Efficient vacuum-state preparation for quantum simulation of strongly interacting local quantum field theories

Thomas D. Cohen ●* and Hyunwoo Oh ●†

Department of Physics and Maryland Center for Fundamental Physics, University of Maryland, College Park, Maryland 20742, USA

We present an efficient approach for preparing ground states in the context of strongly interacting local quantum field theories on quantum computers. The approach produces the vacuum state in a time proportional to the square root of the volume, which is a square-root improvement in speed compared to traditional approaches. The approach exploits a method for traversing the path in parameter space in which the resources scale linearly with a path length suitably defined in parameter space. Errors due to practical limitations are controlled and do not exhibit secular growth along the path. The final accuracy can be arbitrarily improved with an additive cost, which is independent of the volume and grows slower than logarithmically with the overlap between the state produced and the exact ground state. We expect that the method could potentially hold practical value not only within the realm of quantum field theories but also in addressing other challenges involving long path lengths.

Simulating aspects of quantum field theories (QFTs) on classical computers is often hindered by the notorious sign problem [1,2]; quantum computing is a candidate to evade the sign problem and simulate QFTs [3–11]. Studying QFTs on quantum computers requires the preparation of initial states and various state preparation methods have been suggested [12–30]. This Letter focuses on the preparation of vacuum states of lattice field theories (with local interactions) aimed at future quantum computations. The classes of theories considered are ones with finite correlation lengths and gapped spectra. The principal purpose of this Letter is to show that the exact vacuum state of the (lattice) theories can be computed with an arbitrarily high probability with a computational time that scales as the square root of the volume; an explicit algorithm doing this is developed.

In contrast to some problems in quantum information [31,32], in many physics applications it is essential to prepare a physical state a very large number of times: Numerous properties of the state will be studied and the calculation of any of these requires a large number of runs. For these problems, it can be efficient to expend substantial resources the first time(s) the state is prepared in order to subsequently prepare the state more efficiently. The algorithm considered here is of this sort.

One general approach to state preparation is to begin with a system whose vacuum state can be explicitly constructed and then evolve the system (by varying a parameter λ specifying the Hamiltonian) in such a manner that—to a good approximation—the system remains in $|g(\lambda)\rangle$, the ground state of $\hat{H}(\lambda)$. Adiabatic state preparation (ASP) [18,33–39] and methods based on the quantum Zeno effect (QZE) exploiting projection [40–44] are in this category. One can always choose parametrizations where λ is dimensionless.

There are two key issues in optimizing such algorithms. An obvious one is the choice of path through parameter space, since clearly this affects how efficiently a ground state can be produced. The other is finding an efficient manner to traverse the path to find the ground state of interest. This Letter addresses the second issue.

There are important technical restrictions when using the approach outlined here: For the algorithm to be efficient the theory must remain local for all values of λ along the path and that no phase transitions (in the infinite volume theory) are encountered. We recognize that the restriction that no phase transitions are encountered along the path may reduce the applicability of this approach, but also note that this restriction applies generally to all methods based on traversing paths in parameter space including ASP and methods based on the QZE.

Even in the absence of phase transitions, there is a challenge faced by such methods for field theories with large volumes: The natural path through the various ground states is necessarily long—with path lengths scaling with the square root of the volume. It is important to note that while the size of the minimum spectral gap has received significant attention particularly in the context of adiabatic quantum computing [45–49], the total cost of traversing a path in parameter space is also quite sensitive to the path length. This should be clear if one imagines a Hamiltonian trajectory that has been rescaled so the spectral gap is constant along the trajectory. Clearly, longer paths are more expensive to traverse. The approach developed here is designed to be efficient for long paths and should be useful for a variety of problems with long paths, but it is particularly useful for QFTs with large volumes.

We follow a standard approach to define the path length. Denote $|g'(\lambda)\rangle \equiv \frac{d|g(\lambda)\rangle}{d\lambda}$, where $|g(\lambda)\rangle$ is a differentiable function of λ. Phases are chosen conventionally so that

$$\forall \lambda_a, \lambda_b, \quad \text{Im}[\langle g(\lambda_a)|g(\lambda_b)\rangle] = 0. \tag{1a}$$

*cohen@umd.edu
†hyunwooh@umd.edu

L_{λ_a, λ_b}, the path length from a to b (where b is a point further along the path than a), and total path length L are then defined [40–42] as

$$L_{\lambda_a, \lambda_b} \equiv \int_{\lambda_a}^{\lambda_b} d\lambda \|\|g'(\lambda)\rangle\|, \quad L = L_{\lambda_0, \lambda_f}; \quad (1b)$$

λ_0 and λ_f are the initial and final values of λ for the full path. L_{λ_a, λ_b} has a property of the path and is invariant under reparametrization of the path.

