

Massively parallel individually selecting configuration interaction[☆]

P. Stampfuß*, K. Hamacher, W. Wenzel

Universität Dortmund, Institut für Physik, D-44221 Dortmund, Germany

Abstract

We report on the implementation and initial applications of an integral-driven algorithm of the configuration-selecting multi-reference configuration interaction method for massively parallel architectures with distributed memory. The transition-residue based matrix element evaluation allows the treatment of Hilbert spaces of 10^{10} determinants, correlating up to 50 electrons. We demonstrate the scalability of the method for up to 128 nodes on the IBM-SP2 and for up to 256 nodes on the CRAY-T3E and report calculations of the vertical excitation energies of benzene and of members of the first row transition metal dihalide family. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Multi-reference configuration interaction; Parallel computing; Correlation effects

1. Introduction

For many years the multi-reference configuration interaction method (MRCI) [1–3] has been one of the benchmark methods for accurate investigations into the electronic structure of molecules that require an adequate treatment of both dynamical and non-dynamical correlation effects. However, because of its high computational cost, applications of MRCI remain constrained to relatively small systems. For this reason the configuration-selective versions of MRCI, such as the A_K [4], CIPSI [5,6] or MRD-CI [7–9] methods, have become its most widely used versions. In this variant only the most important configurations of the interacting space of a given set of primary configurations are chosen for the variational wavefunction [10], while the energy contributions of the remaining configurations are estimated on the basis of second-order Rayleigh–Schrödinger

perturbation theory [4,5]. The generic lack of extensivity of the MRCI method has at least been partially addressed with a number of a posteriori [11,12] corrections and through direct modification of the CI energy-functional [13–18].

Even within this approximation, the cost of MRCI calculations remains rather high. The development of efficient configuration-selecting CI codes [5,6,19, 20,21–24] is inherently complicated by the sparseness and the lack of structure of the selected state-vector. In order to extend further the applicability of the method, it is thus desirable to employ the most powerful computational architectures available for such calculations. Here we briefly report on implementation and benchmark calculations of a recently developed massively parallel implementation of the MRD-CI method for distributed memory architectures. This implementation constructs the subset of non-zero matrix elements using a residue-based representation of the matrix elements that was originally developed for the distributed memory implementation of MRSDCI [25]. This approach allows us to efficiently evaluate the matrix elements both in the expansion loop as

[☆] Presented at the 5th World Congress of Theoretically Oriented Chemists (WATOC), Imperial College, London, 1–6 August, 1999.

* Corresponding author.




Fig. 1. Schematic representation of the two-particle residue-tree. For each element of the configuration-list (A) all possible two-particle residues are constructed. In the configuration illustrated in (B) each box represents one occupied orbital, the shaded region corresponds to the residue and the two white boxes to the orbital pair. The $(n_e - 2)$ -electron residue configuration is looked up in the residue-tree (C), where an element (D) is added that encodes the orbitals that were removed, information regarding the permutation required and the index of the original configuration in the configuration list. Solid arrows in the figure indicate logical relationships, dotted arrows indicate pointers incorporated in the data structure. The residue-list, along with all elements must be rebuilt once after each expansion loop, the effort to do so is proportional to product of n_c^2 with the number of configurations. The number of matrix elements encoded in a single element of the residue-tree is proportional to the *square of the number of entries* of type (D).

well as during the variational improvement of the coefficients of the selected vectors.

We briefly review the general approach that allows an a priori scalable implementation of the matrix element evaluation of the configuration-selecting MRCI method and its integral-driven implementation. We then demonstrate the scalability of the method for up to 128 nodes of an IBM-SP2 and up to 256 nodes of a CRAY-T3E for Hilbert spaces of dimension up to 3×10^9 of which up to 8×10^6 elements were selected for the variational wavefunction. We finally present results for two of our first applications of our object-oriented program [26], the vertical excitation energies of benzene and two members of the transition metal dihalide family, VF_2 and VCl_2 .

