

Local weak-pairs pseudospectral multireference configuration interaction

Derek Walter, Andrew B. Szilva, Keith Niedfeldt, and Emily A. Carter

Department of Chemistry and Biochemistry, University of California, Los Angeles, California 90095-1569

(Received 8 February 2002; accepted 1 May 2002)

We present a new reduced scaling multireference singles and doubles configuration interaction (MRSDCI) algorithm based upon the combination of local correlation and pseudospectral methods. Taking advantage of the locality of the Coulomb potential, the weak-pairs approximation of Saebø and Pulay is employed to eliminate configurations having simultaneous excitations out of pairs of distant, weakly interacting orbitals. In conjunction with this, the pseudospectral approximation is used to break down the most time-consuming two-electron integrals into a product of intermediate quantities depending on no more than two orbital indices. The resulting intermediate quantities are then used directly in the CI equations to substantially reduce the number of floating point operations required for diagonalization of the Hamiltonian. Additionally, our CI algorithm is based upon the symmetric group graphical approach CI (SGGA-CI) of Duch and Karwowski. For the purpose of developing reduced scaling CI algorithms, this approach has some important advantages. The most important of these advantages are the on-the-fly calculation of integral coupling coefficients and the separation of the spin and spatial parts of the wave function, which simplifies implementation of local correlation approximations. We apply the method to determine a series of binding energies in hydrocarbons and show that the approximate method predicts binding energies that are within a few kcal/mol of those predicted by the analytic nonlocal method. For large molecules, the local pseudospectral method was shown to be over 7 times as fast as the analytic nonlocal method. We also carry out a systematic study on the performance of different basis sets in the weak-pairs method. It was determined that triple- ζ basis sets were capable of recovering only 99.0% of the correlation energy, whereas double- ζ basis sets recovered 99.9% of the correlation energy. © 2002 American Institute of Physics. [DOI: 10.1063/1.1487816]

INTRODUCTION

Electron correlation in *ab initio* quantum chemistry is typically described using one of three principal methods: Møller–Plesset (MP) perturbation theory, coupled cluster (CC) theory, or configuration interaction (CI). In terms of the computational resources required, these methods all scale very poorly with the size of the system under consideration. Methods based on various degrees of Møller–Plesset perturbation theory (designated as MP x , where x is the degree of perturbation theory used) are typically the most economical of these methods and offer the important advantage of being rigorously size extensive. For these reasons, MP x became arguably the most widely used correlation technique. However, MP2 still scales as nN^4 , while MP3 and MP4 restricted to single, double, and quadruple substitutions scale as n^2N^4 and n^3N^3 , respectively. Here, n is the number of orbitals occupied in the reference and N is the number of basis functions. Clearly, the cost of even these relatively inexpensive *ab initio* treatments of electron correlation will become prohibitive for larger systems. Recognizing this fact, various researchers have developed reduced scaling MP x methods. Saebø and Pulay developed reduced scaling MP2, MP3, and MP4 methods based on local correlation techniques.^{1–3} These methods were shown to be quite successful, recovering greater than 99% of the nonlocal MP x correlation energy.

Martinez and Carter implemented a pseudospectral (PS) version of MP2 and MP3.⁴ Murphy *et al.* then developed a MP2 method that simultaneously takes advantage of local correlation techniques and PS methods.^{5,6} This method was shown to be very successful at describing conformational changes and was incorporated into the Jaguar quantum chemistry software package.⁷ Subsequent work by Werner and co-workers has led to a linear scaling MP2 method based on local correlation.^{8–10} Maslen, Lee, and Head-Gordon have developed a local noniterative MP x method that makes use of both atom-centered occupied and virtual orbitals.^{11–14} This “atoms-in-molecules” approach has been quite successfully applied to MP2, MP4, and various perturbative corrections to CC methods. Along slightly different lines, Ayala and Scuseria used a Laplace transform scheme (originally conceived of by Almlöf¹⁵) to recast the MP2 equations entirely in the atomic orbitals (AO) basis.¹⁶ By defining AO domains and neglecting weak interactions, they were able to achieve a linear scaling MP2 method that gives energies differing from the canonical MP2 energies by only microhartrees. In spite of this success, MP x suffers from some important deficiencies. First, compared to CC and CI, MP x generally recovers less correlation energy. Second, MP x is known to perform poorly when the reference is not a true spin eigenfunction.¹⁷ In addition, MP x performs poorly for open-shell systems and therefore cannot be used to describe covalent bond breaking.

Although there are multireference versions of MP_x ,^{18–20} the question of what the zeroth-order Hamiltonian should be does not have an unequivocal answer.

Coupled cluster methods are considered by many to be the most accurate, size extensive correlation methods. Of course, the scaling of CC is even more severe than MP_x . The coupled cluster with single and double substitutions (CCSD) method scales as N^6 , while CCSD with a perturbative triples correction [CCSD(T)] scales as N^7 . As in the case of MP_x , extensive research has been devoted towards reducing the severe scaling of coupled cluster methods. Similar in spirit to the local MP2 methods of Saebø and Pulay, Hampel and Werner developed a fully local CCSD method.²¹ This method was recently augmented to take advantage of integral direct techniques, resulting in a local CCSD program that was applied to treat systems having 300 correlated electrons using 1000 basis functions.²² Additionally, Schütz and Werner developed a local method for treating the perturbative triples correction to CCSD,²³ which was then developed further into a linear scaling perturbative triples technique by Schütz.²⁴ Scuseria and Ayala implemented a linear scaling CCSD method based on the same Laplace transform technology used in their linear scaling MP2 method.²⁵ As is typical, this method was shown to recover more than 99% of the correlation energy. Constans *et al.* subsequently applied this to the perturbative triples correction to CCSD as well.²⁶ Unfortunately, similar to MP2, CC methods are not easily used with open-shell or multireference SCF wave functions. Although there has been substantial progress in the area of using CC with high-spin coupled spin-restricted reference wave functions,^{27–31} the treatment of low-spin and open-shell singlet cases remains problematic and is an active area of research.^{32,33} The difficulty in using CC with multireference wave functions is particularly troubling as it hinders the use of CC for mapping out the potential energy surfaces (PESs) associated with bond cleavage, processes that are central to chemistry. Although the development of multireference CC methods (MRCC) has been an active area of research over the last 20 years,^{34–43} a completely general, computationally practical MRCC method remains elusive.

Compared to CC or MP_x , far less research has gone into the development of reduced scaling CI methods. Saebø and Pulay developed a local single reference CI method.⁴⁴ However, it was used only with closed-shell references, no chemically interesting phenomena were studied, and the method was never fully developed to the same extent that their local MP_x and CC methods were. Martinez and co-workers later implemented PS versions of full CI, doubles CI (DCI), and MRSDCI.^{45–47} Reynolds and co-workers later incorporated local correlation into these^{48,49} methods. However, these local PS CI methods could only be used with closed-shell single determinant reference wave functions. Consequently, these methods were only used to study a few simple torsions in hydrocarbons. Furthermore, there were a number of significant technical issues that severely constrained the long-term growth of these algorithms. Nevertheless, in terms of the accuracy and percentage of correlation energy recovered, the results of Refs. 45–49 were quite encouraging. Spurred on by these results, Walter and Carter recently reported a

weak-pairs MRSDCI algorithm that removed the most troublesome technical difficulties associated with the methods of Refs. 45–49. This method was used to map out the PES for cleavage of the double bond in *trans*-2-butene.⁵⁰ The resulting local CI PES yielded a D_e that was within 1 kcal/mol of the D_e predicted by the analytic method. However, this method included no means of reducing the overall scaling associated with the treatment of the virtual orbitals. Furthermore, the absolute central processing unit (CPU) times were quite poor, the algorithm still being in early stages of development.

The important advantage CI has over CC and MP_x is the ease with which it can be used with just about any style of SCF reference wave function. This makes the CI method quite general and applicable to a wide range of systems and phenomena. In light of this, it is perhaps surprising that CI has not received more attention in the area of reduced scaling *ab initio* methods. The principal drawback of CI methods is their lack of size extensivity, which raises questions about their usefulness in studies on larger systems. As this is the driving force behind the development of reduced scaling *ab initio* methods, this no doubt helps explain the lack of attention received by CI in this area. However, there are a number of different size-extensivity corrections available that can be applied both to SDCI and MRSDCI.^{51–56} In addition, a number of different size extensive and approximately size-extensive MRCI techniques have been proposed.^{57–63} Thus, it seems quite likely that reduced scaling CI methods have a promising future and there is a substantial need for more research in this area.

Here, we present a reduced scaling MRSDCI method that combines local correlation in the internal space (occupied orbitals, denoted as i, j, k , and l) and the PS method in the virtual space (unoccupied orbitals, denoted as a, b, c , and d) for the treatment of the $(ab|cd)$ and $(ij|ab)$ classes of integrals. Our MRSDCI program is based upon the symmetric group graphical approach (SGGA) formalism of Duch and Karwowski.^{64,65} This offers important advantages for the development of reduced scaling CI algorithms, the most important of which are the separation of the spin and spatial parts of the wave function and on-the-fly calculation of integral coupling coefficients. Unlike in most previous discussions of reduced scaling *ab initio* methods, we apply our MRSDCI method to study various bond strengths in a series of hydrocarbons. We show that results comparable to those obtained using a fully analytic method can be obtained at a significantly reduced cost.

