Fermion Monte Carlo without fixed nodes: A Game of Life, death and annihilation in Slater Determinant space

George H. Booth[(a)], Alex J. W. Thom[(a,b)], and Ali Alavi[(a)*]
[(a)] *University of Cambridge, Chemistry Department, Lensfield Road, Cambridge CB2 1EW, U.K. and*
[(b)] *University of California Berkeley, Department of Chemistry, Berkeley, CA 94720 U.S.A.*
(Dated: July 12, 2009)

We have developed a new Quantum Monte Carlo method for the simulation of correlated many-electron systems in Full Configuration-Interaction (Slater Determinant) spaces. The new method is a population dynamics of a set of walkers, and is designed to simulate the underlying imaginary-time Schrödinger equation of the interacting Hamiltonian. The walkers (which carry a positive or negative sign) inhabit Slater determinant space, and evolve according to a simple set of rules which include spawning, death and annihilation processes. We show that this method is capable of converging onto the Full Configuration-Interaction (FCI) energy and wavefunction of the problem, without any *a priori* information regarding the nodal structure of the wavefunction being provided. Walker annihilation is shown to play a key role. The pattern of walker growth exhibits a characteristic plateau once a critical (system-dependent) number of walkers has been reached. At this point, the correlation energy can be measured using two independent methods – a projection formula and a energy shift; agreement between these provides a strong measure of confidence in the accuracy of the computed correlation energies. We have verified the method by performing calculations on systems for which FCI calculations already exist. In addition, we report on a number of new systems, including CO, O_2, CH_4 and NaH – with FCI spaces ranging from 10^9 to 10^{14}, whose FCI energies we compute using modest computational resources.

PACS numbers:

INTRODUCTION

It has long been known that the only major obstacle preventing the exact numerical simulation of many-electron systems via stochastic methods such as Diffusion quantum Monte Carlo (DMC) [1] or the related Green's function Monte Carlo (GFMC) [2] is the Fermion sign problem [3]. This problem stems from the antisymmetry property of many-electron wavefunctions to electron exchange, which leads to wavefunctions which have both positive and negative amplitudes. Since the Schrödinger equation can be viewed as a diffusion equation in imaginary time, its lowest energy solution is in general nodeless and symmetric, and therefore does not satisfy the required Fermion antisymmetry. Stochastic propagation leads exponentially quickly to this undesired solution. One way to prevent this "Boson catastrophe" is to constrain the propagation to disjoint areas of similar sign using the fixed-node approximation [4–7], a procedure which would be exact if the applied nodal boundaries coincided with the exact nodal hypersurface of the ground-state electronic wavefunction. However, in practice, this is not the case, and it has proven extremely difficult to improve the fixed-node surface towards the exact one. The Fermion sign problem is thus recognised as one of the most important unsolved problems of computational theoretical physics and chemistry. A key question is whether the exact nodal hypersurface can *emerge* during the course of a simulation. Such a simulation would therefore not require any *a priori* information regarding the nodes of the exact wavefunction. In this paper we describe a new quantum Monte Carlo (QMC) method in which this highly desirable property is shown to arise, for systems described with basis-sets commonly used in quantum chemistry. Our method is shown to converge on to the Full Configuration-Interaction (FCI) solution, i.e. the exact wavefunction and energy for the basis-set under consideration [8–11]. In addition to reproducing existing FCI calculations, we have used this method to predict the FCI energies of several molecules which have not been reported to date. We also report a benchmark study of the Ne atom (correlating all electrons) in several basis sets up to cc-pCVQZ, which has a FCI basis exceeding 10^{14} determinants. That is over 10,000 times larger than the largest FCI calculation reported to date [12]. Our method unifies QMC and FCI in a way which profoundly extends the scope of both techniques.

There are three ingredients to the new method, which take elements from DMC and FCI: (i) In common with DMC, we perform a long-time integration of the imaginary-time Schrödinger equation; however, in contrast to DMC, this is achieved in a space of Slater determinants. In addition, the propagation step in our algorithm differs from DMC in that it consists purely of population dynamics (i.e. walker birth and death processes). There are no diffusive moves as such. (ii) In

*Electronic address: asa10@cam.ac.uk

common with DMC, the instantaneous wavefunction is represented using "walkers", rather than amplitude coefficients, the latter being the case in FCI. The representation using "walkers" enables us to describe stochastically the FCI wavefunction *without* storing all amplitudes simultaneously. (iii) Each walker carries a positive or negative sign. The key ingredient in our algorithm is walker *annihilation*, in which pairs of walkers of opposite sign which coincide on the same determinant are removed from the simulation. We show that an algorithm based on these three ideas converges on the *exact* solution to the Hamiltonian expressed in the basis, which owing to the Fermion antisymmetry of the Slater determinant basis, will be the lowest Fermionic solution available in this basis. It should be noted that working in Slater determinant space does not circumvent the sign problem, even though it prevents convergence to a Bosonic solution. This is because the FCI wavefunction does not (in general) have strictly non-negative amplitudes in this space. We note that walker annihilation has been previously suggested in the context of nodal release GFMC and DMC [13–16]. However, as noted in these studies, only low-dimensional phase-spaces could be treated owing to the formidable difficulties in achieving cancellation in continuum spaces. Instead, by working in discrete Slater determinant spaces, we find that walker annihilation proves an essential and effective component of a Fermion Monte Carlo method.

There have been several proposals to simulate Fermionic systems in determinantal spaces using auxilliary-field Monte Carlo(AFMC) [17–22]. However, in their exact formulation, auxillary-field methods suffer from exponentially large statistical noise in the limit of large imaginary-time propagation, which is necessary to project out the ground state. In order to stabilise this problem, a phaseless-approximation AFMC method has been proposed [21], with promising energies (i.e. to within a few milli Hartrees of known FCI energies) when applied to molecular systems [22]. However, like the fixed-node approximation, the phaseless approximation is an uncontrolled approximation which may be difficult to improve upon.

Below we give motivation and elaborate on our algorithm, before showing applications of the method to several real physical systems.

MOTIVATION AND DERIVATION OF THE ALGORITHM

In the FCI method, we seek a wavefunction Ψ_0 which satisfies the time-independent Schrödinger equation, $H\Psi_0 = E_0\Psi_0$, as a variationally optimised linear combination of Slater determinants $\{|D_\mathbf{i}\rangle\}$. These are antisymmetric functions in which N orbitals are chosen out of $2M$ spin-orbitals $\{\phi_1, \phi_2, ..., \phi_{2M}\}$ and constructed as follows:

$$|D_\mathbf{i}\rangle \equiv |D_{n_1,n_2,..,n_N}\rangle = a^\dagger_{n_1} a^\dagger_{n_2} .... a^\dagger_{n_N} | \ \rangle \quad (1)$$

$$= \frac{1}{\sqrt{N!}} |\phi_{n_1}\phi_{n_2} ..... \phi_{n_N}|, \quad n_1 < n_2 < ... < n_N \quad (2)$$

$a^\dagger_i$ is a Fermion creation operator for spin-orbital ϕ_i. In this work, the orbitals used are real, canonical Hartree-Fock orbitals. The size of the Slater determinant space (N_{FCI}) is on the order of $\binom{M}{N/2}^2$ for a spin-unpolarised system, a number that grows factorially with M and N (although symmetry restrictions for some systems can reduce this by up to an order of magnitude). The FCI wavefunction is expressed as:

$$\Psi_0^{\mathrm{FCI}} = \sum_\mathbf{i} C_\mathbf{i} |D_\mathbf{i}\rangle \quad (3)$$

where the CI coefficients $\{C_\mathbf{i}\}$ satisfy an eigenvector problem:

$$\sum_\mathbf{j} \langle D_\mathbf{i}|H|D_\mathbf{j}\rangle C_\mathbf{j} = E_0^{\mathrm{FCI}} C_\mathbf{i}. \quad (4)$$