2. Methodology

Virtually all the computational effort of the config-

uration selective CI method is concentrated in two steps. First, the many body field

$$q_i = \langle \phi_i | H | \Psi \rangle \quad (1)$$

must be computed for all non-selected configurations $|\phi_i\rangle$ to assess their importance. Here

$$|\Psi\rangle = \sum_j c_j |\phi_j\rangle \quad (2)$$

designates either the set of all previously selected configurations or a suitably chosen reference set. Secondly, matrix elements of the same form as Eq. (1) must be evaluated repeatedly for all selected configurations with respect to the CI vector in the variational subspace to determine its eigenstates. In configuration-selecting CI these operations are complicated by the lack of structure in the selected Hilbert space even on single-processor machines.

Several parallelization strategies have been advocated to implement complex algorithms on distributed memory machines, where *data locality* becomes a paramount issue. One widely used approach is the use of client-server models, where one central node distributes the data among the client nodes on demand. This model is very versatile and has been used for a number of applications. However, in complex algorithms involving large amounts of data, communication bottlenecks can easily arise as the communication patterns vary widely with the size of the active space, the number of electrons and the size of the orbital basis. In our implementation we have therefore chosen an alternate communication scheme, where all operations and data are distributed a priori among the nodes of the machine according to an organizing principle that ensures an equal distribution of the work among the nodes. The advantage of this approach is that the scalability of the algorithm is a mathematical necessity, however, much effort must be devoted to optimize the organizing principle for each particular application of the scheme.

To compute the matrix elements of the Hamilton operator we exploit an enumeration scheme in which each matrix element between two determinants (or configuration state functions) $|\phi_1\rangle$ and $|\phi_2\rangle$ is associated with the subset of orbitals that occur in both the target and the source determinant. This unique subset of orbitals is called the *transition residue* mediating the matrix element and serves as a sorting




Fig. 2. Schematic representation of the computation of two-particle matrix-elements in the expansion step using the residue-tree. For a given configuration (A) we form all two-particle residues, which are looked up in the residue tree. In the configuration illustrated in (A) each box represents one occupied orbital, the shaded region corresponds to the transition residue and the two white boxes to the orbital pair. The $(n_e - 2)$ -electron residue configuration is looked up in the residue-tree (B). Each orbital pair (C) associated with the residue encodes a matrix element with an element of the configuration list (D). The orbital indices of the required integral are encoded in the orbital pairs in (C), the coefficient of the source configuration is looked up directly in (D). Only one lookup operation is required to compute all matrix elements associated with the given transition residue and only the subset of matrix elements that lead to selected source-configurations are constructed.

criterion to facilitate the matrix element evaluation on distributed memory architectures. For a given many-body state, we consider a tree of all possible transition residues as illustrated in Fig. 1. For each such residue we build a list of *residue-entries*, composed of the orbital-pairs (or orbital for a single-particle residue) which combine with the residue to yield a selected configuration and a pointer to that configuration. For configuration-selecting CI the reduction in the number of selected configurations combined with the large total memory of modern distributed memory machines allows us to build the residue tree for the selected configuration, provided that only the required section of the residue tree is stored in the different stages of the computation.

Once the residue tree is available the evaluation of the matrix elements is very efficient. In the *expansion step*, one must evaluate $q_i = \langle \phi_i | HP | \Psi \rangle$, where P projects on the part of the Hilbert space in which only inactive and active and low orbitals are occupied. In a singles–doubles calculation the allowed excita-

tions are directly enumerated on the basis of the internal residues. The information in the tree enables us to immediately compute *all matrix elements* associated with the given residue. As a result the overall numerical effort scales strictly linear with the number of configurations Φ_i for which matrix elements must be evaluated and the number of non-trivial operations per configuration is proportional to n_e^2 .

In the *iteration phase* the full residue tree for all selected configurations must be built, but a single copy of the tree can be distributed across all nodes. All matrix elements associated with a given transition residue can be locally evaluated if the associated orbital pairs are present on a unique node. We note that the residue tree itself (part B in Fig. 2) is not required at all, only the set of connected orbital pairs is needed. As a result no lookup operations are required in this step and one can simply loop over the locally available section of the orbital pair segments to evaluate all matrix elements that can be constructed for the present orbital sets. Since each matrix element is uniquely identified by its transition residue, the contributions to the many-body field can be simply collected at the end of this step on a single node to perform the Davidson iteration. This mechanism allows a rapid evaluation of all matrix elements while using the available core memory to its fullest extent.

