King's Research Portal

DOI:

[10.1063/1.4986963](https://doi.org/10.1063/1.4986963)

Document Version

Publisher's PDF, also known as Version of record

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Blunt, N. S., Booth, G. H., & Alavi, A. (2017). Density matrices in full configuration interaction quantum Monte Carlo: Excited states, transition dipole moments, and parallel distribution. *Journal of Chemical Physics*, 146(24), [244105]. <https://doi.org/10.1063/1.4986963>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Density matrices in full configuration interaction quantum Monte Carlo: Excited states, transition dipole moments, and parallel distribution

N. S. Blunt, George H. Booth, and Ali Alavi

Citation: *The Journal of Chemical Physics* **146**, 244105 (2017); doi: 10.1063/1.4986963

View online: <http://dx.doi.org/10.1063/1.4986963>

View Table of Contents: <http://aip.scitation.org/toc/jcp/146/24>

Published by the [American Institute of Physics](#)

Articles you may be interested in

[Time-dependent N-electron valence perturbation theory with matrix product state reference wavefunctions for large active spaces and basis sets: Applications to the chromium dimer and all-trans polyenes](#)

The Journal of Chemical Physics **146**, 244102 (2017); 10.1063/1.4986975

[Dealing with the exponential wall in electronic structure calculations](#)

The Journal of Chemical Physics **146**, 194107 (2017); 10.1063/1.4983207

[Strong correlation in incremental full configuration interaction](#)

The Journal of Chemical Physics **146**, 224104 (2017); 10.1063/1.4985566

[Communication: Density functional theory embedding with the orthogonality constrained basis set expansion procedure](#)

The Journal of Chemical Physics **146**, 211101 (2017); 10.1063/1.4984777

[Quantifying uncertainty in the chemical master equation](#)

The Journal of Chemical Physics **146**, 244103 (2017); 10.1063/1.4986762

[Application of the spectral element method to the solution of the multichannel Schrödinger equation](#)

The Journal of Chemical Physics **146**, 244106 (2017); 10.1063/1.4987026

**PHYSICS
TODAY**

**COMPLETELY
REDESIGNED!**

Physics Today Buyer's Guide
Search with a purpose.

Density matrices in full configuration interaction quantum Monte Carlo: Excited states, transition dipole moments, and parallel distribution

N. S. Blunt,^{1,2,a)} George H. Booth,³ and Ali Alavi^{1,2,b)}

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

²University Chemical Laboratory, Lensfield Road, Cambridge CB2 1EW, United Kingdom

³Department of Physics, King's College London, Strand, London WC2R 2LS, United Kingdom

(Received 2 April 2017; accepted 7 June 2017; published online 26 June 2017)

We present developments in the calculation of reduced density matrices (RDMs) in the full configuration interaction quantum Monte Carlo (FCIQMC) method. An efficient scheme is described to allow storage of RDMs across distributed memory, thereby allowing their calculation and storage in large basis sets. We demonstrate the calculation of RDMs for general states by using the recently introduced excited-state FCIQMC approach [N. S. Blunt *et al.*, J. Chem. Phys. **143**, 134117 (2015)] and further introduce calculation of transition density matrices in the method. These approaches are combined to calculate excited-state dipole and transition dipole moments for heteronuclear diatomic molecules, including LiH, BH, and MgO, and initiator error is investigated in these quantities. *Published by AIP Publishing*, [<http://dx.doi.org/10.1063/1.4986963>]

I. INTRODUCTION

The estimation of general expectation values in quantum Monte Carlo (QMC) methods is a challenging task, one that has been achieved with varying levels of success across the range of QMC methods employed and solutions attempted. In diffusion Monte Carlo (DMC)¹ where sampling is performed in real space, the energy can be estimated simply through the use of a mixed estimator, but approaches to calculate the expectation value of operators that do not commute with the Hamiltonian are considerably less straightforward. The estimation of such expectation values remains a difficult task in general.^{1–4}

The full configuration interaction quantum Monte Carlo (FCIQMC) method^{5–7} has much in common with DMC and other projector QMC Carlo methods. However, a crucial feature that separates FCIQMC from some such methods is that sampling is performed in a discrete basis set, usually comprised of basis states that share the antisymmetry of the exact fermionic wave function. While the use of a discrete basis set necessarily leads to a finite basis set error, it also leads to many of the methods' favourable properties. Perhaps most notably, the discrete space allows for efficient annihilation that significantly ameliorates the sign problem.^{8,9} Furthermore, in a discrete space, the so-called replica trick^{10,11} results in an incredibly simple and effective approach for estimating general observables.

The replica trick was first used in FCIQMC by Overy *et al.*¹² to stochastically sample the two-particle reduced density matrix (RDM) in an unbiased manner. This was in contrast to previous efforts to compute the 2-RDM, which suffered from large systematic errors.¹³ Replica-sampled

2-RDMs were subsequently used by Thomas *et al.*, first to compute nuclear forces, dipole moments, and polarizabilities for various small molecules¹⁴ and subsequently to perform orbital optimization in FCIQMC-based extensions to the multi-configurational self-consistent field approach, allowing the use of very large active spaces.¹⁵ RDMs have also been fundamental to a recent study of explicitly correlated approaches in the FCIQMC method.¹⁶ From the range of applications even at this early stage, it is clear that the calculation of RDMs in FCIQMC will be of key importance going forward.

There also exists a separate but important family of methods that construct the 2-RDMs without using the wave function.^{17–19} These methods instead rely on imposing N-representability conditions, with a lower bound on the energy improving as more conditions are enforced. In contrast, FCIQMC directly uses a stochastic sampling of the wave function to construct 2-RDMs, becoming exact as the initiator approximation^{6,20} is improved upon, and in the limit of longer sampling.

Current implementations of the FCIQMC algorithm make use of its stochastic nature to allow efficient large-scale parallelization. The walkers that sample the underlying wave function are distributed among processors, thus allowing the most substantial memory requirements to be distributed.²¹ Unfortunately, this has not been true for the reduced density matrices (RDMs), which in previously described implementations have been stored in their entirety on every processing core. This has limited the scope of applications thus far, especially with the emphasis on large active spaces that the method can effectively sample.

Meanwhile, the current authors recently introduced an excited-state approach within the FCIQMC framework.²² In this approach, multiple FCIQMC simulations are performed simultaneously, one for each state to be sampled. Simulations

^{a)}Electronic mail: nicksblunt@gmail.com

^{b)}Electronic mail: alavi@fkf.mpg.de

representing higher-energy states are orthogonalized against the stochastic snapshots of wave functions in simulations representing lower-energy states, thus preventing collapse to the ground state. It was shown that this approach appears to be possible without incurring detectable systematic errors. This therefore provides a very simple and accurate approach to sampling excited states, which we believe shows significant promise.

In this article, we build upon these recent developments by introducing an extension to the computation of unbiased excited-state reduced density matrices and transition density matrices (TDMs). These latter objects are required in order to calculate transition amplitudes induced by different operators between excited states (such as an optical excitation) and critical in the connection of results to spectroscopic experiments. In order to circumvent the limiting storage of these sampled objects, we will furthermore describe an efficient approach to distribute such objects efficiently across distributed memory architecture.

In Sec. II, we provide a brief overview of the excited-state FCIQMC method and of the calculation of RDMs within the ground-state FCIQMC approach. In Sec. III, the theory underlying the unbiased estimation of excited-state RDMs and TDMs is presented. Section IV then describes an efficient parallelized implementation of RDMs within a large-scale FCIQMC code, and results are presented in Sec. V.

