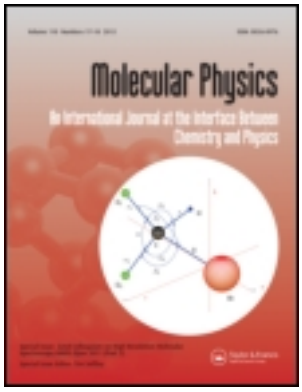
This article was downloaded by: [KIT Library]
On: 29 November 2013, At: 02:37
Publisher: Taylor & Francis
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Molecular Physics: An International Journal at the Interface Between Chemistry and Physics
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/tmph20

A scaling PNO-MP2 method using a hybrid OSV-PNO approach with an iterative direct generation of OSVs†

Gunnar Schmitz [a], Benjamin Helmich [a] & Christof Hättig [a]
[a] Lehrstuhl für Theoretische Chemie, Ruhr Universität Bochum, Germany
Accepted author version posted online: 29 Apr 2013.Published online: 08 May 2013.

To cite this article: Gunnar Schmitz, Benjamin Helmich & Christof Hättig (2013) A scaling PNO-MP2 method using a hybrid OSV-PNO approach with an iterative direct generation of OSVs†, Molecular Physics: An International Journal at the Interface Between Chemistry and Physics, 111:16-17, 2463-2476, DOI: 10.1080/00268976.2013.794314

To link to this article: http://dx.doi.org/10.1080/00268976.2013.794314

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at http://www.tandfonline.com/page/terms-and-conditions

Molecular Physics, 2013
Vol. 111, Nos. 16–17, 2463–2476, http://dx.doi.org/10.1080/00268976.2013.794314

Taylor & Francis
Taylor & Francis Group

Downloaded by [KIT Library] at 02:37 29 November 2013

INVITED ARTICLE

A $\mathcal{O}(\mathcal{N}^3)$ scaling PNO–MP2 method using a hybrid OSV–PNO approach with an iterative direct generation of OSVs†

Gunnar Schmitz*, Benjamin Helmich and Christof Hättig

Lehrstuhl für Theoretische Chemie, Ruhr Universität Bochum, Germany

(*Received 5 February 2013; final version received 28 March 2013*)

We present an implementation of pair natural orbital second-order Møller–Plesset perturbation theory with computational costs that scale only cubically with the system size. The low cost-scaling is achieved by combining a hybrid approach, where the pair natural orbitals are build from orbital-specific virtuals OSVs, with an iterative block Davidson algorithm for solving the equations for OSVs. We thereby avoid a complete diagonalisation of amplitude matrices and the explicit construction of the corresponding exchange integral matrices. This reduces the cost-scaling for the generation of the OSVs and of the pair natural orbitals to $\mathcal{O}(\mathcal{N}^3)$ without *a priori* assumptions about locality. The costs can be further reduced by combining the approach with a local resolution-of-the-identity approximation for the exchange integrals. The errors introduced by these approximations are negligible and do not affect the final accuracy of the correlation energy. Test calculations on a set of organic and inorganic molecules demonstrate that the additional errors are at least one order of magnitude smaller than the truncation error for the pair natural orbital space and that the accuracies for all steps can be controlled by a single threshold parameter. Calculations on glycine chains $(Gly)_n$ with $n = 1, 2, 4, 8, 16, 32$ in the aug-cc-pVTZ basis set reveal an early break even point between the $\mathcal{O}(\mathcal{N}^3)$-scaling and the non-iterative $\mathcal{O}(\mathcal{N}^4)$-scaling implementation that uses a full diagonalisation and with canonical RI-MP2. We present applications to systems, which are of interest in the field of nano machines: a C_{60} fullerene acceptor system with over 8000 basis functions and a foldamer with nearly 8000 basis functions.

Keywords: local MP2; hybrid OSV–PNO approach; iterative OSV generation; local RI

1. Introduction

The enormous development in the field of computer techniques opened the possibility to study increasingly larger systems with quantum chemical methods. Density functional theory (DFT) is today the most widely applied first-principle approach for systems up to 1000 atoms, while wavefunction-based methods are mostly used for more accurate calculations on smaller systems, since DFT provides for relatively low computational costs an accuracy that is often sufficient. Most implementations of semi-local or hybrid functionals, as e.g. PBE [1,2] or B3-LYP [3–6], exhibit computational costs, which scale already for small molecules only with $\mathcal{O}(\mathcal{N}^4)$, where $\mathcal{N}$ is a measure for the system size. For hybrid functionals the effective cost-scaling reduces as for Hartree–Fock only slowly with the system size to a cubic and then for very large systems to a quadratic scaling. Pure semi-local functionals can be combined [7–9] with the resolution-of-the-identity (RI) approximation for the electron repulsion integrals (ERIs) to reduce already the initial cost-scaling to $\mathcal{O}(\mathcal{N}^3)$, which decreases for larger systems quickly to a quadratic or even linear scaling [10]. A problem of DFT methods is that there are only limited possibilities for systematic improvement and that semi-local functionals cannot describe dispersion interactions in a first-principle manner. In contrast to DFT, wavefunction-based methods provide hierarchies as for example the coupled-cluster (CC) and the configuration interaction (CI) approaches, which converge systematically towards the exact solution of the electronic Schrödinger equation. Already second-order Møller–Plesset perturbation theory (MP2), as one of the simplest correlated wavefunction methods, describes dispersion interactions without empirical corrections. Although the costs for MP2 calculations scale as $\mathcal{O}(\mathcal{N}^5)$ with the system size, they are for small and medium-sized molecules relatively low due to a small prefactor. This inspired during the last years also the development of so-called double hybrid density functionals [11,12], which incorporate an MP2-like non-local and orbital-dependent contribution into the exchange-correlation functional. For large systems the computational costs for such functionals will, however, exhibit the same steep scaling as for MP2.

Thus, and since the MP2 contributions occur also in almost any higher-order correlated wavefunction method, it is desirable to develop implementations of MP2 and other wavefunction-based correlation methods, which have

*Corresponding author. Email: Gunnar.Schmitz@rub.de

†Dedicated to Prof. Dr. Werner Kutzelnigg on the occasion of his 80th birthday.

© 2013 Taylor & Francis

Downloaded by [KIT Library] at 02:37 29 November 2013

a lower cost-scaling than the canonical implementations. Many groups are currently working on this field. To achieve a lower scaling they generally exploit the short-ranged nature of dynamical electron correlation to discard negligible contributions or to truncate the wavefunction expansion. For example Ochsenfeld and coworkers successfully developed an approach [13,14], which goes back to the work of Almlöf and Häser [15] and allows the formulation of MP2 in the AO basis. Using refined integral estimates they could devise a linear-scaling AO–MP2 method [13,14]. Werner and coworkers [16] and Head-Gordon *et al.* [17] demonstrated how a linear cost-scaling can be achieved for MP2 and also for several CC methods [18–20]. They combined localised occupied orbitals (LMOs) with projected atomic orbitals (PAOs) as basis for the virtual space; an idea, which goes back to the work of Saebø and Pulay [21,22]. The approach uses, however, correlation domains, which are defined *a priori* without incorporating information on the correlation contribution to the wavefunction, which can limit the error control. Another possibility to represent the virtual orbital space goes back to work of Meyer [23,24], Ahlrichs, Lischka, Staemmler, and Kutzelnigg [25–27] and involves pair natural orbitals (PNOs). The PNO approach has recently been revived by Neese and his coworkers [28] and since then attracted also the interest of other groups [29–32]. PNOs are a linear combination of virtual orbitals, which diagonalise the correlation contribution of the pair *ij* to a one-electron density matrix. Thereby, the PNOs adapt automatically to the requirements for the wavefunction. The error introduced by truncating the PNO expansion is controlled by a single parameter, the threshold for occupation numbers T_{PNO}. The truncated PNO expansion retains nearly smooth potential energy surfaces [28], which is important for the calculation of properties, which do not depend on the energy itself, but on its derivatives. Closely related to PNOs are orbital specific virtuals (OSVs), which have recently been proposed by Yang *et al.* [29,33] and are specific for one occupied orbital *i* and not for a pair *ij*. One possible choice of OSVs are the PNOs for the diagonal pairs *ii*. Recently, it was shown [30,34] that the OSV and PNO approaches can be combined to a hybrid OSV-PNO approach, which avoids the steep $\mathcal{O}(\mathcal{N}^5)$-scaling steps in the construction of PNOs, but retains the compact representation of the double excitation space from the PNO approach and leads to an efficient $\mathcal{O}(\mathcal{N}^4)$-scaling MP2 implementation. The generation of OSVs is an $\mathcal{O}(\mathcal{N}^4)$ step and dominates the overall computational costs. It involves the construction and diagonalisation of the doubles amplitudes t_{ab}^{ii} for the diagonal pairs.

More recently, Riplinger *et al.* [35] implemented a very efficient near linear scaling PAO–PNO–CCSD method, which combines the PAO, OSV and PNO approaches. In a first step negligible electron pairs are screened out, similar as in Ref. [34], based on a cheap estimate for the MP2 pair correlation energies. For all surviving pairs the PNOs are expanded in a pair-specific PAO basis. While very efficient for the description of dynamic electron correlation in the ground state wavefunction, PAOs introduce an *a priori* assumption on the locality of the amplitudes and are problematic, where one needs an unbiased treatment of local and non-local effects. This is for example the case in excitation energy calculations with response methods. Here Helmich and Höttig [32] have recently shown that the PNO representation can be used to obtain a very compact representation of the excitation amplitudes even though they are strongly non-local for charge-transfer states.