We finally note that a local inversion of the accumulated data on the individual nodes allows the formulation of an *integral driven* version of the algorithm. Sorting the locally available data by orbital pairs rather than transition residue allows the selection of all matrix elements that require specific orbital pairs. As a result, when a batch of integrals (in physics notation) is offered to the node, all associate matrix elements can be evaluated directly without further lookup. We have therefore implemented modes of the program in which the integral data is: (a) duplicated on all nodes; (b) split among the nodes; (c) held on disk on the root node; or (d) split among the disks of all nodes. Integral modes (b) and (d) allow the use of very large orbital basis sets.

Going somewhat beyond the standard individually selecting CI, our implementation allows the inclusion of triple and quadruple excitations of the reference configurations. The energy arising from such configurations yields the overwhelming

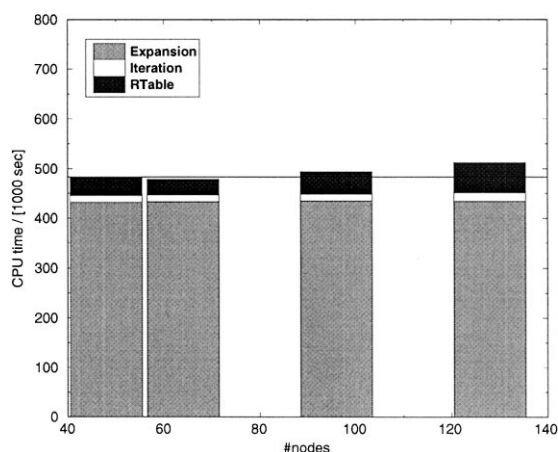


Fig. 3. Total CPU time in seconds for the fully converged calculation of the ground state of the two benchmark calculations described in the text as a function of the number of nodes of the IBM-SP2. A straight line indicates perfect scaling of the computational effort with the number of nodes. The shaded areas in the bars, from top to bottom, indicate the contributions of the matrix element evaluation, logic and the expansion loop.

contribution to the energy difference between FCI and MR-SDCI and is thus of paramount importance for the development approximately extensive versions of the MRCI method [13–16]. Since the number of higher-than-doubly excited configuration rises so quickly with system size, FCI as well as CI-SDTQ calculations are prohibitively expensive for all but the smallest systems. In addition it is possible to modify the treatment of the TQ excitations, such as to provide explicit extensive dressings of the CI matrix elements for incomplete primary spaces. Configuration-selecting CI provides a particularly effective, maybe the only viable, compromise between computational efficiency and accuracy for the treatment of the TQ space.

3. Results

3.1. Scalability

In a truly scalable implementation great care must be taken to divide all work equally across the participating nodes. A remaining non-scalable portion of 1% of the computational effort of a single processor application translates into a 100% overhead if the same

task is distributed across 100 nodes. Our massively parallel algorithm for configuration-selecting MRCI is therefore based on a client-server model that strictly separates the calculation from the communication steps. The latter were chosen to require only global communication directives of the underlying MPI communication library which can be expected to execute efficiently on most modern parallel architectures.

In order to demonstrate the scalability of the implementation we have conducted benchmark, computed the ground state energy of benzene in a cc-pVDZ basis set using active spaces of 6 and 12 active orbitals. The latter calculation was motivated by the desire to test the program for very large Hilbert spaces, but the smaller active space is sufficient to adequately describe e.g. the π - π^* excitations of benzene. The calculation was performed in D_{2h} symmetry at the experimental C_{6v} geometry resulting in Hilbert spaces of up to 3×10^9 determinants of which up to 8×10^6 were selected for the variational subspace.

The most important consideration in the evaluation of the performance of a parallel program is its scalability with the number of processors used for a given calculation. We performed benchmark runs on the 256-node IBM-SP2 of the Karlsruhe supercomputer center. We also tested the program on the on the 256-node and 512-node CRAY-TE3's of the supercomputer center (HLRZ) of the Research Center Jülich and using the maximally available number of processors for standard runs, i.e. 256 on the CRAY-TE3 and 128 on the IBM-SP2, respectively. The runs on the CRAY T3E with its larger data types but smaller core memory per node forced us to use a somewhat smaller threshold than on the IBM-SP2 for the scaling runs in order to be able to finish the calculation even for a small number of processors. Since all data except integrals and state-vector is distributed across the machines the size of the maximally treatable Hilbert space grows significantly with the number of nodes. Unfortunately the T3E consists of two machines of different physical characteristics: the smaller cluster (128 nodes) permits runs ranging from 16 to 64 nodes, the larger one allows runs requiring 65–256 nodes. The interpretation of the scaling data will have to take this “break” into account.

