


Mitigating the Sign Problem through Basis Rotations

Ryan Levy^{*} and Bryan K. Clark[†]

*Institute for Condensed Matter Theory and IQIUST and Department of Physics,
University of Illinois at Urbana-Champaign, Champaign, Illinois 61801, USA*

 (Received 5 July 2019; revised 24 June 2020; accepted 15 April 2021; published 24 May 2021)

Quantum Monte Carlo simulations of quantum many-body systems are plagued by the Fermion sign problem. The computational complexity of simulating Fermions scales exponentially in the projection time β and system size. The sign problem is basis dependent and an improved basis, for fixed errors, leads to exponentially quicker simulations. We show how to use sign-free quantum Monte Carlo simulations to optimize over the choice of basis on large two-dimensional systems. We numerically illustrate these techniques decreasing the “badness” of the sign problem by optimizing over single-particle basis rotations on one- and two-dimensional Hubbard systems. We find a generic rotation which improves the average sign of the Hubbard model for a wide range of U and densities for $L \times 4$ systems. In one example improvement, the average sign (and hence simulation cost at fixed accuracy) for the 16×4 Hubbard model at $U/t = 4$ and $n = 0.75$ increases by $\exp[8.64(6)\beta]$. For typical projection times of $\beta \gtrsim 100$, this accelerates such simulation by many orders of magnitude.

DOI: 10.1103/PhysRevLett.126.216401

The power of quantum computers and the difficulty of simulating quantum many-body ground states stem from similar sources: nontrivial entanglement and paths of interfering signs. Low entanglement states can be solved using the density matrix renormalization group [1,2]. Paths of interfering signs are the source of the fermion sign problem in quantum Monte Carlo (QMC) simulations. Systems, such as bosonic Hamiltonians, which have only positive paths and hence no sign problem, can be simulated efficiently. For some classes of systems with “naive” sign problems, approaches have been found that also allow for efficient simulations. One common technique involves finding a sign-problem-free basis in which to perform the simulation [3–7]. Examples where this has also been successful include Hubbard models at half filling [8,9] and employing a Majorana representation [10]. Note that, while there is always a sign-problem-free basis (i.e., the eigenstate basis), it is often as (or more) difficult to find this basis as doing the simulation. Even in cases where the sign problem cannot be removed, the choice of basis can affect the badness of the sign problem [11]. Interestingly, there has not yet been a significant amount of work in automatically searching for a good basis to help mitigate or remove the sign problem. It is this problem that we approach in this Letter.

Our approach is as follows: given a quantification of the sign problem, search over a class of unitary rotations of the Hamiltonian, minimizing the badness of the sign problem. The key question, then, is the development of an algorithm to efficiently accomplish this. As a proof of principle, we will consider hole-doped Hubbard models, finding a rotation that decreases the badness of the sign problem for $L \times 4$ lattices over a wide range of U/t .

Quantification of sign problem.—Let H be the Hamiltonian represented in a standard basis (i.e., real space) and $H(R) \equiv RHR^\dagger$ be the Hamiltonian rotated by the unitary rotation R . Our choice for quantifying the sign problem of $H(R)$ has two goals: (1) The objective function should be proportional to the cost of a QMC simulation on $H(R)$. (2) We can measure (and hence optimize) this objective function efficiently on large bulk two-dimensional systems.

The natural objective function to maximize is the relative variance of the average sign

$$\langle s \rangle \equiv \frac{1}{N} \sum_i s_i, \quad (1)$$

where s_i is the sign for each Monte Carlo sample i and N is the total number of Monte Carlo samples. While the efficiency of different observables vary, their values $\langle O \rangle = \sum_i s_i O_i / \langle s \rangle$ all involve the average sign. Here O_i is the observable measured on Monte Carlo sample i . Given the variance of $\langle s \rangle$ is $O(1)$, minimizing the relative variance comes down to maximizing $\langle s \rangle$. Because our interest will be in the limit of a large QMC projection time β , our focus will be in optimizing the component of $\langle s \rangle = A\bar{s}^\beta$ which causes $\langle s \rangle$ to decay exponentially with β , i.e., \bar{s} .

