Optimizing large parameter sets in variational quantum Monte Carlo

Eric Neuscamman, C. J. Umrigar, and Garnet Kin-Lic Chan

Phys. Rev. B **85**, 045103 — Published 5 January 2012

DOI: 10.1103/PhysRevB.85.045103

Optimizing large parameter sets in variational quantum Monte Carlo

Eric Neuscamman*, C. J. Umrigar†, and Garnet Kin-Lic Chan‡
*Department of Chemistry, University of California, Berkeley, California 94720
†Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853
‡Department of Chemistry and Chemical Biology, Cornell University, Ithaca, New York 14853
(Dated: November 21, 2011)

We present a technique for optimizing hundreds of thousands of variational parameters in variational quantum Monte Carlo. By introducing iterative Krylov subspace solvers and by multiplying by the Hamiltonian and overlap matrices as they are sampled, we remove the need to construct and store these matrices and thus bypass the most expensive steps of the stochastic reconfiguration and linear method optimization techniques. We demonstrate the effectiveness of this approach by using stochastic reconfiguration to optimize a correlator product state wavefunction with a pfaffian reference for four example systems. In two examples on the two dimensional Fermionic Hubbard model, we study 16 and 64 site lattices, recovering energies accurate to 1% in the smaller lattice and predicting particle-hole phase separation in the larger. In two examples involving an ab initio Hamiltonian, we investigate the potential energy curve of a symmetrically dissociated 4x4 hydrogen lattice as well as the singlet-triplet gap in free base porphin. In the hydrogen system we recover 98% or more of the correlation energy at all geometries, while for porphin we compute the gap in a 24 orbital active space to within 0.02 eV of the exact result. The number of variational parameters in these examples ranges from 4×10^3 to 5×10^5.

PACS numbers:

I. INTRODUCTION

Quantum Monte Carlo (QMC) is a powerful technique for extracting predictions from the electronic Schrödinger equation[1,2]. The variational (VMC) and diffusion (DMC) Monte Carlo methods in particular can produce highly accurate predictions provided that a sufficiently flexible trial wavefunction is available and that the variational parameters of this wavefunction can be optimized. However, VMC and DMC suffer from the major limitation that the most effective stochastic optimization algorithms cannot handle more than a few thousand variational parameters. These algorithms, which include the Newton[3], approximate Newton[4], linear (LM)[5–8] and stochastic reconfiguration (SR)[9,10] methods, are currently constrained by their need to build and store matrices that become unmanageable when the number of variational parameters becomes large. Other stochastic optimization algorithms[11,12] that rely only on stochastic estimates for the energy gradient can treat more variational parameters, but their steepest-descent character makes for less efficient convergence to the energy minimum, especially compared to the LM. In order to make effective use of sophisticated trial wavefunctions such as tensor networks, which can contain millions of variational parameters, it is imperative that more capable optimization methods be developed.

The LM and SR optimization methods reduce to solving either a system of linear equations or a linear eigenvalue problem in which the matrices in question are determined by stochastic sampling. The essential difficulty in this approach is that the dimension of these matrices is equal to the number of variational parameters, preventing their construction when there are more than a few thousand variables. Here we propose solving the central linear algebra problems of these optimization methods by using iterative Krylov subspace algorithms, which do not require the matrices to be built explicitly. Instead, these solvers require that one evaluate matrix-vector products, which we will show to be far less computationally expensive than actually building the relevant matrices. In VMC this approach is made particularly efficient by the strategy of operating by the matrices during the sampling process, as each sampled configuration contributes an outer product to the overall matrix, and outer products are particularly easy to operate by.

In this paper we will demonstrate this approach by using the conjugate gradient (CG) iterative solver to improve the SR method. We also derive a method for improving the LM using the generalized Davidson solver, although we will present numerical results only for SR (a computer implementation for the LM is underway). We begin by developing the theory for the accelerated SR and LM and also for the particular wavefunction ansatz that we employ. After developing the theory, we present numerical results for the SR method in four examples: (a) the Fermionic Hubbard model on a 4x4 lattice, (b) phase separation behavior in the 8x8 Fermionic Hubbard model, (c) the potential energy curve of a symmetrically dissociated 4x4 hydrogen lattice, and (d) the singlet-triplet gap of free base porphin. Note that the numerical studies carried out here are primarily concerned with the optimization problem. A detailed examination of the physics of these examples is left for future investigations.