One needs to model the computational cost so that comparisons between different methods are straightforward. One important cost is the relative time needed by various schemes to complete the calculation on the same equipment; this cost is denoted as C. C depends on how the calculation is realized. A useful and concrete way to envision it is to consider a hybrid between analog and digital quantum computers; the system undergoes continuous time evolution according to a Hamiltonian but quantum gates can control the circumstance for which this time evolution occurs and gated operations on a state are allowable. This avoids issues associated with the cost to achieve the desired accuracy using Trotterization while still allowing algorithms which require quantum control. Since rescaling the Hamiltonian by A reduces the time by $1/A$, to make predictions of the relative cost concrete, we assume that $H(\lambda)$ has always been scaled to the maximum that the equipment can accommodate.

Schemes directly based on the QZE [50] generically have

$$C \sim T \gtrsim \frac{L^2}{\Delta}, \quad (2)$$

where Δ is the typical spectral gap along the path. Theorems for an upper bound on the time needed for ASP with fixed error [45–47] are consistent with Eq. (2).

As will be shown, the path length is proportional to the square root of the volume of the system,

$$\frac{L_{V_2}}{L_{V_1}} = \sqrt{\frac{V_2}{V_1}}. \quad (3)$$

The costs of schemes based on QZE or ASP then scale linearly with the volume. In this Letter, a scheme with $C \sim \sqrt{V}$ is developed.

Let us derive Eq. (3). The locality of the interactions and the finiteness of the correlation lengths imply that at large volumes

$$\langle g'(\lambda_0)|g'(\lambda_0)\rangle = -\frac{\partial^2 \langle g(\lambda_1)|g(\lambda_0)\rangle}{\partial(\lambda_1)^2}\bigg|_{\lambda_1 = \lambda_0}$$
$$= \left(\frac{V}{\xi^d}\right) h(\lambda_0) \times \left[1 + O\left(\frac{\xi^d}{V}\right)\right], \quad (4a)$$

where h is independent of the volume at large V, ξ is the correlation length, and d is the dimension of space. More generally,

$$\frac{\partial^n \langle g(\lambda_0 + \Delta\lambda/2)|g(\lambda_0 - \Delta\lambda/2)\rangle}{\partial(\Delta\lambda)^n}\bigg|_{\Delta\lambda=0}$$
$$= -\frac{1 + (-1)^n}{2} \frac{\langle g'(\lambda_0)|g'(\lambda_0)\rangle^{\frac{n}{2}} n!}{2^{n/2}(n/2)!}\left[1 + O\left(\frac{\xi^d}{V}\right)\right], \quad (4b)$$

implying

$$\langle g(\lambda_1)|g(\lambda_0)\rangle = e^{-L_{\lambda_0, \lambda_1}^2/2}\left[1 + O\left(\frac{\xi^d}{V}\right)\right] \quad \text{and}$$

$$\|\|g'(\lambda_0 + \Delta\lambda)\rangle\| = \|\|g'(\lambda_0)\rangle\| \times \left[1 + O\left(\frac{\xi^d}{V}\right)^{\frac{1}{2}}\right], \quad (4c)$$

which in turn implies that $L \sim (\frac{V}{\xi^d})^{\frac{1}{2}}$.

The scheme proposed here differs fundamentally from ASP [where the system passes through every ground state from the initial one, $|g(\lambda_0)\rangle$, to the final one, $|g(\lambda_f)\rangle$]. It also differs fundamentally from conventional QZE schemes where only discrete points along the path, specified by λ_i, are considered, but adjacent points on the path are close: $|\langle g(\lambda_{i+1})|g(\lambda_i)\rangle|^2 \approx 1$.

