

Semistochastic Heat-Bath Configuration Interaction Method: Selected Configuration Interaction with Semistochastic Perturbation Theory

Sandeep Sharma,^{*,†} Adam A. Holmes,^{†,‡} Guillaume Jeanmairet,^{§,||} Ali Alavi,^{§,⊥} and C. J. Umrigar[‡]

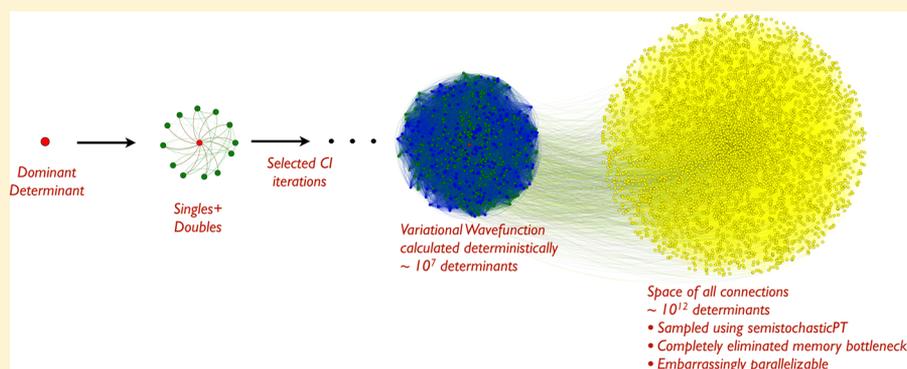
[†]Department of Chemistry and Biochemistry, University of Colorado Boulder, Boulder, Colorado 80302, United States

[‡]Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853, United States

[§]Max Planck Institute for Solid State Research, Heisenbergstraße 1, 70569 Stuttgart, Germany

^{||}Sorbonne Universités, UPMC Univ Paris 06, CNRS, Laboratoire PHENIX, Case 51, 4 place, F-75005 Paris, France

[⊥]Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, United Kingdom



ABSTRACT: We extend the recently proposed heat-bath configuration interaction (HCI) method [Holmes, Tubman, Umrigar, *J. Chem. Theory Comput.* **2016**, *12*, 3674], by introducing a semistochastic algorithm for performing multireference Epstein–Nesbet perturbation theory, in order to completely eliminate the severe memory bottleneck of the original method. The proposed algorithm has several attractive features. First, there is no sign problem that plagues several quantum Monte Carlo methods. Second, instead of using Metropolis–Hastings sampling, we use the Alias method to directly sample determinants from the reference wave function, thus avoiding correlations between consecutive samples. Third, in addition to removing the memory bottleneck, semistochastic HCI (SHCI) is faster than the deterministic variant for many systems if a stochastic error of 0.1 mHa is acceptable. Fourth, within the SHCI algorithm one can trade memory for a modest increase in computer time. Fifth, the perturbative calculation is embarrassingly parallel. The SHCI algorithm extends the range of applicability of the original algorithm, allowing us to calculate the correlation energy of very large active spaces. We demonstrate this by performing calculations on several first row dimers including F_2 with an active space of (14e, 108o), Mn–Salen cluster with an active space of (28e, 22o), and Cr_2 dimer with up to a quadruple- ζ basis set with an active space of (12e, 190o). For these systems we were able to obtain better than 1 mHa accuracy with a wall time of merely 55 s, 37 s, and 56 min on 1, 1, and 4 nodes, respectively.

I. INTRODUCTION

Many methods, for example, coupled cluster and Møller–Plesset perturbation theory, can accurately and efficiently treat the electronic correlation of single-reference (weakly correlated) systems. In particular, a coupled cluster with singles, doubles, and perturbative triples (CCSD(T)) is very accurate for such systems and is often referred to as the “gold standard” of quantum chemistry. However, these methods fail catastrophically when applied to multireference (strongly correlated) systems, such as molecules undergoing chemical reactions or systems containing transition metal atoms with partially filled d or f orbitals.

One common approach for tackling such multireference problems is to abandon the Hartree–Fock wave function and

instead use a multideterminantal reference wave function obtained by correlating a subset of orbitals around the Fermi surface. Examples include the complete active space (CAS) method, in which all possible occupancies of orbitals within the active space are included, and the restricted/generalized active space (RAS/GAS) methods,^{1–3} in which further restrictions are placed on the occupancies of the active orbitals in order to reduce the size of the Hilbert space. The CAS method is limited to about 16 active electrons and orbitals. Other possibilities include highly accurate but approximate methods such as the density matrix renormalization group (DMRG),^{4–6} full

Received: October 20, 2016

Published: March 6, 2017

configuration interaction quantum Monte Carlo (FCIQMC),^{7,8} and its semistochastic improvement (S-FCIQMC),⁹ which routinely treat up to about 40–50 active orbitals. A well-chosen active space often results in a reference wave function that contains qualitatively correct physics. However, quantitative accuracy requires one to take into account the dynamical correlation by allowing excitations into inactive-space orbitals. Common methods for including a dynamical correlation include multireference configuration interaction (MRCI) and its size-consistent variants,^{10–12} various flavors of multi-reference perturbation theory,^{13–16} and multireference coupled cluster theories.^{17–21} The accuracy of these methods is often limited by the fact that only a relatively small number of active space orbitals can be used in the reference wave function because the cost of enlarging the active space increases exponentially with the number of orbitals.

Although the number of determinants in a CAS scales combinatorially with the number of active electrons and orbitals, many of these determinants are “configurational deadwood,” and do not contribute appreciably to the reference wave function.²² The so-called *selected configuration interaction* (SCI) methods,^{22–37} which have been in use for more than four decades, take advantage of this fact and generate a reference wave function by selecting only important determinants, rather than including all determinants in the CAS. A subset of these methods improves upon the variational energy by employing a perturbative correction to the energy using multireference Epstein–Nesbet perturbation theory. We refer to these methods as *selected configuration interaction plus perturbation theory* (SCI+PT) methods. The first such method was called *configuration interaction perturbing a multiconfigurational zeroth-order wave function selected iteratively* (CIPSI).²³

The focus of this paper is a newly introduced SCI+PT method called *heat-bath configuration interaction* (HCI). HCI³⁸ distinguishes itself from other SCI+PT techniques by employing an algorithm that greatly improves the efficiency of both the variational and perturbative steps. Although it is more efficient than other SCI+PT methods, HCI, in its original formulation, is limited by a memory bottleneck because it stores in memory all the determinants that contribute to the perturbative correction (see the end of section II for more details).

In this paper, we introduce a semistochastic implementation of multireference Epstein–Nesbet perturbation theory, and use it to overcome the memory bottleneck of HCI. This method has several attractive properties. First, it does not have a sign problem that plagues quantum Monte Carlo methods. Second, instead of using the Metropolis–Hastings method, we use the Alias method to sample the variational wave function directly, so the samples are all uncorrelated. Third, in addition to removing the memory bottleneck, semistochastic HCI (SHCI) is often faster than the deterministic variant if a stochastic error of 0.1 mHa is acceptable. Fourth, within the SHCI algorithm one can trade memory for a modest increase in computer time. Fifth, the perturbative calculation is embarrassingly parallel.

In section II we review the improvements made in the original HCI algorithm that make it much more efficient than other SCI+PT algorithms. In section III, we present our stochastic perturbation theory which removes the memory bottleneck of the original HCI algorithm, and then our semistochastic perturbation theory which is more efficient in terms of computer time than either the deterministic or the stochastic variants. In section IV we provide various implementation details of both the variational and the

perturbative parts of our algorithm. We then demonstrate the utility of the stochastic and semistochastic methods by applying them in section V to various diatomic molecules including F₂ with an active space of (14e, 108o), Mn–Salen cluster with an active space of (28e, 22o), and Cr₂ dimer with up to a quadruple- ζ basis set with an active space of (12e, 190o), obtaining energies that are accurate to better than 1 mHa with very modest computer resources. Finally, in section VI, we conclude and discuss future research directions.