II. THEORY

A. Accelerated Stochastic Reconfiguration

The SR method[9,10] can be viewed as an approximate imaginary time evolution in a specially chosen subspace Ω of the full Hilbert space. For a wavefunction $|\Psi(\alpha_1, \alpha_2, \ldots)\rangle$ with n_v variational parameters $\boldsymbol{\alpha}$, this subspace is spanned by the wavefunction and its $\boldsymbol{\alpha}$-derivatives,

$$\Omega = \text{span}\left(\ |\Psi^0\rangle,\ |\Psi^1\rangle,\ |\Psi^2\rangle,\ \ \ldots\ \right), \tag{1}$$

where $|\Psi^0\rangle \equiv |\Psi\rangle$ and $|\Psi^i\rangle \equiv \partial|\Psi\rangle/\partial\alpha_i$ for $i > 0$. The strategy of the SR method is to minimize the wavefunction's energy by repeatedly operating by $T = 1 - \tau H$ (the imaginary time evolution operator $e^{-\tau H}$ expanded to first order), where τ is a small number and H is the Hamiltonian. After each application of T, the result is projected into Ω to produce a new wavefunction of the form $|\Psi'\rangle = \sum_i x_i |\Psi^i\rangle$, in which the coefficients $\boldsymbol{x}$ are given by

$$\langle\Psi^i| \left(1 - \tau H\right) |\Psi\rangle = \sum_j^{n_v} \langle\Psi^i|\Psi^j\rangle x_j. \tag{2}$$

Finally, because τ and therefore x_i/x_0 are small, the new wavefunction $|\Psi'\rangle$ can be closely approximated by $|\Psi(\alpha_1', \alpha_2', \ldots)\rangle$, where $\alpha_i' = \alpha_i + x_i/x_0$. To summarize, one solves the linear equation given in Eq. (2) and updates $\boldsymbol{\alpha}$ accordingly, after which the subspace Ω is redefined for the new wavefunction. This entire procedure is repeated until the energy of the wavefunction has converged.

Previously, the SR overlap matrix $S_{ij} = \langle\Psi^i|\Psi^j\rangle$ was constructed explicitly. Here we will avoid building $\boldsymbol{S}$ entirely, relying instead on the CG algorithm to solve Eq. (2). This method proceeds iteratively, using information gained from a series of matrix-vector multiplications to successively refine an approximation to the solution $\boldsymbol{x}$ in a space of orthonormal conjugate vectors. The iteration proceeds until an arbitrary accuracy is achieved and typically converges in a number of steps far smaller than the dimension of the matrix. To see the advantages of using CG, consider the following expressions showing how the overlap matrix was previously constructed through stochastic sampling,

$$\frac{S_{ij}}{\langle\Psi|\Psi\rangle} = \sum_{\boldsymbol{n}} \frac{|\Psi_{\boldsymbol{n}}|^2}{\langle\Psi|\Psi\rangle} \left(\frac{\Psi^i_{\boldsymbol{n}}}{\Psi_{\boldsymbol{n}}}\right)^* \left(\frac{\Psi^j_{\boldsymbol{n}}}{\Psi_{\boldsymbol{n}}}\right), \tag{3}$$

$$|\Psi\rangle = \sum_{\boldsymbol{n}} \Psi_{\boldsymbol{n}} |\boldsymbol{n}\rangle, \tag{4}$$

$$|\Psi^i\rangle = \sum_{\boldsymbol{n}} \Psi^i_{\boldsymbol{n}} |\boldsymbol{n}\rangle. \tag{5}$$

Here a resolution of the identity $\sum_{\boldsymbol{n}} |\boldsymbol{n}\rangle\langle\boldsymbol{n}|$ has been inserted, creating a summation over all possible system configurations $|\boldsymbol{n}\rangle$. By multiplying and dividing by $|\Psi_{\boldsymbol{n}}|^2$ the summation has been formulated so that it can be evaluated stochastically by sampling from the distribution $|\Psi_{\boldsymbol{n}}|^2/\langle\Psi|\Psi\rangle$. (In a stochastic evaluation, the sum over all system configurations and the factor $|\Psi_{\boldsymbol{n}}|^2/\langle\Psi|\Psi\rangle$ in Eq. (3) are replaced by a sum over Monte Carlo samples drawn from this distribution.) However, building $\boldsymbol{S}$ stochastically using Eq. (3) takes at least $O(n_s n_v^2)$ time, where n_s is the number of samples and n_v is the number of variational parameters. Using the CG algorithm we may avoid this cost by instead evaluating matrix-vector products of the form $\boldsymbol{Sz}$. As with the expression for constructing $\boldsymbol{S}$, this expression can be evaluated by stochastic sampling if we insert a resolution of the identity,

