optimization problem in Eq. (26) with τ set to $\tau = 2$. Figure 14 shows the error as a function of the degree of the Chebyshev polynomial used. The error decreases exponentially at first, levels off, then reaches zero as the number of parameters equals the number of points. Importantly, the error levels off at a value of $\approx 10^{-10}$, which will be completely negligible for a realistic simulation. While the error for the same number of Chebyshev polynomials is generally larger than if one implemented the even and odd pieces separately and then added them using LCU, the cost of performing LCU leads to over an order of magnitude larger gate count. For this reason, converting $\mathsf{F}(x)$ into an even function to avoid the cost of LCU results in the best precision for a given gate cost.

The final improvement studied exploits the fact that, because the eigenvalues of the $\hat{x}_{\rm sh}$ operator are known exactly, we must only reproduce the function $\mathsf{F}(x)$ at those points. After fixing $\tau = 2$, the parameter η and the Chebyshev coefficients are determined according to the optimization problem in Eq. (27). Looking at Fig. 14, we find the error using this method results in smaller errors than the previous method of sampling all x values. In a similar way, the error decreases exponentially, approaches a constant of $\approx 10^{-10}$, then becomes zero when the number of Chebyshev polynomials equals the number of points the function $\mathsf{F}(x)$ is sampled. This method results in the smallest error per gate count, and is used to produce the rest of the results in this section.

We now study how the precision varies with the number of qubits n_q. Figure 15 shows the error of the prepared state as a function of n_q for different degree polynomial approximations, using $\sigma_x/x_{\rm max} = 0.2$. Our results indicate that the precision is independent of n_q, except in the cases where the small number of sample points results in an exactly prepared Gaussian state. From this we learn that the number of Chebyshev polynomials required to achieve some desired precision is independent of n_q. For a fixed with σ_x, the gate cost required to implement a Gaussian is therefore $O(n_q \log(1/\epsilon))$. This result will be important when comparing the gate cost with exact state preparation methods.

The final component of our precision study is to understand how the precision scales as the width $\sigma_x/x_{\rm max}$ of the wave packet is varied. Figure 16 shows the error of the prepared state as a function of the number of Chebyshev polynomials used for different values of $\sigma_x/x_{\rm max}$, using $n_q = 5$ qubits. We find that the error decreases exponentially for all values of the width, with more sharply peaked Gaussian states requiring more Chebyshev polynomials to achieve the same precision as more broadly peaked states. This intuitive result is similar

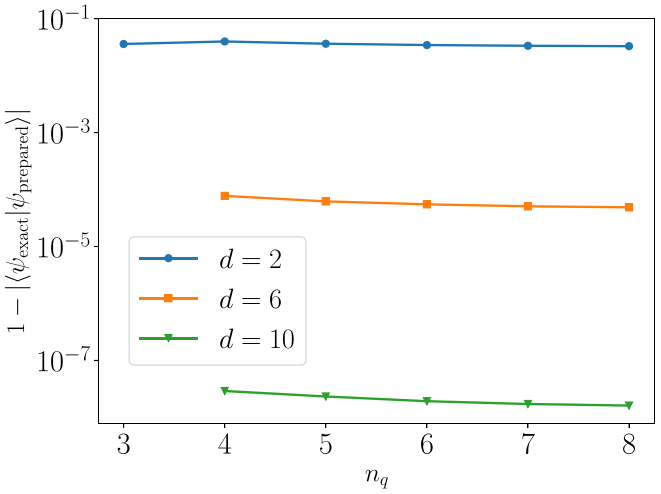

FIG. 15. Error of the prepared Gaussian state using QETU as a function of the number of qubits n_q used to represent the state. The state was prepared by setting $\tau = 2$ and solving the optimization problem in Eq. (27). Different colored points show different degree Chebyshev expansions. The results for $n_q = 3$ using $d = 6, 10$ are not shown because the error was zero. Except for small values of n_q where the state is prepared exactly, the error is independent of n_q. Results are shown for $\sigma_x/x_{\rm max} = 0.2$.

to the fact that the cost of approximating the shifted error function increases as the energy gap Δ is decreased.

We conclude this section with a comparison of the gate cost when preparing a Gaussian state using QETU to using an exact state preparation procedure. For this comparison, we compile the resulting quantum circuits using a universal gate-set consisting of CNOT, R_z, and R_x gates using QISKIT

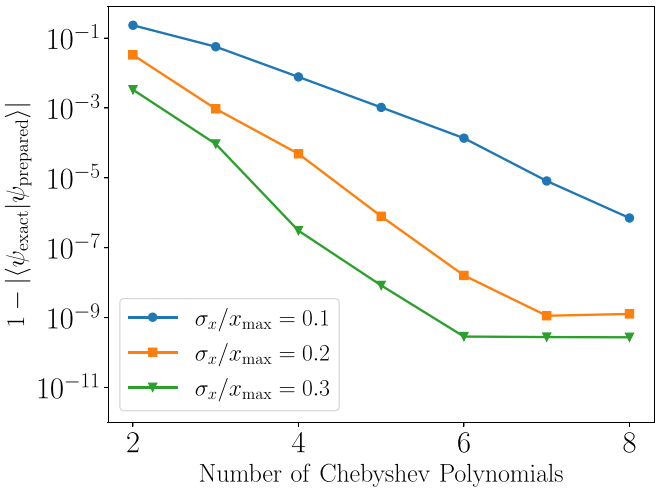

FIG. 16. Error of the prepared Gaussian state using QETU as a function of the number of Chebyshev polynomials used to construct the Gaussian filter operator. The state was prepared by setting $\tau = 2$ and solving the optimization problem in Eq. (27). Different colored points show different values of the wave packet width $\sigma_x/x_{\rm max}$. In all cases, the error decreases exponentially. More sharply peaked wave packets require more Chebyshev polynomials to achieve the same level of precision. Results are shown for $n_q = 5$.

[125]. Before discussing precise gate counts, we argue the expected scaling for both methods. Starting with QETU, one controlled call to the $e^{-i\tau\hat{x}_{\rm sh}}$ operator requires $O(n_q)$ CNOT and R_z gates. The number of R_x gates is simply equal to the number of Chebyshev polynomials included in the approximation. Exact state preparation methods on the other hand require $O(2^{n_q})$ CNOT, R_z and R_x gates. Even though QETU has a better asymptotic scaling, because of the relatively large coefficient in the overall QETU cost, we expect that for small values of n_q, exact state preparation will be less costly than using QETU. The question we now answer is at what value of n_q does the cost of QETU become cheaper than exact state preparation methods.

For the gate count comparison, we compare the number of gates and neglect the scaling from the γ factor for the QETU cost. We do this because algorithms with shorter depths for a single run, rather than total gates required, when using NISQ and early fault tolerant devices are preferred. For the gate count comparison, because the cost of R_z and R_x gates are similar, we chose to compare the number of total rotation gates, as well as the number of CNOT gates. It is important to note that the gate count using QETU for a fixed number of calls to the $e^{-i\tau\hat{x}_{\rm sh}}$ is independent of value of $\sigma_x/x_{\rm max}$. What is important, however, is the precision one can achieve for that particular value of $\sigma_x/x_{\rm max}$. The cost for exact state preparation methods is independent of $\sigma_x/x_{\rm max}$.

Figure 17 shows the gate counts required to prepare the Gaussian state, using both QETU and exact state preparation methods, as a function of n_q. The top and bottom plots show the number of CNOT and rotation gates, respectively. For QETU, the gate count is shown for two, three, and four calls to the $e^{-i\tau\hat{x}_{\rm sh}}$ operator. As expected, the scaling for the number of CNOT and rotation gates is linear in n_q when using QETU, and exponential in n_q for exact preparation methods. Comparing first the CNOT count, looking at the top plot in Fig. 17, we notice that for $n_q = 5$, the CNOT count when using QETU starts to be cheaper than exact state preparation methods.

We now compare the number of rotation gates for both methods. Looking at the bottom plot in Fig. 17, we see that already for $n_q = 2, 3$, QETU requires less rotation gates than exact state preparation methods. As already shown in Fig. 16, for certain values of $\sigma_x/x_{\rm max}$, using only two to three Cheybshev polynomials can achieve a subpercent precision on the prepared state. Additionally, we showed in Fig. 15 that for fixed degree Chebyshev approximations, the precision is independent of the value of n_q.

Taken together, all of these results imply that, depending on the desired error tolerance, value of $\sigma_x/x_{\rm max}$, and whether noisy or error-corrected qubits are used, it makes sense to consider using QETU for values of n_q as small as $n_q \gtrsim 2$–5. Note that it is in principle possible to reduce the cost of exact state preparation methods by dropping gates with small rotation angles below some threshold. The gate count savings and the error introduced will require a dedicated study, which will be interesting to perform in future work.

We conclude by discussing how the cost of our method compares to the cost of the Kitaev-Webb (KW) algorithm for preparing Gaussian states [126]. While the KW algorithm can implement Gaussian states with a cost polynomial in the number of qubits, the algorithm comes with a large overall prefactor due to the need to perform arithmetic on the quantum computer. A detailed study comparing the gate cost of using the KW algorithm to the gate cost of exact state preparation methods was performed in Ref. [127]. It was found that, due to a large prefactor, state preparation using the KW algorithm for a single Gaussian wave packet was more expensive than exact state preparation methods up to $n_q \approx 14$ qubits. This, combined with the fact that the cost of KW scales more quickly than linear in n_q, implies that using QETU to prepare a one-dimensional Gaussian wave function will also be cheaper than using the KW algorithm for any value of n_q.

