28 September 2001

ELSEVIER

Chemical Physics Letters 346 (2001) 177–185

CHEMICAL PHYSICS LETTERS

www.elsevier.com/locate/cplett

Multi-reference weak pairs local configuration interaction: efficient calculations of bond breaking

Derek Walter, Emily A. Carter *

Department of Chemistry and Biochemistry, University of California, P.O. Box 951569, Los Angeles, CA 90095-1569, USA

Received 4 June 2001; in final form 23 July 2001

Abstract

We present a new local multi-reference singles and doubles configuration interaction (MRSDCI) algorithm. The method presented here eliminates configurations if they involve simultaneous excitations out of widely separated internal orbitals and is therefore based on the weak pairs approximation of Saebø and Pulay. Although the resulting truncated CI expansions have only about 50% as many CSFs as the non-local MRSDCI, we show that they can recover over 99% of the correlation energy. Additionally, we show for the first time that they can accurately describe bond dissociation. © 2001 Elsevier Science B.V. All rights reserved.

1. Introduction

To obtain quantitative accuracy in a quantum chemistry calculation one must use some post self-consistent-field method (SCF) to treat electron correlation. Configuration interaction (CI) is perhaps the oldest and most well known method for doing this [1,2]. Indeed, CI is extremely simple in concept. Briefly, CI requires one to build a set of n-electron configuration state functions (CSFs or spin-adapted Slater determinants) by replacing occupied molecular orbitals (MOs) in some SCF reference function with virtual (unoccupied) orbitals. The Hamiltonian is then diagonalized in the resulting basis of n-electron configurations. Configuration interaction can be easily used in conjunction with multi-configuration SCF (MCSCF) wavefunctions to incorporate electron correlation into wavefunctions which properly describe the cleavage of chemical bonds. This makes the CI method quite general and applicable to a wide range of systems and phenomena.

With the introduction of direct-CI [3] and graphical methods, [4–6] it has become possible to routinely carry out large-scale CI calculations with a few million CSFs on small molecules, resulting in very accurate electronic energies. For the case of CI restricted to single and double excitations from the reference (SDCI), Siegbahn [7] made a particularly important contribution in describing how to treat the electrons occupying the virtual orbitals in a very efficient manner. Unfortunately, in terms of disk storage, memory, and CPU time, CI still scales very poorly with the size of the system, making the treatment of larger systems difficult if not impossible. The main disk storage requirement for a CI calculation is storage of all the electron repulsion integrals (ERIs), which scales as N^4, where N is the number of basis functions. The memory and CPU requirements, on the other

* Corresponding author. Fax: +1-310-267-0319.
E-mail address: eac@chem.ucla.edu (E.A. Carter).

PII: S0009-2614(01)00966-6

hand, are far more severe and are the factors which truly limit the size of a CI calculation that one may do. For optimum efficiency, one would like to retain two copies of the vector of CI expansion coefficients (the CI vector) in high-speed memory [8]. This is in fact a rather demanding requirement and will quickly overwhelm typical machines. By sequentially reading in blocks of the CI vector, one can partially overcome this requirement. However, the problem of the CPU time requirement still remains. For the most common type of CI calculation, SDCI, the CPU time scales as N^6. For CI including all possible excitations out of the reference (full CI), the method scales as $N!$. Thus, the treatment of large systems using the CI method necessitates the development of specialized CI algorithms which are less resource intensive than standard CI. In particular, CPU time and memory consumption must be reduced.

A great deal of research has been devoted to developing CI methods which scale better with the system size but still retain a high level of accuracy. Saebø and Pulay (SP) used orbital localization to truncate the number of CSFs used in the CI expansion [9], reducing both the memory and CPU requirements. Martinez et al. later applied the pseudospectral approximation [10,11] for computing ERIs to full CI [12], doubles CI [13] and multi-reference SDCI (MRSDCI) [14]. Through the direct use of intermediate pseudospectral ERI quantities in the CI equations, these methods reduced the CPU time required for a CI calculation. However, they offered no relief for the memory requirements. Reynolds et al. [15,16] later incorporated localization into the pseudospectral CI methods, yielding a method which addressed both the CPU time and memory requirements. However, with the exception of the study of Martinez et al. on pseudospectral MRSDCI, none of the reduced scaling CI methods were developed for multi-reference wavefunctions. This is a rather serious limitation as it precludes the use of these techniques for describing bond dissociation, perhaps the most important phenomenon in chemistry.