II. HEAT-BATH CONFIGURATION INTERACTION

We begin by describing the HCI algorithm in its original formulation,³⁸ emphasizing the key innovations that make it much more efficient than other SCI+PT methods. In the following discussion the indices i, j, \dots will be used for determinants in the variational space \mathcal{V} and the indices a, b, \dots will be used for determinants in \mathcal{C} , the space of determinants that are connected to \mathcal{V} but not in \mathcal{V} . Similar to other SCI+PT methods, HCI has two stages: (1) a variational stage, in which a variational wave function is obtained as a linear combination of a set of determinants chosen by an iterative procedure, and (2) a perturbative stage, in which the second-order correction to the variational energy is computed using multireference Epstein–Nesbet perturbation theory,^{39,40} but each stage is much faster than in other SCI+PT methods.

II.A. Variational Stage. At the start of the algorithm, \mathcal{V} consists of an initial set of determinants, usually just the Hartree–Fock determinant. Then, at each iteration, new determinants are added to \mathcal{V} , chosen using a parameter ϵ_1 , as follows. The initial wave function is the ground state of the Hamiltonian in \mathcal{V} , $|\Psi_0\rangle = \sum_i c_i |D_i\rangle$. At each iteration the following is done.

1. Add to the variational space \mathcal{V} all determinants D_a in the space of connections \mathcal{C} , such that

$$|H_{ai}c_i| > \epsilon_1 \quad (1)$$

for at least one determinant D_i in the current \mathcal{V} .

2. Calculate the lowest eigenvalue E_0 with eigenvector $|\Psi_0\rangle = \sum_i c_i |D_i\rangle$ of the Hamiltonian in \mathcal{V} .

The iterations are terminated when the number of new determinants is less than a threshold, for example, 1% of the current size of \mathcal{V} , or when a maximum number of iterations is reached. Since the values of c_i tend to be larger in the initial iterations when there are few determinants in \mathcal{V} , ϵ_1 is set during the first few iterations to be larger than its final value.

HCI takes advantage of the fact that the double excitation matrix elements depend only on the four orbitals whose occupancy is changing. Step 1 is performed efficiently by storing the double excitation matrix elements in order of decreasing magnitude, so that no time is wasted on determinants that do not meet the cutoff in eq 1. For details, we refer the reader to the original HCI paper.³⁸ Thus, we see that HCI identifies new determinants to add to \mathcal{V} in a manner that is more efficient than other SCI methods in two ways:

- HCI uses a selection criterion which is *inexpensive to evaluate* for each determinant, namely eq 1. In contrast, other SCI methods use a criterion based on a perturbative expression which is more expensive; for example, CIPSI²³ uses the magnitude of the coefficient of the first-order correction to the wave function, namely

$$\left| \frac{\sum_i H_{ai}c_i}{E_0 - H_{aa}} \right| > \epsilon_1.$$

- HCI evaluates its selection criterion (eq 1) only for those doubly excited determinants which will be added to \mathcal{V} ! By comparison, other SCI methods iterate through *all* candidate determinants $\{D_a\}$ (determinants for which there exists at least one nonzero matrix element H_{ai} with $D_i \in \mathcal{V}$), evaluating their expensive selection criteria for each one.

The simplification in HCI is possible because it was demonstrated³⁸ that variation in the perturbative expression for the coefficients is dominated by variation in the largest-magnitude term in the numerator, since the matrix elements $\{H_{ai}\}$ and coefficients $\{c_i\}$ span many orders of magnitude. The minor deviation from optimality in the choice of the most important determinants is by far outweighed by the fact that many more determinants can be included because the selection criterion of HCI allows the variational and perturbative steps to be performed at a much reduced computational cost.

II.B. Perturbative Stage. The variational wave function is used to define the zeroth order Hamiltonian H_0 and the perturbation V ,

$$H_0 = \sum_{i,j} H_{ij}|D_i\rangle\langle D_j| + \sum_a H_{aa}|D_a\rangle\langle D_a|$$

$$V = H - H_0 \quad (2)$$

It can easily be verified that $|\Psi_0\rangle$ is the ground state of H_0 with eigenvalue E_0 . Using the partitioning in eq 2, the first-order correction of the wave function $|\Psi_1\rangle$ and the second-order energy correction ΔE_2 can be written as

$$|\Psi_1\rangle = \frac{1}{E_0 - H_0} V|\Psi_0\rangle$$

$$= \sum_a \frac{\sum_i H_{ai}c_i}{E_0 - E_a} |D_a\rangle \quad (3)$$

and

$$\Delta E_2 = \langle \Psi_0 | V | \Psi_1 \rangle$$

$$= \sum_a \frac{(\sum_i H_{ai}c_i)^2}{E_0 - E_a} \quad (4)$$

where $E_a = H_{aa}$. It is worth noting that the expression for the total energy, $E_0 + \Delta E_2$ is identical to that for the mixed estimator of the energy used in quantum Monte Carlo calculations, provided that the projected wave function is replaced by the perturbed wave function.

This expression in eq 4 is expensive to calculate, as it requires a summation over many small terms. Instead, HCI includes only those terms in the sum that contribute substantially,

$$\Delta E_2 \approx \sum_a \frac{(\sum_i^{(\epsilon_2)} H_{ai}c_i)^2}{E_0 - E_a} \quad (5)$$

where $\sum_i^{(\epsilon_2)}$ denotes a sum in which all terms in the sum that are smaller in magnitude than ϵ_2 are discarded; that is, $\sum_i^{(\epsilon_2)} H_{ai}c_i$ includes only terms for which $|H_{ai}c_i| > \epsilon_2$.

Once again, since the double excitation matrix elements are stored in order of decreasing magnitude, no time is spent on the doubly excited terms that do not contribute to the sum. The parameter ϵ_2 is kept much smaller than the parameter ϵ_1 because discarding small amplitude determinants can lead to

significant errors in the calculation of dynamical correlation. In the original HCI paper,³⁸ for each ϵ_1 , several values of ϵ_2 were used, and the energy for $\epsilon_2 = 0$ was obtained by extrapolation. In this paper, a single value of ϵ_2 is used, that is sufficiently small to recover the $\epsilon_2 = 0$ limit to a precision that is better than 1 mHa.

It was shown in the previous publication³⁸ that the above algorithm is highly efficient and can be used to obtain sub-milliHartree accuracy for challenging problems such as all-electron (48e, 42o) Cr_2 with the small Ahlrichs double- ζ basis⁴¹ in a few minutes on a single computer core. However, since the contributions from all i in eq 5 must be summed and then squared, the efficient deterministic approach to computing the perturbative correction requires storing the partial sums $\{\sum_i^{(\epsilon_2)} H_{ai}c_i\}$ for all a for which $|H_{ai}c_i| > \epsilon_2$ which creates a severe memory bottleneck.

To see this, we note that the N_v determinants in the variational space are connected to $O(n^2v^2N_v)$ determinants in the perturbative space with nonzero Hamiltonian matrix elements, where n is the number of electrons and v is the number of virtual orbitals. For a relatively conservative number of $n = 12$, $v = 50$, and $N_v = 10^7$, the perturbative space will contain over 10^{12} determinants, requiring over 10 terabyte memory. The original HCI algorithm reduces this storage requirement by orders of magnitude by storing only determinants D_a for which $|H_{ai}c_i| > \epsilon_2$ for at least one determinant $D_i \in \mathcal{V}$. But it should be recognized that using a large ϵ_2 to truncate the perturbative space can lead to significant errors in the perturbative correction. If on the other hand ϵ_2 is reduced to decrease the error in the perturbative correction the number of determinants in the perturbative space C rapidly increases as shown in Figure 1, making the calculation infeasible.