Fig. 3 shows the total computational effort (excluding the time to read the integral file) of the aforementioned scaling runs on the IBM-SP2 as a function of

Table 1

Total CPU times for the benchmark calculations described in the text on the IBM-SP2 and the CRAY T3E as a function of the number of nodes. Given is the time in s/node for the expansion and convergence of a single state in each calculation. The fractional computational loss between two test runs is defined as the ratio of the CPU-times per node divided by the perfect speedup factor given by the ratio of the nodes. The loss-data in the table always refer to successive entries. The calculation for benzene on the IBM-SP2 employed 12 active orbitals that on the CRAY-T3E used a realistic active space of six orbitals. Note that the sensible limit for the latter calculations lies around 64 nodes, where less than 5 min are required to converge the calculation

Number of nodes	Time (s/node)	Loss (%)
IBM-SP2 12 active orbitals		
16		
32		
48	10 122	
64	7510	0
96	5160	3
128	4012	4
CRAY-T3E 6 active orbitals		
16		
32	319	
64	172	8
65	244	
128	147	20
256		

Table 2

Vertical π - π^* excitation energies of benzene in configuration-selecting CI in cc-pVDZ basis in a 6 active orbital CAS space in comparison with experimental results (data as in Ref. [28] and multi-reference perturbation theory. State-averaged approximate natural orbital from a BW-MRPT calculation were used for all states of one symmetry, the calculations were performed in D_{2h} symmetry at the experimental C_{6v} geometry

D_{6h}	D_{2h}	MRD-CI	EXP	MRMP	CASPT2
Singlets					
1^1B_{2u}	1^1B_{3u}	4.93 (+0.03)	4.90	4.71	4.58
1^1B_{1u}	1^1B_{2u}	6.30 (+0.10)	6.20	5.83	5.89
1^1E_{1u}	2^1B_{2u}	7.16 (+0.18)	6.94	6.93	6.52
1^1E_{2g}	1^1B_{1g}	7.45 (-0.35)	7.80	7.74	7.68
Triplets					
1^3B_{1u}	1^3B_{2u}	4.01 (+0.06)	3.95	3.90	3.85
1^3E_{1u}	1^3B_{3u}	4.78 (+0.02)	4.76	4.44	4.38
1^3B_{1u}	1^3B_{2u}	5.40 (-0.20)	5.60	5.07	5.22
1^3E_{2g}	1^3A_g	7.13 (+0.30)	6.83	6.96	6.90

the number of nodes. In these plots, the computational effort for all logic-steps sections are subsumed in one category, the expansion loop and the iteration loop constitute the other main components of the program.

We find almost perfect scaling from 48 to 128 nodes for the IBM-SP2. The total computational effort in the expansion loop, which dominates the overall computational effort, is constant to within 0.4% in going from the smallest to the largest number of nodes. In contrast, the effort associated with logic and communication grows somewhat with the number of nodes. This is to be expected, since the communication cost grows with the number of nodes and a total of 3.7/9.1 GB of data have to be transmitted across the machine for the small and large residue tables, respectively. The overall speedup factor from 64 to 128 nodes is 1.86 (see (see Table 1)). On the CRAY-T3E we find a similar situation: For calculations with the more realistic six-orbital this calculation requires less than three minutes total turnaround time on 128 nodes and the residue table in the iteration step is spread so thinly that it becomes impossible to balance. This explains the somewhat large loss of 20% efficiency in going from 65 to 128 nodes. Note that the data for 64 (65) nodes were obtained on the small (large) cluster of the T3E described above. The time differences are indicative of the relative performance of these two machines.

3.2. Applications

The above data demonstrate that the residue-driven matrix element evaluation yields an efficient and scalable algorithm on distributed memory architectures. Unfortunately our program has not yet been integrated into a complete massively parallel quantum chemistry code. As a result integral and orbital generation for the benchmark tests, as well as for the two test applications we will report here, were performed on single node workstations and only the resulting integral files were sent to the supercomputer centers for the MRCI runs. This cumbersome procedure still complicates the application of the program, the aforementioned integration of the configuration-selecting MRCI module into a complete package is presently under way.