THEORY

Local configuration interaction

In a recent publication,⁵⁰ a detailed description of our weak-pairs CI method was given. Here, we briefly reiterate the essential points. Since the Coulomb potential is short ranged, the motion of widely separated electrons is weakly correlated. This fact can be exploited to greatly reduce the amount of effort required in a CI calculation. First introduced by Saebø and Pulay, the weak-pairs approximation³ neglects or approximates the contribution of CSFs (configuration state functions; the combination of a spatial orbital configuration

pattern and a spin eigenfunction properly antisymmetrized) having simultaneous excitations out of widely separated orbitals. In the local MP x method of Saebø and Pulay the internal pairs were characterized as being negligible, weak, or strong. The strong pairs were treated without approximation and the negligible pairs were neglected. However, the weak pairs were treated approximately⁷ by either using a low level of theory or omitting groups of terms that should approximately cancel. In a similar vein, the local MP2 method of Hetzer *et al.* treats weak pairs approximately using multipole expansions.⁸ In contrast to these methods, we simply neglect all of the weak-pairs interactions. Since in large molecules there are a large number of weak pairs, it would certainly be desirable to use an approximate technique for treating the weak pairs as it would substantially improve the accuracy of the method. Furthermore, it would allow the treatment of systems in which dispersion plays a substantial role. Nevertheless, the total neglect of weak pairs amounts to a loss of only a few percent of the correlation energy and this error will largely cancel upon computing an energy difference for most systems.

Of course, in order to meaningfully discuss the notion of widely separated orbitals, the molecular orbitals (MOs) must be localized. Here, the doubly occupied orbitals are localized using the Boys algorithm.⁶⁶ The generalized valence bond (GVB) active orbitals are localized naturally as a by-product of the GVB orbital optimization.^{67,68}

In order to determine the weak-pairs orbitals, spheres are associated with each occupied MO. Pairs of MOs whose spheres do not overlap are then considered to be weak pairs. The spheres are determined in the following fashion. First, a Mulliken population analysis is done on each internal MO to determine which atoms contribute the most to that orbital. For each MO, the atoms in the molecule are then sorted according to their contribution to the MO in question. A group of atoms with dominant contributions to the MO are then selected from the top of the list such that the sum of their contributions to the MO exceed some threshold. Since the orbitals are normalized to unity, this number is usually something like 0.8. The center of the sphere is then determined by finding the contribution-weighted center of mass of the most heavily contributing atoms. The radius of the sphere for the MO is then set equal to a constant, α , multiplied by the maximum separation of any two atoms in the list of heavily contributing atoms. The parameter α and the occupation threshold are included in order to allow one to tune the strength of the local CI calculation. Larger values of the occupation threshold and α will lead to CI calculations that approach the nonlocal versions. The variation in the percentage of the correlation energy recovered as a function of the occupation threshold and α has been described previously.⁵⁰

The use of the weak-pairs approximation in large systems can lead to CI expansions requiring only 40% as much computational effort as the nonlocal CI expansions yet still recover more than 99% of the MRSDCI correlation energy when basis sets of double- ζ quality are used.⁵⁰ Furthermore, in the case of *trans*-2-butene, it yielded smooth potential energy surfaces for cleavage of the C–C double bond and predicted the D_e to within 0.9 kcal/mol of the analytic value.

It should be noted that the local CI method presented here does not make use of any form of localized virtuals to partially truncate the virtual space. However, for a given internal CSF having holes in a set of weak-pair orbitals, the entire set of excitations to virtuals is eliminated. Thus, the savings in the weak-pairs method is derived from reducing the number of times the virtual space must be treated.

Symmetric group approach to CI

Before discussing the details of the symmetric group approach to CI, a few introductory remarks regarding modern CI calculations should be made. Large-scale CI is based on direct methods in which small bits of the Hamiltonian are formed and then used directly in diagonalization algorithms that only require one to form the product of the Hamiltonian and a trial vector.^{69,70} The partial pieces of the Hamiltonian are constructed using the expression

$$\mathbf{H}_{\kappa\nu} = \sum_{ij} A_{ij}^{\kappa\nu}(i|\hat{h}|j) + \sum_{ijkl} A_{ijkl}^{\kappa\nu}(ij|kl), \quad (1)$$

where i, j, k, l are orbitals, κ and ν are CSFs, and the $A_{ij}^{\kappa\nu}$ and $A_{ijkl}^{\kappa\nu}$ are known as integral coupling coefficients.

In the development of reduced scaling methods, one would like to take advantage of as many symmetries as possible to reduce the calculational effort required. Unfortunately, the use of localized internal orbitals makes it impossible to take advantage of spatial point group symmetry in our local CI technique. For typical systems studied by CI, the remaining symmetry that can be exploited is spin symmetry. Although CI methods based upon Slater determinants are very efficient and relatively easy to implement,⁷¹ Slater determinants are not native spin eigenfunctions, except in the case of a closed-shell singlet and open shells coupled to maximum spin. For this reason, we have chosen to build our reduced scaling CI method within a CSF-based formalism rather than a Slater-determinant-based approach.

Configuration-state-function-based approaches to CI come in two general flavors: those based on the unitary group and those based on the symmetric group. Large-scale CI methods were first implemented in terms of the unitary group approach⁷² (UGA) and its graphical form (GUGA).^{73–75} In this approach, the spin and spatial parts of the wave functions are entangled. Integral coupling coefficients are calculated for a specific pair of CSFs using complex formulas in which a series of factors are concatenated for every occupied orbital. The process is not well suited to on-the-fly computation, and coupling coefficients are generally precomputed and stored on disk in a lengthy sequential access file. Clearly, this will introduce a strong dependence on low-speed disk storage. However, the principal disadvantage of this situation is that it inconveniently forces one to organize a CI calculation according to the sequence of coupling coefficients on the file. If one wants to carry out a specialized or truncated CI calculation, one must first sort or prune the coupling coefficient file. Finally, it should be noted that many GUGA programs are restricted in the number of

internal orbitals they can handle, making them inappropriate for use in the development of algorithms designed to treat larger systems.

Because, in the symmetric group approach (SGA), the spin and spatial parts of the wave function are separated, the SGA tends to be much simpler compared to the UGA. Consequently, many electron wave functions were first formulated in terms of the SGA.^{76–78} However, large-scale CI algorithms^{64,65} based on the SGA did not emerge until well after the development of algorithms based on the UGA. For the purpose of developing reduced scaling CI algorithms, the SGA offers numerous advantages over the UGA. In the SGA, matrices for irreducible representations of the symmetric group form a block of coupling coefficients for all pairs of spin couplings associated with a given pair of orbital occupations. The algorithms for computing these matrices^{79,80} are extremely rapid. In fact, they are rapid enough that the matrices can be recomputed on-the-fly as they are needed. In addition to removing the dependence on low-speed disk storage, this allows one to easily reorganize or truncate a CI calculation with minimal overhead. Of course, the previous statement is in general true regardless of the specific needs one may have in reorganizing a CI calculation. In local correlation techniques, however, the decision on whether or not to retain a particular CSF is made solely on the basis of the *spatial* aspects of the CSF. The separation of the spin and spatial parts of the CSF in the SGA makes the elimination of all CSFs corresponding to the different spin couplings associated with a given occupation pattern natural and trivial. In the UGA, such simplifications are not possible because of the entanglement of the spin and spatial parts of the wave function.

As was mentioned previously, coupling coefficient codes based on the GUGA approach are often limited in the number of internal orbitals they can handle. By definition, this is a quantity that is dependent on the size of the system. The algorithms for computing the irreducible representations of the symmetric group used in the SGA, on the other hand, are limited by the maximum number of open shells in the CI wave function. This is a quantity that depends on the spin multiplicity of the molecule and the choice of reference wave function, not the size of the system. For the overwhelming majority of CI calculations imaginable, this is a quantity that is easily handled by the algorithms for computing the symmetric group representation matrices. Thus, in addition to being simpler and well suited for use with local correlation techniques, the SGA appears to be intrinsically more appropriate for the treatment of larger systems.

Pseudospectral methods

In *ab initio* quantum chemistry, the PS method^{81–83} is used to decompose two-electron integrals into a product of intermediate quantities depending on no more than two orbital indices:

$$(ij|kl) \approx \sum_g Q_{ig} A_{kl}(r_g) R_{gj}, \quad (2)$$

$$A_{kl}(r_g) = \int \frac{\phi_k(r)\phi_l(r)}{|r-r_g|} dr. \quad (3)$$

Here, the variable g indexes a series of grid points and the matrix \mathbf{R} holds the values of the basis functions at those grid points. The matrix \mathbf{Q} is obtained as the solution to an equation of the form $\mathbf{Q}\tilde{\mathbf{R}} = \tilde{\mathbf{S}}$. The matrices $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{S}}$ are augmented versions of \mathbf{R} and the analytic overlap matrix containing an extra set of “dealiasing” functions. These dealiasing functions are introduced with the aim of minimizing the error arising from spurious components outside the basis set contained in the products $A_{kl}(r_g)\phi_j(r_g)$.⁸⁴ It has been shown that the accuracy of the PS approximation can be substantially enhanced through the use of an appropriate set of dealiasing functions.⁸⁵

The utility of the PS approximation becomes apparent when two-electron integrals must be contracted with arrays of coefficients. Such operations are extremely common in *ab initio* quantum chemistry. Indeed, they arise in SCF calculations, MP x , CC, and CI methods. In CI, the dominant computational effort comes from the contraction of the $(ij|ab)$ and $(ab|cd)$ integrals with the vector of CI expansion coefficients, which we shall denote as \mathbf{C} and refer to as the CI vector. Here, i and j are internal orbitals and a , b , c , and d are indexing virtual orbitals.

