Sign problem in full configuration interaction quantum Monte Carlo: Linear and sublinear representation regimes for the exact wave function

James J. Shepherd,[1] Gustavo E. Scuseria,[1] and James S. Spencer[2]

[1]*Department of Chemistry and Department of Physics and Astronomy, Rice University, Houston, Texas 77005-1892, USA*

[2]*Department of Materials and Department of Physics, Imperial College London, Exhibition Road, London SW7 2AZ, United Kingdom*

We investigate the sign problem for full configuration interaction quantum Monte Carlo (FCIQMC), a stochastic algorithm for finding the ground-state solution of the Schrödinger equation with substantially reduced computational cost compared with exact diagonalization. We find k-space Hubbard models for which the solution is yielded with storage that grows sublinearly in the size of the many-body Hilbert space, in spite of using a wave function that is simply a linear combination of states. The FCIQMC algorithm is able to find this sublinear scaling regime without bias and with only a choice of the Hamiltonian basis. By means of a demonstration we solve for the energy of a 70-site half-filled system (with a space of 10^{38} determinants) in 250 core hours, substantially quicker than the $\sim 10^{36}$ core hours that would be required by exact diagonalization. This is the largest space that has been sampled in an unbiased fashion. The challenge for the recently developed FCIQMC method is made clear: Expand the sublinear scaling regime while retaining exact-on-average accuracy. We comment upon the relationship between this and the scaling law previously observed in the initiator adaptation (i-FCIQMC). We argue that our results change the landscape for the development of FCIQMC and related methods.

I. INTRODUCTION

Exact methods for solving the Schrödinger equation are used at the forefront of understanding in condensed matter physics [1–4] and in molecular quantum chemistry [5–7]. However, exact parametrizations of the many-body wave function for a general system of interacting fermions scale exponentially with the system size, i.e., $O(e^N)$. Quantum Monte Carlo (QMC) techniques attempting to determine these parameters are hindered by their values being either positive or negative, causing more pronounced variability than a set of parameters with a single sign. Unconstrained algorithms [8] therefore also scale exponentially. Constrained algorithms, in which the sign of the wave-function parameters are fixed (e.g., by a trial wave function), can be polynomially scaling, but at the cost of a bias [9–12]. The (unsolved) challenge to find a constraint on the signs of a wave function, while still reproducing the exact result, is usually termed the fermion sign problem. For some QMC methods and specific systems this constraint can be imposed exactly, abstracting away the fundamental complexity of the problem. Overcoming the sign problem is vital for the accurate treatment of real systems.

Here, we investigate the sign problem of a recently developed QMC method developed for use in finite molecular basis sets: full configuration interaction QMC (FCIQMC) [5]. This is the direct (ground-state) QMC analog of exact diagonalization, finding the exact lowest-energy solution for a finite Hilbert space with an exponential number of states using a walker-based algorithm, where the Hilbert space is a set of Slater determinants and grows exponentially with the number of fermions. Since it does not impose the signs of the wave function in advance, this method does in general have a sign problem [13], and the cost of the storage of the exact-on-average wave function has been shown to scale linearly in the size of the Hilbert space for a series of atomic systems [14–16]. Inspired by recent interest in high-throughput data driven informatics [17], here we study 378 systems with a plateau using a high-throughput approach. In so doing, we considerably extend the information available about this method.

As shown below, we discover a regime of the k-space one-dimensional (1D) Hubbard model where the amount of information required to store the ground-state wave vector has sublinear scaling with the size of the Hilbert space states. This is achieved during the simulation in the presence of a sign problem, assuming a linear wave-function ansatz, without requiring any information or bias beyond the Hamiltonian. This regime is smoothly connected to the more typical linear-scaling regime including situations in which there are more walkers required than the number of states. We use this to build a conceptual map based on the system's parameter space (Hubbard U and size of Hilbert space) for the regions of scaling for this promising QMC technique. We relate these findings back to FCIQMC and its initiator adaptation, providing concrete insight for the development of these methods. We discuss whether, in light of this, QMC for the FCI problem could receive routine use for the treatment of correlated electrons in many more realistic contexts.