In principle, it is also possible to use QETU to also prepare multidimensional Gaussian states. Because of the probabilistic nature of using QETU, one will have to ensure that the γ factor does not decrease exponentially with the dimension of

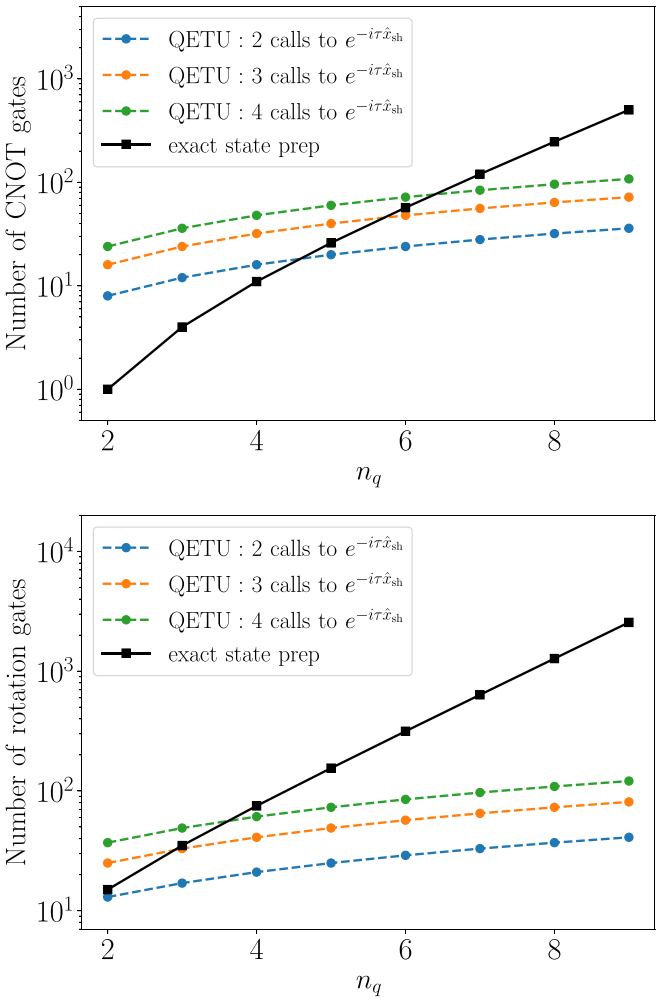

FIG. 17. (top) Number of CNOT gates required to prepare a Gaussian state using both QETU and exact state preparation methods as a function of the number of qubits used to represent the state n_q. The factor of γ is not included in the QETU count. The black squares with a solid line show the count for exact state preparation method. The different colored circles with dashed lines show the cost of QETU for different number of calls to the $e^{-i\tau\hat{x}_{\rm sh}}$ operator. Exact state preparation methods become more expensive around $n_q = 5$. (bottom) Same as top but for the number of rotation gates required. Exact state preparation becomes more expensive around $n_q = 2$.

the Gaussian state being prepared. This would be nontrivial to achieve, given the nonunitary nature of the Gaussian filter transformation (22). It will be interesting to explore this application of QETU in a future work.

The final study we perform is on how γ depends on n_q, the number of Chebyshev polynomials, and the width $\sigma_x/x_{\max}$. Recall that γ is defined to be the magnitude of the final state prepared using QETU. Because the initial state guess and implemented operator $F(\cos(\tau\hat{x}_{\rm sh}/2))$ are known exactly, we can directly evaluate γ in the limit where we reproduce $\mathsf{F}(x)$ exactly. Doing so gives

$$\gamma = |\langle\psi_{\rm init}|\mathsf{F}(\cos(\tau\hat{x}_{\rm sh}/2))^2|\psi_{\rm init}\rangle| = \frac{c^2}{2^{n_q}}\sum_{j=0}^{2^{n_q}-1} e^{-x_j^2/\sigma_x^2}, \quad (47)$$

where going to the second line we used $|\psi_{\rm init}\rangle = 2^{-n_q/2}\sum_{j=0}^{2^{n_q}-1}|x_j\rangle$ and $F(\cos(\tau\hat{x}_{\rm sh}/2)) = c\,e^{-\hat{x}^2/(2\sigma_x^2)}$. If we replace F with the approximation F with error ϵ, the above expression for γ is expected to be correct up to $O(\epsilon)$ corrections. In the limit of $n_q \to \infty$, the value of γ is simply the area under the Gaussian curve. We therefore expect γ to approach a constant as n_q is increased. The value of this constant will depend on $\sigma_x/x_{\max}$, and is given by $c^2\int_{-x_{\max}}^{x_{\max}} dx\, e^{-x^2/\sigma_x^2}$. From this we expect that γ decreases as the ratio $\sigma_x/x_{\max}$ decreases.

Using this expression for γ, we can identify a scenario where γ can be prohibitively small. Suppose that one uses a small value of n_q to implement a sharply peaked Gaussian state. It is possible that one only samples at values of x where the function $e^{-x^2/(2\sigma_x^2)}$ is exponentially suppressed, resulting in an exponentially small value of γ. However, because sampling a sharply peaked function a small number of times will introduce large digitization errors, this situation is not likely to occur in practice; in a realistic scenario, one will sample the wave packet using a large enough value of n_q to avoid a small value of γ.

Figure 18 shows the dependence of γ^{-1} on n_q and $\sigma_x/x_{\max}$. All results were calculated using a degree 18 even Chebyshev approximation. For the data shown, the largest error from the finite Chebyshev approximation of the state produced using QETU was $\approx 10^{-9}$. The top plot in Fig. 18 shows γ^{-1} as a function of n_q for different values of $\sigma_x/x_{\max}$. Using $n_q = 2$ for $\sigma_x/x_{\max} = 0.1$ results in $\gamma^{-1} \approx 10^5$ because the wave packet is only sampled at values where the value is exponentially suppressed. This problem can be avoided by increasing n_q and is less severe for larger values of $\sigma_x/x_{\max}$. We notice that by $n_q = 5$ the value of γ^{-1} levels out. The bottom plot in Fig. 18 shows γ^{-1} as a function of $\sigma_x/x_{\max}$ for $n_q = 5$. We see that γ is directly proportional to $\sigma_x/x_{\max}$, with more sharply peaked states having smaller values of γ. For a sharply peaked value of $\sigma_x/x_{\max} = 0.1$, we find $\gamma^{-1} \approx 13$.

The main value of our scrupulous numerical investigations in this section is the demonstrated ability to prepare wave packets with exponential precision for arbitrary values of the distribution parameters. This required several optimization procedures and could not have been trivially predicted based on the construction proposed in Sec. II C. Our result proves an important point that, despite the presence of $\arccos(x)$ in the QETU approximation, it can be used for representing functions of Hermitian operators, with exponential precision.

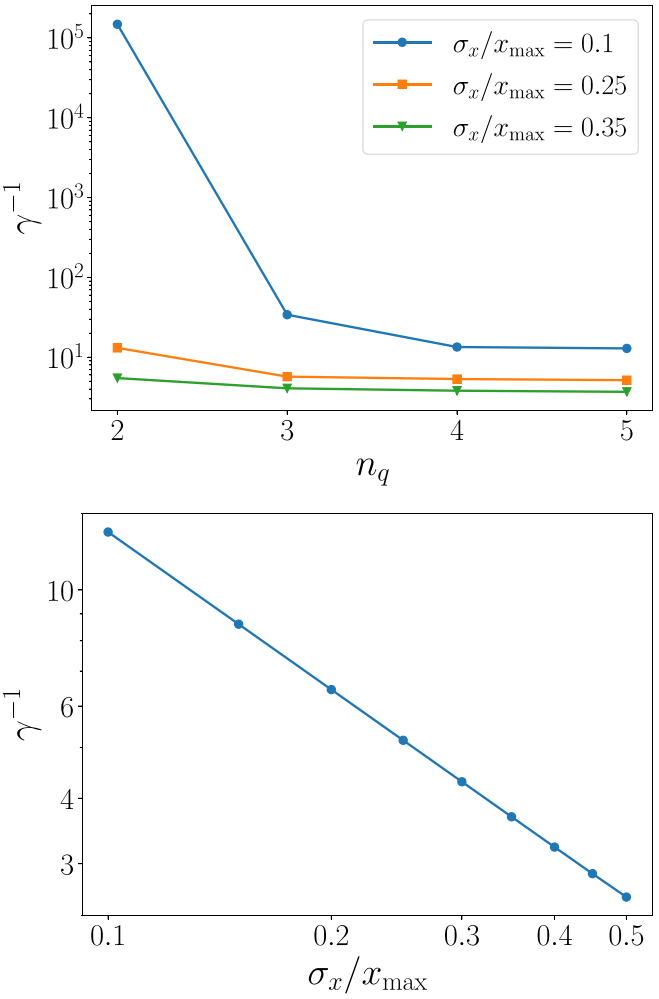

FIG. 18. Dependence of γ on various parameters when constructing the Gaussian state by setting $\tau = 2$ and solving the optimization problem in Eq. (27). (top) γ^{-1} as a function of n_q. Different colored points indicate different values of $\sigma_x/x_{\max}$. The value of γ^{-1} is large for $n_q = 2$ and $\sigma_x/x_{\max} = 0.1$ due to sampling the sharply peaked Gaussian only at points where the value is exponentially suppressed. (bottom) γ^{-1} as a function of $\sigma_x/x_{\max}$ for $n_q = 5$. The value of γ is proportional to $\sigma_x/x_{\max}$.