II. BACKGROUND THEORY

A. The excited-state FCIQMC method

We begin with a brief overview of the recently introduced excited-state FCIQMC method.²² The ground-state FCIQMC algorithm is a projector QMC method that achieves a stochastic sampling of the ground-state wave function. This is accomplished by repeated application of a projection operator, $\hat{P}$, to some initial wave function. In FCIQMC, $\hat{P}$ is defined by

$$\hat{P} = 1 - \Delta\tau(\hat{H} - S1), \quad (1)$$

where $\hat{H}$ is the Hamiltonian, $\Delta\tau$ is a small imaginary-time step, and S is a diagonal shift parameter, varied slowly throughout the simulation to allow some control over the wave function normalization. We work in a discrete basis set, usually of Slater determinants, $\{|D_i\rangle\}$, and the projection operator can then be expressed as $\hat{P} = \sum_{ij} P_{ij}|i\rangle\langle j|$. The FCIQMC method then primarily consists of an algorithm to sample elements of $\mathbf{P}\psi$ (and through iteration therefore of $(\mathbf{P})^m\psi$), for some initial ψ , in an extremely efficient and unbiased manner. In the limit of large m , a stochastic sampling of the ground state is achieved. Detailed introductions to the FCIQMC algorithm, by which this is achieved, exist elsewhere.^{5,8,21}

The excited-state FCIQMC method consists of a small extension to the ground-state approach, which prevents collapse to the ground-state through a simple orthogonalization procedure. Suppose the exact ground-state wave function was known and denoted as $|\Phi^0\rangle$, then the first excited-state could be converged upon by repeated application of the operator

$$\hat{P}_{\text{first excited}} = [1 - |\Phi^0\rangle\langle\Phi^0|]\hat{P}. \quad (2)$$

A single iteration of a QMC algorithm for the first-excited state could then be to apply $\hat{P}$ via FCIQMC's spawning rules, followed by applying $1 - |\Phi^0\rangle\langle\Phi^0|$ to the resulting wave function. This latter projection exactly removes the ground-state contribution, and the next-lowest energy state would be converged upon instead. In practice, this is not possible because the exact ground state is unknown. However, the basic FCIQMC method provides a procedure to stochastically sample this state. Therefore, by performing two FCIQMC simulations simultaneously, the first-excited state may be sampled in the second simulation by orthogonalizing against the instantaneous FCIQMC-sampled ground-state in the first simulation, a simple task to perform computationally. While one may worry that orthogonalizing against stochastic snapshots of a state is not equivalent to orthogonalizing against an exact wave function, we have been unable to detect any discrepancy whatsoever resulting from this difference, and this simple procedure appears unbiased beyond any achievable level of statistical accuracy.²²

The excited-state FCIQMC algorithm to sample the lowest eigenstates of $\hat{H}$ can therefore be summarized as follows. Perform multiple simulations simultaneously, one for each state to be sampled. Let the (unnormalized) wave function sampled by simulation n and at imaginary-time τ be denoted as $|\Psi^n(\tau)\rangle$. Then the evolution sampled by the various simulations is given by

$$|\Psi^n(\tau + \Delta\tau)\rangle = \hat{O}^n(\tau + \Delta\tau)\hat{P}|\Psi^n(\tau)\rangle, \quad (3)$$

where

$$\hat{O}^n(\tau) = 1 - \sum_{m < n} \frac{|\Psi^m(\tau)\rangle\langle\Psi^m(\tau)|}{\langle\Psi^m(\tau)|\Psi^m(\tau)\rangle}. \quad (4)$$

A single iteration therefore first consists of applying the FCIQMC ground-state projection operator $\hat{P}$ to each simulation, which is achieved by the usual FCIQMC algorithm, and includes performing all annihilation. Only then is orthogonalization performed, with each simulation being orthogonalized against *all lower-energy* states. This fully describes the excited-state adaptation to the FCIQMC algorithm, which can be performed at essentially linear cost with the number of desired states, since the orthogonalization step requires close to negligible computational overhead.

B. Unbiased RDM estimation in FCIQMC

The second-order (symmetric) reduced density matrix (2-RDM) for a state $|\Phi^n\rangle$ is defined as

$$\Gamma_{pq,rs}^n = \langle\Phi^n|a_p^\dagger a_q^\dagger a_s a_r|\Phi^n\rangle, \quad (5)$$

where p, q, r , and s denote spin-orbital labels. In the rest of this article, we will only be concerned with states representing the eigenstates of the Hamiltonian, $\hat{H}$. We therefore let $|\Phi^n\rangle$ for $n > 0$ denote the *exact* n th excited state of $\hat{H}$, while $|\Phi^0\rangle$ denotes the exact ground state, and the determinant expansion of a state $|\Phi^n\rangle$ is expressed as

$$|\Phi^n\rangle = \sum_i \phi_i^n |D_i\rangle. \quad (6)$$

ϕ_i^n will be assumed to be real throughout, although complex FCIQMC extensions are straightforward.⁷ The RDM for

energy eigenstate $|\Phi^n\rangle$ can then be expressed as

$$\Gamma_{pq,rs}^n = \sum_{ij} \phi_i^n \phi_j^n \langle D_i | a_p^\dagger a_q^\dagger a_s a_r | D_j \rangle. \quad (7)$$

Let FCIQMC coefficients be denoted as C_i^n . Then the FCIQMC wave function for state n is $|\Psi^n\rangle = \sum_i C_i^n |D_i\rangle$, with C_i^n denoting the (signed) number of walkers on determinant $|D_i\rangle$. The FCIQMC and excited-state FCIQMC algorithms are unbiased at large enough walker numbers to a very good approximation, meaning that the average value of C_i^n throughout the FCIQMC simulation should approach the exact coefficient, ϕ_i^n . A more precise statement is that, once convergence has been reached,

$$E[C_i^n] = \phi_i^n, \quad (8)$$

where $E[\dots]$ denotes the expectation value, equivalent to the long time-averaged walker population. We ignore for now that the FCIQMC wave function is not normalized.

A simple approach to sample an (unnormalized) RDM within FCIQMC would be through

$$\Gamma_{pq,rs}^n \approx \sum_{ij} C_i^n C_j^n \langle D_i | a_p^\dagger a_q^\dagger a_s a_r | D_j \rangle. \quad (9)$$

However, some thought quickly shows that this is not correct: the expectation of this expression does not give back the desired exact 2-RDM because $E[C_i^n C_j^n] \neq \phi_i^n \phi_j^n$ unless C_i^n and C_j^n are uncorrelated with a zero covariance, which they will not be. This is particularly clear for diagonal elements: expressing $C_i^n = \phi_i^n + \delta_i^n$, then $E[(C_i^n)^2] = (\phi_i^n)^2 + (\delta_i^n)^2$, which is always greater than the desired exact value for a non-zero δ_i^n . Therefore, averaging this expression over many iterations will result in a biased 2-RDM estimate. If one could average all FCIQMC coefficients over the entire simulation, then the associated errors, δ_i^n , could be reduced to a point where any bias is negligible, but this is not feasible due to the memory requirements of storing all averaged coefficients rather than the far smaller set of instantaneously occupied determinants. In an earlier implementation, Booth *et al.* attempted a scheme where the FCIQMC coefficients are partially averaged over a reduced number of iterations, but an unacceptable bias was still found to occur.¹³

This problem can be resolved by a replica sampling approach, whereby a second identical FCIQMC simulation is performed concurrently, starting from separate random number generator (RNG) seeds. The two FCIQMC simulations are performed for each state entirely independently, guaranteeing that the two wave function estimates will be uncorrelated. Let coefficients from the first and second replica simulations be denoted by $C_i^{1,n}$ and $C_i^{2,n}$, respectively. Then because the two simulations are uncorrelated, $E[C_i^{1,n} C_j^{2,n}] = \phi_i^n \phi_j^n$, and so

$$\Gamma_{pq,rs}^n = E \left[\sum_{ij} C_i^{1,n} C_j^{2,n} \langle D_i | a_p^\dagger a_q^\dagger a_s a_r | D_j \rangle \right]. \quad (10)$$

While storing and averaging all coefficients in the FCIQMC wave function is not feasible, averaging all elements of $\Gamma_{pq,rs}^n$ is usually much more reasonable. In our previous implementation, we have therefore stored the entire 2-RDM in memory for each parallel process in a 2-dimensional non-sparse array. While possible for many small single-particle basis sets, this

quickly becomes unreasonable, and Sec. IV is therefore dedicated to introducing a sparse and distributed implementation. We lastly remark that the above estimate for the 2-RDM in Eq. (10) can be normalized in a straightforward fashion at the end of a calculation by enforcing that the trace of the 2-RDM must be equal to $N(N-1)/2$, as necessarily found in the exact 2-RDM.