As QZE-based schemes, our approach uses discrete points along the path. However, instead of large overlaps between neighboring points on the path, ideally the points are chosen so that

$$|\langle g(\lambda_{i+1})|g_i(\lambda_i)\rangle|^2 = 1/2, \quad L_{\lambda_i, \lambda_{i+1}} = \sqrt{\log(2)}; \quad (5)$$

the second form is valid for large volumes. We assume that in initial studies, the set of points satisfying Eq. (5) has been found (to a good approximation).

An idealized version of the algorithm invokes an oracle $\hat{U}(\lambda)_\gamma$, parametrized by λ and γ, where γ can take the value of a or b. $\hat{U}(\lambda)_\gamma$ acts on a state combining $|\psi\rangle$, the state describing the system of interest, with two ancillary qubits $|\chi\rangle_a$ and $|\chi\rangle_b$. The action of $\hat{U}(\lambda)_\gamma$ is given by

$$\hat{U}(\lambda)_a = |g(\lambda)\rangle\langle g(\lambda)| \otimes \hat{I}_a \otimes \hat{I}_b$$
$$+ [\hat{I}_\psi - |g(\lambda)\rangle\langle g(\lambda)|] \otimes \hat{\sigma}_a^x \otimes \hat{I}_b,$$
$$\hat{U}(\lambda)_b = |g(\lambda)\rangle\langle g(\lambda)| \otimes \hat{I}_a \otimes \hat{I}_b$$
$$+ [\hat{I}_\psi - |g(\lambda)\rangle\langle g(\lambda)|] \otimes \hat{I}_a \otimes \hat{\sigma}_b^x, \quad (6)$$

where $\hat{\sigma}_a^x$ and $\hat{\sigma}_b^x$ are Pauli operators acting on the ancillas, $\hat{I}_a$ and $\hat{I}_b$ are identity operators acting on the ancillas, while $\hat{I}_\psi$ is the identity acting on the state of the system.

$\hat{U}_a(\lambda)$ flips $|\chi\rangle_a$, if the system is in an excited state and leaves it if the system is in the ground state; $\hat{U}_b(\lambda)$ does the same for $|\chi\rangle_b$. Highly accurate approximations of the oracle can be constructed using, for example, Kitaev's phase estimation scheme [51] or a variant of the rodeo algorithm or similar approaches [19,20,52,53]. It is straightforward to show that when $|e\rangle$ is a superposition of excited states of $H(\lambda)$ and $|\alpha|^2 + |\beta|^2 = 1$,

$$\hat{U}(\lambda)_b\{[\alpha|g(\lambda)\rangle + \beta|e\rangle] \otimes |\chi\rangle_a \otimes \hat{H}^{\text{ad}}|\downarrow\rangle_b\}$$
$$= [\alpha|g(\lambda)\rangle - \beta|e\rangle] \otimes |\chi\rangle_a \otimes \hat{H}^{\text{ad}}|\downarrow\rangle_b, \quad (7)$$

where $\hat{H}^{\text{ad}}$ is the Hadamard operator; the effect of $\hat{U}(\lambda)_b$ is to reflect any state in the physical Hilbert space around $|g(\lambda)\rangle$.

Consider an idealized case: The oracle is implemented exactly and the points λ_i are known exactly. Then a simple scheme acts to move a system in the ground state at point λ_i to the ground state at point λ_{i+1}. The scheme requires that $|\chi\rangle_a$ begins as $|\uparrow\rangle_a$ while $|\chi\rangle_b$ begins as $\hat{H}^{\text{ad}}|\downarrow\rangle_b$; the full state is $|\psi_i\rangle = |g(\lambda_i)\rangle \otimes |\uparrow\rangle_a \otimes \hat{H}^{\text{ad}}|\downarrow\rangle_b$.
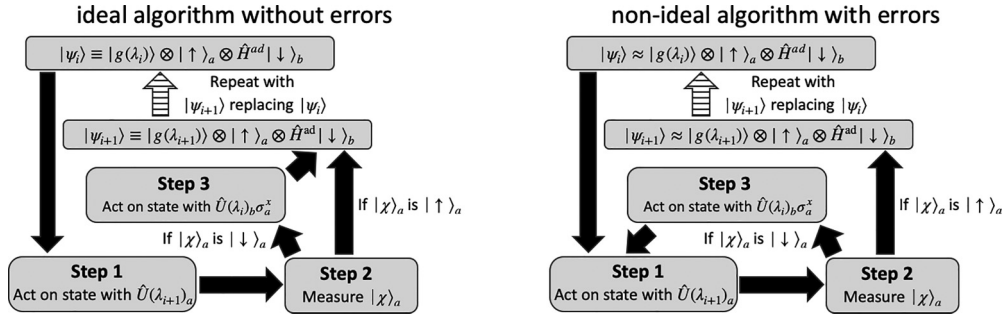
FIG. 1. A flowchart for one cycle in the idealized algorithm, which assumes no errors (left), and a realistic version that can handle errors (right).