Here, we present a new multi-reference local SDCI program. Unlike previous implementations of local CI, it is not restricted to excitations out of a single reference and can therefore be used to study bond breaking. Furthermore, the new local CI presented here is based upon the symmetric group graphical approach (SGGA) CI of Duch and Karkowski [17,18]. In this formalism, the spatial part of the wavefunction is separated from the spin part of the wavefunction. This has a number of advantages for doing local CI, where the CSFs are selected solely on the basis of the spatial features of the MOs. The most important of these advantages is the rapid on-the-fly calculation of integral coupling coefficients for all the different spin couplings associated with a given pair of spatial configurations. This allows one to quickly and efficiently reorganize a CI calculation to accommodate the elimination of configurations associated with local CI. This will be discussed in more detail below. In addition, single point energies and potential energy surfaces calculated using our new multi-reference local CI method will be discussed. It will be shown that our new CI method accurately describes bond dissociation using significantly fewer resources than standard MRSDCI.

2. Theory

2.1. Local configuration interaction

Because the Coulomb potential between electron pairs is short ranged, the motion of widely separated electrons is only very weakly correlated. This fact can be used to greatly reduce the effort required in methods used to describe electron correlation. One way of doing this is the weak pairs approximation, first introduced by SP [19]. The weak pairs approximation in the CI method neglects or approximates the contributions of CSFs having simultaneous single excitations out of widely separated MOs, as the electrons occupying these MOs are weakly correlated. Of course, in order to meaningfully discuss the notion of widely separated orbitals, the occupied MOs should be localized. In this study and the work of SP, this is done using the Boys method [20].

In previous implementations of weak pairs CI by SP [19] and Reynolds et al. [15], the weak pairs

were determined by calculating the distance between the centroids of pairs of localized molecular orbitals. In the weak pairs CI of Reynolds et al., if this distance was beyond a specified cut-off radius, the pair was characterized as making a negligible contribution to the correlation energy. CSFs having simultaneous excitations out of these orbitals were then eliminated from the CI expansion while simultaneous excitations out of all other pairs of orbitals were allowed. In the weak pairs CI of SP, on the other hand, pairs of orbitals were divided into three classes. Like Reynolds et al., there were categories for totally negligible pairs and strongly correlated pairs. However, there was also a class for pairs of MOs separated by an intermediate distance. The contributions from excitations out of these pairs of MOs were treated approximately using second-order perturbation theory.

Here, we take a different approach to determining the weak pairs. A Mulliken analysis is first done on each localized internal orbital to determine the atoms upon which the orbital most heavily resides. The atoms are then sorted from greatest to least according to this criteria. Next, a group of atoms is selected from the top of this list such that the sum of their contributions to the total charge exceeds some threshold. Generally, the orbitals are normalized to unity so this threshold is a number like 0.75, for example. The maximum distance between any two atoms in this group, $r_{\max}$, is then calculated, as is a charge weighted average position, r_c, for the atoms in the group. Finally, a sphere of radius $\alpha r_{\max}$ centered at r_c is associated with the orbital. Here, α is an adjustable parameter on the order of 1.0 chosen to adjust the radius slightly from $r_{\max}$. Weak-pairs are then defined as pairs of orbitals whose spheres do not overlap. Compared to the previously used method of determining weak pairs, this method has the advantage of more accurately taking into account the spatial extent of the individual orbitals. Additionally, the scheme can be seamlessly applied to virtual orbitals for use in a CI algorithm where the virtual orbital space is truncated. This possibility will be explored in future research.