E_0^{FCI} is the lowest energy solution available in this basis, an upper bound to the exact energy. Owing to the fact that the off-diagonal Hamiltonian matrix elements are not all of the same sign, the CI coefficients can be positive or negative. The "sign" structure of the FCI wavefunction is given by a vector whose components are $\mathrm{sign}(C_\mathbf{i})$ or 0 if $C_\mathbf{i} = 0$. A trial wavefunction has the correct sign structure only if the sign of every component matches those of this vector (up to an overall sign, since $-\Psi_0^{\mathrm{FCI}}$ is an equally valid solution with the same energy). In practice, it turns out to be impossible to predict $\mathrm{sign}(C_\mathbf{i})$ without a knowledge of $C_\mathbf{i}$, and this, in essence, is the manifestation of the Fermion sign problem in the discrete Slater determinant space. In FCI, the CI coefficients are obtained via a (non-stochastic) iterative diagonalisation method. Whilst it is recognised that the FCI method is the most robust method to treat electron correlation, its scope is greatly limited by the prohibitive computational requirements (especially storage) of such iterative diagonalisation in the full space of Slater determinants. For example the Davidson [23] method or the preconditioned conjugate-gradient method [24] requires at least two vectors of length equal to the FCI space to be stored, and often many more. The largest molecular FCI calculation to date is the N_2 molecule [12], and has $\approx 10^{10}$ determinants in D_{2h} symmetry. As Nicholas Handy has observed [25], "unless something unexpected happens, it is unlikely that size will be much exceeded". Inspired by the DMC method, rather than attempt a direct diagonalisation of the FCI Hamiltonian, we propose to simulate the imaginary-time Schrödinger equation by performing a stochastic population dynamics on an evolving set of "walkers" which live and propagate in Slater determinant

space. However, we do not impose any prior knowledge on the signs of the Slater determinants.

The imaginary-time Schrödinger equation provides the fundamental starting point of all "projector" techniques. It states that the imaginary-time derivative of any wavefunction Ψ is given simply by the Hamiltonian acting on $-\Psi$, namely:

$$\frac{\partial \Psi}{\partial \tau} = -H\Psi \tag{5}$$

Given a starting wavefunction $\Psi(\tau = 0)$, the wavefunction $\Psi(\tau)$ for arbitrary τ is proportional to:

$$\Psi(\tau) \propto e^{-\tau H}\Psi(\tau = 0) \tag{6}$$

In order to project onto the Fermionic ground-state, $\Psi(\tau = 0)$ must chosen to be a fully-antisymmetric function with an overlap with the ground state. Usually a single determinant function with the correct symmetry, such as the Hartree-Fock determinant $D_{\mathbf{0}}$, suffices. The long-time limit of Eq. (6) then projects out the components of $D_{\mathbf{0}}$ on excited-states, leaving the Fermionic ground-state, Ψ_0:

$$\Psi_0 = \lim_{\tau\to\infty} e^{-\tau(H-E_0)}D_{\mathbf{0}} \tag{7}$$

where the constant of proportionality $e^{+\tau E_0}$ has been introduced to keep the ground-state contribution from decaying to zero. The aspiration of all Monte Carlo projector methods is to realise this long-time limit by performing a stochastic integration of Eq. (5). In normal Diffusion Monte Carlo, however, the integration is not performed on an antisymmetrised space, and therefore admits solutions with Bosonic character, which rapidly grow to dominate unless constrained in some manner, e.g. by the fixed-node approximation. An integration in a pure antisymmetric subspace would not suffer from this effect, and this is may be a help in reducing the sign problem for Fermion systems [26, 27]. Our first aim, therefore, is to develop an analogue to Eq. (5) in a Slater determinant basis, in a form which can be integrated stochastically.

To this end, let us define a matrix K, whose elements are the matrix elements between Slater determinants of the Hamiltonian, with the Hartree-Fock energy E_{HF} subtracted from the diagonal elements:

$$K_{\mathbf{ij}} \equiv \langle D_{\mathbf{i}}|K|D_{\mathbf{j}}\rangle = \langle D_{\mathbf{i}}|H|D_{\mathbf{j}}\rangle - E_{\mathrm{HF}}\delta_{\mathbf{ij}} \tag{8}$$

The diagonal matrix-elements of K are therefore all positive (or zero), whilst the off-diagonal elements are simply the off-diagonal matrix elements of the Hamiltonian. With this definition, the lowest energy eigenvalue of K is $E_0^{\mathrm{FCI}} - E_{\mathrm{HF}}$, which is the correlation energy, E_{corr} for the problem. The matrix elements of H can be computed in terms of the one-electron and two-electron integrals using standard methods, which are reviewed in Appendix A.

Writing

$$\Psi(\tau) = \sum_{\mathbf{i}} C_{\mathbf{i}}(\tau)|D_{\mathbf{i}}\rangle \tag{9}$$

and substituting into Eq. (5), we obtain a set of coupled linear first-order differential equations for the CI coefficients in terms of the K matrix:

$$-\frac{dC_{\mathbf{i}}}{d\tau} = \sum_{\mathbf{j}}(K_{\mathbf{ij}} - S\delta_{\mathbf{ij}})C_{\mathbf{j}} \tag{10}$$

$$= (K_{\mathbf{ii}} - S)C_{\mathbf{i}} + \sum_{\mathbf{j}\neq\mathbf{i}} K_{\mathbf{ij}}C_{\mathbf{j}} \tag{11}$$

where an arbitrary "energy shift", S, has been introduced into the diagonal terms, whose role will be population control, to be discussed later. It is evident that if we instantaneously have a vector $\mathbf{C}$ of amplitudes whose components $C_{\mathbf{i}}$ satisfy

$$\sum_{\mathbf{j}} K_{\mathbf{ij}}C_{\mathbf{j}} = SC_{\mathbf{i}} \tag{12}$$

then $d\mathbf{C}/d\tau = 0$, making this vector stationary. In addition, by virtue of Eq. (12), this vector is an eigenstate of the K matrix (and hence of the H), with eigenvalue S. It follows that if S equals the correlation energy E_{corr}, then the stationary state is the ground state of the H matrix. Furthermore, if we start from an arbitrary set of $C_{\mathbf{i}}$ amplitudes (which does not satisfy Eq. (12)) and integrate the above set of coupled differential equations, the long-time solution leads to the ground-state eigenvector. This can be proven by decomposing the starting vector into the eigenstates of H, and noting that the components on the excited states would decay exponentially with time, leaving the ground-state eigenvector.

A direct numerical integration of Eq. (11) requires the full set of $C_{\mathbf{i}}$ coefficients to be available at each time step, which is prohibitive. Instead, in the spirit of DMC, let us consider a population of N_w walkers; each walker α is located on a determinant $\mathbf{i}_\alpha$, and has a sign $s_\alpha = \pm 1$. We now *define* the $C_{\mathbf{i}}$ amplitude on determinant $|D_{\mathbf{i}}\rangle$ to be proportional to the *signed* sum of walkers ($N_{\mathbf{i}}$) ,

$$C_{\mathbf{i}} \propto N_{\mathbf{i}} = \sum_{\alpha} s_\alpha \delta_{\mathbf{i},\mathbf{i}_\alpha} \tag{13}$$

($\delta_{\mathbf{i},\mathbf{i}_\alpha}$ is the discrete Kronecker delta, and equals one if $\mathbf{i}_\alpha = \mathbf{i}$, and is otherwise zero). According to Eq. (13), $N_{\mathbf{i}}$ can be positive or negative. However, the total number of walkers N_w is defined to be the sum over the absolute values of the $N_{\mathbf{i}}$:

$$N_w = \sum_{\mathbf{i}} |N_{\mathbf{i}}| \tag{14}$$

and is always positive.