II. THEORY

FCIQMC, as all projector-based methods, exploits the fact that $|\Psi(\tau)\rangle = e^{-\tau \hat{H}} |\Psi(0)\rangle$ tends to the ground-state solution of the imaginary time Schrödinger equation in the limit of $\tau \to \infty$ if $|\Psi(0)\rangle$ has a nonzero overlap with the ground state. The configuration interaction wave-function ansatz $|\Psi\rangle = \sum_i c_i |D_i\rangle$ is used, where $\{|D_i\rangle\}$ is a set of Slater determinants of size N_{det} formed from N_{elec} electrons and N_{orb} spin orbitals. A first-order Euler finite-difference approximation to the imaginary-time Schrödinger equation gives

$$c_i(\tau + \delta\tau) = c_i(\tau) - \delta\tau \sum_j H_{ij} c_j(\tau), \qquad (1)$$
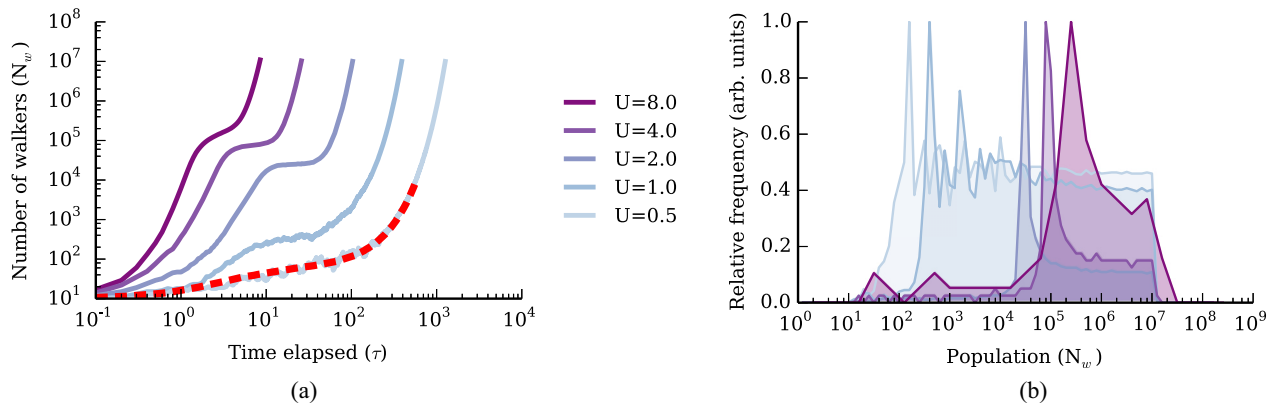
FIG. 1. (Color online) These two panels describe how we are able to identify plateaus. In (a), the decrease in U is shown to obscure the plateau. The bold, black (red online), dashed annotation indicates the $U = 0.5$ population growth averaged over 100 random number seeds. In (b), we show that the population at the plateau corresponds to the maximum value of histogram of the population (although, in general, variable bin widths were used). These two panels share a common key.

where $H_{ij} = \langle D_i | \hat{H} - S | D_j \rangle$ and an energy offset ("shift") S has been introduced in order to conserve normalization. For a sufficiently small time step [13] $\delta\tau$ the coefficients tend to the ground state of the Hamiltonian matrix. Although FCIQMC is essentially a stochastic version of the power method [13,18], the core algorithm has inspired a wide range of developmental advances and new methods [19].

The wave-function coefficients are discretized by representing them with a set of signed walkers [20]. In each time step, the set of walkers is considered each in turn and Eq. (1) sampled according to unbiased rules [14]. Since the simulation allows the sign of a site to change, because the off-diagonal matrix elements are of different signs, the arithmetic that occurs on a site to determine its overall sign involves a process termed annihilation, removing pairs of walkers which belong to the same determinant but have opposite signs. The annihilation step preserves the expected distribution of walkers and crucially prevents the growth of exponential noise and collapse onto the sign-problem-free ground state of the matrix defined by $\tilde{H}_{ij} = H_{ij}\delta_{ij} - (1 - \delta_{ij})|H_{ij}|$ [13].