A special case arises for orbitals localized entirely on one atom, such as core orbitals. For these orbitals, a default sphere radius should be chosen. It should be noted, however, that the CI energies are insensitive to this default radius and it can be taken to have a fairly small value. This is simply a reflection of the fact that core electrons are held close to the nucleus and their motion is strongly correlated only with other electrons localized on the same atom.

2.2. SGGA approach

Large-scale CI is carried out using direct methods [3] in which the full Hamiltonian is never constructed. Rather, the Hamiltonian is diagonalized using procedures which only require one to multiply the Hamiltonian by a trial vector [21]. To do this, matrix elements or partial matrix elements of the Hamiltonian are constructed using the expression

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

where i, j, k, l are orbitals, μ and ν are CSFs, and the $A_{ij}^{\mu\nu}$ and $A_{ijkl}^{\mu\nu}$ are known as integral coupling coefficients. In general, there are two styles of large scale CI available today. The first relies on Slater determinants as the CSFs [22]. The principal advantage of these methods is that integral coupling coefficients can be computed very inexpensively in terms of both computational and mathematical effort. This, of course, makes CI algorithms based on Slater determinants very rapid and relatively easy to implement. However, except in the case of closed shell singlets and open shells coupled to maximum spin, Slater determinants are not native spin eigenfunctions. Because of the poor scaling of the CI method, CI algorithms intended for large systems should include as few CSFs as possible. Thus, it is advantageous to use only spin-adapted CSFs in such a CI method, ruling out the use of Slater determinants.

Spin adapted many-electron wavefunctions were first formulated in terms of the symmetric group [23–25]. However, large-scale CI methods based on spin-adapted CSFs were first implemented in terms of the unitary group approach (UGA) [26] and its graphical form (GUGA) [4–6]. In this approach, the spin and spatial parts of the wavefunction are intertwined. Consequently,

integral coupling coefficients are calculated separately for all the different spin couplings associated with a given configuration. Furthermore, the procedure for doing this is rather complicated and does not lend itself well to on-the-fly computation. Thus, coupling coefficients are calculated once beforehand and stored on disk in a sequential access file. There are several disadvantages to this arrangement. First, it introduces a strong dependence on low speed disk storage. Second, one must carry out the CI in a fashion dictated by the sequence of coupling coefficients on the file, generally forcing one to sort the coupling coefficient file to carry out a specialized CI calculation. In order to carry out a truncated CI calculation, one must first prune the coupling coefficient file. Finally, the GUGA-style programs available for generating coupling coefficients are typically quite limited in the number of internal orbitals they can handle, making them inappropriate for use in the development of CI programs designed to treat larger systems. It should be noted that the pseudospectral CI method of Martinez et al. [12–14] and the local pseudospectral CI method of Reynolds et al. [15,16] were implemented in terms of the GUGA formalism and therefore suffered from the difficulties described above.

Subsequent to the development of GUGA-style CI programs, large-scale CI algorithms based on the symmetric group began to appear [17,18]. In addition to offering a formalism which is much simpler compared to the UGA, the symmetric group approach (SGA) offers special advantages for local CI. The key feature of the SGA is the separation of the spin and spatial parts of the wavefunction. As described above, in local CI CSFs are eliminated only according to the spatial aspects of the MOs. In the SGA, this elimination can be done en masse for all the spin couplings associated with a given orbital configuration. In the UGA, this is not possible because the spin and spatial parts of the wavefunction are entangled. For the configurations not eliminated from the CI, all coupling coefficients for all spin functions associated with an orbital occupation are computed simultaneously in the SGA. These coupling coefficients take the form of symmetric group representation matrices, and extremely efficient algorithms for computing these matrices have been developed by Duch [27,28]. In fact, the efficiency of the algorithms is such that coupling coefficients can be recalculated as they are needed. Thus, one no longer needs to rely on a coupling coefficient file which is stored on disk, eliminating the need to sort such a file and allowing one to organize the CI calculation in any way one desires. Clearly, the SGA offers both a more efficient and more elegant means of carrying out local CI.