C. Practical considerations for sampling RDMs

In this section, we describe some additional details and efficiency improvements regarding the practical calculation of 2-RDMs in FCIQMC which are relevant for the discussion on their distributed storage in Sec. IV. More details on many of these practical aspects can be found in Ref. 12.

In general, we choose not to explicitly include all terms in Eq. (10) exactly. Doing so would involve an expensive $O(N^2 M^2)$ operation for each occupied determinant, where N is the number of electrons and M is the number of spatial orbitals. Instead we exploit the fact that the FCIQMC spawning procedure already involves randomly generating and sampling from the distribution of up-to-double excitations and that $\{|D_i\rangle, |D_j\rangle\}$ pairs also only contribute to Eq. (10) if within a double excitation of each other. The FCIQMC spawning procedure is therefore used to stochastically sample contributions to $\Gamma_{pq,rs}^n$, with contributions carefully weighted to avoid biases.

However, spawning attempts in FCIQMC are not performed for diagonal elements. Therefore, contributions to diagonal elements of the 2-RDM,

$$\Gamma_{pq,pq}^n = E \left[\sum_{\{p,q\} \in i} C_i^{1,n} C_i^{2,n} \right], \quad (11)$$

are added in exactly and explicitly. The summation here is performed over all determinants for which spin-orbitals p and q are both occupied. As a result of Brillouin's theorem, single excitations from the Hartree-Fock determinant in a canonical basis are also never generated, and all contributions to $\Gamma_{pq,rs}^n$ involving the Hartree-Fock determinant are therefore included explicitly. Doing this also greatly reduces stochastic noise in the 2-RDM estimate, since the Hartree-Fock determinant usually has a large amplitude. While it is possible that zero Hamiltonian matrix elements could result in the simulation "missing" further contributions to the RDM sampling, we have not found this to be a problem in practice.

Finally we note that instead of simply adding in contributions to the 2-RDM at every iteration, a more sophisticated block-averaging scheme is used. In this, the FCIQMC coefficients for the occupied determinants are averaged over "blocks" of many iterations. Contributions to diagonal elements of the 2-RDM and contributions involving the Hartree-Fock determinant are only added in at the end of each averaging block, using the FCIQMC coefficients averaged over the block length. All other contributions are added in on each iteration. This block averaging scheme was originally devised to help reduce biases in Eq. (9) before the replica sampling was in use,¹³ as discussed above. However, it also significantly reduces the computational cost

associated with calculating diagonal and Hartree–Fock contributions to $\Gamma_{pq,rs}^n$ since these contributions are now only calculated infrequently. The calculation of diagonal elements is particularly expensive, with an $\mathcal{O}(N^2)$ cost per determinant, and so this averaging scheme is greatly beneficial computationally and provides identical results, despite not being required to avoid biased sampling when using the replica approach.

III. EXCITED-STATE AND TRANSITION DENSITY MATRICES

A. Excited-state density matrices

The calculation of excited-state RDMs ($\Gamma_{pq,rs}^n$ with $n > 0$) in FCIQMC can be performed in exactly the same manner as for the ground state. The RDM estimator, Eq. (10), depends only on the walker coefficients, $C^{1,n}$ and $C^{2,n}$. Coefficients for $n > 0$ are stored and processed in exactly the same way as for $n = 0$, and RDM sampling is therefore unchanged.

Two replica simulations must now be performed for each state being sampled, whereas the initial presentation of excited-state FCIQMC used only one such simulation for each state. However, this does not cause any difficulties beyond the additional computational effort required. One effectively performs two excited-state FCIQMC simulations concurrently, with walker coefficients in the first labeled $C^{1,n}$ and those in the second labeled $C^{2,n}$, as above. Orthogonalization is then performed equally in both simulations, i.e., via the operators

$$O^{R,n} = I - \sum_{m < n} \frac{C^{R,m} (C^{R,m})^\dagger}{(C^{R,m})^\dagger C^{R,m}}, \quad R \in \{1, 2\}, \quad (12)$$

with R denoting the replica label and n denoting the state to which the operator is applied in the excited-state FCIQMC procedure. This choice ensures that all simulations with $R = 1$ are uncorrelated from all those with $R = 2$, and therefore avoids sampling biases in the RDM estimates, exactly analogously to the ground-state RDM estimation in Eq. (10). We note that one could modify the orthogonalization operator to make use of both replicas as follows:

$$O^n \stackrel{?}{=} I - \sum_{m < n} \frac{C^{1,m} (C^{2,m})^\dagger}{(C^{1,m})^\dagger C^{2,m}}. \quad (13)$$

The motivation behind this is the desire to remove a theoretical non-linear bias in the operator since $E[(C^{1,m})^\dagger C^{2,m}] = \langle \Phi^m | \Phi^m \rangle$, while $E[(C^{1,m})^\dagger C^{1,m}] \neq \langle \Phi^m | \Phi^m \rangle$. However, in practice, this approach introduces correlations between walkers from the two replicas that biases RDM estimates, and as noted previously, a bias from the operator in Eq. (12) has not been detected in practice.

B. Transition density matrices

A two-body transition density matrix (TDM) is defined by

$$\Gamma_{pq,rs}^{nm} = \langle \Phi^m | a_p^\dagger a_q^\dagger a_s a_r | \Phi^n \rangle, \quad (14)$$

and the equivalent one-body TDM is defined by

$$\gamma_{p,q}^{nm} = \langle \Phi^m | a_p^\dagger a_q | \Phi^n \rangle, \quad (15)$$

$$= \frac{1}{N-1} \sum_a \Gamma_{pa,qa}^{nm}. \quad (16)$$

These objects are of significance in spectroscopy and other processes where transitions are induced between eigenstates due to some operator. For instance, optical transitions are driven by the electric field, and therefore the transition dipole moment between states m and n can be calculated as

$$t_{nm} = \sum_{pq} \gamma_{p,q}^{nm} \langle p | \hat{r} | q \rangle, \quad (17)$$

where $|p\rangle$ denotes a single-particle spin orbital and $\hat{r}$ denotes the position operator. From this, using Fermi's golden rule, transition probabilities can be obtained. Note that the TDM (and hence, e.g., transition dipole moment) is only determined up to an arbitrary sign, and physical observables must depend on the square modulus of this quantity.

One may suppose that TDMs can be estimated using similar estimators as introduced above, by

$$\Gamma_{pq,rs}^{nm} = E \left[\sum_{ij} C_i^{1,m} C_j^{2,n} \langle D_i | a_p^\dagger a_q^\dagger a_s a_r | D_j \rangle \right], \quad (18)$$

or equivalently by

$$\Gamma_{pq,rs}^{nm} = E \left[\sum_{ij} C_i^{2,m} C_j^{1,n} \langle D_i | a_p^\dagger a_q^\dagger a_s a_r | D_j \rangle \right]. \quad (19)$$

However, because the FCIQMC wave functions are not normalized, the above objects also require explicit normalization. Non-transition (symmetric) 2-RDMs are normalized by noting that their trace must equal $N(N-1)/2$. TDMs are not so simple to correct since their traces equal 0 due to the orthogonality of the states. Instead they must be normalized directly by calculating the norm of the FCIQMC wave functions, which should equal 1.

Defining the normalization of FCIQMC wave function by $E[|\Psi^{R,m}\rangle] = A^{R,m} |\Phi^{R,m}\rangle$, it is anticipated that one way that the normalization of the TDMs could occur is with the following scheme:

$$\Gamma_{pq,rs}^{nm}[1] \equiv \frac{E[\langle \Psi^{1,m} | a_p^\dagger a_q^\dagger a_s a_r | \Psi^{2,n} \rangle]}{\sqrt{E[\langle \Psi^{1,m} | \Psi^{2,m} \rangle] E[\langle \Psi^{1,n} | \Psi^{2,n} \rangle]}} \quad (20)$$

$$= \frac{A^{1,m} A^{2,n}}{\sqrt{A^{1,m} A^{2,m} A^{1,n} A^{2,n}}} \Gamma_{pq,rs}^{nm}, \quad (21)$$

$$\Gamma_{pq,rs}^{nm}[2] \equiv \frac{E[\langle \Psi^{2,m} | a_p^\dagger a_q^\dagger a_s a_r | \Psi^{1,n} \rangle]}{\sqrt{E[\langle \Psi^{1,m} | \Psi^{2,m} \rangle] E[\langle \Psi^{1,n} | \Psi^{2,n} \rangle]}} \quad (22)$$

$$= \frac{A^{2,m} A^{1,n}}{\sqrt{A^{1,m} A^{2,m} A^{1,n} A^{2,n}}} \Gamma_{pq,rs}^{nm}. \quad (23)$$

However, this would only be sufficient if both replicas are always normalized equally, $A^{1,m} = A^{2,m}$, $\forall m$, otherwise non-linear expressions of stochastic quantities would result in a bias in the results. In general, the normalization of FCIQMC simulations can only be controlled to a certain extent, and attempting to enforce a constant normalization too aggressively can

lead to additional biases.²³ Relying on the assumption that $A^{1,m} \approx A^{2,m}$ would likely be unreliable, we proceed with a different approach which avoids this.

The product of the above two expressions gives the square of the desired result,

$$\Gamma_{pq,rs}^{nm}[1] \times \Gamma_{pq,rs}^{nm}[2] = (\Gamma_{pq,rs}^{nm})^2. \quad (24)$$

However, this expression is not practical for calculating individual elements of $\Gamma_{pq,rs}^{nm}$ since the required square root operation is highly non-linear and introduces biases. This is most debilitating for very small elements of $\Gamma_{pq,rs}^{nm}$, where the stochastic samples, $\Gamma_{pq,rs}^{nm}[1]$ and $\Gamma_{pq,rs}^{nm}[2]$, have large stochastic relative errors and therefore can have opposite signs, resulting in imaginary components of the TDM.

We therefore do not attempt to estimate individual elements of $\Gamma_{pq,rs}^{nm}$ but instead consider the final desired physical estimates. Such estimates are calculated as traces of the underlying density matrices, such as $\text{Tr}(\hat{O} \hat{\Gamma}^{nm})$ for some operator $\hat{O}$. Because the trace is a linear operation,

$$\text{Tr}(\hat{O} \hat{\Gamma}^{nm}) = \sqrt{\text{Tr}(\hat{O} \hat{\Gamma}^{nm}[1]) \times \text{Tr}(\hat{O} \hat{\Gamma}^{nm}[2])}. \quad (25)$$

This estimator does not suffer from a serious non-linear bias because only a single square root operation is taken once all averaging and all summations have been performed. This final quantity will have a small relative error. This is in contrast to the above approach where a square root operation is performed for *every* element of the density matrix, many of which will have large relative errors. Furthermore, since observables correspond to squared contractions of the TDM, generally this final square root operation will not generally be required, and truly unbiased computation can be made directly for the square of the transition operator.

To perform this unbiased sampling according to Eq. (25), two arrays must be stored and averaged for each TDM being calculated, with averaging being performed according to the expectation values in Eqs. (20) and (22). Physical quantities, such as transition dipole moments, can then be calculated according to Eq. (25), and errors can be estimated using the standard propagation of error formula. The algorithmic sampling of $\Gamma_{pq,rs}^{nm}[1]$ and $\Gamma_{pq,rs}^{nm}[2]$ is identical to that of $\Gamma_{pq,rs}^n$ for non-transition RDMs. This includes the use of an array to average walker coefficients over blocks, in order to avoid calculating expensive diagonal density matrix elements at every iteration, as discussed in Sec. II C. Because the boundaries of these blocks depend on *both* FCIQMC simulations contributing to a density matrix (see Ref. 12 for details), separately averaged walker coefficients must be stored for *each* density matrix sampled. This increases memory requirements, but these requirements are usually not too severe in practice. If such memory requirements do become limiting, then it may be possible to devise a different block-averaging scheme that avoids this greater memory requirement.

IV. PARALLEL AND SPARSE IMPLEMENTATION OF DENSITY MATRIX ARRAYS

In previous implementations, 2-RDMs in FCIQMC simulations were stored in their entirety for each compute process, in non-sparse two-dimensional arrays. Even accounting for

the symmetries $\Gamma_{ij,kl} = -\Gamma_{ji,kl} = -\Gamma_{ij,lk}$ and $\Gamma_{ii,kl} = \Gamma_{ij,kk} = 0$, and using spin symmetry, this still leaves a total of $(M(M-1)/2)^2$ elements to be held in total for each combination of spin labels ($\Gamma_{\alpha\alpha,\alpha\alpha}$, $\Gamma_{\alpha\beta,\alpha\beta}$, and $\Gamma_{\alpha\beta,\beta\alpha}$ for restricted Hartree–Fock (RHF) orbital bases), where M is the number of spatial orbitals. While feasible for small orbital basis sets, this quickly becomes unmanageable. For example, in Sec. V, we consider LiH and BH molecules in aug-cc-pVQZ basis sets, both of which contain 126 spatial orbitals. Assuming 64-bit integers, this gives a total memory requirement of around 500 MB per process and per RDM desired. Noting that a separate array must be held for each excited state RDM calculated, and two for each TDM, as well as arrays for both instantaneous and averaged estimates when desired, this quickly becomes impossible.

One option is to output elements to a file on a disk drive, but this would be unacceptable due to high access times. Another option is to not attempt to store such density matrices, but rather add their contributions to estimators on-the-fly. However, there are many situations where density matrices are required for post-FCIQMC calculations.^{15,16,24} We therefore now describe a parallel and sparse implementation, which we have implemented in our FCIQMC program, NECI.^{21,25} The approach taken is similar to the sparse and parallel implementation of the FCIQMC walker array, already in use and described previously.²¹

We now choose to store all RDM elements in a sparse two-dimensional array. This array contains all density matrices simultaneously, both for the ground state and excited state RDMs and for all TDMs. The process by which sampled elements are added to this array over the course of a single iteration is described in Fig. 1. First, a double excitation is generated by the FCIQMC spawning algorithm, from a determinant $|D\rangle$ to $\hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_l \hat{a}_k |D\rangle$, which leads to a sampling of a 2-RDM element $\Gamma_{ij,kl}$, as described in Sec. II B. In practice, a contribution is calculated for every density matrix involving the state from which FCIQMC spawning occurred. Thus, a single FCIQMC spawning leads a contribution in multiple density matrix samples.

Because a sparse implementation is desired, it is necessary to store the spin orbital labels, i, j, k , and l , in addition to the value $\Gamma_{ij,kl}$. To this end, we define a row label $R(i, j)$, a column label $C(k, l)$, and a total combined label $T(i, j, k, l)$ according to

$$R(i, j) = (i - 1)M + j, \quad (26)$$

$$C(k, l) = (k - 1)M + l, \quad (27)$$

$$T(i, j, k, l) = (R(i, j) - 1)M^2 + C(k, l). \quad (28)$$

To save memory, the total label, $T(i, j, k, l)$, is stored as a single integer in the array of RDM contributions rather than the individual spin orbital labels. This encoding does not make use of all symmetries but is efficient to decode, allowing individual orbital labels to be accessed when required. For a single-particle basis with M spin orbitals, the largest 2-RDM contribution index is $T(M, M, M, M) = M^4$. This allows indexing for an M of more than 10^4 , far larger than any FCIQMC calculation to date.

FIG. 1. Algorithm used to add an RDM element to the RDM spawning array, after a double excitation occurs in the FCIQMC spawning step.

Importantly, we note that certain RDM elements are generated very many times in an iteration. It is clearly inefficient and memory intensive to store these repeated elements separately throughout the spawned RDM array. We therefore ensure that repeated elements are always added into the same position. To achieve this, an efficient method is required to check if an element already exists in the array, and if so, where. We therefore also store a hash table to the RDM array. When a 2-RDM element $\Gamma_{ij,kl}$ is generated, a hash function is applied to i, j, k , and l to calculate an associated hash value (the hash function applied is the same one used for the parallel distribution of determinants in FCIQMC, as presented in Ref. 21). This hash value specifies which row of the hash table in which to search for this RDM element. If the element already exists in the RDM array, then a reference to that element will exist in that row of the hash table, and so the position of the existing element may be quickly looked up. This avoids having to perform expensive and regular sort operations on the RDM array to prevent the proliferation of repeated elements. In practice, each “row” of the hash table is a linked list implemented with pointers, whose size can increase dynamically as more hash clashes occur.

A mapping is also required to distribute density matrix elements across process threads within a distributed memory. In FCIQMC implementations, determinants have been distributed using a hash function of their occupied orbital list, to attempt to evenly distribute walkers across processors, balancing the load. Here, we simply split all possible RDM elements across processes evenly using the row label, $R(i, j)$. We find this adequate for calculations performed thus far, and it makes the printing of RDM arrays in a predefined order (as required

in by many quantum chemistry packages) efficient and simple.

Finally, we consider efficient parallel distribution of the memory requirements. For a simulation performed on N_{proc} processors, the total RDM spawning array is split evenly into N_{proc} sections. When an RDM element is generated, its processor label is calculated, and the element is simply added into the earliest available position for the assigned process. In this way, it is then simple to distribute RDM elements to their correct processors at the end of an iteration via an `MPI_Alltoallv`, when using a Message Passing Interface (MPI) implementation. Before communication, RDM elements belonging to any MPI process can be held by any other process. However, for large RDMs, only a small fraction of elements are sampled in any one iteration, and so the memory requirements remain small. Then, once elements have been sent to their assigned processes, they can be averaged and accumulated over the entire simulation, with the greater memory requirements of this averaging efficiently distributed across processing cores.

Beyond allowing RDM arrays to be distributed across processing cores, there are further benefits to this approach. In particular, the full sparsity of the exact 2-RDMs is automatically taken advantage of entirely for a completely general system. This is because RDM elements are only added to the array once sampled, and only non-zero elements are ever sampled. Indeed, this will also be true for higher-order density matrices, including the full N -body density matrix, if desired. In contrast, taking full advantage of symmetries in the previously used scheme required separate computer code for separate systems, which often became difficult to maintain.

There are undoubtedly some overheads in this scheme. In particular, the fact that a hash lookup is performed for each RDM element sampled can be expensive. This is especially true since our implementation of the hash table is built from linked lists, which are slow to search compared to contiguous memory lookups. However, as described in Sec. II C, a block averaging scheme is used where diagonal elements, and elements involving the reference, are calculated very infrequently. These account for the majority of RDM elements sampled in practice, and so long as these elements are included infrequently, the overhead is small. There are also some memory overheads, in particular the requirement to store a hash table for each RDM array. However, this is easily outweighed by the parallel distribution and automatic exploitation of sparsity in all matrices. This is demonstrated in Sec. V, where 2-RDMs are sampled in single-particle basis sets which would have not been possible in the previous scheme.

V. RESULTS

As an initial test of these ideas, we consider the calculation of dipole moments, transition dipole moments, and oscillator strengths for low-lying states of small diatomic molecules. These quantities are of great importance for understanding various properties of molecular systems. The oscillator strength in particular is required to explain optical spectra,

as it determines the probabilities of absorption and emission of photons coupling different electronic states. Nonetheless, dipole moments are challenging to calculate accurately, even for small molecules, because they are very sensitive to the quality of the wave function and single-particle basis set used, generally requiring many diffuse orbitals for an accurate description, with far greater basis set sensitivity than the energy.²⁶

We therefore begin by considering the LiH and BH molecules in aug-cc-pVDZ, aug-cc-pVTZ, and aug-cc-pVQZ, containing 32, 69, and 126 spatial orbitals, respectively. The aug-cc-pVQZ basis 2-RDM was unobtainable in the previous RDM implementation, despite the small molecular size. We then consider the MgO molecule in an aug-cc-pVDZ basis set. We note that while the calculation of dipole moments only requires the 1-RDM, for these calculations we obtain the 1-RDM by contracting the 2-RDM, which we also use to calculate the energy using the estimator

$$(E_{\text{RDM}})_n = \frac{\text{Tr}[\hat{H} \hat{\Gamma}^n]}{\text{Tr}[\hat{\Gamma}^n]}. \quad (29)$$

Therefore, the following is a good test of the newly introduced ideas, as well as providing further insight into the effect of the initiator adaptation for different estimators and excited states.

The dipole moment for the state $|\Phi^n\rangle$ is defined by

$$\mu_n = \sum_{pq} \gamma_{p,q}^n \langle p | \hat{r} | q \rangle, \quad (30)$$

while a transition dipole moment, t_{nm} , is defined by Eq. (17), and the corresponding oscillator strength by

$$f_{nm} = \frac{2}{3} \Delta E_{nm} |t_{nm}|^2, \quad (31)$$

for an energy gap of ΔE_{nm} between states $|\Phi^n\rangle$ and $|\Phi^m\rangle$.

For all simulations, the initial restricted Hartree–Fock (RHF) calculation was performed by PySCF.²⁷ Integrals from PySCF were then passed to our FCIQMC program, NECI, for the main calculation, which output one- and two-body density matrices. These were then contracted with integrals from PySCF to calculate final dipole moment estimates. Energy estimates were calculated on-the-fly in NECI.

The five lowest energy states were calculated for LiH and BH, and the four lowest states of MgO, considering only states with $M_s = 0$ and using the A_1 irreducible representation (irrespective) of the C_{2v} point group. Also, time-reversal symmetrized functions²⁸ were used as the many-particle basis states, therefore restricting the total spin quantum number, S , to be even, and thus removing triplet states. In all cases, the FCIQMC simulation time step was varied in the initial iterations so as to prevent “bloom” events, where many walkers can be created in a single spawning event (which often leads to large initiator error).

We also note that in generating excitations for the walker spawning step, we use an approach that greatly improves efficiency compared to the uniform sampling used in early FCIQMC results.⁵ In this approach, the pair of orbital labels from which electrons are excited, (i, j) , are chosen uniformly, while the orbitals excited to, (a, b) , are selected with probabilities drawn from a Cauchy-Schwarz distribution, namely,

$p(ab|ij) \propto \sqrt{\langle ia|ia \rangle \langle jb|jb \rangle}$.²⁹ Another approach to select connections efficiently was considered by Holmes *et al.*³⁰ but not used here.

All simulations used the semi-stochastic adaptation to reduce stochastic errors.^{31,32} For the LiH molecule, the deterministic space consisted of all configurations up to and including double excitations from the Hartree–Fock determinant. For the BH and MgO molecules, the deterministic space was formed from the 10^4 most populated configurations across all wave functions sampled, once the simulations were deemed to have largely converged, using the approach described in Ref. 32.

A. LiH

Simulations on LiH were performed using between 1.25×10^4 and 10^6 walkers per simulation (i.e., for each state and replica sampled), in order to converge initiator error for all states. Density matrices were typically averaged over 10^5 iterations, once convergence was deemed to have been reached for all states and all estimators. These entire simulations were then repeated five times with different initial RNG seeds, and the results were averaged in order to calculate error estimates.

Figure 2 shows initiator convergence for LiH in the aug-cc-pVQZ basis set, for the lowest five energy eigenstates, and for four different estimators: dipole moments, transition dipoles moments, and energies calculated from both the RDM-based energy estimator, Eq. (29), and from a trial wave function-projected estimator,

$$(E_{\text{Trial}})_n = \frac{\langle \Psi_{\text{Trial}}^n | \hat{H} | \Psi_{\text{Trial}}^n \rangle}{\langle \Psi_{\text{Trial}}^n | \Psi_{\text{Trial}}^n \rangle}. \quad (32)$$

Here, $|\Psi_{\text{Trial}}^n\rangle$ is a trial wave function designed to have a large overlap with the exact state $|\Phi^n\rangle$. We have discussed the use of such trial wave function estimators in excited-state FCIQMC in Ref. 22. To generate $|\Psi_{\text{Trial}}^n\rangle$, we calculate the configuration interaction singles and doubles (CISD) wave functions for the lowest fifteen energy states. Then, once convergence of all FCIQMC simulations is deemed to have been reached, we assign each simulation one trial wave function by choosing the CISD solution with the largest overlap in each case. The reason for obtaining more CISD solutions than FCIQMC simulations is that CISD solutions can have a different energy ordering to FCI solutions. Averaging of each E_{Trial} estimate was performed from roughly the same point that RDM sampling began, and so both RDM and trial energy estimates are obtained from a similar number of iterations, usually 10^5 .

The initiator-FCIQMC estimates in Fig. 2 are all plotted relative to their values at the largest walker population considered, $N_w = 10^6$. Here, convergence has been largely reached in all cases, and so the figures effectively plot initiator error against walker population. Reassuringly, initiator error in energy estimates is incredibly small for both estimators and for all states. Indeed, the largest error at the smallest walker population tested is less than ~ 0.5 mE_h for E_{Trial} .

Interestingly, initiator error in E_{RDM} is much smaller than in E_{Trial} . This is a trend that we have often observed, although

FIG. 2. Initiator error convergence for the five lowest energy states of LiH in an aug-cc-pVQZ basis, at an internuclear distance of 1.5957 Å as the number of walkers in each distribution is increased. Results are shifted relative to their values at the largest walker population considered, therefore approximately representing the initiator error. (a) Dipole moments. (b) Transition dipole moments from the ground state. (c) Energy calculated from a trial estimator, E_{Trial} . (d) Energy calculated from the RDM estimator, E_{RDM} . N_w denotes the number of walkers for each state and replica sampled. Simulations were typically averaged over 5 simulations to obtain error bars.

exceptions do occur (and in the limit of an exact $|\Psi_{\text{Trial}}^m\rangle$, the initiator error is zero). Initiator error in the E_{RDM} energies are variational in all cases within stochastic errors, while it is not strictly enforced (though common) for this to also be the case for E_{Trial} . For RDM-based energy estimates, this variability is effectively ensured by the Hylleraas-Undheim-McDonald theorem,^{33,34} which is expected to approximately hold for FCIQMC-sampled wave functions. Initiator error is larger for excited states, as previously observed.²² This is expected due to the more multi-configurational nature of excited states. Random errors however are larger in the RDM-based energy estimates, which is expected due to the fact that two uncorrelated simulations (from the two replicas) contribute to this quantity. However, error bars are extremely small in all cases here, always being smaller than $10^{-2} mE_h$.

The calculation of dipole moments provides a more interesting test, due to their greater dependence on more highly excited determinants and diffuse single-particle orbitals. The relative initiator error is much larger, particularly for certain excited states (i.e., μ_2 and μ_3). The transition dipole moments considered involve transitions from the ground ($n = 0$) state to excited ($n > 1$) states. Because they always involve the ground state, it is to be expected that they have smaller relative initiator and stochastic errors, compared to the corresponding non-transition dipole moment (i.e., t_{0n} compared to μ_n). This expectation is borne out in the results, with initiator and stochastic errors in t_{0n} often being ~ 5 times smaller than for μ_n . For the calculation of dipole moments from FCIQMC-sampled RDMs, relative stochastic errors are clearly much larger than for energies, and so the use of the semi-stochastic adaptation is of great importance here, whereas its use can be somewhat unnecessary in small ground-state energy calculations.

Clearly, the accurate calculation of dipole moments is more challenging than energies, requiring larger walker populations to obtain similar relative errors. However, this is not uniquely a feature of the initiator approximation in FCIQMC

but is equally true in other approximate methods, where properties such as the dipole moment are far more sensitive to the basis set and quality of the wavefunction than ground state energetics. That we are able to observe that systematic convergence of these quantities, with respect to a single simulation parameter, is reassuring.

Table I gives final results for the aug-cc-pVXZ basis sets, with $X = 2, 3, 4$. Results in the small $X = 2$ basis were fully converged at the smallest walker populations considered, $N_w = 1.25 \times 10^4$, as confirmed by comparison to FCI results from the PySCF program. As expected, dipole moments vary quite substantially with basis set, particularly for the second, third, and fourth excited states, demonstrating the importance of large basis sets with diffuse functions. Errors in brackets denote stochastic error bars, not initiator error, which is larger. However, given the careful convergence of initiator error, as shown in Fig. 2, we expect dipole moments to be converged to around $10^{-3} ea_0$ in most cases and energies to be converged *substantially* beyond chemical accuracy.

B. BH

Figure 3 shows results for BH in the aug-cc-pVTZ basis set and at an internuclear distance of 1.2324 Å, demonstrating similar initiator convergence plots to those in Fig. 2. Here, results are used between 1.25×10^4 and 2×10^6 walkers per simulation. RDM estimators and E_{Trial} were averaged over 5×10^4 iterations, once convergence was achieved for all states and estimators. Here, instead of using CISD solutions as trial wave functions for E_{Trial} , a slightly different approach was used: a “trial space” was defined as consisting of the 2×10^3 most populated configurations across all simulations, once convergence had been approximately reached. Trial wave functions were then obtained as the eigenstates of $\hat{H}$ within this subspace. This is similar to the approach to generate the deterministic space, as described above,³² and allows important basis states to be picked, while allowing an inexpensive calculation to determine each $|\Psi_{\text{Trial}}^m\rangle$.

TABLE I. Final converged estimates for the LiH molecule at an internuclear distance of 1.5957 Å. Results are for the five lowest energy states in A_1 irrespective of the C_{2v} point group, with $M_S = 0$ and $S = \text{even}$ quantum numbers (which happen to all be $^1\Sigma^+$ states). $n = 0$ refers to the ground state, and $n > 1$ to excited states. Numbers in parentheses denote stochastic error, not initiator error. Energy gaps (ΔE_{1n}) were calculated using RDM-based energy estimates, Eq. (29). Integrals were generated using the PySCF program.²⁷ In the small aug-cc-pVDZ, all results were verified against exact FCI results obtained from PySCF (not shown here).

Basis	State n	Energy gap (ΔE_{0n})	Dipole moment (μ_n)	Transition dipole moment (t_{0n})	Oscillator strength (f_{0n})
aug-cc-pVDZ	0 ($^1\Sigma^+$)	...	-2.325 137 2(2)	...	...
	1 ($^1\Sigma^+$)	0.130 434(1)	2.019 47(4)	0.965 189(7)	0.081 007(1)
	2 ($^1\Sigma^+$)	0.214 979 9(6)	-3.354 3(9)	0.374 71(1)	0.020 123(1)
	3 ($^1\Sigma^+$)	0.229 077(4)	5.083 2(8)	0.091 26(8)	0.001 271(2)
	4 ($^1\Sigma^+$)	0.246 350(3)	-0.295 8(3)	0.560 74(2)	0.051 639(4)
aug-cc-pVTZ	0 ($^1\Sigma^+$)	...	-2.306 440(9)	...	...
	1 ($^1\Sigma^+$)	0.132 458(3)	2.025 41(7)	0.935 38(2)	0.077 262(4)
	2 ($^1\Sigma^+$)	0.216 705(6)	-3.794(1)	0.411 46(2)	0.024 459(2)
	3 ($^1\Sigma^+$)	0.230 621(2)	5.533(1)	0.070 42(8)	0.000 762(2)
	4 ($^1\Sigma^+$)	0.246 520(2)	-0.623 5(7)	0.693 170(7)	0.078 966(2)
aug-cc-pVQZ	0 ($^1\Sigma^+$)	...	-2.301 68(3)	...	...
	1 ($^1\Sigma^+$)	0.132 943(7)	2.018 8(1)	0.926 58(4)	0.076 093(7)
	2 ($^1\Sigma^+$)	0.217 616(7)	-3.696(2)	0.398 4(1)	0.023 03(1)
	3 ($^1\Sigma^+$)	0.231 229(2)	6.211(2)	0.108 3(2)	0.001 809(6)
	4 ($^1\Sigma^+$)	0.242 846(9)	-1.998(2)	0.620 1(5)	0.062 24(9)

Results contain the same features as observed for LiH. Initiator error in the energy estimates is extremely small in all cases, particularly for estimates obtained from contraction of the RDM, and initiator convergence always occurs variationally. Stochastic error bars are larger for E_{RDM} , as well as for excited states, but are always extremely small. For dipole moments, similar trends also occur. Initiator and stochastic relative errors for the dipole moment are very small for the ground and first excited states (μ_0 and μ_1) and for the corresponding transition dipole moment (t_{01}) even at small walker populations. However, results for higher excited states contain larger errors, although we once again observe that errors in t_{0n} are smaller than errors in μ_n for each n , presumably because of

the involvement of the ground state, which is well converged at lower walker populations, in each of the transition dipole moments considered.

Table II shows final results in aug-cc-pVXZ basis sets, for $X = 2, 3, 4$. Results for $X = 2$ used 2×10^5 walkers per simulation, while results for $X = 3$ and $X = 4$ results used 2×10^6 walkers per simulation. The expected strong dependence of dipole moments on the basis set is once again observed. This is particularly true for the second, third, and fourth excited states ($n = 2, 3, 4$). We note that these three states also contained the largest initiator error at small walker populations, as seen in Fig. 3. This is probably not a coincidence since the initiator approximation will inevitably result in a poorer description

FIG. 3. Initiator convergence for the five lowest energy states of BH in an aug-cc-pVTZ basis, at an internuclear distance of 1.2324 Å. Results are shifted relative to their values at the largest walker population considered, therefore approximately representing the initiator error. (a) Dipole moments. (b) Transition dipole moments from the ground state. (c) Energy calculated from a trial estimator, E_{Trial} . (d) Energy calculated from the RDM estimator, E_{RDM} . N_w denotes the number of walkers for each state and replica sampled. Simulations were typically averaged over 5 simulations to obtain error bars.

TABLE II. Final converged estimates for the BH molecule at an internuclear distance of 1.2324 Å. Results are for the five lowest energy states in A_1 irrespective of the C_{2v} point group, with $M_S = 0$ and $S = \text{even}$ quantum numbers. $n = 0$ refers to the ground state and $n > 1$ refers to excited states. Numbers in parentheses denote stochastic error not initiator error. Energy gaps (ΔE_{0n}) were calculated using RDM-based energy estimates, Eq. (29). Integrals were generated using the PySCF program.²⁷

Basis	State n	Energy gap (ΔE_{0n})	Dipole moment (μ_n)	Transition dipole moment (t_{0n})	Oscillator strength (f_{0n})
aug-cc-pVDZ	0 ($^1\Sigma^+$)	...	0.528 082(7)	...	...
	1 ($^1\Delta$)	0.216 230(3)	-0.189 83(3)	0.0	0.0
	2 ($^1\Sigma^+$)	0.237 27(1)	-1.414 6(5)	0.934 78(3)	0.138 22(1)
	3 ($^1\Sigma^+$)	0.257 587(4)	-0.321 9(3)	0.210 2(1)	0.007 590(9)
	4 ($^1\Sigma^+$)	0.282 665(1)	3.545 9(1)	0.447 25(4)	0.037 696(7)
aug-cc-pVTZ	0 ($^1\Sigma^+$)	...	0.545 61(2)	...	...
	1 ($^1\Delta$)	0.211 482(6)	-0.192 71(7)	0.0	0.0
	2 ($^1\Sigma^+$)	0.238 668(8)	-1.294 3(5)	0.885 08(5)	0.124 64(1)
	3 ($^1\Sigma^+$)	0.253 574(4)	-0.497 3(6)	0.145 4(2)	0.003 58(1)
	4 ($^1\Sigma^+$)	0.283 481(2)	3.408 8(2)	0.357 40(7)	0.024 141(9)
aug-cc-pVQZ	0 ($^1\Sigma^+$)	...	0.549 14(6)	...	...
	1 ($^1\Delta$)	0.210 59(2)	-0.196 8(3)	0.0	0.0
	2 ($^1\Sigma^+$)	0.238 76(3)	-1.268(3)	0.870 4(3)	0.120 6(1)
	3 ($^1\Sigma^+$)	0.252 61(3)	-0.504(3)	0.139(1)	0.003 27(7)
	4 ($^1\Sigma^+$)	0.282 89(1)	3.288 9(9)	0.313 8(1)	0.018 57(2)

of highly excited regions of the wave function, presumably including excitations into high-energy diffuse functions, which appear important for accurate calculation of dipole moments for these particular states. Despite larger initiator error than for energy estimates, there is still a substantial undersampling of the space here, using 2×10^6 walkers for a space size of $\sim 7 \times 10^9$ for the aug-cc-pVQZ basis, even for this small molecule, with benefits of Monte Carlo sampling typically increasing with system size.

C. MgO

To study a more challenging problem, we consider the calculation of energies and dipole moments for the MgO molecule, at its ground state equilibrium separation of 1.749 Å, and with 4 core electrons frozen at the Hartree–Fock level. Thus, a total of 16 electrons are correlated in 48 spatial orbitals. Enforcing $M_s = 0$, using A_1 irrespective of the C_{2v} point group, and working with time-reversal symmetrized functions²⁸ (to

FIG. 4. Initiator convergence for dipole moments (left) and energies (right), for MgO in an aug-cc-pVDZ basis set, at an internuclear distance of 1.749 Å, and with 4 core electrons frozen. The four lowest-energy states are considered in A_1 irrespective of C_{2v} and with $S = \text{even}$ enforced (all $^1\Sigma^+$ states). Energies are calculated from both RDM (E_{RDM}) and trial wave function (E_{Trial}) based estimates and become equal to good accuracy at large walker number, N_w . Dipole moments appear mostly converged at $N_w = 3.2 \times 10^7$ except for μ_1 . Error bars are only available for $N_w < 10^6$ but are small by this point and should only decrease in magnitude for larger walker populations.

enforce $S = \text{even}$) result in a space size of roughly 1.8×10^{16} basis functions. This is a large space, particularly given the challenges of converging initiator error in excited-state dipole moments, as seen already.

Figure 4 presents initiator convergence for walker populations (per state and per replica), N_w , ranging from 2.5×10^4 to 3.2×10^7 . The ground state and first three excited states are calculated. For $N_w \leq 4 \times 10^5$, error bars are calculated by averaging over 5 repeated calculations with varying RNG seeds. Due to the expensive nature of calculations, repeats were not performed for $N_w > 4 \times 10^5$, and so error bars were not obtained. However, these error bars should mostly only decrease with increasing N_w and are already small at $N_w = 4 \times 10^5$. Therefore, at the largest walker populations considered, stochastic error should be much smaller than initiator error.

Initiator profiles of both E_{RDM} and E_{Trial} estimators are presented in Fig. 4. At convergence, these should clearly become equal. By $N_w = 3.2 \times 10^7$, this is the case to much better than 1 mE_h accuracy. As previously found, convergence is monotonic in all cases and E_{RDM} usually results in smaller initiator error.

Convergence of dipole moments is also shown. Here, relative initiator error is once again larger than for energies, and convergence is non-monotonic. Because of this non-monotonic behavior, combined with the challenging nature of the system, our confidence in the accurate convergence of these values is somewhat less than for LiH and BH results. We cannot rule out the possibility of sudden further convergence at higher N_w values. However we believe any significant deviations unlikely, although it is clear that μ_1 in particular is not fully converged on the scale shown.

Table III presents FCIQMC energies and dipole moments, using $N_w = 3.2 \times 10^7$, and with energies taken from the E_{RDM} estimator. For comparison, coupled cluster results are shown, using both singles and doubles (CCSD) and singles, doubles, and triples (CCSDT). These values were calculated using the NWChem package,³⁵ with the equation-of-motion (EOM-CCSD and EOM-CCSDT) variants used for excited states. As expected, energies obtained from CCSDT are accurate compared to FCIQMC values, even for excited states. Meanwhile, dipole moments show greater differences particularly for the $n = 3$ state. For this state, EOM-CCSD and EOM-CCSDT values also greatly differ, with a flipped dipole moment

resulting from EOM-CCSD. These results are consistent with those observed in FCIQMC in regions of large initiator error—the relative error in dipole moments is much greater than in energies. We again expect that this is primarily due to the increased dependence on highly excited determinants, and such configurations have particularly large amplitudes in excited states. CCSD and CCSDT appear to be unable to describe the wave function with sufficient accuracy in this region of configuration space, for this system, and for these challenging states.