The sign problem arises in FCIQMC because the signs of the FCI coefficients are not known in advance. If they were known, it is postulated that the sign problem could be removed by factorization [21–23]. In contrast, the matrix $\tilde{H}$ has coefficients that are all of the same sign in much the same way as a bosonic wave function has the same sign in its value for all particle coordinates. This is the determinant space analog of the real-space bosonic solution for FCIQMC [24].

A typical simulation contains four distinct phases [25]. Initially the shift is held constant (typically to a mean-field energy) and the population of walkers grows exponentially. The population spontaneously stops growing and enters the *plateau* phase at a system-dependent population, during which the ground-state sign structure emerges. The population spontaneously begins to grow again at an exponential rate, albeit slower than before. The shift is then varied to keep the population approximately constant. Above the plateau, statistics can be accumulated that are demonstrably from the exact solution [14]; the postplateau population is a stochastic

representation of the exact wave function. The first three phases can be seen in Fig. 1(a).

III. PLATEAU DETERMINATION

The plateau is therefore a very powerful conceptual feature of FCIQMC. Phenomenologically, the plateau provides an unambiguous signal of how hard the sign problem is because it represents the minimum storage cost for an on-average exact representation of the FCI vector [13,14]. Computationally, this number of walkers determines the dominant scaling bottleneck, for both memory and computer time, of the method since each Monte Carlo iteration loops over this list. The stochastic sampling of the propagator will also contribute to the noise of the simulation, but this is premultiplied by the length of the main vector.

Crucially therefore (and uniquely in projector QMC methods) the plateau provides an unambiguous measure of the sign problem in FCIQMC: By comparing the plateau height against the number of determinants in the Hilbert space, we obtain a measure for how "hard" a system is for FCIQMC.

In order to study plateau heights, it is important to establish a unique and reproducible definition accounting for variations seen in Fig. 1(a), where plateaus are obscured by becoming "shoulder" shaped or being overwhelmed by stochastic noise. The plateau can be thought of as the walker population that the simulation spends the most time at. To find this, the relative frequency that a certain population window appears in the simulation is computed and the maximum of this distribution found. The histogram of the logarithm of the population rather than that of the population is used in order to handle the exponential growth in population. The plateau signal is shown for various values of U in Fig. 1(b). The disadvantage of this approach is that for some values of U, this can lead to *overestimation* of the plateau for some runs as multiple peaks compete. This can be circumvented by changing the bin width, and this must sometimes be interpreted manually. This procedure is discussed in more detail in the Supplemental Material, where each plateau can also be verified by inspection [27].
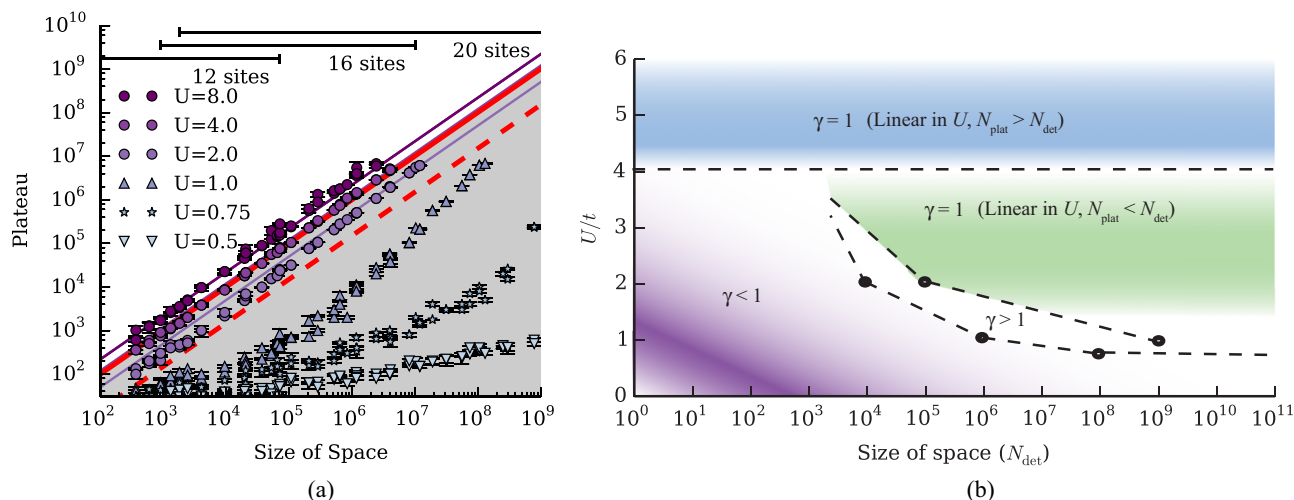
FIG. 2. (Color online) Analysis of all plateaus found in this study. In (a), plateau heights are plotted against the sizes of space. In (b) the scaling trends are summarized in a diagram relating it to system parameters. In (a), the following bold, black (red online) annotations have been made. The solid line indicates the $N_{\text{plat}} = N_{\text{det}}$, where we are storing one walker per determinant (on average) across the simulation and the storage becomes comparable to FCI (unshaded region). The dashed line indicates where $U = 1.0$ would be if it retained the same plateau height scaling as high U values. The size of space was calculated by a Monte Carlo method and takes into account momentum symmetry [26].