Here we report the results of three calculations of vertical excitation energies. To demonstrate capabilities and limitations of the algorithm we have

Table 3

Vertical excitation energies of VX_2 ($X = F, Cl$) in $D_{2\infty}$ geometry with $R = 176 \text{ \AA}$ and $r = 1.81 \text{ \AA}$ in a DZP ANO basis in configuration-selecting CI in comparison with adiabatic term energies in DFT [50]. All energies (in eV) are relative to the respective ground-state assignments of these calculations, indicated by 0.00, respectively

	MRD-CI T_v	DFT T_c
VF_2		
$^4\Pi_g$	0.00	1.12
$^4\Sigma_g$	0.05	0.00
$^2\Delta_g$	0.94	2.20
$^2\Pi_g$	1.83	4.17
VCl_2		
$^4\Pi_g$	0.00	0.81
$^4\Sigma_g$	0.29	0.00
$^2\Delta_g$	1.01	1.82
$^2\Pi_g$	2.11	3.35

computed the vertical $\pi-\pi^*$ excitation energies of benzene as a suitable reference molecule. Table 2 summarizes the results of our calculation in a cc-pVDZ basis [27] using a six-electron active space in comparison with multi-reference perturbation theory calculations [28,29] using 12 orbital active spaces. MRD-CI by virtue of the individual selection of the most important configurations should be able to variationally treat the most important configuration that cannot be treated adequately in perturbation theory, while diagrammatic MRPT has to resort to using large active spaces. This expectation is fulfilled for the lower excitations in the spectrum where we obtained very good agreement with experimental observations, but the quality of the agreement degrades for higher excitation energies. This may be in part due to the partial loss of full orthogonality of the wavefunctions of the higher excitations, aggravated by the fact that our determinant based program computes both singlets and triplets in a single calculation. Comparison with the spin-adapted version of the program will allow us to evaluate this source of error. A further source of error may be the use of approximate state-averaged natural orbitals [30] as opposed to CASSCF orbital in the calculations reported here.

As a second application we have computed the low-lying vertical excitation spectrum of two model compounds of the family of transition metal dihalides.

The accurate quantitative description of the electronic structure and bonding of many transition metal compounds remains a challenge even to modern quantum chemical methods. The family of the first row transition metal dihalides (MX_2 , $M = Sc, \dots Zn$, and $X = F, Cl, Br$) represent a simple model series that nevertheless exhibits key problems associated with the theoretical description of transition metals in their chemical environment. Because of their interesting properties many experimental studies have been undertaken to characterize this family [31–41]. Despite equally active theoretical interest [42–50] many questions regarding their theoretical description remain unresolved. A number of theoretical investigations that at least partially account for the orbital interaction and the electron correlation effects [43,48] have cast the assignments of the spectra based on ligand field theory (LFT) in doubt. A recent broad-based investigation of this series using density functional theory (DFT) [50] supported some of these reassignments but showed significant qualitative and quantitative discrepancies to experiment and prior theoretical investigations for bond-lengths, dissociation and excitation energies. As a result we have computed the low-lying vertical excitations for two model members of this family, VF_2 and VCl_2 at their experimental geometries. We used a DZP quality ANO basis set [51] and state-averaged approximate MR-BWPT natural orbitals as described above for the benzene calculation. Table 3 summarizes our results for the *vertical* excitation energies in comparison to the *adiabatic* term energies in DFT [50]. Although not directly comparable, the data shows that both the assignment of the ground state as well as that of excited states in MRCI is at variance with that of DFT. While the energy difference between the $^4\Pi_g$ and the $^4\Sigma_g$ in VF_2 are close to the resolution of the selecting CI method, the results for VCl_2 clearly indicate a discrepancy in the ground state assignment at this level of theory. These results are of particular interest, as for VF_2 and VCl_2 they reverse the assignment of the ground states by DFT to the original LFT proposition. However, further calculations in larger basis sets and the possible influence of relativistic effects should be considered. Such calculations, along with the investigations of other members of the MX_2 family, are presently under way.