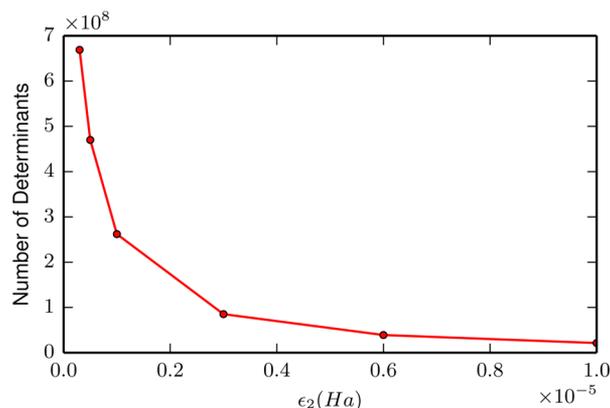


Figure 1. Demonstration that the number of determinants in the perturbative space C increases rapidly as ϵ_2 is reduced. These calculations, for the C_2 dimer with a QZ basis, used $\epsilon_1 = 2 \times 10^{-4}$ Ha in the variational calculation resulting in 403 071 variational determinants. All the calculations presented in the results section use an $\epsilon_2 = 10^{-8}$ Ha, whereas the smallest ϵ_2 used in this graph is 3×10^{-7} Ha.

In the following section we show that this memory bottleneck can be completely eliminated, without having to increase ϵ_2 , by using a stochastic or semistochastic version of the perturbation theory.

III. STOCHASTIC MULTIREFERENCE PERTURBATION THEORY

We first discuss the stochastic method for computing the perturbative correction before discussing the more efficient semistochastic method.

III.A. Stochastic PT. We write the perturbative correction in a slightly different form than presented in eq 5 to highlight the fact that it is a bilinear function of the coefficients of the zeroth-order state.

$$\Delta E_2 = \sum_a \frac{1}{E_0 - E_a} \left(\sum_{ij}^{(\epsilon_2)} H_{ai} H_{aj} c_i c_j \right) \quad (6)$$

We compute the expected value of this expression stochastically by employing N_s samples, each of which consists of N_d determinants $\{D_i\}$ sampled from \mathcal{V} with probability

$$p_i = \frac{|c_i|}{\sum_i |c_i|} \quad (7)$$

Any given sample will contain N_d^{diff} distinct determinants D_i with some number of repeats w_i , such that

$$\sum_i^{N_d^{\text{diff}}} w_i = N_d$$

The number of repetitions w_i is distributed according to the well-known multinomial distribution. The mean and second moment, for $i \neq j$, of this distribution are

$$\langle w_i \rangle = p_i N_d \quad (8)$$

$$\langle w_i w_j \rangle = p_i p_j N_d (N_d - 1) \quad (9)$$

where $\langle \cdot \rangle$ denotes the expectation value of a quantity evaluated for a sample of N_d determinants, a notation we will use hereafter.

Using these expressions, the unbiased estimate of the second-order perturbation can be calculated from the sampled wave function as follows,

$$\begin{aligned} \Delta E_2 &= \sum_a \frac{1}{E_0 - E_a} \left[\sum_{ij}^{N_v} H_{ai} H_{aj} c_i c_j \right] \\ &= \sum_a \frac{1}{E_0 - E_a} \left[\sum_{i \neq j}^{N_v} H_{ai} H_{aj} c_i c_j + \sum_i H_{ai}^2 c_i^2 \right] \\ &= \left\langle \sum_a \frac{1}{E_0 - E_a} \left[\sum_{i \neq j}^{N_d^{\text{diff}}} \frac{w_i w_j c_i c_j H_{ai} H_{aj}}{\langle w_i w_j \rangle} + \sum_i \frac{w_i c_i^2 H_{ai}^2}{\langle w_i \rangle} \right] \right\rangle \\ &= \left\langle \sum_a \frac{1}{E_0 - E_a} \left[\sum_{i \neq j}^{N_d^{\text{diff}}} \frac{w_i w_j c_i c_j H_{ai} H_{aj}}{p_i p_j N_d (N_d - 1)} + \sum_i \frac{w_i c_i^2 H_{ai}^2}{p_i N_d} \right] \right\rangle \\ &= \frac{1}{N_d (N_d - 1)} \left\langle \sum_a \frac{1}{E_0 - E_a} \left[\left(\sum_i^{N_d^{\text{diff}}} \frac{w_i c_i H_{ai}}{p_i} \right)^2 \right. \right. \\ &\quad \left. \left. + \sum_i^{N_d^{\text{diff}}} \left(\frac{w_i (N_d - 1)}{p_i} - \frac{w_i^2}{p_i^2} \right) c_i^2 H_{ai}^2 \right] \right\rangle \quad (10) \end{aligned}$$

where for brevity we have suppressed the superscript (ϵ_2) on the i and j sums, though of course we will always use a nonzero ϵ_2 value for efficiency.

Going from the 2nd to the 3rd line in eq 10, we replace the sum over the states in \mathcal{V} by a sum over the sample, so in order to have an unbiased expectation value we divide the two terms by $\langle w_i w_j \rangle$ and $\langle w_i \rangle$, respectively. In going from the 3rd to the 4th line we use eqs 8 and 9.

In practice, the exact average in eq 10 will be replaced by an average over N_s samples. For any $N_d \geq 2$ we obtain an unbiased estimate of the second-order correction to the energy, and this estimate can be made progressively more precise by averaging over a large number of samples N_s . Each batch contains an independently chosen set of N_d^{diff} determinants and thus there is no autocorrelation between consecutive batches. This is in sharp contrast to discrete-space quantum Monte Carlo methods, such as the FCIQMC method^{7,8} and its semi-stochastic improvement,⁹ for which the autocorrelation time increases both with system size and the size of the basis to the point that it can become difficult to accurately estimate the statistical error. This drawback of the FCIQMC method is ameliorated but not eliminated by using the more efficient sampling method of ref 42.

We note that the expression in eq 10 is evaluated in much the same way as the deterministic evaluation of the perturbative correction using a single batch, the main difference being that the N_v variational determinants have been replaced by the much smaller subset of N_d^{diff} distinct sampled determinants and that an additional summation is needed to ensure that the result is unbiased. Note that for each sample, the summation over a in eq 10 is restricted to only those determinants in \mathcal{C} that have a nonzero Hamiltonian matrix element with the N_d^{diff} determinants used to sample the zeroth-order wave function.

Figure 2a shows that the CPU time per sample increases nearly linearly with N_d , the number of determinants in the sample, for the C_2 and F_2 molecules. As shown in section IV, the scaling contains two terms: one that scales linearly with N_d and another that scales as $N_d \log(N_d)$ (ref 43). Figure 2b shows the CPU time necessary to reach a standard deviation of less than 0.1 mHa versus the number of determinants in the sampled wave function N_d . There is a rapid initial decrease followed by a much shallower decrease beyond about $N_d = 200$. Consequently, it is desirable to use as large a value of N_d as memory allows. Another consideration is that N_s needs to be large enough to get a reasonable estimate of the statistical error, and since the computer time is approximately $\propto N_d N_s$, it sometimes makes sense to use a smaller N_d than available memory allows. In all the calculations presented in section V, we have used $N_d = 200$, even though the computer time could be greatly reduced by using a larger N_d for the larger systems.

It is worth mentioning that the memory bottleneck can also be removed without recourse to the stochastic method. This can be achieved by dividing the N_v determinants in \mathcal{V} into N_b batches, each containing on average N_d determinants ($N_b = N_v / N_d$), and computing the contribution from all pairs of batches independently. In section IV we will see that the leading cost of performing the calculation for each pair of batches is $\propto N_d$ and so the cost of performing the entire calculation containing N_b batches scales as $N_d N_b^2 \propto N_v^2 / N_d$. Thus, for $N_v > N_d$, the cost of the deterministic approach scales quadratically with N_v . In contrast, the cost of the stochastic approach scales sublinearly with N_v because additional low weight determinants are sampled only infrequently. The price to be paid is that the computed energy has a stochastic error. Still, for large N_v and small ϵ_2 , the N_s required to get a statistical error of 1 mHa in the stochastic method is much smaller than the N_b^2 required in

