



ELSEVIER

14 February 1997

Chemical Physics Letters 265 (1997) 660–666

**CHEMICAL
PHYSICS
LETTERS**

Removal of the bottleneck in local correlation methods

Gregg Reynolds, Emily A. Carter

Department of Chemistry and Biochemistry, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90095-1569, USA

Received 25 November 1996; in final form 3 December 1996

Abstract

We present preliminary results from a new hybrid pseudospectral local correlation method in which we combine the pseudospectral approximation to two-electron integral evaluation with local truncation of the virtual space to accelerate singles and double excitation configuration interaction (SDCI). The combination of the two leads even in small molecules to overall speed-ups of factors of 3–5 and reduces the main disk space requirement by factors of 9–20 while retaining at least 98% of the correlation energy of the exact SDCI calculations. These preliminary results bode well for the advantages of further development of pseudospectral local correlation methods.

1. Introduction

For many common chemical processes, the self-consistent field (SCF) description of the electronic structure is insufficient. One common way of obtaining an estimate of the correlation energy is singles and doubles configuration interaction (SDCI). SDCI, however, has a number of limitations that prevent calculations from being performed on large molecules. These limitations are separate from the problem of size-consistency, which increasingly affects the accuracy of energies predicted by SDCI as molecular size increases. However, the same obstacles addressed here prevent calculations on large molecules in Davidson- [1] or Pople-corrected [2] SDCI; they also plague many other similarly formulated correlation energy methods, e.g., third- (and higher) order Møller–Plesset perturbation theory, coupled-cluster theory, and the coupled electron pair approximation.

In performing calculations on larger molecules with conventional (spectral) SDCI, one faces three main problems. First, the overall cost of the method

scales as N^6 , where N is the number of basis functions. Secondly, the amount of memory needed to store the CI vector increases as $n^2N_v^2$, where n is the number of occupied orbitals and N_v is the number of unoccupied or virtual orbitals. Finally, the amount of disk space required to store the all-virtual two-electron integrals, $(ab|cd)$, increases as N_v^4 . All three of these limitations combine to restrict the size of molecules one can treat with SDCI.

We explore here the combination of two approximations which have been used separately to address limitations of SDCI: the pseudospectral approximation [3] and local truncation of the virtual space [4]. The pseudospectral (PS) approximation reduces the amount of disk space required for the two-electron integrals (electron repulsion integrals or ERIs), especially $(ab|cd)$, increases the speed with which they are processed, but does not affect the size of the CI vector. Truncation of the virtuals decreases the size of the CI vector but has no effect on the disk space needed to store the ERIs. Also, in all previous local spectral methods, the treatment of the $(ab|cd)$ was a bottleneck. The savings in truncation of the virtual

space results from using localized correlating orbitals (projected atomic orbitals or PAOs) which are orthogonal to the occupied orbital space. To perform the (ab|cd) part of the calculation in the PAO basis requires an N^3 atomic orbital (AO) to PAO transformation which would dominate the local calculation. Thus, instead the (ab|cd) contributions are typically done in the full AO basis, whose orbitals are not orthogonal to the occupied orbitals, which precludes employing truncation of the virtuals or any savings derived from it. The combination of the PS approximation with truncation of the virtuals might alleviate this problem.

Two previous papers have applied localization and the PS approximations in correlation methods, but neither addressed the problem of handling (ab|cd). Friesner and co-workers developed the first local PS method, local PS MP2 [5]. The MP2 energy expression, however, does not involve (ab|cd) and thus their paper did not directly address the issue of how to handle these integrals. Similarly, recently Reynolds and Carter [6] have completed a study of local PS SDCl, but the most important part of localization – truncation of the virtuals – was not implemented. Here we will use truncation of the virtuals in the PAO basis to accelerate all integrals except (ab|cd), which will be pseudospectrally calculated in the AO basis to avoid the transformation, performance, and disk storage bottlenecks mentioned above, to show how a PS treatment of (ab|cd) can greatly speed up local SDCl.