Naively, though, this appears to be difficult to optimize, as computing $\langle s \rangle$ itself has a sign problem. It turns out, though, that locally optimizing $\langle s(R) \rangle$ over a class of unitary rotations R is significantly easier than computing it. In particular, $\overline{s(R)} = \exp[-\Delta E(R)]$ [11–13], where $\Delta E(R) \equiv E(R) - \overline{E(R)}$ is the gap between the ground state $E(R)$ of the fermionic Hamiltonian $H(R)$ and the

ground state $\overline{E(R)}$ of an effective bosonic version of the Hamiltonian $|H(R)\rangle$ [14–16], defined as

$$\begin{aligned} |H(R)\rangle_{ij} &= -H(R)_{ij} \quad \text{if } i \neq j \quad \text{and} \quad H(R)_{ij} > 0 \\ &= H(R)_{ij} \quad \text{otherwise,} \end{aligned} \quad (2)$$

with $\overline{E(R)} = \langle \bar{\Psi} | \overline{H(R)} | \bar{\Psi} \rangle$, where $|\bar{\Psi}\rangle$ is the ground state of $\overline{H(R)}$. Note that this definition ensures that the propagator used in a QMC simulation, $|G\rangle \equiv 1 - \tau|H\rangle$ (where τ is a small constant), contains strictly positive matrix elements. While $E(R)$ is hard to compute (there is a sign problem in evaluating it), its value is independent of the choice of the unitary rotation, i.e., $E(R) = E$. On the other hand $\overline{E(R)}$ is sign-free and hence computable by Monte Carlo simulations on large systems. Maximizing $\overline{s(R)}$ then comes down to maximizing $\overline{E(R)}$.

Note this metric is closely related to, but not identical to having a ground state with only positive amplitudes or a Hamiltonian with no positive off-diagonal matrix elements. For example, the Heisenberg model on a bipartite lattice has no sign problem (i.e., no exponentially large relative variance) in a QMC simulation, but does have positive off-diagonal Hamiltonian matrix elements and a nonpositive ground state (although both these can be removed by applying the Marshall sign rule [17]). $\Delta E(I)$ where I is the identity rotation is correctly zero in this case.

Unitary rotation.—Here we describe the parametrization of our unitaries. While a unitary rotation exists that removes the sign problem (i.e., the eigenstate basis), it is often harder to find than the ground state itself. Instead, we focus on unitary rotations that have simple representations. Two such classes are unitary rotations from shallow-depth quantum circuits and basis rotations on the set of single-particle orbitals used to represent the Hamiltonian: i.e., given a Hamiltonian written as $H = \sum_{ij} t_{ij} c_i^\dagger c_j + \sum_{ijkl} V_{ijkl} c_i^\dagger c_j^\dagger c_k c_l$, we can write $b_j \equiv \sum_k U_{jk} c_k$, where U is an $N \times N$ unitary matrix for a system of N sites. U can be parametrized in various ways, including as the eigenvalues of an orthogonal matrix H or as $U = e^A$, where A is a skew-symmetric matrix, i.e., $A = -A^T$. In our numerical examples, we focus on single-particle basis rotations, although the approach we describe works also with quantum circuits.

Optimization.—Our goal now is to maximize $\overline{E(R)}$. In this section we demonstrate a series of techniques to optimize this quantity that we benchmark on single-particle basis rotations of the Hubbard model under periodic boundary conditions,

$$H = -t \sum_{\sigma(i,j)} c_{j\sigma}^\dagger c_{i\sigma} + \text{H.c.} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \quad (3)$$

with $\langle i, j \rangle$ denoting nearest neighbors.