A flowchart for the idealized algorithm is on the left in Fig. 1. The first step is to act on $|\psi_i\rangle$ with $\hat{U}(\lambda_{i+1})_a$. Next, $|\chi\rangle_a$ is measured. Half of the measurements will yield $|\uparrow\rangle_a$, indicating that the system is in the ground state of $\hat{H}(\lambda_{i+1})$; nothing more needs to be done: The system has advanced to the next point along the path. If the measurement yields $|\downarrow\rangle_a$, the system is in an excited state of $\hat{H}(\lambda_{i+1})$; this excited state is in the two-dimensional vector space spanned by $|g(\lambda_i)\rangle$ and $|g(\lambda_{i+1})\rangle$; the state is $\sqrt{2}|g(\lambda_i)\rangle - |g(\lambda_{i+1})\rangle$, as seen Fig. 2. When $|\chi\rangle_a$ is measured to be $|\downarrow\rangle_a$, the third step is implemented: $\hat{U}(\lambda_i)_b\sigma_a^x$ acts on the state. As seen in Fig. 2, this reflection yields $|g(\lambda_{i+1})\rangle$. Regardless of whether $|\chi\rangle_a$ is measured to be $|\uparrow\rangle_a$ or $|\downarrow\rangle_a$, after the cycle the system is in the ground state of $\hat{H}(\lambda_{i+1})$, having called the oracle at most two times. In this idealized scenario, a path of length L would be traversed by calling the oracle $3L/\sqrt{4\log(2)}$ times (up to a small statistical correction) when $L \gg 1$. If each call of the oracle took a fixed time, then the total time would scale linearly with L, and the cost of state preparation scales as $\sqrt{V}$.

This ideal case is unrealistic: There will always be errors arising from imperfect information, imperfect devices, or the inability to construct a perfect oracle that runs in finite time (even with perfect devices and perfect information). There have been previous schemes [40–42] proposed for the traversal of long paths that mitigate the effects of such imperfections by increasing the resources at each stage so that the errors could be kept below the final desired value at all points along the path; such schemes lead to additional costs and they scaled asymptotically with L no better than $L \log L$. A key insight in
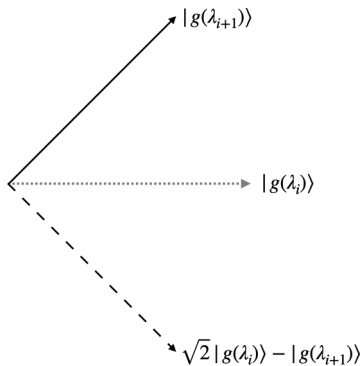

FIG. 2. The two-dimensional space spanned by $|g(\lambda_i)\rangle$ and $|g(\lambda_{i+1})\rangle$.

this work is that it is sufficient that there is no secular growth in the error with increasing L; at the end of the path, all one needs is admixtures of excited states that are either strictly bounded independent of L or given in terms of a sufficiently narrow distribution whose width does not grow with L. If so, one can then use an "afterburner" to produce the ground state with whatever accuracy is needed: Vacuum states are calculable with a cost that scales as $V^{1/2}$ without a logarithmic correction.

The flow chart on the right in Fig. 1 is a variant of the algorithm that is useful for nonideal conditions and allows long paths to be traversed without a secular growth of errors. The scheme makes a single change from the ideal case: After step three, the algorithm returns to step one rather than completing the cycle. If everything were ideal, this change would be redundant: The system would be in the ground state of λ_{i+1} and steps one and two would simply go on to the next cycle. However, under nonideal conditions it suppresses the secular growth of errors with only a modest multiplicative increase in cost relative to the ideal case.