$$\sum_j^{n_v} \frac{S_{ij}}{\langle\Psi|\Psi\rangle} z_j = \sum_j^{n_v} \sum_{\boldsymbol{n}} \frac{|\Psi_{\boldsymbol{n}}|^2}{\langle\Psi|\Psi\rangle} \left(\frac{\Psi^i_{\boldsymbol{n}}}{\Psi_{\boldsymbol{n}}}\right)^* \left(\frac{\Psi^j_{\boldsymbol{n}}}{\Psi_{\boldsymbol{n}}}\right) z_j. \tag{6}$$

By interchanging the order of summations we can rewrite this product as

$$\sum_j^{n_v} \frac{S_{ij}}{\langle\Psi|\Psi\rangle} z_j = \sum_{\boldsymbol{n}} \frac{|\Psi_{\boldsymbol{n}}|^2}{\langle\Psi|\Psi\rangle} \left(\frac{\Psi^i_{\boldsymbol{n}}}{\Psi_{\boldsymbol{n}}}\right)^* \left(\sum_j^{n_v} \frac{\Psi^j_{\boldsymbol{n}}}{\Psi_{\boldsymbol{n}}} z_j\right), \tag{7}$$

which can be evaluated in $O(n_s n_v)$ time provided that the derivative ratios $\Psi^i_{\boldsymbol{n}}/\Psi_{\boldsymbol{n}}$ have been pre-computed and stored, which is not difficult as the storage can be trivially divided between the different processors. In practice, we

also add a small diagonal shift when multiplying by the overlap matrix to remove linear dependencies that can arise from redundant variational parameters or stochastic error. While the CG algorithm does require multiple matrix-vector products to be evaluated, we will see that the number of such products will be much smaller than n_v, greatly improving the efficiency of the SR method.

B. Accelerated Linear Method

The linear method (LM), formulated by Nightingale for linear parameters[5] and later extended to optimize nonlinear parameters[6–8], works in the same subspace Ω as the SR method but typically converges more rapidly to the energy minimum. It can be viewed as an approximate Newton method with a built in stabilization[7] and often converges even more rapidly than the Newton method. Note that the LM should not be confused with similarly named linear *scaling* techniques, where the goal is typically to reduce the cost of evaluating Slater determinant ratios in systems with large numbers of electrons. Here we are concerned with reducing the cost of optimizing large numbers of variational parameters, and the way in which we take advantage of iterative Krylov methods differs significantly from recent work[13] employing them in order to reduce the cost of determinant ratio evaluations.

Instead of using imaginary time evolution, the LM optimizes $|\Psi\rangle$ by finding the eigenstate of lowest energy in the Ω subspace. This eigenstate can be found by solving the following generalized eigenvalue problem,

$$\sum_j \langle \Psi^i | H | \Psi^j \rangle x_j = E \sum_k \langle \Psi^i | \Psi^k \rangle x_k, \tag{8}$$

where we now take $\boldsymbol{x}$ to be the coefficients of the desired eigenvector. Once these coefficients are found, the variables $\boldsymbol{\alpha}$ can be updated to their new values in the same manner as in SR, though care must be taken to check that for large parameter changes the resulting parameters give an energy that is not higher outside of statistical errors. If they do not, the step can be scaled down by using a line search, or, rotated and scaled down by adding a diagonal shift. In practice, it is essential to modify the update in order to make it orthogonal to the original wavefunction, a procedure that can be performed using the information resulting from a single matrix-vector multiply involving the overlap matrix.

As with SR, the eigenvector $\boldsymbol{x}$ can be found without explicitly building the matrices $\boldsymbol{H}$ and $\boldsymbol{S}$ by using a Krylov subspace method, in this case the generalized Davidson algorithm[14]. As with CG, it is sufficient to evaluate the matrix-vector products of $\boldsymbol{H}$ and $\boldsymbol{S}$ with arbitrary trial vectors. For $\boldsymbol{S}$, this product can be performed efficiently as explained above. For $\boldsymbol{H}$, the difficulty of the multiplication depends on the complexity of the system's Hamiltonian, but for the relatively general case of the non-relativistic Born-Oppenheimer Hamiltonian in an orthonormal one-particle basis, an efficient evaluation is possible. If we assume a fixed particle number, we may use a matrix factorization such as the Cholesky decomposition[15] to express this Hamiltonian as