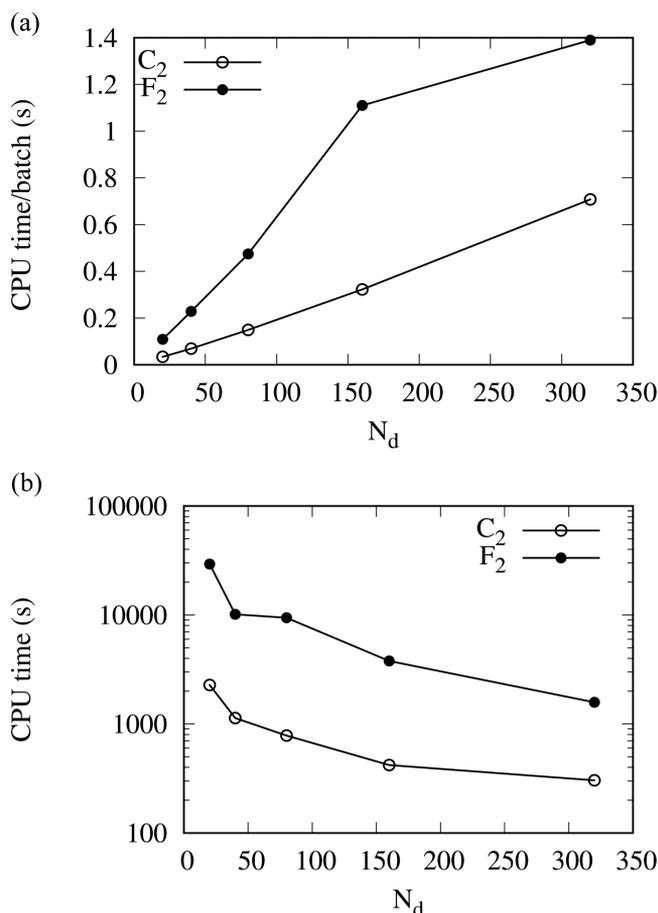


Figure 2. (a) Demonstration of the near-linear scaling of the CPU time per batch for the perturbative calculation as a function of the number of determinants, N_d , sampled in each batch. The open and the filled circles are for the C_2 and F_2 molecules respectively with cc-pVQZ basis sets. (b) CPU time in seconds required to reach a standard deviation of 0.1 mHa for various values of N_d . Note the initial rapid decrease in CPU time followed by a more gradual decrease at larger N_d values.

the deterministic calculation, making the stochastic approach the more efficient choice.

III.B. Semistochastic PT. Although the stochastic method eliminates the memory bottleneck, for large variational spaces and large basis sets it takes many stochastic samples, N_s , to reduce the stochastic error to the desired value. The computer time can be reduced by using a semistochastic method in which the perturbative calculation is split into two steps. In the first step a deterministic perturbative calculation is performed using a relatively loose threshold ϵ_2^d to obtain the deterministic second-order correction $\Delta E_2^D[\epsilon_2^d]$. The error incurred due to the use of large ϵ_2^d in the deterministic calculation can be corrected stochastically by calculating the difference between the second-order energies obtained from a tight threshold ϵ_2 and a loose threshold ϵ_2^d :

$$\Delta E_2[\epsilon_2] = (\Delta E_2^S[\epsilon_2] - \Delta E_2^S[\epsilon_2^d]) + \Delta E_2^D[\epsilon_2^d] \quad (11)$$

where $\Delta E_2^S[\epsilon_2]$ and $\Delta E_2^S[\epsilon_2^d]$ are the second-order energies calculated with ϵ_2 and ϵ_2^d , respectively, using the stochastic method. The key point here is that both the $\Delta E_2^S[\epsilon_2]$ and $\Delta E_2^S[\epsilon_2^d]$ are calculated using the same set of sampled N_d variational determinants and thus there is substantial cancellation of stochastic error, and almost no increase in

memory or computer time. The value of ϵ_2^d affects the statistical error of the energy for given computer time, but not the expectation value of the energy. In the results section we show that using the semistochastic method for Cr_2 dimer can speed up the calculation by more than a factor of 2. This speed up can be even larger if the available computer memory permits a smaller ϵ_2^d to be used. In the limit that one can afford to use $\epsilon_2^d = \epsilon_2$, the stochastic noise is completely eliminated. Thus, the semistochastic method gives us the ability to go smoothly from the fully deterministic to the fully stochastic algorithm.

IV. IMPLEMENTATION

Here we briefly describe the implementation and the leading-order cost of the various steps of the algorithm. In the variational stage there are three main operations, identifying the significant determinants to be included in the variational space, building the Hamiltonian matrix, and diagonalizing the matrix. The cost of identifying important determinants is $O(kN_v \ln(N_v) + kN_v \ln(N_p))$, where N_v is the number of determinants in the variational space, k is the average number of H_{ai} elements for determinants D_i in \mathcal{V} that satisfy eq 1, and $N_p = kN_v$ is the total number of new determinants that satisfy the criterion in eq 1. The two terms in the cost function result from generating kN_v determinants and then doing a binary search of the list of the N_v existing variational determinants and the N_p newly generated determinants before including the determinant just generated in the newly generated determinant list.

In the current implementation we store all the nonzero elements of the Hamiltonian in memory using a list of lists (LIL) sparse storage format. In LIL format for each row we store a list containing the column index and the value of the nonzero Hamiltonian matrix elements.

The determinant labels are bit-packed strings that represent the occupancies of the up-spin (α) and the down-spin (β) orbitals. To build the Hamiltonian efficiently, we first generate a list of all unique β strings and associated with each β string we store a list of all determinants in \mathcal{V} that have that β string. We also generate a list of all unique α strings with $N_\alpha - 1$ electrons and associated with each α string we store a list of all determinants in \mathcal{V} that give the α string on removing one α electron. Here, N_α is the number of α electrons in our system. Determinants that are related to each other by double or single α excitations have the same β string, and all the pairs of determinants that are related to each other with the remaining possible single or double excitations have the same α string with $N_\alpha - 1$ electrons. Hence to find the connected determinants, only the determinants in these two lists need to be considered rather than the entire set of N_v determinants in \mathcal{V} . Once the Hamiltonian is generated the Davidson algorithm is used to diagonalize it and the most expensive step there is, the Hamiltonian wave function multiplication which costs $O(kN_v)$. Despite the fact that the Hamiltonian is sparse, building it is the most expensive part of the variational step, and storing it is currently the largest memory bottleneck in the code. In the future we intend to implement the *direct* method^{22,44} for carrying out Hamiltonian wave function multiplication which does not require storing the Hamiltonian and can take less computer time as well.

The stochastic perturbation step has two major components: sampling N_d determinants from the list of N_v variational determinants, and identifying the determinants in \mathcal{C} that are connected to these N_d determinants and computing their contribution to the perturbative correction, eq 10. The N_d

Table 1. Ground State SHCI Energies of the C_2 , N_2 , O_2 , NO , and F_2 Molecules with Bond Lengths of 1.2425, 1.0977, 1.2075, 1.1508, and 1.4119 Å, Respectively, and DZ, TZ, and QZ Basis Sets^a

molecule	basis	sym	ϵ_1 (Ha)	N_v	energy (Ha)		wall time (sec)		
					var	total	var	PT	total
C_2	DZ	$^1A_{1g}$	5×10^{-4}	28566	-75.7217	-75.7286(2)	1	2	3
C_2	TZ	$^1A_{1g}$	3×10^{-4}	142467	-75.7738	-75.7846(3)	7	4	11
C_2	QZ	$^1A_{1g}$	2×10^{-4}	403071	-75.7894	-75.8018(4)	36	10	46
N_2	DZ	$^1A_{1g}$	5×10^{-4}	37593	-109.2692	-109.2769(1)	1	2	3
N_2	TZ	$^1A_{1g}$	3×10^{-4}	189080	-109.3608	-109.3748(6)	10	4	14
N_2	QZ	$^1A_{1g}$	2×10^{-4}	499644	-109.3884	-109.4055(9)	44	10	54
O_2	DZ	$^3A_{1g}$	5×10^{-4}	52907	-149.9793	-149.9878(2)	2	3	5
O_2	TZ	$^3A_{1g}$	3×10^{-4}	290980	-150.1130	-150.1307(8)	17	6	23
O_2	QZ	$^3A_{1g}$	2×10^{-4}	770069	-150.1541	-150.1748(9)	72	23	95
NO	DZ	2B_1	5×10^{-4}	48305	-129.5881	-129.5997(3)	2	3	5
NO	TZ	2B_1	3×10^{-4}	227004	-129.6973	-129.7181(9)	18	7	25
NO	QZ	2B_1	2×10^{-4}	606381	-129.7311	-129.7548(9)	93	23	116
F_2	DZ	$^1A_{1g}$	5×10^{-4}	68994	-199.0913	-199.1001(7)	2	3	5
F_2	TZ	$^1A_{1g}$	3×10^{-4}	395744	-199.2782	-199.2984(9)	8	6	14
F_2	QZ	$^1A_{1g}$	2×10^{-4}	1053491	-199.3463	-199.3590(9)	129	22	151
Natural Orbitals									
F_2	DZ	$^1A_{1g}$	1×10^{-3}	16824	-199.0871	-199.0994(4)	0	3	3
F_2	TZ	$^1A_{1g}$	5×10^{-4}	141433	-199.2787	-199.2972(7)	7	6	13
F_2	QZ	$^1A_{1g}$	5×10^{-4}	221160	-199.3355	-199.3590(9)	28	27	55