The equation for treating the $(ij|ab)$ is

$$(\mathbf{Z})_\kappa = \sum_{vijac} A_{ijac}^{\kappa\nu} (ij|ac)(\mathbf{C})_\nu + \sum_{vijac} A_{iajc}^{\kappa\nu} (ia|jc)(\mathbf{C})_\nu, \quad (4)$$

where \mathbf{Z} is a vector that contains the product of the Hamiltonian and the CI vector. As was noted by Siegbahn,⁸⁶ the coupling coefficients can be broken down into an internal part and an external part. The external part of the coupling coefficients is quite simple, usually amounting only to a sign change or a factor of the square root of two. As was first noted by Meyer,⁸⁷ this external coupling coefficient can be easily absorbed into the \mathbf{C} and \mathbf{Z} vectors.⁸⁸ Denoting the internal parts of the coupling coefficients as B_J and B_K and assuming the external parts are absorbed into \mathbf{C} and \mathbf{Z} , Eq. (4) can be rewritten as

$$\mathbf{Z}_{ab}^\lambda = \sum_{\mu c} B_J^{\lambda\mu} (ij|ac) \mathbf{C}_{cb}^\mu + \sum_{\mu c} B_K^{\lambda\mu} (ia|jc) \mathbf{C}_{cb}^\mu, \quad (5)$$

where λ and μ label internal CSFs and the variable \mathbf{C}_{ab}^μ is the element of the CI vector corresponding to the CSF with internal part μ and having external orbitals a and b occupied. Note here the implicit sums over allowed i, j values (within the sum over μ). This is the exact equation presented previously by Martinez and Carter, and the discussion of the scaling that follows is similar to theirs.⁴⁷ In terms of floating point operations, the calculation of each term contributing to the \mathbf{Z}_{ab}^λ will scale formally like XN_v^3 where X is the total number of interacting λ and μ and N_v is the number of virtual orbitals. The factor X is in many cases quite substantial, and can cause the treatment of the $(ij|ab)$ to be the most time-consuming portion of the diagonalization process. In fact, for systems with four or five heavy atoms treated using

standard double- ζ polarization (DZP) basis sets, the $(ij|ab)$ can consume as much as 60% of the CPU time depending on the complexity of the internal space.

The PS approximation for the two-electron integrals can be inserted straightforwardly into Eq. (5), leading to the expression

$$\mathbf{Z}_{ab}^{\lambda} = \sum_g Q_{ag} \sum_{\mu} A_{ij}(r_g) B J_{ij}^{\lambda\mu} \sum_c R_{gc} \mathbf{C}_{cb}^{\mu} + \sum_{ig} A_{ia}(r_g) \sum_{\mu} Q_{jg} B K_{ij}^{\lambda\mu} \sum_c R_{gc} \mathbf{C}_{cb}^{\mu}. \quad (6)$$

If the index μ can run from 1 to P and the total number of grid points is denoted as N_g , then the sums over c in both terms of Eq. (6) require $P N_g N_v^2$ floating point operations. The resulting quantity $(RC)_{gb}^{\mu}$ can be precomputed and stored. The sum over μ in both terms of Eq. (6) becomes more complicated, but can be much better understood by noting that fixing the values λ and μ also fixes the values of i and j . Thus, the result of the middle sum of the first term in Eq. (6) can be safely written as $(ABJRC)_{gb}^{\lambda}$ and the cost of computing it will scale as $X N_g N_v$. Because in the integral $(ia|jc)$ the indices i and j are not adjacent, the result of the middle sum of the second term of Eq. (6) must be written as $(QBKRC)_{igb}^{\lambda}$. Although the presence of the extra index i does not increase the cost of computing these terms (still scaling as $X N_g N_v$), they must still be binned according to i . The final sum over g in the first term of Eq. (6) will scale straightforwardly like $P N_g N_v^2$. The final sum of the second term of Eq. (6), however, scales slightly worse owing to the presence of the extra index i . Although the index i does not range over all the internal orbitals, it typically includes a nontrivial fraction of them. Depending on the number of internal orbitals and the complexity of the internal space, the effect of the index i can be substantial. Denoting the effect of this extra index as β , the scaling of the sum over i and g in the second term of Eq. (6) is $\beta P N_g N_v^2$.

The dominant terms in the scaling of the PS treatment of the $(ij|ab)$ finally amount to $(2 + \beta) P N_g N_v^2$. Here, the 2 arises from the calculation of the $(RC)_{gb}^{\mu}$ [which should be done only once and used in both the first and second term of Eq. (6)] and the sum over g in the first term of Eq. (6). The β arises from the sum over i, g in the second term of Eq. (6). The final scaling of the analytic treatment of the $(ij|ab)$ comes to $2 X N_v^3$, as readily seen from Eq. (5). Thus, the scaling advantage from the PS treatment of the $(ij|ab)$ is $2 X N_v / [(2 + \beta) P N_g]$. The magnitude of this scaling advantage varies greatly depending on the system under consideration. For small systems with only a few atoms, N_g is typically at least an order of magnitude larger than N_v . This typically overwhelms the effect of X/P , leading to a PS *disadvantage* for such systems. For small systems, one will only save in the treatment of the $(ij|ab)$ if extremely large basis sets are used to partially cancel out the effect of N_g and if the internal space is sufficiently complicated so that X is large. As the size of the system grows, N_g is still roughly an order of magnitude larger than N_v , but the ratio N_g/N_v begins to decrease. More importantly, the value of X increases rapidly, especially in the case of multiconfigurational reference wave

functions. For example, the values of X for ethane and butane each treated using three reference configurations are 15 336 and 120 528, respectively. Consequently, one starts to see marked savings in the treatment of the $(ij|ab)$ as the size of the system grows.

The analysis of the $(ab|cd)$ is far more straightforward compared to the $(ij|ab)$. The analytic expression for the $(ab|cd)$ is

$$\mathbf{Z}_{ab}^{\lambda} = \sum_{cd} (ac|bd) \mathbf{C}_{cd}^{\lambda}. \quad (7)$$

Here, we use the same notation and make the same assumptions regarding the internal parts of the coupling coefficients as in Eq. (6). This calculation will scale as $P N_v^4$. The PS version of Eq. (7) is given by

$$\mathbf{Z}_{ab}^{\lambda} = \sum_g Q_{ag} \sum_d A_{bd}(r_g) \sum_c R_{gc} \mathbf{C}_{cd}^{\lambda} \quad (8)$$

and will scale as $3 P N_g N_v^2$. The PS scaling advantage in the treatment of the $(ab|cd)$ is therefore $N_v^2/3 N_g$. As we have noted previously, N_g is roughly an order of magnitude larger than N_v . Therefore, one can easily see that in terms of floating point operations, one starts to save in the PS treatment of the $(ab|cd)$ even for fairly small systems. It should be noted that the savings in the PS $(ab|cd)$ are derived from a different source than in the case of the $(ij|ab)$. For the $(ab|cd)$, the savings are derived from decoupling the c and d indices in the two-electron integral. In the $(ij|ab)$, however, the savings come from decoupling the λ and μ indices in the coupling coefficients. The former is far more potent.

Although counting floating point operations can provide insight into the cost of a calculation, the main bottleneck in modern computers is memory access. For smaller systems where there are relatively few floating point operations, this fact makes it difficult to predict the amount of time a calculation will consume as memory accesses may be the bottleneck. However, as the number of virtual orbitals increases, the ratio of floating point operations to memory accesses will decrease in the treatment of both the $(ij|ab)$ and $(ab|cd)$. Since the floating point operations are the bottleneck, the count of floating point operations will then become an accurate predictor of the cost of a calculation. Thus, the formulas given above are expected to be most accurate when N_v is large.