(In Ref. [98] this issue was avoided by constructing a polynomial approximation to the error function which, in turn, was approximating the cosine-transformed original step function.)

V. GROUND-STATE PREPARATION COST FOR A GENERAL GAUGE THEORY

In this section, we discuss how the cost of ground-state preparation using QETU is expected to scale with the parameters of a general lattice pure gauge theory in d spatial dimensions. Understanding the cost in terms of a total gate count requires detailed information regarding the specific formulation of the lattice gauge theory in question, as well as the algorithm used to implement $e^{-i\tau H}$. To keep the discussion in this section as general as possible, we consider the cost in terms of the number of calls to the time evolution circuit. For ease of notation, throughout this section, we denote energies

of the unscaled Hamiltonian of the general lattice gauge theory by E_i.

We consider a pure gauge theory on a hypercubic spatial lattice with N_s sites in each dimension, and a lattice spacing a between neighboring sites. The d-dimensional volume is given by $V = (aN_s)^d$. To truncate the infinite-dimensional Hilbert space of this bosonic theory, we represent each lattice site using n_q qubits. We consider a general gauge theory sharing the same qualitative features as QCD. First, the theory is assumed to have a mass gap between the vacuum and the first-excited state. Second, we assume that the theory is asymptotically free, i.e., that as the lattice spacing a goes to zero, the bare coupling $g(a)$ also goes to zero. The true, continuous, infinite volume theory is recovered while simultaneously taking $(aN_s)^d \to \infty$, $n_q \to \infty$, and $a \to 0$ (while appropriately adjusting g according to the chosen renormalization scheme). We work in the units of $\hbar = c = 1$.

One important consideration is that, in a lattice gauge theory, one does not directly choose the lattice spacing a. One can only choose the value of the bare gauge coupling g. The two parameters are related via the renormalization group; one should view the gauge coupling as a function of the lattice spacing, i.e., $g = g(a)$. What this means in practice is that one can only calculate dimensionless quantities in a lattice gauge theory. For example, if one considers an energy E, one can only calculate the dimensionless value aE. Once the lattice spacing a has been determined using some renormalization scheme, the dimensionful energy E can be extracted. However, this consideration can be avoided for our discussion using the fact that the cost of QETU depends only on ratios of energies. The explicit lattice spacing dependence cancels, and we can consider directly energy differences $E_1 - E_0$ and $E_{\max} - E_0$ in physical units.

While the explicit dependence on the lattice spacing cancels, the energies still have implicit dependence on a. The energies also have an implicit dependence on the physical volume V and the number of qubits per lattice site n_q. To better understand this dependence, it is useful to recall that the lattice acts as both a high- and low-energy regulator of our quantum field theory. Using a finite lattice spacing introduces an energy cutoff $\approx 1/a$, while using a finite volume provides a low-energy cutoff of $\approx 1/V$. In addition, using a finite value for n_q provides another high-energy cutoff, denoted as $\Lambda_{n_q} \sim 2^{n_q}$. Assuming one has properly renormalized the theory, as these regulators are removed, i.e., sending $a \to 0$, $V \to \infty$, and $\Lambda_{n_q} \to \infty$, one should obtain the correct physical energies. We proceed with the discussion assuming access to (early) fault-tolerant quantum computers, where finite-volume errors ε_V, finite lattice spacing errors ε_a, and finite n_q errors ε_{n_q} are controlled.

We now discuss the scaling of $E_1 - E_0$ and $E_{\max} - E_0$ with N_s, a and n_q. The energy difference $E_1 - E_0$ in a continuous theory with a mass gap has a finite value. Therefore, as one removes the regulators, the gap $E_1 - E_0$ is expected to approach a constant. However, $E_{\max} - E_0$ diverges in this limit. We now consider how quickly this term grows as we remove each of the regulators. Starting with the volume, it is helpful to think from the perspective of fixed a and n_q. Increasing the volume is achieved by increasing the number of sites N_s. In the simplest possible case of a free theory, each additional site increases the maximum possible energy by the maximum energy of a single site. For a local gauge theory, the maximum energy is expected to increase in a similar way when using more lattice sites.[3] The maximum energy of our interacting theory is therefore expected to grow linearly with the total number of sites N_s^d. For fixed a, this is equivalent to growing linearly with the volume V. Turning now to the lattice spacing, recall that a finite value of a imposes a high-energy cutoff given by $\approx 1/a$. We therefore expect the maximum energy to scale in the same way as $1/a$. Lastly, it is known that the maximum energy of a bosonic theory generally increases exponentially with n_q [118,128].

In terms of these energy gaps, we know $\Delta \sim (E_1 - E_0)/(E_{\max} - E_0)$. Combining our scaling arguments, we find

$$\Delta^{-1} = O\left(\frac{E_{\max} - E_0}{E_1 - E_0}\right) = O\left(a^{-1} N_s^d 2^{n_q \alpha}\right), \tag{48}$$

where the parameter $\alpha \in \mathbb{R}_{>0}$ is both theory and formulation dependent. Numerical studies of this scaling were performed for a compact U(1) lattice gauge theory in two spatial dimensions in Sec. IV A.

We now turn our attention to how the overlap γ between the initial guess and exact ground state scales. For this discussion, we assume that the vacuum is translationally invariant, which is true for theories in the standard model. There are many ways to prepare this initial guess, including adiabatic state preparation [129–131], variational methods [66,71,72,132–134], and direct preparation [135–145]. We now argue that preparing such an initial guess wave function using direct state preparation will lead to exponential scaling in the cost of preparing the ground state with QETU.

Consider a situation where one prepares regions of the lattice using exact state preparation, where for simplicity we assume that each region contains N_r sites in each dimension. The number of these regions is given by $(N_s/N_r)^d$. The gate cost of preparing the state of a single region, each with overlap γ_i, scales as $2^{(N_r)^d}$, with the total overlap given by $\gamma \sim (\gamma_i)^{(N_s/N_r)^d}$. Any choice of N_r that breaks the exponential scaling in one of these costs necessarily introduces an exponential scaling in the other. Even though each overlap γ_i can in principle be improved using amplitude amplification from γ_i to $\sqrt{\gamma_i}$ [102,146], the overall γ parameter still decreases exponentially as one increases the number of lattice sites. Creating a trial state using direct state preparation methods is therefore inefficient. In light of this argument, we studied how γ scales with the volume if one uses a simple adiabatic state preparation procedure to prepare the initial guess wave function $|\psi_{\text{init}}\rangle$ for the U(1) formulation we considered. Our results indicate the scaling of γ using adiabatic state preparation is highly dependent on the gauge coupling g, ranging from nearly constant in the volume to decreasing exponentially with the volume. We conclude by noting that it is possible that more sophisticated adiabatic preparation procedures, or the use of

[3]Note that, because gauge fixing generally leads to some nonlocality in the resulting Hamiltonian, this argument does not necessarily apply to such formulations. If one does use a nonlocal Hamiltonian, understanding how $E_{\max} - E_0$ scales with the number of sites will likely require a dedicated study.

variational state preparation, could result in a better scaling of γ^{-1} with the volume. It will be interesting to explore this in future work.

Another component of the cost is how the number of Trotter steps used to approximate the time evolution circuit scales with the system size. As one increases N_s or n_q, the number of noncommuting terms in the Hamiltonian increases, and we therefore expect the number of steps required to maintain a constant precision to increase accordingly. However, as discussed in Sec. IV A, scaling the Hamiltonian such that the spectrum is in the range $[0, \pi]$ can be equivalently viewed as scaling the Trotter step size. Because the maximum energy generally grows with N_s and n_q, the effective Trotter step size will decrease with N_s and n_q. If this decrease in the Trotter step size results in a smaller error that the increase from the additional noncommuting terms, the overall Trotter error would actually *decrease* with N_s and n_q. As shown in Sec. IV A, this was indeed found to be the case for the U(1) gauge theory we considered. Due to the similarity of Hamiltonian structure of general lattice gauge theories, it is possible that this trend will also be present for other lattice gauge theories.

Lastly, we point out an observation that may help dampen the scaling of Δ^{-1} as one increases N_s and n_q. The main idea is that in general, an initial guess with good overlap with the ground state will generally have smaller overlap with excited states, with the overlap continuing to decrease for higher excited states. As a simple example, consider the quantum harmonic oscillator, with operators sampled using n_q qubits. One possible initial guess is a constant wave function for all x. Because the nth excited state has n nodes, the overlap of this initial guess and higher excited states will decrease with n, due to the fact that one sums more highly oscillatory functions. Let us denote the overlap of the highest energy state and the initial guess as $\gamma_{n_q,\min}$, with associated energy $E_{2^{n_q}-1}$. Now suppose that one increases the number of qubits to $m_q > n_q$. This increases the size of the Hilbert space, but the key point is that only states with higher energies are added. Because these higher-energy states also have more oscillations, the overlap the initial guess has with these additional states are all smaller than $\gamma_{n_q,\min}$. Depending on the error threshold, it is conceivable that, for a small enough value of $\gamma_{n_q,\min}$, one can simply ignore the states with energy larger than $E_{2^{n_q}-1}$. With regards to using QETU, this implies we only have to divide our energy gap by $E_{2^{n_q}-1} - E_0$ instead of $E_{2^{m_q}-1} - E_0$. Because this argument is true for any $m_q > n_q$, the amount one must scale down the physical energy gap eventually becomes independent of n_q.