^aThe variational cutoff, ϵ_1 , the number of determinants in the variational space, N_v , the variational energy, var, and the total energy, total, are shown. The statistical error in the last digit is shown in parentheses. All perturbative calculations were performed with an $\epsilon_2 = 10^{-8}$ Ha and $N_d = 200$. The last three rows show that the SHCI calculations converge with a much looser ϵ_1 threshold when MP2 natural orbitals, rather than canonical HF orbitals, are used. The final three columns show the wall time in seconds required to perform the variational and the stochastic perturbative parts of the calculation. All the results obtained from perturbation theory (PT) agree within error bars with the results published previously using FCIQMC.⁴⁷ Each calculation was performed on a single node (see text for details).

determinants are sampled using the Alias method,^{45,46} which has an initial one-time memory cost of $O(N_v)$, and a subsequent cost of $O(N_d)$ each time a sample is drawn. The Alias method was used by some of us⁴² for efficiently sampling determinants in the S-FCIQMC method. The computer time for identifying the connected determinants along with their contributions is $O(n^2 v^2 N_d \log(n^2 v^2 N_d) + n^2 v^2 N_d \log(N_v))$, while the memory required is $O(n^2 v^2 N_d + N_v)$. Since the minimum required value of N_d is just two, the memory requirement for the stochastic perturbation theory is smaller than that of other parts of the calculation.

We have parallelized the entire code using hybrid OMP/MPI (open multiprocessing/message passing interface) programming to make full use of the symmetric multiprocessor (SMP) architecture of most modern computers. A separate MPI process is initiated on each computer node and then each process forks into several threads (one for each computational core) on the node. The variational wave function is replicated on each node but a single copy is shared among the different threads on a node. As mentioned previously, the most memory intensive data structure is the LIL used to store the sparse Hamiltonian matrix. The rows of the Hamiltonian are distributed in round-robin fashion among the nodes. Of course, only the nonzero elements are stored. With this strategy the storage of the Hamiltonian and the computation of the Hamiltonian wave function multiplication is distributed approximately evenly between the different nodes and threads.

The perturbative step is embarrassingly parallel and no special strategy is needed to parallelize this step.

V. BENCHMARKS

We perform frozen core calculations (including all the Hartree–Fock virtual orbitals in the active space) on a series of first row dimers including C_2 , N_2 , O_2 , NO , and F_2 with cc-pVDZ, cc-pVTZ, and cc-pVQZ basis sets. (In the interest of brevity, these three basis sets will sometimes be abbreviated as DZ, TZ, and QZ, respectively.) Although the active spaces used for the first row diatomics have large Hilbert spaces, they do not exhibit strong correlation in their equilibrium geometry and traditional methods such as CCSD(T) are inexpensive and reliable. Thus, we also perform frozen-core calculations on the Cr_2 dimer using cc-pVDZ, cc-pVTZ, and cc-pVQZ bases, which have active spaces containing (12e, 68o), (12e, 118o), and (12e, 190o) respectively. Cr_2 is well-known for being very strongly correlated. Most multireference methods can use no more than the minimal active space and therefore fail to get even qualitatively correct dissociation curves. Finally, we also perform calculations on the Mn–Salen model complex which is a prototypical strongly correlated inorganic molecule containing open shell d -orbitals giving rise to nearly degenerate singlet and triplet ground states. For all the systems we obtain energies that are accurate to 1 mHa for the chosen basis; for the first-row dimers we compare to i-FCIQMC energies,⁴⁷ for the Cr_2 dimer we perform internal convergence tests since there are no

good approximations to the FCI energies in the literature, and for Mn-Salen we compare to DMRG energies.⁴⁸

V.A. First Row Diatomics. In the variational calculations we start with a value of ϵ_1 during the first few iterations that is larger than its final value because the values of c_i tend to be larger in the initial iterations when there are few determinants in \mathcal{V} . For example, for the cc-pVQZ basis set, we successively reduce the value of ϵ_1 from 10^{-3} to 5×10^{-4} to 3×10^{-4} and 2×10^{-4} Ha, and perform three iterations at each value. The cost of performing the first iteration at a value of ϵ_1 is larger than that for subsequent iterations because relatively few new determinants are introduced after the first iteration.

Table 1 shows benchmark calculations on the first row dimers using a single node containing two Intel Xeon E5-2680 v2 processors of 2.80 GHz each and 128 gigabyte memory. Among these calculations F_2 had the largest active space containing 14 electrons in 108 orbitals (14e, 108o) with a Hilbert space containing over 10^{20} determinants. On a single node it required less than 3 min to get the energy converged to better than 1 mHa. It is not possible to perform the calculations for the larger systems and basis sets on a single node with the original algorithm because the cost of storing all the determinants in the space of connections \mathcal{C} that contribute to the perturbative corrections is prohibitive. Interestingly, not only is the memory requirement of the stochastic method smaller than that of the original deterministic algorithm, but even the computer time required to obtain sub-mHa accuracy is smaller, for the systems for which the deterministic algorithm is feasible.

As expected, these calculations can be done even more efficiently, if the Hartree–Fock orbitals are replaced by natural orbitals from some approximate correlated theory. For example the last three rows of Table 1 show that the calculations on the F_2 dimer, for all three basis sets, can be run with a larger ϵ_1 resulting in about a factor of 3 speedup when MP2 natural orbitals are used.

V.B. Cr_2 Dimer. The Cr_2 dimer is well-known to be a very challenging system for most electronic structure methods. We perform frozen core calculations by including all the virtual orbitals in the active space with a bond length of 1.68 Å using the cc-pVDZ-DK, cc-pVTZ-DK, and cc-pVQZ-DK basis sets. The relativistic effects are included using the second-order Douglas–Kroll–Hess Hamiltonian. For all the basis sets calculations are performed using natural orbitals obtained by first performing a short unconverged FCIQMC calculation. For the TZ basis, calculations are also performed using approximate natural orbitals obtained from an SHCI variational wave function. The active spaces with the DZ, TZ, and QZ basis sets contained (12e, 68o), (12e, 118o), and (12e, 190o), respectively, with the largest Hilbert space containing more than 10^{21} determinants.

Figure 3 shows the number of determinants and the sum of the squares of the coefficients of the variational wave function for each excitation level relative to the dominant determinant for the three basis sets with $\epsilon_1 = 8 \times 10^{-5}$ Ha. The wave function contains determinants with excitation orders all the way up to the maximum possible of 12. Hence CI expansions that are truncated at double or even quadruple excitations are far from adequate.