There are several causes of errors. A key one is that one cannot project perfectly in finite time. Imperfect projections lead to "false positives" where an excited state component of a wave function is identified as the ground state. A second factor is the inability to prepare points that satisfy Eq. (5) exactly. Additionally, there can be "false negatives" in projection: Uncertainty in knowledge of the ground-state energy can lead to the projection indicating that the system is not in the ground state even when it is.

One might anticipate that these errors could accumulate along the path, causing the amplitudes of ground states to diminish along the path. However, a simple heuristic argument suggests that the error does not grow in a secular manner as the path is traversed.

Denote the maximum error introduced in any stage of the algorithm as ϵ; it represents the maximum amplitude introduced for non-ground-state components at any stage of the algorithm (assuming it had started in the ground state). Given that each stage concludes with the oracle $\hat{U}(\lambda)_a$, the error introduced in the nth iteration will be mitigated by the $(n + 1)$th iteration. Of course, the $(n + 1)$th iteration will also introduce an error (which will be less than or equal to ϵ). The same situation happens for the next iterations and consequently the error converges with a geometric series: The error will not grow secularly.
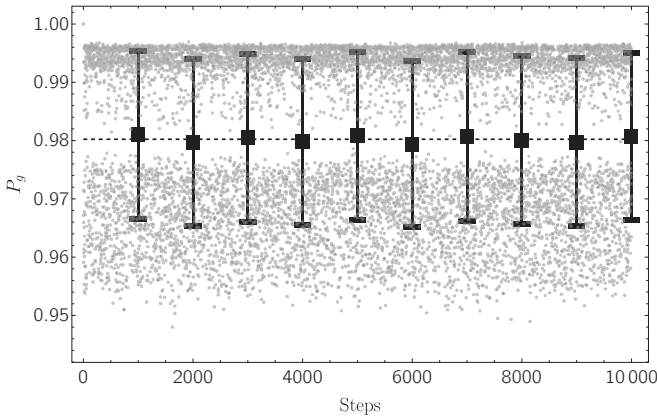
FIG. 3. A numerical simulation of the algorithm given for the algorithm in Fig. 1 using the toy model discussed in the text. P_g represents the ground-state probability (the square of the inner product of the true intermediate ground state with the one produced by the algorithm). The dots represent the ground-state probabilities at each step, while the rectangles represent averages taken over every 1000 steps (with the standard deviations indicated by the error bars). The dashed line is the average of all steps. Note that the scale of P_g is highly compressed, ranging from 0.94 to 1.00.

To test this heuristic reasoning, we implemented the algorithm numerically in a "toy model." In constructing this model, we exploited the fact that what matters is the ground state itself (along with excited states) rather than the Hamiltonian that leads to the states. Therefore, for simplicity, the model consists of "ground" and "excited" states constructed (essentially randomly) directly as linear combinations of an orthogonal basis of 1024 levels (this number was chosen to be large enough to represent a generic system but small enough to keep the calculation tractable). The inner product between the ith and $(i+1)$th ground state was chosen to be close to $1/\sqrt{2}$ for each step. To simulate false positive errors we assumed that the projection was allowed through excited states with a probability of 0.1 or less (randomly chosen with a uniform distribution); this is extremely conservative (a single "superiteration" of the rodeo projection has a maximum rate of less than 0.05 and is typically much less [52]). We also chose a 0.05 probability of a false negative (which is also quite conservative if one has reasonable knowledge of the ground-state energies) and assumed that $\cos^{-1}$ of the inner product between the ith and $(i+1)$th ground states was within 10% of $\frac{\pi}{4}$ (which again should be quite achievable).

The results of a simulation using this toy model are shown in Fig. 3: In that figure, P_g, the ground-state probabilities (the square of the intermediate ground-state component of the state produced at a given step along the path) are nearly unity for all steps (with a minimum P_g of greater than 0.947). Although errors fluctuate as additional steps along the path are taken, the errors do not grow in a secular manner: The characteristic distribution of P_g for the steps 9000–10 000 is qualitatively the same as steps 1–1000, with both the average and standard deviations nearly identical. This numerical check strongly suggests that the conclusion of the heuristic argument is sound.