CALCULATIONAL DETAILS

All of the SCF reference wave functions used in this study were obtained using the Jaguar version 4.1 quantum chemistry package.⁷ For single reference cases, the wave function used was either Hartree–Fock (HF) or restricted open-shell Hartree–Fock (ROHF), of course. For multireference cases, the wave function used was either generalized valence bond perfect-pairing⁸⁹ (GVB-PP) or generalized valence bond with restricted configuration interaction^{90,91} (GVB-RCI). Briefly, the GVB-PP wave function uses two configurations for a particular pair of electrons. One configuration places the pair of electrons in an orbital that is bonding in nature while the other places the pair in an orbital that is

correspondingly antibonding. If the pair of electrons in question are taking part in bonding, then such a wave function will provide a proper zeroth-order description of the cleavage of the bond. The GVB-RCI wave function differs from the GVB-PP wave function in that it also includes a configuration placing one electron in the bonding orbital and one electron in the antibonding orbital. It has been shown that such wave functions are competitive with complete active space SCF (CASSCF) wave functions.⁹² In general, the GVB-PP and GVB-RCI wave functions are denoted as GVB ($x/2x$) and GVB-RCI ($x/2x$), where x is the number of pairs of electrons and $2x$ is the number of orbitals used to describe those pairs of electrons. We note in passing that although we only use GVB style reference functions here, there is nothing intrinsic to the method presented here that restricts it to use with such references. All molecular geometries were optimized at the GVB-PP level of theory.

As was noted previously by Martinez and Carter,⁴⁷ the use of the pseudospectral approximation leads to a slightly asymmetric Hamiltonian. For example, in the case of the $(ab|cd)$ [Eq. (7)], this can be easily understood by noting that Eq. (2) is not necessarily symmetric with respect to interchange of the indices i and j . To deal with these slight asymmetries, Martinez and Carter used a modified version of the Davidson matrix diagonalization algorithm⁹³ proposed by Rettrup.⁹⁴ Here, we do the same. The Hamiltonian diagonalization algorithm was considered converged when the energy difference between consecutive iterations was less than 10^{-9} hartrees.

All analytic one- and two-electron integrals were computed using the HERMIT program of Helgaker *et al.*⁹⁵ The pseudospectral units \mathbf{Q} and \mathbf{R} were computed using Jaguar while the $A_{ki}(r_g)$ were computed using the McMurchie–Davidson algorithm.⁹⁶ The grids used were the “coarse” grids of Jaguar.

The bond dissociation energies (BDEs) reported here are all 298 K binding enthalpies. They are determined by first computing the 0 K D_e as the difference between the CI energy for the fully bound molecule and the CI energy for a “supermolecule” in which the fragments are separated by roughly 100 bohrs. To convert our 0 K D_e s to 298 K binding enthalpies, zero-point vibrational energies and enthalpy corrections are required. Zero-point energies were obtained from HF or ROHF analytic vibrational frequencies computed by Jaguar. No scaling corrections are applied to the frequencies. The enthalpy corrections are taken directly from Jaguar and are also computed using the analytic vibrational frequencies.

There are a number of sources of error present in our calculations that we do not account for. The basis set superposition error (BSSE) will have the effect of slightly increasing our BDEs. As mentioned previously, CI is not rigorously size extensive. Although there are size-extensivity corrections available for MRSDCI,⁵⁶ accurate corrections are not trivial to implement and we have not yet done so. Nevertheless, the error due to size extensivity is typically a few percent of the correlation energy and will largely cancel upon determining a D_e . Perhaps the biggest source of error in our calculations will be from basis set incompleteness. In gen-

TABLE I. Percentage of SDCI correlation energy in n -pentane recovered by the weak-pairs method as a function of the basis set. In the weak-pairs calculations, the occupation threshold was 0.8 and the parameter α was 1.7.

Basis set	E_{corr} (hartrees)	% E_{corr} weak pairs
6-31G ^a	0.4502353	99.8
6-31G** ^a	0.6441177	99.9
6-31G**+ + ^a	0.6776429	99.9
6-311G ^b	0.5459467	97.7
6-311G** ^b	0.7706377	98.3
6-311G**+ + ^b	0.7719486	98.4
cc-pVDZ ^c	0.6698297	99.9
aug-cc-pVDZ ^c	0.6873657	99.8
cc-pVTZ(-f) ^c	0.7798551	99.0

^aReference 100.

^bReference 101.

^cReference 102.

eral, this will have the consequence of lowering our BDEs because of an underestimation of correlation effects in the fully bound molecule. While the errors mentioned here are significant, the goal of this study is not necessarily to produce the most accurate theoretical thermochemistry. Rather, the main goal is to reproduce the analytical results using the local and pseudospectral approximations. However, comparison with experimental thermochemistry is still useful in gauging the performance of the method.

Experimental thermochemistry was obtained from either the tables of BDEs in the CRC Handbook of Chemistry and Physics⁹⁷ or estimated from the heats of formation contained in the CRC Handbook of Chemistry and Physics.⁹⁷ All calculations were performed on Compaq DEC ES40 machines.

RESULTS AND DISCUSSION

Basis sets in local correlation

Table I shows the variation in the percentage of the correlation energy recovered by the weak-pairs SDCI for n -pentane as a function of the basis set used. The reference function used was closed-shell HF. The local CI calculations used an occupation threshold of 0.8 and a value for α of 1.7 to define the sphere radius. This led to local CI calculations having 66.7% as many CSFs as the nonlocal CI calculations for all of the basis sets used. It should be noted that the percentage of the CSFs used in a weak-pairs CI is not very sensitive to the choice of the basis. Rather, it depends mainly on the choice of α and the occupation threshold. For example, if the occupation threshold is chosen to be 0.8 and α is taken to be 1.0 then this leads to weak-pairs CI wave functions for n -pentane having 46.3% as many CSFs as the nonlocal CI expansions for all of the basis sets listed in Table I. Similarly, for an occupation threshold of 0.65 and an α value of 1.35 the weak-pairs CI expansions have 64.0% as many CSFs as the nonlocal CI expansions for all basis sets listed in Table I. These results strongly suggest that the determination of the weak pairs is largely independent of the basis set.

In all cases, double- ζ basis sets performed very well, recovering at least 99.8% of the correlation energy. Compared to the double- ζ basis sets, the triple- ζ basis sets recov-

TABLE II. Comparison of absolute energies in hartrees for medium-sized molecules. The term "local" refers to the weak-pairs method. In the weak-pairs calculations, the occupation threshold was 0.8 and the parameter α was 1.7.

System and reference wave function	Total energies (hartrees)				
	GVB	Analytic CI	PS CI	Local CI	PS/local CI
C ₂ H ₅ -C ₂ H ₅ RCI(1/2)	-157.355 552	-157.986 976	-157.987 596	-157.974 215	-157.974 885
CH ₃ CH ₂ CH ₂ -CH ₃ RCI(1/2)	-157.355 391	-157.986 882	-157.987 601	-157.974 096	-157.974 853
CH ₃ -OC(CH ₃)CH ₂ RCI(1/2)	-231.041 295	-231.800 580	-231.800 561	-231.785 748	-231.785 772
H-CH ₂ CHCHCHCH ₂ RCI(1/2)	-194.021 147	-194.725 617	-194.725 454	-194.699 409	-194.699 592
H-CH ₂ CHCHCHCH ₂ ^a RCI(1/2)	-193.985 419	-194.611 466	-194.611 160	-194.605 716	-194.605 430
CH ₃ NH-NHCH ₃ RCI(1/2)	-189.309 732	-189.971 730	-189.971 313	-189.956 850	-189.956 434
C ₂ H ₅ -CH ₂ NH ₂ RCI(1/2)	-173.349 773	-173.998 960	-173.999 103	-173.985 468	-173.985 623

^aDone with 6-31G** basis; all others with 6-311G** basis.

ered roughly 0.1 hartrees more correlation energy. This suggests that triple- ζ basis sets may be more preferable for use in calculations that require chemical accuracy. Unfortunately, the performance of the local CI deteriorates substantially in triple- ζ basis sets. For example, the 6-31G basis set recovers 99.8% of the correlation energy while the 6-311G basis set recovers only 97.7% of the correlation energy. Although this error might at first appear to be inconsequential, in fact it can seriously affect the accuracy of a calculation. The essential difference between the double- ζ and the triple- ζ basis sets is the inclusion of an additional uncontracted long-ranged function to describe the valence electrons in the triple- ζ basis sets. This additional function is apparently allowing for a better description of long-ranged correlation effects that one neglects in local CI calculations. The obvious solution is to extend the sizes of the spheres used to determine the weak-pairs orbitals. Unfortunately, extending them much further eliminates almost all of the savings from the local CI.

The inclusion of polarization functions in the 6-31G and 6-311G basis sets lead to the recovery of an additional 0.194 and 0.225 hartrees in the *n*-pentane correlation energy, respectively. In the case of the 6-311G** basis set, the fraction of the correlation energy recovered was 98.3%. This represents an improvement of 0.6% over the 6-311 basis. This effect can be traced to the exponents used in the *d* polarization functions. For the carbon atom, the 6-311G** basis set includes a set of *d* polarization functions having an exponent of 0.6266. The most long-ranged uncontracted *p* valence function in the 6-311G** basis has an exponent of 0.146. Thus, the *d* polarization functions are much shorter-ranged than the uncontracted valence basis functions. The net effect is that the correlation energy is increased without further impeding the performance of the local CI.

In the case of correlation-consistent basis sets, the correlation-consistent polarized valence double- ζ (cc-pVDZ) basis set recovered 99.9% of the correlation energy while the

correlation-consistent polarized valence triple- ζ (cc-pVTZ) basis set recovered only 99.0% of the correlation energy. Similar to results obtained using Pople-style basis sets, the weak-pairs method appears to perform worse when more extended basis sets are used. However, the observed degradation in performance is not as severe as in the case of the Pople basis sets. This is possibly due to the fact that the correlation-consistent basis sets are optimized for use in correlated methods.