We optimize $\overline{E(R)}$ via three separate techniques. First, we optimize $\overline{E(R)}$ via exact diagonalization using finite

differences. When diagonalization is too costly, we turn to the second technique, the projector quantum Monte Carlo (PQMC) method [18,19]. The PQMC method is a form of QMC simulation which propagates walkers for a projection time β giving the energy

$$E(\bar{R}; \beta) = \frac{\langle \Psi_{\text{init}} | \overline{H(R)} \exp[-\beta \overline{H(R)}] | \Psi_{\text{init}} \rangle}{\langle \Psi_{\text{init}} | \exp[-\beta \overline{H(R)}] | \Psi_{\text{init}} \rangle}, \quad (4)$$

where walkers are sampled from an initial wave function $|\Psi_{\text{init}}\rangle$. In a PQMC simulation, by taking large β , we can compute $\overline{E(R)} \equiv \lim_{\beta \rightarrow \infty} E(\bar{R}; \beta)$ in polynomial time; derivatives can be computed using finite differences, giving an algorithm that scales linearly in the number of parameters and otherwise is similar in cost to a sign-free QMC simulation of the same rotated Hamiltonian (i.e., polynomial in system size).

Parameters are then updated using the optimization scheme described in [20] inspired by [21]. For each of the unique entries v_i of our unitary parametrization, the next parameter is determined by

$$v_{i+1} = v_i + \alpha \gamma \frac{|\partial \overline{E} / \partial v_i|}{\partial \overline{E} / \partial v_i}, \quad (5)$$

where α is a random number chosen between zero and one, and γ controls the size of the random step.

Although the cost of the direct PQMC approach is polynomial, it can still be expensive, especially since this rotation can result in a quadratic number of terms per (matrix) row (which is a significant increase from the linear number of terms in the Hubbard model but standard for QMC simulations on materials [22]). To overcome this expense, we develop and benchmark a further approximation to our algorithm that significantly accelerates the optimization. Instead of directly optimizing $\overline{E(R)}$, we optimize

$$\overline{E_V(R)} \equiv \langle \Psi_{u(R)} | \overline{H(R)} | \Psi_{u(R)} \rangle, \quad (6)$$

where $\Psi_{u(R)} \equiv 1/2^N$ is the uniform wave function in the basis R . This is a strict upper bound in energy to $\overline{H(R)}$ and is equivalent to $\overline{E(R; \beta = 0)}$ in the PQMC simulation when the initial wave function is chosen as $|\Psi_{u(R)}\rangle$. This approximation turns out to be reasonable as the variational energy $\overline{E_V(R)}$ tracks the ground-state energy $\overline{E(R)}$ and so pushes the parameters in the correct direction (see Fig. 2). This approximation can be made even faster by choosing a relatively fixed few random states [$O(10)$] and using these states to optimize $\overline{E_V(R)}$. We further utilize the Jax library [23] and implement the derivative $d\overline{E_V(R)}/dR$ using automatic differentiation. Because these derivatives (and the energy) are computed from only $O(10)$ Monte Carlo

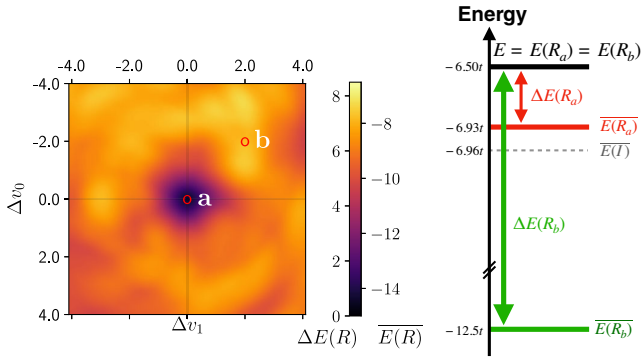



FIG. 1. Left: $\Delta E(R)$ and $\overline{E(R)}$ for an eight-site Hubbard chain at $n = 0.5$ and $U/t = 1$ in the plane of the first two parameters; the final plot is interpolated. Dark (light) colors indicate mitigation (increased) sign problem, respectively. Color bar shows both ΔE (which measures the badness of the sign problem but cannot be computed efficiently) as well as $\overline{E(R)}$, which can be efficiently computed for large systems. Right: depiction (not to scale) of the energies of two particular rotations R_a and R_b . Despite having the same energy E when there is a sign problem, by optimizing for the largest sign-free $\overline{E(R)}$, we minimize the gap $\Delta E(R)$, which exponentially improves the average sign $\langle s \rangle$ by $\exp[\beta(E(R) - E(T))] \approx \exp[0.03\beta]$ for projection time β .

samples, the statistical error is large, but the correlations introduced by using the same configurations better controls the path generated by the gradients.