It is possible that these arguments can be applied in a similar way to lattice gauge theories in order to argue that the scaling with N_s and n_q could be milder than previously argued. In the best case scenario where the dependence on N_s and n_q vanishes for some values of the parameters, Δ would only depend on physical energies E_0, E_1, and some E_{n^*}, where E_{n^*} is the highest-energy state that must be filtered out using QETU. In this way, Δ is no longer explicitly dependent on the lattice spacing, only through finite lattice spacing errors. If such a scenario is true, then Δ would become independent of N_s, n_q, and a. The cost would then be some overall large prefactor, dependent on the value of Δ where it becomes independent of N_s, n_q, and a, multiplied by the cost of implementing $e^{-i\tau H}$ using Trotter methods. Our preliminary numerical investigations showed that the validity of such a hypothesis is heavily dependent on both the initial guess state and the value of the coupling constant. We leave further investigations for future work.

Importantly, the above-described method of dampening the scaling with N_s or n_q is specific to the QETU approach, and would not be possible if one instead used the Hamiltonian input model, as in Ref. [105]. In the Hamiltonian input model, one performs repeated calls to a block encoding of H and uses the quantum eigenvalue transformation to implement the projector. Already at the stage of constructing the block encoding circuit for H, this method requires the Hamiltonian to be scaled so that $||H|| \leqslant 1$. Scaling down the Hamiltonian, and therefore the gap, by a factor of $E_{\max} - E_0$ is thus unavoidable in this scenario. On the contrary, in the QETU case, where one performs repeated calls to the unitary operator $e^{-i\tau H}$, there are no fundamental obstacles for constructing the time-evolution circuit for some H with a large spectral norm. The spectrum of H is shifted solely to avoid the problems associated with the periodic nature of the matrix function which QETU implements. Consequently, this technical difference between the Hamiltonian and time-evolution input models could result in the cost of QETU having better asymptotic scaling with N_s or n_q.

VI. DISCUSSION AND CONCLUSION

In this work, we performed an extensive study of the QETU algorithm and its applications to state preparation in simulations of quantum field theory. By modifying the original algorithm in Ref. [98], we were able to achieve significant cost savings when the time evolution circuit is implemented both exactly as well as approximately using Trotter methods.

We applied our improved procedure to prepare a ground state in a particular lattice formulation of U(1) gauge theory in two spatial dimensions. To avoiding the costly controlled calls to the time evolution circuit, we based our circuits on the control-free version of the QETU algorithm. The considered control-free implementation of QETU generalizes to any Hamiltonian of the form $H = H_x + H_p$, where the bases of the kinetic piece H_p and the potential piece H_x are related by a Fourier transformation; a form that is common to many types of lattice field theories and their different formulations [41–44,118,122,128].

We studied how the cost of the QETU-based state preparation algorithm scales with parameters of our physical system. In particular, we discussed how scaling down the spectrum of the physical Hamiltonian ensues the scaling of the energy gap Δ and placed upper bounds on this parameter for arbitrary system sizes. Next, we discovered that scaling down the Hamiltonian spectrum leads to unexpected positive consequences; namely, it makes the Trotter error decrease as the number of sites N_p or the number of qubits per site n_q are increased. This behavior is due to the fact that scaling down the Hamiltonian spectrum is equivalent to scaling down the Trotter step size. Additionally, using a simple adiabatic state preparation procedure, we found that the volume scaling of γ^{-1} is highly sensitive to the gauge coupling g considered, ranging from nearly constant in the volume to decreasing

exponentially with the volume. These studies lay the basis for the further studies on applications of QETU-based state preparation techniques for alternative formulations of gauge theories.

We followed our numerical results with a discussion of asymptotic costs of using QETU for state preparation in general lattice field theories with properties similar to QCD. We also argued that the scaling of Δ with the number of sites or qubits per site can be reduced by exploiting the fact that, for a good initial guess, highly excited states will have a negligible overlap, and do not need to be filtered out.

In this work, we also developed an alternative application of the QETU algorithm for the preparation of Gaussian states. The main idea is to use QETU to implement the Gaussian filter operator $e^{-\hat{x}^2/(2\sigma_x^2)}$. We argued that a naïve application of QETU to this problem results in the error decreasing only polynomially with the degree of the Chebyshev expansion. We showed, however, that one could instead achieve an exponential scaling, for any value of the width of the Gaussian state, by performing simple modifications to the QETU procedure. With these improvements, we presented a procedure that allows one to prepare Gaussian states while avoiding the costly step of adding the even and odd pieces using LCU. By performing a gate cost analysis, we showed that preparing Gaussian states with QETU using our improved procedure outperforms exact state preparation methods for as few as $n_q \gtrsim 2$–5 qubits, depending on whether one is using noisy or error-corrected qubits.

This work leads naturally to a number of additional interesting applications. While we used QETU to prepare the vacuum state of a pure gauge theory, it is in principle possible to extend QETU to prepare hadronic states in QCD, another important step towards simulating QCD. For concreteness, consider the task of preparing the quantum state of the pion. While the pion is not the ground state of the QCD Hamiltonian, it is, however, the lowest energy state with the quantum numbers of the pion. Additionally, the gap between the pion state and the next excited state with the desired quantum numbers is roughly twice the pion mass. If one prepares an initial state $|\psi_{\text{init}}\rangle$ with the quantum numbers of the pion, e.g., using the high-quality interpolating fields that have been developed for use in lattice QCD (see, e.g., Ref. [147] for a pedagogical introduction) then one can construct a filter operator using the QCD Hamiltonian and isolate the pion state. Because these interpolating fields are not unitary but Hermitian operators, a dedicated study to how best implement these operators will be required. It is possible, in principle, to use QETU to construct interpolating fields. An interesting question to ask is if one needs to apply the interpolating field to the interacting vacuum prepared using, e.g., QETU, or if one can simply apply the interpolating field to any state with the quantum numbers of the vacuum. If one could avoid preparing the interacting ground state while avoiding γ scaling exponentially poorly with the number of sites, this would result in a significant cost reduction.

Another interesting follow-up is exploring the use of QETU for preparing multidimensional Gaussian states, which are relevant for lattice gauge theories as well. The general form of such a state is $e^{-\sum_{ij} c_{ij} x_i x_j}$. A simple procedure for constructing this state using QETU is to set $\mathsf{f}(x) = e^{-x}$ and use $e^{i \sum_{ij} c_{ij} \hat{x}_i \hat{x}_j}$ as a building block in the QETU circuit. One important consideration using this method, or any variation, is to ensure that the γ parameter does not decrease exponentially as one increases the dimension of the Gaussian. If this can be avoided, it will be interesting to compare the cost of QETU to the Kitaev-Webb algorithm for preparation of such states.

Our studies provide the foundation for further investigations on the applicability of QETU in the context of highly efficient preparation of various functions of Hermitian operators. In our following work [148], we investigate whether the QETU algorithm can be utilized for efficiently implementing block encodings, which paves the way toward synergizing QETU with state-of-the-art simulation strategies [90,149].

ACKNOWLEDGMENTS

C.F.K. would like to thank S. Meinel for access to their computing cluster for some of the calculations. C.F.K. would also like to thank D. Grabowska for useful discussions regarding adiabatic state preparation for the U(1) formulation considered in the paper. We thank N. Modi and S. Hariprakash for access to their code and C. Bauer for useful discussions. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award No. DE-SC0020347. M.K. acknowledges support from the DOE grant PH-HEP24-QuantISED, B&R KA2401032/34. This work was supported by DOE Office of Advanced Scientific Computing Research (ASCR) through the ARQC program (N.G.). Support is also acknowledged from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator.