Adding diffuse functions to the basis set only incrementally improves the amount of correlation energy recovered in *n*-pentane. Compared to the 6-311G** basis, the 6-311G**++ basis only recovers an additional 0.001 hartrees of correlation energy. Compared to cc-pVDZ, the augmented cc-pVDZ (aug-cc-pVDZ) basis set recovered an additional 0.018 hartrees of correlation energy. Accordingly, the percentage of the correlation energy recovered by the local CI was largely unaffected. It should be pointed out that this very well may not be the case for systems where diffuse functions contribute substantially to the correlation energy.

Although the 6-311G** basis set appears to recover the smallest fraction of the correlation energy, we have elected to use this basis set in our subsequent binding energy calculations for two reasons. First, the use of any quantum chemistry method to compute binding energies is largely predicated on cancellation of errors. For approximate methods, this is even more true. Clearly, using the 6-311G** basis set will apply the most stringent test of the ability of the method presented here to adequately cancel errors. Second, the magnitude of the correlation energy recovered by the 6-311G** basis set is much greater than that for the double- ζ basis sets. For the analytic calculations, this fact will lead to greater overall accuracy and facilitate comparison with experimental thermochemistry.

TABLE III. Analytic and pseudospectral MRSDCI 298-K bond dissociation energies^a (BDEs) for small molecules.

Bond	Reference	Method	BDE (kcal/mol)	
			This work	Expt. ^b
CH ₃ —CH ₃	RCI (1/2)	Analytic	88.7	87.9±0.5
		PS	89.1	
C ₂ H ₅ —H	RCI (1/2)	Analytic	98.7	101.0±0.4
		PS	98.9	
NH ₂ —H	RCI (4/8)	Analytic	101.1	108.1±0.3
		PS	101.2	
HO—H	RCI (1/2)	Analytic	109.9	118.9±1.0
		PS	109.7	
HO—H	RCI (4/8)	Analytic	110.5	118.9±1.0
		PS	110.5	
CH ₃ —NH ₂	RCI (1/2)	Analytic	79.3	85.4±0.4
		PS	79.7	
CH ₃ —NH ₂	Bond: RCI(1/2) Lone pair: RCI(1/2)	Analytic	80.7	85.4±0.5
		PS	80.9	
CH ₃ —OH	RCI (1/2)	Analytic	86.2	92.4±0.1
		PS	86.5	
CH ₃ —OH	Bond: RCI (1/2) Lone pairs: GVB (2/4)	Analytic	86.4	92.4±0.1
		PS	86.6	
CH ₂ =CH ₂	RCI (2/4)	Analytic	161.9	174±1.4
		PS	162.0	
PS average error			0.2	

^aDone with 6-311G** basis.^bReference 97.

Absolute errors in total energies

Shown in Table II are the absolute energies computed at various levels of theory for a series of hydrocarbons in their equilibrium geometries. Here, one can see that the PS CI energies differ from the analytic CI energies by only a few tenths of a millihartree. This result is in line with those previously obtained by Martinez *et al.*^{45–47} For the purpose of computing energy differences, such an error is almost entirely negligible in most cases. Unfortunately, the errors incurred by the weak-pairs approximation are far more substantial. This increase in the error is of course due to the complete neglect of the weak-pairs interactions. However, the percentage of the correlation energy recovered by the local correlation method used here is relatively constant. Excluding the case of H-CH₂CHCHCH₂, which is discussed in detail below, the local correlation method recovers on average (97.9±0.08)% of the correlation energy. This stability in the fraction of the correlation recovered is an important reason why the method presented here can produce fairly accurate energy differences. In this regard, the method presented here relies considerably on cancellation of errors, as is typical in quantum chemistry methods.

Comparison of analytic and approximate BDEs and CPU times

Analytic and PS BDEs for a series of small hydrocarbons are listed in Table III. Clearly, the PS method is performing extremely well. Compared to the analytic BDEs, the

PS BDEs are no more than 0.5 kcal/mol different. This result is a direct consequence of the fact that the raw MRSDCI total energies used to compute the BDEs (not shown) differ from the analytic MRSDCI total energies by only a few millihartrees as mentioned previously. Because the systems in Table III are so small, the local correlation technique leads to CI expansions that are essentially as large as the nonlocal CI expansions. Therefore, we show no BDEs obtained using the weak-pairs method. Additionally, most of the systems in Table III are small enough such that the CPU times are not particularly meaningful and we therefore do not bother to show those either.

Table IV displays BDEs in a series of hydrocarbons having 4–5 main group atoms. As in the case of the small systems, the PS BDEs are no greater than 0.4 kcal/mol different from the analytic BDEs and this is again a consequence of the similarity of the analytic and PS MRSDCI total energies. Although the weak-pairs correlation method appears to perform rather well, the errors introduced are on average greater than those incurred by the PS method. For C₂H₅-CH₂NH₂, CH₃-OC(CH₃)CH₂, and CH₃NH-NHCH₃ errors of 0.5, 0.7, and 1.8 kcal/mol are introduced by the weak-pairs method, respectively. Although these errors are significant, they are still acceptable. In the case of the H-CH₂CHCHCH₂ BDE calculated using the 6-311G** basis set, on the other hand, an error of 8.4 kcal/mol is introduced by the local correlation method. This is a rather large error and can be attributed to the spheres associated with the GVB natural orbitals being too small to force the inclusion of CSFs needed to adequately describe the correlation of the C–H bond. When the sizes of these spheres are doubled, the weak-pairs BDE is determined to be 85.4 kcal/mol, in excellent agreement with nonlocal results. When less extended basis sets are used, it is expected that the correlation effects captured by a correlation method will be more localized. It is therefore not surprising that when a 6-31G** basis set is used, the error introduced by the weak-pairs method with unmodified spheres for the GVB orbitals is only 2 kcal/mol compared to the nonlocal method. For all other systems studied, the local MRSDCI error was only 0.1 kcal/mol. The combined local PS method gives results very similar to the local method. This is expected, as the PS method gives results very similar to the analytic method. The largest difference on going from the local MRSDCI method to the local PS MRSDCI is observed for CH₃-OC(CH₃)CH₂ and is 0.8 kcal/mol. For all other molecules, this difference is less than 0.4 kcal/mol. Thus, it appears that the limiting factor in the accuracy of the approximate CI methods is the effectiveness of the weak-pairs approximation. However, it is likely that the error only arises when more extended basis sets are used and even in these cases the error seems to cancel to a large extent.

Also shown in Table IV are the total CPU times per iteration and the CPU times required to treat the contribution of various integrals to the **Z** vector. It should be noted that the total CPU time is not simply the sum of the times for the (*ab|cd*), (*ij|ab*), (*ij|kl*), and (*ai|jk*) integrals. Rather, the total CPU time is equal to the total amount of time required to compute the **Z** vector taking into account the contributions

TABLE IV. Analytic and pseudospectral CPU times and 298-K BDEs for medium-sized molecules. The term *local* refers to the weak-pairs method. In the weak-pairs calculations, the occupation threshold was 0.8 and the parameter α was 1.7.

System and reference wave function	Method	CPU times Per iteration (s)				Total	BDE (kcal/mol)	
		$(ij ab)$	$(ab cd)$	$(ij kl)$	$(ai jk)$		This work	Expt. ^a
C ₂ H ₅ -C ₂ H ₅ RCI(1/2)	Analytic	2248	5613	466	726	10975	88.2	87.8±0.9
	PS	911	890	466	726	4965	88.6	
	Local	1300	4640	307	478	8295	88.1	
	PS/local	642	694	307	478	3654	88.5	
CH ₃ CH ₂ CH ₂ -CH ₃ RCI(1/2)	Analytic	2258	6035	483	730	11365	87.2	89.1±0.8
	PS	1030	1016	483	730	5008	87.6	
	Local	1363	4459	324	442	7945	87.1	
	PS/local	811	667	324	442	3639	87.5	
CH ₃ -OC(CH ₃)CH ₂ RCI(1/2)	Analytic	6627	9980	1041	1266	22057	81.6	66.3
	PS	1483	1406	1041	1266	7958	81.7	
	Local	3892	7757	595	730	15302	82.3	
	PS/local	1065	908	595	730	5339	81.5	
H-CH ₂ CHCHCHCH ₂ RCI(1/2)	Analytic	3557	8018	855	1083	16165	85.5	82.9±3.1
	PS	1357	1361	855	1083	7329	85.3	
	Local	1704	7185	487	568	10683	77.1	
	PS/local	874	1056	487	568	4657	76.9	
H-CH ₂ CHCHCHCH ₂ ^b RCI(1/2)	Analytic	1904	3932	520	661	8460	86.3	82.9±3.1
	PS	958	870	520	661	4332	86.3	
	Local	882	2438	291	344	4825	84.3	
	PS/local	585	621	291	344	2692	84.3	
CH ₃ NH-NHCH ₃ RCI(1/2)	Analytic	1653	4722	391	540	8653	57.1	62.7±3.8
	PS	666	649	391	540	3512	57.3	
	Local	824	3021	226	320	7188	55.3	
	PS/local	431	503	226	320	3654	55.5	
C ₂ H ₅ -CH ₂ NH ₂ RCI(1/2)	Analytic	1909	4455	433	657	9083	80.3	79.3±1.9
	PS	784	768	433	657	4074	80.4	
	Local	1086	3561	275	415	6654	79.8	
	PS/local	522	577	275	415	2444	79.9	

^aReference 97.^bDone with 6-31G** basis; all others with 6-311G** basis.

of all possible classes of integrals. As mentioned previously, the $(ij|ab)$ and $(ab|cd)$ integrals are the usual bottlenecks in CI calculations and are treated using the PS approximation. For systems with four heavy atoms, the PS approximation reduces the cost of treating the $(ij|ab)$ by slightly more than a factor of 2. As the size of the system increases, the PS approximation should become more effective for the $(ij|ab)$. Accordingly, for CH₃-OC(CH₃)CH₂ the PS approximation reduces the cost of treating the $(ij|ab)$ by a factor of 4.5. For the $(ab|cd)$, the PS approximation is even more effective. For the smallest systems listed in Table IV, the PS approximation reduces the cost of treating the $(ab|cd)$ by roughly a factor of 6. For the larger systems in Table IV, this increases to a factor of 7.