In this work we demonstrate how the scaling of a hybrid OSV–PNO implementation can be reduced by replacing the full diagonalisation of the amplitude matrices t_{ab}^{ii} for the OSV generation by an iterative direct diagonalisation, which avoids an explicit construction of the amplitudes t_{ab}^{ii}. Furthermore, a local RI approximation is introduced to reduce the prefactor. The paper is organised as follows: first the theory of PNOs, OSVs, and PNO–MP2 is briefly reviewed. In Section 3 the implementation of the iterative direct OSV construction is described to show how the $\mathcal{O}(\mathcal{N}^3)$-scaling is achieved. In Section 4, we analyse the accuracy of the approach in comparison to PNO–MP2 and the performance of the implementation in comparison to OSV–PNO–MP2 with a full diagonalisation and RI–MP2. Illustrative calculations demonstrate the scope of our new implementation including calculations with more than 8000 basis functions, which would be difficult to perform with canonical RI–MP2. Finally, we summarise in Section 5 our findings and conclusions and comment on future perspectives.

2. Theory

2.1. Pair natural orbitals

Pair natural orbitals (in the following denoted by overbars) are, as mentioned above, linear combinations of virtual orbitals $|a\rangle$

$$|\bar{a}\rangle = \sum_a |a\rangle d_{a\bar{a}}^{ij}, \tag{1}$$

chosen such that they diagonalise the contribution of the pair *ij* to the virtual-virtual block of a one-electron density matrix D_{ab}^{ij}

$$\sum_{ab} d_{a\bar{a}}^{ij} D_{ab}^{ij} d_{b\bar{b}}^{ij} = n_{\bar{a}}^{ij} \delta_{\bar{a}\bar{b}}. \tag{2}$$

The eigenvalues can be interpreted as occupation numbers $n_{\bar{a}}^{ij}$. For MP2 the density matrices are, in a spin-free closed-shell formalism, obtained as:

$$D_{ab}^{ij} = 2 \sum_c \left(u_{ac}^{ij} t_{bc}^{ij} + u_{ca}^{ij} t_{cb}^{ij} \right) \tag{3}$$

Downloaded by [KIT Library] at 02:37 29 November 2013

with

$$t_{ab}^{ij} = -\frac{(ia|jb)}{\varepsilon_a - \varepsilon_i + \varepsilon_b - \varepsilon_j} \tag{4}$$

and $u_{ab}^{ij} = 2t_{ab}^{ij} - t_{ba}^{ij}$. Here and in the following, ε_p denote orbital energies for canonical Hartree–Fock orbitals. Indices i, j are used for occupied orbitals and a, b for virtual orbitals. From Equation (2) it is obvious that PNOs for a pair ij are orthogonal to each other, which is, however, in general not the case for PNOs of different pairs. The respective overlap matrices are obtained by

$$S_{\bar{a}\bar{b}}^{ij,kl} = \sum_c d_{c\bar{a}}^{ij} d_{c\bar{b}}^{kl}. \tag{5}$$

With localised occupied orbitals only the PNOs located in the same region of space as the pair of occupied orbitals ij have large occupation numbers [23,24] and the occupation numbers converge rapidly to zero if sorted in descending order. Truncating the PNO spaces to orbitals with occupation numbers above a given threshold T_{PNO} leads to a very compact parameterisation of the double excitation space.

2.2. MP2 in a PNO formalism

For a closed-shell system the MP2 correlation energy can be written as:

$$E_{\mathrm{MP2}} = \sum_{i \le j} \frac{2}{1+\delta_{ij}} \sum_{ab} L_{ab}^{ij} t_{ab}^{ij} \tag{6}$$

with $L_{ab}^{ij} = 2K_{ab}^{ij} - K_{ba}^{ij}$ where K_{ab}^{ij} are the exchange integral matrices $K_{ab}^{ij} = (ia|jb)$. Restricting the double excitations to the PNO spaces leads to

$$E_{\text{PNO-MP2}} = \sum_{i \le j} \frac{2}{1+\delta_{ij}} \sum_{\bar{a}\bar{b}} L_{\bar{a}\bar{b}}^{ij} t_{\bar{a}\bar{b}}^{ij}, \tag{7}$$

where, with LMOs the amplitudes have to be determined as the solution of the linear equation

$$\begin{aligned} 0 = K_{\bar{a}\bar{b}}^{ij} &+ \sum_{\bar{c}} \left(f_{\bar{a}\bar{c}} t_{\bar{c}\bar{b}}^{ij} + t_{\bar{a}\bar{c}}^{ij} f_{\bar{c}\bar{b}} \right) \\ &- \sum_k \sum_{\bar{c}\bar{d}} \left(f_{ik} S_{\bar{a}\bar{c}}^{ij,kj} t_{\bar{c}\bar{d}}^{kj} S_{\bar{d}\bar{b}}^{kj,ij} + S_{\bar{a}\bar{c}}^{ij,ik} t_{\bar{c}\bar{d}}^{ik} S_{\bar{d}\bar{b}}^{ik,ij} f_{kj} \right). \end{aligned} \tag{8}$$

2.3. Orbital specific virtuals

Closely related to PNOs are OSVs, which were introduced by Yang *et al.* [29] and are specific for individual occupied orbitals i. One possible choice for OSVs are the PNOs for the diagonal pairs ii:

$$D_{ab}^{ii} = 4 \sum_c t_{ac}^{ii} t_{cb}^{ii} = \sum_{\tilde{c}} d_{a\tilde{c}}^{ii} n_{\tilde{c}}^{ii} d_{b\tilde{c}}^{ii}. \tag{9}$$

Here and in the following we denote OSVs with a tilde. Instead of D_{ab}^{ii} one can just diagonalise t_{ab}^{ii}, which has the same eigenvectors and eigenvalues given by $\tau_{\tilde{a}_i} = \frac{1}{2}\sqrt{n_{\tilde{a}}^{ii}}$. Similar as for PNOs, also the OSV expansion can be truncated by discarding orbitals with occupation numbers $n_{\tilde{c}}^{ii}$ below a threshold T_{OSV}. The virtual space for off-diagonal pairs ij is in the OSV approach constructed by merging together the OSV spaces for the orbitals i and j. This leads to a double excitation space, which is much less compact than with PNOs. But the computational overhead for the construction of OSVs is much smaller than for PNOs. The dominating step is the computation of the exchange integrals K_{ab}^{ii} needed for t_{ab}^{ii}, which are evaluated within the RI approximation:

$$K_{ab}^{ii} = \sum_Q B_{Q,ai} B_{Q,bi}, \tag{10}$$

$$B_{Q,ai} = \sum_P [\mathbf{V}^{-1/2}]_{QP} (Q|ai), \tag{11}$$

and features an $\mathcal{O}(ON^3)$-scaling if no integral screening or local restrictions of the auxiliary basis is applied. Here O and N refer to the number of occupied orbitals and the number of basis functions, respectively. For a detailed discussion of the relation between OSVs and PNOs the reader is referred to Refs. [30,33,34].

2.4. Hybrid OSV–PNO approach

In a previous work [34], it was demonstrated that the advantages of the OSV and PNO expansions can be combined by using the double excitation space obtained with the truncated OSV expansion for the computation of the density matrices for the PNOs. If the OSV selection threshold, T_{OSV}, is chosen one order of magnitude smaller than the PNO threshold T_{PNO}, the errors introduced by this approximation do not to affect the accuracy in comparison to a full PNO–MP2 approach.

In a first step we compute a SOS–MP2-like estimate for the pair correlation energies in the OSV basis

$$\mathcal{E}_{ij} \approx c_{OS} \left(1+\delta_{ij}\right) \sum_{\tilde{a}_i \tilde{b}_j} \frac{2K_{\tilde{a}_i\tilde{b}_j}^{ij} K_{\tilde{a}_i\tilde{b}_j}^{ij}}{f_{ii} - f_{\tilde{a}_i\tilde{a}_i} + f_{jj} - f_{\tilde{b}_j\tilde{b}_j}} \tag{12}$$

with the diagonal elements of the Fock matrix in the LMO f_{ii} and OSV $f_{\tilde{a}_i\tilde{a}_i}$ basis. To compute the exchange integrals $K_{\tilde{a}_i\tilde{b}_j}^{ij}$ in the OSV basis, the intermediates $B_{Q,ai}$ are transformed to the OSV basis and then combined to $K_{\tilde{a}_i\tilde{b}_j}^{ij}$ with

Downloaded by [KIT Library] at 02:37 29 November 2013

$\mathcal{O}(\mathcal{N}^3)$-scaling costs. Test calculations revealed that with $c_{\mathrm{OS}} = 1.5$ the estimates from Equation (12) agree typically within a factor of 2 with the estimates for the pair correlation energy in the PNO basis, which have been proposed in Ref. [28] for screening the pair list *ij*. In the hybrid OSV–PNO approach, we thus discard already at this stage all pairs with $|\mathcal{E}_{ij}| < 0.5 \cdot T_{\mathrm{pair}}$. If T_{pair} is set such that only pairs are discarded, which cannot have PNOs with eigenvalues above T_{PNO}, this screening does not affect the final result, but leads for large systems to a linearly scaling number of significant pairs *ij*. For these pairs *ij* orthonormal and quasicanonical pair specific basis sets (in the following called pre-PNOs) $\tilde{a}_{ij}$ are build similar as in the pure OSV approach [29]. First, the OSV sets of *i* and *j* are merged and orthonormalised against each other. Then, the Fock matrix is set up in this basis and diagonalised to obtain a quasicanonical basis $\tilde{a}_{ij}$ with $f_{\tilde{a}_{ij},\tilde{b}_{ij}} = \delta_{\tilde{a}_{ij},\tilde{b}_{ij}} \varepsilon_{\tilde{a}_{ij}}$. Afterwards, the PNOs are generated analogously to the canonical case by diagonalising the density matrices:

$$D^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}} = 2 \sum_{\tilde{c}_{ij}} \left(u^{ij}_{\tilde{a}_{ij}\tilde{c}_{ij}} t^{ij}_{\tilde{b}_{ij}\tilde{c}_{ij}} + u^{ij}_{\tilde{c}_{ij}\tilde{a}_{ij}} t^{ij}_{\tilde{c}_{ij}\tilde{b}_{ij}} \right), \tag{13}$$

expressed with doubles amplitudes in the pre-PNO basis:

$$t^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}} = \frac{K^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}}}{f_{ii} - \varepsilon_{\tilde{a}_{ij}} + f_{jj} - \varepsilon_{\tilde{b}_{ij}}}. \tag{14}$$

The exchange integrals $K^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}}$ are obtained similar as in Equation (10) from three-index intermediates $B_{Q,\tilde{a}_{ij}i}$. The calculation of these three-index intermediates is, after the computation of the OSVs, the second most expensive step for the construction of PNOs in the hybrid OSV–PNO approach. In the regime of a linear scaling number of significant pairs *ij*, this step scales only cubically, however, exhibits a large prefactor. Altogether, the hybrid OSV–PNO approach as discussed above has an $\mathcal{O}(\mathcal{N}^4)$ scaling, which is also achieved with the OSV–MP2 method proposed by Yang *et al.* [33], but it retains the very compact representation of the double excitation space of the PNO methods. This reduces the costs of the iterative solution of the MP2 Equation (8) and is also important for higher order methods.