$$H = \sum_\mu \sum_{pqrs} L^\mu_{pq} R^\mu_{rs} a^\dagger_p a_q a^\dagger_r a_s, \tag{9}$$

where the operator $a^\dagger_p$ (a_p) is the fermionic creation (destruction) operator for the pth spin orbital, the index μ has a range of $O(n_o^2)$ (n_o is the number of orbitals/one-particle basis functions), and the indices p, q, r, s each have range n_o. In practice, the range of μ can often be taken to be much smaller than n_o^2 while still representing H with sufficient accuracy. By inserting an identity operator in the center of the Hamiltonian, the matrix-vector product on the left hand side of Eq. (8) can be written as

$$\begin{aligned} &\frac{1}{\langle \Psi | \Psi \rangle} \sum_j \langle \Psi^i | H | \Psi^j \rangle x_j \\ &= \sum_{j \boldsymbol{n} \mu p q r s} \frac{|\Psi_{\boldsymbol{n}}|^2}{\langle \Psi | \Psi \rangle} L^\mu_{pq} R^\mu_{rs} \frac{\langle \Psi^i | a^\dagger_p a_q | \boldsymbol{n} \rangle}{\Psi^*_{\boldsymbol{n}}} \frac{\langle \boldsymbol{n} | a^\dagger_r a_s | \Psi^j \rangle}{\Psi_{\boldsymbol{n}}} x_j \\ &= \sum_{\boldsymbol{n}} \frac{|\Psi_{\boldsymbol{n}}|^2}{\langle \Psi | \Psi \rangle} \sum_{pq} Q^*_{\boldsymbol{n} q p i} \sum_\mu L^\mu_{pq} \sum_{rs} R^\mu_{rs} \sum_j Q_{\boldsymbol{n} r s j} x_j, \end{aligned} \tag{10}$$

where we have defined the intermediate tensor $Q_{\boldsymbol{n} r s j} = \langle \boldsymbol{n} | a^\dagger_r a_s | \Psi^j \rangle / \Psi_{\boldsymbol{n}}$. For the wavefunction presented in the next section, for which $n_v \sim n_o^2$ due to the pairing matrix and long range pair correlators, this intermediate can be evaluated in $O(n_o^4)$ time for a given configuration $|\boldsymbol{n}\rangle$. If we sample the configurations $|\boldsymbol{n}\rangle$ from the distribution $|\Psi_{\boldsymbol{n}}|^2 / \langle \Psi | \Psi \rangle$,

we see that the entire matrix-vector product can be evaluated in $O(n_s n_o^4)$ time by performing the summations in the last line of Eq. (10) from right to left. For comparison, the cost of explicitly constructing the Hamiltonian matrix in the Ω subspace using our wavefunction is $O(n_s n_o^6)$. Thus by introducing an iterative Krylov solver, the cost of performing the LM for a non-relativistic Born-Oppenheimer Hamiltonian in an orthonormal one-particle basis can be reduced by a factor of the system size squared.

C. Wavefunction Ansatz

For our variational ansatz, we use a product of a correlator product state (CPS) tensor network[16,17] and a pfaffian pairing wavefunction[18–21]. As discussed in Ref.[22], the CPS ansatz can be expressed as a product of correlators acting on a reference wavefunction. Here we take the same approach, but with a pfaffian as the reference rather than a Slater determinant. The wavefunction is written as

$$|\Psi\rangle = \prod_p \hat{C}_p \left(\sum_{i<j} f_{ij} a_i^\dagger a_j^\dagger \right)^{N/2} |0\rangle, \tag{11}$$

where the operators $\hat{C}_p$ are correlators, $\boldsymbol{f}$ is the pairing matrix, N is the number of electrons, and $|0\rangle$ is the vacuum. The indices i, j range over all spin orbitals, so our pairing function creates both singlet and triplet pairs, unlike the more restrictive antisymmetrized geminal power[23,24]. Two typical types of correlators are long range pairs and $n \times n$ square plaquettes. In each case, both the spin $\uparrow$ and $\downarrow$ versions of the spatial orbitals are included in a correlator, so the number of spin orbitals (variational parameters) is 4 (2^4) for a pair correlator and $2n^2$ (2^{2n^2}) for an $n \times n$ plaquette. Note that when we employ both squares and long range pairs, we do not include pairs that are contained within a square as such pairs are redundant.