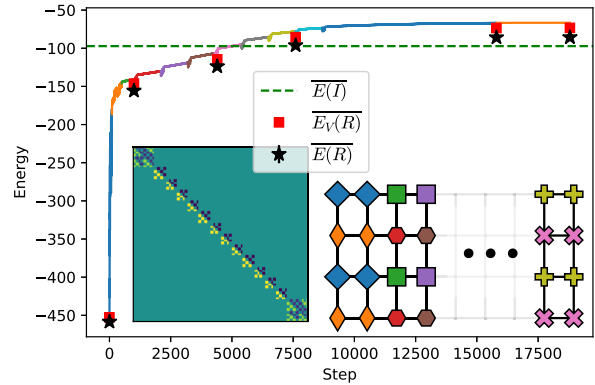
Results.—1D: We perform initial tests optimizing one-dimensional Hubbard chains for $N = \{3, 4, 8\}$. We evaluate $\Delta E(R)$ via exact diagonalization and perform derivative-free optimization using Nelder-Mead optimization [24,25]. In these runs, R is parametrized by the eigenvectors of a real Hermitian matrix. Our optimization completely removes the sign problem at $N = 3$ at $2/3$ doping and $N = 4$ Hubbard model at half filling for all values of U (see Supplemental Material [26] for the specific matrices); *a posteriori*, we discovered these rotations are known in the continuous-time QMC literature [29]. Interestingly at $N = 4$ (see Supplemental Material [26]), our optimization procedure finds a basis that does not make all the off-diagonal signs of $|G|$ positive despite still removing the sign problem (although by hand G can be made positive). We also include an example for $N = 8$, where we improved but were unable to completely remove the sign problem (see Fig. 1).

2D: Next we consider a variety of width-four Hubbard models. First we implement projector optimization for the $n = 0.5$, $U/t = 1$, 4×4 Hubbard model. We parametrize R as the matrix exponential of a real skew-symmetric matrix and start from a random unitary that has a bad sign problem (worse than not rotating at all).

Figure 2(a) shows the improvement in $\overline{E(R)}$ during optimization. By rounding the final unitary and removing noise, we obtain a sparse unitary rotation that rotates each site operator into linear combinations of four-operator plaquettes [see inset of Fig. 2(a)]. The rotated operators



(a)



(b)

FIG. 2. (a) Projector optimization of 4×4 Hubbard model with $\gamma = 0.01$ showing $\overline{E(R)}$ (blue) and $\overline{E_V(R)}$ (red). (b) Optimization of a 16×4 Hubbard model showing both noisy values of $\overline{E_V(R)}$ used in optimization (colored lines), as well as accurate values of $\overline{E_V(R)}$ (red squares) and $\overline{E(R)}$ (stars). Each color indicates a change in hyperparameters of optimization (see Supplemental Material [26]). Left insets: rounded optimal unitary rotations at the end of optimization, in a matrix representation. Right insets: lattices illustrating the new basis. The rotated basis is made up of linear combinations of two or four operators, marked by their same color and shape.

do not map to their counterpart in the original basis; i.e., operator $b_{(i,j)}$ does not contain $c_{(i,j)}$. We find that $\Delta E(R) = 0.137$ and achieve an increase in the average sign, which goes as $\exp[6.27\beta]$ at large β .

We test our optimized rotation for a range of doping n and interaction strengths U in Fig. 3. While only optimized for a single U and n , we find this rotation mitigates the sign problem over the real-space basis for the 4×4 Hubbard model for all $0.125 < n < 1.0$ when $U/t \leq 4$. The ability for a single optimized rotation to be applicable to a larger regime of phase space shows the versatility of our approach.

We then turn to width-four cylinders with periodic boundary conditions. We perform optimization on the 16×4 Hubbard model with $n = 0.5$ and $U/t = 2$, using automatic differentiation of $\overline{E_V(R)}$ and varying step sizes




FIG. 3. Comparing the results of projector optimization of a 4×4 Hubbard model at $U/t = 1$ and $n = 0.5$ doping (star) to no basis rotation (I) at different U and doping values. PQMC simulations are done in continuous time with a resampling rate of $T = 0.1$ to a projection time of $\beta = 20$. Inset: relative $\Delta E(R)/\Delta E(I)$ between the optimal rotation and no basis rotation ($R = I$).