For the purpose of carrying out CI calculations on large molecules, the SGA offers an additional advantage. Coupling coefficient codes based on the UGA are generally limited in the number of orbitals they can handle. In the SGA, however, the limiting factor is ultimately the maximum number of open shells. For SDCI, this maximum is four greater than the largest number of open shells in any of the references. This is a quantity which depends on the spin multiplicity and the choice of reference wavefunction, not the size of the molecule. Furthermore, for the overwhelming majority of CI calculations imaginable, this is a quantity which is easily handled by the available algorithms for computing the symmetric group representation matrices.

3. Calculational details

We have developed a new FORTRAN 90 MRSDCI program based on the SGGA–CI algorithms of Duch and Karwowski [17,18]. Integral coupling coefficients are calculated as they are needed during the diagonalization process using the algorithms of Duch [27,28]. To minimize the memory requirement, dynamic memory allocation is used throughout the program. Following the recommendation of Siegbahn [7], our code takes advantage of the very simple structure of the external space. However, the treatment of the external space has not yet been vectorized in a manner described by Saunders and van Lenthe [29]. We expect a substantial enhancement in the efficiency of our program once this is done. In most cases, some of the highest energy orbitals were eliminated from the virtual space to keep the total number of CSFs down to a few million.

The Hamiltonian matrix was diagonalized using the Davidson method [8]. A starting guess for this diagonalization procedure was obtained by diagonalizing a small Hamiltonian formed in the basis of all CSFs containing only internal orbitals.

The SCF reference wavefunctions and optimized geometries were computed either with the HONDO [30] or JAGUAR [31] quantum chemistry packages. Molecular integrals in the MO basis were computed using HONDO. All calculations were done on a Compaq DEC ES40 machine.

4. Results and discussion

4.1. Calibration of the method

As described in Section 2.1, our local CI scheme depends on three parameters. First, there is a charge threshold used to determine the group of atoms upon which the localized internal orbitals most heavily reside. Second, there is a parameter α which serves to scale the radii of the spheres associated with each orbital. Finally, there is a default sphere radius for orbitals localized entirely on one atom. To determine how these parameters effect the CI energies, we carried out a series of calibration calculations on the *trans*-4-octene molecule. Because this molecule is fairly long compared to a typical bond length, we expect the local CI approximation to perform very well and lead to substantial savings in CPU time. For these studies, a Hartree–Fock (HF) reference and the DZP basis of Dunning [32] were used. The 68 lowest energy virtual orbitals were used as correlating orbitals, leading to 2370753 CSFs for non-local SDCI. Table 1 shows the variation in the percentage of correlation energy recovered as the charge threshold is varied in local SDCI. For a charge threshold of 0.30, the local CI expansion contains roughly one-third the number of CSFs as the non-local CI, yet it still recovers a respectable 96.45% of the correlation energy. Increasing the threshold to 1.00 yields a CI expansion having slightly more than half the number of CSFs of the non-local CI, and almost 99% of the correlation energy is recovered. It should be noted that as the threshold is increased to 1.00 the number of CSFs does not approach the number of CSFs used in non-local CI. This is a consequence of the fact that the Boys localized orbitals are often localized entirely on only a few atoms.

Table 1
Percentage of correlation energy recovered as a function of charge threshold

Threshold	No. of CSFs	Time per iteration (s)	Correlation energy (%)
0.30	789345	2800	96.45
0.60	928065	3000	98.88
0.85	978929	3100	98.88
1.00	1464449	4200	98.89

The default radius is 2.0 bohr and α is 1.0.