As the number of references increases, the treatment of the $(ai|jk)$ is also expected to become a bottleneck. Since these integrals contain three internal indices, the cost of treating them should be greatly reduced by the weak-pairs approximation. Although the $(ij|kl)$ are never bottlenecks, the treatment of these integrals should also be greatly reduced by the weak-pairs approximation. This can be easily understood for both classes of integrals by noting that these integrals

link configurations having different internal CSFs. Thus, the cost of treating these integrals should depend approximately quadratically on the number of internal CSFs. The $(ab|cd)$ integrals, on the other hand, only link CSFs having the same internal part and the cost of treating them can therefore only depend linearly on the number of internal CSFs. If the weak-pairs approximation reduces the number of internal CSFs by a factor of 2, for example, then the cost of treating the $(ij|kl)$ and $(ai|jk)$ will be reduced roughly fourfold while the cost of the $(ab|cd)$ will be reduced only twofold. Of course, the situation is complicated, since the $(ai|jk)$ link CSFs differing in their number of electrons in the internal space by one and the $(ij|kl)$ link CSFs having the same number of electrons in the internal space. Nevertheless, the argument serves to explain why the weak-pairs approximation is more effective for these integrals.

For the small systems listed in Table IV, the weak-pairs approximation reduces the cost of treating the $(ij|kl)$ and $(ai|jk)$ by less than a factor of 2. However, the savings grow slightly as the size of the system increases. Given the extended orbital spheres used here to accommodate the 6-311G** basis set, these results are consistent with those

TABLE V. Weak-pairs pseudospectral MRSDCI CPU times and 298 K BDEs for large molecules. The term *local* refers to the weak-pairs method. In the weak-pairs calculations, the occupation threshold was 0.8 and the parameter α was 1.7. The theoretical BDEs shown below were obtained using the combined weak-pairs pseudospectral method.

System and reference wave function	Method	CPU times per iteration (s)					BDE (kcal/mol)	
		$(ij ab)$	$(ab cd)$	$(ij kl)$	$(ai jk)$	Total	This work	Expt. ^a
CH ₃ CH ₂ CH ₂ -CH ₂ CH ₂ CH ₃ RCI(1/2)	Analytic	40 432	84 564	5 302	5 477	148 925	85.7	88.0±1.0
	PS/local	3 845	3 810	2 525	2 500	20 532		
Ph-H ^b RCI(1/2)	Analytic	5 396	14 203	1 389	1 604	26 408	99.1	113.0±0.7
	PS/local	1 238	1 258	978	1 118	7 322		
Ph-H ^{b,c} RCI(1/2)	Analytic	2 495	4 303	860	957	10 428	108.3	113.0±0.7
	PS/local	978	918	610	688	4 609		
C ₂ H ₅ -NNC ₂ H ₅ RCI(1/2)	Analytic	12 504	32 340	3 226	3 463	59 358	52.7	50.0
	PS/local	1 799	1 503	1 503	1 605	11 643		
Ph-Ph ^{b,c} RCI(1/2)	Analytic ^d	153 807	235 473	33 910	40 523	463 713	113.3	118.1±1.4
	PS/local	15 883	17 542	27 345	16 044	111 559		

^aReference 97.

^bPh refers to the aromatic phenyl group.

^cDone with 6-31G** basis; all others with 6-311G** basis.

^dEstimated times, as described in text.

reported previously by Walter and Carter.⁵⁰ The cost of treating the $(ij|ab)$ and $(ab|cd)$ are also reduced by the weak-pairs approximation. For the $(ij|ab)$, the savings are similar to those seen for the $(ai|jk)$ and $(ij|kl)$. This is not surprising as the $(ij|ab)$ also link CSFs having different internal parts. As expected, the savings from the weak-pairs approximation in treating the $(ab|cd)$ are much smaller.

The savings from the combined weak-pairs PS method is expected to be equal to the product of the savings from the weak-pairs approximation and the savings from the PS approximation. This fact was previously noted by Reynolds *et al.*⁴⁸ For example, in the case of the $(ab|cd)$ integrals for C₂H₅-C₂H₅, the PS and weak-pairs approximations reduce the effort required by a factor of 6.3 and 1.2, respectively. The product of these savings, 7.6, is close to the actual observed factor of 8.1. Similar agreement is observed for the other systems listed in Table IV.

Table V lists the analytic and weak-pairs PS CPU times for a series of large molecules having 6 or more heavy atoms. In all cases except C₆H₆, the analytic calculations were too costly to carry to completion as these CI calculations typically involved over 30 000 000 CSFs. In the case of biphenyl, it was too costly to even carry out one iteration of the CI calculation and the analytic times are estimated as follows. As discussed previously, the cost of treating the $(ab|cd)$ integrals should be proportional to the product of P and N_v^4 . The proportionality constant was determined from an average of results for which analytic CPU times were available. Similar procedures were used for the $(ij|ab)$, $(ij|kl)$, and $(ai|jk)$ integrals. The treatment of the $(ij|ab)$ integrals should be proportional to the product of X and N_v^3 . For the $(ij|kl)$ and $(ai|jk)$ integrals, the number of pairs of interacting doubly excited internal CSFs was taken to be proportional to the number of internal orbitals, N_i , raised to the power of the number of internal indices in the integral. Thus, the $(ij|kl)$ calculation was taken to scale as $N_i^4 N_v^2$ and

the $(ai|jk)$ were taken to scale as $N_i^3 N_v^2$. The total time per iteration reported for biphenyl is simply the sum of the times for the $(ij|ab)$, $(ab|cd)$, $(ij|kl)$, and $(ai|jk)$ integrals. Of course, this neglects the contribution of some interactions, so this estimated time will be slightly lower than the actual CPU time required. Nevertheless, one can verify from the other entries in Table V that this will indeed be a very good estimation of the total CPU time required for one iteration.

The CPU times in Table V are very encouraging. The local PS method is capable of reducing the cost of the $(ij|ab)$ by over a factor of 10 and the $(ab|cd)$ by over a factor of 20. For the $(ij|kl)$ and $(ai|jk)$ integrals, the local

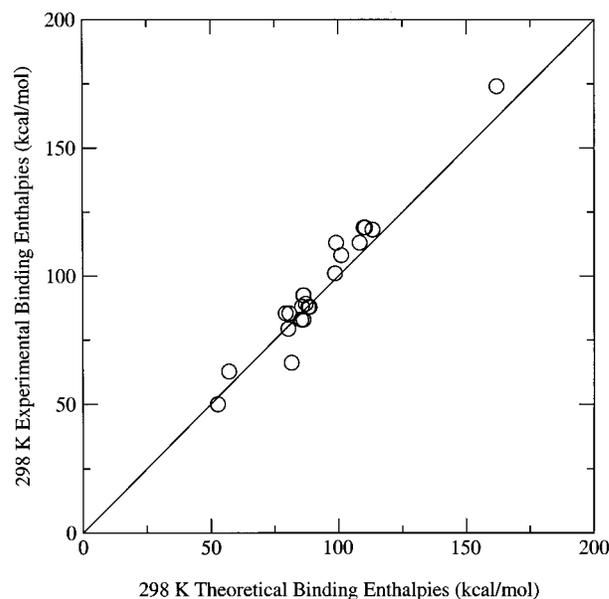


FIG. 1. Experimental BDEs versus theoretical BDEs. The theoretical BDEs are the analytic BDEs of Tables III and IV and the local/PS BDEs of Table V.

PS method reduces the cost by more than a factor of 2. Here, we see that the savings from both the weak-pairs and PS approximations grow rapidly with the size of the system. For the biphenyl molecule, the calculation was carried out using a 6-31G** basis. Obviously, for this basis the ratio of the size of the virtual space to the size of the internal space is not as great as it is for the 6-311G** basis. Consequently, the savings seen for biphenyl are somewhat less than they are for the other molecules in Table V.