III. RESULTS

Here we demonstrate the accelerated SR method by applying it to four example systems in conjunction with our CPS-pfaffian ansatz. In each system, we restrict our sampling to configurations with the correct total number of electrons and the correct total S_z. We are currently constructing a computer implementation of the accelerated LM and will present results for it in a future publication.

A. 4x4 Fermionic Hubbard Model

In our first example we studied a 4x4 Fermionic Hubbard model at half filling with periodic boundary conditions, which was chosen as it is an exactly soluble system that contains many of the challenging features of the general 2D Fermionic Hubbard model. In one set of calculations, two translationally invariant 3x3 correlators were used (no additional symmetries were enforced), one anchored on each sublattice, giving a wavefunction with a total of 524,784 variational parameters (524,288 are from the CPS part of the wavefunction). In Figure 1 and Table I, we compare our results to the exact energies, seeing that for all ratios U/t our relative error is less than 1%. The exact result was computed using our own unpublished exact diagonalization software.

In a separate set of calculations, we investigated the convergence of the ground state energy with respect to the number of parameter update iterations. In this case we used as our correlators all 2×2 squares and all long range pairs. As seen in Figure 2, SR converged in roughly one fifth as many iterations as steepest descent, even when only ten CG iterations were used to compute each update step. We also see that there is no improvement when increasing the number of CG iterations to 100, showing that very few CG iterations are necessary compared to the number of variational parameters (5,488 in this case). One should note that in the case of an ab initio Hamiltonian, the cost of a few hundred CG iterations will be negligible compared to the cost of computing the gradient or the left hand side of Eq. (2). Thus in the ab initio case, the cost per parameter update iteration of SR and steepest descent will be essentially the same, greatly favoring the faster convergence of SR.

B. 8x8 Fermionic Hubbard Model

We have also applied our method to test for phase separation in the 2D Hubbard model, the exact nature of which remains an interesting and unresolved problem in solid state physics. To do so we studied an 8x8 lattice with twist-averaged boundary conditions (TABC)[25–27] (we used 12 randomly chosen twists), $U = 4$, and $t = 1$. We used translationally invariant 2x2 and long range pair correlators, again using separate correlators for each sublattice. To check whether the system phase separates, we computed the hole energy $e_h(h)$ employed in Ref.[27], which will display a minimum at the critical hole density h_c if phase separation occurs. The hole energy is defined as

$$e_h(h) \equiv \frac{e(1-h) - e(1)}{h}, \tag{12}$$

where h is the hole density and $e(x)$ is the energy per site for a particle density x. As seen in Figure 3, our approach predicts that the system will phase separate at a critical hole density $0.14 < h_c < 0.15$. This result provides a qualitative corroboration of the Constrained-Path Auxiliary Field QMC (CP-AFQMC)[27] results of Zhang et al, who predicted phase separation at $h_c = 0.1$ for the 8x8 lattice with TABC and $U/t = 4$. At half filling, our twist-averaged energy per site of -0.855 ± 0.002 is in close agreement with the CP-AFQMC result of -0.856.

C. 4x4 Hydrogen Lattice

As an example of a strongly correlated problem involving an ab initio Hamiltonian, we have studied a 4x4 square lattice of hydrogen atoms in the STO-3G orbital basis[28] at various nearest-neighbor distances. While this one particle basis is by no means large enough to capture all the details of electron correlation, it is sufficient to capture the particularly challenging strong correlations that arise during the simultaneous breaking of multiple bonds. As this system has open boundary conditions, we did not use translationally invariant correlators. Instead, we used all 2x2 and long range pair correlators, which results in a wavefunction with 4,048 variational parameters. As seen in Figure 4 and Table II, the results closely match those of the exact wavefunction in our chosen one particle basis, as obtained by full configuration interaction (FCI) using the Molpro[29] quantum chemistry package. At the H-H distance with the worst error, our approach still captures 98% of the correlation energy, which we define as the energy difference between the restricted Hartree Fock (RHF) and exact wavefunctions.

D. Free Base Porphin

As our final example, we computed the singlet-triplet gap of free base porphin in the 6-31G orbital basis[30]. This system was chosen as an important quantum chemical problem for which exact results in the active space are available for comparison. For both the singlet and triplet wavefunctions, the 1s and σ bonding orbitals resulting from a restricted Hartree Fock calculation were treated as a closed shell determinant, while the 24 out-of-plane 2p orbitals from the RHF solution were localized by the Pipek-Mezey[31] scheme to form an active space containing the remaining 26 electrons. This active space was treated with our CPS-pfaffian ansatz, with the correlators taken to be all pairs as well as those shown in Figure 5, for a total of 9,064 variational parameters. Holding the core orbitals frozen, we computed an active space singlet-triplet gap of 1.77eV, which compares very favorably with the converged spin-adapted density matrix renormalization group[32] result of 1.75eV.