Table 2 shows the variation in the percentage of the correlation energy recovered as a function of the radius multiplier α. Unlike the charge threshold, the correlation energy recovered smoothly approaches 100% as α is increased. For an α of 1.3 there are roughly half as many CSFs as the non-local CI, and 99.57% of the correlation energy is recovered. When a similar number of CSFs were used in the CI expansion by increasing the charge threshold (Table 1), the percentage of correlation energy recovered was less than 99%. Thus, systematically increasing α is a more efficient method of adjusting the number of CSFs to maximize the amount of correlation energy recovered. This result is hardly surprising. For an orbital localized mainly on a few atoms but having a small amount of charge on a handful of stray atoms, increasing

Table 2
Percentage of correlation energy recovered as a function of radius multiplier α

α	No. of CSFs	Time per iteration (s)	Correlation energy (%)
0.30	789345	2800	96.45
0.80	826337	3000	97.70
1.00	978929	3100	98.88
1.30	1219377	3800	99.57
1.60	1404337	4100	99.85
2.00	1677153	4800	99.95
2.50	1899105	5700	99.99
3.00	2042449	6100	99.99

The default radius is 2.0 bohr and charge threshold was 0.8.

the charge threshold to 1.00 will cause these stray atoms to be included in the sphere associated with such an orbital. This will lead to the inclusion of CSFs having simultaneous excitations out of the orbital in question and orbitals localized near the stray atoms. Clearly, these CSFs will not make a substantial contribution to the correlation energy. On the other hand, increasing α simply extends the radius of the sphere surrounding the atoms upon which the orbital is localized. The additional CSFs included in the CI expansion by doing this will clearly make a greater contribution to the correlation energy than the additional CSFs included by increasing the charge threshold.

Studies to determine the effect of the default radius on the percentage of the correlation energy recovered were also carried out. We find that increasing the default radius from 0.05 to 4.00 bohr lead to the recovery of only an additional 0.03% of the correlation energy in *trans*-4-octene. This result can be explained by noting that the default radius is used to define a radius for spheres associated with orbitals localized entirely on just one atom. Generally, these are core orbitals. As the core electrons are held very close to the nucleus, the motion of these electrons is strongly correlated only with other electrons localized on the same atom. So long as the default radius is not zero, CSFs having simultaneous excitations out of core orbitals and orbitals localized on the same atom will be included in the CI expansion. The strong localization of the core orbitals makes this sufficient to describe the correlation of the core electrons.

4.2. Scaling of the method with the size of the system

Fig. 1 shows the CPU times per iteration for both non-local and local CI calculations on a series of saturated alkanes ranging from methane to dodecane. All calculations were done using a HF reference and a 6-31G basis set [33]. For the local CI calculations α was 1.3 and the charge threshold was 0.8. This choice of parameters lead to recovery of more than 99.6% of the non-local CI correlation

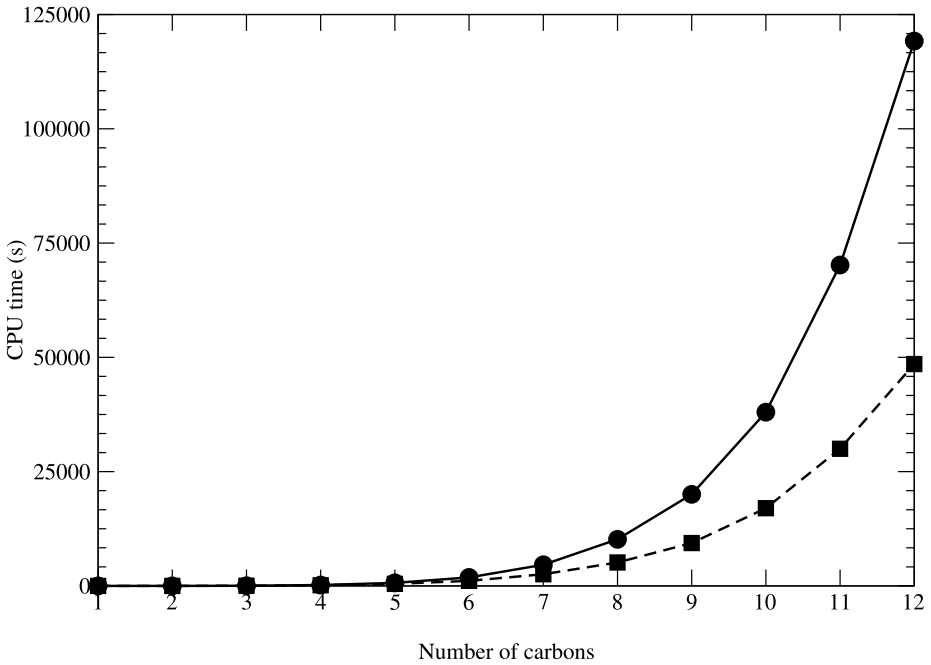

Fig. 1. CPU times per iteration for CI calculations on saturated alkanes. Solid line with circles, conventional non-local CI; dashed line with squares, weak pairs SDCI.

