Efficient linear-scaling second-order Møller-Plesset perturbation theory: The divide–expand–consolidate RI-MP2 model

Pablo Baudin, Patrick Ettenhuber, Simen Reine, Kasper Kristensen, and Thomas Kjærgaard

Citation: The Journal of Chemical Physics **144**, 054102 (2016); doi: 10.1063/1.4940732
View online: http://dx.doi.org/10.1063/1.4940732
View Table of Contents: http://aip.scitation.org/toc/jcp/144/5
Published by the American Institute of Physics

Articles you may be interested in

Fast linear scaling second-order Møller-Plesset perturbation theory (MP2) using local and density fitting approximations
The Journal of Chemical Physics **118**, 8149 (2003); 10.1063/1.1564816

The molecular gradient using the divide-expand-consolidate resolution of the identity second-order Møller-Plesset perturbation theory: The DEC-RI-MP2 gradient
The Journal of Chemical Physics **145**, 024106 (2016); 10.1063/1.4956454

The Laplace transformed divide-expand-consolidate resolution of the identity second-order Møller-Plesset perturbation (DEC-LT-RIMP2) theory method
The Journal of Chemical Physics **146**, 044103 (2017); 10.1063/1.4973710

Stochastic multi-reference perturbation theory with application to the linearized coupled cluster method
The Journal of Chemical Physics **146**, 044107 (2017); 10.1063/1.4974177

Quartic scaling MP2 for solids: A highly parallelized algorithm in the plane wave basis
The Journal of Chemical Physics **146**, 104101 (2017); 10.1063/1.4976937

Nuclei-selected atomic-orbital response-theory formulation for the calculation of NMR shielding tensors using density-fitting
The Journal of Chemical Physics **145**, 234108 (2016); 10.1063/1.4972212

Efficient linear-scaling second-order Møller-Plesset perturbation theory: The divide–expand–consolidate RI-MP2 model

Pablo Baudin,[1,a)] Patrick Ettenhuber,[1] Simen Reine,[2] Kasper Kristensen,[1] and Thomas Kjærgaard[1,b)]
[1]*qLEAP Center for Theoretical Chemistry, Department of Chemistry, Aarhus University, Langelandsgade 140, DK-8000 Aarhus C, Denmark*
[2]*Centre for Theoretical and Computational Chemistry, Department of Chemistry, University of Oslo, P.O. Box 1033, N-1315 Blindern, Norway*

(Received 15 October 2015; accepted 13 January 2016; published online 1 February 2016)

The Resolution of the Identity second-order Møller-Plesset perturbation theory (RI-MP2) method is implemented within the linear-scaling Divide-Expand-Consolidate (DEC) framework. In a DEC calculation, the full molecular correlated calculation is replaced by a set of independent fragment calculations each using a subset of the total orbital space. The number of independent fragment calculations scales linearly with the system size, rendering the method linear-scaling and massively parallel. The DEC-RI-MP2 method can be viewed as an approximation to the DEC-MP2 method where the RI approximation is utilized in each fragment calculation. The individual fragment calculations scale with the fifth power of the fragment size for both methods. However, the DEC-RI-MP2 method has a reduced prefactor compared to DEC-MP2 and is well-suited for implementation on massively parallel supercomputers, as demonstrated by test calculations on a set of medium-sized molecules. The DEC error control ensures that the standard RI-MP2 energy can be obtained to the predefined precision. The errors associated with the RI and DEC approximations are compared, and it is shown that the DEC-RI-MP2 method can be applied to systems far beyond the ones that can be treated with a conventional RI-MP2 implementation. © *2016 AIP Publishing LLC.* [http://dx.doi.org/10.1063/1.4940732]

I. INTRODUCTION

The resolution-of-the-identity (RI) approximation[1–5] has emerged as an important tool to reduce the computational cost of second-order Møller-Plesset perturbation theory[6] (MP2) and related methods. The RI approximation was first applied to MP2 by Feyereisen *et al.*,[7] and subsequently implemented by different groups.[8–18] For reviews, the reader is referred to Refs. 19 and 20. In RI-MP2, also denoted density-fitting MP2 (DF-MP2), the four-center electron repulsion integrals (ERIs) are decomposed into a sum involving only two-center and three-center ERIs. The RI technique reduces the computational cost as well as the required memory considerably by removing the linear dependencies of the original atomic orbital (AO) product basis while maintaining reliable accuracies for practical chemical applications. The development of optimized auxiliary basis sets[21–25] has reduced the error introduced by the RI approximation, and the error is often judged to be so small that RI-MP2 is the recommended method for calculating energies of MP2 quality. It should also be noted that MP2 with Cholesky decomposition of the ERIs is in many ways similar to RI-MP2.[26–28]