APPENDIX A: CALCULATING THE PHASES IN THE QETU CIRCUIT

In this Appendix, we review the procedure for calculating the phases $\{\varphi_j\}$ used in the QETU circuit. Note that multiple conventions for defining the phases $\{\varphi_j\}$ exist, and throughout this work we use the so-called *W-convention*. Here we mostly follow the pedagogical discussion given in Ref. [102]. To calculate the phases $\{\varphi_j\}$, one needs to combine a number of techniques, including quantum signal processing (QSP) [84], qubitization [85], and QETU [98]. QSP is the theory of the unitary representation of scalar polynomials $P(x)$. Put more concretely, QSP tells us that, for some set of symmetric phase factors $(\phi_0, \phi_1, \ldots, \phi_1, \phi_0) \in \mathbb{R}^{d+1}$, one can construct an even Chebyshev polynomial $g(x, \{\phi_j\})$ of degree d in the following way:

$$g(x, \{\phi_j\}) = \mathrm{Re}[\langle 0| e^{i\phi_0 Z} e^{i \arccos(x) X} e^{i\phi_1 Z} e^{i \arccos(x) X} \cdots e^{i\phi_1 Z} e^{i \arccos(x) X} e^{i\phi_0 Z} |0\rangle], \tag{A1}$$

where X and Z are the usual Pauli matrices. Qubitization can then be used to lift the above SU(2) representation to matrices of arbitrary dimensions. Finally, the quantum eigenvalue transformation for unitary matrices (QETU) tells us that, if

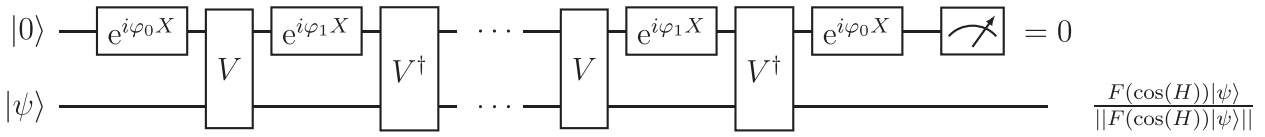

FIG. 19. Control-free QETU circuit diagram. The top qubit is the control qubit, and the bottom register is the state we apply a matrix function to. Here V is the operator given in Eq. (C1) that implements simultaneous controlled forward and backward time evolution. By applying alternating calls to V and $V^\dagger$, upon measuring the ancillary qubit to be in the zero state, one prepares the normalized quantum state $F(\cos(H))|\psi\rangle/||F(\cos(H))|\psi\rangle||$, where $F(\cos(H))$ can be a general even polynomial. For symmetric phase factors $\{\varphi_j\} = (\varphi_0, \varphi_1, \ldots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$, then $F(\cos(H))$ is an even polynomial of degree d. The probability of measuring the control qubit in the zero state is $p = ||F(\cos(H))|\psi\rangle||^2$.

we choose $\{\phi_j\}$ such that $g(x, \{\phi_j\}) = F(x)$, the circuit in Fig. 1 implements a block encoding of $F(\cos(H/2))$, where $\varphi_j = \phi_j + (2 - \delta_{j0})\pi/4$ (see Appendix B in Ref. [98]). Combining these techniques leads to a procedure for calculating the phases $\{\phi_j\}$ completely in terms of SU(2) matrices, which we describe now.

Because we have n_{Ch} independent phases, to produce exactly the Chebyshev polynomial $F(x)$ we must sample at n_{Ch} values, which are taken to be the positive roots of the Chebyshev polynomial $T_{2n_{\text{Ch}}}(x)$ given by $x_j = \cos(\pi \frac{2k-1}{4n_{\text{Ch}}})$. If we define the functional $\mathcal{F}(\{\phi_j\})$ as

$$\mathcal{F}(\{\phi_j\}) = \frac{1}{n_{\text{Ch}}} \sum_{j=1}^{n_{\text{Ch}}} |g(x_j, \{\phi_j\}) - F(x_j)|^2, \tag{A2}$$

then the phases $\{\phi_j\}$ that produce $F(x)$ can be found by solving $\mathcal{F}(\{\phi_j\}) = 0$. It has been found in Ref. [116] that using a quasi-Newton method with a particular initial guess of

$$\{\phi_j^{(0)}\} = \left(\frac{\pi}{4}, 0, 0, \ldots, 0, 0, \frac{\pi}{4}\right) \tag{A3}$$

one can robustly find the symmetric phase factors for values of $d \sim 10\,000$. An example code to solve for the Chebyshev coefficients c_k and the associated phase factors has been implemented in QSPPACK [117]. Through numerical studies, we find that the cost of finding the phase factors scales roughly quadratically with the number of phases.

APPENDIX B: OPTIMAL CHOICE OF TIME STEP $\delta\tau$

In this Appendix, we discuss the optimal choice of the time step $\delta\tau$ that minimizes the total number of calls to the Trotter circuit N_{tot}, given by

$$N_{\text{tot}} = \frac{1}{b\Delta\delta\tau} \log\left(\frac{a}{\epsilon - c(\delta\tau)^p}\right). \tag{B1}$$

We start by showing that $dN_{\text{tot}}/d\delta\tau$ always has a zero. From there, we present a perturbative solution to the value of $\delta\tau$ that minimizes N_{tot}. We conclude by showing that $d^2N_{\text{tot}}/d\delta\tau^2$ is positive definite, which shows that the extreme value is always a minimum.

The first derivative is given by

$$\frac{dN_{\text{tot}}}{d\delta\tau} = \frac{1}{b\Delta\delta\tau^2}\left[c\,p\frac{\delta\tau^p}{\epsilon - c\delta\tau^p} - \log\left(\frac{a}{\epsilon - c(\delta\tau^*)^p}\right)\right]. \tag{B2}$$

For $\delta\tau = 0$, $dN_{\text{tot}}/d\delta\tau = -\log(a/\epsilon)$. Because $a > \epsilon$ in general, the first derivative is negative at $\delta\tau = 0$. This, combined with the fact that the function diverges to $+\infty$ when $\delta\tau \to (\epsilon/c)^{1/p}$ from the left, implies the first derivative always has a zero for some value of $\delta\tau$.

Setting the first derivative to zero gives

$$0 = c\,p\frac{(\delta\tau^*)^p}{\epsilon - c(\delta\tau^*)^p} - \log\left(\frac{a}{\epsilon - c(\delta\tau^*)^p}\right). \tag{B3}$$

To obtain approximate analytic values for $\delta\tau^*$, one can expand the logarithm. If we define $x = \frac{c}{\epsilon}\delta\tau^p$, then the equation becomes

$$0 = p\frac{x}{1-x} + \log\left(\frac{\epsilon}{a}\right) + \log(1-x). \tag{B4}$$

By expanding the log to lowest order, we obtain the following approximate equation for x

$$0 = p\frac{x}{1-x} + \log\left(\frac{\epsilon}{a}\right), \tag{B5}$$

which can be solved analytically to give the following expression

$$\delta\tau^* \approx \left[\frac{\epsilon}{c}\left(1 - \frac{p}{p + \log\left(\frac{a}{\epsilon}\right)}\right)\right]^{1/p}. \tag{B6}$$

Because $a > \epsilon$ in general, we see that this approximate solution for $\delta\tau^*$ is slightly smaller than the maximum value of $(\epsilon/c)^{1/p}$. Using the values $\epsilon = 10^{-3}$, $p = 1$, $a = 1$, $c = 0.1$, the approximate expression in Eq. (B6) gives a value of $\delta\tau^*$ which differs from the exact value obtained by solving Eq. (B4) numerically by only 3.2%.

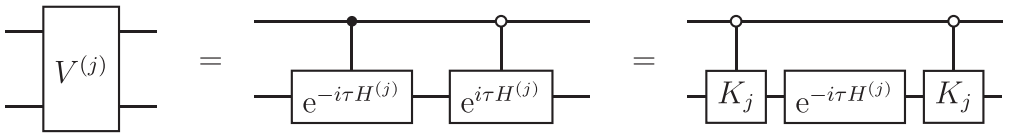

FIG. 20. Circuit relations demonstrating how the relation $K_j e^{-i\tau H^{(j)}} K_j = e^{i\tau H^{(j)}}$ can be leveraged to implement $V^{(j)}$ while only having to control the K_j Pauli operators.

The second derivative is given by

$$\frac{d^2N_{\rm tot}}{d\delta\tau^2} = \frac{1}{b\Delta(\delta\tau)^3}\left[cp(\delta\tau)^p\frac{(p-3)\epsilon+3c(\delta\tau)^p}{[\epsilon-c(\delta\tau)^p]^2} +2\log\left(\frac{a}{\epsilon-c(\delta\tau)^p}\right)\right]. \tag{B7}$$

The goal is to show that when $\delta\tau = \delta\tau^*$, the second derivative is positive. This will be done by implicitly solving Eq. (B4) to replace the log term, which gives

$$\frac{d^2N_{\rm tot}}{d\delta\tau^2}\Big|_{\delta\tau=\delta\tau^*} = \frac{1}{b\Delta(\delta\tau^*)^3}\left[cp(\delta\tau^*)^p\frac{(p-3)\epsilon+3c(\delta\tau^*)^p}{[\epsilon-c(\delta\tau^*)^p]^2} +2c\,p\frac{(\delta\tau^*)^p}{\epsilon-c(\delta\tau^*)^p}\right]. \tag{B8}$$

By combining the fractions, we obtain the following result

$$\frac{d^2N_{\rm tot}}{d\delta\tau^2}\Big|_{\delta\tau=\delta\tau^*} = \frac{c\,p(\delta\tau^*)^p}{b\Delta(\delta\tau^*)^3}\frac{1}{[\epsilon-c(\delta\tau^*)^p]^2} \times[\epsilon(p-1)+c(\delta\tau^*)^p]. \tag{B9}$$

Now, for any $p \geqslant 1$, and assuming all parameters are positive, we see $\frac{d^2N_{\rm tot}}{d\delta\tau^2}|_{\delta\tau=\delta\tau^*} > 0$.

APPENDIX C: CONTROL-FREE IMPLEMENTATION OF GROUND-STATE PREPARATION IN LATTICE FIELD THEORIES

In this Appendix, we show how to prepare the ground state of the U(1) gauge theory described in Sec. III using the control-free version of QETU.