energy in all cases. For 1–4 carbon atoms the local CI method leads to very little reduction in computational effort. In the case of butane, for example, the computational effort was reduced by only 23%. As the size of the alkane increases beyond butane, the savings from the weak pairs approximation become more substantial. For octane and dodecane, the local CI calculations take only 50% and 40% as long as the non-local CI calculations, respectively. The decrease in computational effort with increasing molecular size can be easily understood in terms of the increasing number of weak pairs in larger molecules. From butane to dodecane this number increases from 36 to 791.

In almost all cases, the virtual orbital space is much larger than the occupied orbital space. However, since in SDCI only two electrons may occupy the virtual space, the treatment of the virtuals may be simplified dramatically [7]. Although this leads to considerable savings in the time required to treat the virtuals, the large size of the virtual space still demands the overwhelming majority of the computational effort in a CI calculation. In the weak pairs approximation, simultaneous excitations to all virtuals out of widely separated occupied orbitals are neglected. The savings are therefore derived from reducing the number of times the virtual orbital space must be treated. Although the code presented here takes advantage of the relatively simple structure of the virtual space, the treatment has not been optimized according to the prescriptions of Duch [17] and Saunders and van Lenthe [29]. Once this is done, it is expected that the overall CPU times will be reduced further. However, the savings of the weak pairs CI relative to non-local CI should be unaffected since the weak pairs approximation simply changes the number of times the virtuals must be treated. Similarly, the savings of the local CI relative to non-local CI will not depend on the size of the basis set as this will only increase the total number of virtuals.

4.3. Dissociation of trans-butene

Unlike previous implementations of weak pairs CI [9,15,16], our code is not limited to the use of a single closed shell reference. This opens up the possibility of using our weak pairs CI method to study bond breaking. Here, we use our method to examine the cleavage of the C–C double bond in *trans*-2-butene. A GVB-RCI wavefunction was used as the reference [34,35]; both electron pairs of the C–C double bond were correlated as GVB pairs. This reference wavefunction, which contains nine spatial configurations, provides a qualitatively proper description of the bond dissociation. The 6-31G** basis set of Pople was used [33], and the 42 lowest energy virtual orbitals were used as correlating orbitals. This resulted in 4193202 CSFs for non-local CI. Since the *trans*-2-butene molecule is relatively small, we do not expect the weak pairs approximation to lead to a huge savings in computational effort. However, the point of this set of calculations is not to achieve the maximum reduction in computational effort. Rather, it is to show that the weak pairs CI method can give results comparable to non-local CI.

The potential energy surface (PES) for cleavage of the *trans*-2-butene double bond computed at various levels of theory is shown in Fig. 2. At the GVB-RCI level, the D_e is determined to be 153.0 kcal/mol. Augmenting the GVB-RCI wavefunction with non-local SDCI leads to a D_e of 165.1 kcal/mol. In both cases, the PESs describe smooth dissociation into fragments. When the PES is computed using local CI with an α of 1.0 and a charge threshold of 0.8, the D_e is determined to be 164.2 kcal/mol. Compared to the non-local CI, this is an error of less than 1 kcal/mol. At the equilibrium bond length 99.4% of the correlation energy is recovered and at a bond length of 13.0 bohr 99.7% of the correlation energy is recovered. At all other points, more than 99% of the correlation energy is recovered. Consequently, this local CI surface is almost superimposable upon the non-local CI surface. It should be pointed out that this local CI only reduces the computational effort compared to the non-local CI by roughly 35%. However, the parameters for the non-local CI were chosen to be similar to parameters that one would likely use in calculations on larger molecules. The results shown here indicate that these parameters would likely lead to results comparable to the results obtained from non-local CI calculations (with no CI vector truncation).