(see Supplemental Material [26]), shown in Fig. 2(b). By rounding and removing noise of the final unitary we obtain a sparse rotation that improves the sign by $\exp[12.86(5)\beta]$. We observe this rotation also works for a variety of doping and U values; e.g., $n = 0.5$ $U/t = 6$ and $n = 0.75$ $U/t = 4$ have improvements of the average sign by $\exp[3.78(6)\beta]$ and $\exp[8.64(6)\beta]$, respectively.

Upon analyzing the unitary, we find a clear structure that can be used for any width-four cylinder. The single-particle orbitals in the bulk of the lattice become a linear combination over each column of width four, leading to four creation (annihilation) operators (with implicit spin indices)

$$\begin{pmatrix} b_{(n,4)} \\ b_{(n,3)} \\ b_{(n,2)} \\ b_{(n,1)} \end{pmatrix}^{(\dagger)} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & -1 & 0 & -1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ 1 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} c_{(n,4)} \\ c_{(n,3)} \\ c_{(n,2)} \\ c_{(n,1)} \end{pmatrix}^{(\dagger)}. \quad (7)$$

At the left and right edges of the system, despite periodic boundaries, two consecutive columns are rotated as

$$\begin{pmatrix} b_{(2,4)} \\ \vdots \\ b_{(2,1)} \\ b_{(1,4)} \\ \vdots \\ b_{(1,1)} \end{pmatrix}^{(\dagger)} = \frac{1}{2} \begin{pmatrix} 0 & -1 & 0 & -1 & 0 & 1 & 0 & 1 \\ 1 & 0 & -1 & 0 & -1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 & 0 & -1 & 0 & 1 \\ 1 & 0 & 1 & 0 & -1 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 & 0 & -1 & 0 & -1 \\ 1 & 0 & -1 & 0 & 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 & 0 & 1 & 0 & -1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} c_{(2,4)} \\ \vdots \\ c_{(2,1)} \\ c_{(1,4)} \\ \vdots \\ c_{(1,1)} \end{pmatrix}^{(\dagger)}. \quad (8)$$

We illustrate this rotation in the inset of Fig. 2(b). Similar to the structure of the 4×4 rotation, the new operators are linear combinations of nearby sites but lack the original site with respect to the old basis. However, the bulk is made of two-site pairs of sites rather than four-site plaquettes, and the edge plaquettes are not spaced apart in the \hat{x} direction.

Testing this for a larger system, we observe an average sign reduction over not rotating on a 32×4 cylinder at $n = 0.5$ doping of $\exp[22.5(4)\beta]$ and $\exp[16.0(3)\beta]$ for $U/t = 4$ and $U/t = 6$, respectively. These energy differences $\overline{E(I)} - \overline{E(R)}$ per site at $U = 4$ are increasing with respect to system size (from $L = 8$ to $L = 32$), suggesting in the thermodynamic limit that our rotation provides a better lower bound to the true fermionic energy than identity rotation (details can be found in the Supplemental Material [26]).

Finally, we can take the edge rotation found in Eq. (8) and apply it both to the 4×4 and 8×8 lattices. For the 4×4 case, we find that the rotation improves the sign more not than rotating, but not as well as our optimized 4×4 unitary, i.e., $\exp[5.40\beta]$ compared to $\exp[6.27\beta]$ for $n = 1$ $U/t = 1$, respectively. When using this rotation to tile an 8×8 unitary, we find that we improve the sign only for $U/t \lesssim 1$; e.g., at $U/t = 1$ we find an improvement of $\exp[-3.0(1)\beta]$ and $\exp[-3.9(1)\beta]$ for $n = 0.5$ and $n = 0.875$, respectively.

In addition to the above results, we include in the Supplemental Material [26] an optimization of a 6×6 lattice to show that the original PQMC optimization of $\overline{E(R)}$ scales beyond what is accessible by exact diagonalization. We also include a study on 4×4 systems of how the rotation affects the average sign in a PQMC simulation with respect to β and the benefit of introducing annihilation [22,30] on top of the rotation.