Table 2 shows the variational and the total energies for each of the three basis sets. Although the variational energies are far from convergence, the total energies including the perturbative corrections converge rapidly as ϵ_1 is reduced. Also shown in the

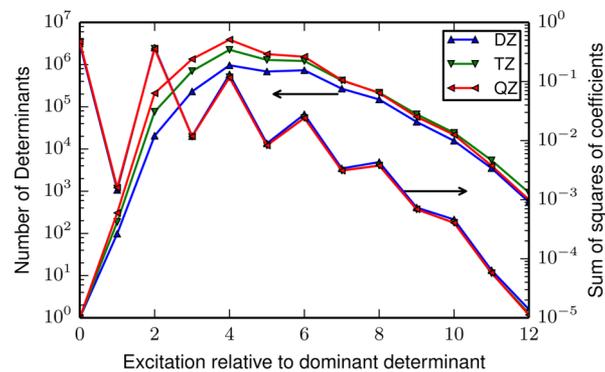


Figure 3. The number of determinants and the sum of the squares of the coefficients of the variational wave function versus excitation level relative to the dominant determinant for each of the three basis sets, with $\epsilon_1 = 8 \times 10^{-5}$ Ha. Determinants of all excitation orders are important.

table are the computer times for those calculations that employed the computer architecture described above. The cost of the perturbative calculation for a constant statistical error is smaller for the semistochastic variant compared to the stochastic variant of the method. This is particularly important when a small statistical error is required since the stochastic error of the perturbative calculation decreases as $1/\sqrt{N_s}$, where N_s is the number of samples used in the stochastic perturbative calculation. For example, for the TZ basis with $\epsilon_1 = 7 \times 10^{-5}$ Ha, if a statistical error of 0.1 mHa is required the computer time is 82800 s for stochastic PT but only $394 + 32800 = 33194$ s for semistochastic PT. The efficiency of the semistochastic method can be further improved by reducing the value of ϵ_2^d . For large systems it is most efficient to use the smallest value of ϵ_2^d that is permitted by computer memory, though for small systems a larger value of ϵ_2^d can be optimal.

The computed total energies for the smallest ϵ_1 values for the DZ and QZ bases is within 1 mHa of their respective extrapolated energies but not for the TZ basis. Hence, as described above, we also calculated energies using approximate natural orbitals obtained from an SHCI wave function. These calculations are denoted by TZ' (although the basis is the same) and they are much better converged.

It is noteworthy that the cost of the perturbative correction relative to the variational calculations decreases as the size of the variational space \mathcal{V} increases. In fact, the CPU time required to reach a fixed statistical error is rather insensitive to the size of the variational space for a given basis set. This is because as the variational space is increased, the perturbative correction becomes smaller, and because the additional variational determinants have relatively few connections that satisfy the ϵ_2 threshold. This indicates that much larger calculations could be performed if the memory and computer time of the variational step were reduced, potentially by using the direct CI method.⁴⁴

V.C. Mn-Salen. The calculations for the first-row dimers and the Cr_2 dimer were for the exact frozen core energies. We now demonstrate that SHCI can be used to calculate the active space energy of a prototypical strongly correlated molecule like Mn-Salen ($MnClO_3N_2C_8H_{10}$) (see Figure 4) very quickly. Mn-Salen derivatives such as Jacobson's catalyst are used to catalyze enantioselective epoxidation of olefins. Despite their widespread use and importance, the mechanism of the catalysis reaction is not known and has spawned a series of theoretical

Table 2. Ground State SHCI Energies of the Cr₂ Molecule with Bond Length 1.68 Å with DZ, TZ, and QZ Basis Sets^a

basis	Sym	ϵ_1 (Ha)	ϵ_2^d (Ha)	N_v	energy (Ha)		wall time (sec)			
					var	total	var	PT _{det}	PT(1)	no. nodes
DZ	¹ A _{1g}	8×10^{-5}		3114163	-2099.4692	-2099.4875(5)	395	0	983	1
DZ	¹ A _{1g}	50×10^{-5}	5×10^{-6}	210421	-2099.4344	-2099.4851(1)				
DZ	¹ A _{1g}	20×10^{-5}	5×10^{-6}	832196	-2099.4560	-2099.4864(1)				
DZ	¹ A _{1g}	8×10^{-5}	5×10^{-6}	3114163	-2099.4692	-2099.48754(3)	395	100	246	1
DZ	¹ A _{1g}	6×10^{-5}	5×10^{-6}	4708713	-2099.4724	-2099.48782(7)				1
DZ	¹ A _{1g}	5×10^{-5}	5×10^{-6}	6114463	-2099.4741	-2099.48788(7)				1
DZ	¹ A _{1g}	4×10^{-5}	5×10^{-6}	8390964	-2099.4760	-2099.48809(3)				1
DZ	¹ A _{1g}	extrapolated				-2099.4887(2)				
TZ	¹ A _{1g}	8×10^{-5}		6268840	-2099.5051	-2099.5276(7)	593	0	783	4
TZ	¹ A _{1g}	7×10^{-5}		7651680	-2099.5070	-2099.5283(6)	862	0	828	4
TZ	¹ A _{1g}	6×10^{-5}		9666032	-2099.5090	-2099.5283(6)	1024	0	795	4
TZ	¹ A _{1g}	50×10^{-5}	5×10^{-6}	379428	-2099.4637	-2099.52183(1)				
TZ	¹ A _{1g}	20×10^{-5}	5×10^{-6}	1561960	-2099.4887	-2099.52532(5)				
TZ	¹ A _{1g}	10×10^{-5}	5×10^{-6}	4496674	-2099.5017	-2099.5277(1)				
TZ	¹ A _{1g}	8×10^{-5}	5×10^{-6}	6268840	-2099.5051	-2099.5283(1)	593	250	233	4
TZ	¹ A _{1g}	7×10^{-5}	5×10^{-6}	7651680	-2099.5070	-2099.52851(4)	862	319	313	4
TZ	¹ A _{1g}	6×10^{-5}	5×10^{-6}	9687009	-2099.5090	-2099.52884(4)	1024	394	328	4
TZ	¹ A _{1g}	5×10^{-5}	5×10^{-6}	12759006	-2099.5113	-2099.52936(7)				
TZ	¹ A _{1g}	extrapolated				-2099.5312(15)				
TZ'	¹ A _{1g}	50×10^{-5}	5×10^{-6}	365104	-2099.4650	-2099.52831(1)				
TZ'	¹ A _{1g}	20×10^{-5}	5×10^{-6}	1549370	-2099.4899	-2099.52961(2)				
TZ'	¹ A _{1g}	10×10^{-5}	5×10^{-6}	4429824	-2099.5028	-2099.53073(4)				
TZ'	¹ A _{1g}	8×10^{-5}	5×10^{-6}	6185301	-2099.5062	-2099.53104(3)				
TZ'	¹ A _{1g}	7×10^{-5}	5×10^{-6}	7553766	-2099.5081	-2099.53124(5)				
TZ'	¹ A _{1g}	6×10^{-5}	5×10^{-6}	9510287	-2099.5101	-2099.53135(6)				
TZ'	¹ A _{1g}	5×10^{-5}	5×10^{-6}	12479803	-2099.5123	-2099.53155(4)				
TZ'	¹ A _{1g}	extrapolated				-2099.5325(4)				
QZ	¹ A _{1g}	8×10^{-5}		9516339	-2099.5246	-2099.5553(13)	1584	0	3578	4
QZ	¹ A _{1g}	50×10^{-5}	5×10^{-6}	497747	-2099.4782	-2099.5562(3)				
QZ	¹ A _{1g}	20×10^{-5}	5×10^{-6}	2285120	-2099.5064	-2099.5560(3)				
QZ	¹ A _{1g}	10×10^{-5}	5×10^{-6}	6768521	-2099.5208	-2099.55649(3)				
QZ	¹ A _{1g}	8×10^{-5}	5×10^{-6}	9516339	-2099.5246	-2099.55670(7)	1584	372	1429	4
QZ	¹ A _{1g}	6×10^{-5}	5×10^{-6}	14812200	-2099.5290	-2099.55668(8)				
QZ	¹ A _{1g}	5×10^{-5}	5×10^{-6}	19481471	-2099.5315	-2099.55682(8)	4960	2751	1087	4
QZ	¹ A _{1g}	extrapolated				-2099.5571(3)				