A POPULATION DYNAMICS ALGORITHM

We introduce a population dynamics algorithm which simulates the set of coupled differential equations in Eq. (11). The algorithm consists of three steps performed at each timestep whose length is $\delta\tau$:

(i) *The spawning step*: for each walker α (located on $D_{\mathbf{i}_\alpha}$), we select a (coupled) determinant $D_{\mathbf{j}}$ with normalised probability $p_{\rm gen}(\mathbf{j}|\mathbf{i}_\alpha)$ and attempt to spawn a child there with probability

$$p_s(\mathbf{j}|\mathbf{i}_\alpha) = \frac{\delta\tau|K_{\mathbf{i}_\alpha\mathbf{j}}|}{p_{\rm gen}(\mathbf{j}|\mathbf{i}_\alpha)} \tag{15}$$

If a spawning event is successful, (i.e. if p_s exceeds a uniformly chosen random number between 0 and 1), then the sign of the child is determined by the sign of $K_{\mathbf{i}_\alpha\mathbf{j}}$ and the sign of the parent: it is the same sign as the parent if $K_{\mathbf{i}_\alpha\mathbf{j}} < 0$, and opposite to the parent otherwise. Our method to compute the generation probabilities $p_{\rm gen}$ is given in Appendix B. If $p_s > 1$, then multiple copies of walkers are spawned on $\mathbf{j}$ (namely with probability 1, $\lfloor p_s \rfloor$ walkers are spawned, and with probability $p_s - \lfloor p_s \rfloor$ an additional walker is spawned).

In general, the number of newly-spawned walkers (N_s) is much smaller than number of parent walkers (N_w), since the time step $\delta\tau$ is such that the probability to spawn is quite low. Typically we find $N_s \sim 10^{-4} N_w$.

(ii) *The diagonal death/cloning step*: for each (parent) walker compute

$$p_d(\mathbf{i}_\alpha) = \delta\tau(K_{\mathbf{i}_\alpha\mathbf{i}_\alpha} - S) \tag{16}$$

If $p_d > 0$, the walker dies with probability p_d, and if $p_d < 0$ the walker is cloned with probability $|p_d|$. The death event happens immediately, and such a parent does not participate in the following (annihilation) step to be described shortly. Cloning events are quite rare, and only occur for $S > 0$, and even then only on determinants for which $\langle D_{\mathbf{i}}|K|D_{\mathbf{i}}\rangle < S$. In simulations where we desire to grow the number of walkers rapidly, a positive value of S is adopted, and this can lead to cloning events. However, more often, the value of S is negative (as it tries to match the correlation energy), and in such cases there can be no cloning events at all.

Both the spawning step and the diagonal death step can be done without reference to other walkers. Therefore, these two steps are "embarrassingly parallel", and can be performed without communication overhead on a parallel machine.

(iii) *The annihilation step*: In this (final) part of the algorithm, we run over all (newly-spawned, cloned and surviving parent) walkers, and annihilate pairs of walkers of opposite sign which are found to be on the same determinant. Each time an annihilation event occurs, the corresponding pair is removed from the list of walkers, and the total number of walkers N_w reduced by two. By keeping sorted lists of walkers, it is possible to do the search for possible annihilation events using binary searches, with the result that the annihilation step can be done with $\mathcal{O}\left[N_s \ln(N_s N_w)\right]$ effort.

At the end of the annihilation step, the lists of surviving newly-spawned walkers and parents are merged. The merged list remains sorted, and becomes the main list of walkers for the next timestep[28]. The annihilation can achieved in a memory efficient manner, keeping only one copy of the main list of walkers. In a one-electron basis with $2M$ orbitals, [29], in discrete spaces (with a finite spread of eigenvalues), repeated application of each walker needs $\lceil 2M/32 \rceil + 1$ four-byte integers of storage [30] to encode the occupation number information of the determinant it lives on, as well as the sign of the walker. For a simulation with N_w walkers, therefore, the main walker list requires $4N_w(\lceil 2M/32 \rceil + 1)$ bytes of RAM. As will be seen in the applications, the number of walkers required to achieve convergence is typically smaller than the size of the FCI space, often by a significant factor. Therefore the memory requirements of this algorithm are less severe than that of an conventional FCI calculation.

It should be noted that with the above algorithm all (symmetry-allowed) determinants in the FCI space are accessible, since no restrictions are placed on the spawning step as to which determinants can be generated: given a walker on some determinant, any determinant connected to it can be chosen to be spawned upon. It follows that the entire space of (symmetry-related) determinants can eventually be reached starting from any one determinant. The fact that the entire space is accessible enables the algorithm to converge onto the FCI wavefunction, which in general has non-zero amplitudes on all such determinants. Of course, it is also a straightforward matter to impose a truncation in the CI space, for example, by excitation level, by simply not accepting any spawning events at determinants beyond the specified truncation level. In this way, one can perform calculations equivalent to truncated CI, such as CISDTQ, etc. In large calculations, this approach can be used to prepare equilibrated or near-equilibrated ensembles of walkers at truncated excitation levels, before attempting a simulation in the full CI space.

A second point to note is that in the current algorithm each walker attempts to spawn only once per timestep. It is possible to construct a modified algorithm in which each walker at each time step systematically attempts to spawn at *all* connected determinants (but with a different spawning probability, given by $p_s(\mathbf{j}|\mathbf{i}_\alpha) = \delta\tau|K_{\mathbf{i}_\alpha\mathbf{j}}|$). Obviously, with this modification, the computational cost per time step is much greater, although it might be expected that a faster convergence is achieved. In tests, however, this version of the spawning algorithm proved to be less efficient overall than the one we outlined above.

A further point to note concerns time-step error. The above algorithm results in changes in populations on each determinant proportional to $\delta\tau$. The underlying propagator is $G = 1 - \delta\tau(H - S)$. As has been pointed out in [29], in discrete spaces (with a finite spread of eigenvalues), repeated application of such propagators starting from any initial state converges onto the exact ground-state as long as $\delta\tau \leq 2/(E_{max} - S)$, where E_{max} is the largest eigenvalue of H and $S \approx E_0$. Longer timesteps lead to propagators in which excited-state contributions do not decay with time. In a non-stochastic implementation of this propagator, convergence to the ground-state is therefore guaranteed as long as the time-step does not exceeed this upper bound value. It is possible to estimate E_{max} as the energy of the most highly excited determinant available in a given one-electron basis, with E_{max} increasing with the size of the basis.

In practice, for the present stochastic algorithm, we find that there is an additional constraint on the time-step which makes it to be smaller than the above value, namely that is it desireable that p_s should not greatly exceed unity, since this results in large walker blooms (i.e. multiple copies of walkers being suddenly spawned on a determinant), which can lead to inefficient sampling. On the other hand, a small timestep leads to slow evolution and requires proportionately longer simulations, which is also inefficient. In practice, for the systems under consideration in this study, we find that that values for $\delta\tau$ in the range $10^{-4} - 10^{-3}$ a.u. (with the smaller values being used for larger basis sets), provides a reasonable balance between these two sources of inefficiency, and that we have explicitly verified that the correlation energies we compute are insensitive of the value of the time-step as long as they are in this range. In the remaining sections we discuss some additional features of our simulation methodology.

Constant S and Constant N_w simulations

If the shift S is kept constant at a value above $E_{\rm corr}$, then in general the number of walkers increases exponentially. It is therefore desirable to be able to do "constant" N_w simulation. This can be achieved by periodically (every A steps) adjusting the shift S according to the following simple prescription

$$S(\tau) = S(\tau - A\delta\tau) - \frac{\zeta}{A\delta\tau} \ln \frac{N_w(\tau)}{N_w(\tau - A\delta\tau)} \qquad (17)$$

where $N_w(\tau)$ is the total number of walkers on timestep τ, and ζ is a damping parameter, i.e. if N_w has grown over the current update cycle, the shift is reduced (made more negative), whilst if N_w has decreased, the shift is increased (made more positive). Thus in "constant" N_w mode, the shift S varies to attempt to keep walker number constant. The damping coefficient ζ, prevents undesirable large fluctuations in S. In constant S mode the number of walkers varies, typically increasing if $S > E_{\rm corr}$. If the distribution of walkers is correct (i.e. distributed according to the ground-state amplitude), then the value of S which maintains the number of walkers is equal to the correlation energy of the problem. Conversely, in constant N_w mode, the shift will fluctuate about the correct correlation energy when the distribution of walkers is correct. In the examples we report below, we chose $A = 5 - 10$, and $\zeta = 0.05 - 0.1$.