3. Implementation

To achieve a scaling not higher than $\mathcal{O}(\mathcal{N}^3)$ in the OSV generation, it is necessary to avoid a full construction of the matrices with the amplitudes t^{ii}_{ab} and integrals K^{ii}_{ab} in a canonical virtual basis. One way to achieve this goal is to use an iterative diagonalisation method, which does not require the full amplitude matrices to compute the eigenpairs with the largest eigenvalues, only an implementation of matrix vector products of the form $\sigma^k = \mathbf{T}\mathbf{b}^k$. To solve the OSV eigenvalue problem, we developed a modified block Davidson method, which is described in 3.2.

3.1. Computation of the matrix–vector products

To avoid in the matrix vector products the evaluation of the integral matrices K^{ii}_{ab}, we first need to rewrite the expression for the amplitudes, which in the canonical basis and within the RI approximation are obtained as:

$$\begin{aligned} t^{ii}_{ab} &= -\frac{K^{ii}_{ab}}{\varepsilon_a - \varepsilon_i + \varepsilon_b - \varepsilon_i} \\ &= -\frac{\sum\limits_{P_i Q_i} (ia|Q_i) \left[{}^{(ii)}V^{-1}\right]_{P_i Q_i} (ib|P_i)}{\varepsilon_a - \varepsilon_i + \varepsilon_b - \varepsilon_i}. \end{aligned} \tag{15}$$

In the last equation, P_i, Q_i can refer either to the full or a local auxiliary basis, specific for the occupied orbital *i*. To obtain a factorisable expression we use a numerical Laplace transformation as proposed by Almöf [15]:

$$\begin{aligned} t^{ii}_{ab} \approx &- \sum_{z}^{n_L} \omega_z e^{(-\varepsilon_a + \varepsilon_i - \varepsilon_b + \varepsilon_i)t_z} \\ &\times \sum_{Q_i P_i} (ia|Q_i) \left[{}^{(ii)}V^{-1}\right]_{Q_i P_i} (ib|P_i). \end{aligned} \tag{16}$$

Above, n_L is the number of the quadrature points t_z and ω_z the integration weights. n_L is usually small (<10) and does not scale with the system size. For the implementation we transform the occupied indices into the LMO basis and rewrite the last equation as:

$$\begin{aligned} t^{ii}_{ab} = &- \sum_{z}^{n_L} \omega_z \sum_{\nu} C^z_{\nu a} \sum_{Q_i} (Q_i|\nu i^z) \sum_{P_i} \left[{}^{(ii)}V^{-1}\right]_{Q_i P_i} \\ &\times \sum_{\mu} (P_i|\mu i^z) C^z_{\mu b} \end{aligned} \tag{17}$$

with modified MO coefficients defined as $C^z_{\mu a} = C_{\mu a} e^{-\varepsilon_a t_z}$ and $C^z_{\mu i} = \sum_j C_{\mu j} e^{\varepsilon_j t_z} T_{ij}$, where T_{ij} contains the LMO coefficients in the canonical MO basis. The sparsity of the AO integrals $(Q|\mu\nu)$ is exploited via the usual Schwarz screening condition $|(Q|\mu\nu)| \leq \sqrt{(\mu\nu|\mu\nu)(Q|Q)}$ to precompute the half-transformed three-index integrals $(Q_i|\nu i^z)$ with $\mathcal{O}(\mathcal{N}^3)$-scaling costs.

The multiplication of an amplitude matrix with a trial OSV vector d^k_b can now be expressed as a sequence of matrix–vector multiplications:

$$\begin{aligned} \sigma^k_a = \sum_c t^{ii}_{ac} b^k_c = &- \sum_{z}^{n_L} \omega_z \sum_{\nu} C^z_{\nu a} \sum_{Q_i} (Q_i|\nu i^z) \\ &\times \sum_{P_i} \left[{}^{(ii)}V^{-1}\right]_{Q_i P_i} \sum_{\mu} (P_i|\mu i^z) \sum_c C^z_{\mu c} d^k_c, \end{aligned} \tag{18}$$

Downloaded by [KIT Library] at 02:37 29 November 2013

$n_{cur} = 1$
$\alpha = 0.1$
• generate m orthonormal start vectors $\mathbf{b}^k$
for iter $= 1, ..., $ maxiter
  for $k = n_{cur}, ..., m$
    $\sigma_a^k = \sum_b t_{ab}^{ii} b_b^k$
  end for
  • calculate/update $\bar{T}_{kl} = \mathbf{b}^k \cdot \sigma_b^l$
  • calculate eigenvalues $\tau_{\tilde{a}}$ and -vectors $v_k^{\tilde{a}}$ of $\bar{\mathbf{T}}$
  • set N_{OSV} to number of eigenvalues $\tau_{\tilde{a}} > \alpha \cdot \frac{1}{2}\sqrt{T_{OSV}}$
  if $N_{OSV} = 0$ **then**
    • generate new start vectors $\mathbf{b}^k$
    goto next cycle
  end if
  for $\tilde{a} = 1, 2, ..., N_{OSV}$
    $\mathbf{d}^{\tilde{a}} = \sum_k b^k v_k^{\tilde{a}}$
    $\mathbf{r}^{\tilde{a}} = \sum_k \sigma^k v_k^{\tilde{a}} - \tau_{\tilde{a}} d^{\tilde{a}}$
  end for
  ! test convergence:
  converged$_{\tilde{a}} = (\tau_{\tilde{a}} - \tau_{\tilde{a}}^{(old)} < \frac{\alpha}{2}\sqrt{T_{OSV}})$ and $(\|\mathbf{r}^{\tilde{a}}\| < T_{res})$
  if for all $(\tilde{a} = 1, ..., N_{OSV})$ converged$_{\tilde{a}}$ **then**
    • algorithm converged. exit loop.
  else
    $n_{cur} = m + 1$
    for $\tilde{a} = 1, 2, ..., N_{OSV}$
      if not converged$_{\tilde{a}}$ **then**
        $m = m + 1$
        • $\mathbf{b}^m \leftarrow \mathbf{r}^{\tilde{a}}$ and add 5% white noise to $\mathbf{b}^m$
        • orthonormalize $\mathbf{b}^m$ against $\mathbf{b}^n$ with $n < m$
      end if
    end for
  end if
end for

Figure 1. Modified block Davidson algorithm for the solution of the OSV eigenvalue problem.

which all have at most a quadratic cost-sclaing. This leads to an overall scaling of $\mathcal{O}(\mathcal{N}^3)$ for the computation of the OSVs for all occupied orbitals.

3.2. *Block Davidson algorithm for the OSV eigenvalue problem*

The modified block Davidson algorithm [36], which we developed for the solution of the OSV eigenvalue problem is sketched in Figure 1. Like in principle every iterative diagonalisation technique, the Davidson algorithm needs start vectors to initialise the search for the requested eigenvectors and eigenvalues of a matrix. The eigenvalue problem is first solved in the subspace spanned by the start vectors, which is then in each iteration extended by additional vectors until a user defined accuracy for the eigenvectors and eigenvalues is reached. In difference to common applications of the block Davidson algorithm, we do not know in advance the number of required eigenvectors (i.e. OSVs), but want to solve for all eigenpairs with eigenvalues that are larger than or equal to a selection threshold.

The iterative diagonalisation converges fast if the start vectors are a good approximation for the eigenvectors of interest. The OSV spaces of spatially close LMOs have usually a large overlap. We, therefore, sort the LMOs such that LMOs with spatially close charge centers are also close in index space and use the converged OSVs of the previous LMO $i - 1$ as start vectors for LMO i. For the first LMO, where such vectors are not available, the trial vectors are chosen as unit vectors with the non-zero elements selected according to the largest diagonal elements in a very rough approximation for the diagonal amplitudes:

$$t_{aa}^{ii} \approx -\sum_z \omega_z \sum_\mu \left(C_{\mu a}^z\right)^2 \sum_P (P|\mu i^z)^2. \tag{19}$$

In each iteration, the trial and sigma vectors are used to construct the amplitude matrix projected onto the current subspace $\bar{T}_{kl} = \sum_k \mathbf{b}^k \cdot \sigma^l$, which is then diagonalised as $\bar{T}_{kl} = \sum_{\tilde{a}} V_{k\tilde{a}} \tau_{\tilde{a}} V_{l\tilde{a}}$. For all eigenvalues with $\tau_{\tilde{a}} \geq \alpha \cdot \frac{1}{2}\sqrt{T_{\mathrm{OSV}}}$ we construct then the approximate eigenvectors in the full virtual orbital space and their remaining residuals as

$$\mathbf{d}^{\tilde{a}} = \sum_k \mathbf{b}^k V_{k\tilde{a}}, \tag{20}$$

$$\mathbf{r}^{\tilde{a}} = \sum_k \sigma^k V_{k\tilde{a}} - \tau_{\tilde{a}} \mathbf{d}^{\tilde{a}}, \tag{21}$$

and construct from the residual vectors with a norm larger than a threshold T_{res} new trial vectors. It turned out that $\alpha = 0.1$ and $T_{\mathrm{res}} = \min(10 \cdot \sqrt{T_{\mathrm{OSV}}}, 10^{-4})$ are sufficient to ensure that the algorithm converges all significant eigenvectors accurately enough (vide infra). If in the first iteration no significant eigenvectors are found the start vectors have been inappropriate for the current LMO, which happens if a jump occurred in the sorted list of LMOs. In this case new start vectors are created by the same recipe as for the first LMO.