^aFor each basis set, calculations were performed with both the stochastic and the semistochastic methods described in section III. The various columns in the table have the same meaning as those in Table 1, but there are three new columns. The column labeled ϵ_2^d gives the value of the ϵ_2 used to perform the deterministic calculation in the semistochastic variant of the method. The column labeled PT_{det} shows the time required for the deterministic part of the PT calculation for the semistochastic method, and the column labelled PT(1) shows the time it would take to perform the stochastic perturbative step if we terminated the calculation after obtaining an uncertainty of 1 mHa. The final column labeled no. nodes shows the number of computer nodes used. The extrapolated energies are obtained from a quadratic fit to the energies for $\epsilon_1 \leq 20 \times 10^{-5}$ Ha. The relatively slow convergence of the TZ basis is due to an integrals file that uses natural orbitals from a short, unconverged, FCIQMC calculation. The energies in the TZ' rows were obtained using the same basis, but with approximate natural orbitals obtained from an SHCI variational wavefunction. They converge much faster. The wall times are not shown for some calculations because these were done on a different computer architecture so a meaningful comparison is not possible.

studies.^{49–59} Recently, some of us performed DMRG-SCF calculations⁴⁸ on the model cluster with the cc-pVDZ basis set using an active space of (28e, 22o). The initial orbitals were obtained by using the HOMO–13 to LUMO+7 canonical Hartree–Fock orbitals, which were subsequently optimized using the DMRG-SCF method. Here we perform the SHCI calculations on the converged orbitals obtained at the end of the converged DMRG-SCF calculations. The results in Table 3

show that both the singlet and the triplet energies converge to better than 1 mHa accuracy in only 37 s on a single node.

VI. CONCLUSIONS

We have introduced stochastic and semistochastic implementations of multireference Epstein–Nesbet perturbation theory, for computing the expectation value of the perturbative correction to the variational energy of a multideterminant

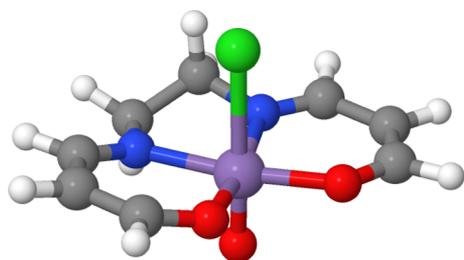


Figure 4. Geometry of the chlorine containing neutral oxo-Mn(salen) cluster was optimized with (10e, 10o)-CASSCF and a 6-31G* basis by Ivanic et al.⁴⁵ The Mn, Cl, N, O, C, and H atoms are shown in purple, green, blue, red, gray and white, respectively.

Table 3. Comparison of the DMRG and SHCI Energies ($E + 2251$) Ha of the Singlet and Triplet States of Mn-Salen^a

sym	N_v	energy (Ha)			wall time (s)
		var	PT	DMRG ref (48)	
¹ A	232484	-0.7880	-0.7980(7)	-0.7991	37
³ A	208334	-0.7910	-0.7994(8)	-0.8001	32

^aDMRG was performed with an $M = 2000$ and SHCI was performed with $\epsilon_1 = 2 \times 10^{-4}$ Ha, $\epsilon_2 = 1 \times 10^{-8}$ Ha and $N_d = 200$. The wall time needed to perform the SHCI calculation on a single computer node is shown in the final column (see text for additional details).

wave function without storing all the contributing determinants. In addition to completely removing the memory bottleneck, the semistochastic algorithm is faster than the fully deterministic algorithm for most systems if a stochastic noise of 0.1 mHa is acceptable.

Our method is capable of efficiently computing the correlation energies of very large active spaces, as we have demonstrated by computing the energies of the challenging, multireference systems Mn-Salen (28e, 22o) and Cr_2 (12e, 19o). For all systems studied we obtained correlation energies accurate to within 1 mHa. In the case of the first-row dimers and Mn-Salen we compared to FCIQMC and DMRG energies in the literature. For Cr_2 there are no published values, but one of the positive features of our method is that one can reliably check the convergence within the method itself.

Having removed the memory bottleneck in the perturbative step, the largest memory requirement comes from storing the Hamiltonian in the variational space. The next step is to create an efficient method for obtaining the variational wave function without storing the Hamiltonian. Other research directions include the optimization of the orbitals within the CAS space, and the calculation of excited states.

AUTHOR INFORMATION

Corresponding Author

*E-mail: sanshar@gmail.com.

ORCID

Sandeep Sharma: 0000-0002-6598-8887

Funding

S.S. acknowledges the startup package from the University of Colorado. A.A.H. and C.J.U. were supported in part by NSF Grant ACI-1534965.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The calculations made use of the facilities of the Max Planck Society's Rechenzentrum Garching. C.J.U. thanks Garnet Chan for valuable discussions.

REFERENCES

- Olsen, J.; Roos, B. O.; Jorgensen, P.; Jensen, H. J. A. Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces. *J. Chem. Phys.* **1988**, *89*, 2185.
- Malmqvist, P. A.; Rendell, A.; Roos, B. O. The restricted active space self-consistent-field method, implemented with a split graph unitary group approach. *J. Phys. Chem.* **1990**, *94*, 5477–5482.
- Ma, D.; Li Manni, G.; Gagliardi, L. The generalized active space concept in multiconfigurational self-consistent field methods. *J. Chem. Phys.* **2011**, *135*, 044128.
- White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **1992**, *69*, 2863.
- White, S. R. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1993**, *48*, 10345.
- Chan, G. K.-L.; Sharma, S. The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.* **2011**, *62*, 465–481.
- Booth, G. H.; Thom, A. J.; Alavi, A. Fermion monte carlo without fixed nodes: A game of life, death, and annihilation in slater determinant space. *J. Chem. Phys.* **2009**, *131*, 054106.
- Cleland, D.; Booth, G. H.; Alavi, A. Communications: Survival of the fittest: Accelerating convergence in full configuration-interaction quantum monte carlo. *J. Chem. Phys.* **2010**, *132*, 041103.
- Petruzielo, F. R.; Holmes, A. A.; Changlani, H. J.; Nightingale, M. P.; Umrigar, C. J. Semistochastic projector monte carlo method. *Phys. Rev. Lett.* **2012**, *109*, 230201.
- Werner, H.-J.; Reinsch, E. A. The self-consistent electron pairs method for multiconfiguration reference state functions. *J. Chem. Phys.* **1982**, *76*, 3144.
- Knowles, P. J.; Werner, H.-J. Internally contracted multi-configuration-reference configuration interaction calculations for excited states. *Theor. Chim. Acta* **1992**, *84*, 95.
- Shamasundar, K. R.; Knizia, G.; Werner, H.-J. A new internally contracted multi-reference configuration interaction method. *J. Chem. Phys.* **2011**, *135*, 054101.
- Andersson, K.; Malmqvist, P. A.; Roos, B. O.; Sadlej, A. J.; Wolinski, K. Second-order perturbation theory with a CASSCF reference function. *J. Phys. Chem.* **1990**, *94*, 5483–5488.
- Fink, R. F. The multi-reference retaining the excitation degree perturbation theory: A size-consistent, unitary invariant, and rapidly convergent wavefunction based ab initio approach. *Chem. Phys.* **2009**, *356*, 39–46.
- Angeli, C.; Cimraglia, R.; Evangelisti, S.; Leininger, T.; Malrieu, J.-P. Introduction of n-electron valence states for multireference perturbation theory. *J. Chem. Phys.* **2001**, *114*, 10252.
- Hirao, K. Multireference Moeller-Plesset method. *Chem. Phys. Lett.* **1992**, *190*, 374–380.
- Lyakh, D. I.; Musiał, M.; Lotrich, V. F.; Bartlett, R. J. Multireference nature of chemistry: the coupled-cluster view. *Chem. Rev.* **2012**, *112*, 182–243.
- Evangelista, F. A.; Hanauer, M.; Köhn, A.; Gauss, J. A sequential transformation approach to the internally contracted multireference coupled cluster method. *J. Chem. Phys.* **2012**, *136*, 204108.
- Sharma, S.; Jeanmairet, G.; Alavi, A. Quasi-degenerate perturbation theory using matrix product states. *J. Chem. Phys.* **2016**, *144*, 034103.
- Sharma, S.; Alavi, A. Multireference linearized coupled cluster theory for strongly correlated systems using matrix product states. *J. Chem. Phys.* **2015**, *143*, 102815.
- Jeanmairet, G.; Sharma, S.; Alavi, A. Stochastic multi-reference perturbation theory with application to the linearized coupled cluster method. *J. Chem. Phys.* **2017**, *146*, 044107.