Simulation procedure and the projected energy

The simulations are started by placing a single walker on the HF determinant. The simulation is then run for a start-up period, in which the shift S is kept at a constant value (of zero or small positive number). This results in the number of walkers initially growing very rapidly. When a desired number of walkers is reached, we start to adjust the shift S, as given by Eq. (17). It is observed that the value of the shift initially falls, and then stabilises around a value about which it oscillates. If the number of walkers is above a critical number (to be discussed later), this value is the correlation energy of the system. It is desirable to have another measure of the correlation energy. This is provided by the following projection:

$$E(\tau) = \frac{\langle D_{\mathbf{0}}|He^{-\tau H}|D_{\mathbf{0}}\rangle}{\langle D_{\mathbf{0}}|e^{-\tau H}|D_{\mathbf{0}}\rangle} \qquad (18)$$

$$= E_{\rm HF} + \sum_{\mathbf{j}\neq\mathbf{0}} \langle D_{\mathbf{j}}|H|D_{\mathbf{0}}\rangle \frac{C_{\mathbf{j}}(\tau)}{C_{\mathbf{0}}(\tau)} \qquad (19)$$

$$= E_{\rm HF} + \sum_{\mathbf{j}\neq\mathbf{0}} \langle D_{\mathbf{j}}|H|D_{\mathbf{0}}\rangle \frac{N_{\mathbf{j}}(\tau)}{N_{\mathbf{0}}(\tau)} \qquad (20)$$

where $N_{\mathbf{j}}(\tau)$ is the number of walkers (defined by Eq. (13)) on $D_{\mathbf{j}}$ at timestep τ, and N_0 is this number on $D_{\mathbf{0}}$ (which we assume to be the Hartree-Fock determinant). In addition, only the singles and double excitations are connected to $D_{\mathbf{0}}$ (the singles being connected only in case of open-shell systems), it is only the populations at these excitations that need to be evaluated for the purposes of

computing the energy, i.e.:

$$E(\tau) = E_{\mathrm{HF}} + \sum_{\mathbf{j}\in\{\text{Singles, Doubles}\}} \langle D_{\mathbf{j}}|H|D_{\mathbf{0}}\rangle \frac{N_{\mathbf{j}}(\tau)}{N_{\mathbf{0}}(\tau)} \quad (21)$$

In the long-time limit, it is evident that $E(\tau)$ converges onto the exact ground-state energy:

$$E_0 = \lim_{\tau\to\infty} E(\tau) \quad (22)$$

as long as the ratio $N_{\mathbf{j}}/N_{\mathbf{0}}$ equals $C_{\mathbf{j}}/C_{\mathbf{0}}$ of the FCI wavefunction.

The two measures of the correlation energy, i.e. S and $E(\tau)$, are to a large extent independent of each other, since S depends on the total number of walkers, whereas $E(\tau)$ depends only on the population of walkers on the Hartree-Fock determinant and the singles and double excitations, which typically consititute only a small fraction of the total walker population. In the limit of large walker populations, the variations in S and $E(\tau)$ become largely uncorrelated.

AN ILLUSTRATIVE EXAMPLE

Let us consider stretched N_2, which is a strongly multiconfigurational system. For the purposes of illustration of how the exact wavefunction emerges during the course of the simulation, we consider here a very small basis, with 379 Slater determinants, which allows a convenient visualisation of the evolution of the walker distribution as the simulation proceeds. The behaviour of the shift, and of the projected energy, are shown in Fig. 1. The instantaneous energy for each timestep (Eq. 20) converges on the FCI energy within 10,000 steps. In this simulation, we started varying the shift at step 10,000, and the shift very rapidly decreases from zero towards the FCI energy, and stabilises at the correct value. The two independent measures of the correlation energy agree to better than 1 mH. In real applications, where the exact FCI energy is not known, the co-incidence of the these two measures provides a strong degree of confidence in the obtained correlation energies. Equally remarkable is the convergence of the walker distribution onto the exact wavefunction: An animation of the time evolution of this simulation (as well as other systems) is also provided as a movie held on our website [31]. In the third panel of Fig. 1 we also show the behaviour of the simulation if the annihilation step is removed. In this case, the shift does not stabilise at the correct value, and the instantaneous energy fluctuates wildly. Without annihilation, the walker distribution cannot converge onto to correct ground-state distribution; there is no "interaction" between the positive and negative walkers, and in fact the distribution of the positive walkers among the determinants increasingly matches the distribution of negative walkers as the simulation proceeds. In other words, in common with the classic sign problem, the signal to noise ratio decreases exponentially quickly. In addition, because the shift goes to the wrong (and much too negative) value, the rate of death processes becomes very large, leading to indiscriminate death of walkers, essentially irrespective of which determinant they occupy. The walker distribution is not able to build a signal on the determinants with significant weight. This example shows the crucial role played by the annihilation step. In fact the annihilation step allows a symmetry-breaking to occur, in which one of the two allowed solutions, either $+\Psi$ or $-\Psi$, is settled upon. Without providing any information regarding the sign structure of the FCI wavefunction, our algorithm is capable of settling down onto the correct wavefunction. In a real sense, the many-body wavefunction has *emerged* during the course of the simulation. The breaking of the $\pm\Psi$ symmetry is a necessary condition for this emergence. An analogy can be made with phase-transitions in classical spin-systems. The ordered phase (e.g. ferromagnetic or antiferromagnetic) of a spin-system is is degenerate with respect to overall sign-change of all spins. For the system to acquire net magnetisation (or staggered magentisation) below the critical point, the symmetry must be spontaneously broken, and the ordered phase, which has a symmetry lower than that present in the Hamiltonian can then emerge.

This example demonstrates that the method in principle works exactly. The question now is: it is capable of working for much larger FCI spaces? To answer this question, we have studied a number of molecular systems (Ne, H_2O, C_2, $N_2\ldots$) in a variety of basis sets, to which we now turn.

RESULTS

We next consider a range of molecules, treated in more realistic Dunning basis sets[32]. We also perform a benchmark study of the all-electron Ne atom in several families of basis sets, namely cc-pVXZ, aug-cc-pVXZ and cc-pCVXZ (X=D,T,Q). The size of spaces considered range from $\sim 10^6$ to $\sim 10^{14}$.

In spaces of such size, the pattern of walker growth generally changes. Keeping the shift fixed, initially the walker number grows exponentially. However, at a system specific value, this exponential growth attenuates, apparently suddenly, and the number of walkers hits a plateau. After a period of essentially zero growth in overall walker number, the growth picks up again, after which point we start adjusting the shift. An example of this is shown for the all-electron water molecule treated with a cc-pVDZ basis in Fig. 2.