2. Implementation

Here we discuss details of how these methods were implemented. First, we used the Foster–Boys [7] method to localize the occupied orbitals. The PAOs were formed following the procedure of Sæbø and Pulay [4]; the PAOs result from applying a projector, $\mathbf{P} = \mathbf{I} - \frac{1}{2}\mathbf{DS}$, to the AOs, where \mathbf{S} is the AO overlap matrix and \mathbf{D} is the Hartree–Fock (HF) density matrix. Truncation of the virtuals involves restricting excitations to a limited set of correlating orbitals (a ‘domain’) for each excited configuration (CSF). The savings come from being able to replace N_v in scaling arguments with \bar{N}_v , the average of the domain sizes. \bar{N}_v depends mostly on the basis set

used and is independent of molecular size. We employed a simple hierarchical method of constructing the domains for the CSFs. First, we constructed atomic domains (sets of PAOs associated with a given atom) by placing each PAO into the domain of that atom on which its dominant basis function is centered. An MO domain for each occupied MO was then constructed by taking the union of all the atomic domains of the AOs that contributed significantly (had a normalized coefficient of 0.15 or greater in absolute value) to the MO. Domains for singles CSFs were just the domain of the occupied MO from which the electron is excited (‘hole’ orbital), and domains for doubles CSFs were the union of the domains of the two hole orbitals.

The PS approximation [3] was conducted in the following way. Essentially the PS approximation speeds up a calculation by performing part of the integration of a two-electron integral on a grid:

$$\begin{aligned} (\text{ab|cd}) &= \int \frac{\varphi_a^*(\mathbf{r}_1)\varphi_b(\mathbf{r}_1)\varphi_c^*(\mathbf{r}_2)\varphi_d(\mathbf{r}_2)}{|\mathbf{r}_2 - \mathbf{r}_1|} d\mathbf{r}_1 d\mathbf{r}_2 \\ &= \int \varphi_a^*(\mathbf{r}_1)\varphi_b(\mathbf{r}_1) d\mathbf{r}_1 \int \frac{\varphi_c^*(\mathbf{r}_2)\varphi_d(\mathbf{r}_2)}{|\mathbf{r}_2 - \mathbf{r}_1|} d\mathbf{r}_2 \\ &= \int \varphi_a^*(\mathbf{r}_1)\varphi_b(\mathbf{r}_1) A^{\text{cd}}(\mathbf{r}_1) d\mathbf{r}_1 \\ &\approx \sum_{\mathbf{g}} Q_a(\mathbf{r}_{\mathbf{g}}) R_b(\mathbf{r}_{\mathbf{g}}) A^{\text{cd}}(\mathbf{r}_{\mathbf{g}}) \\ &= \sum_{\mathbf{g}} Q_{a\mathbf{g}} R_{b\mathbf{g}} A_{\mathbf{g}}^{\text{cd}}. \end{aligned}$$

The integration over the Coulomb singularity is done analytically first to form the Coulomb potential, $A^{\text{cd}}(\mathbf{r}_1)$. Then the integral over the second coordinate is performed approximately on a grid using the $A_{\mathbf{g}}^{\text{cd}}$ and the basis functions’ values on the grid, $R_{b\mathbf{g}}$; and then the result is transformed back into function (or spectral) space by $Q_{a\mathbf{g}}$ (the details on forming these quantities can be found in Ref. [8]). The PS approximation’s dominant disk storage requirement is the Coulomb potential on the grid, $A_{\mathbf{g}}^{\text{cd}}$, which scales roughly as N^3 , a savings of N over conventional storage of ERIs. In SDCl and similar correlation methods, the processing of PS quantities can also be accelerated by a factor of N [9].

One particularly important point when using a non-orthogonal basis, such as the PAOs, is the elimination of linear dependence within the domains.

Until we did this correctly our local SDCI calculations *did not converge*. After much experimentation, we followed the first option suggested by Hampel and Werner [10] to eliminate dependence from the domains and their advice on forming an orthonormal diagonalizing basis and updating the CI vector. This approach is similar to that used by Friesner et al. in their local PS MP2 code [11]. As these details are technical and do not affect our results or conclusions, their discussion will be delayed to a future publication [12].

3. Computational details

In this Letter we performed calculations on three molecules – the anti conformer of ethane (C_2H_6), the cis conformer of glyoxal (OCHCHO), and the equilibrium conformer of glycine (NH_2CH_2COOH) – all using the 6-31G** bases of Pople and co-workers [13]. All SDCI calculations were conducted using the HF optimized geometries as performed by GAUSSIAN 92 [14]. The integral and SCF calculations were performed using the gvb2p5 suite of codes [15]. In the CI calculations, no orbitals were kept as frozen core orbitals. PS quantities were calculated using the ‘coarse’ grids (roughly 70 points per atom) of the program PSGVB version 1.0 of Friesner and co-workers [16]. CI coupling coefficients were generated by GUGA2 [17]. Non-local spectral and PS SDCI calculations were performed using MARTCI developed by Martinez and Carter. In all the calculations, the contributions from $(ab|cd)$ were handled in the full AO basis. In the non-local SDCI calculations, the non- $(ab|cd)$ contributions were done in the MO basis; in the local calculations, only truncation of the virtuals was implemented and the PAO basis was used for all terms excepting $(ab|cd)$. In the PS calculations, both local and non-local PS quantities were used only to form contributions from $(ab|cd)$.