IV. PLATEAU ANALYSIS

We now consider the 1D translationally invariant (k-point) one-band Hubbard model, for the parameter range $U = 0.5$, 0.75, 1.0, 2.0, 4.0, and 8.0 for $N_s = 12, 14, 16, 18, 20,$ and 22 sites per simulation cell. We explore a wide range of doping levels (4 to $2N_s - 4$), where the spin is always unpolarized, for consistent simulation parameters [28]. All calculations were performed with the HANDE QMC code [29]. We focus on 1D systems because shell-filling effects were anticipated to make interpretation substantially more difficult in higher dimensions. Although this system is only of one dimension, the range of parameters encompasses a wide range of correlation regimes.

All of the plateaus we have found are plotted in Fig. 2(a). We can use this to probe the different scaling laws, based on $N_{\text{plat}} = \beta N_{\text{det}}^{\gamma}$, where the exponent γ is defined by the tangent to the curve at a given N_{det}. Some of the trends described below are slight and to aid readers a larger version of the graph can be found in Supplemental Material.

A. Linear in N_{det} and in U

This is the conventional scaling regime that has been previously observed. Three diagonal-running parallel lines ($N_{\text{det}} \propto N_{\text{plat}}$) fit the data from $U = 2.0, 4.0,$ and 8.0 at high N_{det}. The behavior of the gradient with U is consistent with the plateau being linear in U (as shown in Ref. [13]). Overall, therefore, $\gamma = 1$ and $\beta \propto U$. We note in passing that these trends are remarkably consistent as doping and the particle number are changed.

The bold red line, almost coincident with most of the $U = 4.0$ data set, shows the line of $N_{\text{det}} = N_{\text{plat}}$ where on average we store the same number of integers as the size of the space. The gray shading indicates where we would expect, therefore, to store less information than the full wave vector in order to obtain the solution via an exact diagonalization (or FCI);

above this line in the unshaded region the memory requirement is comparable to FCI. Although this is true for storage, the computational time is still expected to be linear in the size of the space, and this (being the upper limit of the scaling here) is better than many diagonalization routines.

B. Sublinear in N_{det}, nonlinear in U

At sufficiently small system sizes, sublinear scaling ($\gamma < 1$) is observed for all U except $U = 8.0$. The region of this reduced scaling depends on U, and extends to larger system sizes for smaller U. The lowest measurable exponent observed is $\gamma \approx 0.1$ for $U = 0.5$. The reduced exponent is surprising for two reasons. The first is that the FCI wave function is apparently representable with storage that is sublinear with the size of the space. The second is that the projector algorithm here is able to find this minimal representation with no additional information than the Hamiltonian, and in particular no biasing.