We first describe the general procedure for implementing the control-free version of QETU originally described in Ref. [98]. Suppose that one has access to an oracle that implements a controlled call to both forward and backward time evolution simultaneously,

$$V = \begin{pmatrix} e^{i\tau H} & 0 \\ 0 & e^{-i\tau H}\end{pmatrix}. \tag{C1}$$

Instead of a controlled call to $U = e^{-i\tau H}$, as in the original QETU circuit, the oracle V can then be used as a building block. The control-free QETU theorem [98] assumes the access to a circuit implementing V for an n-qubit Hermitian operator H. It states that for any even real polynomial $F(x)$ of degree d satisfying $|F(x)| \leqslant 1 \;\forall\; x \in [-1,1]$, one can find a sequence of symmetric phase factors $\{\varphi_j\} = (\varphi_0, \varphi_1, \ldots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$, such that the circuit in Fig. 19 denoted by $\mathcal{U}$ satisfies $(\langle 0|\otimes \mathbb{1}_n)\mathcal{U}(|0\rangle\otimes\mathbb{1}_n) = F(\cos(H))$ [98].

We now describe a general procedure for preparing V assuming one has access to $U = e^{-i\tau H}$. The procedure for implementing V involves grouping the terms of H into l groups $H = \sum_{j=1}^{l} H^{(j)}$ such that each term in $H^{(j)}$ anticommutes with the Pauli operator K_j, i.e., $K_j H^{(j)} K_j = -H^{(j)}$. We now show that, once this grouping is found, only the K_j operators must be controlled. We now define

$$V^{(j)} = \begin{pmatrix} e^{i\tau H^{(j)}} & 0 \\ 0 & e^{-i\tau H^{(j)}}\end{pmatrix}, \tag{C2}$$

such that $V = \prod_{j=1}^{l} V^{(j)}$. Using the key relation $K_j e^{-i\tau H^{(j)}} K_j = e^{i\tau H^{(j)}}$, each $V^{(j)}$ will be implemented using the circuit identity in Fig. 20, in which only the K_j operators need to be controlled. If one finds l such groups, instead of having to control each term in $e^{-i\tau H}$, the control-free version of QETU requires only an additional $O(l)$ controlled operations.

We now explain how to prepare V for a general theory where the Hamiltonian is a sum of a kinetic and potential piece. Furthermore, we assume that the basis in which the kinetic piece is diagonal is related to the basis in which the potential piece is diagonal by a Fourier transformation. Under these assumptions, the Hamiltonian is written as $H = H_x + \mathcal{F}^\dagger H_p \mathcal{F}$, where H_x and H_p are the potential and kinetic components, respectively, and $\mathcal{F}$ is the Fourier transform. Note that H_x and H_p as written are diagonal matrices. A single Trotter step is implemented as $U(\delta\tau) = \mathcal{F}^\dagger e^{-i\delta\tau H_p}\mathcal{F} e^{-i\delta\tau H_x}$. The procedure we show below can be used for any system of this type, including the U(1) gauge theory described in Sec. III.

Because the controlled implementation of a product of unitary operators is equal to the product of the controlled versions of the individual operators, we only need to consider a single Trotter step; the total control-free time evolution operator can then be built from combining multiple Trotter steps. First, we show that the Fourier transform does not need to be controlled. We denote by $U_{\rm ctrl}(\delta\tau)$ the controlled version of a single Trotter step. Explicit evaluation gives

$$\begin{aligned} U_{\rm ctrl}(\delta t) &= (\mathbb{1}\otimes|0\rangle\langle 0| + \mathcal{F}^\dagger e^{-i\delta t H_p}\mathcal{F}\otimes|1\rangle\langle 1|)(\mathbb{1}\otimes|0\rangle\langle 0| + e^{-i\delta t H_x}\otimes|1\rangle\langle 1|) \\ &= (\mathcal{F}^\dagger\mathcal{F}\otimes|0\rangle\langle 0| + \mathcal{F}^\dagger e^{-i\delta t H_p}\mathcal{F}\otimes|1\rangle\langle 1|)(\mathbb{1}\otimes|0\rangle\langle 0| + e^{-i\delta t H_x}\otimes|1\rangle\langle 1|) \\ &= (\mathcal{F}^\dagger\otimes\mathbb{1})(\mathbb{1}\otimes|0\rangle\langle 0| + e^{-i\delta t H_p}\otimes|1\rangle\langle 1|)(\mathcal{F}\otimes\mathbb{1})(\mathbb{1}\otimes|0\rangle\langle 0| + e^{-i\delta t H_x}\otimes|1\rangle\langle 1|), \end{aligned} \tag{C3}$$

where we used $\mathcal{F}^\dagger\mathcal{F} = \mathbb{1}$. From this we see that the Fourier transform does not need to be controlled.

Next, using the following circuit identity:

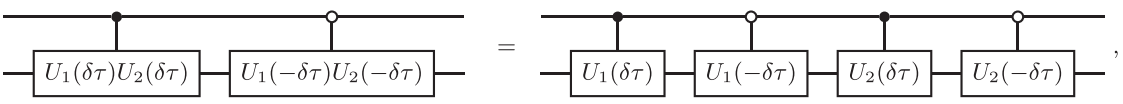

and replacing $U_1(\delta\tau) = U_x(\delta\tau) \equiv e^{-i\delta t H_x}$ and $U_2(\delta\tau) = \mathcal{F}^\dagger U_p(\delta\tau)\mathcal{F} \equiv \mathcal{F}^\dagger e^{-i\delta t H_p}\mathcal{F}$, we find

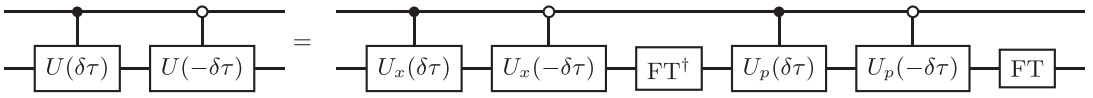

Note that H_x and H_p are both diagonal matrices. If we can find a general procedure for decomposing a general diagonal unitary matrix into l groups $H^{(j)}$ that all anticommute with the same K_j Pauli operator, we have a control-free implementation of QETU for Hamiltonians of this form. We now present a general method applicable to arbitrary Hamiltonians; applying this method to the specialized case of diagonal Hamiltonians is straightforward.

A general Hamiltonian $\tilde{H}$ acting n qubits can be decomposed into a sum of 4^n Pauli strings containing $\mathbb{1}, X, Y, Z$ matrices.[4] The method for choosing the $H^{(j)}$ groups and their associated K_j Pauli strings contains two main steps. The first step considers only Pauli strings containing $\mathbb{1}$ and Z matrices. The process starts by defining $H_Z^{(1)}$ as the sum of all Pauli strings with a Z gate acting on the first qubit. The associated $K_1^{(Z)}$ operator is simply an X gate acting on the first qubit. From the remaining pool of Pauli strings we now define $H_Z^{(2)}$ as a sum of all terms with a Z gate acting on the second qubit, and $K_2^{(Z)}$ is an X gate acting on the second qubit. By repeating this method for all n qubits, we construct n such groups $H_Z^{(j)}$ and their associated $K_j^{(Z)}$ operators (if the Hamiltonian $\tilde{H}$ is diagonal, one can stop here and implement the control-free version of QETU).

The second step focuses on the remaining Pauli strings that contain only $\mathbb{1}$, X, Y matrices. In a similar way, we start by defining $H_{XY}^{(1)}$ as the sum of all Pauli strings with either an X or Y gate acting on the first qubit. Because Z anticommutes with both X and Y, the associated $K_1^{(XY)}$ operator is simply a Z gate acting on the first qubit. From the remaining pool of operators, we can now choose $H_{XY}^{(2)}$ as a sum of all terms with either an X or Y gate acting on the second qubit, and $K_2^{(XY)}$ will be a Z gate acting on the second qubit. After repeating this for all n qubits, we have n groups n groups $H_{XY}^{(j)}$ and their associated $K_j^{(XY)}$ operators. Using this method, the operator V needed to implement the control-free version of QETU for arbitrary Hamiltonians $\tilde{H}$ can therefore be implemented with the same cost as $e^{i\tau\tilde{H}}$ plus an extra $O(n)$ CNOT gate. For many physical systems of interest, this additional cost is negligible compared with implementing the uncontrolled time evolution circuit.