Conclusion.—We have introduced a new approach for finding a basis that reduces the sign problem in bulk systems of arbitrary dimensions. For fixed accuracy and projection time, this improves the calculation speed of a ground-state QMC simulation by many orders of magnitude; a similar approach will also improve finite-temperature calculations. Our method uses a QMC method to maximize the energy of the “bosonic” version of the rotated Hamiltonian $|H(R)\rangle$; this maximization corresponds to maximizing the average sign of a QMC calculation in that rotated basis. In many ways, our sign problem mitigation approach is similar in flavor to variational Monte Carlo method [31–37] and many of the tricks from that community will be applicable here. Although we used the PQMC approach in this Letter, any QMC method that works in an orthogonal basis (world line Monte Carlo method, path integral ground state, etc.) have the same sign problem and are equally improved by the rotation. In addition, we developed a variational optimization approach that instead optimizes the variational energy of the uniform state $|\Psi_{u(R)}\rangle$ with respect to $|H(R)\rangle$ over a small number of

configurations, for a significant speed-up and at polynomial scaling. This heuristic works because of the empirical observation that this variational energy and the projector energy track each other closely up to some (slowly varying) fixed offset.

To illustrate the method, we optimized the single-particle basis of 4×4 and 16×4 Hubbard model at $n = 0.5$ doping and $U/t = 1$ and $U/t = 2$, respectively. This resulted in an exponential increase in average sign, by approximately $\exp[6.278\beta]$ and $\exp[12.86(5)\beta]$, leading to significant speed-ups for a typical $\beta \in [10, 100]$. The rotations we found are sparse, generalize to a wider range of U and n , and the latter rotation generalizes to all width-four cylinders.

A uniting theme among these optimized rotations is their structure of rotating into linear combinations of a small number of nearby sites. The new basis in both cases seem relatively insensitive to both doping and U/t , with each $b_{(i,j)}$ not including the matching $c_{(i,j)}$. In addition, a four-operator plaquette is apparent in both geometries, although in slightly different forms. Other techniques beyond QMC simulations may also benefit from this insight, although its unclear if it extends beyond the Hubbard model itself. The automatic optimization of our method should be useful for other sign-plagued models to potentially turn up other sparse bases.

We acknowledge Luiz Santos for helping to initiate this work, trying to find a sign-problem mitigating unitary, early collaboration on this project, and useful discussions and insights related to this work. B.K.C. thanks David Ceperley, as well as Shiwei Zhang, Miles Stoudenmire, and Giuseppe Carleo for conversations. R.L. thanks Shivesh Pathak for useful conversations and Yubo Yang for visualization assistance. We acknowledge support from the Department of Energy (DOE) Award No. DE-SC0020165. This project is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (Grants No. OCI-0725070 and No. ACI-1238993) and the State of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. This work also made use of the Illinois Campus Cluster, a computing resource that is operated by the Illinois Campus Cluster Program (ICCP) in conjunction with the National Center for Supercomputing Applications (NCSA) and which is supported by funds from the University of Illinois at Urbana-Champaign.

Note added.—Recently, a number of related works appeared on the arXiv [38–40].