We did not find an efficient preconditioner for the amplitude matrices. Thus, we use the residual vectors directly as new trial vectors after they have been orthonormalised onto the previous basis vectors. To ensure that no significant eigenvectors are missed, e.g. because of a symmetry blocking of the amplitude matrix we add before the orthonormalisation 5% white noise.

3.3. *Local RI approximation*

Similar as the RI or density fitting approximation in local correlation methods [37] and in the PNO–CEPA method reported in Ref. [28] can be done in a local orbital- or pair-specific auxiliary basis, also the hybrid OSV–PNO approach can be combined with a local RI to reduce further the computational efforts. To conserve the smoothness of potential energy surfaces it is desirable to select the orbital-specific auxiliary basis sets at the level of individual shells of fitting

functions instead of complete atomic auxiliary basis sets [37]. Important for an accurate approximation of the exchange integrals (*ai*|*bi*) for the diagonal pairs *ii* are only the auxiliary functions in the same region of space as the orbital products *ai*. This leads to the suggestion to use an overlap criterion to select the local RI basis. The 3-index overlap integrals (*Qpi*) are, however, not suitable for this purpose, because they can vanish even though *Q* and *pi* have contributions in the same region of space. Better choices would be the overlap integrals for the squared orbitals (*QQppii*) or the absolute values $\int |Qpi| d\tau$, but the exact evaluation of these quantities is too costly. A cheaper approach is based on the Schwarz condition for the three-index overlap integrals:

$$\frac{|(Qpi)|}{\sqrt{(QQ)}} \leq \frac{\sqrt{(QQii)}}{\sqrt{(QQ)}} = \mathcal{S}_{Qi}, \tag{22}$$

where normalised molecular orbitals p are assumed. $\mathcal{S}_{Qi}$ has the advantage that it only becomes zero if also (*QQppii*) and $\int |Qpi| d\tau$ vanish, but it can be computed with almost the same (linearly scaling) costs as (*Qpi*). For the calculation of the OSVs we select for every LMO i a local RI basis $\mathcal{Q}_i = \left\{ Q \,|\, \mathcal{S}_{Qi} \geq T_{RI} \right\}$ based on the overlap criterion $\mathcal{S}_{Qi}$ and a suitably chosen threshold T_{RI} (vide infra). For the calculation of the PNOs and the exchange integrals in the equation for the amplitudes (8) and the PNO–MP2 energy (7) we use a pair-specific RI basis selected as:

$$\mathcal{Q}_{ij} = \left\{ Q \,|\, \mathcal{S}_{Qi} \cdot f_{ij} \geq T_{\mathrm{RI}} \text{ or } \mathcal{S}_{Qj} \cdot f_{ij} \geq T_{\mathrm{RI}} \right\} \tag{23}$$

Here, f_{ij} is a bias factor. For $f_{ij} = 1$ the set $\mathcal{Q}_{ij}$ becomes the union of the LMO-specific sets $\mathcal{Q}_i$ and $\mathcal{Q}_j$. This would have the disadvantage that the largest RI basis sets are obtained for distant pairs, which make only small contributions to the correlation energy. For these pairs the LMO-specific domains have not much in common with each other. It is thus reasonable to chose f_{ij} as a function of the pair energy estimates $\mathcal{E}_{ij}$ from Equation (12) to adjust the auxiliary basis sets to the importance of a pair. Test calculations revealed that $f_{ij} = 10 \times \mathcal{E}_{ij}^{1/3}$ gives a good accuracy to cost performance. To compute the exchange integrals $K^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}}$ within the local RI approximation we transform the three-index integrals $(Q_{ij}|\mu j)$ into the pre-PNO basis of the pair *ij* and solve the linear equation

$$V_{P_{ij}Q_{ij}} C_{Q_{ij}\tilde{a}_{ij}j} = \left(Q_{ij} \middle| \tilde{a}_{ij} j \right). \tag{24}$$

The solution vector is then contracted with the three-index integrals:

$$K^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}} = \sum_{Q_{ij}} C_{Q_{ij}\tilde{a}_{ij}j} \left(Q_{ij} | \tilde{b}_{ij} i \right). \tag{25}$$

The computation of the integrals $K^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}}$ is one of the most time-determining steps in the $\mathcal{O}(\mathcal{N}^3)$-scaling implementation of the OSV–PNO hybrid approach, but for large systems its cost contribution becomes less important as these steps scale linearly with the system size in the asymptotic limit.

4. Results and discussion

4.1. Computational details

The canonical MP2 calculations have been carried out with the ricc2 [38] module of the TURBOMOLE program package [39]. For the HF and DFT calculations the dscf [40] module has been used. The PNO–MP2 calculations have been done with an experimental PNO–MP2 code, which is part of the development version of the TURBOMOLE package. We applied the frozen core approximation for all orbitals with eigenvalues below $-3\,E_h$, with the exception of the calculations on glycine chains, Section 4.3, and $(GaAs)_n$ clusters, Section 4.2.1. Calculations have been done with the SVP [41], cc-pVDZ, cc-pVTZ [42,42], aug-cc-pVTZ [43], and cc-pwCVDZ [44] orbital basis set in combination with the optimised auxiliary basis sets reported in Refs. [45–47] for the RI approximation in the MP2 calculations. For the $(GaAs)_n$ clusters, the same structures as in Ref. [47] were used. The structures for the glycine chains are part of Ref. [34]. If not otherwise mentioned all other structures were optimised on the level of DFT/PBE/SVP.

4.2. Accuracy of the $\mathcal{O}(\mathcal{N}^3)$-scaling implementation

For the $\mathcal{O}(\mathcal{N}^3)$-scaling implementation we introduced an iterative OSV generation and a local RI approximation, which were not part of the previous implementation. In the following sections we analyse their performance and demonstrate that they do not affect the overall accuracy at all but retain the full accuracy of PNO–MP2 if thresholds are properly adjusted.

4.2.1. Iterative OSV generation

To test the accuracy of the iterative OSV generation, we compare the relative error due to the iterative diagonalisation

$$\Delta_{\text{it.-OSV}} = \left| \frac{E^{\mathcal{O}(\mathcal{N}^4)}_{\text{PNO-MP2}} - E^{\mathcal{O}(\mathcal{N}^3)}_{\text{PNO-MP2}}}{E_{\text{RI-MP2}}} \right|, \tag{26}$$

with the relative error due to the PNO truncation

$$\Delta_{\text{PNO}} = \left| \frac{E_{\text{RI-MP2}} - E_{\text{PNO-MP2}}}{E_{\text{RI-MP2}}} \right|, \tag{27}$$

Downloaded by [KIT Library] at 02:37 29 November 2013

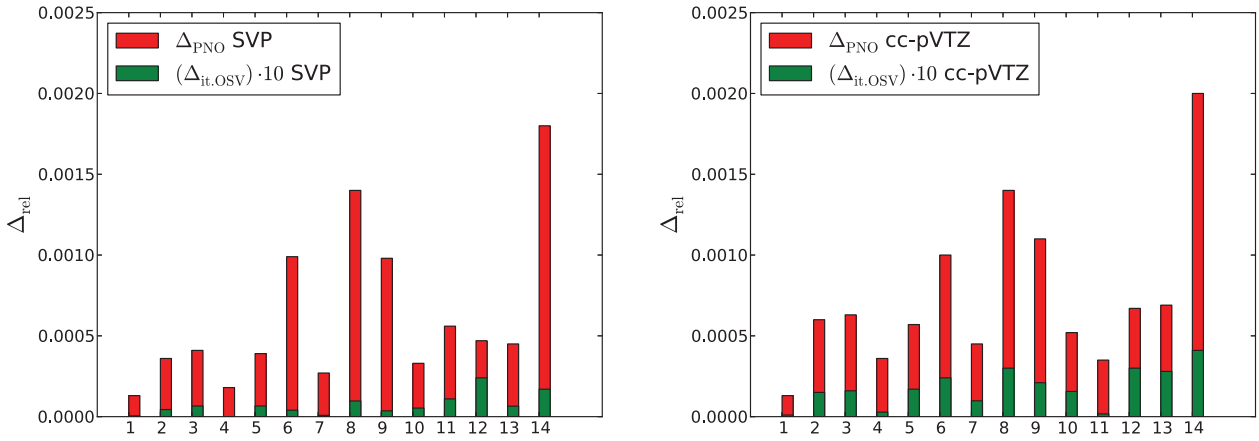


Figure 2. Relative errors due to the iterative OSV generation compared to the PNO truncation error for 14 organic molecules in the SVP and cc-pVTZ basis sets. Note, that for better visibility the errors due to iterative solution of the OSV equations has been scaled by a factor of 10.