There is strong evidence that the errors do not grow as the path is traversed. Assuming that this is true, the inner product

of the state produced at the end of the long path, and the true ground state will be close to unity. If one wishes the final state to be extremely close to the true ground state, one can use an "afterburner." It involves the last step of the traversal—which is likely not very close to satisfying Eq. (5)—and thus may require several rounds of projection and reflection but is guaranteed to converge. It simply uses a projector with a false positive rate which is exponentially smaller than the one used to traverse the path. This ensures that when the last step converges one has the true ground state with exponential accuracy. As was shown in Ref. [52], the cost in time grows more slowly than logarithmically with the false positive rate when using the rodeo algorithm. Thus, for long paths the total cost is dominated by traversing the path and not the afterburner.

The total cost for long paths depends on the cost of each projection as well as the number of projections needed. Since typical projection methods require a time proportional to the inverse of the spectral gap, Δ, the total cost is given by $C \sim V^{\frac{1}{2}} \langle \frac{1}{\Delta} \rangle$, where $\langle \Delta^{-1} \rangle$ is the average of the inverse spectral gap along the path.

Properties of the path need to be determined in the initial traversals of the path so that subsequent traversals can be efficient: Reasonable knowledge of the values of $\Delta\lambda$ to reach the next point approximately satisfying Eq. (5), the ground-state energy (to suppress the false negative rate), and the spectral gap (to efficiently project) are needed at each point on the path. The algorithm outlined above is applicable to all problems with long path lengths. However, if determining the properties of the system along the path during initial studies is too expensive, the algorithm may be impractical.

Fortunately, for local field theories with large volumes relevant quantities change very slowly from λ_i to λ_{i+1}; one can extract the values at relatively few points and then smoothly interpolate. To see this, recall that $\Delta\lambda_i = \lambda_{i+1} - \lambda_i \propto V^{-\frac{1}{2}}$ from Eq. (5). Moreover, at a large volume, the spectral gap $\Delta(\lambda)$ depends only on λ and not the volume. Therefore, $\Delta\lambda_{i+N}$ can be approximated as $\Delta\lambda_i$ up to $N \propto V^{\frac{1}{2}}$. Similarly, Eqs. (4a)–(4c) imply that $\Delta(\lambda_{i+N})$ is well approximated as $\Delta(\lambda_i)$ until $N \propto V^{\frac{1}{2}}$. However, since the ground-state energy $E_0(\lambda)$ is proportional to the volume, $E_0(\lambda)$ must be approximated as a Taylor series; an nth-order Taylor series at λ_i provides an accurate estimate of $E_0(\lambda_{i+N})$ until $N \propto V^{\frac{1}{2}-\frac{1}{n}}$.

In this Letter, a highly efficient state preparation method on quantum computers was presented, particularly tailored for use with QFTs with large volumes. It was demonstrated that errors do not accumulate secularly. The results here are of interest from a theoretical quantum information perspective as they indicate that the cost of preparing the ground state of QFTs scales as the $V^{\frac{1}{2}}$. The approach may also prove to be valuable practically both in the context of QFTs and perhaps in other problems with long path lengths. However, practical implementations of the approach are unlikely in the short term; the approach relies on controlling the buildup of errors and is unlikely to be viable without the development of fault-tolerant quantum computers.

[1] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations, Phys. Rev. Lett. **94**, 170201 (2005).

[2] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Sign problem in the numerical simulation of many-electron systems, Phys. Rev. B **41**, 9301 (1990).

[3] T. Byrnes and Y. Yamamoto, Simulating lattice gauge theories on a quantum computer, Phys. Rev. A **73**, 022328 (2006).

[4] S. P. Jordan, K. S. M. Lee, and J. Preskill, Quantum algorithms for quantum field theories, Science **336**, 1130 (2012).

[5] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, Rev. Mod. Phys. **86**, 153 (2014).

[6] E. Zohar and M. Burrello, Formulation of lattice gauge theories for quantum simulations, Phys. Rev. D **91**, 054506 (2015).

[7] C. W. Bauer, Z. Davoudi, A. B. Balantekin, T. Bhattacharya, M. Carena, W. A. de Jong, P. Draper, A. El-Khadra, N. Gemelke, M. Hanada *et al.*, Quantum simulation for high-energy physics, PRX Quantum **4**, 027001 (2023).

[8] D. Beck, J. Carlson, Z. Davoudi, J. Formaggio, S. Quaglioni, M. Savage, J. Barata, T. Bhattacharya, M. Bishof, I. Cloet *et al.*, Quantum information science and technology for nuclear physics. Input into U.S. long-range planning, arXiv:2303.00113.