During this plateau phase, the rate of walker birth is exactly matched by the combined rate of walker death and walker annihilation. Before reaching the plateau, the rate of annihilation is small and therefore the overall rate

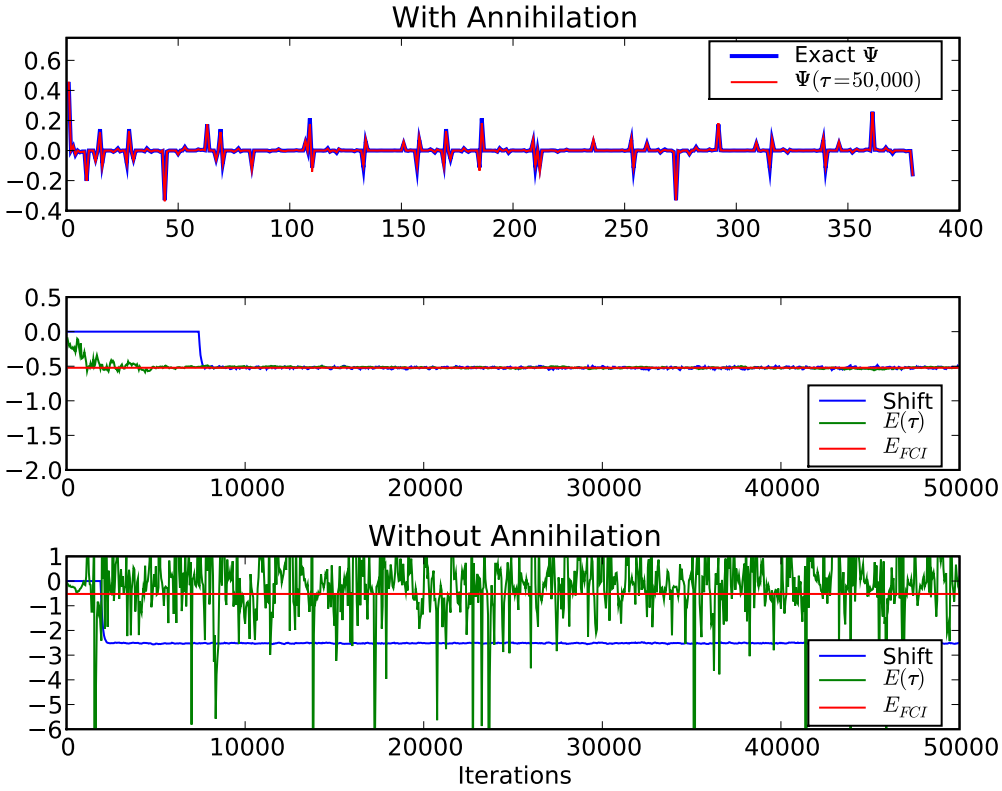

FIG. 1: Comparison of sampling of determinants compared to full diagonalisation result for a stretched N_2 molecule in a space of 379 determinants. Ψ indicates the normalised walker number on each of the determinants after 50,000 iterations, comparing it to the exact wavefunction. The shift and E_0 values for each iteration are shown in the lower plot.

of walker growth is positive. At the plateau, sufficient numbers of walkers begin to annihilate each other, and this leads to the attenuation of walker number growth. Very significantly, owing to walker annihilation, *we observe that during the plateau phase, the sign structure of the FCI wavefunction is converged upon*, in every case we have studied. As long as we have a larger number of walkers in the system than the number defined by the plateau position, N_c, then when we allow the shift to vary (in constant N_w mode), it converges onto the correlation energy. This N_c value is therefore an important, system-dependent parameter which indicates the required sampling of the space. If this sampling can be achieved, then the correct correlation energy for the system can be confidently converged upon to the desired accuracy. This N_c value in all cases is smaller than the size of the complete space, often substantially so. We can therefore describe a parameter f_c which gives us the relative number of walkers required, compared to the number of determinants in the full space ($f_c = \frac{N_c}{N_{\mathrm{FCI}}}$). f_c is one measure of the difficulty the method has in achieving convergence. Where f_c is small (or even zero), convergence can be achieved with a relatively small number of walkers.

It turns out that the value of N_c (for a given system) is remarkably insensitive to the initial conditions, and method of equilibration, of the walkers. Thus, we have found that we need the same number of walkers to attain the plateau if we start the simulation with one walker and let the number grow rapidly throughout the entire FCI space, or if we gradually enlarge the space, e.g. via excitation level. This indicates that N_c is an intrinsic parameter which characterises the system, for a given one-electron basis.

In Fig. 3 we show the convergence of the averaged projected energy $\langle E(\tau)\rangle$ as the simulation proceeds, in the case of the water molecule. Also included, is the absolute error $|\langle E(\tau)\rangle - E_{\mathrm{FCI}}|$ from the known FCI energy. It can be seen that there is an exponential convergence to an accuracy of under 0.1mH of the energy. To give an idea of the computational cost of the method, the CPU time taken to converge all-electron cc-pVDZ water to this accuracy with our current code is three hours using a quad-core (2.66GHz) PC.

Apart from the water molecule, we have tested our al-

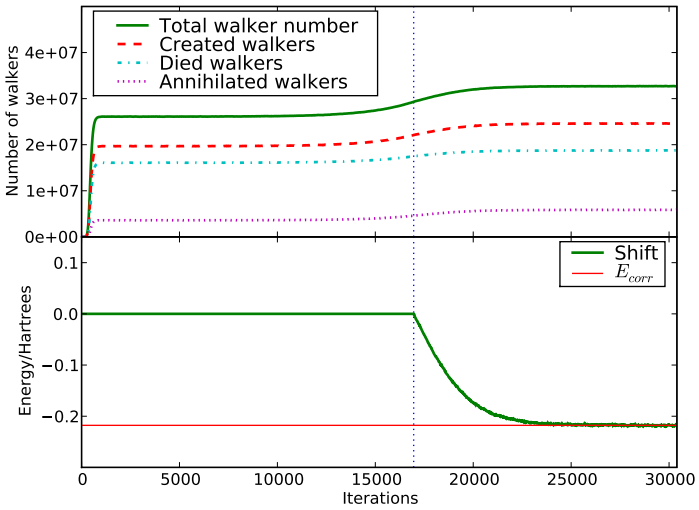

FIG. 2: A typical trend of walker growth for H_2O in a cc-pVDZ basis set in constant shift mode ($S = 0$) until iteration 16,970 , showing the appearance of a well-defined plateau in N_w, in this case at $N_c = 26 \times 10^6$ walkers. For this system, $N_{\mathrm{FCI}} = 451 \times 10^6$. The number of walkers at the plateau is therefore about 6% of the FCI space. Also shown on the plot are the number of walkers created, died and annihilated per $A = 10$ iterations. It is evident that at the plateau, the combined rate of death and annihilation matches the birth rate from the spawning. The plateau gradually gives way to a growth phase in walker number, which increases exponentially. Once in the growth phase, the shift is then allowed to vary according to Eq. (17), (in this example on iteration number 16,970). The number of walkers then rapidly stabilises. The lower plot shows that the value for the shift which stabilises the walker growth is exactly the correlation energy for the system.

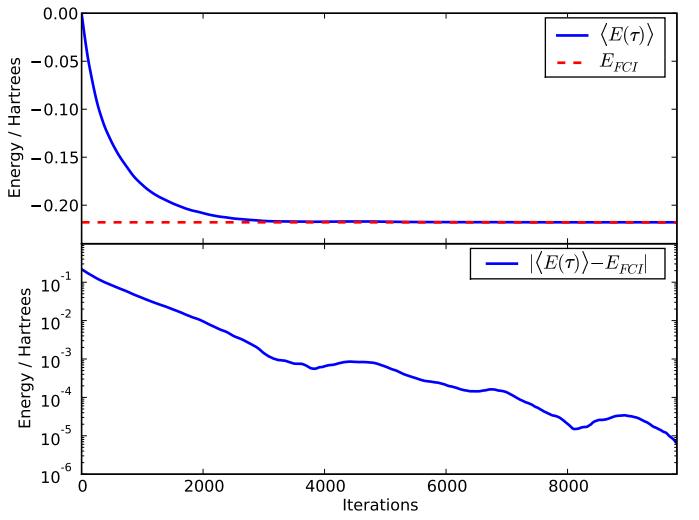

FIG. 3: Convergence of the energy estimator averaged over all previous iterations, showing an exponential convergence to the E_{FCI} for a cc-pVDZ water system at equilibrium geometry. The lower plot shows the absolute difference between the energy and the exact energy on a logarithmic scale. $\delta\tau = 1 \times 10^{-3}$ a.u. In this example, the simulation was started by placing a single walker at the Hartree-Fock determinant and allowing the population to grow throughout the entire space until the post-plateau growth phase has been reached.

gorithm in other systems for which FCI results are available for comparison, shown in Table I, namely Ne atom, C_2 and N_2 at equilibrium and at stretched geometry. As can be seen, the computed energies are to within 1 mH of the published FCI results for all systems considered, confirming the ability of the method to reproduce the FCI energies across a broad range of systems.