One should note from the CPU times per iteration when compared to corresponding quantities for similarly sized molecules and bases in the previous publications of Sæbø and Pulay [4], Friesner and co-workers [5], and Hampel and Werner [10] that our code is not completely optimized. The relative speed-ups provided by localization and the PS ap-

proximation, however, should not be adversely affected by this.

4. Results

Table 1 contains the combined CPU and system timings per iteration for various routines for a series of local and non-local spectral and PS SDCI calculations on three molecules: ethane, glyoxal, and glycine. The timings are broken down by the class of integral treated in a particular contribution; $(i|h|j)$ denote one-electron integrals. $i, j, k,$ and l represent occupied molecular orbitals; $a, b, c,$ and d represent correlating MO or PAO correlating orbitals. In general, our data exemplify that the effect the truncation of the virtuals approximation exhibits depends on the number of correlating orbital indices involved in the integrals that make up the contribution. The two routines that dominate the non-local spectral calculations above and conventional SDCI calculations in general are $(ab|cd)$ and $(ab|ij)/(ai|bj)$. Together both these routines take up over 90% of the CPU time per iteration in all the spectral nonlocal calculations with $(ab|cd)$ taking roughly three times longer than $(ab|ij)/(ai|bj)$. Thus, localization’s acceleration of all the other routines combined will have a more or less negligible effect, as eliminating them entirely would result in a speed-up factor of only 1.1. For this reason, the routines for $(i|h|j)$, $(ij|kl)$, $(ai|jk)$, and $(ai|bc)$ were not optimized and may be slower locally in some cases. Localization alone accelerated the $(ab|ij)/(ai|bj)$ contributions by a factor of 2.35 in glyoxal to a factor of 4.9 in glycine. However, it does not, as discussed above, affect the effort expended with $(ab|cd)$. Thus, assuming that $(ai|bj)/(ab|ij)$ and $(ab|cd)$ take 92% of the CPU time and that $(ab|cd)$ takes three times more CPU time than $(ai|bj)/(ab|ij)$, the greatest speed increase we could expect in these molecules, i.e. if localization eliminated the effort in all other routines, would be a factor of 1.4. Since localization cannot eliminate the other parts of the calculation, speed increases due only to localization are somewhat smaller, ranging from a factor of 1.15 in ethane to a factor of 1.28 in glycine. These are comparable to the speed-up of 1.17 reported by Hampel and Werner [10] for local spectral SDCI on pentane with somewhat more basis

Table 1
CPU times per iteration for SDCI calculations^a on ethane, glyoxal, and glycine

HC segment	SPNL	PSNL	Speed-up	SP-L	Speed-up	PS-L	Speed-up
ethane (60 basis functions)							
(i hj) and (ijkl)	0.17	0.18	0.94	0.38	0.45	0.37	0.46
(i hla)	0.73	0.73	1.00	1.03	0.71	1.01	0.72
(ai jk)	1.86	1.86	1.00	2.37	0.78	2.35	0.79
(ab ij) and (ai bj)	23.85	23.87	1.00	9.55	2.50	9.53	2.50
(ai bc)	3.22	3.22	1.00	1.13	2.85	1.12	2.88
(ab cd)	87.41	24.68	3.54	87.39	1.00	25.06	3.49
total	117.24	54.54	2.15	101.85	1.15	39.44	2.97
glyoxal (70 basis functions)							
(i hj) and (ijkl)	2.16	2.19	0.99	2.66	0.81	2.66	0.81
(i hla)	2.51	2.50	1.00	4.02	0.62	4.10	0.61
(ai jk)	11.56	11.56	1.00	11.81	0.98	11.61	1.00
(ab ij) and (ai bj)	140.56	140.35	1.00	59.77	2.35	59.67	2.36
(ai bc)	12.03	11.93	1.01	4.01	3.00	3.91	3.08
(ab cd)	387.03	96.58	4.01	384.46	1.01	97.35	3.98
total	555.85	265.11	2.10	466.73	1.19	179.3	3.10
glycine (100 basis functions)							
(i hj) and (ijkl)	19.35	19.20	1.01	10.63	1.82	10.58	1.83
(i hla)	12.06	11.84	1.02	9.56	1.26	9.54	1.26
(ai jk)	66.17	63.51	1.04	46.49	1.42	45.70	1.45
(ab ij) and (ai bj)	1013.30	1011.26	1.00	208.26	4.87	207.35	4.89
(ai bc)	78.67	79.35	0.99	18.09	4.35	17.29	4.55
(ab cd)	3106.83	619.00	5.02	3058.92	1.02	606.79	5.12
total	4296.38	1804.16	2.38	3351.95	1.28	897.25	4.79