As a demonstration of this method, we work through a two qubit example considering a diagonal Hamiltonian. A general two qubit diagonal Hamiltonian can be written as

$$\tilde{H} = a_0\mathbb{1}\otimes\mathbb{1} + a_1\mathbb{1}\otimes Z + a_2 Z\otimes\mathbb{1} + a_3 Z\otimes Z \tag{C4}$$

for general coefficients a_i. The term associated with a_0 is a global phase and can be ignored for this argument. Using the previously described method, the groups $H^(j)$ are given by

$$H^{(1)} = a_2 Z\otimes\mathbb{1} + a_3 Z\otimes Z, \tag{C5}$$

$$H^{(2)} = a_1\mathbb{1}\otimes Z. \tag{C6}$$

For this choice of groupings, the associated K_j operators are

$$9K_1 = X\otimes\mathbb{1}, \tag{C7}$$

$$K_2 = \mathbb{1}\otimes X. \tag{C8}$$

With these choices, the circuit for the V operator is

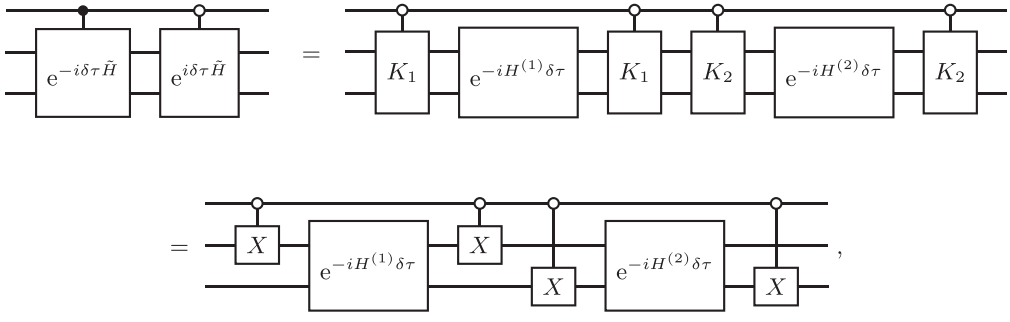

which requires an additional four CNOT gates, instead of controlling on the entire time evolution operator $e^{-i\delta\tau\tilde{H}}$.

[4]Note that, in a practical setting, to implement the exponential of an arbitrary Hamiltonian, one first splits the Hamiltonian into groups of commuting Pauli strings. Constructing the V operator for control-free QETU will then require using our method for each set of commuting Pauli strings.

[1] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
[2] S. Lloyd, Science **273**, 1073 (1996).
[3] S. Wiesner, arXiv:quant-ph/9603028.
[4] D. S. Abrams and S. Lloyd, Phys. Rev. Lett. **79**, 2586 (1997).
[5] C. Zalka, Proc. R. Soc. London, Ser. A **454**, 313 (1998).
[6] T. Byrnes and Y. Yamamoto, Phys. Rev. A **73**, 022328 (2006).
[7] J. Kogut and L. Susskind, Phys. Rev. D **11**, 395 (1975).
[8] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Phys. Rev. Lett. **129**, 051601 (2022).
[9] J. C. Halimeh and P. Hauke, Phys. Rev. Lett. **125**, 030503 (2020).
[10] J. C. Halimeh, H. Lang, J. Mildenberger, Z. Jiang, and P. Hauke, PRX Quantum **2**, 040311 (2021).
[11] J. C. Halimeh, L. Homeier, C. Schweizer, M. Aidelsburger, P. Hauke, and F. Grusdt, Phys. Rev. Res. **4**, 033120 (2022).
[12] J. C. Halimeh, H. Lang, and P. Hauke, New J. Phys. **24**, 033015 (2022).
[13] H. Lamm, S. Lawrence, and Y. Yamauchi, arXiv:2005.12688.
[14] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, PRX Quantum **2**, 010323 (2021).
[15] E. Zohar, J. I. Cirac, and B. Reznik, Phys. Rev. Lett. **109**, 125302 (2012).
[16] D. Banerjee, M. Dalmonte, M. Müller, E. Rico, P. Stebler, U.-J. Wiese, and P. Zoller, Phys. Rev. Lett. **109**, 175302 (2012).
[17] V. Kasper, T. V. Zache, F. Jendrzejewski, M. Lewenstein, and E. Zohar, Phys. Rev. D **107**, 014506 (2023).
[18] P. Hauke, D. Marcos, M. Dalmonte, and P. Zoller, Phys. Rev. X **3**, 041018 (2013).
[19] S. Kühn, J. I. Cirac, and M.-C. Bañuls, Phys. Rev. A **90**, 042305 (2014).
[20] K. Stannigel, P. Hauke, D. Marcos, M. Hafezi, S. Diehl, M. Dalmonte, and P. Zoller, Phys. Rev. Lett. **112**, 120406 (2014).
[21] J. R. Stryker, Phys. Rev. A **99**, 042301 (2019).
[22] R. Anishetty, M. Mathur, and I. Raychowdhury, J. Math. Phys. **50**, 053503 (2009).
[23] I. Raychowdhury and J. R. Stryker, Phys. Rev. Res. **2**, 033039 (2020).
[24] I. Raychowdhury and J. R. Stryker, Phys. Rev. D **101**, 114502 (2020).
[25] M. Mathur, I. Raychowdhury, and R. Anishetty, J. Math. Phys. **51**, 093504 (2010).
[26] M. Mathur, I. Raychowdhury, and T. P. Sreeraj, J. Math. Phys. **52**, 113505 (2011).
[27] I. Raychowdhury, Prepotential formulation of lattice Gauge theories, Ph.D. thesis, University of Calcutta, 2013.
[28] A. Alexandru, P. F. Bedaque, S. Harmalkar, H. Lamm, S. Lawrence, and N. C. Warrington (NuQS Collaboration), Phys. Rev. D **100**, 114501 (2019).
[29] E. Zohar and M. Burrello, Phys. Rev. D **91**, 054506 (2015).
[30] E. Zohar, A. Farace, B. Reznik, and J. I. Cirac, Phys. Rev. A **95**, 023604 (2017).
[31] Y. Ji, H. Lamm, and S. Zhu (NuQS Collaboration), Phys. Rev. D **102**, 114513 (2020).
[32] A. Alexandru, P. F. Bedaque, R. Brett, and H. Lamm, Phys. Rev. D **105**, 114508 (2022).
[33] E. J. Gustafson, H. Lamm, and F. Lovelace, Phys. Rev. D **109**, 054503 (2024).
[34] D. C. Hackett, K. Howe, C. Hughes, W. Jay, E. T. Neil, and J. N. Simone, Phys. Rev. A **99**, 062341 (2019).
[35] M. Kreshchuk, W. M. Kirby, G. Goldstein, H. Beauchemin, and P. J. Love, Phys. Rev. A **105**, 032418 (2022).
[36] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, Entropy **23**, 597 (2021).
[37] D. B. Kaplan, E. Katz, and M. Ìnsal, J. High Energy Phys. 05 (2003) 037.
[38] A. J. Buser, H. Gharibyan, M. Hanada, M. Honda, and J. Liu, J. High Energy Phys. 09 (2021) 034.
[39] J. Bender and E. Zohar, Phys. Rev. D **102**, 114517 (2020).
[40] J. F. Unmuth-Yockey, Phys. Rev. D **99**, 074502 (2019).
[41] C. W. Bauer and D. M. Grabowska, arXiv:2111.08015.
[42] D. B. Kaplan and J. R. Stryker, Phys. Rev. D **102**, 094515 (2020).
[43] J. F. Haase, L. Dellantonio, A. Celi, D. Paulson, A. Kan, K. Jansen, and C. A. Muschik, Quantum **5**, 393 (2021).
[44] C. W. Bauer, I. D'Andrea, M. Freytsis, and D. M. Grabowska, arXiv:2307.11829.
[45] N. Klco, M. J. Savage, and J. R. Stryker, Phys. Rev. D **101**, 074512 (2020).
[46] P. W. Shor, Phys. Rev. A **52**, R2493 (1995).
[47] A. Y. Kitaev, in *Quantum Communication, Computing, and Measurement*, edited by O. Hirota, A. S. Holevo, and C. M. Caves (Springer, Boston, MA, 1997), pp. 181–188.
[48] A. Y. Kitaev, Ann. Phys. (NY) **303**, 2 (2003).
[49] A. Steane, Proc. R. Soc. London, Ser. A **452**, 2551 (1996).
[50] D. A. Lidar and T. A. Brun, *Quantum Error Correction* (Cambridge University Press, 2013).
[51] D. Litinski, Quantum **3**, 128 (2019).
[52] B. J. Brown, K. Laubscher, M. S. Kesselring, and J. R. Wootton, Phys. Rev. X **7**, 021029 (2017).
[53] H. Bombin, C. Dawson, R. V. Mishmash, N. Nickerson, F. Pastawski, and S. Roberts, PRX Quantum **4**, 020303 (2023).
[54] A. Rajput, A. Roggero, and N. Wiebe, npj Quantum Inf. **9**, 41 (2023).
[55] A. J. Landahl and B. C. Morrison, arXiv:2110.10280.
[56] S. P. Jordan, K. S. M. Lee, and J. Preskill, Quant. Inf. Comput. **14**, 1014 (2014).
[57] L. García-Álvarez, J. Casanova, A. Mezzacapo, I. L. Egusquiza, L. Lamata, G. Romero, and E. Solano, Phys. Rev. Lett. **114**, 070502 (2015).
[58] S. P. Jordan, K. S. M. Lee, and J. Preskill, arXiv:1404.7115.
[59] S. P. Jordan, H. Krovi, K. S. M. Lee, and J. Preskill, Quantum **2**, 44 (2018).
[60] A. Hamed Moosavian and S. Jordan, Phys. Rev. A **98**, 012332 (2018).
[61] A. H. Moosavian, J. R. Garrison, and S. P. Jordan, arXiv:1911.03505.
[62] N. Klco and M. J. Savage, Phys. Rev. A **102**, 012612 (2020).
[63] N. Klco and M. J. Savage, Phys. Rev. A **102**, 012619 (2020).
[64] S. Harmalkar, H. Lamm, and S. Lawrence (NuQS Collaboration), arXiv:2001.11490.
[65] A. N. Ciavarella, S. Caspar, H. Singh, and M. J. Savage, Phys. Rev. A **107**, 042404 (2023).
[66] C. Kokail *et al.*, Nature (London) **569**, 355 (2019).
[67] A. Roggero, A. C. Y. Li, J. Carlson, R. Gupta, and G. N. Perdue, Phys. Rev. D **101**, 074038 (2020).
[68] D. B. Kaplan, N. Klco, and A. Roggero, arXiv:1709.08250.
[69] Z. Davoudi, N. Mueller, and C. Powers, Phys. Rev. Lett. **131**, 081901 (2023).