C. Nonlinear in N_{det} and in U

In the intermediate region between these two regimes, there is a polynomial region in N_{det} ($\gamma < 1$) as the scaling law seems to return to the original linear scaling regime (with no shift) and crosses through $\gamma = 1$. This is most prominently seen by careful examination of $U = 2.0$, which only slightly deviates from the conventional scaling laws. The return to $\gamma = 1$ appears around $N_{\text{det}} = 10^5$. The dashed red line shows where we would expect the $U = 1.0$ scaling to be if the linear scaling with N_{det} and U continued, which our data set never reaches. Nonetheless, the limiting scaling at high N_{det} seems to be linear in N_{det} and U.

D. Sign problem diagram

These scaling relationships are summarized with respect to the system parameters N_{det} and U in Fig. 2(b). The tie lines

we have plotted are made by hand, and are estimates limited by the breadth of our data set. In particular, sharp lines should be considered as estimates and not definitive. This provides a comprehensive summary of the known information about the sign problem, and scaling, in FCIQMC.

We also observe that as the system size is raised, the method returns to linear scaling in the size of the space and exponential scaling in the particle number. This is interesting because it seems as it is a reverse to what might be expected to happen. As the system gets closer to the thermodynamic limit, quantities such as the energy become extensive (i.e., scale linearly with N_{elec}). As the correlation length is increasingly well contained within the simulation cell, we would expect the problem to become easier and of improved scaling (due to self-averaging).

V. CONCLUSION AND DISCUSSIONS

Our principal conclusion is the discovery of a regime of the k-space Hubbard model where the exact ground state can be stored with sublinear representation cost. This exact-on-average representation requires no prior knowledge beyond on-the-fly access to the Hamiltonian matrix elements, and we believe this finding to be widely significant [30]. By means of a practical demonstration of the significance of this regime, we can find the ground-state energy for the half-filled 70-site system for $U = 0.1$ in 250 core hours [$E = -87.418\,564(7)t$]. The size of space is 10^{38} determinants, and this is the largest unbiased simulation to date. By way of comparison, we estimate exact diagonalization would take 10^{36} core hours, based on known scaling laws and calculations from smaller system sizes using the algorithm implemented in ALPS [31,32]. This poses the question: How many more, larger, systems are available for study that have simply not yet been attempted?

The low-scaling regime, occurring in a greater range of systems at low U, seems coincident with the weakly interacting, or weakly correlated, regime. Although this is a tempting conclusion to draw, this is not a link that we have the scope to explore in detail here. This is in part due to the sign problem being representation dependent. In particular, the 1D Hubbard model is a toy model, not only because it is already solvable at polynomial cost [33,34], but also because a transformation to the real space basis set leaves it sign problem free for FCIQMC [13]. Nevertheless, where there exist large expanses of the Slater determinant space that are redundant, FCIQMC should be able to find them, but that sparsity must exist to be found.

Finding such representations is greatly facilitated by our study here. This is first and foremost because we demonstrate the potential benefits to be found, but also for the resources this study provisions for the development of FCIQMC. We start a database of plateaus, semiautomated plateau height analysis, and a practical understanding as to what might happen to the plateau or sign problem with further development. We hope

that these concerns are placed at the forefront of FCIQMC development. One such route of promise and significance is the discussion of symmetry breaking and restoration in the context of QMC techniques [35–38]. Another is the adaptation of the core algorithm to other Fock space QMC methods [18,39–41], but in the wider context of other quantum chemical methods it is important to know whether there is a sign problem at all [40,42–45].