Downloaded by [KIT Library] at 02:37 29 November 2013

where $E_{\text{RI-MP2}}$ is the RI–MP2 energy without the PNO approximation, $E_{\text{PNO-MP2}}$ the PNO–MP2 energy without using the hybrid OSV–PNO approach, and $E^{\mathcal{O}(\mathcal{N}^4)}_{\text{PNO-MP2}}$ and $E^{\mathcal{O}(\mathcal{N}^3)}_{\text{PNO-MP2}}$ are, respectively, the results obtained with the $\mathcal{O}(\mathcal{N}^4)$- and $\mathcal{O}(\mathcal{N}^3)$-scaling implementations without using the local RI approximation. We evaluated these errors for a test set of 14 organic molecules including acetone (1), an adenine-thymine dimer (2), azobenzene (3), benzene (4), cyclohexane (5), 1,3-dibutylphenol (6), 8, 8-dipentyl-tetradecane (7), $(\text{glycine})_4$ (8), hexadecane (9), norbornane (10), pentacosane (11), Phe-Val-Asn (12), triphenyl-methanol (13), and the trityl dimer (14). For these calculations the SVP and cc-pVTZ basis sets have been used

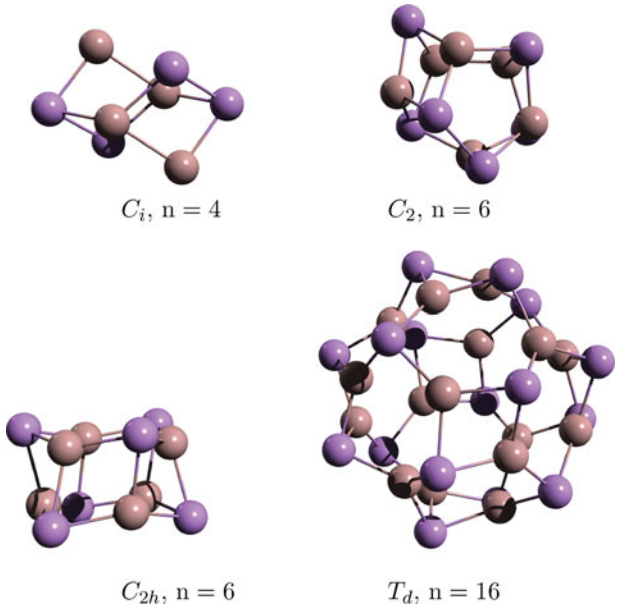


Figure 3. Structures of the four $(\text{GaAs})_n$ clusters.

Table 1. Relative PNO truncation error for the different $(\text{GaAs})_n$ clusters and the relative error due to the iterative diagonalisation for the cc-pwCVDZ basis set for $T_{\text{PNO}} = 10^{-8}$ and $T_{\text{OSV}} = 10^{-9}$.

Molecules	$\Delta_{\text{it.-OSV}}$	Δ_{PNO}
$(\text{GaAs})_4$ C_i	1.2×10^{-6}	6.5×10^{-4}
$(\text{GaAs})_6$ C_2	3.1×10^{-5}	8.5×10^{-4}
$(\text{GaAs})_6$ C_{2h}	3.0×10^{-5}	8.4×10^{-4}
$(\text{GaAs})_{16}$ T_d	8.2×10^{-5}	1.2×10^{-3}

and the PNO truncation threshold was set to $T_{\text{PNO}} = 10^{-8}$ and the threshold for the OSV selection to $T_{\text{OSV}} = 0.1 \cdot T_{\text{PNO}}$ as proposed in Ref. [34]. As a second test set we studied the $(\text{GaAs})_n$ clusters with $n = 4, 6, 16$ (see Figure 3), for which we performed analogous calculations in the cc-pwCVDZ basis set to test the performance of the OSV–PNO hybrid approach for the correlation of a semi-core $3d$ shell. Figure 2 shows a comparison of the errors for the organic test set. The results demonstrate that with the settings for α and T_{res} proposed in section 3.2 the error introduced by the iterative diagonalisation is at least one, but in most cases actually two, orders of magnitude smaller than the PNO truncation error and does not affect the overall accuracy. Table 1 illustrates that this is also the case for the $(\text{GaAs})_n$ clusters, where in addition to the valence $4s4p$ also the semi-core $3d$ electrons have been correlated.

4.2.2. *Local RI approximation*

To study how the local RI approximation affects the accuracy, the absolut value of the relative error

$$\Delta^{rel}_{RI} = \left| \frac{E^{\text{Full}}_{\text{RI}} - E^{\text{Local}}_{\text{RI}}}{E_{\text{RI-MP2}}} \right| \tag{28}$$

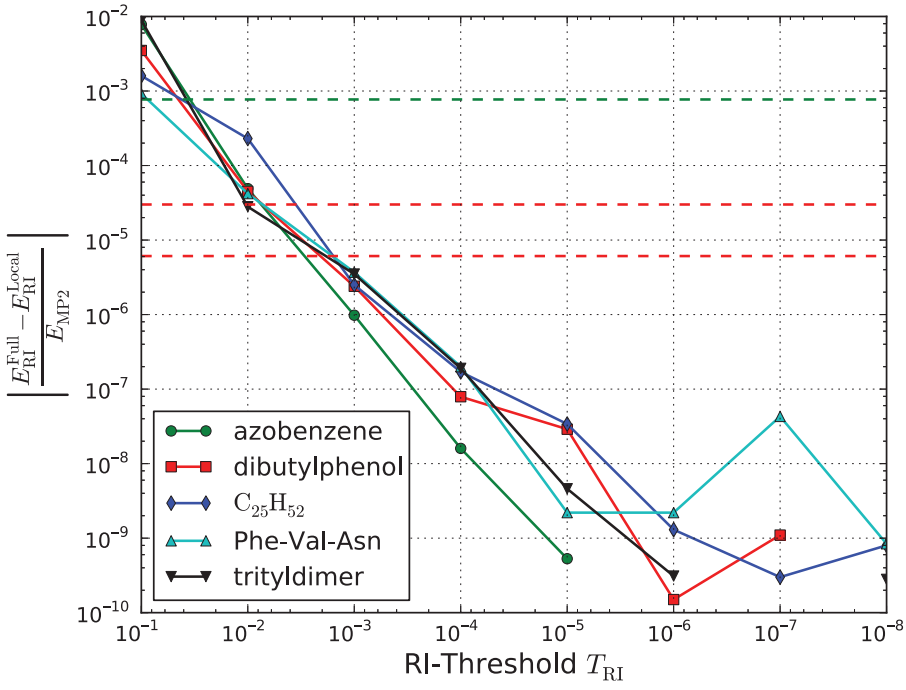


Figure 4. Error due to the local RI approximation for a set of 5 molecules in the cc-pVTZ basis set. The green dashed line indicates the average of the PNO truncation error and the red dashed lines indicate the range of the error due to the iterative OSV generation by the average error plus and minus one standard deviation.

has been calculated for the basis sets SVP, cc-pVDZ, cc-pVTZ and aug-cc-pVTZ and plotted against the local RI threshold T_{RI} for $T_{PNO} = 10^{-8}$. As an example such a plot is shown in Figure 4 for the cc-pVTZ basis set. The dashed green line indicates the average of the PNO truncation error and the red dashed lines indicate the range of the error due to the iterative OSV generation by the average of this error plus and minus the standard deviation. In the range of 10^{-1} to 10^{-6} the error Δ_{RI}^{rel} decreases smoothly until it tends to oscillate a bit below $T_{RI} = 10^{-6}$, due to numerical noise. The exceptional large increase in the error for Phe-Val-Asn at $T_{RI} = 10^{-7}$ is caused by one additional OSV, which fell accidentally above the threshold T_{OSV} and lead to a slightly larger correlation energy. Although, this effect can clearly be seen on the logarithmic plot, it is not a problem for practical calculations, because the effect is several orders of magnitude smaller than the PNO truncation error. The plots for the other basis sets, which are not shown, are similar, although there is of course a small basis set dependence. In general, the local RI error for a given threshold T_{RI} is smaller for larger auxiliary basis sets since in a larger basis the number of auxiliary functions in the same region of space is increased and more functions fulfill the overlap criterion.

For an optimal cost to accuracy ratio, the error caused by the local RI approximation should be just below the error due to the iterative OSV generation. For the examples in Figure 4 with $T_{PNO} = 10^{-8}$ this criterion is met with $T_{RI} = 10^{-3}$. To find a general relation between T_{PNO} and T_{RI} we carried out calculations on the five molecules used for Figure 4 with the SVP, cc-pVDZ, cc-pVTZ and aug-cc-pVTZ basis set and varied T_{PNO} and T_{RI}. As an example Figure 5 shows the average relative PNO truncation and local RI errors for the trityl dimer in the SVP and cc-pVTZ basis set as functions of the thresholds T_{PNO} and T_{RI}. The dashed lines show the results for power laws $a \times T^b$ fitted to the relative errors averaged over the five molecules and four basis sets. Demanding that the error due the local RI approximation is at least two orders of magnitude below the PNO truncation error, leads to the relation

$$\log_{10} T_{RI} \leq 0.5 + 0.5 \log_{10} T_{PNO} \quad , \qquad (29)$$

which we propose to adjust T_{RI} automatically to T_{PNO}.

4.3. *Performance in comparison to the $\mathcal{O}(\mathcal{N}^4)$-scaling implementation and canonical RI–MP2*

To study the performance of the $\mathcal{O}(\mathcal{N}^3)$-scaling OSV–PNO–MP2 implementation compared to the $\mathcal{O}(\mathcal{N}^4)$-scaling implementation and canonical RI–MP2 we calculated the correlation energy of glycine chains $(Gly)_n$ with $n = 1, 2, 4, 8, 16, 32$ in the aug-cc-pVTZ basis set with $T_{RI} = 10^{-3}$ and PNO thresholds of 10^{-6}, 10^{-7} and 10^{-8}. Theses choices for T_{PNO} are sufficient for practical MP2 calculations [34]. In Figure 6 the wall time is plotted against the number of basis functions for a PNO

Downloaded by [KIT Library] at 02:37 29 November 2013

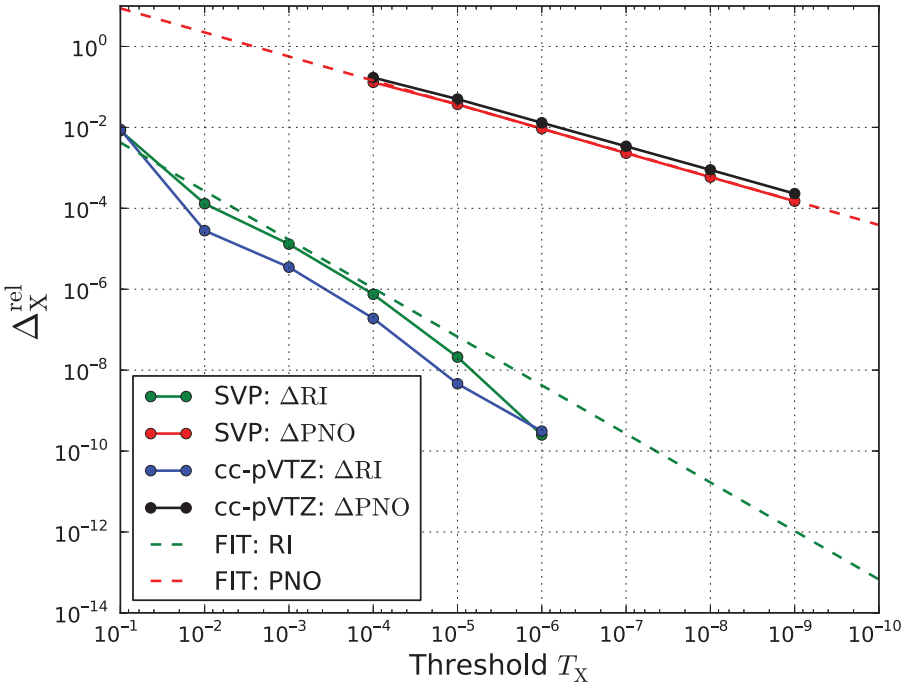


Figure 5. Comparison of the PNO truncation error and the local RI error as functions of the respective threshold (T_X). As an example the results for the SVP and cc-pVTZ basis set are plotted for the trityl dimer (solid lines). The dashed lines show the results for power laws fitted to average errors for the five molecules and two basis sets.