^a SPNL, PSNL, SP-L, and PS-L refer respectively to results from non-local spectral, non-local PS, local spectral, and local PS calculations. Speed-ups are all measured relative to the non-local spectral results.

functions (130 compared to our range of 60–100). Thus, truncation of the virtuals by itself can only accelerate a spectral SDCI calculation modestly.

What does affect the speed of the calculation quite significantly is the use of AO PS routines to calculate (ab|cd). The PS approximation accelerates the (ab|cd) routine from a factor of 3.5 in ethane to a factor of 5 in glycine. Such improvements in the dominant CPU time expenditure in the calculation lead to the non-local PS calculations being at least a factor of 2 faster than their spectral counterparts in all three molecules. Furthermore, since localization accelerates the next most time-intensive routine, (ab|ij)/(ai|bj), the combination of the localization and PS approximations is even more advantageous. The overall speed increases then for PS local calculations compared to their spectral non-local counterparts ranges from a factor of 3.0 in ethane to a factor of 4.8 in glycine.

The other advantage of the PS approximation is the use of disk space. Table 2 shows the size of files, in megabytes, needed to hold the quantities necessary for forming the (ab|cd) contributions. For spectral calculations, this is just the size of the AO two-electron integral file; for PS calculations, three

Table 2
Disk space usage (in megabytes) for spectral and pseudospectral quantities

Molecule	(ab cd) ^a	A_g^{cd} ^b	R_{gb} ^b	Q_{ag} ^b
ethane	103.7	9.4	0.8	0.8
glyoxal	192.1	12.7	0.8	0.8
glycine	800.1	36.2	1.1	1.1

^a Spectral two-electron integrals in the AO basis.

^b Pseudospectral quantities that taken together approximate (ab|cd): (ab|cd) $\approx \sum_g Q_{ag} R_{gb} A_g^{cd}$. The sum of the disk space requirements for A , R , and Q is thus the total disk space needed for pseudospectral formation of (ab|cd). See text for more details.

files are needed, one each for the Q_{ag} , R_{ga} , and A_g^{ab} . We see for all three cases the combined storage space needed for the PS quantities is much smaller than that for the AO (ab|cd). The space saved by using PS quantities ranges from a factor of 9 in ethane to over a factor of 20 in glycine. Since the disk space needed for the spectral calculation is almost 1 GB for 100 basis functions and has a worse scaling behavior than its PS counterpart, this will mean that there will be a range of cases where a spectral SDCI calculation will be prohibitive due to required disk space but a PS SDCI calculation can be performed. And the PS calculation, whether local or not, can be done many times faster with comparable accuracy, as demonstrated here (below) and previously by Martinez and Carter [9].

Finally, we examine the recovery of the SDCI correlation energy in the approximate SDCI calculations. In Table 3, we see that the predicted correlation energies for all approximations are in good agreement for the three molecules studied. The PS calculations recover slightly more correlation energy than their spectral counterparts. This is not a problem though, since due to the approximate integrals used in a PS calculation, PS SDCI calculations are not variational. The overall discrepancy in energy between any PS calculation and its spectral counterpart is less than $1 \text{ m}E_h$ in all cases, which is within both chemical accuracy and the total accuracy of the method. The local calculations all recover at least 98% of the non-local spectral SDCI correlation energy. This is in good agreement with behavior of local methods reported by Sæbø and Pulay [4], Friesner and co-workers [5], and Hampel and Werner [10]. This small reduction has been generally deemed as acceptable in the literature and has also been

shown to be due in part to reduction in basis set superposition error [18]. Thus, while being between 3.0 and 4.8 times faster, the local PS calculations provide correlation energies that are not significantly different from non-local spectral SDCI correlation energies.