There is considerable variation in f_c among these systems. For the Ne atom, and all-electron H_2O, $f_c \ll 0.1$, whereas for C_2 and N_2, $f_c \sim 0.5$, i.e. comparable though still smaller than the FCI space. The behaviour of the two N_2 systems is particularly instructive. As the N_2 bond is stretched from its equilibrium geometry, the description of the problem changes from one which is essentially single-reference, to one which is highly multi-reference with significant contributions up to hextuple excitations. This makes an accurate description of the binding curve of nitrogen particularly challenging for many methods. Surprisingly, the results from these two systems show a similar f_c, indicating that the ease with which the method achieves convergence is not simply related to the dominance of the Hartree-Fock determinant, and that the number of significant determinants is not the key factor in its efficiency. This feature can hopefully be exploited in the study of multiconfigurational systems. The fact that in all cases f_c is less than one, sometimes substantially so, means that one needs fewer walkers than determinants to achieve converge the calculation. . At any one time step, only a fraction of the space is occupied, but as long as f_c has been exceeded, averaged over time each determinant is correctly sampled according to its amplitude. We have exploited this property to study systems for which the FCI space is extremely large.

In Table II, we report the energies of several molecules, including CO, CH_4, O_2 and NaH – for which we have not been able to find published FCI energies, and therefore serve as predictions which could be verified by explicit FCI calculations. Included in this list are the open-shell systems CN($^2\Sigma^+$) and O_2($^3\Sigma_g^-$) . As can be seen from the table, the size of FCI spaces for these systems range from a few million to over 10^{11}. Particularly striking are the variations in f_c. The worst case is for CH_4, where we find an $f_c = 0.898$, whilst for NaH $f_c = 3 \times 10^{-4}$. Remarkably, in this case, only 64 million walkers were needed to converge the energy for a space of 205×10^9 determinants. Results from coupled-cluster with perturbative triple excitations, CCSD(T), have been included for comparison [38], and as can be seen our correlation energies are generally slightly larger (with the notable exception of CH_4), although since CCSD(T) is not variational, not too much can be read into this result. We do note that the agreement between the two is generally better for closed-shell systems than open-shell ones.

An important question with the present method is how the plateau varies for a given system as the basis set is improved, since we typically want to do calculations in substantial basis sets, which improve accuracy and allow extrapolation to the complete basis-set limit [39]. In Table III, we present a study of the Ne atom in various basis-sets, to investigate the fraction f_c required to achieve the plateau. Several trends are discernible. Within a given type of basis set, the value of f_c is fairly constant, but there are marked variations between different families. Thus cc-pVXZ (denoted VXZ in Table III, where X=D,T,Q) family have $f_c \approx 5 \times 10^{-4}$, whereas the aug-cc-pVXZ (AVXZ) family have a larger $f_c \approx 1.5 \times 10^{-3}$, whilst the cc-pCVXZ (CVXZ) have a very favourable $f_c \approx 2 \times 10^{-5}$. The latter are also give the most accurate energies. Clearly inclusion of core electrons with CVXZ basis sets is very favourable for the present method. Extrapolation to the complete basis-set limit [39] yields a correlation energy of 393.0 mH, which is within 2.5 mH of the "exact" non-relativistic result [40]. Such extrapolated energies are of course not variational, and are known to yield errors of about 5 mH, consistent with what is observed here.

These results illustrate that the memory requirements of the method are substantially lower than a conventional FCI calculation. For example, the all-electron Ne atom in a cc-CVQZ basis required $N_c = 2.2 \times 10^9$ walkers, which amounts to 62 Gbytes of RAM. By contrast, a FCI calculation of the same system would require a minimum of two vectors of length $N_{\rm FCI} = 131 \times 10^{12}$, i.e. over 2×10^6 Gbytes, a vastly greater amount of storage.

Since in general N_c scales linearly with N_{FCI}, this implies that the present method has exponential scaling with N and M. The method, therefore, does not constitute a "solution to the sign problem", for which a polynomial scaling is required. It is perhaps best thought of as an alternative method to FCI, with a smaller prefactor (proportional to f_c), which in some cases is substantially so. It remains to be seen if variations on the present algorithm can be found in which either the prefactor can be further supressed, or even better, if the scaling can be improved upon.

CONCLUSIONS

We have described a new Quantum Monte Carlo algorithm in Slater determinant space, which we show is able to converge on the FCI energy of the system under consideration without any *a priori* information, as long as a system-dependent critical number of walkers is exceeded. The exact FCI wavefunctions emerge spontaneously once this critical number has been reached. Walker annihilation is shown to play a key role in this process. The computational requirements of the method are largely determined by the value of this critical number. If it can be reduced for a given system, e.g. by a judicious choice of orbitals obtained via orthonormal transformations of the canonical Hartree-Fock orbitals, then the computa-

System	*(N,M)*	$N_{\rm FCI}/10^6$	$N_c/10^6$	f_c	$E_{\rm total}$	$E_{\rm FCI}$	Reference
Ne: aug-cc-pVDZ	(8,22)	6.69	0.21	0.031	-128.7094*9*	-128.709476	[33]
C_2: cc-pVDZ	(8,26)	27.9	15.0	0.538	-75.729*9*	-75.729853	[34]
H_2O: cc-pVDZ	(10,24)	451	26	0.058	-76.2418*6*	-76.241860	[35]
N_2-eqm: cc-pVDZ	(10,26)	541	270	0.499	-109.2764*9*	-109.276527	[33]
N_2-stretched: cc-pVDZ	(10,26)	541	345	0.637	-108.966*9*	-108.96695	[36]

TABLE I: Results for systems with FCI comparisons. The geometries for the N_2 molecule were eqm: $2.068a_0$, stretched: $4.2a_0$, and C_2: 1.27273Å. The geometry for the water molecule was taken from [35]. The working space includes all point group symmetry of the molecule from D_{2h} or the largest available subset thereof. All systems had core electrons frozen apart from H_2O. $N_{\rm FCI}$ is the size of the FCI space in the D_{2h} point group (C_{2v} for H_2O). The digit in italics for $E_{\rm total}$, represents the first uncertain digit. N_c is the number of walkers required to achieve the plateau. $f_c = N_c/N_{\rm FCI}$.

System	*(N,M)*	$N_{\rm FCI}/10^6$	$N_c/10^6$	f_c	$E_{\rm total}$	$E_{\rm CCSD(T)}$
Be: cc-V5Z	(4,91)	2.11	0	0	-14.64638(2)	-14.64629
CN: cc-pVDZ	(9,26)	246	173	0.704	-92.4938(3)	-92.49164
HF: cc-pCVDZ	(10,23)	283	0.998	0.0035	-100.27098(3)	-100.27044
CH_4: cc-pVDZ	(8,33)	419	377	0.898	-40.38752(1)	-40.38974
CO: cc-pVDZ	(10,26)	1,080	777	0.719	-113.05644(4)	-113.05497
H_2O: cc-pCVDZ	(10,28)	2,410	47.4	0.0196	-76.28091(3)	-76.28028
O_2: cc-pVDZ	(12,26)	5,409	2,651	0.490	-149.9875(2)	-149.98562
NaH: cc-pCVDZ	(12,32)	205,300	63.8	0.00031	-162.6090(1)	-162.60901

TABLE II: Predicted FCI results. The geometries of the molecules were (in Ångstroms): CN(1.1941), HF(0.91622), CH_4(r_{CH}=1.087728), CO(1.1448), H_2O(r_{OH}=0.975512, θ=110.565) [35], O_2(1.2074) and NaH(1.885977). CN and O_2 orbitals were constructed from a restricted open-shell HF calculation with a spin-multiplicity of two and three respectively. CN, CH_4, CO and O_2 had frozen core electrons. The number in brackets represents the error in the previous digit, obtained through a Flyvbjerg-Petersen blocking analysis [37] of $E(\tau)$.