For large molecular systems, conventional RI-MP2 implementations encounter a scaling wall, both memory- and time-wise. The correlation effects described by the RI-MP2 method are local, and the steep scaling of the method with the system size — $O(N^5)$ where N is a measure of the system size — is therefore unphysical. In the last decades, many groups have been developing alternative implementations of RI-MP2 in order to reduce the fifth order scaling using the locality of correlation effects (see Ref. 20 for a review of new developments within MP2 theory). One approach consists in compressing the number of wave function parameters required to describe the correlation energy by using projected atomic orbitals (PAOs) and pair natural orbitals (PNOs)[12,29] or orbital specific virtual orbitals (OSVs).[30] Ochsenfeld and co-workers have developed an AO-MP2 algorithm, where the orbital energy denominator is eliminated from the conventional molecular orbital (MO)-based MP2 energy expression by means of a Laplace transformation[31,32] to obtain a formulation in terms of AO integrals. Since AOs are local by construction, efficient integral screening may be performed.[33–35] Pioneer work in that direction was done by Ayala and Scuseria.[36] Another commonly used approach for carrying out approximations in a MP2 calculation relies on a physical fragmentation of the molecular system and performing standard canonical calculations for each of the fragments before collecting the information for the full system. This category of methods includes the divide and conquer (DC),[37] the Fragment Molecular Orbital (FMO) methods,[38–41] the Molecular Tailoring Approach (MTA),[42] and the systematic molecular fragmentation approach.[43]

Yet another category of methods for obtaining linear-scaling MP2 energies relies on a partitioning of the orbital

a)Electronic mail: pablo.baudin@chem.au.dk
b)Electronic mail: tkjaergaard@chem.au.dk

 © 2016 AIP Publishing LLC

space rather than on a physical fragmentation of the molecule. This category contains the cluster in a molecule (CIM) method,[44,45] the incremental scheme,[46] and our recently proposed local correlation method, the Divide-Expand-Consolidate (DEC) scheme.[47–52] In a DEC calculation, the correlation energy is expressed in terms of local molecular orbitals,[53–64] and the full molecular calculation is replaced by a set of independent fragment calculations. The number of fragments scales linearly with the system size, rendering the method linear-scaling and massively parallel. The linear-scaling and parallel performance of the DEC-MP2 model has been demonstrated recently.[65]

In this paper, the DEC scheme is applied in connection with the RI-MP2 method. The overall linear-scaling properties of the DEC-MP2 scheme are not affected when the RI approximation is applied in the fragment calculations. However, the DEC-RI-MP2 method has a reduced computational cost as well as reduced memory requirements compared to the DEC-MP2 method. The resulting DEC-RI-MP2 method thus provides an efficient linear-scaling and massively parallel algorithm for the calculation of MP2 energies with rigorous error control.

In a DEC-MP2 calculation, the error introduced compared to a canonical MP2 calculation is controlled by the fragment optimization threshold (FOT).[47,48] The canonical MP2 result is thus systematically approached when the FOT is tightened. Compared to the canonical MP2 energy, the DEC-RI-MP2 model contains an intrinsic DEC error governed by the FOT as well as an error associated with the RI approximation. In this work, we compare these two errors to analyze the performance of the DEC-RI-MP2 model. We also show that the DEC-RI-MP2 algorithm can be applied to systems that are much larger than the ones that can be treated using a standard RI-MP2 implementation.

The paper is organized as follows. In Section II, we present the basic equations of the DEC model and introduce the RI approximation. In Section III, we present numerical results and perform a detailed error and performance analysis for the DEC-RI-MP2 model. The parallel performance is discussed in Section IV, while Section V contains some conclusive remarks.