4. Discussion

Since the computational effort of accurate benchmark techniques, such as MRCI and its variants, rises rapidly with the molecular size, only the use of the most powerful computational architectures ensures their continued relevance to the field. Because massively parallel architectures with distributed memory will yield the highest computational throughput in the foreseeable future, it is worthwhile to pursue the use of these machines for quantum chemical benchmark calculations. We hope that the development of the scalable implementation of one of the most popular variants of the MRCI method family on such architectures is a useful step in this direction. The present implementation allows the treatment of Hilbert spaces and systems that are larger than those that can be treated on traditional architectures, while significantly reducing the turnaround time for more moderate applications. With the ability to routinely treat Hilbert space exceeding 10 billion determinants many questions that require a delicate balance of dynamical and non-dynamical correlation effects, e.g. in transition metal chemistry, become amenable to the MRCI method.

A spin-adapted implementation of the residue-based MRCI algorithm, both for MR-SDCI [25] and its configuration-selecting variant is presently under way. In addition, it is worthwhile to investigate approximations, such as multi-reference second-order Brillouin–Wigner perturbation theory [4,52,53], that eliminate the selected variational subspace in MRD-CI altogether. We note that the selection step in MRD-CI scales with $n_c^2 N^2$ (where N is the number of orbitals), as opposed to the iteration step which principally scales as $n_c^2 N^4$. A variant of MR-BWPT was already used in this work to generate approximate natural orbitals for the excited state calculations [30]. In combining these techniques we hope to be able to partition a large part of the orbital for exclusive treatment in BW-MRPT, similar to the CIPSI [6] method. This in turn allows the selection of only the required excitation integrals in a partial four-index transformation and ultimately the use of an integral-direct algorithm.

Acknowledgements

This work was supported by the DFG Grant KEI-

164/11-2 and computational resources at the HLRZ Jülich and the HRZ Karlsruhe.

References

- [1] B.O. Roos, Chem. Phys. Lett. 15 (1972) 153.
- [2] B.O. Roos, P.E.M. Siegbahn, The direct configuration interaction method, in: H.F. Schaefer III (Ed.), *Methods of Electronic Structure Theory*, Plenum, New York, 1994, p. 189.
- [3] I. Shavitt, in: H.F. Schaefer III (Ed.), *Modern Theoretical Chemistry*, Plenum, New York, 1977.
- [4] Z. Gershgorin, I. Shavitt, Int. J. Quantum Chem. 2 (1968) 751.
- [5] B. Huron, J. Malrieu, P. Rancurel, J. Chem. Phys. 58 (1973) 5745.
- [6] C. Agnelli, M. Persico, Theor. Chim. Acta 98 (1998) 117.
- [7] R.J. Bunker, S. Peyerimhoff, Theor. Chim. Acta 12 (1968) 183.
- [8] R.J. Bunker, S.D. Peyerimhoff, Theor. Chim. Acta 35 (1974) 33.
- [9] R.J. Bunker, S.D. Peyerimhoff, Theor. Chim. Acta 39 (1975) 217.
- [10] F. Illas, J. Rubio, J.M. Ricart, P.S. Bagus, J. Chem. Phys. 95 (1991) 1877.
- [11] S.R. Langhoff, E.R. Davidson, Int. J. Quantum Chem. 8 (1974) 61.
- [12] W. Butscher, S. Shih, R.J. Bunker, S.D. Peyerimhoff, Chem. Phys. Lett. 52 (1977) 457.
- [13] J. Cizek, J. Chem. Phys. 45 (1966) 4256.
- [14] R.J. Bartlett, I. Shavitt, Chem. Phys. Lett. 50 (1977) 190.
- [15] R. Gdanitz, R. Ahlrichs, Chem. Phys. Lett. 143 (1988) 413.
- [16] P. Szalay, R.J. Bartlett, J. Chem. Phys. 103 (1995) 3600.
- [17] J.P. Daudey, J.-L. Heully, J.P. Malrieu, J. Chem. Phys. 99 (1993) 1993.
- [18] J.P. Daudey, J.-L. Heully, J.P. Malrieu, J. Chem. Phys. 99 (1993) 1993.
- [19] R.J. Bunker, S.D. Peyerimhoff, *New Horizons in Quantum Chemistry*, Reidel, Dordrecht, 1983.
- [20] J.L. Whitten, M. Hackmeyer, J. Chem. Phys. 51 (1969) 5548.
- [21] S. Evangelisti, J.P. Daudey, J.P. Malrieu, Chem. Phys. 75 (1983) 91.
- [22] R.J. Harrison, J. Chem. Phys. 94 (1991) 5021.
- [23] S. Krebs, R.J. Bunker, J. Chem. Phys. 103 (1995) 5613.
- [24] M. Hanrath, B. Engels, New algorithms for an individually selecting mr-ci program, CP 225 (1997) 197.
- [25] F. Stephan, W. Wenzel, J. Chem. Phys. 108 (1998) 1015.
- [26] M.M. Steiner, W. Wenzel, J.W. Wilkins, K.G. Wilson, *Computers Phys.* 11 (1997) 467.
- [27] D.E. Woon, T.H. Dunning, J. Chem. Phys. 98 (1993) 1358.
- [28] B.O. Roos, K. Andersson, M.P. Fülscher, Chem. Phys. Lett. 5 (1992) 192.
- [29] T. Hashimoto, S. Nakano, Hirao, J. Chem. Phys. 104 (1996) 6244.
- [30] W. Wenzel, M.M. Steiner, J. Chem. Phys. 108 (1998) 4714.
- [31] R.A. Kent, J.D. McDonald, J.L. Margrave, JPC 70 (1960) 874.
- [32] C.E. Moore, *Atomic Energy Levels*, NSRDS-NBS#35, vol. 2, US GPO, Washington, DC, 1971.