.

tional effort required to compute the correlation energy will be similarly reduced. We are currently investigating such ideas.

Once this critical number of walkers has been reached, the correlation energy can be computed with confidence. Agreement between two estimators of the correlation energy provides additional support for the values obtained. We have reproduced existing FCI energies, as well as reporting new systems with very large FCI spaces. Favourable memory requirements, as well as ease of parallelisation, are attributes of this algorithm which should enable yet larger systems to be tackled in the near future.

In common with the FCI method, the only systematic error in the calculation of the correlation energy arises through basis-set incompleteness, which can however be systematically improved via explicit calculation on larger basis sets, together with extrapolation to the complete basis-set limit. This is marked contrast with the uncontrolled approximation of fixed-node QMC where the error introduced is hard to reduce systematically.

The present method provides a synthesis of Quantum Monte Carlo and quantum chemistry. Single and multi-reference problems can both be tackled, and the difficulty of the procedure does not seem to be closely tied to this classification of systems. This suggests that future effort will be primarily focused on multi-reference systems which provide a sterner challenge to many other methods. In addition to this, several extensions of the present technique can be envisaged which should allow much larger systems to be treated. These include CASSCF methodologies, perturbation theory extensions, and the use of configuration state functions, which are all currently under investigation.

APPENDIX A

Two determinants $|D_{\mathbf{i}}\rangle$ and $|D_{\mathbf{j}}\rangle$ are said to be coupled if and only if $\langle D_{\mathbf{i}}|H|D_{\mathbf{j}}\rangle \neq 0$. In our applications, H is a molecular Hamiltonian with one-electron and two-electron terms:

$$H = \sum_i^N h_i + \sum_{i<j} \frac{1}{r_{ij}}, \tag{23}$$

$$h_i = -\frac{1}{2}\nabla_i^2 + v_{ext}(\mathbf{r}_i) \tag{24}$$

Determinants which differ by three or more (orthonormal) spin-orbitals are therefore uncoupled, and we need only consider pairs of determinants which differ by two or fewer spin-orbitals, i.e. determinants which are double or single excitations of each other. Such matrix elements

Basis Set	Orbitals	$N_{\rm FCI}/10^6$	$N_c/10^6$	$f_c/10^{-3}$	$E_{\rm corr}$
VDZ	14	0.502	0	0	0.19211(4)
CVDZ	18	9.19	0	0	0.23365(3)
AVDZ	23	142	0.248	1.7	0.21510(3)
VTZ	30	2540	0.506	0.199	0.28341(9)
CVTZ	43	116,000	2.3	0.0198	0.33628(2)
AVTZ	46	235,000	338	1.43	0.2925(4)
VQZ	55	1.51×10^6	681	0.451	0.3347(10)
CVQZ	84	119×10^6	2200	0.0185	0.3691(1)
Extrap.					0.3930
Exact					0.3905

TABLE III: Ne atom results in Dunning[32] basis sets, giving correlation energies in Hartrees. All electrons (10) were correlated over all excitation levels. The "exact" result is the non-relativistic infinite-nuclear mass corrected experimental value, calculated in reference [40]. The extrapolated result for comparison is obtained using the technique by Halkier *et al.* [39] using the cc-pCVTZ and cc-pCVQZ correlation energies.

can be computed using the Slater-Condon rules [41–43] as follows. In the case of double excitations, let $D_{\mathbf{j}}$ be the following double excitation of $D_{\mathbf{i}}$:

$$|D_{\mathbf{j}}\rangle = a_r^\dagger a_s^\dagger a_q a_p |D_{\mathbf{i}}\rangle \tag{25}$$

Then

$$\langle D_{\mathbf{i}}|H|D_{\mathbf{j}}\rangle = \langle rs||pq\rangle \equiv \langle rs|pq\rangle - \langle rs|qp\rangle \tag{26}$$

where the (2-electron) 4-index integrals are defined by

$$\langle rs|pq\rangle = \int d\mathbf{r}_1 d\mathbf{r}_2 \phi_r^*(\mathbf{r}_1)\phi_s^*(\mathbf{r}_2)\frac{1}{r_{12}}\phi_p(\mathbf{r}_1)\phi_q(\mathbf{r}_2) \tag{27}$$

Similarly, in the case of single excitations, let $|D_{\mathbf{j}}\rangle = a_r^\dagger a_p |D_{\mathbf{i}}\rangle$, then

$$\langle D_{\mathbf{i}}|H|D_{\mathbf{j}}\rangle = \langle r|h|q\rangle + \sum_k{}' \langle rk||qk\rangle \tag{28}$$

where the sum over k extends over the $N-1$ spin-orbitals common to $D_{\mathbf{i}}$ and $D_{\mathbf{j}}$.

The diagonal matrix elements are given by:

$$\langle D_{\mathbf{i}}|H|D_{\mathbf{i}}\rangle = \sum_{p\in\mathbf{i}}\langle p|h|p\rangle + \frac{1}{2}\sum_{p,q\in\mathbf{i}}\langle pq||pq\rangle \tag{29}$$

In the present work, all the necessary integrals (4-index and 2-index) were generated from restricted Hartree-Fock orbitals using a modified version of QChem [38].

APPENDIX B

In our algorithm, it is necessary to be able to generate all single or double excitations of any determinant in such a way that its generation probability, $p_{\rm gen}(\mathbf{j}|\mathbf{i})$ is computable, non-zero and normalised. We employed the following simple strategy (which is probably not optimal from a sampling perspective; but this consideration is left for a future study).

Let us consider the generation of $|D_{\mathbf{j}}\rangle = a_r^\dagger a_s^\dagger a_q a_p |D_{\mathbf{i}}\rangle$, which involves the selection of the occupied pair (p,q) and the unoccupied pair (r,s) (with respect to $D_{\mathbf{i}}$). Then

$$p_{\rm gen}(\mathbf{j}|\mathbf{i}) = p_{\rm gen}(r,s|p,q)p_{\rm gen}(p,q) \tag{30}$$

where $p_{\rm gen}(p,q)$ is the probability to select the orbital pair (p,q) in $D_{\mathbf{i}}$, and $p_{\rm gen}(r,s|p,q)$ is the probability to select the orbitals pair (r,s) given that we have selected (p,q). For an N electron system, we select the occupied pair (p,q) with uniform probability, i.e.:

$$p_{\rm gen}(p,q) = \binom{N}{2}^{-1} = 2[N(N-1)]^{-1} \tag{31}$$

We further write:

$$\begin{aligned} p_{\rm gen}(r,s|p,q) &= p_{\rm gen}(r|s,p,q).p_{\rm gen}(s|p,q) \\ &\quad + p_{\rm gen}(s|r,p,q).p_{\rm gen}(r|p,q) \end{aligned} \tag{32}$$

In other words, we select s with probability $p_{\rm gen}(s|p,q)$, and r with probability $p_{\rm gen}(r|s,p,q)$, before computing the probability that we *could* have picked the unoccupied orbital r first, $p_{\rm gen}(r|p,q)$, followed by s, to obtain the same excitation. Generally, $p_{\rm gen}(r|s,p,q) \neq p_{\rm gen}(s|r,p,q)$. The advantage of this approach is that we can combine spin and symmetry information, which is usually available in the form of the irreducible representation spanned by each spatial orbital. For example, in the generation of r, given (s,p,q), the irreducible representation (Γ_r) of r is dictated to be the direct product of the irreps of (s,p,q):

$$\Gamma_r = \Gamma_s \otimes \Gamma_p \otimes \Gamma_q \tag{33}$$

[In the case of Abelian groups this uniquely determines Γ_r. In this study of atoms and homonuclear diatomics,

we used D_{2h} as such a group]. Having selected s, we then determine the irreducible representation of r, and then explicitly count how many unoccupied orbitals of this irreducible representation are available for selection (imposing, in addition, that the net change in spin-polarisation is zero). The reciprocal of this number gives $p_{\rm gen}(r|s,p,q)$. This completes the generation algorithm for double excitations.