II. THEORY

A. The divide-expand-consolidate energy expression

The MP2 correlation energy $E_{\text{corr}}^{\text{MP2}}$ for a closed shell molecule may be expressed as[66]

$$E_{\text{corr}}^{\text{MP2}} = \sum_{ijab} t_{ij}^{ab} \left(2g_{aibj} - g_{biaj}\right), \tag{1}$$

where t_{ij}^{ab} are the MP2 doubles amplitudes. In this article, the indices i,j (a,b) refer to occupied (virtual) *localized* real HF orbitals, and g_{aibj} is a 4-center ERI in the local molecular orbital (MO) basis using the Mulliken notation. Assigning each orbital to an atomic site (given by the nuclear positions $P,Q,\ldots$), the summation over two occupied orbitals in Eq. (1) may be replaced by a summation over atomic sites and pair sites,

$$E_{\text{corr}}^{\text{MP2}} = \sum_{P} E_P + \sum_{P>Q} \Delta E_{PQ}. \tag{2}$$

The atomic fragment energy E_P and the pair interaction energy ΔE_{PQ} are defined according to

$$E_P = \sum_{ij\in\underline{P}} \sum_{ab} t_{ij}^{ab}(2g_{aibj} - g_{biaj}), \tag{3}$$

$$\Delta E_{PQ} = \sum_{\substack{i\in\underline{P}\\ j\in\underline{Q}}} \sum_{ab} t_{ij}^{ab}(2g_{aibj} - g_{biaj}) + \sum_{\substack{i\in\underline{Q}\\ j\in\underline{P}}} \sum_{ab} t_{ij}^{ab}(2g_{aibj} - g_{biaj}), \tag{4}$$

where $\underline{P}$ denotes the set of occupied orbitals assigned to atomic site P.

Using a local HF orbital basis, the free summations over virtual orbitals ab in the atomic fragment energy E_P may be restricted. This is justified by the fact that the integrals g_{aibj} (with $ij \in \underline{P}$) vanish for virtual orbitals spatially far from the atomic site P. We introduce the notation $[\overline{P}]$ for the set of virtual orbitals spatially close to atomic site P in the sense that the integrals g_{aibj} are non-vanishing (to the desired precision), and the virtual summations in Eq. (3) may thus be restricted to the $[\overline{P}]$ space. A similar replacement leading to the union of orbital spaces is introduced for the pair interaction energy ΔE_{PQ},[48] and Eqs. (3) and (4) may thus be written as

$$E_P = \sum_{ij\in\underline{P}} \sum_{ab\in[\overline{P}]} t_{ij}^{ab}(2g_{aibj} - g_{biaj}), \tag{5}$$

$$\Delta E_{PQ} = \sum_{ab\in[\overline{P}]\cup[\overline{Q}]} \Bigg(\sum_{\substack{i\in\underline{P}\\ j\in\underline{Q}}} t_{ij}^{ab}(2g_{aibj} - g_{biaj}) + \sum_{\substack{i\in\underline{Q}\\ j\in\underline{P}}} t_{ij}^{ab}(2g_{aibj} - g_{biaj}) \Bigg). \tag{6}$$

The details on how each local molecular orbital is assigned to atomic sites and how the $[\overline{P}]$ space is obtained is not the subject of this paper, and the reader is referred to Ref. 48. Here we just note that the $[\overline{P}]$ space is determined in a black box manner such that the error of E_P is smaller than the FOT, which is the central threshold that defines the precision of a DEC calculation. Naturally, in the limit where the space $[\overline{P}]$ includes all virtual orbitals for all atomic sites P, the conventional MP2 correlation energy is recovered. The approximations introduced in Eqs. (5) and (6) are summarized in Section II C along with other approximations of the DEC-RI-MP2 scheme.

The set of orbitals used to evaluate the fragment energy is denoted the *energy orbital space* (EOS), i.e., for atomic fragment P in Eq. (5) this corresponds to $\underline{P} \cup [\overline{P}]$. In order to describe coupling effects between the amplitudes in the EOS and the amplitudes outside the EOS, the MP2 equations (Eq. (7)) have to be solved in an extended space, and $[\underline{P}] \cup [\overline{P}]$ denoted the *amplitude orbital space* (AOS), where $[\underline{P}]$ includes $\underline{P}$ as well as additional occupied orbitals involved in the coupling mechanism. As for the $[\overline{P}]$ space, the practical determination of $[\underline{P}]$ is described in Ref. 48.