IV. CONCLUSIONS

We have shown that by using the conjugate gradient iterative solver, it is possible to optimize hundreds of thousands of variational parameters with the stochastic reconfiguration algorithm in the context of variational Monte Carlo. In addition, we have presented an outline showing how the generalized Davidson solver can be used to provide a similar improvement for the linear method. Using our accelerated stochastic reconfiguration algorithm, we demonstrated that a CPS-pfaffian wavefunction ansatz is capable of treating a number of challenging two dimensional systems that display both weakly and strongly correlated physics. Together, these advances provide a powerful new method for modeling both quantum chemical and solid state systems.

V. ACKNOWLEDGMENTS

This work was supported by NSF grants CHE-1004603 and DMR-0908653 and by the Miller Institute for Basic Research in Science.

[1] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).
[2] M. P. Nightingale and C. J. Umrigar, editors, *Quantum Monte Carlo Methods in Physics and Chemistry*, NATO ASI Ser. C 525, Kluwer, Dordrecht, 1999.
[3] C. J. Umrigar and C. Filippi, Phys. Rev. Lett. **94**, 150201 (2005).
[4] S. Sorella, Phys. Rev. B **71**, 241103 (2005).
[5] M. P. Nightingale and V. Melik-Alaverdian, Phys. Rev. Lett. **87**, 043401 (2001).
[6] J. Toulouse and C. J. Umrigar, J. Chem. Phys. **126**, 084102 (2007).
[7] J. Toulouse and C. J. Umrigar, J. Chem. Phys. **128**, 174101 (2008).
[8] C. J. Umrigar, J. Toulouse, C. Filippi, S. Sorella, and R. G. Hennig, Phys. Rev. Lett. **98**, 110201 (2007).
[9] S. Sorella, Phys. Rev. B **64**, 024512 (2001).
[10] M. Casula1, C. Attaccalite, and S. Sorella, J. Chem. Phys. **121**, 024512 (2004).
[11] A. W. Sandvik and G. Vidal, Phys. Rev. Lett. **99**, 220602 (2007).
[12] A. Harju, B. Barbiellini, S. Siljamäki, R. M. Nieminen, and G. Ortiz, Phys. Rev. Lett. **79**, 1173 (1997).
[13] K. Ahuja, B. K. Clark, E. de Sturler, D. M. Ceperley, and J. Kim, SIAM J. Sci. Comput. **33**, 1837 (2011).
[14] R. B. Morgan, J. Comput. Phys. **89**, 241 (1990).
[15] N. H. F. Beebe and J. Linderberg, Int. J. Quantum Chem. **12**, 683 (1977).
[16] H. J. Changlani, J. M. Kinder, C. J. Umrigar, and G. K.-L. Chan, Phys. Rev. B. **80**, 245116 (2009).
[17] F. Mezzacapo, N. Schuch, M. Boninsegni, and J. I. Cirac, New J. Phys. **11**, 083026 (2009).
[18] M. Bajdich, L. Mitas, G. Drobný, L. K. Wagner, and K. E. Schmidt, Phys. Rev. Lett. **96**, 130201 (2006).
[19] M. Bajdich, L. Mitas, L. K. Wagner, and K. E. Schmidt, Phys. Rev. B **77**, 115112 (2008).
[20] M. Wimmer, arXiv:1102.3440v2 (2011).
[21] J. Bouchaud, A. Georges, and C. Lhuillier, J. Phys. (France) **49**, 553 (1988).
[22] E. Neuscamman, H. Changlani, J. Kinder, and G. K.-L. Chan, Phys. Rev. B **84**, 205132 (2011).
[23] A. C. Hurley, J. Lennard-Jones, and J. A. Pople, Proc. R. Soc. London, Ser. A **220**, 446 (1953).
[24] A. J. Coleman, J. Math. Phys. **6**, 1425 (1965).
[25] D. Poilblanc, Phys. Rev. B **44**, 9562 (1991).
[26] C. Gros, Phys. Rev. B **53**, 6865 (1996).
[27] C.-C. Chang and S. Zhang, Phys. Rev. B **78**, 165101 (2008).
[28] W. J. Hehre, R. F. Stewart, and J. A. Pople, J. Chem. Phys. **51**, 2657 (1969).
[29] H.-J. Werner, P. J. Knowles, R. Lindh, F. R. Manby, and M. Schütz, MOLPRO, version 2006, a package of ab initio programs, see http://www.molpro.net.
[30] W. J. Hehre, R. Ditchfield, and J. A. Pople, J. Chem. Phys. **56**, 2257 (1972).
[31] J. Pipek and P. G. Mezey, J. Chem. Phys. **90**, 4916 (1989).
[32] S. Sharma and G. K.-L. Chan, in preparation .