The case of C₆H₅-H (Ph-H) is clearly another example of the 6-311G** basis set failing in the local correlation method. The use of a more localized 6-31G** basis set improves the results substantially.

Comparison of theoretical and experimental BDEs

Tables III–V list experimental BDEs alongside the theoretical values computed here. Figure 1 shows the experimental BDEs plotted versus the theoretical BDEs. The correlation shown in Fig. 1 is in general quite good. However, the points do tend to lie above the diagonal, indicating that the theoretical BDEs are lower than the experimental BDEs. Excluding the outlier cases of CH₃-OC(CH₃)CH₂ and C-H BDEs obtained with a 6-311G** basis set, the BDEs shown in Fig. 1 are lower than the theoretical BDEs by an average of 3.5 kcal/mol. As pointed out previously, this discrepancy is very likely due to the inability of the basis sets used to fully describe the electron correlation in the bound molecules. The errors are more severe for cases involving lone pairs and multiple bonds where electron correlation is more important.

As mentioned briefly above, Fig. 1 includes one troublesome outlier for the CH₃-OC(CH₃)CH₂ BDE. The experimental value for this BDE is reported to be 66.1 kcal/mol. This BDE is 26.3 kcal/mol smaller than the well known CH₃-OH BDE of 92.4 kcal/mol. Given the similarity of the bonds, one might expect these BDEs to be much more similar. Here, we report a CH₃-OC(CH₃)CH₂ BDE of 81.6 kcal/mol. Based on our other results, it is very likely that this BDE is 4–6 kcal/mol lower than the actual BDE. Thus, we estimate the actual value of the CH₃-OC(CH₃)CH₂ BDE to be close to 86 kcal/mol, which is much more in line with the CH₃-OH BDE. This gives us additional confidence in our theoretical BDEs and casts some doubt upon the experimental CH₃-OC(CH₃)CH₂ BDE.

CONCLUSIONS AND FUTURE OUTLOOK

We have presented here a reduced scaling MRSDCI method that simultaneously takes advantage of the pseudospectral and weak-pairs approximation. Unlike most previous studies on reduced scaling *ab initio* methods, we have applied our method to study homolytic bond cleavage in a series of hydrocarbons. These are applications for which multireference methods are well suited, but much of the previous work in reduced scaling *ab initio* quantum chemistry has been based on methods in which the extension to the multireference case is very difficult both theoretically and computationally. The results of this study are very encouraging. The BDEs obtained from the approximate CI methods

are in reasonable agreement with both the analytic BDEs and the experimental BDEs. Additionally, the computational effort required in the approximate CI methods is significantly less than in the analytic CI and the savings grow with the size of the system.

There are a number of improvements we intend to make to the method presented here. We have only treated the $(ij|ab)$ and $(ab|cd)$ integrals pseudospectrally. Another candidate for PS treatment is the $(ia|bc)$ integrals. These integrals can only link configurations having two virtual orbitals occupied to those having one virtual orbital occupied. As a result, they never consume much CPU time compared to the $(ij|ab)$ and $(ab|cd)$. The previous work of Martinez and Carter⁴⁷ implemented PS versions of the $(ia|bc)$ primarily as a means of reducing disk storage requirements. Of course, conserving disk space is no longer a big concern. However, treating the $(ia|bc)$ using the PS approximation would still reduce the CPU time requirement somewhat. One could also consider the possibility of treating integrals involving three or four internal indices pseudospectrally. However, the use of the PS approximation for such integrals could drastically effect the MRSDCI energies and one should proceed along these lines with caution.

The results given here strongly suggest that the standard quantum chemistry basis sets may not be the most appropriate for use in local correlation methods. The problem arises from the use of long-ranged uncontracted valence functions, and reduces the fraction of correlation energy recovered by the local correlation methods by a few percent. This result opens up the possibility of trying to develop basis sets which strike a better balance between recovering correlation energy and limiting the weak-pairs method.

Although, the programming for the most expensive parts of the CI algorithm has been optimized, further optimization remains to be done. This is especially true for the integrals that individually do not consume much CPU time but taken as a whole are significant. Optimization of these parts of the code is progressing.

We feel the most appealing possibility for extending the method presented here is to apply the local correlation approximation to the virtual space. As suggested by Saebø and Pulay,^{2,98} this can be done by excluding configurations having excitations to virtuals localized far from the “hole” orbitals. For closed-shell systems using single reference wave functions, this has been previously implemented by Reynolds and Carter.⁴⁹ Unfortunately the extension is rather difficult, as it requires the use of a nonorthogonal, localized set of virtual orbitals. Furthermore, combining truncation of the virtual space with the pseudospectral approximation appears quite challenging since the virtual space becomes a function of the internal space. Nevertheless, methods based on truncation of the virtual orbitals are well suited to integral direct approaches, which can allow the treatment of systems much larger than those discussed here.²² Thus, truncation of the virtual orbitals is an active subject of our ongoing research.⁹⁹

ACKNOWLEDGMENTS

D.W. gratefully acknowledges Wlodzislaw Duch for providing a copy of his SGGA-CI program for benchmarking

purposes and Trygve Helgaker for providing his HERMIT analytic integral program. This work was supported by the National Science Foundation via Grant No. CHE-9970033 to E.A.C. and via a graduate research fellowship to D.W.