[70] N. Wiebe, D. Berry, P. Høyer, and B. C. Sanders, J. Phys. A: Math. Theor. **43**, 065203 (2010).

[71] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Nat. Commun. **5**, 4213 (2014).

[72] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, New J. Phys. **18**, 023023 (2016).

[73] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, Nat. Commun. **10**, 3007 (2019).

[74] A. Shlosberg, A. J. Jena, P. Mukhopadhyay, J. F. Haase, F. Leditzky, and L. Dellantonio, Quantum **7**, 906 (2023).

[75] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth *et al.*, Phys. Rep. **986**, 1 (2022).

[76] D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, Mater. Theory **6**, 2 (2022).

[77] A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Nat. Rev. Phys. **3**, 625 (2021).

[78] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Phys. Rev. A **95**, 042308 (2017).

[79] J. R. McClean, Z. Jiang, N. C. Rubin, R. Babbush, and H. Neven, Nat. Commun. **11**, 636 (2020).

[80] T. Takeshita, N. C. Rubin, Z. Jiang, E. Lee, R. Babbush, and J. R. McClean, Phys. Rev. X **10**, 011004 (2020).

[81] N. Yoshioka, T. Sato, Y. O. Nakagawa, Y. Y. Ohnishi, and W. Mizukami, Phys. Rev. Res. **4**, 013052 (2022).

[82] N. Yoshioka, H. Hakoshima, Y. Matsuzaki, Y. Tokunaga, Y. Suzuki, and S. Endo, Phys. Rev. Lett. **129**, 020502 (2022).

[83] D. W. Berry, A. M. Childs, and R. Kothari, in *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science, Berkeley, CA, USA* (IEEE, Piscataway, NJ, 2015), pp. 792–809.

[84] G. H. Low and I. L. Chuang, Phys. Rev. Lett. **118**, 010501 (2017).

[85] G. H. Low and I. L. Chuang, Quantum **3**, 163 (2019).

[86] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, SIAM J. Comput. **52**, FOCS18-250 (2023).

[87] G. H. Low and N. Wiebe, arXiv:1805.00675.

[88] M. Kieferová, A. Scherer, and D. W. Berry, Phys. Rev. A **99**, 042314 (2019).

[89] D. W. Berry, A. M. Childs, Y. Su, X. Wang, and N. Wiebe, Quantum **4**, 254 (2020).

[90] D. Motlagh and N. Wiebe, arXiv:2308.01501.

[91] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, arXiv:1806.01838.

[92] S. Hariprakash, N. S. Modi, M. Kreshchuk, C. F. Kane, and C. W. Bauer, arXiv:2312.11637.

[93] M. Rhodes, M. Kreshchuk, and S. Pathak, arXiv:2405.10416.

[94] R. D. Somma, New J. Phys. **21**, 123025 (2019).

[95] K. Choi, D. Lee, J. Bonitati, Z. Qian, and J. Watkins, Phys. Rev. Lett. **127**, 040505 (2021).

[96] Z. Qian, J. Watkins, G. Given, J. Bonitati, K. Choi, and D. Lee, arXiv:2110.07747.

[97] L. Lin and Y. Tong, PRX Quantum **3**, 010318 (2022).

[98] Y. Dong, L. Lin, and Y. Tong, PRX Quantum **3**, 040305 (2022).

[99] H. H. S. Chan, D. Muñoz-Ramo, and N. Fitzpatrick, arXiv:2303.06161.

[100] D. M. Grabowska, C. Kane, B. Nachman, and C. W. Bauer, arXiv:2208.03333.

[101] C. Kane, D. M. Grabowska, B. Nachman, and C. W. Bauer, arXiv:2211.10497.

[102] L. Lin, arXiv:2201.08309.

[103] D. Poulin and P. Wocjan, Phys. Rev. Lett. **102**, 130503 (2009).

[104] Y. Ge, J. Tura, and J. I. Cirac, J. Math. Phys. **60**, 022202 (2019).

[105] L. Lin and Y. Tong, Quantum **4**, 372 (2020).

[106] T. D. Cohen and H. Oh, arXiv:2305.19952.

[107] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Phys. Rev. Lett. **114**, 090502 (2015).

[108] R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, New J. Phys. **18**, 033032 (2016).

[109] R. Babbush, D. W. Berry, Y. R. Sanders, I. D. Kivlichan, A. Scherer, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, Quantum Sci. Technol. **3**, 015006 (2018).

[110] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Phys. Rev. X **8**, 041015 (2018).

[111] W. Kirby, M. Motta, and A. Mezzacapo, Quantum **7**, 1018 (2023).

[112] G. H. Low and I. L. Chuang, arXiv:1707.05391.

[113] L. N. Trefethen, *Approximation Theory and Approximation Practice, Extended Edition* (Society for Industrial and Applied Mathematics, Philadelphia, 2019).

[114] M. Grant and S. Boyd, CVX: Matlab software for disciplined convex programming, version 2.1, http://cvxr.com/cvx (2014).

[115] S. Diamond and S. Boyd, J. Mach. Learn. Res. **17**, 1 (2016).

[116] Y. Dong, X. Meng, K. B. Whaley, and L. Lin, Phys. Rev. A **103**, 042419 (2021).

[117] Qsppack, https://github.com/qsppack/QSPPACK.

[118] N. Klco and M. J. Savage, Phys. Rev. A **99**, 052335 (2019).

[119] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).

[120] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Phys. Rev. D **104**, 094519 (2021).

[121] A. M. Childs and N. Wiebe, Quant. Inf. Comput. **12**, 901 (2012).

[122] J. Bender, P. Emonts, E. Zohar, and J. I. Cirac, Phys. Rev. Res. **2**, 043145 (2020).

[123] L. K. Kovalsky, F. A. Calderon-Vargas, M. D. Grace, A. B. Magann, J. B. Larsen, A. D. Baczewski, and M. Sarovar, Phys. Rev. Lett. **131**, 060602 (2023).

[124] D. Cheung, P. Høyer, and N. Wiebe, J. Phys. A: Math. Theor. **44**, 415302 (2011).

[125] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, arXiv:2405.08810.

[126] A. Kitaev and W. A. Webb, arXiv:0801.0342.

[127] P. Deliyannis, M. Freytsis, B. Nachman, and C. W. Bauer, arXiv:2109.10918.

[128] S. P. Jordan, K. S. Lee, and J. Preskill, arXiv:1112.4833.

[129] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, arXiv:quant-ph/0001106.

[130] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, Science **292**, 472 (2001).

[131] D. Aharonov and A. Ta-Shma, in *Proceedings of the Thirty-Fifth Annual ACM Symposium on the Theory of Computing*

(Association for Computing Machinery, New York, 2003), pp. 20–29.
[132] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, npj Quantum Inf. **5**, 75 (2019).
[133] N. Gomes, A. Mukherjee, F. Zhang, T. Iadecola, C.-Z. Wang, K.-M. Ho, P. P. Orth, and Y.-X. Yao, Adv. Quantum Technol. **4**, 2100114 (2021).
[134] J. Liu, Z. Li, H. Zheng, X. Yuan, and J. Sun, Mach. Learn.: Sci. Technol. **3**, 045030 (2022).
[135] J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, Phys. Rev. Lett. **92**, 177902 (2004).
[136] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Phys. Rev. Lett. **93**, 130502 (2004).
[137] N. Khaneja and S. J. Glaser, Chem. Phys. **267**, 11 (2001).
[138] S. S. Bullock and G. K. Brennen, J. Math. Phys. **45**, 2447 (2004).
[139] H. N. S. Earp and J. K. Pachos, J. Math. Phys. **46**, 082108 (2005).
[140] B. Drury and P. Love, J. Phys. A: Math. Theor. **41**, 395305 (2008).
[141] R. Iten, R. Colbeck, I. Kukuljan, J. Home, and M. Christandl, Phys. Rev. A **93**, 032318 (2016).
[142] J. A. Cortese and T. M. Braje, arXiv:1803.01958.
[143] X. Sun, G. Tian, S. Yang, P. Yuan, and S. Zhang, arXiv:2108.06150.
[144] X.-M. Zhang, M.-H. Yung, and X. Yuan, Phys. Rev. Res. **3**, 043200 (2021).
[145] X.-M. Zhang, T. Li, and X. Yuan, Phys. Rev. Lett. **129**, 230504 (2022).
[146] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, in *Contemporary Mathematics*, edited by S. J. Lomonaco and H. E. Brandt (American Mathematical Society, Providence, 2002), Vol. 305.
[147] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice* (Springer, Berlin, 2010), Vol. 788.
[148] S. Hariprakash, N. S. Modi, M. Kreshchuk, C. F. Kane, and C. W. Bauer, arXiv:2312.11637.
[149] Z. M. Rossi and I. L. Chuang, Quantum **6**, 811 (2022).