TABLE I: Total ground state energies in unites of t of a periodic 4x4 Hubbard lattice at half filling with different ratios U/t. Statistical errors in the final digit are given in parentheses. See Section III A.

U/t	CPS-Pfaffian	Exact
1	-1.29871(1)	-1.299602
2	-1.12340(2)	-1.126098
4	-0.84539(4)	-0.851366
6	-0.65326(4)	-0.659514
8	-0.52440(4)	-0.529305
10	-0.43689(2)	-0.439313
12	-0.37281(2)	-0.374514
14	-0.32434(1)	-0.325925
16	-0.28697(1)	-0.288241

TABLE II: Total ground state energies in E_h for the 4x4 Hydrogen lattice at various nearest neighbor distances R in the STO-3G basis. Statistical errors in the final digit are given in parentheses. See Section III C.

R (Å)	RHF	CPS-Pfaffian	FCI
0.6	-3.304481	-3.4541(4)	-3.460659
1.0	-7.534453	-7.7765(2)	-7.785104
1.4	-7.396180	-7.8914(1)	-7.903634
1.8	-6.669009	-7.6767(1)	-7.684200
2.2	-5.997164	-7.53726(3)	-7.539753
2.6	-5.438428	-7.48167(3)	-7.486036
3.0	-5.196679	-7.46921(1)	-7.470067

Figure Captions

FIG. 1: Relative energy errors for the CPS-Pfaffian ansatz on a periodic 4x4 Hubbard lattice. Statistical errors are smaller than the symbol size and lines are guides to the eye. See Section III A.

FIG. 2: Convergence of the total energy of the 4x4 Hubbard model at $U = 4$ and $t = 1$ with respect to the number of parameter update iterations for steepest descent and stochastic reconfiguration. The same random seed and a step size of $\tau = 0.01$ were used for both steepest descent and stochastic reconfiguration (for steepest descent the step taken was $-\tau$ times the energy gradient). The initial guess for the wavefunction is an unrestricted Hartree-Fock solution with slightly randomized molecular orbitals. See Section III A.

FIG. 3: The hole energy $e_h(h)$ on an 8x8 Hubbard lattice with twist-averaged boundary conditions, $U = 4$, and $t = 1$. The presence of a minimum implies that our ansatz predicts phase separation in the 2D Hubbard model. See Section III B.

FIG. 4: Total energies, in Hartrees, of a 4x4 hydrogen lattice. Statistical errors are smaller than the symbol size and lines are guides to the eye. See Section III C.

FIG. 5: In treating free base porphin, we used the following correlators between out-of-plane 2p orbitals: (i) all long range orbital pairs, (ii) the 3-orbital elbows highlighted by rectangles, and (iii) the 5-orbital rings highlighted by pentagons. See Section III D.

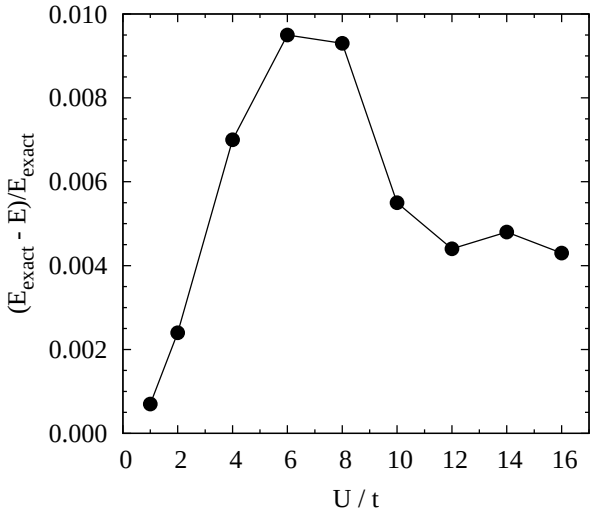
0.010
0.008
0.006
0.004
0.002
0.000
0
2
4
6
8
10
12
14
16
U / t
(E_exact - E)/E_exact