- (22) Ivanic, J.; Ruedenberg, K. Identification of deadwood in conformation spaces through general direct conformation interaction. *Theor. Chem. Acc.* **2001**, *106*, 339.
- (23) Huron, B.; Malrieu, J.; Rancurel, P. Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions. *J. Chem. Phys.* **1973**, *58*, 5745–5759.
- (24) Buenker, R. J.; Peyerimhoff, S. D. Individualized configuration selection in CI calculations with subsequent energy extrapolation. *Theor. Chim. Acta* **1974**, *35*, 33–58.
- (25) Evangelisti, S.; Daudey, J.-P.; Malrieu, J.-P. Convergence of an improved CIPSI algorithm. *Chem. Phys.* **1983**, *75*, 91–102.
- (26) Harrison, R. J. *J. Chem. Phys.* **1991**, *94*, 5021.
- (27) Steiner, M. M.; Wenzel, W.; Wilson, K. G.; Wilkins, J. W. The efficient treatment of higher excitations in ci calculations: A comparison of exact and approximate results. *Chem. Phys. Lett.* **1994**, *231*, 263.
- (28) Wenzel, W.; Steiner, M.; Wilson, K. G. Multireference basis-set reduction. *Int. J. Quantum Chem.* **1996**, *60*, 1325.
- (29) Neese, F. A spectroscopy oriented configuration interaction procedure. *J. Chem. Phys.* **2003**, *119*, 9428.
- (30) Abrams, M. L.; Sherrill, C. D. Important configurations in configuration interaction and coupled-cluster wave functions. *Chem. Phys. Lett.* **2005**, *412*, 121.
- (31) Bytautas, L.; Ruedenberg, K. A-priori identification of configurational deadwood. *Chem. Phys.* **2009**, *356*, 64.
- (32) Evangelista, F. A. A driven similarity renormalization group approach to quantum many-body problems. *J. Chem. Phys.* **2014**, *140*, 054109.
- (33) Knowles, P. J. Compressive sampling in configuration interaction wavefunctions. *Mol. Phys.* **2015**, *113*, 1655.
- (34) Schriber, J. B.; Evangelista, F. A. Communication: An adaptive configuration interaction approach for strongly correlated electrons with tunable accuracy. *J. Chem. Phys.* **2016**, *144*, 161106.
- (35) Tubman, N. M.; Lee, J.; Takeshita, T. Y.; Head-Gordon, M.; Whaley, K. B. A deterministic alternative to the full configuration interaction quantum monte carlo method. *J. Chem. Phys.* **2016**, *145*, 044112.
- (36) Liu, W.; Hoffmann, M. R. iCI: Iterative CI toward full CI toward full CI. *J. Chem. Theory Comput.* **2016**, *12*, 1169.
- (37) Caffarel, M.; Applegate, T.; Giner, E.; Scemama, A. Using CIPSI nodes in diffusion Monte Carlo. 2016; <https://arxiv.org/abs/1607.06742v2>.
- (38) Holmes, A. A.; Tubman, N. M.; Umrigar, C. J. Heat-bath configuration interaction: An efficient selected ci algorithm inspired by heat-bath sampling. *J. Chem. Theory Comput.* **2016**, *12*, 3674.
- (39) Epstein, P. S. The Stark effect from the point of view of schrodinger's quantum theory. *Phys. Rev.* **1926**, *28*, 695.
- (40) Nesbet, R. K. Configuration interaction in orbital theories. *Proc. R. Soc. London, Ser. A* **1955**, *230*, 312.
- (41) Schäfer, A.; Horn, H.; Ahlrichs, R. "fully optimized contracted gaussian basis sets for atoms Li to Kr. *J. Chem. Phys.* **1992**, *97*, 2571–2577.
- (42) Holmes, A. A.; Changlani, H. J.; Umrigar, C. J. Efficient heat-bath sampling in fock space. *J. Chem. Theory Comput.* **2016**, *12*, 1561.
- (43) The logarithmic corrections come from having to do binary searches to check whether a determinant is already present in the variational space, or, in the space of connected determinants that have already been generated.
- (44) Knowles, P.; Handy, N. A new determinant-based full configuration interaction method. *Chem. Phys. Lett.* **1984**, *111*, 315–321.
- (45) Walker, A. J. An efficient method for generating discrete random variables with general distributions. *ACM Trans. on Math. Software (TOMS)* **1977**, *3*, 253–256.
- (46) Kronmal, R. A.; Peterson, A. V., Jr On the alias method for generating random variables from a discrete distribution. *Am. Stat.* **1979**, *33*, 214–218.
- (47) Cleland, D.; Booth, G. H.; Overy, C.; Alavi, A. Taming the First-Row Diatomics: A Full Configuration Interaction Quantum Monte Carlo Study. *J. Chem. Theory Comput.* **2012**, *8*, 4138–4152.
- (48) Sharma, S.; Knizia, G.; Guo, S.; Alavi, A. Combining internally contracted states and matrix product states to perform multireference perturbation theory. *J. Chem. Theory Comput.* **2016**, *13*, 488.
- (49) Ivanic, J.; Collins, J. R.; Burt, S. K. Theoretical Study of the Low Lying Electronic States of oxoX(salen) (X = Mn, Mn²⁺, Fe, and Cr³⁺) Complexes. *J. Phys. Chem. A* **2004**, *108*, 2314–2323.
- (50) Teixeira, F.; Mosquera, R. A.; Melo, A.; Freire, C.; Cordeiro, M. N. D. S. Principal component analysis of Mn(salen) catalysts. *Phys. Chem. Chem. Phys.* **2014**, *16*, 25364–25376.
- (51) Stein, C. J.; Reiher, M. Automated Selection of Active Orbital Spaces. *J. Chem. Theory Comput.* **2016**, *12*, 1760.
- (52) Bogaerts, T.; Wouters, S.; Van Der Voort, P.; Van Speybroeck, V. Mechanistic Investigation on Oxygen Transfer with the Manganese-Salen Complex. *ChemCatChem* **2015**, *7*, 2711–2719.
- (53) Ma, D.; Li Manni, G.; Gagliardi, L. The generalized active space concept in multiconfigurational self-consistent field methods. *J. Chem. Phys.* **2011**, *135*, 044128.
- (54) Sears, J. S.; Sherrill, C. D. The electronic structure of oxo-Mn(salen): single-reference and multireference approaches. *J. Chem. Phys.* **2006**, *124*, 144314.
- (55) Wouters, S.; Bogaerts, T.; Van Der Voort, P.; Van Speybroeck, V.; Van Neck, D. Communication: DMRG-SCF study of the singlet, triplet, and quintet states of oxo-Mn(salen). *J. Chem. Phys.* **2014**, *140*, 241103.
- (56) Linde, C.; Åkermark, B.; Norrby, P.-O.; Svensson, M. Timing Is Critical: Effect of Spin Changes on the Diastereoselectivity in Mn(salen)-Catalyzed Epoxidation. *J. Am. Chem. Soc.* **1999**, *121*, 5083–5084.
- (57) Abashkin, Y. G.; Collins, J. R.; Burt, S. K. (Salen)Mn(III)-Catalyzed Epoxidation Reaction as a Multichannel Process with Different Spin States. Electronic Tuning of Asymmetric Catalysis: A Theoretical Study. *Inorg. Chem.* **2001**, *40*, 4040–4048.
- (58) Abashkin, Y. G.; Burt, S. K. (Salen)Mn(III) Compound as a Nonpeptidyl Mimic of Catalase: DFT Study of the Metal Oxidation by a Peroxide Molecule. *J. Phys. Chem. B* **2004**, *108*, 2708–2711.
- (59) Abashkin, Y. G.; Collins, J. R.; Burt, S. K. (Salen)Mn(III)-Catalyzed Epoxidation Reaction as a Multichannel Process with Different Spin States. Electronic Tuning of Asymmetric Catalysis: A Theoretical Study. *Inorg. Chem.* **2001**, *40*, 4040–4048.