threshold of 10^{-8}. The results reveal an early break even point between the cubically and the quartically scaling implementations between 700 and 800 basis functions that depends only slightly on T_{PNO}. The break even points with the canonical RI-MP2 implementation are considerably later, i.e. for $(\mathrm{Gly})_n$ and the three values of T_{PNO} at about 830, 1000 and 1360 basis functions.

To analyse the effective scaling, we fitted the results for the timings to power laws of the form $10^B \cdot N_{bf}^A$, where N_{bf} is the number of basis functions, A the effective scaling, and 10^B the prefactor. The results are listed in Table 2 and show that the computational costs for the three implementations scale indeed with $\mathcal{O}(\mathcal{N}^3)$, $\mathcal{O}(\mathcal{N}^4)$, and $\mathcal{O}(\mathcal{N}^5)$ or even better. Interesting are the timings for the largest molecule, $(\mathrm{Gly})_{32}$, where the orbital basis contains 8188 basis functions. The time for a canonical RI-MP2 calculation extrapolated with the power law would be over 200 days. With the $\mathcal{O}(\mathcal{N}^4)$-scaling implementation the calculation took 20 days and with the cubically scaling implementation less than 6 days, in both cases for $T_{\mathrm{PNO}} = 10^{-8}$. This demonstrates impressively the new possibilities opened by the cubically scaling OSV–PNO–MP2 implementation. Taking a closer look at the plot, one might recognises that the slope for the $\mathcal{O}(\mathcal{N}^3)$-scaling implementation between $(\mathrm{Gly})_{16}$ and $(\mathrm{Gly})_{32}$ (or 4140 and 8188 basis functions) is slightly larger than between $(\mathrm{Gly})_4$ and $(\mathrm{Gly})_{16}$. This increase is related to a switch to a multipass algorithm for the evaluation of the Laplace-transformed three-index integrals $(Q|\mu i^z)$, which for $(\mathrm{Gly})_{32}$ could no longer be stored for all LMOs at once with the available disc space, because we had at this step not yet exploited the local RI approximation for storing the integrals.

Calculations on essentially linear chain-like molecules are suitable for analysing the scaling of an implementation, but provide only limited insights into the typical performance in practical applications.

To illustrate how the local RI affects this step we performed calculations on the trityldimer in the cc-pVTZ basis, and compared the timings for the evaluation of the exchange integrals with an analogue calculation without using the local RI approximation. With $T_{\mathrm{RI}} = 10^{-3}$ the average number of auxiliary functions was 1335, compared to 3978 functions in the full basis. In the full RI basis the evaluation of $K^{ij}_{\tilde{a}_{ij}\tilde{b}_{ij}}$ took 2 hour and 38 min; 1 hour and 23 min of this for the construction of the B intermediates in the pre-PNO basis from $(Q|\mu i)$. With $T_{\mathrm{RI}} = 10^{-3}$, the computation of the exchange integrals took only 45 min where 19 min were spend to solve the equations for the fit coefficients, $C_{Q_{ij}\tilde{a}_{ij}j}$, in Equation (24). Moreover the local RI also reduces for larger systems significantly the disk demands, in particular for the Laplace-transformed integrals.

An example for a molecule with a bulky three-dimensional structure is somatostatin, a small peptide consisting of 14 amino acids (Ala-Gly-Cys-Lys-Asn-Phe-

Downloaded by [KIT Library] at 02:37 29 November 2013

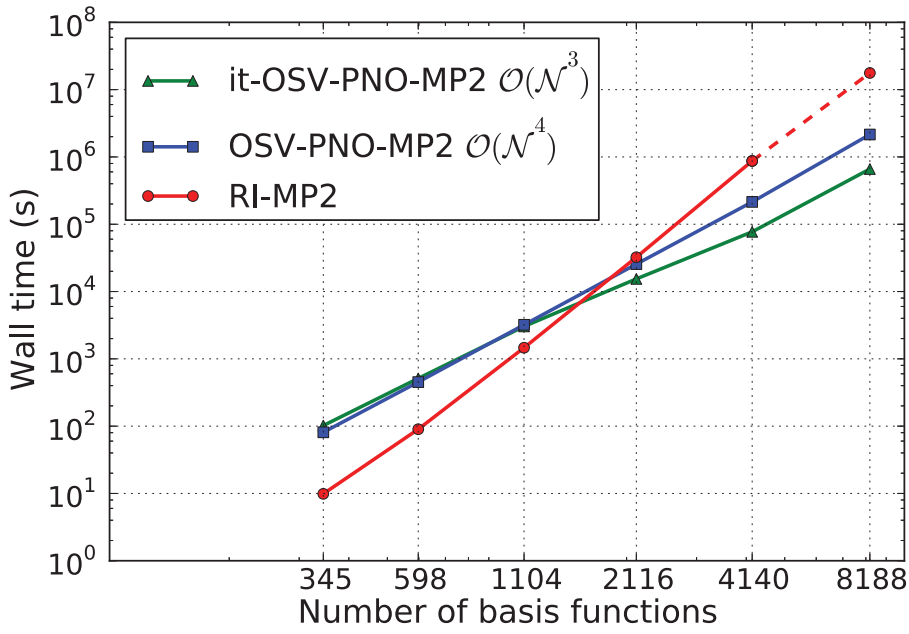


Figure 6. Wall time for the $\mathcal{O}(\mathcal{N}^3)$, $\mathcal{O}(\mathcal{N}^4)$ and canonical $\mathcal{O}(\mathcal{N}^5)$ implementation using the aug-cc-pVTZ basis with a PNO threshold of 10^{-8} and a RI threshold of $T_{\mathrm{RI}} = 10^{-3}$ for glycine chains $(\mathrm{Gly})_n$ with $n = 1, 2, 4, 8, 16, 32$ (the dashed line has been extrapolated using the power law in Table 2).

Table 2. Results for the power laws of the form $10^B \cdot N_{bf}^A$ fitted to the wall time (in seconds) for $(\mathrm{Gly})_n$ with the three MP2 implementation and different values for the PNO truncation threshold T_{PNO}.

T_{PNO}	Wall(MP2)	Wall($\mathcal{O}(\mathcal{N}^3)$)	Wall($\mathcal{O}(\mathcal{N}^4)$)
10^{-6}	$10^{-10.8} \cdot N_{bf}^{4.60}$	$10^{-5.4} \cdot N_{bf}^{2.75}$	$10^{-7.0} \cdot N_{bf}^{3.32}$
10^{-7}	”	$10^{-5.1} \cdot N_{bf}^{2.72}$	$10^{-6.7} \cdot N_{bf}^{3.26}$
10^{-8}	”	$10^{-4.7} \cdot N_{bf}^{2.67}$	$10^{-6.2} \cdot N_{bf}^{3.18}$

Phe-Trp-Lys-Thr-Phe-Thr-Ser-Cys). Its structure is shown in Figure 7. For this molecule, we computed the OSV–PNO–MP2 energy in the cc-pVQZ basis set (9573 basis functions) with $T_{\mathrm{PNO}} = 10^{-8}$ and $T_{\mathrm{RI}} = 10^{-4}$. The calculation took on a machine with 48 GB memory on a single core of an Intel Xeon® X5670 CPU with 2.93 GHz about 17 days and 19 hours. This shows that calculations with nearly 10,000 basis functions in more complex systems are not out of the scope, even without having yet applied any parallelisation.

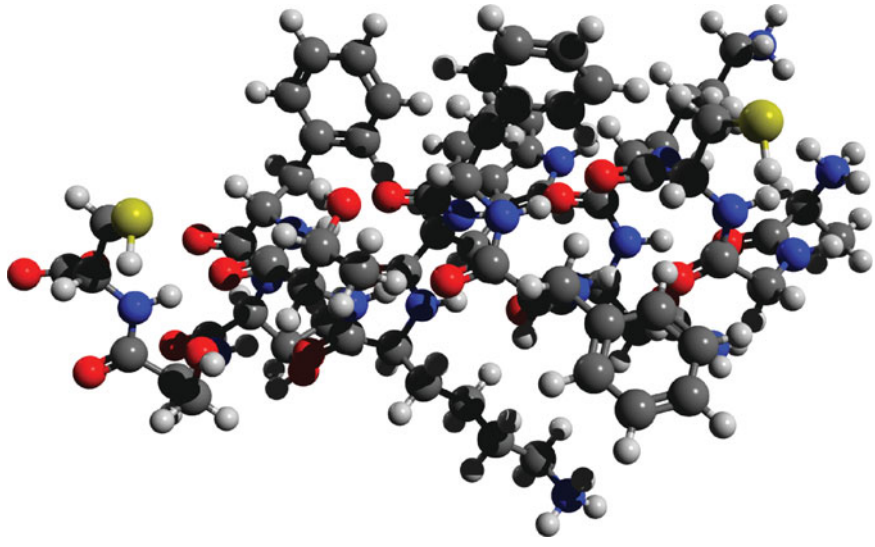


Figure 7. Structure of somatostatin as obtained with the molden sequence builder [48].