FIG. 1:

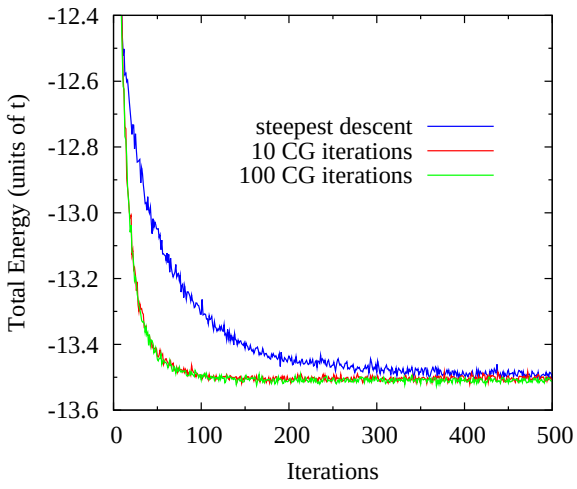
-12.4
-12.6
-12.8
-13.0
-13.2
-13.4
-13.6
Total Energy (units of t)
0
100
200
300
400
500
Iterations
steepest descent
10 CG iterations
100 CG iterations

FIG. 2:

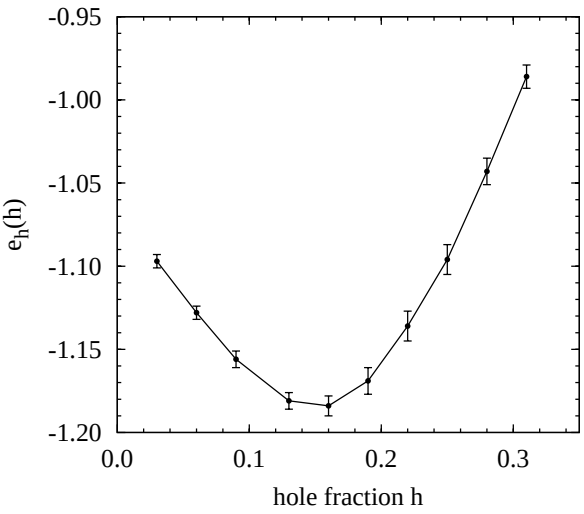

FIG. 3:

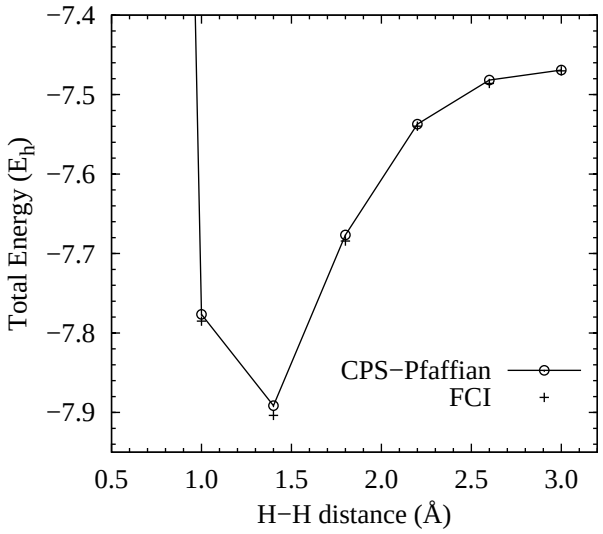
−7.4
−7.5
−7.6
−7.7
−7.8
−7.9
Total Energy (E_h)
0.5
1.0
1.5
2.0
2.5
3.0
H−H distance (Å)
CPS−Pfaffian
FCI

FIG. 4:

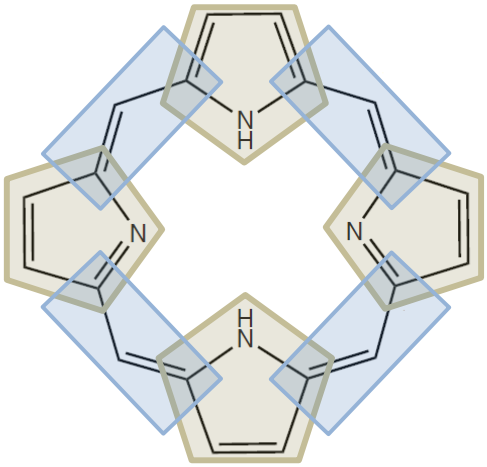

FIG. 5: