

Cholesky decomposition within local multireference singles and doubles configuration interaction

Tsz S. Chwee¹ and Emily A. Carter^{2,a)}¹*Department of Chemistry, Princeton University, Princeton, New Jersey 08544, USA*²*Department of Mechanical and Aerospace Engineering and Program in Applied and Computational Mathematics, Princeton University, Princeton, New Jersey 08544-5263, USA*

(Received 3 November 2009; accepted 21 January 2010; published online 16 February 2010)

A local multireference singles and doubles configuration interaction method in which Cholesky vectors are used in place of conventional two-electron integrals has been developed (CD-LMRSDCI). To reduce the overall cost associated with our linear scaling LMRSDCI method presented earlier [T. S. Chwee *et al.*, *J. Chem. Phys.* **128**, 224106 (2008)], we adopt a two-pronged approach. First, localized orthogonal virtual orbitals, introduced by Subotnik *et al.* [*J. Chem. Phys.* **123**, 114108 (2005)], are substituted for nonorthogonal projected atomic orbitals. This obviates the need for contraction with overlap matrices and simplifies our working formalism. In addition, we restructure the rate-limiting step of our LMRSDCI algorithm to be driven by the search for two-electron integrals instead of configuration state functions. The shift necessitates a flexible way of processing the four-indexed two-electron integrals, which is facilitated by use of two-indexed Cholesky vectors. Our restructured LMRSDCI method is an order of magnitude faster and has greatly reduced storage requirements so that we are able to apply it to molecules containing up to 50 heavy atoms. However, generation of the Cholesky vectors and their subsequent transformation to the molecular orbital (MO) basis is not linear scaling. Together with assembling the MO integrals from the Cholesky vectors, these now constitute the rate-limiting steps in our method.

© 2010 American Institute of Physics. [doi:10.1063/1.3315419]

I. INTRODUCTION

Despite the prevalent use of *ab initio* electronic structure methods, such as Møller–Plesset perturbation and coupled cluster (CC) theories, configuration interaction (CI) remains one of the most flexible and accurate methods for describing electron correlation in atoms and molecules. Full CI (FCI) calculations provide the benchmarks for these other techniques. However, the associated $O(N!)$ scaling renders FCI inaccessible to all but very small molecules. This prohibitive algorithmic scaling is typically reduced by truncating the CI expansion at second order, as justified by the presence of at most two-body operators in the Hamiltonian. The most accurate version of truncated second order CI is multireference singles and doubles CI (MRSDCI), but in its conventional form it still scales very poorly [$O(N^6)$] and furthermore requires modifications to retain size extensivity.^{1–10} Although multireference CC methods are under active development,^{11–23} MRSDCI is generally the *ab initio* technique of choice for molecules and phenomena exhibiting near degeneracies or open-shell character, such as transition metal complexes and breaking/forming chemical bonds.

In order to apply MRSDCI to ever larger molecules, the overall scaling and cost must be reduced. This is possible by mapping conventional MRSDCI onto a local correlation framework whereby localized orbitals are used to span the occupied and virtual subspaces of a molecule. By exploiting locality, cutoff thresholds can be applied to reduce the over-

all computational overhead. The incorporation of local correlation treatments into MRSDCI (LMRSDCI) was described in earlier work^{24–29} and is briefly summarized here.

The effective dimension of the CI expansion in the basis of N -electron configuration state functions (CSFs)/determinants can be reduced by exploiting the short-ranged nature of electron correlation effects, as pioneered by Saebo and Pulay (SP).^{30–35} SP's weak pairs (WP) approximation³⁰ exploits the observation that the motion of two electrons occupying spatially distant orbitals is only weakly correlated. As such, configurations involving simultaneous excitations out of spatially well-separated occupied orbitals are not retained in the CI expansion within the WP approximation. Here excitations to the full virtual space are used to correlate those electron pairs deemed nonweak. Substantially greater savings is achieved by using SP's truncation of virtuals (TOV) approximation,³⁰ in which they recognized that the best correlating orbitals for a specified occupied orbital are spatially close to it. As such, one can define prior to a CI calculation a set of correlating orbitals for each occupied orbital. Only excitations to this truncated virtual space are allowed. Taken together, the WP and TOV approximations ensure that the total number of correlated electron pairs will increase in a linear fashion asymptotically, while their associated virtual domain will have a constant value for some preset threshold.

Consequently, the total number of CSFs included in the CI expansion will scale linearly with system size while CI Hamiltonian matrix elements [consisting largely of two-

^{a)}Electronic mail: eac@princeton.edu.

electron repulsion integrals (ERIs)]^{36,37} involving CSFs that do not satisfy the WP/TOV approximations will be excluded. When local truncations are used in conjunction with integral prescreening techniques,^{38–41} the number of required ERIs can be made to scale in a linear fashion. For example, Schütz and Werner³⁹ used the prescreening test quantity Thr_{AO} , defined in Eq. (1) for leaving out numerically small integrals ($\mu\nu|\rho\sigma$) prior to their evaluation and subsequent transformation to the MO basis,

$$Thr_{AO} = \max(D_{\nu\sigma}^{\max}, D_{\mu\rho}^{\max}) |(\mu\nu|\mu\nu)|^{1/2} |(\rho\sigma|\rho\sigma)|^{1/2},$$

$$D_{\nu\sigma}^{\max} = \max_{ij \in P} (L_{\sigma i} L_{\nu j}). \quad (1)$$

$D_{\nu\sigma}^{\max}$ is computed by taking the maximum product of localized molecular orbital coefficients of the respective basis functions ν and σ . The domain P is further limited by the WP approximation to include only orbitals i and j , which are in close proximity and as such $D_{\nu\sigma}^{\max}$ will be sparse when a localized basis is used.

In our previous work on linear scaling LMRSDCI,²⁹ occupied orbitals of the internal space were localized using the Pipek–Mezey functional⁴² while the virtual (external) space was spanned by localized, nonorthogonal projected atomic orbitals (PAOs).³⁰ By applying local truncation approximations in tandem with integral screening to weed out numerically small integrals, we were able to achieve $O(N)$ scaling for the rate limiting step in our calculations. However, the large associated prefactor restricted the application of our method to systems with no more than 25 heavy atoms. The aim of the present work is to reduce the prefactor and thus extend the reach of our method towards larger molecules.

The rest of the paper proceeds as follows. First, the necessary background on CI within a local correlation framework is given to point out existing computational bottlenecks. This is followed by a detailed description of the restructuring of our previous LMRSDCI implementation to work with orthonormal virtual orbitals and a CSF-driven algorithm employing Cholesky vectors to efficiently process two-electron integrals in noncanonical order on demand. The accuracy and performance of our restructured Cholesky decomposition (CD)-LMRSDCI method are summarized in the results section, while the final section provides concluding remarks.

II. THEORY

A. CI

The method of CI recovers electron correlation by expanding the exact wave function in the basis of N -electron CSFs, which are spin-adapted linear combinations of antisymmetrized products of some chosen one-particle basis. A series of such configurations may be generated by promoting electrons from occupied to unoccupied orbitals in the reference configurations. Unfortunately, the length of the CI expansion grows very quickly and diagonalization of the CI matrix in the basis of these CSFs becomes computationally prohibitive except for very small systems. A more tractable approach involves retaining only configurations generated

from the reference configurations via single and double excitations, i.e., the MRSDCI expansion given below^{43,44}

$$\Psi^{CI} = \sum_R c_R \Psi_R + \sum_{i,a} c_i^a \Psi_i^a + \sum_{ij,ab} c_{ij}^{ab} \Psi_{ij}^{ab}, \quad (2)$$

where $\{\Psi_R\}$ contains the set of reference configurations while $\{\Psi_i^a\}$ and $\{\Psi_{ij}^{ab}\}$ are singly and doubly excited configurations generated by promoting electrons in the reference configurations from internal orbitals $\{i, j\}$ to external orbitals $\{a, b\}$, respectively.

Given that the Hamiltonian has only one- and two-body operators, in an orthonormal basis, only configurations that differ from the reference configurations by at most double excitations will have nonzero matrix elements, so truncation at MRSDCI recovers the majority of the electron correlation.

The matrix formulation of the MRSDCI method requires large matrix diagonalizations to solve the eigenvalue problem $\mathbf{HC} = \mathbf{EC}$. In spite of a truncated CI expansion, the total number of terms may amount to billions, even for modestly sized systems. Standard methods for diagonalizing symmetric $N \times N$ matrices typically scale as $O(N^3)$ and are therefore ill suited, since the storage and computational time requirements would overwhelm current computational capacities. Instead, the Davidson diagonalization algorithm⁴⁵ is routinely used for this purpose, whereby a trial CI vector C^k (k th iteration), containing approximate CI expansion coefficients, is updated via the correction vector

$$\delta C_i^k = \frac{E^k C_i^k - \sigma_i^k}{H_{ii} - E^k}, \quad \text{where} \quad \sigma_i^k = \sum_j H_{ij} C_j^k. \quad (3)$$

Here H_{ij} represents the CI Hamiltonian matrix element between CSFs i and j , consisting primarily of two-electron integrals while C_j is the expansion coefficient of CSF j in the CI vector. σ^k is the Hamiltonian matrix-CI vector product and E^k refers to the approximate eigenvalue in the k th iteration. It is possible to avoid the explicit construction and storage of the CI Hamiltonian matrix since the corresponding matrix elements can be generated from required one/two-electron integrals and coupling coefficients (this is known as “direct CI”).⁴⁶ Due to the excessive length of the CI expansion, the σ vector is computed piecewise by matrix multiplication between matching portions of the Hamiltonian matrix and the CI vector. In practice, this further translates into matching the two-electron integrals/coupling coefficients with the appropriate components of the CI vector (expansion coefficients of each CSF). Constructing the σ vector in Eq. (3) is the centerpiece of the Davidson diagonalization algorithm and is routinely the most expensive step in a conventional MRSDCI calculation.

B. LMRSDCI

Local correlation approximations can be used to reduce the overall cost of MRSDCI and most other post-Hartree–Fock (HF) methods.^{24–35,39,47–59} In particular, the effective dimensions of both the Hamiltonian matrix and CI vector can be reduced via local truncation when a localized basis is chosen to span the occupied and virtual subspace of a molecule.

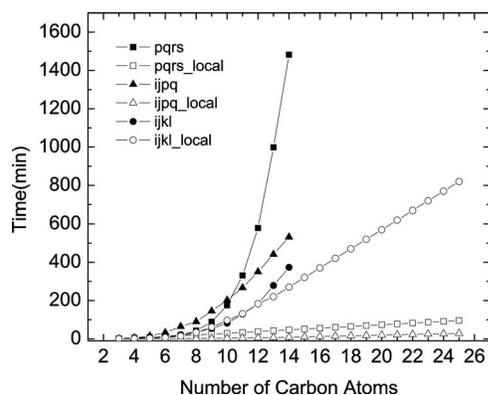


FIG. 1. CPU time (minutes) for constructing components of the σ vector that utilize three different representative integral classes, in local SDCI (open symbols) compared to nonlocal (conventional) SDCI (filled symbols). $\{i, j, k, \text{ and } l\}$ and $\{p, q, r, \text{ and } s\}$ refer to internal (occupied) and external (virtual) orbitals, respectively. The test cases here are linear alkanes, C_nH_{2n+2} , within a 6-31G** basis set.

While the localization of virtual orbitals is fraught with difficulties, various schemes to localize occupied orbitals are available.^{42,60–62} In our previous work, the virtual subspace was spanned by PAOs, formed by projecting the occupied space out of the AO basis. Despite the localized nature of the PAOs, they constitute a nonorthogonal set. As such, they complicate the MRSDCI formalism due to extra contractions with overlap matrices, e.g., as shown in Eq. (4), for the $(ij|kl)$ contribution to the σ vector

$$\sigma_{(ik,rs)}^\lambda = \sum_{\mu,j,l,p,q} A_{ijkl}^{\lambda\mu}(ij|kl) S_{rp} C_{(jl,pq)}^\mu S_{qs}, \quad (4)$$

where

$$(ij|kl) = \int \frac{1}{r_{12}} \varphi_i(1) \varphi_j(1) \varphi_k(2) \varphi_l(2) dr_1 dr_2. \quad (5)$$

Here $\{i, j, k, \text{ and } l\}$ and $\{p, q, r, \text{ and } s\}$ refer to internal and external (nonorthogonal) orbitals, respectively. λ and μ refer to CSFs involving double excitations to a common virtual domain from orbitals $\{i, k\}$ and $\{j, l\}$, respectively, that are occupied in the reference configuration. A is the coupling coefficient while the matrix S contains overlap integrals involving the set of external orbitals in the common virtual domain.

As seen in Fig. 1, constructing the portion of the σ vector in Eq. (4) that uses the $(ij|kl)$ integrals constitutes the rate limiting step in our previous LMRSDCI algorithm, where other integral classes are shown alongside for comparison.

Thus reducing the overall cost of the calculation requires a decrease in the computational overhead associated with Eq. (4). In the present work, we achieve this via two major changes to our scheme. First, by replacing the nonorthogonal PAOs with localized orthogonal virtual orbitals (LOVOs), we eliminate the need to contract integrals with overlap matrices and reduce the scaling prefactor accordingly. In particular, in an orthogonal basis Eq. (4) simplifies to

$$\sigma_{(ik,r's')}^\lambda = \sum_{\mu,j,l} A_{ijkl}^{\lambda\mu}(ij|kl) C_{(jl,r's')}^\mu, \quad (6)$$

where now $\{r', s'\}$ are orthogonalized, external orbitals in the common virtual domain of λ and μ . This simplification in sigma vector construction occurs for many integral classes containing internal orbital indices.^{28,29} Subotnik *et al.*⁶³ recently outlined how to generate such LOVOs for the purpose of electron correlation treatment within a local correlation setting, which we exploit fruitfully here.

Second, our previous integral-driven approach to σ vector construction used a disproportionately large amount of time (e.g., more than 90% for $C_{25}H_{52}$, the largest molecule considered) to search for the CSFs λ and μ that match up with a given $(ij|kl)$ integral. Our restructured scheme for constructing the σ vector is instead driven by the search for the correct $(ij|kl)$ integral given a pair of CSFs, λ and μ . In a configuration-driven setting, the burden rests on retrieving the correct $(ij|kl)$ integral given matching CSFs λ and μ . This shift from retrieval of the correct CSFs in the previous algorithm to retrieval of the correct integrals in the new algorithm requires a flexible way to process the four-indexed $(ij|kl)$ integrals since they will not be accessed in canonical order. To facilitate this, here we replace conventional four-indexed ERIs with two-indexed Cholesky vectors which permit us to build up the required $(ij|kl)$ directly via the corresponding Cholesky vectors.^{64–66}

In our current implementation, the migration from an integral-driven approach to the configuration-driven approach is only done for the sigma vector components containing the $(ij|kl)$ integrals, as they formed the bottleneck of our previous algorithm. Changes to the algorithm are summarized in pseudocodes provided in the Appendix. Construction of all other parts of the sigma vector is still calculated using the original integral-driven algorithm since those operations are not rate limiting. However, a Cholesky vector representation is adopted for all ERIs, which reduces overall storage requirements. This also simplifies our overall formalism by eliminating the use of test densities and Cauchy–Schwarz inequalities to prescreen the ERIs. LOVOs are used throughout to span the virtual space to simplify our working formalism, as discussed above.

C. Cholesky vectors

Representation of two-electron integrals by Cholesky vectors has been analyzed by various authors.^{64–70} Incomplete CD of the two-electron integral matrix, V , exploits the inherent linear dependence of basis sets used in electronic structure calculations and permits a more compact representation of four-indexed ERIs in terms of easier-to-process two-indexed Cholesky vectors as

$$V_{\nu\rho\sigma} = (\mu\nu|\rho\sigma) = \sum_m^n L_{\mu\nu}^m L_{\rho\sigma}^m, \quad (7)$$

where m indexes the elements in the Cholesky vector $L_{\mu\nu}$ of

dimension n in the atomic orbital (AO) basis. In our work, the construction of the Cholesky vectors begins by calculating the diagonal elements in V using conventional means to calculate the ERIs and sorting them into descending order. Then the steps shown in Eqs. (8a)–(8c) are carried out, whereby for clarity, $\{i, j, \text{ and } k\}$ are temporarily used in place of compound indices $\{\mu\nu\}$. Thus, the ERI $(\mu\nu|\rho\sigma)$ is represented by V_{ij} and is equivalent to $V_{\mu\nu\rho\sigma}$.

Starting with $i=1$, i.e., for the largest diagonal element we set

$$L_i^i = (V_{ii})^{1/2}. \quad (8a)$$

Then, for each i ,

$$L_j^i = \left(V_{ij} - \sum_{k=1}^{i-1} L_i^k L_j^k \right) \frac{1}{L_i^i}, \quad j = i+1, i+2, \dots, N, \quad (8b)$$

where N is the last index. The required column of ERIs, V_{ij} , is calculated on-the-fly using conventional means. In the following step, the remaining diagonal elements, V_{jj} , are updated according to

$$V_{jj} = (V_{jj} - (L_j^i)^2), \quad j = i+1, i+2, \dots, N \quad (8c)$$

and Eqs. (8a)–(8c) are repeated with the subsequent i .

Applying complete CD to the positive semidefinite ERI matrix represents additional work on top of evaluating the ERIs and presents no opportunities for overall reduction in computational effort. However, if CD is terminated after a preset decomposition threshold is reached (i.e., when the largest remaining diagonal element, V_{ii} , is smaller than the decomposition threshold), all fully formed Cholesky vectors are accurate to within machine precision while the remaining vectors incur an error no larger than the preset threshold.⁶⁶ A related issue is that a numerically stable algorithm requires pivoting among the diagonal elements such that the largest V_{ii} is used to start every cycle of Eqs. (8a)–(8c). The final number of Cholesky vectors may be controlled prior to the start of CD by leaving out columns of integrals in the two-electron integral matrix if their diagonal value falls below a threshold. In our work, we used a diagonal screening threshold of 1.0d-9, which is an order of magnitude below the lowest decomposition threshold, to ensure numerical stability during CD. Other than this prescreening of diagonal elements, no other numerical screening of ERIs was carried out. The maximum rank of the largest Cholesky vector is similarly determined by the CD threshold. When appropriate thresholds are applied to CD, significant savings in computational effort relative to conventional evaluation of two-electron integrals can be realized. Our implementation of the CD essentially follows Ref. 68. Similar to the transformation of two-electron integrals calculated in the AO basis, the Cholesky vectors also have to be transformed stepwise to the MO basis,

$$T_{iv}^m = \sum_{\mu} C_{\mu}^i L_{\mu\nu}^m, \quad (8d)$$

$$T_{ij}^m = \sum_{\nu} C_{\nu}^j T_{iv}^m.$$

$L_{\mu\nu}^m$ and T_{ij}^m are Cholesky vector elements in the AO and MO basis, respectively, while T_{iv}^m is an intermediate quantity between the two steps. C_{μ}^i refers to the MO coefficient of basis function μ in orbital i . This transformation is a less cumbersome procedure relative to the conventional, stepped four-indexed transformation of ERIs. Use of Cholesky vectors engenders an increased flexibility in the processing of ERIs as they may be generated from the stored Cholesky vectors via

$$(ij|kl) = \sum_{m=1}^n T_{ij}^m T_{kl}^m. \quad (9)$$

A flexible ERI construction is precisely what we desire since our rearranged algorithm with a configuration-driven σ vector update places the computational burden on the procedure to retrieve the correct integrals, once the corresponding partner CSFs are built. To further reduce computation time, the newly assembled ERIs are stored on disk for retrieval in subsequent passes instead of being rebuilt each iteration via the Cholesky vectors. Storage of these ERIs on disk is done not just for the $(ij|kl)$ integrals but for all MO ERIs with values above the screening threshold of 1.0d-9, as described in our earlier work.²⁹ Despite the tandem storage of the final MO ERIs along with the Cholesky vectors, there is nevertheless a reduction of secondary storage required compared to using regular ERIs. The relative savings are derived from the stepwise transformation of ERIs from the AO to the MO basis whereby the number of Cholesky intermediates (two-indexed) is substantially smaller than the number of partially transformed ERIs in the conventional form (four-indexed).

By making these three changes (using LOVOs, converting our algorithm involving construction of the σ vector portions that use $(ij|kl)$ type integrals from integral-driven to CSF-driven, and consequently using Cholesky vectors for all ERIs), we are able to speed up the rate-limiting step of our algorithm by more than two orders of magnitude compared to our previous implementation. This speedup, coupled with reduced integral storage, makes it possible to apply our new method to much larger molecules. However, the CD is done over the entire molecule to retain accuracy and as such it does not scale linearly. As a result, our restructured code scales between $O(N^2)$ and $O(N^3)$, but with a much smaller prefactor so as to make MRSDCI calculations on much larger molecules possible for the first time. Further computational savings is possible with the recently introduced atomic CD scheme to approximate ERIs.^{71–74} Investigations on the reduction in computational effort and its attendant accuracy limitations when applied to our LMRSDCI method are currently underway.

III. COMPUTATIONAL DETAILS

For our benchmark calculations, we used a series of linear hydrocarbon chains C_nH_{2n+2} ($n=3-50$) that were structurally optimized at the HF/6-31G** level with a constrained “linear,” staggered backbone motif. While larger hydrocar-

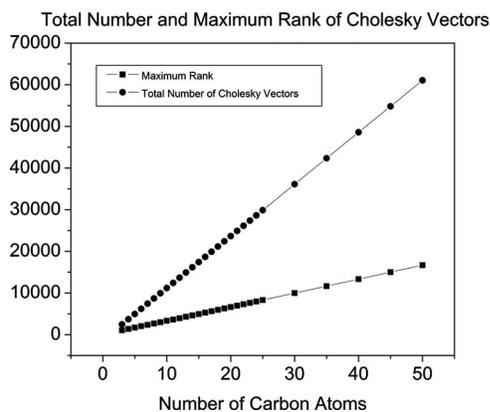


FIG. 2. Total number and maximum rank of Cholesky vectors for C_nH_{2n+2} ($n=3-50$) using a 6-31G** basis set. The CD threshold was set to 1.0d-7 while the screening threshold for the diagonal elements was set to 1.0d-9.

bons actually prefer a more globular structure to maximize dispersion interactions, this will result in variations in the correlating domain for occupied orbitals which are unrelated to our local correlation studies. A few other chemically varied systems were also included to show the effects of π electrons, heteroatoms, and electron delocalization on the performance of our LMRSDCI method. In addition, an MRSDCI calculation on trans-2-dodecene, $C_{12}H_{24}$, was carried out to verify accuracy in the multireference case. The reference wave function for the latter case was generated from a preceding complete active space self-consistent-field (CASSCF) calculation using a 6-31G** basis set. The active space consists of the σ , σ^* , π , and π^* orbitals (four electrons in four orbitals, denoted 4e/4o). All HF and CASSCF wave functions were generated using the MOLCAS 6.4 quantum chemistry package.⁷⁵ In the subsequent LMRSDCI calculations, orbital domains (corresponding to “spheres” and “cylinders”) are assigned to each localized orbital before applying local truncations. The procedure for generating these orbital domains was described previously.²⁸ All single and double excitations out of all occupied orbitals, subject to local truncations, are allowed, i.e., we do not employ a frozen core approximation, although our code is set up to do so.

For numerical stability during CD, full pivoting was employed. The total number of Cholesky vectors was controlled by prescreening the diagonal elements using a threshold of 1.0d-9 while the maximum numerical rank of the set of Cholesky vectors was controlled by the CD threshold of 1.0d-7.

IV. RESULTS

A. CD of two-electron integral matrix

Incomplete CD was performed on the two-electron integral matrix corresponding to the test system of C_nH_{2n+2} ($n=3-50$) using a 6-31G** basis set. As discussed above, the total number of Cholesky vectors was determined prior to commencement of the CD via screening precalculated diagonal elements in the ERI matrix. Entire columns of two-electron integrals containing the diagonal element were left out if the latter fell below the preset threshold of 1.0d-9. Furthermore, the diagonal elements were updated with the

TABLE I. Savings in storage of data elements in the ERI matrix when using a Cholesky vector representation for the two-electron integrals within a 6-31G** basis set. a denotes the total number of data elements in the Cholesky-decomposed integral matrix using a CD threshold of 1.0d-7. b denotes the total number of unique two-electron integrals—approximated as $[N^*(N+1)/2]$, where N is defined as $bf^{**}(bf+1)/2$ and bf refers to the number of contracted Gaussian basis functions. c denotes the percentage of stored data elements in (a) with respect to unique two-electron integrals.

C_nH_{2n+2}	a	b	$c=[a/b] \times 100\%$
C_3H_8	2 383 809	6 681 340	35.68
C_4H_{10}	4 775 408	18 638 565	25.62
C_5H_{12}	7 995 033	42 140 790	18.97
C_6H_{14}	12 118 620	82 953 640	14.61
C_7H_{16}	16 849 844	148 014 615	11.38
C_8H_{18}	22 534 200	245 433 090	9.18
C_9H_{20}	28 840 413	384 490 315	7.51
$C_{10}H_{22}$	36 344 052	575 639 415	6.31
$C_{11}H_{24}$	43 943 940	830 505 390	5.29
$C_{12}H_{26}$	52 888 525	1 161 885 115	4.55

completion of each Cholesky vector and the CD was stopped once the largest remaining diagonal element in the partially decomposed two-electron integral matrix fell below the CD threshold. In this way, the CD threshold determines the maximum rank of the largest Cholesky vector. Figure 2 reveals that the total number of Cholesky vectors and the number of elements in the largest Cholesky vector each scale linearly with system size, leading to an overall $O(N^2)$ scaling in the total number of stored data elements in the Cholesky-decomposed ERI matrix.

We illustrate the savings due to a Cholesky vector representation of ERIs in Table I. The reduction in the number of stored data elements (for conventional two-electron integrals, these are the explicit numerical values of each electron integral while in the Cholesky vector representation, these refer to the elements of each Cholesky vector) is progressively pronounced as we go to larger systems, with >95% of the storage requirement removed by using a CD, even for a molecule with only 12 heavy atoms. Further savings could be derived by storing only numerically significant components of each Cholesky vector. However, this would have required a supporting auxiliary vector to keep track of indices and we have not implemented this in our present work.

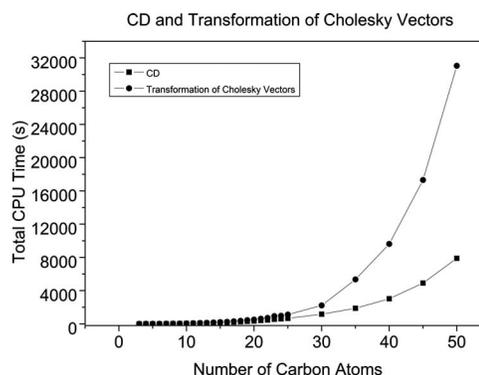


FIG. 3. Total CPU time for CD of the AO two-electron integrals and transformation of the Cholesky vectors to an MO basis for linear alkanes C_nH_{2n+2} within a 6-31G** basis set.

TABLE II. Total energies (hartrees) for C_nH_{2n+2} at different CD thresholds using a 6-31G** basis set.

Molecule	(MR)SDCI/ CD-(MR)SDCI ^a	CD-SDCI/ (1.0d-6) ^b	CD-SDCI (1.0d-7) ^c	CD-SDCI (1.0d-8) ^d
C_4H_{10}	-157.607 192 6	-157.607 188 0	-157.607 195 9	-157.607 193 8
C_6H_{14}	-235.795 784 0	-235.795 761 3	-235.795 789 1	-235.795 784 0
C_8H_{18}	-313.972 066 9	-313.972 048 6	-313.972 063 1	-313.972 066 1
$C_{10}H_{22}$	-392.138 700 6	-392.138 690 5	-392.138 700 5	-392.138 700 5
$C_{12}H_{26}$	-470.297 503 5	-470.297 487 0	-470.297 507 0	-470.297 502 1
$C_{14}H_{30}$	-548.449 777 9	-548.449 761 2	-548.449 778 2	-548.449 777 5
$C_{12}H_{24}^e$	-468.628 197 2	-468.628 102 2	-468.628 160 9	-468.628 195 5

^a(MR) SDCI and CD-(MR)SDCI calculations with a CD threshold of 1.0d-12 for the latter.

^bCD-SDCI calculations with local correlation treatment and a CD threshold of 1.0d-6.

^cCD-SDCI calculations with local correlation treatment and a CD threshold of 1.0d-7.

^dCD-SDCI calculations with local correlation treatment and a CD threshold of 1.0d-8.

^eMRSDCI calculation with a (4e/4o) active space.

While full CD is inherently $O(N^3)$, the effective numerical rank of the two-electron integral matrix is reduced via the CD decomposition threshold. As such, the cubic scaling of full CD is not explicitly manifested. The subsequent transformation of the Cholesky vectors from the AO to the MO basis involves contraction over two AO indices instead of the usual four when using conventional ERIs. In our work, the transformation is broken down into two half steps Eq. (8d) and represents substantial simplification and savings in coding, storage, and computation time. The effective scaling of both the evaluation and transformation of the Cholesky vectors to the MO basis lies between $O(N^2)$ and $O(N^3)$ and the total central processing unit (CPU) time taken for both procedures are summarized in Fig. 3.

B. Accuracy of Cholesky vectors

In Tables II–IV, we examine several aspects of using a Cholesky vector representation for the two-electron integrals in our local correlation work. SDCI calculations were performed for the series of linear alkane chains, C_nH_{2n+2} (Table III) and other more chemically varied species (Table IV). The latter includes hex-3-ene, hex-3-yne, cytosine, 1,1,1-trifluorobutane, and phenol, along with an LMRSDCI example for trans-6-dodecene ($C_{12}H_{24}$), which represents the largest system that our implementation of nonlocal MRSDCI (exact MRSDCI, i.e., with no screening or local correlation

approximations) can handle.

The effects of using different CD thresholds on the converged total energies are summarized in Table II. When a strict CD threshold of 1.0d-12 is applied, total energies from CD-(MR)SDCI calculations are identical to regular (MR)SDCI runs. We find that a CD threshold of 1.0d-7 is sufficient to retain accuracy at the 5 μ H level for the CD-SDCI calculations while the CD-MRSDCI result shows a deviation of less than 40 μ H.

In Table III, the Cholesky vectors are used in conjunction with LOVOs within a local correlation framework (CD-LSDCI) to study linear alkane chains. We find that up to 98.5% of the correlation energy with respect to nonlocal, non-CD SDCI can be recovered. Selected molecules containing π electrons, moderate electron delocalization, and heteroatoms are included in Table IV. Extensively delocalized π systems will require more centers to be included in the local correlation domains while the choice of a common basis type to treat heteroatoms (N, O, and F) containing more electrons than C is expected to lower the proportion of correlation energy recovered relative to exact SDCI. In all cases, at least 98.0% of the SDCI correlation energy was recovered. The least correlation energy recovery was obtained for cytosine ($C_4H_5N_3O$), which is indicative of the challenges associated with the combined effects of electron delocalization and heteroatoms.

TABLE III. Total energies (hartrees) for linear alkanes using a 6-31G** basis set.

Molecule	(MR)SDCI/ CD-(MR)SDCI ^a	CD-(MR)SDCI/ (1.0d-7) ^b	CD-(MR)LSDCI (1.0d-7) ^c	%correl.E. ^d
C_4H_{10}	-157.607 192 6	-157.607 195 9	-157.603 573 4	99.0%
C_6H_{14}	-235.795 784 0	-235.795 789 1	-235.788 806 6	98.7%
C_8H_{18}	-313.972 066 9	-313.972 063 1	-313.962 747 0	98.6%
$C_{10}H_{22}$	-392.138 700 6	-392.138 700 5	-392.126 605 4	98.5%
$C_{12}H_{26}$	-470.297 503 5	-470.297 507 0	-470.281 013	98.5%
$C_{14}H_{30}$	-548.449 777 9	-548.449 778 2	-548.434 394 0	98.5%

^a(MR) SDCI and CD-(MR)SDCI calculations with a CD threshold of 1.0d-12 for the latter.

^bCD-(MR)SDCI with a CD threshold of 1.0d-7.

^cCD-L(MR)SDCI with local correlation treatment and a CD threshold of 1.0d-7.

^dPercent correlation energy recovered in CD-(MR)LSDCI (fourth column) relative to (MR)SDCI (second column).

TABLE IV. Total energies (hartrees) for hex-3-ene, hex-3-yne, cytosine, 1,1,1-trifluorobutane, phenol, and trans-6-dodecene using a 6-31G** basis set.

Molecule	(MR)SDCI/ CD-(MR)SDCI ^a	CD-(MR)SDCI/ (1.0d-7) ^b	CD-(MR)LSDCI (1.0d-7) ^c	%correl.E. ^d
C ₆ H ₁₂	-234.970 619 3	-234.970 617 4	-234.959 024	98.5%
C ₆ H ₁₀	-233.741 315 0	-234.741 319 1	-233.734 576	99.1%
C ₄ H ₅ N ₃ O	-393.602 890 1	-393.602 895 8	-393.580 444	98.0%
C ₄ H ₇ F ₃	-454.827 785 3	-454.827 780 2	-454.809 208	98.1%
C ₆ H ₅ OH	-306.403 795 5	-306.403 801 1	-306.389 650	98.3%
C ₁₂ H ₂₄ ^e	-468.628 197 2	-468.628 160 9	-468.615 102 1	98.9%

^a(MR)SDCI and CD-(MR)SDCI calculations with a CD threshold of 1.0d-12 for the latter.

^bCD-(MR)LSDCI with a CD threshold of 1.0d-7.

^cCD-L(MR)SDCI with local correlation treatment and a CD threshold of 1.0d-7.

^dPercent correlation energy recovered in CD-(MR)LSDCI (fourth column) relative to (MR)SDCI (second column).

^eMRSDCI calculation with a (4e/4o) active space.

C. σ vector construction in CD-LSDCI

The construction of the σ vector during the Davidson diagonalization is organized as follows. In the first step, two-electron integrals in the MO basis belonging to different classes are assembled from the transformed Cholesky vectors and stored on disk. The ERIs were segregated into $(ij|kl)$, $(ij|ka)$, $(ij|ab)$, $(ia|bc)$, and $(ab|cd)$ classes where $\{i, j, k$, and $l\}$ refer to internal orbitals while $\{a, b, c$, and $d\}$ denote external orbitals (the LOVOs). In each case, only two-electron integrals which pass the screening threshold of 1.0d-9 were stored or used to construct the relevant portions of the σ vector. For the σ vector components utilizing the $(ij|kl)$ integrals, the configuration-driven approach outlined earlier is used, while the other components are calculated with the original integral-driven algorithm. Figure 4 shows the total CPU time associated with constructing the σ vector for three representative integral classes $(ij|kl)$, $(ij|ab)$, and $(ab|cd)$.

Significantly, the previous bottleneck (see Fig. 1) associated with calculating the portion of the sigma vector involving the $(ij|kl)$ integrals has disappeared and we observe a speedup of more than 100 times (for C₂₅H₅₂) for the previously rate limiting step. The switch to a configuration-driven scheme for constructing the σ vector portion that

utilizes the $(ij|kl)$ integrals is largely responsible for the major reduction in computational overhead, although the speedup is further enhanced by the use of LOVOs. We found that the reduction in time due to use of the latter is about 10% in the best case scenario.

Instead, the new bottleneck involves assembling the MO integrals from the Cholesky vectors in the first step, prior to σ vector construction. As the total number and maximum rank of each Cholesky vector scales with system size, leading to an overall $O(N^{2-3})$ scaling in computational effort, the cost of assembling the integrals will tend to dominate at larger system size. The cost for assembling the $(ij|ka)$ integrals from the Cholesky vectors (not shown in Fig. 4) is less than twice the cost of assembling the $(ij|kl)$ integrals for the largest molecule considered (C₅₀H₁₀₂), and therefore does not constitute a rate-limiting step.

The total CPU time for our current method utilizing Cholesky vectors (CD-LMRSDCI) is compared against our previous implementation in Fig. 5. The reduction in time is especially significant for the lower members of the alkane series, while the cost of CD/transformation of Cholesky vectors and assembling the integrals starts to dominate toward larger systems.

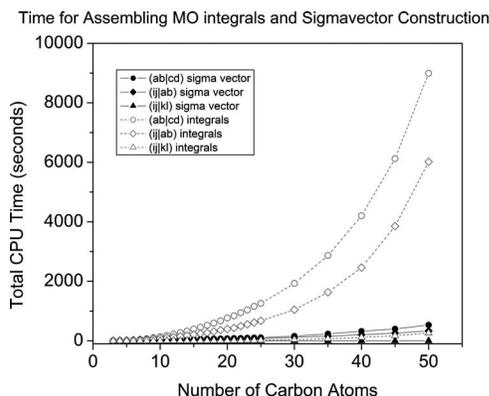


FIG. 4. CPU time for assembling two-electron integrals in the MO basis from transformed Cholesky vectors and subsequent construction of the σ vector, for linear alkanes C_nH_{2n+2} within a 6-31G** basis set.

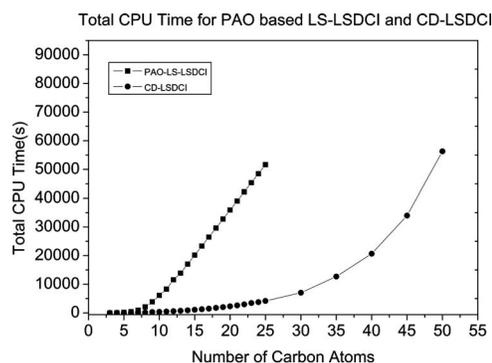


FIG. 5. Total CPU time (seconds) for our previous nonorthogonal PAO-based linear scaling LSDCI and the current implementation of LSDCI using Cholesky vectors, a CSF-driven scheme for the $(ij|kl)$ integrals, along with localized orthogonal virtual orbitals. The test set again is linear alkanes C_nH_{2n+2} within a 6-31G** basis set.

V. CONCLUSIONS

We have restructured our LMRSDCI algorithm and code so that it can be used to provide a correlated wave function description of molecules containing up to 50 heavy atoms. The bottleneck in our previous PAO-based LMRSDCI method involves the construction of the portion of the Hamiltonian matrix-CI vector product that uses the all-internal $(ij|kl)$ type integrals, as shown in (4). We were able to reduce the overall cost of our method by more than an order of magnitude, by the following algorithmic changes.

First, the most time-consuming step in our previous approach involved finding CSFs λ and μ which match up with the $(ij|kl)$ integral in an integral-driven formulation. To reduce the overhead associated with this step, the process was restructured to follow a configuration-driven approach, i.e., the calculation of the σ vector components containing all-internal integrals is now driven by the search for $(ij|kl)$ integrals needed when given a pair of CSFs λ and μ . This requires a flexible way of processing the four-indexed ERIs. To facilitate this, two-indexed Cholesky vectors are used to represent the ERIs such that required integrals in the MO basis are built from these Cholesky vectors. Once assembled from these vectors, the two-electron integrals are stored on disk for reuse in subsequent passes. In addition, we simplified our working formalism by switching to an orthogonal set of virtual orbitals so as to obviate the need for contraction with metric matrices. To this end, we replaced the nonorthogonal PAOs with localized, orthogonal virtual orbitals as proposed by Subotnik *et al.* For the largest molecule in our work, $C_{50}H_{102}$, generation of the LOVOs constitutes less than 3% of the total computation time.

The rate limiting steps in our current method are the incomplete CD of the ERI matrix and the subsequent transformation of the Cholesky vectors to the molecular orbital basis, together with assembling the ERIs from these Cholesky vectors for ERIs containing at least two external indices. Unfortunately, incomplete CD is inherently not a linear scaling procedure, since the total number and the maximum rank of the Cholesky vectors both scale with system size. Nevertheless, a parallel implementation of CD, which scales well with system size, has been reported⁷⁶ and density fitting schemes may be useful in reducing the current overhead involved during the generation of Cholesky vectors.^{71-74,77-84} This offers a path forward to even further reductions in cost of LMRSDCI calculations on large molecules.

ACKNOWLEDGMENTS

We are grateful to the National Science Foundation for support of this work. One of the authors (T.S.C.) thanks the Agency for Science, Technology and Research A*STAR for funding.

APPENDIX: CSF-DRIVEN ALGORITHM TO CONSTRUCT THE σ VECTOR

We summarize in the following pseudocodes A and B the restructuring of the algorithm carried out to migrate from the original integral-driven scheme to a configuration-driven one during the construction of the σ vector using $(ij|kl)$ type integrals. $\{i, j, k, \text{ and } l\}$ are used to index internal orbitals while CSF_{λ} and CSF_{μ} are vectors, representing distinct CSFs. Each vector contains information that specifies the occupation number of individual molecular orbitals along with the spin eigenfunction that together defines a CSF.

In the original scheme (A), the first step involves the retrieval of $(ij|kl)$ type integrals subject to numerical screening and local truncation approximations. The integral in turn determines the matching CSF_{λ} and CSF_{μ} which remain to be found. The matching CSF_{λ} and CSF_{μ} may only differ in the distribution of *singly occupied* internal orbitals $\{i, j, k, \text{ and } l\}$ while the occupation patterns in other molecular orbitals are otherwise constrained to be identical. Once the matching CSFs are found, they are processed along with the $(ij|kl)$ integral to update the σ vector. We have found that the constrained search for the matching CSFs such that each has the required occupation pattern in the prefixed indices $\{i, j, k, \text{ and } l\}$ take up the majority of our CPU time and hence the motivation for the switch to a configuration-driven mode to construct the σ vector.

The alternative scheme (B) begins with an initial call for an unconstrained search for some CSF_{λ} where no restrictions are placed on the occupation pattern of the internal orbitals. Then it is processed further to locate the number and distribution of singly occupied orbitals in the internal space. For a fixed pair of singly occupied internal orbitals (e.g., i and k among the singly occupied orbitals) in CSF_{λ} , the required CSF_{μ} is built up such that its occupation pattern is identical with CSF_{λ} other than the orbitals $\{i, j, k, \text{ and } l\}$ where $\{j \text{ and } l\}$ are unconstrained orbital indices in the internal space. The required $(ij|kl)$ integral is then assembled from the relevant Cholesky vectors (T_{ij}, T_{kl}) and further written to disk for reuse in later passes. The assembled integral and partner CSFs are then processed to update the σ vector.

1. Pseudocode A

```

Do i=1, num_internal
  Do j=1, i-1
    Do k=1, j-1
      Do l=1, k-1
        ! Perform numerical screening/local truncations
        and retrieve the (ij|kl) integrals from disk.
        ! Call routine to build CSF_lambda and CSF_mu with
        constraints in occupation pattern at the i, j, k,
        and l indices.
        ! Process CSF_lambda, CSF_mu, and (ij|kl) integrals.
        Call sigma vector routine H*C.
      Enddo
    Enddo
  Enddo
Enddo

```

2. Pseudocode B

! Call routine to perform unconstrained search for CSF_λ with no restriction on the occupation pattern of internal orbitals. Discard any CSF_λ with no singly occupied orbitals.

! Search and process the number of singly occupied orbitals in CSF_λ. Perform loops over all distinct pairs of singly occupied orbitals in CSF_λ.

Do i=1, num_singles_in_CSF_λ

 Do k=i+1, num_singles_in_CSF_λ

 ! Perform local truncations

 ! The required CSF_μ differs from CSF_λ at the indices *j* and *l*. Find all possible CSF_μ by looping over distinct *j* and *l*.

 Do j=1, num_internal_

 ! Check (*j*.ne. *i*.and. *j*.ne. *k*)

 Do l=j+1, num_internal

 ! Check (*l*.ne. *i*.and. *l*.ne. *k*)

 ! Perform local truncations/numerical screening

 ! Build CSF_μ. The occupation pattern is

 identical to CSF_λ other than

 ! {*i*, *j*, *k*, and *l*}

 ! Assemble (*ij|kl*) integral from relevant

 Cholesky vectors *T_{ij}* and *T_{kl}* in the first pass.

 Write integral to disk for reuse in later passes.

 ! Process CSF_λ, CSF_μ, and (*ij|kl*) integrals.

 Call σ vector routine H^{*}C.

 Enddo

 Enddo

Enddo

Enddo

¹ S. R. Langhoff and E. R. Davidson, *Int. J. Quantum Chem.* **8**, 61 (1974).

² E. R. Davidson and D. W. Silver, *Chem. Phys. Lett.* **52**, 403 (1977).

³ P. E. M. Siegbahn, *Chem. Phys. Lett.* **55**, 386 (1978).

⁴ P. J. Bruna, S. D. Peyerimhoff, and R. J. Buenker, *Chem. Phys. Lett.* **72**, 278 (1980).

⁵ P. G. Burton, R. J. Buenker, P. J. Bruna, and S. D. Peyerimhoff, *Chem. Phys. Lett.* **95**, 379 (1983).

⁶ R. Ahlrichs, P. Scharf, and C. Ehrhardt, *J. Chem. Phys.* **82**, 890 (1985).

⁷ R. J. Gdanitz and R. Ahlrichs, *Chem. Phys. Lett.* **143**, 413 (1988).

⁸ R. Fink and V. Staemmler, *Theor. Chim. Acta* **87**, 129 (1993).

⁹ R. J. Gdanitz, *Int. J. Quantum Chem.* **85**, 281 (2001).

¹⁰ A. Venkatnathan, A. B. Szilva, D. Walter, R. J. Gdanitz, and E. A. Carter, *J. Chem. Phys.* **120**, 1693 (2004).

¹¹ P. G. Szalay and R. J. Bartlett, *Chem. Phys. Lett.* **214**, 481 (1993).

¹² P. G. Szalay and R. J. Bartlett, *J. Chem. Phys.* **103**, 3600 (1995).

¹³ W. D. Laidig and R. J. Bartlett, *Chem. Phys. Lett.* **104**, 424 (1984).

¹⁴ I. Hubač, J. Pittner, and P. Čársky, *J. Chem. Phys.* **112**, 8779 (2000).

¹⁵ I. Hubač and S. Wilson, *Adv. At., Mol., Opt. Phys.* **34**, 4259 (2001).

¹⁶ J. Paldus and X. Li, *J. Chem. Phys.* **118**, 6769 (2003).

¹⁷ J. Pittner, *J. Chem. Phys.* **118**, 10876 (2003).

¹⁸ S. Chattopadhyay, D. Pahari, and D. Mukherjee, *J. Chem. Phys.* **120**, 5968 (2004).

¹⁹ X. Li and J. Paldus, *J. Chem. Phys.* **120**, 5890 (2004).

²⁰ M. Musiał and R. J. Bartlett, *J. Chem. Phys.* **121**, 1670 (2004).

²¹ J. Pittner, X. Li, and J. Paldus, *Mol. Phys.* **103**, 2239 (2005).

²² O. Demel and J. Pittner, *J. Chem. Phys.* **124**, 144112 (2006).

²³ F. A. Evangelista, W. D. Allen, and H. F. Schaefer, *J. Chem. Phys.* **127**, 024102 (2007).

²⁴ G. Reynolds, T. J. Martinez, and E. A. Carter, *J. Chem. Phys.* **105**, 6455 (1996).

²⁵ G. Reynolds and E. A. Carter, *Chem. Phys. Lett.* **265**, 660 (1997).

²⁶ D. Walter and E. A. Carter, *Chem. Phys. Lett.* **346**, 177 (2001).

²⁷ D. Walter, A. B. Szilva, K. Niedfeldt, and E. A. Carter, *J. Chem. Phys.* **117**, 1982 (2002).

²⁸ D. Walter, A. Venkatnathan, and E. A. Carter, *J. Chem. Phys.* **118**, 8127 (2003).

²⁹ T. S. Chwee, A. B. Szilva, R. Lindh, and E. A. Carter, *J. Chem. Phys.* **128**, 224106 (2008).

³⁰ P. Pulay, *Chem. Phys. Lett.* **100**, 151 (1983).

³¹ S. Sæbø and P. Pulay, *Chem. Phys. Lett.* **113**, 13 (1985).

³² P. Pulay and S. Sæbø, *Theor. Chim. Acta* **69**, 357 (1986).

³³ S. Sæbø and P. Pulay, *J. Chem. Phys.* **86**, 914 (1987).

³⁴ S. Sæbø and P. Pulay, *J. Chem. Phys.* **88**, 1884 (1988).

³⁵ S. Sæbø and P. Pulay, *Annu. Rev. Phys. Chem.* **44**, 213 (1993).

³⁶ J. C. Slater, *Phys. Rev.* **34**, 1293 (1929).

³⁷ E. U. Condon, *Phys. Rev.* **36**, 1121 (1930).

³⁸ M. Häser and R. Ahlrichs, *J. Comput. Chem.* **10**, 104 (1989).

³⁹ M. Schütz and H. J. Werner, *J. Chem. Phys.* **114**, 661 (2001).

⁴⁰ D. S. Lambrecht, B. Doser, and C. Oschsenfeld, *J. Chem. Phys.* **123**, 102 (2005).

⁴¹ D. S. Lambrecht and C. Oschsenfeld, *J. Chem. Phys.* **123**, 184101 (2005).

⁴² J. Pipek and P. G. Mezey, *J. Chem. Phys.* **90**, 4916 (1989).

⁴³ R. J. Buenker and S. D. Peyerimhoff, *Theor. Chim. Acta* **35**, 33 (1974).

⁴⁴ R. J. Buenker, S. D. Peyerimhoff, and W. Butscher, *Mol. Phys.* **35**, 771 (1978).

⁴⁵ E. R. Davidson, *J. Comput. Phys.* **17**, 87 (1975).

⁴⁶ B. O. Roos, *Chem. Phys. Lett.* **15**, 153 (1972).

⁴⁷ R. B. Murphy, M. D. Beachy, R. A. Friesner, and M. N. Ringnalda, *J. Chem. Phys.* **103**, 1481 (1995).

⁴⁸ C. Hampel and H. J. Werner, *J. Chem. Phys.* **104**, 6286 (1996).

⁴⁹ P. E. Maslen and M. Head-Gordon, *J. Chem. Phys.* **109**, 7093 (1998).

⁵⁰ G. Hetzer, P. Pulay, and H. J. Werner, *Chem. Phys. Lett.* **290**, 143 (1998).

⁵¹ A. El Azhary, A. Rauhut, P. Pulay, and H. J. Werner, *J. Chem. Phys.* **108**, 5185 (1998).

⁵² M. Schütz, G. Hetzer, and H. J. Werner, *J. Chem. Phys.* **111**, 5691 (1999).

⁵³ R. A. Friesner, R. B. Murphy, M. D. Beachy, M. N. Ringnalda, W. T. Pollard, B. D. Dunietz, and Y. Cao, *J. Phys. Chem. A* **103**, 1913 (1999).

⁵⁴ M. S. Lee, P. E. Maslen, and M. Head-Gordon, *J. Chem. Phys.* **112**, 3592 (2000).

⁵⁵ H. J. Werner, F. R. Manby, and P. J. Knowles, *J. Chem. Phys.* **118**, 8149 (2003).

⁵⁶ M. Schütz, H. J. Werner, R. Lindh, and F. R. J. Manby, *Chem. Phys.* **121**, 737 (2004).

⁵⁷ E. A. Carter and D. Walter, *Encyclopedia of Computational Chemistry*, edited by P. V. R. Schleyer, N. L. Allinger, T. Clark, J. Gasteiger, P. A. Kollman, H. F. Schaefer III, and P. R. Schreiner (Wiley, Chichester, 2004).

⁵⁸ Y. Jung, A. Sodt, P. M. W. Gill, and M. Head-Gordon, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 6692 (2005).

⁵⁹ S. Li, J. Shen, W. Li, and Y. Jiang, *J. Chem. Phys.* **125**, 074109 (2006).

⁶⁰ S. F. Boys, *Rev. Mod. Phys.* **32**, 296 (1960).

⁶¹ C. Edmiston and K. Ruedenberg, *Rev. Mod. Phys.* **35**, 457 (1963).

⁶² F. Aquilante, T. B. Pedersen, A. Sánchez de Merás, and H. Koch, *J. Chem. Phys.* **125**, 174101 (2006).

⁶³ J. E. Subotnik, A. D. Dutoi, and M. Head-Gordon, *J. Chem. Phys.* **123**, 114108 (2005).

⁶⁴ N. H. F. Beebe and J. Linderberg, *Int. J. Quantum Chem.* **12**, 683 (1977).

⁶⁵ I. Røeggen and E. Wisløff-Nielsen, *Chem. Phys. Lett.* **132**, 154 (1986).

⁶⁶ D. W. O'neal and J. Simons, *Int. J. Quantum Chem.* **36**, 673 (1989).

⁶⁷ S. Wilson, *Comput. Phys. Commun.* **58**, 71 (1990).

⁶⁸ H. Koch, A. Sánchez de Merás, and T. B. Pedersen, *J. Chem. Phys.* **118**, 9481 (2003).

⁶⁹ T. B. Pedersen, A. M. J. Sánchez de Merás, and H. Koch, *J. Chem. Phys.* **120**, 8887 (2004).

⁷⁰ F. Aquilante, T. B. Pedersen, and R. Lindh, *J. Chem. Phys.* **126**, 194106 (2007).

⁷¹ F. Aquilante, R. Lindh, and T. B. Pedersen, *J. Chem. Phys.* **127**, 114107 (2007).

⁷² F. Aquilante, L. Gagliardi, T. B. Pedersen, and R. Lindh, *J. Chem. Phys.* **130**, 154107 (2009).

⁷³ J. Boström, F. Aquilante, T. B. Pedersen, and R. Lindh, *J. Chem. Theory Comput.* **5**, 1545 (2009).

⁷⁴ T. B. Pedersen, F. Aquilante, and R. Lindh, *Theor. Chim. Acta* **124**, 1 (2009).

⁷⁵ G. Karlström, R. Lindh, P.-Å. Malmqvist, B. O. Roos, U. Ryde, V. Veryazov, P.-O. Widmark, M. Cossi, B. Schimmelpfennig, P. Neogady, and L. Seijo, *Comput. Mater. Sci.* **28**, 222 (2003).

⁷⁶ I. Røeggen and T. Johansen, *J. Chem. Phys.* **128**, 194107 (2008).

- ⁷⁷ S. F. Boys and I. Shavitt, University of Wisconsin, Report No. WIS-AF-13, 1959.
- ⁷⁸ M. W. Feyereisen, G. Fitzgerald, and A. Komornicki, *Chem. Phys. Lett.* **208**, 359 (1993).
- ⁷⁹ O. Vahtras, J. E. Almlöf, and M. W. Feyereisen, *Chem. Phys. Lett.* **213**, 514 (1993).
- ⁸⁰ A. P. Rendell and T. J. Lee, *J. Chem. Phys.* **101**, 400 (1994).
- ⁸¹ K. Eichkorn, O. Treutler, H. Ohm, M. Häser, and R. Ahlrichs, *Chem. Phys. Lett.* **240**, 283 (1995).
- ⁸² K. Eichkorn, F. Weigend, O. Trutler, and R. Ahlrichs, *Theor. Chim. Acta* **97**, 119 (1997).
- ⁸³ C. Hättig and F. Weigend, *J. Chem. Phys.* **113**, 5154 (2000).
- ⁸⁴ A. Sodt, J. E. Subotnik, and M. Head-Gordon, *J. Chem. Phys.* **125**, 194109 (2006).