[9] C. W. Bauer, Z. Davoudi, N. Klco, and M. J. Savage, Quantum simulation of fundamental particles and forces, Nat. Rev. Phys. **5**, 420 (2023).

[10] L. Funcke, T. Hartung, K. Jansen, and S. Kühn, Review on quantum computing for lattice field theory, PoS Proc. Sci. **430**, 228 (2023).

[11] A. Di Meglio, K. Jansen, I. Tavernelli, C. Alexandrou, S. Arunachalam, C. W. Bauer, K. Borras, S. Carrazza, A. Crippa, V. Croft *et al.*, Quantum computing for high-energy physics: State of the art and challenges. Summary of the QC4HEP working group, arXiv:2307.03236.

[12] A. N. Ciavarella, S. Caspar, H. Singh, and M. J. Savage, Preparation for quantum simulation of the (1+1)-dimensional O(3) nonlinear σ model using cold atoms, Phys. Rev. A **107**, 042404 (2023).

[13] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Scalable circuits for preparing ground states on digital quantum computers: The Schwinger model vacuum on 100 qubits, arXiv:2308.04481.

[14] A. Roggero, C. Gu, A. Baroni, and T. Papenbrock, Preparation of excited states for nuclear dynamics on a quantum computer, Phys. Rev. C **102**, 064624 (2020).

[15] A. N. Ciavarella and I. A. Chernyshev, Preparation of the SU(3) lattice Yang-Mills vacuum with variational quantum methods, Phys. Rev. D **105**, 074504 (2022).

[16] N. Klco and M. J. Savage, Minimally entangled state preparation of localized wave functions on quantum computers, Phys. Rev. A **102**, 012612 (2020).

[17] A. H. Moosavian, J. R. Garrison, and S. P. Jordan, Site-by-site quantum state preparation algorithm for preparing vacua of fermionic lattice field theories, arXiv:1911.03505.

[18] A. J. Buser, T. Bhattacharya, L. Cincio, and R. Gupta, State preparation and measurement in a quantum simulation of the O(3) sigma model, Phys. Rev. D **102**, 114514 (2020).

[19] I. Stetcu, A. Baroni, and J. Carlson, Projection algorithm for state preparation on quantum computers, Phys. Rev. C **108**, L031306 (2023).

[20] K. Choi, D. Lee, J. Bonitati, Z. Qian, and J. Watkins, Rodeo algorithm for quantum computing, Phys. Rev. Lett. **127**, 040505 (2021).

[21] L. Lumia, P. Torta, G. B. Mbeng, G. E. Santoro, E. Ercolessi, M. Burrello, and M. M. Wauters, Two-dimensional $\mathbb{Z}_2$ lattice gauge theory on a near-term quantum simulator: Variational quantum optimization, confinement, and topological order, PRX Quantum **3**, 020320 (2022).

[22] A. Hamed Moosavian and S. Jordan, Faster quantum algorithm to simulate fermionic quantum field theory, Phys. Rev. A **98**, 012332 (2018).

[23] C. F. Kane, N. Gomes, and M. Kreshchuk, Nearly-optimal state preparation for quantum simulations of lattice gauge theories, arXiv:2310.13757.

[24] C. Ball and T. D. Cohen, Boltzmann distributions on a quantum computer via active cooling, Nucl. Phys. A **1038**, 122708 (2023).

[25] D. Lee, J. Bonitati, G. Given, C. Hicks, N. Li, B.-N. Lu, A. Rai, A. Sarkar, and J. Watkins, Projected cooling algorithm for quantum computation, Phys. Lett. B **807**, 135536 (2020).

[26] Z. Davoudi, N. Mueller, and C. Powers, Towards quantum computing phase diagrams of gauge theories with thermal pure quantum states, Phys. Rev. Lett. **131**, 081901 (2023).

[27] F. G. S. L. Brandão and M. J. Kastoryano, Finite correlation length implies efficient preparation of quantum thermal states, Commun. Math. Phys. **365**, 1 (2019).

[28] W. A. de Jong, K. Lee, J. Mulligan, M. Płoskoń, F. Ringer, and X. Yao, Quantum simulation of nonequilibrium dynamics and thermalization in the Schwinger model, Phys. Rev. D **106**, 054508 (2022).