5. Discussion

In order to avoid the prohibitive amount of disk space needed for large calculations with local spectral methods, one can turn to direct calculation [19] of at least the (ab|cd). Here instead of storing these quantities, a direct algorithm would have them calculated on the fly whenever they were needed. As conventional local spectral correlation methods already have any (ab|cd) contributions calculated in the original AO basis, there is no further penalty for partial transformation of the AO integrals as there would be for any other contributions. This might put spectral local correlation methods on a superior footing in terms of disk space usage to the local PS method presented here but would incur a performance penalty for having to calculate all the AO ERIs every iteration. Such a direct handling of the (ab|cd) will eventually become necessary to allow correlation calculations on large molecules to be done by local spectral methods. It should be noted, though, that direct calculation of the (ab|cd) is not limited to local spectral methods. Indeed one could also directly calculate the A_g^{cd} used in PS treatments of the (ab|cd). This would have two distinct advantages over the direct spectral treatment above. First, there are far fewer A_g^{cd} than (ab|cd) needed to con-

Table 3
Correlation energies (in E_h) for SDCI calculations

Molecule	Quantity	SPNL ^b ('exact')	PSNL ^b	SP-L ^b	PS-L ^b
ethane	correlation energies	-0.32077	-0.3216	-0.31532	-0.31622
	% recovery ^a	100.00	100.26	98.30	98.58
glyoxal	correlation energies	-0.55462	-0.55534	-0.54926	-0.54999
	% recovery	100.00	100.13	99.03	99.16
glycine	correlation energy	-0.71035	-0.71068	-0.69924	-0.69961
	% recovery	100.00	100.05	98.44	98.49

^a % Recovery is the recovery of the SDCI correlation energy.

^b Definitions are given in footnote a of Table 1.

struct the external exchange term. This can be seen in the much reduced disk space required to store these quantities above. Secondly, each A_g^{cd} is essentially a one-electron integral, closely related to the nuclear–electron attraction integrals. Thus, not only are there fewer A_g^{cd} , but each of them is more easily calculated. Therefore, while a PS direct treatment of the $(ab|cd)$ would be less efficient than its non-direct counterpart due to the need to repeatedly calculate the A_g^{cd} , direct local PS correlation methods would allow calculations on molecules impossible to treat with non-direct methods and would be superior in performance to direct spectral local methods.

Though the current PS treatment of the $(ab|cd)$ contributions significantly accelerates the local calculation, it is not the most efficient way of using the PS quantities. Here, as noted above, the spectral and PS $(ab|cd)$ routines were organized very similarly. This is quite different from the approach recommended by Martinez and Carter [9] in their PS MRSDCI algorithm where the $(ab|cd)$, $(ai|bc)$, and $(ai|bj)/(ab|ij)$ were all done pseudospectrally in the same routine (in a factorized manner) to allow maximum use of common intermediate quantities. They also calculated all other integrals spectrally to achieve a balance between enhancing performance, which favors using all PS quantities calculated on a sparse grid, and gaining accuracy, which favors the use of dense grids or spectral integrals. Though the formation of common intermediates in a non-orthogonal basis (Martinez and Carter's code uses MO basis integrals) is more complicated since the $(ai|bj)/(ab|ij)$ terms require multiplications by the overlap matrix while $(ab|cd)$ and $(ai|bc)$ do not, it is expected that a similar factorized treatment will be employed in our final code. This should enhance the performance and disk space utilization even further than our current formulation and will be the subject of a future publication [12].

To conclude, we have combined localization of the virtuals (as developed previously by Sæbø and Pulay [4] and implemented in other contexts by Friesner and co-workers [5] and Hampel and Werner [10]) with the PS treatment, as developed by Friesner [3], of the $(ab|cd)$ integrals for the first time in a correlation method. The combination of these two approximations significantly accelerates the two most CPU-intensive contributions in SDCI, those from

$(ai|bj)/(ab|ij)$ and $(ab|cd)$, and greatly reduces the disk space needed by the $(ab|cd)$. Our local PS method recovers over 98% of the full SDCI correlation energy, in agreement with the performance of previous local correlation methods, but achieves these results 3–5 times more quickly than similarly programmed local spectral SDCI. Our continuing development of local PS correlation methods should lead to even further performance improvements over local spectral methods. We have not performed calculations on particularly large molecules because our code has not been optimized yet and because the spectral SDCI benchmarks would have taken prohibitive amounts of CPU time and disk space. Both approximations, though, improve in performance as the size of the molecule increases. This can be seen even in the small range of molecules examined here. Thus, not only will local PS methods allow calculations on molecules impossible spectrally, but the performance advantage will also increase as one treats larger molecules. Finally, though these results have been reported only for SDCI, the improvements are in principle transferable to many other similarly constructed correlation methods.