Downloaded by [KIT Library] at 02:37 29 November 2013

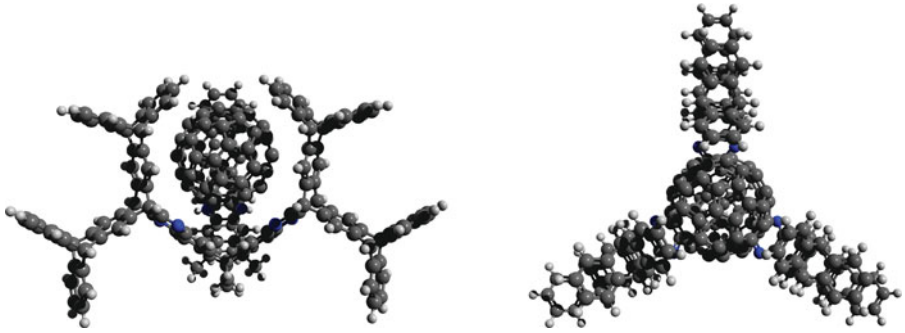


Figure 8. C_{60} fullerene acceptor system with 320 atoms.

Downloaded by [KIT Library] at 02:37 29 November 2013

Since it has been shown that the cubically scaling implementation of OSV–PNO–MP2 with an iterative solution of the OSV equations and the local RI approximation provides the same accuracy as PNO–MP2 without these and the OSV approximation, the term PNO–MP2 will now be used synonymously for the new implementation in the remaining sections of this paper.

4.4. Dispersion interaction in large systems

Since DFT with semi-local functionals is the most widely applied approach for large systems, but fails to describe dispersion interactions if not combined with empirical corrections or, as e.g. in the case of double hybrid functionals with a wavefunction-like orbital-dependent correlation contribution, it is particulary interesting to apply the new MP2 implementation to large systems, where these interactions are essential. If the performance achieved with PNO–MP2 for such systems is reasonable, the arsenal of computational chemistry to treat large systems can be complemented with this tool. Therefore, we carried out calculations on the C_{60} fullerene acceptor system, shown in Figure 8, and a foldamer system, shown in Figure 9. The fullerene acceptor system was recently studied in Ref. [49], where the binding energy was calculated at the DFT/B97-D/TZVPP level [50,51]. A binding energy of $-185\,\mathrm{kJ/mol}$ was predicted, which solely arises from Grimme's dispersion correction [52,53] in the B97-D functional. We computed with the implementation described in Section 3 the binding energy employing a cc-pVTZ basis set (8064 functions) on the level of PNO–MP2 and its spin-scaled variants [54,55] SCS–PNO–MP2 and SOS–PNO–MP2 with $T_{\mathrm{PNO}} = 10^{-7}$ and $T_{\mathrm{RI}} = 10^{-4}$. Moreover, we used the PNO–MP2 implementation to calculate the MP2-like correlation contribution for the pure double hybrid functional B2-PLYP [11] and with Grimme's empirical dispersion correction [56]. The longest calculation was the B2-PLYP computation, which took on a machine with an Intel Xeon® X5675 CPU with 3.07 GHz and 48 GB memory on a single core about 22 days and 6 hours.[1] The computation of the MP2-like energy contributions for the B2-PLYP functional took slightly longer than corresponding PNO–MP2 calculations with a Hartree–Fock reference. This is related to the different reference wave functions. For the Kohn-Sham orbitals from B2-PLYP the orbital energy differences in the denominators of the amplitudes are smaller than with Hartree–Fock orbitals which results in larger OSV and PNO occupation numbers. This leads for the fullerene acceptor system to approximately 5% more OSVs (73,999 compared to 70,041) with B2-PLYP orbitals.

The counterpoise corrected and uncorrected binding energies are listed in Table 3. As canonical MP2, also

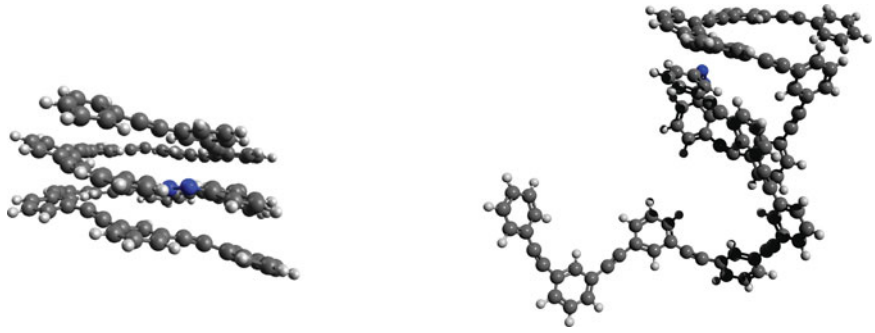


Figure 9. Foldamer in the folded and unfolded state.

Table 3. Counterpoise corrected $E_{\text{bind}}^{\text{CP}}$ and uncorrected $E_{\text{bind}}^{\text{UC}}$ binding energies for the fullerene acceptor system in the cc-pVTZ basis.

Method	$E_{\text{bind}}^{\text{CP}}$ (kJ/mol)	$E_{\text{bind}}^{\text{UC}}$ (kJ/mol)
HF	228.1	205.8
PNO-MP2	−371.3	−465.8
SCS-PNO-MP2	−235.8	−329.6
SOS-PNO-MP2	−168.0	−261.5
PNO-B2-PLYP	−38.9	−79.2
PNO-B2-PLYP-D	−217.3	−257.6

Table 4. Folding energy E_{fold} for the foldamer in the basis cc-pVTZ and cc-pVQZ.

Method	$E_{\text{fold}}^{\text{TZ}}$ (kJ/mol)	$E_{\text{fold}}^{\text{QZ}}$ (kJ/mol)	$E_{\text{fold}}^{\text{CBS}}$ (kJ/mol)
HF	83.2	86.3	–
PNO–MP2	−287.1	−287.5	−290.1
SCS–PNO–MP2	−168.2	−167.1	−168.5
SOS–PNO–MP2	−108.7	−106.9	−107.7
PNO–B2-PLYP	−38.1	−33.2	−32.1
PNO–B2-PLYP-D	−174.5	−169.6	−168.5

PNO–MP2 overestimates the dispersion interaction between π systems. The spin-component scaling in SCS–MP2 corrects for this overestimation [54,57] and, thus, the results from the SCS–PNO–MP2 calculation are more reliable. The PNO–B2-PLYP-D calculation predicts a value for the binding energy, which is only 8% smaller than the SCS–PNO–MP2 result. Because both SCS–MP2 and B2-PLYP-D are well known to provide usually quite accurate results, we conclude that the binding energy should be in the range of −217.3 to −235.8 kJ/mol.

The dispersion interaction is also essential for the folding energy of the photo-switchable molecule shown in Figure 9. The structures of the folded and unfolded states were obtained from a molecular dynamics (MD) simulation of 200 ps using a NVT ensemble and the GROMOS96 45a3 force field [58] with acetonitril as solvent. The system was synthesised by Khan *et al.* [59] and studied with MD techniques by Braun *et al.* in Ref. [60]. In the current work, we calculated the gas phase folding energy for this system at the same levels of theory as for the fullerene acceptor system using again $T_{\text{PNO}} = 10^{-7}$ and $T_{\text{RI}} = 10^{-4}$. The calculations were carried out in the cc-pVTZ and the cc-pVQZ basis sets. To correct for the remaining basis set error in the correlation energy and thereby reduce the, in this case intramolecular, basis set superposition error we extrapolated from the results for the two basis sets to the basis set limit using Helgaker's two point formula [61]:

$$E_X = E_{\text{lim}} - AX^{-3} \quad . \tag{30}$$

The extrapolated basis set limit for the correlation energy was then combined with the HF energy from the cc-pVQZ basis set. The gas phase folding energies for the cc-pVTZ and cc-pVQZ basis and the extrapolated basis set limits are listed in Table 4. The results indicate that the folding is solely driven by the dispersion interaction. Hartree–Fock, where this interaction is absent, predicts the folded state to be more than 80 kJ/mol higher in energy than the unfolded state. As for the fullerene acceptor system PNO–MP2 overestimates the dispersion interaction and thus the folding energy. The SCS–PNO–MP2 and PNO–B2-PLYP-D results are again in good agreement. Within the number of digits shown in Table 4 the extrapolated basis set limits for the two methods are actually identical and predict a folding energy of −168.5 kJ/mol. For the foldamer system the longest calculation took on a machine with an Intel Xeon® X5670 CPU with 2.93 GHz and 48 GB memory about 12 days and 21 hours. For this calculation the orbital basis (cc-pVQZ) contained 7790 basis functions.

5. Conclusions

We presented a $\mathcal{O}(\mathcal{N}^3)$-scaling local PNO–MP2 implementation, which combines the hybrid OSV–PNO approach with an iterative solution of the OSV equations and a local RI approximation. The additional approximations introduce only negligible errors so that the accuracy is still controlled by a single parameter, the threshold for truncating the PNO expansion, T_{PNO}. This was demonstrated with calculations on organic and inorganic molecules. The approach is combined with a local RI approximation. By using as selection criterion an upper bound for the overlap between basis function products and an auxiliary function, the local RI basis adapts automatically to the requirements of the wavefunction. To adjust the selection threshold T_{RI} automatically with T_{PNO}, we performed calculations with different T_{PNO} and T_{RI} to derive a simple relation between the two thresholds. Calculations on glycine chains $(\text{Gly})_n$ with $n = 1$, 2, 4, 8, 16, 32 show an early break even point with the $\mathcal{O}(\mathcal{N}^4)$-scaling implementation reported in Refs. [34] and with canonical RI–MP2. A calculation on $(\text{Gly})_{32}$ demonstrates the new possibilities opened up with the $\mathcal{O}(\mathcal{N}^3)$ scaling PNO–MP2 method: for this system the computational time for a canonical RI-MP2 calculation is extrapolated to be over 200 days while with the $\mathcal{O}(\mathcal{N}^3)$-scaling implementation the calculation takes less than 6 days.