The appearance of a sublinear regime is interesting because it mirrors some evidence that the initiator adaptation has this scaling in its wave-function representation [16,46]. The initiator adaptation (i-FCIQMC) imposes a population dependence on H_{ij}, zeroing some elements that are considered outside the currently well-sampled space. This greatly enlarges the size of systems that can be sampled, up to 10^{108} determinants to date [47], but at the cost of a systematically improvable bias. In this context, therefore, i-FCIQMC is an approximate (but systematically improvable) method that *expands* the sublinear scaling rather than this scaling being unique to this approximation. This strongly implies there are further improvements that can be made to expand this reduced scaling still further. We would welcome further investigation of the mechanism behind this observation; analysis of the sign problem in i-FCIQMC likely requires a very detailed analysis of the initiator error, for which one method is outlined in Ref. [48].

To the wider community in QMC methods, we hope this shows that FCIQMC provides interesting phenomenology and therefore something else to offer beyond FCI-quality energies. The analysis we present here argues that a sign problem that is easy to detect is potentially more informative than an error that is unquantifiable. It also demonstrates that FCIQMC does have the *potential* to solve large systems, which is surely required for its application in condensed matter physics, provided that its sign problem can be controlled. This puts emphasis back onto understanding and solving the sign problem, which is also a more universal effort.

[1] L. Capriotti, D. J. Scalapino, and S. R. White, Phys. Rev. Lett. **93**, 177004 (2004).

[2] M. Daghofer, A. Moreo, J. A. Riera, E. Arrigoni, D. J. Scalapino, and E. Dagotto, Phys. Rev. Lett. **101**, 237004 (2008).

[3] T. A. Maier, D. Poilblanc, and D. J. Scalapino, Phys. Rev. Lett. **100**, 237001 (2008).

[4] S. Yan, D. A. Huse, and S. R. White, Science **332**, 1173 (2011).

[5] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi, Nature (London) **493**, 365 (2013).

[6] Y. Kurashige, G. K.-L. Chan, and T. Yanai, Nat. Chem. **5**, 660 (2013).

[7] S. Sharma, K. Sivalingam, F. Neese, and G. K.-L. Chan, Nat. Chem. **6**, 927 (2014).

[8] D. M. Ceperley and B. J. Alder, Phys. Rev. Lett. **45**, 566 (1980).

[9] S. Zhang and M. H. Kalos, Phys. Rev. Lett. **67**, 3074 (1991).

[10] S. Zhang, J. Carlson, and J. E. Gubernatis, Phys. Rev. Lett. **74**, 3652 (1995).

[11] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).

[12] M. Capello, F. Becca, M. Fabrizio, and S. Sorella, Phys. Rev. Lett. **99**, 056402 (2007).

[13] J. S. Spencer, N. S. Blunt, and W. M. C. Foulkes, J. Chem. Phys. **136**, 054110 (2012).

[14] G. H. Booth, A. J. W. Thom, and A. Alavi, J. Chem. Phys. **131**, 054106 (2009).

[15] G. H. Booth and A. Alavi, J. Chem. Phys. **132**, 174104 (2010).

[16] D. M. Cleland, G. H. Booth, and A. Alavi, J. Chem. Phys. **134**, 024112 (2011).

[17] A. Jain, G. Hautier, C. J. Moore, S. Ping Ong, C. C. Fischer, T. Mueller, K. A. Persson, and G. Ceder, Comput. Mater. Sci. **50**, 2295 (2011).

[18] F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale, and C. J. Umrigar, Phys. Rev. Lett. **109**, 230201 (2012).

[19] These include CCMC [40], SQMC [18], DMQMC [41], FCIQMC for excited states [51], and MSQMC [39].

[20] It has recently been shown that using walkers with noninteger weights can be more efficient [18].

[21] M. Kolodrubetz and B. K. Clark, Phys. Rev. B **86**, 075109 (2012).

[22] A. Roggero, A. Mukherjee, and F. Pederiva, Phys. Rev. B **88**, 115138 (2013).

[23] A. Mukherjee and Y. Alhassid, Phys. Rev. A **88**, 053622 (2013).

[24] M. H. Kolodrubetz, J. S. Spencer, B. K. Clark, and W. M. C. Foulkes, J. Chem. Phys. **138**, 024110 (2013).

[25] Here, we have concentrated on features which are salient to our investigation; for full details refer to, e.g., Refs. [5,13,14,18,46,48,52,53].