VI. CONCLUSION

We have presented the first calculation of excited-state density matrices and transition density matrices using the excited-state FCIQMC method, and subsequently performed accurate RDM-based calculations of energies and dipole moments in small heteronuclear molecules. Initiator convergence has been investigated in some detail for various estimators and excited states, giving further insight into the initiator approach in more general situations than previously considered.

To allow the storage of multiple density matrices in large single-particle basis sets, we have described and implemented an efficient algorithm to store sampled density matrix elements in a sparse format, one which automatically makes full use of symmetries and is efficiently distributed in a massively parallel simulation. While currently only used for sampling one- and two-body density matrices, this approach will trivially allow efficient storage of higher body density matrices too, which could be used to calculate important entanglement entropy measures, among many other potential applications.

Although the applications here were to small systems, the calculation of RDMs in FCIQMC will perhaps be key in allowing this QMC approach to be extended to significantly larger problems: this QMC approach has recently been used to develop stochastic versions of the complete active space self-consistent field (CASSCF) method,^{15,36} in which 1- and 2-RDMs are required to perform the orbital optimization step. RDMs are also required in the explicitly correlated $[2]_{R12}$ approach of Torheyden and co-workers,^{37,38} which can be used with FCIQMC to greatly reduce basis set errors,¹⁶ also allowing larger systems to be studied for a given computational cost and accuracy. Given the clear importance of density matrices in FCIQMC currently, we hope that these developments and insights will be of great value going forward.

ACKNOWLEDGMENTS

N.S.B. acknowledges support during this work from Trinity College, Cambridge. G.H.B. gratefully acknowledges the Royal Society for funding via a university research fellowship, and the support of the Air Force Office of Scientific Research (AFOSR) through Grant Award No. FA9550-16-1-0256. A.A. acknowledges support from the EPSRC under Grant No. EP/J003867/1. We also thank Qiming Sun for helpful conversations regarding the PySCF program.

TABLE III. Energies and dipole moments for MgO in an aug-cc-pVDZ basis set, at an internuclear distance of $1.749 \text{ \AA}$, and with 4 core electrons frozen at the Hartree–Fock level. The four lowest-energy states are considered in A_1 irrespective of C_{2v} and with $S = \text{even}$ enforced (all $^1\Sigma^+$ states). Error bars on FCIQMC results are not given but are smaller than the order to which results are presented. FCIQMC energies are taken from the RDM-based estimates, E_{RDM} . CCSD and CCSDT values were obtained from NWChem.³⁵

State n	Energy/ E_h			Dipole moment (μ_n)/ ea_0		
	CCSD	CCSDT	FCIQMC	CCSD	CCSDT	FCIQMC
0 ($^1\Sigma^+$)	−274.632	−274.651	−274.654	2.590	2.398	2.382
1 ($^1\Sigma^+$)	−274.531	−274.559	−274.564	1.811	2.008	2.289
2 ($^1\Sigma^+$)	−274.480	−274.514	−274.517	0.297	0.847	1.154
3 ($^1\Sigma^+$)	−274.440	−274.478	−274.480	−0.366	0.529	1.198

- ¹W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- ²R. Gaudoin and J. M. Pitarke, *Phys. Rev. Lett.* **99**, 126406 (2007).
- ³J. Casulleras and J. Boronat, *Phys. Rev. B* **52**, 3654 (1995).
- ⁴M. C. Per, I. K. Snook, and S. P. Russo, *Phys. Rev. B* **86**, 201107 (2012).
- ⁵G. H. Booth, A. J. W. Thom, and A. Alavi, *J. Chem. Phys.* **131**, 054106 (2009).
- ⁶D. M. Cleland, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **132**, 041103 (2010).
- ⁷G. H. Booth, A. Gruneis, G. Kresse, and A. Alavi, *Nature* **493**, 365 (2012).
- ⁸J. S. Spencer, N. S. Blunt, and W. M. C. Foulkes, *J. Chem. Phys.* **136**, 054110 (2012).
- ⁹M. H. Kolodrubetz, J. S. Spencer, B. K. Clark, and W. M. Foulkes, *J. Chem. Phys.* **138**, 024110 (2013).
- ¹⁰S. Zhang and M. H. Kalos, *J. Stat. Phys.* **70**, 515 (1993).
- ¹¹N. S. Blunt, T. W. Rogers, J. S. Spencer, and W. M. C. Foulkes, *Phys. Rev. B* **89**, 245124 (2014).
- ¹²C. Overy, G. H. Booth, N. S. Blunt, J. J. Shepherd, D. Cleland, and A. Alavi, *J. Chem. Phys.* **141**, 244117 (2014).
- ¹³G. H. Booth, D. Cleland, A. Alavi, and D. P. Tew, *J. Chem. Phys.* **137**, 164112 (2012).
- ¹⁴R. E. Thomas, D. Opalka, C. Overy, P. J. Knowles, A. Alavi, and G. H. Booth, *J. Chem. Phys.* **143**, 054108 (2015).
- ¹⁵R. E. Thomas, Q. Sun, A. Alavi, and G. H. Booth, *J. Chem. Theory Comput.* **11**, 5316 (2015).
- ¹⁶J. A. F. Kersten, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **145**, 054117 (2016).
- ¹⁷*Reduced-Density-Matrix Mechanics: With Application to Many-Electron Atoms and Molecules*, Advances in Chemical Physics, edited by D. A. Mazziotti (Wiley, New York, 2007), Vol. 134.
- ¹⁸A. W. Schlimgen, C. W. Heaps, and D. A. Mazziotti, *J. Phys. Chem. Lett.* **7**, 627 (2016).
- ¹⁹D. A. Mazziotti, *Phys. Rev. Lett.* **117**, 153001 (2016).
- ²⁰D. M. Cleland, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **134**, 024112 (2011).
- ²¹G. H. Booth, S. D. Smart, and A. Alavi, *Mol. Phys.* **112**, 1855 (2014).
- ²²N. S. Blunt, S. D. Smart, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **143**, 134117 (2015).
- ²³W. A. Vigor, J. S. Spencer, M. J. Bearpark, and A. J. W. Thom, *J. Chem. Phys.* **142**, 104101 (2015).
- ²⁴S. Sharma, T. Yanai, G. H. Booth, C. J. Umrigar, and G. K.-L. Chan, *J. Chem. Phys.* **140**, 104112 (2014).
- ²⁵See <https://github.com/ghb24/NECL.STABLE> for “Neci github web page.”
- ²⁶S. Green, *Adv. Chem. Phys.* **25**, 179 (1974).
- ²⁷Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. McClain, S. Sharma, S. Wouters, and G. K.-L. Chan, e-print [arXiv:1701.08223](https://arxiv.org/abs/1701.08223) [physics.chem-ph] (2017).
- ²⁸Y. G. Smeyers and L. Doreste-Suarez, *Int. J. Quantum Chem.* **7**, 687 (1973).
- ²⁹S. D. Smart, G. H. Booth, and A. Alavi, “Excitation generation in full configuration interaction quantum Monte Carlo based on Cauchy-Schwarz distributions” (unpublished).
- ³⁰A. A. Holmes, H. J. Changlani, and C. J. Umrigar, *J. Chem. Theory Comput.* **12**, 1561 (2016).
- ³¹F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale, and C. J. Umrigar, *Phys. Rev. Lett.* **109**, 230201 (2012).
- ³²N. S. Blunt, S. D. Smart, J. A. F. Kersten, J. S. Spencer, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **142**, 184107 (2015).
- ³³E. Hylleraas and B. Undheim, *Z. Phys.* **65**, 759 (1930).
- ³⁴J. K. L. McDonald, *Phys. Rev.* **43**, 830 (1933).
- ³⁵M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, H. J. J. van Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus, and W. A. de Jong, *Comput. Phys. Commun.* **181**, 1477 (2010).
- ³⁶G. Li Manni, S. D. Smart, and A. Alavi, *J. Chem. Theory Comput.* **12**, 1245 (2016).
- ³⁷M. Torheyden and E. F. Valeev, *J. Chem. Phys.* **131**, 171103 (2009).
- ³⁸L. Kong and E. F. Valeev, *J. Chem. Phys.* **135**, 214105 (2011).