- ¹ P. Pulay and S. Saebø, *Theor. Chim. Acta* **69**, 357 (1986).
- ² S. Saebø and P. Pulay, *J. Chem. Phys.* **86**, 914 (1987).
- ³ S. Saebø and P. Pulay, *Annu. Rev. Phys. Chem.* **44**, 213 (1993).
- ⁴ T. J. Martinez and E. A. Carter, *J. Chem. Phys.* **100**, 3631 (1994).
- ⁵ R. B. Murphy, M. D. Beachy, R. A. Friesner, and M. N. Ringnalda, *J. Chem. Phys.* **103**, 1481 (1995).
- ⁶ R. B. Murphy, W. T. Pollard, and R. A. Friesner, *J. Chem. Phys.* **106**, 5073 (1997).
- ⁷ Jaguar 4.1 (Schrödinger, Inc., Portland, Oregon, 2000).
- ⁸ G. Hetzer, P. Pulay, and H.-J. Werner, *Chem. Phys. Lett.* **290**, 143 (1998).
- ⁹ M. Schütz, G. Hetzer, and H.-J. Werner, *J. Chem. Phys.* **111**, 5691 (1999).
- ¹⁰ G. Hetzer, M. Schütz, H. Stoll, and H. J. Werner, *J. Chem. Phys.* **113**, 9443 (2000).
- ¹¹ P. E. Maslen and M. Head-Gordon, *Chem. Phys. Lett.* **283**, 102 (1998).
- ¹² P. E. Maslen and M. Head-Gordon, *J. Chem. Phys.* **109**, 7093 (1998).
- ¹³ M. S. Lee, P. E. Maslen, and M. Head-Gordon, *Chem. Phys. Lett.* **112**, 3592 (2000).
- ¹⁴ P. E. Maslen, M. S. Lee, and M. Head-Gordon, *Chem. Phys. Lett.* **319**, 205 (2000).
- ¹⁵ J. Almlöf, *J. Chem. Phys.* **96**, 489 (1991).
- ¹⁶ P. Y. Ayala and G. E. Scuseria, *J. Chem. Phys.* **110**, 3660 (1999).
- ¹⁷ N. C. Handy, P. J. Knowles, and K. Somasundram, *Theor. Chim. Acta* **68**, 87 (1985).
- ¹⁸ K. Andersson and B. O. Roos, *Int. J. Quantum Chem.* **45**, 591 (1993).
- ¹⁹ R. B. Murphy and R. P. Messmer, *Chem. Phys. Lett.* **183**, 443 (1991).
- ²⁰ P. Celani and H.-J. Werner, *J. Chem. Phys.* **112**, 5546 (2000).
- ²¹ C. Hampel and H.-J. Werner, *J. Chem. Phys.* **104**, 6286 (1996).
- ²² M. Schütz and H.-J. Werner, *J. Chem. Phys.* **114**, 661 (2001).
- ²³ M. Schütz and H.-J. Werner, *Chem. Phys. Lett.* **318**, 370 (2000).
- ²⁴ M. Schütz, *J. Chem. Phys.* **113**, 9986 (2000).
- ²⁵ G. E. Scuseria and P. Y. Ayala, *J. Chem. Phys.* **111**, 8330 (1999).
- ²⁶ P. Constans, P. Y. Ayala, and G. E. Scuseria, *J. Chem. Phys.* **113**, 10 451 (2000).
- ²⁷ D. Jayatilaka and T. J. Lee, *J. Chem. Phys.* **98**, 9734 (1993).
- ²⁸ T. D. Crawford, T. J. Lee, and H. F. Schaefer, *J. Chem. Phys.* **107**, 9980 (1997).
- ²⁹ P. J. Knowles, C. Hampel, and H.-J. Werner, *J. Chem. Phys.* **99**, 5219 (1993).
- ³⁰ P. Neogrády and M. Urban, *J. Chem. Phys.* **100**, 3706 (1994).
- ³¹ P. G. Szalay and J. Gauss, *J. Chem. Phys.* **107**, 9028 (1997).
- ³² A. Balkova and R. J. Bartlett, *J. Chem. Phys.* **193**, 364 (1992).
- ³³ X. Li and J. Paldus, *J. Chem. Phys.* **103**, 6536 (1995).
- ³⁴ R. J. Bartlett, *J. Phys. Chem.* **93**, 1697 (1989).
- ³⁵ R. J. Bartlett, in *Modern Electronic Structure Theory*, edited by D. R. Yarkony (World Scientific Publishing, Singapore, 1995), p. 1047.
- ³⁶ M. Urban, I. Cernusak, V. Kello, and J. Noga, in *Electron Correlation in Atoms and Molecules*, edited by S. Wilson (Plenum, New York, 1987), p. 117.
- ³⁷ P. Piecuch, N. Oliphant, and L. Adamowicz, *J. Chem. Phys.* **99**, 1875 (1993).
- ³⁸ P. Piecuch and L. Adamowicz, *J. Chem. Phys.* **102**, 898 (1995).
- ³⁹ A. Banerjee and J. Simons, *J. Chem. Phys.* **76**, 4548 (1982).
- ⁴⁰ W. D. Laidig and R. J. Bartlett, *Chem. Phys. Lett.* **104**, 424 (1984).
- ⁴¹ K. B. Ghose, P. Piecuch, and L. Adamowicz, *J. Chem. Phys.* **103**, 9331 (1995).
- ⁴² K. Kowalski and P. Piecuch, *Chem. Phys. Lett.* **344**, 165 (2001).
- ⁴³ T. V. Voorhis and M. Head-Gordon, *J. Chem. Phys.* **115**, 7814 (2001).
- ⁴⁴ S. Saebø and P. Pulay, *Chem. Phys. Lett.* **113**, 13 (1985).
- ⁴⁵ T. J. Martinez, A. Mehta, and E. A. Carter, *J. Chem. Phys.* **97**, 1876 (1992).
- ⁴⁶ T. J. Martinez and E. A. Carter, *J. Chem. Phys.* **98**, 7081 (1993).
- ⁴⁷ T. J. Martinez and E. A. Carter, *J. Chem. Phys.* **102**, 7564 (1995).
- ⁴⁸ G. Reynolds, T. J. Martinez, and E. A. Carter, *J. Chem. Phys.* **105**, 6455 (1996).
- ⁴⁹ G. Reynolds and E. A. Carter, *Chem. Phys. Lett.* **265**, 660 (1997).
- ⁵⁰ D. Walter and E. A. Carter, *Chem. Phys. Lett.* **346**, 177 (2001).
- ⁵¹ J. A. Pople, J. S. Binkley, and R. Seeger, *Int. J. Quantum Chem.* **S11**, 149 (1977).
- ⁵² E. R. Davidson and D. W. Silver, *Chem. Phys. Lett.* **52**, 403 (1978).
- ⁵³ S. R. Langhoff and E. R. Davidson, *Int. J. Quantum Chem.* **8**, 61 (1974).
- ⁵⁴ P. E. M. Siegbahn, *Chem. Phys. Lett.* **58**, 421 (1978).
- ⁵⁵ L. Meissner, *Chem. Phys. Lett.* **146**, 204 (1988).
- ⁵⁶ W. Duch and G. H. F. Diercksen, *J. Chem. Phys.* **101**, 3018 (1994).
- ⁵⁷ P. J. A. Ruttink, J. H. van Lenthe, R. Zwaans, and G. C. Groenenboom, *J. Chem. Phys.* **94**, 7212 (1991).
- ⁵⁸ R. Gdanitz and R. Ahlrichs, *Chem. Phys. Lett.* **143**, 413 (1988).
- ⁵⁹ P. G. Szalay and R. J. Bartlett, *Chem. Phys. Lett.* **214**, 481 (1993).
- ⁶⁰ R. Fink and V. Staemmler, *Theor. Chim. Acta* **87**, 129 (1993).
- ⁶¹ J.-P. Daudey, J.-P. Heully, and J.-P. Malrieu, *J. Chem. Phys.* **99**, 1240 (1993).
- ⁶² N. Ben Amor and D. Maynau, *Chem. Phys. Lett.* **286**, 211 (1998).
- ⁶³ L. Adamowicz and J.-P. Malrieu, *J. Chem. Phys.* **105**, 9240 (1996).
- ⁶⁴ W. Duch and J. Karwowski, *Comput. Phys. Rep.* **2**, 93 (1985).
- ⁶⁵ W. Duch and J. Karwowski, *Int. J. Quantum Chem.* **22**, 783 (1982).
- ⁶⁶ S. F. Boys, *Quantum Theory of Atoms, Molecules, and the Solid State* (Academic, New York, 1966), p. 263.
- ⁶⁷ J. P. Hay, W. J. Hunt, and W. A. Goddard III, *J. Am. Chem. Soc.* **94**, 8293 (1972).
- ⁶⁸ W. A. Goddard III and R. C. Ladner, *J. Am. Chem. Soc.* **93**, 6750 (1971).
- ⁶⁹ B. O. Roos, *Chem. Phys. Lett.* **15**, 153 (1972).
- ⁷⁰ E. R. Davidson, *Comput. Phys. Commun.* **53**, 49 (1989).
- ⁷¹ N. C. Handy, *Chem. Phys. Lett.* **74**, 280 (1980).
- ⁷² J. Paldus, *Theoretical Chemistry: Advances and Perspectives* (Academic, New York, 1974), p. 131.
- ⁷³ I. Shavitt, *Int. J. Quantum Chem., Symp.* **11**, 131 (1977).
- ⁷⁴ I. Shavitt, *Int. J. Quantum Chem., Symp.* **12**, 5 (1978).
- ⁷⁵ I. Shavitt, *Chem. Phys. Lett.* **63**, 421 (1979).
- ⁷⁶ T. Yamanouchi, *Proc. Phys. Math. Soc. Jpn.* **18**, 623 (1936).
- ⁷⁷ M. Kotani, A. Amemiya, E. Ishiguro, and T. Kimura, *Table of Molecular Integrals* (Maruzen, Tokyo, 1963).
- ⁷⁸ W. A. Goddard III, *Phys. Rev.* **157**, 73 (1965).
- ⁷⁹ W. Duch, *Int. J. Quantum Chem.* **27**, 59 (1985).
- ⁸⁰ W. Duch, *GRMS or Graphical Representation of Model Spaces* (Springer, Berlin, New York, 1986).
- ⁸¹ S. A. Orszag, *Stud. Appl. Math.* **51**, 253 (1972).
- ⁸² R. A. Friesner, *Annu. Rev. Phys. Chem.* **42**, 341 (1991).
- ⁸³ T. J. Martinez and E. A. Carter, *Modern Electronic Structure Theory Part II* (World Scientific, Singapore, 1995), p. 1132.
- ⁸⁴ R. A. Friesner, *J. Chem. Phys.* **85**, 1462 (1986).
- ⁸⁵ C. C. Tazartes, C. R. Anderson, and E. A. Carter, *J. Comput. Chem.* **19**, 1300 (1998).
- ⁸⁶ P. E. M. Siegbahn, *J. Chem. Phys.* **72**, 1647 (1979).
- ⁸⁷ W. Meyer, *J. Chem. Phys.* **64**, 2901 (1976).
- ⁸⁸ V. R. Saunders and J. H. van Lenthe, *Mol. Phys.* **48**, 923 (1983).
- ⁸⁹ F. W. Bobrowicz and W. A. Goddard III, *Methods of Electronic Structure Theory* (Plenum, New York, 1977), p. 79.
- ⁹⁰ R. B. Murphy, R. A. Friesner, and W. A. Goddard III, *J. Chem. Phys.* **101**, 2986 (1994).
- ⁹¹ L. B. Harding and W. A. Goddard III, *J. Am. Chem. Soc.* **97**, 6293 (1975).
- ⁹² R. B. Murphy and R. P. Messmer, *J. Chem. Phys.* **97**, 4974 (1992).
- ⁹³ E. R. Davidson, *J. Comput. Phys.* **17**, 87 (1975).
- ⁹⁴ S. Rettrup, *J. Comput. Phys.* **45**, 100 (1982).
- ⁹⁵ HERMIT integral program, T. U. Helgaker (unpublished).
- ⁹⁶ L. E. McMurchie and E. R. Davidson, *J. Comput. Phys.* **26**, 218 (1978).
- ⁹⁷ D. R. Lide, *Handbook of Chemistry and Physics* (CRC, Boca Raton, FL, 2001).
- ⁹⁸ S. Saebø and P. Pulay, *J. Chem. Phys.* **88**, 1884 (1988).
- ⁹⁹ A. Venkatnathan, D. Walter, and E. A. Carter (unpublished).
- ¹⁰⁰ W. J. Hehre, R. Ditchfield, and J. A. Pople, *J. Chem. Phys.* **56**, 2257 (1972).
- ¹⁰¹ R. Krishnan, J. S. Binkley, R. Seeger, and J. A. Pople, *J. Chem. Phys.* **72**, 650 (1980).
- ¹⁰² T. H. Jr. Dunning, *J. Chem. Phys.* **90**, 1007 (1989).