- [33] F. Ramondo, V. Rossi, L. Bencivenni, *J. Mol. Struct.* 64 (1988) 513.
- [34] F. Ramondo, L. Bencivenni, S.N. Cesaro, *J. Mol. Struct.* 192 (1989) 83.
- [35] V.P. Spiridonov, A.G. Gershikov, V.S. Lyutsarev, *J. Mol. Struct.* 79 (1990) 220.
- [36] O.V. Boltalina, A.Y. Borshchevskii, L.N. Sidorov, *Russ. J. Phys. Chem.* 65 (1991) 466.
- [37] O.V. Boltalina, A.Y. Borshchevskii, L.N. Sidorov, *Russ. J. Phys. Chem.* 66 (1992) 1223.
- [38] A.J. Ross, R. Bacis, A.J. Bouvier, S. Churassy, J.C. Coste, P. Crozet, I. Russier, *J. Mol. Spectrosc.* 158 (1993) 27.
- [39] M.P. Barnes, R.T. Carter, N.M. Lakin, J.M. Brown, *J. Chem. Soc. Faraday Trans.* 89 (1993) 3205.
- [40] D.L. Hildenbrand, *JPC* 103 (1995) 2634.
- [41] D.L. Hildenbrand, *High Temp. Mater. Sci.* 35 (1996) 151.
- [42] S.Y. Shashkin, W.A. Goddard, *JPC* 90 (1986) 225.
- [43] C.W. Bauschlicher, B.O. Roos, *J. Chem. Phys.* 91 (1989) 4785.
- [44] M. Kaupp, P.R. Schleyer, H. Stroll, H. Preuss, *JACS* 113 (1991) 6012.
- [45] R.J. Deeth, *J. Chem. Soc. Dalton Trans.* (1993) 1061.
- [46] H. Chermette, F. Rogemond, D.R. Salahub, *CPL* 219 (1994) 228.
- [47] M. Hassett, C.J. Marsden, *J. Mol. Struct.* 346 (1995) 249.
- [48] B.O. Roos, K. Andersson, M.P. Fülcher, P.-A. Malmqvist, L. Serrano-Andres, K. Pierloot, M. M'erchan, *Adv. Chem. Phys.* 93 (1996) 219.
- [49] D.R. Salahub, C. Blanchet, H.A. Durante, *J. Chem. Phys.* 106 (1997) 8778.
- [50] S.G. Wang, W.H.E. Schwartz, *J. Chem. Phys.* 109 (1998) 7252.
- [51] R. Pou-Amerigo, M. Merchan, I. Nebot-Gil, P.O. Widmark, B. Roos, *Theor. Chim. Acta* 92 (1995) 149.
- [52] W. Wenzel, K.G. Wilson, *Phys. Rev. Lett.* 68 (1992) 800.
- [53] W. Wenzel, *Int. J. Quantum Chem.* 70 (1998) 613.