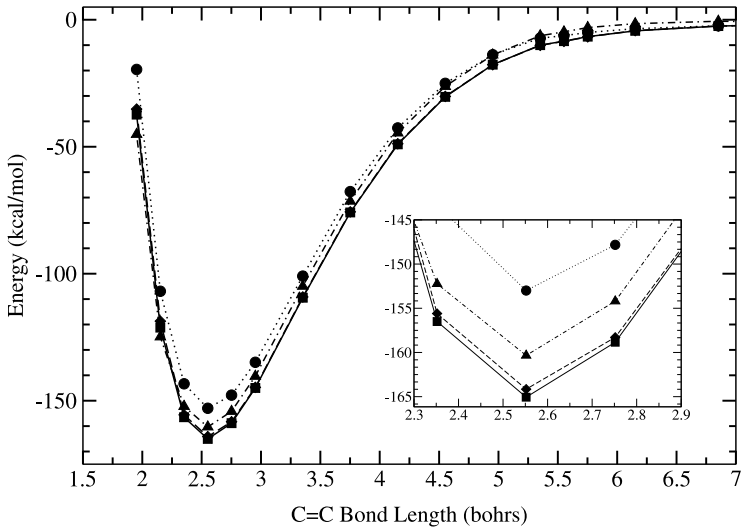

Fig. 2. PES for cleavage of the double bond in *trans*-2-butene computed at various levels of theory. The inset shows the region immediately surrounding the equilibrium bond length. Dotted line with circles, GVB-RCI; alternating dots and dashes with triangles, local SDCI with $\alpha = 0.6$ and charge threshold = 0.75; dashes with diamonds, local SDCI with $\alpha = 1.0$ and charge threshold = 0.8; solid line with squares, conventional non-local SDCI.

When the PES is computed using local CI with an α of 0.6 and a charge threshold of 0.75, the D_e is determined to be 160.3 kcal/mol. Although this is a marked improvement over the GVB-RCI D_e, an error of 4.7 kcal/mol relative to the non-local CI is introduced. The local CI parameters chosen for these calculations correspond to a relatively severe truncation of the CI expansion. As a result, the fraction of the correlation energy recovered is reduced compared to the case described above. At the equilibrium geometry, for example, only 93.6% of the correlation energy is recovered. Nevertheless, the PES computed at this level of theory is quite smooth, and the error in the D_e is moderate. Thus, the multi-reference local CI method appears to be fairly robust.

5. Conclusions

Here, we have presented a new multi-reference weak pairs local SDCI program. It was shown that for large molecules the weak pairs CI method could recover over 99% of the correlation energy with only 40% of the computational effort. For the first time, we have used the local CI method with a multi-reference wavefunction to study bond cleavage. In the case of the *trans*-2-butene double bond it was shown that the local CI method yielded a D_e which was within 1 kcal/mol of the non-local CI D_e. Furthermore, the local CI PES for cleavage of this bond is almost indistinguishable from the non-local CI PES and is just as smooth. When a severe truncation in the CI vector is made, the error in the D_e grows, but the PES retains a high degree of smoothness. We note that our local CI program is built upon the SGGA, in which the spin part of the wavefunction is separated from the spatial part of the wavefunction. This offers a much more elegant formalism for carrying out local CI and is also better suited to the treatment of large systems. The overall speed of the current method can be improved greatly by vectorizing the treatment of the virtual space, which is a subject of our ongoing work. Once this is done, the program will be particularly well suited to a pseudospectral implementation. Additionally, we plan to explore the use of local correlation methods for treating the virtual space in multi-reference CI. Similar to the current work, it is expected that the SGGA will provide a solid foundation for achieving this.