[29] C.-F. Chen, M. J. Kastoryano, F. G. S. L. Brandão, and A. Gilyén, Quantum thermal state preparation, arXiv:2303.18224.

[30] Z. Holmes, G. Muraleedharan, R. D. Somma, Y. Subasi, and B. Şahinoğlu, Quantum algorithms from fluctuation theorems: Thermal-state preparation, Quantum **6**, 825 (2022).

[31] P. W. Shor, Polynomial time algorithms for prime factorization and discrete logarithms on a quantum computer, SIAM J. Sci. Stat. Comput. **26**, 1484 (1997).

[32] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96 (ACM Press, New York, 1996).

[33] T. Albash and D. A. Lidar, Adiabatic quantum computation, Rev. Mod. Phys. **90**, 015002 (2018).

[34] K. Wan and I. H. Kim, Fast digital methods for adiabatic state preparation, arXiv:2004.04164.

[35] E. A. Coello Pérez, J. Bonitati, D. Lee, S. Quaglioni, and K. A. Wendt, Quantum state preparation by adiabatic evolution with custom gates, Phys. Rev. A **105**, 032403 (2022).

[36] A. N. Ciavarella, S. Caspar, M. Illa, and M. J. Savage, State preparation in the Heisenberg model through adiabatic spiraling, Quantum **7**, 970 (2023).

[37] B. Chakraborty, M. Honda, T. Izubuchi, Y. Kikuchi, and A. Tomiya, Classically emulated digital quantum simulation of the Schwinger model with a topological term via adiabatic state preparation, Phys. Rev. D **105**, 094503 (2022).

[38] L. K. Kovalsky, F. A. Calderon-Vargas, M. D. Grace, A. B. Magann, J. B. Larsen, A. D. Baczewski, and M. Sarovar, Self-healing of Trotter error in digital adiabatic state preparation, Phys. Rev. Lett. **131**, 060602 (2023).

[39] D. Aharonov and A. Ta-Shma, Adiabatic quantum state generation, SIAM J. Comput. **37**, 47 (2007).

[40] S. Boixo, E. Knill, and R. D. Somma, Fast quantum algorithms for traversing paths of eigenstates, arXiv:1005.3034.

[41] S. Boixo, E. Knill, and R. Somma, Eigenpath traversal by phase randomization, Quantum Inf. Comput. **9**, 833 (2009).

[42] H.-T. Chiang, G. Xu, and R. D. Somma, Improved bounds for eigenpath traversal, Phys. Rev. A **89**, 012314 (2014).

[43] G. A. Paz-Silva, A. T. Rezakhani, J. M. Dominy, and D. A. Lidar, Zeno effect for quantum computation and control, Phys. Rev. Lett. **108**, 080501 (2012).

[44] A. Luis, Quantum-state preparation and control via the Zeno effect, Phys. Rev. A **63**, 052112 (2001).

[45] S. Boixo and R. D. Somma, Necessary condition for the quantum adiabatic approximation, Phys. Rev. A **81**, 032308 (2010).

[46] N. Wiebe and N. S. Babcock, Improved error-scaling for adiabatic quantum evolutions, New J. Phys. **14**, 013024 (2012).

[47] S. Jansen, M.-B. Ruskai, and R. Seiler, Bounds for the adiabatic approximation with applications to quantum computation, J. Math. Phys. **48**, 102111 (2007).

[48] J. E. Avron and A. Elgart, Adiabatic theorem without a gap condition, Commun. Math. Phys. **203**, 445 (1999).

[49] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, arXiv:quant-ph/0001106.

[50] A straightforward calculation shows that for long paths, a QZE scheme with perfect projections and a negligible cost to (re)create the initial state yields an expected number of measurements, N_p, equal to $N_p = L^2$. The energy-time uncertainty principle implies that the time to project the state of interest is of order $1/\Delta$, where Δ is the spectral gap; Eq. (2) follows.

[51] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, arXiv:quant-ph/9511026.

[52] T. D. Cohen and H. Oh, Optimizing the rodeo projection algorithm, Phys. Rev. A **108**, 032422 (2023).

[53] M. Bee-Lindgren, Z. Qian, M. DeCross, N. C. Brown, C. N. Gilbreth, J. Watkins, X. Zhang, and D. Lee, Rodeo algorithm with controlled reversal gates, arXiv:2208.13557.