*rlevy3@illinois.edu

†bkclark@illinois.edu

- [1] U. Schollwöck, *Ann. Phys. (Amsterdam)* **326**, 96 (2011).
- [2] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [3] T. Nakamura, *Phys. Rev. B* **57**, R3197 (1998).
- [4] K. Okunishi and K. Harada, *Phys. Rev. B* **89**, 134422 (2014).
- [5] A. Honecker, S. Wessel, R. Kerkdyk, T. Pruschke, F. Mila, and B. Normand, *Phys. Rev. B* **93**, 054408 (2016).
- [6] F. Alet, K. Damle, and S. Pujari, *Phys. Rev. Lett.* **117**, 197203 (2016).
- [7] S. Wessel, I. Niesen, J. Stapmanns, B. Normand, F. Mila, P. Corboz, and A. Honecker, *Phys. Rev. B* **98**, 174432 (2018).
- [8] L. Wang, Y.-H. Liu, M. Iazzi, M. Troyer, and G. Harcos, *Phys. Rev. Lett.* **115**, 250601 (2015).
- [9] Z.-X. Li and H. Yao, *Annu. Rev. Condens. Matter Phys.* **10**, 337 (2019).
- [10] Z.-X. Li, Y.-F. Jiang, and H. Yao, *Phys. Rev. B* **91**, 241117 (R) (2015).
- [11] M. H. Kolodrubetz, J. S. Spencer, B. K. Clark, and W. M. C. Foulkes, *J. Chem. Phys.* **138**, 024110 (2013).
- [12] J. Spencer, N. Blunt, and W. Foulkes, *J. Chem. Phys.* **136**, 054110 (2012).
- [13] D. Ceperley, in *Recent Progress in Many-Body Theories* (Springer, New York, 1981), pp. 262–269.
- [14] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, *Phys. Rev. B* **41**, 9301 (1990).
- [15] N. Hatano and M. Suzuki, *Phys. Lett.* **163A**, 246 (1992).
- [16] M. Troyer and U.-J. Wiese, *Phys. Rev. Lett.* **94**, 170201 (2005).
- [17] W. Marshall and R. E. Peierls, *Proc. R. Soc. A* **232**, 48 (1955).
- [18] N. Trivedi and D. M. Ceperley, *Phys. Rev. B* **41**, 4552 (1990).
- [19] D. M. Ceperley and B. J. Alder, *Phys. Rev. Lett.* **45**, 566 (1980).
- [20] D. Luo and B. K. Clark, *Phys. Rev. Lett.* **122**, 226401 (2019).
- [21] J. Lou and A. W. Sandvik, *Phys. Rev. B* **76**, 104432 (2007).
- [22] G. H. Booth, A. J. W. Thom, and A. Alavi, *J. Chem. Phys.* **131**, 054106 (2009).
- [23] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of PYTHON+NUMPY programs (2018), <http://github.com/google/jax>.
- [24] S. G. Johnson, The NLOPT nonlinear-optimization package, <https://github.com/stevengj/nlopt>.
- [25] J. A. Nelder and R. Mead, *Computer Journal (UK)* **7**, 308 (1965).
- [26] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.126.216401> for other explicit rotations mentioned in the text, further details of the 4×4 , 6×6 , width-4 Hubbard models, which includes Refs. [27,28].
- [27] E. Dagotto, A. Moreo, F. Ortolani, D. Poilblanc, and J. Riera, *Phys. Rev. B* **45**, 10741 (1992).
- [28] J. P. F. LeBlanc *et al.*, *Phys. Rev. X* **5**, 041041 (2015).
- [29] H. Shinaoka, Y. Nomura, S. Biermann, M. Troyer, and P. Werner, *Phys. Rev. B* **92**, 195126 (2015).
- [30] M. H. Kalos and F. Pederiva, *Phys. Rev. Lett.* **85**, 3547 (2000).

- [31] C. J. Umrigar, J. Toulouse, C. Filippi, S. Sorella, and R. G. Hennig, *Phys. Rev. Lett.* **98**, 110201 (2007).
- [32] E. Neuscamman, C. J. Umrigar, and Garnet Kin-Lic Chan, *Phys. Rev. B* **85**, 045103 (2012).
- [33] D. Kochkov and B. K. Clark, [arXiv:1811.12423](https://arxiv.org/abs/1811.12423).
- [34] L. Otis and E. Neuscamman, *Phys. Chem. Chem. Phys.* **21**, 14491 (2019).
- [35] M. P. Zaletel and F. Pollmann, *Phys. Rev. Lett.* **124**, 037201 (2020).
- [36] R. Haghshenas, M. J. O'Rourke, and G. K. Chan, *Phys. Rev. B* **100**, 054404 (2019).
- [37] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [38] G. Torlai, J. Carrasquilla, M. T. Fishman, R. G. Melko, and M. P. A. Fisher, *Phys. Rev. Research* **2**, 032060(R) (2020).
- [39] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, *Sci. Adv.* **6**, eabb8341 (2020).
- [40] J. Klassen, M. Marvian, S. Piddock, M. Ioannou, I. Hen, and B. Terhal, *SIAM J. Comput.* **49**, 1332 (2020).