[26] G. Booth, Ph.D. thesis, University of Cambridge, 2010.

[27] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevB.90.155130 which contains finer details on the histogram method for determining plateau height and population graphs for every system considered and includes Refs. [54–56].

[28] In general, we fix the initial shift to the correlation energy zero, the time step to be $\delta\tau = 0.001$ a.u., and start with an initial population of ten positive walkers on the restricted Hartree-Fock determinant. The plateaus are somewhat sensitive to the initial population and $\delta\tau$, but our general conclusions are not.

[29] HANDE QMC code available by request and in preparation, www.hande.org.uk.

[30] We disclaim in passing that we have not considered stochastic error, deliberately. For a discussion of stochastic error in the polynomial-scaling diffusion Monte Carlo, see Ref. [57].

[31] A. Albuquerque, F. Alet, P. Corboz *et al.*, J. Magn. Magn. Mater. **310**, 1187 (2007).

[32] B. Bauer, L. D. Carr, H. G. Evertz *et al.*, J. Stat. Mech.: Theor. Exp. (2011) P05001.

[33] E. H. Lieb and F. Y. Wu, Phys. Rev. Lett. **20**, 1445 (1968).

[34] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).

[35] R. E. Thomas, C. Overy, G. H. Booth, and A. Alavi, J. Chem. Theory Comput. **10**, 1915 (2014).

[36] H. Shi, C. A. Jiménez-Hoyos, R. Rodríguez-Guzmán, G. E. Scuseria, and S. Zhang, Phys. Rev. B **89**, 125129 (2014).

[37] H. Shi and S. Zhang, Phys. Rev. B **88**, 125132 (2013).

[38] C. A. Jiménez-Hoyos, T. M. Henderson, T. Tsuchimochi, and G. E. Scuseria, J. Chem. Phys. **136**, 164109 (2012).

[39] S. Ten-no, J. Chem. Phys. **138**, 164126 (2013).

[40] A. J. W. Thom, Phys. Rev. Lett. **105**, 263004 (2010).

[41] N. S. Blunt, T. W. Rogers, J. S. Spencer, and W. M. C. Foulkes, Phys. Rev. B **89**, 245124 (2014).

[42] S. Y. Willow and S. Hirata, J. Chem. Phys. **140**, 024111 (2014).

[43] D. Neuhauser, E. Rabani, and R. Baer, J. Chem. Theory Comput. **9**, 24 (2013).

[44] S. Y. Willow, K. S. Kim, and S. Hirata, J. Chem. Phys. **138**, 164111 (2013).

[45] S. Y. Willow, K. S. Kim, and S. Hirata, J. Chem. Phys. **137**, 204122 (2012).

[46] D. Cleland, G. H. Booth, and A. Alavi, J. Chem. Phys. **132**, 041103 (2010).

[47] J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi, Phys. Rev. B **85**, 081103 (2012).

[48] J. J. Shepherd, G. H. Booth, and A. Alavi, J. Chem. Phys. **136**, 244101 (2012).

[49] A. J. W. Thom (unpublished).

[50] http://www.imperial.ac.uk/ict/services/teachingandresearchservices/highperformancecomputing.

[51] G. H. Booth and G. K.-L. Chan, J. Chem. Phys. **137**, 191102 (2012).

[52] G. H. Booth, D. Cleland, A. J. W. Thom, and A. Alavi, J. Chem. Phys. **135**, 084104 (2011).

[53] G. H. Booth, S. D. Smart, and A. Alavi, Mol. Phys. **112**, 1855 (2014).

[54] J. J. Shepherd, L. R. Schwarz, R. E. Thomas, G. H. Booth, D. Frenkel, and A. Alavi, arXiv:1209.4023.

[55] http://dx.doi.org/10.6084/m9.figshare.1101537; http://dx.doi.org/10.6084/m9.figshare.1096136.

[56] http://dx.doi.org/10.6084/m9.figshare.1106864.

[57] N. Nemec, Phys. Rev. B **81**, 035119 (2010).