The new implementation has been applied to study a C_{60} fullerene acceptor system with over 8000 basis function. The system was earlier investigated with DFT-D in Ref. [49], where it was observed that a binding only occurs due to the dispersion interaction, which could only be taken into account via Grimme's empirical dispersion correction. With the new PNO-MP2 implementation we were able to

Downloaded by [KIT Library] at 02:37 29 November 2013

compute the binding energy with a correlated wavefunction method. SCS-PNO–MP2 and PNO–B2-PLYP-D calculations predict a binding energy in the range of −217.3 to −235.8 kJ/mol, which is somewhat larger than the DFT-D result of −185 kJ/mol reported in Ref. [49]. A good agreement of SCS–PNO–MP2 and PNO–B2-PLYP-D was also observed in calculations on a foldamer system, where both methods lead to an extrapolated basis set limit result of −168.5 kJ/mol for the gas phase folding energy.

If combined with a parallelisation for distributed memory systems using e.g. the message passing interface (MPI) standard, the proposed hybrid OSV–PNO approach will allow to carry out MP2 calculations with more than 10,000 basis functions within less than a day and turn such calculations to a standard task. The PNO–B2-PLYP calculations illustrate that the new PNO-MP2 implementation also extents the applicability of double hybrid density functionals with an MP2-like correlation contribution. The $\mathcal{O}(\mathcal{N}^3)$-scaling direct solution of the OSV equations opens up also for other methods, as for example the random phase approximation [62], a new possibility to reduce the computational costs without the need to introduce a priori determined domains. The possibilities to use OSV-like expansions for different orbital space in explicitly correlated wavefunction methods have already been discussed in Ref. [34] and an obvious next step will be the combination of PNO–F12 methods with the iterative direct solution of the equations for OSVs and complementary auxiliary orbitals. Here, even greater savings can be expected due to the large size of complementary auxiliary basis sets and the small number of significant orbital- and pair-specific functions.

Downloaded by [KIT Library] at 02:37 29 November 2013

Acknowledgements

We thank R. Schmid for providing the structure of the C_{60} fullerene acceptor system and M. Böckmann for the structures of the foldamer. Financial support by the Deutsche Forschungsgemeinschaft through grant no. HA/2588/7 and the TURBOMOLE GmbH is gratefully acknowledged. C.H. thanks D.P. Tew for helpful discussions about the local RI approximations.

Note

1. The evaluation of the 3-index integrals was parallelised over 10 cores. Therefore we extrapolated for better comparison the time for this step on a single core.

References

[1] J.P. Perdew, and Y. Wang, Phys. Rev. B **45**, 13244 (1992).
[2] J.P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
[3] A.D. Becke, J. Chem. Phys. **98**, 1372 (1993).
[4] C. Lee, W. Yang, and R.G. Parr, Phys. Rev. B **37**, 785 (1988).
[5] S.H. Vosko, L. Wilk, and M. Nusair, Can. J. Phys. **58** (8), 1200 (1980).
[6] P.J. Stephens, F.J. Devlin, C.F. Chabalowski, and M.J. Frisch, J. Phys. Chem. **98** (45), 11623 (1994).
[7] F.P.B. II and J.E. Bloor, J. Chem. Phys. **55** (11), 5178 (1971).
[8] J.L. Whitten, J. Chem. Phys. **58** (10), 4496 (1973).
[9] K. Eichkorn, O. Trutler, H. Öhm, M. Häser, and R. Ahlrichs, Chem. Phys. Lett., **242**, 652 (1995).
[10] M. Sierka, A. Hogekamp, and R. Ahlrichs, J. Chem. Phys. **118** (20), 9136 (2003).
[11] S. Grimme, J. Chem. Phys. **124**, 034108 (2006).
[12] K. Sharkas, J. Toulouse, and A. Savin, J. Chem. Phys. **134** (6), 064113 (2011).
[13] D.S. Lambrecht, B. Doser, and C. Ochsenfeld, J. Chem. Phys. **123**, 184102 (2005).
[14] B. Doser, D.S. Lambrecht, and C. Ochsenfeld, Phys. Chem. Chem. Phys. **10**, 3335 (2008).
[15] M. Häser and J. Almlöf, J. Chem. Phys. **96**, 489 (1992).
[16] M. Schütz, G. Hetzer, and H.-J. Werner, J. Chem. Phys. **111**, 5691 (1999).
[17] P.E. Maslen, M.S. Lee, and M. Head-Gordon, Chem. Phys. Lett. **319**, 205 (1999).
[18] M. Schütz and H.-J. Werner, J. Chem. Phys. **114**, 661 (2001).
[19] M. Schütz and H.-J. Werner, Chem. Phys. Lett., **318**, 370 (2000).
[20] J.E. Subotnik, A. Sodt, and M. Head-Gordon, J. Chem. Phys. **125** (7), 074116 (2006).
[21] S. Saebø and P. Pulay, Chem. Phys. Lett., **113**, 13 (1985).
[22] S. Saebø and P. Pulay, Ann. Rev. Phys. Chem. **44**, 213 (1993).
[23] W. Meyer, Int. J. Quantum Chem. **5**, 441 (1971).
[24] W. Meyer, J. Chem. Phys. **58** (3), 1017 (1973).
[25] R. Ahlrichs, H. Lischka, V. Staemmler, and W. Kutzelnigg, J. Chem. Phys. **62**, 1225 (1975).
[26] R. Ahlrichs, F. Driessler, H. Lischka, V. Staemmler, and W. Kutzelnigg, J. Chem. Phys. **62**, 1235 (1975).
[27] R. Ahlrichs, F. Keil, H. Lischka, W. Kutzelnigg, and V. Staemmler, J. Chem. Phys., **63**, 455 (1975).
[28] F. Neese, F. Wennmohs, and A. Hansen, J. Chem. Phys. **130**, 114108 (2009).
[29] J. Yang, Y. Kurashige, F.R. Manby, and G.K.L. Chan, J. Chem. Phys. **134** (4), 044123 (2011).
[30] C. Krause and H.-J. Werner, Phys. Chem. Chem. Phys. **14**, 7591 (2012).
[31] D.P. Tew and C. Hättig, Int. J. Quantum Chem. **113** (3), 224 (2013).
[32] B. Helmich and C. Hättig, J. Chem. Phys. **135** (21), 214106 (2011).
[33] J. Yang, G.K.-L. Chan, F.R. Manby, M. Schütz, and H.-J. Werner, J. Chem. Phys. **136**, 144105 (2012).
[34] C. Hättig, D.P. Tew, and B. Helmich, J. Chem. Phys. **136**, 204105 (2012).
[35] C. Riplinger and F. Neese, J. Chem. Phys. **138** (3), 034106 (2013).
[36] E.R. Davidson, J. Chem. Phys. **17**, 87 (1975).
[37] H.-J. Werner, F.R. Manby, and P.J. Knowles, J. Chem. Phys. **118** (18), 8149 (2003).
[38] C. Hättig and F. Weigend, J. Chem. Phys. **113**, 5154 (2000).
[39] "TURBOMOLE V6.4 2012, a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH, 1989-2007, TURBOMOLE GmbH, since 2007; <http://www.turbomole.com.>2012.
[40] R. Ahlrichs, M. Bär, H. Horn, and C. Kölmel, Chem. Phys. Lett. **162**, 165 (1989).
[41] A. Schäffer, H. Horn, and R. Ahlrichs, J. Chem. Phys. **97**, 2571 (1992).
[42] D.E. Woon and T.H.J. Dunning, J. Chem. Phys. **98**, 1358 (1993).
[43] T.H.J. Dunning, J. Chem. Phys. **90**, 1007 (1989).
[44] K. Peterson and T.J. Dunning, J. Chem. Phys. **117**, 10548 (2002).

[45] F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs, Chem. Phys. Lett. **294**, 143 (1998).
[46] C. Hättig, Phys. Chem. Chem. Phys. **7**, 59 (2005).
[47] C. Hättig, G. Schmitz, and J. Koßmann, Phys. Chem. Chem. Phys. **14**, 6549 (2012).
[48] G. Schaftenaar and J.H. Noordik, J. Comput. Aided Mol. Des. **14**, 123 (2000).
[49] S. Henne, B. Bredenkötter, A.A.D. Baghi, R. Schmid, and D. Volkmer, Dalton Trans. **41**, 5995 (2012).
[50] F. Furche and R. Ahlrichs, J. Chem. Phys. **117**, 7433 (2002).
[51] A. Schäfer, C. Huber, and R. Ahlrichs, J. Chem. Phys. **100**, 5829 (1994).
[52] S. Grimme, J. Comput. Chem. **27** (15), 1787 (2006).
[53] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, J. Chem. Phys. **132**, 154104 (2010).
[54] S. Grimme, J. Chem. Phys. **118**, 9095 (2003).
[55] Y. Jung, R.C. Lochan, A.D. Dutoi, and M. Head-Gordon, J. Chem. Phys. **121**, 9793 (2004).
[56] T. Schwabe and S. Grimme, Phys. Chem. Chem. Phys. **9**, 3397 (2007).
[57] Y. Jung and M. Head-Gordon, Phys. Chem. Chem. Phys. **8**, 2831 (2006).
[58] L.D. Schuler, X. Daura, and W.F. van Gunsteren, J. Comp. Chem. **22**, 1205 (2001).
[59] A. Khan, C. Kaiser, and S. Hecht, Angew. Chem. Int. Ed. **45** (12), 1878 (2006).
[60] S. Braun, M. Böckmann, and D. Marx, ChemPhysChem, **13** (6), 1440 (2012).
[61] A. Halkier, T. Helgaker, P. Jøsen, W. Klopper, H. Koch, J. Olsen, and A.K. Wilson, Chem. Phys. Lett. **286** (34), 243 (1998).
[62] H. Eshuis, J. Yarkony, and F. Furche, J. Chem. Phys. **132** (23), 234114 (2010).

Downloaded by [KIT Library] at 02:37 29 November 2013