The same method can be used to compute the generation probability of a single excitation, though it is much simpler. Given $|D_{\mathbf{j}}\rangle = a_r^\dagger a_p |D_{\mathbf{i}}\rangle$, we need:

$$p_{\rm gen}(r,p) = p_{\rm gen}(r|p).p_{\rm gen}(p) \tag{34}$$

We select p uniformly out the occupied orbitals of $D_{\mathbf{i}}$, (i.e. $p_{\rm gen}(p) = N^{-1}$) and then select r such that $\Gamma_r = \Gamma_p$. We then count the number of unoccupied orbitals of this irreducible representation that are available with the same spin, and compute $p_{\rm gen}(r|p)$ as the reciprocal of this number. The probability of choosing to create a double excitation is given by P_d, where $0 < P_d < 1$ and hence the probability of choosing to generate a single excitation is $1-P_d$. P_d is chosen to approximately reflect the relative number of double excitations compared to single excitations. To obtain the final $p_{\rm gen}(\mathbf{j}|\mathbf{i})$, we need to multiply by P_d if it is a double excitation, or $1-P_d$ if it is a single excitation to maintain normalization of the probabilities.

[1] N. Metropolis and S. Ulam, J. Am. Stat. Assoc. **44**, 335 (1941).
[2] M. H. Kalos, Phys. Rev. **128**, 1791 (1962).
[3] K. E. Schmidt and M. H. Kalos, *Applications of the Monte Carlo method in Statistical Physics* (Springer-Verlag, 1984), vol. 36.
[4] J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975).
[5] D. J. Klein and H. M. Pickett, J. Chem. Phys. **64**, 4811 (1976).
[6] J. B. Anderson, J. Chem. Phys. **65**, 1421 (1976).
[7] J. B. Anderson, Int. J. Quantum. Chem. **15**, 109 (1979).
[8] P. Löwdin, Phys. Rev. **97**, 1474 (1955).
[9] N. C. Handy, Chem. Phys. Lett. **74**, 280 (1980).
[10] P. Saxe, H. F. Schaefer, III, and N. C. Handy, Chem. Phys. Lett. **79**, 202 (1981).
[11] P. J. Knowles and N. C. Handy, Chem. Phys. Lett. **111**, 315 (1984).
[12] E. Rossi, G. L. Bendazzoli, S. Evangelisti, and D. Maynau, Chem. Phys. Lett. **310**, 530 (1999).
[13] D. M. Ceperley and B. J. Alder, Phys. Rev. Lett. **45**, 566 (1980).
[14] D. M. Arnow, M. H. Kalos, M. A. Lee, and K. E. Schmidt, J. Chem. Phys. **77**, 5562 (1982).
[15] D. M. Ceperley and B. J. Alder, J. Chem. Phys. **81**, 5844 (1984).
[16] D. F. Coker and R. O. Watts, Mol. Phys. **58**, 1113 (1986).
[17] G. Sugiyama and S. E. Koonin, Ann. Phys. (NY) **168**, 1 (1986).
[18] G. H. Lang, C. W. Johnson, S. E. Koonin, and W. E. Ormand, Phys. Rev. C **48**, 1518 (1993).
[19] N. Rom, D. M. Charutz, and D. Neuhauser, Chem. Phys. Lett. **270**, 282 (1997).
[20] S. Jacobi and R. Baer, J. Chem. Phys. **120**, 43 (2003).
[21] S. Zhang and H. Krakauer, Phys. Rev. Lett. **90**, 136401 (2003).
[22] W. A. Al-Saidi, S. Zhang, and H. Krakauer, J. Chem. Phys. **124**, 224101 (2006).
[23] E. R. Davidson, J. Comput. Phys. **17**, 87 (1975).
[24] Y. Saad, *Iterative Methods for Sparse Linear Systems* (SIAM, Seattle, 2004).
[25] D. C. Clary, S. M. Colwell, and H. F. Schaefer, eds., *Molecular Quantum Mechanics: selected papers of N. C. Handy* (2004).
[26] A. J. W. Thom and A. Alavi, J. Chem. Phys. **123**, 204106 (2005).
[27] A. Alavi and A. J. W. Thom, Lect. Notes. Phys. **703**, 685 (2006).
[28] The position of a walker in the list of walkers is determined by the string of n_b integers which define the determinant on which the walker is located.
[29] N. Trivedi and D. M. Ceperley, Phys. Rev. B **41**, 4552 (1990).
[30] A determinant is uniquely specified by a binary string of length $2M$, with 0 to signify an empty spin-orbital and a 1 to specify a filled spin orbital. Such a (bit) string requires $n_b = \lceil 2M/32 \rceil$ 32-bit integers to encode. The sign array is also a 4-byte integer array. Therefore the storage requirement for one walker is $n_b + 1$ 32-bit integers. As an optional extra, the energy of the walker, $\langle D_{\mathbf{i}}|H|D_{\mathbf{i}}\rangle$ can be stored as well, at the cost of an additional 8 bytes per walker. This means that this energy need not be generated every time at the "death/cloning" step, and saves an $\mathcal{O}\left[N^2\right]$ operation per walker each iteration.
[31] Website: http://www-alavi.ch.cam.ac.uk/FciMC.php.
[32] T. H. Dunning, Jr., J. Chem. Phys. **90**, 1007 (1989).
[33] J. Olsen, O. Christiansen, H. Koch, and P. Jørgensen, J. Chem. Phys. **105**, 5082 (1996).
[34] M. L. Leininger, W. D. Allen, H. F. Schaefer, III, and C. D. Sherrill, J. Chem. Phys. **112**, 9213 (2000).
[35] J. Olsen, P. Jørgensen, H. Koch, A. Balkova, and R. J. Bartlett, J. Chem. Phys. **104**, 8007 (1995).
[36] G. K.-L. Chan, M. Kállay, and J. Gauss, J. Chem. Phys. **121**, 6110 (2004).
[37] H. Flyvbjerg and H. G. Petersen, J. Chem. Phys. **91**, 461 (1989).
[38] Y. Shao, L. Fusti-Molnar, Y. Jung, J. Kussmann, C. Ochsenfeld, S. T. Brown, A. T. B. Gilbert, L. V. Slipchenko, S. V. Levchenko, D. P. O'Neill, et al., Phys. Chem. Chem. Phys. **8**, 3172 (2006).
[39] A. Halkier, T. Helgaker, P. Jørgensen, W. Klopper, H. Koch, J. Olsen, and A. K. Wilson, Chem. Phys. Lett. **286**, 243 (1998).
[40] S. J. Chakravorty, S. R. Gwaltney, F. A. Parpia, C. F. Fischer, and E. R. Davidson, Phys. Rev. A **47**, 3649 (1993).
[41] J. C. Slater, Phys. Rev. **34**, 1293 (1929).
[42] E. U. Condon, Phys. Rev. **36**, 1121 (1930).
[43] J. C. Slater, Phys. Rev. **36**, 1109 (1931).