Acknowledgements

D.W. acknowledges Niri Govind, Stuart Watson, Thorsten Kluener, and Frank Starrost for illuminating discussions and technical assistance. D.W. also thanks Wlodiszlaw Duch for providing a copy of his SGGA-CI program for benchmarking purposes. This work was supported by the National Science Foundation via grant #CHE-9970033 to E.A.C. and via a graduate research fellowship to D.W.

References

[1] I. Shavitt, Mol. Phys. 94 (1998) 3.
[2] I. Shavitt, in: Methods of Electronic Structure Theory, Plenum Press, New York, 1977, p. 180.
[3] B.O. Roos, Chem. Phys. Lett. 15 (1972) 153.
[4] I. Shavitt, Int. J. Quantum Chem. Symp. 11 (1977) 131.
[5] I. Shavitt, Int. J. Quantum Chem. Symp. 12 (1978) 5.
[6] I. Shavitt, Chem. Phys. Lett. 63 (1979) 421.
[7] P.E.M. Siegbahn, J. Chem. Phys. 72 (1979) 1647.
[8] E.R. Davidson, J. Comput. Phys. 17 (1975) 87.
[9] S. Saebø, P. Pulay, Chem. Phys. Lett. 113 (1985) 13.
[10] S.A. Orszag, Stud. Appl. Math. 51 (1972) 253.
[11] R.A. Friesner, Ann. Rev. Phys. Chem. 42 (1991) 341.
[12] T.J. Martinez, A. Mehta, E.A. Carter, J. Chem. Phys. 97 (1992) 1876.
[13] T.J. Martinez, E.A. Carter, J. Chem. Phys. 98 (1993) 7081.
[14] T.J. Martinez, E.A. Carter, J. Chem. Phys. 102 (1995) 7564.
[15] G. Reynolds, T.J. Martinez, E.A. Carter, J. Chem. Phys. 105 (1996) 6455.
[16] G. Reynolds, E.A. Carter, Chem. Phys. Lett. 265 (1997) 660.
[17] W. Duch, J. Karwowski, Comput. Phys. Rep. 2 (1985) 93.
[18] W. Duch, J. Karwowski, Int. J. Quantum Chem. 22 (1982) 783.
[19] S. Saebø, P. Pulay, Ann. Rev. Phys. Chem. 44 (1993) 213.
[20] S.F. Boys, in: Quantum Theory of Atoms, Molecules, and the Solid State, Academic Press, New York, 1966, p. 263.
[21] E.R. Davidson, Comput. Phys. Commun. 53 (1989) 49.
[22] N.C. Handy, Chem. Phys. Lett. 74 (1980) 280.
[23] T. Yamanouchi, Proc. Phys. Math. Soc. Jpn. 18 (1936) 623.
[24] M. Kotani, A. Amemiya, E. Ishiguro, T. Kimura, Table of Molecular Integrals, Maruzen, Tokyo, 1963.
[25] W.A. Goddard III, Phys. Rev. 157 (1965) 73.
[26] J. Paldus, in: Theoretical Chemistry: Advances and Perspectives, Academic Press, New York, 1974, p. 131.
[27] W. Duch, Int. J. Quantum Chem. 27 (1985) 59.
[28] W. Duch, GRMS or Graphical Representation of Model Spaces, Springer, Berlin, 1986.
[29] V.R. Saunders, J.H. van Lenthe, Mol. Phys. 48 (1983) 923.
[30] M. Dupuis, F. Johnston, A. Marquez, HONDO 8.5 from CHEM-Station, IBM Corporation, Neighborhood Road, Kingston, NY, 12401, 1994.
[31] JAGUAR 4.0, Schrödinger, Portland, Oregon, 2000.
[32] T.H. Dunning Jr., J. Chem. Phys. 53 (1970) 2823.
[33] W.J. Hehre, R. Ditchfield, J.A. Pople, J. Chem. Phys. 56 (1972) 2257.
[34] L.B. Harding, W.A. Goddard III, J. Am. Chem. Soc. 97 (1975) 6293.
[35] R.B. Murphy, R.A. Friesner, W.A. Goddard III, J. Chem. Phys. 101 (1994) 2986.