Acknowledgements

This work was supported by the Office of Naval Research. EAC is also grateful to the Camille and Henry Dreyfus and Alfred P. Sloan Foundations for support via Teacher-Scholar and Research Fellow Awards, respectively. GGR expresses his gratitude to the National Science Foundation for a predoctoral fellowship. Finally we are indebted to Antonio J.R. da Silva and Todd J. Martinez for many helpful discussions.

References

- [1] S.R. Langhoff and E.R. Davidson, *Int. J. Quantum Chem.* 8 (1974) 61.
- [2] J.A. Pople, R. Secger and R. Krishnan, *Int. J. Quantum Chem. Quantum Chem. Symp.* 11 (1977) 149.
- [3] R.A. Friesner, *Ann. Rev. Phys. Chem.* 42 (1991) 341.
- [4] S. Sæbø and P. Pulay, *J. Chem. Phys.* 86 (1987) 914; 88 (1988) 1884.
- [5] R.B. Murphy, M.D. Beachy, R.A. Friesner and M.N. Ringnalda, *J. Chem. Phys.* 103 (1995) 1481.

- [6] G. Reynolds and E.A. Carter, *J. Chem. Phys.* 105 (1996) 6455.
- [7] S.F. Boys, in: *Quantum theory of atoms, molecules, and the solid state*, ed. P.-O. Löwdin (Academic Press, New York, 1966) p. 263.
- [8] M.N. Ringnalda, M. Belhadj and R.A. Friesner, *J. Chem. Phys.* 93 (1990) 3397.
- [9] T.J. Martinez and E.A. Carter, *J. Chem. Phys.* 102 (1995) 7564.
- [10] C. Hampel and H.J. Werner, *J. Chem. Phys.* 104 (1996) 6286.
- [11] R.B. Murphy, private communication.
- [12] G. Reynolds and E.A. Carter, to be published.
- [13] W.J. Hehre, R. Ditchfield and J.A. Pople, *J. Chem. Phys.* 56 (1972) 2257.
- [14] M.J. Frisch, G.W. Trucks, M. Head-Gordon, P.M.W. Gill, M.W. Wong, J.B. Foresman, B.G. Johnson, H.B. Schlegel, M.A. Robb, E.S. Replogle, R. Gomperts, J.L. Andres, K. Raghavachari, J.S. Binkley, C. Gonzalez, R.L. Martin, D.J. Fox, D.J. DeFrees, J. Baker, J.J.P. Stewart and J.A. Pople, GAUSSIAN 92, Revision B (Gaussian, Inc., Pittsburgh, PA, 1992).
- [15] R.A. Bair, W.A. Goddard III, A.F. Voter, L.G. Rappé, L.G. Yaffe, F.W. Bobrowicz, W.R. Wadt, P.J. Hay and W.J. Hunt, gvb2p5 program (unpublished). See R.A. Bair, PhD Thesis, California Institute of Technology (1980); F.W. Bobrowicz and W.A. Goddard III, in: *Methods of electronic structure theory*, ed. H.F. Schaefer (Plenum Press, New York, 1977) p. 77; W.J. Hunt, P.J. Hay and W.A. Goddard III, *J. Chem. Phys.* 57 (1972) 738.
- [16] M.N. Ringnalda, J.M. Langlois, B.H. Greely, T.V. Russo, R.P. Muller, B. Marten, Y. Won, R.E. Donnelly, Jr., W.T. Pollard, G.H. Miller, W.A. Goddard III and R.A. Friesner, PS-GVB v1.00 (Schrödinger, Inc., 1993).
- [17] P.E.M. Siegbahn, unpublished.
- [18] S. Sæbø, W. Tong and P. Pulay, *J. Chem. Phys.* 98 (1993) 2170.
- [19] J. Almlöf, K. Faegri and K. Korsell, *J. Comput. Chem.* 3 (1982) 385.