Equation (5) requires amplitudes and integrals in the local basis. However, in practice, the solution of the amplitude

equation in the AOS is simplified by transforming the subset of localized HF orbitals to a pseudocanonical basis,[67] which is defined by diagonalizing the local Fock matrix blocks F_{ij} ($ij \in [\underline{P}]$) and F_{ab} ($ab \in [\overline{P}]$). The amplitudes are then constructed directly in a pseudocanonical basis according to

$$t_{IJ}^{AB} = -g_{AIBJ}(\epsilon_A + \epsilon_B - \epsilon_I - \epsilon_J)^{-1}, \tag{7}$$

where I, J (A, B) are occupied (virtual) MO indices in the pseudocanonical basis, and ϵ_I, ϵ_J, ϵ_A, and ϵ_B are the diagonal Fock matrix elements in the pseudocanonical basis. The amplitudes are later transformed to the local basis where the energy is evaluated using Eq. (5). This transformation is necessary since the occupied summation restriction in Eq. (5) is only defined in the local basis.

The DEC algorithm described above scales quadratically with system size due to the number of pair calculations in Eq. (2). However, the pair energies describe dispersion interactions decaying with the inverse pair distance to the sixth power. Distant pairs with small energy contributions may therefore be neglected without affecting the precision of the total correlation energy, which already contains an error of size FOT for each atomic fragment. In the simplest approach, a distance-based cutoff can be used to determine which pairs to include.[48] A more elaborate scheme based on approximate pair energy contributions will be described in a forthcoming paper.

The DEC-MP2 algorithm may be summarized as follows:

1. Determine localized occupied and virtual HF molecular orbitals.
2. For each atomic fragment P, determine the optimized orbital spaces $[\overline{P}]$ and $[\underline{P}]$ (the AOS) as detailed in Ref. 48. This step also provides all atomic fragment energies E_P. Each fragment calculation is carried out using the following procedure:
   (a) Transform the local HF orbitals of the AOS into a pseudocanonical basis in order to generate the doubles amplitudes using Eq. (7).
   (b) Transform the integrals and amplitudes back to the local HF basis and extract their EOS contributions.
   (c) Evaluate the fragment energy using Eq. (5).
3. Pair screening: Use pair energy estimates to screen away pairs with negligible contributions (detailed in a forthcoming paper) to get a list of the important pair fragments.
4. For each important pair fragment PQ, calculate ΔE_{PQ} according to steps 2(a)-2(c) above, where Eq. (6) is used in step 2(c).
5. Add up the fragment energies to obtain the total DEC-MP2 energy using Eq. (2).

B. The resolution of the identity within a DEC framework

The resolution of the identity can be expressed in its standard V approximation[5] using the symmetric decomposition,[68]

$$g_{aibj} \approx \sum_{\alpha\beta}(ai|\alpha)(\alpha|\beta)^{-1}(\beta|bj),$$
$$g_{aibj} \approx \sum_{\alpha\beta\gamma}(ai|\alpha)(\alpha|\gamma)^{-1/2}(\gamma|\beta)^{-1/2}(\beta|bj) = \sum_{\gamma} C_{ai}^{\gamma} C_{bj}^{\gamma}, \tag{8}$$

where $(ai|\alpha)$ is a 3-center ERI, $(\alpha|\beta)$ is a 2-center ERI, and $C_{ai}^{\gamma} = \sum_{\alpha}(ai|\alpha)(\alpha|\gamma)^{-1/2}$. We let $\{\alpha, \beta, \gamma\}$ refer to auxiliary AO indices, while $\{\mu, \nu\}$ are used for standard AO indices. We can directly apply the DEC-MP2 algorithm summarized in Section II A to the DEC-RI-MP2 model by calculating 2-electron integrals using Eq. (8).

Algorithm 1 shows in detail how the EOS amplitudes and EOS integrals entering Eqs. (5) and (6) are determined in the DEC-RI-MP2 method. Here, the i, j, a, b indices refer to localized HF orbitals in the EOS, while the I, J, A, B indices label pseudocanonical orbitals in the AOS. We use $C_{\mu I}$ ($C_{\mu A}$) to denote elements of the transformation matrix from the AOs to the occupied (virtual) pseudocanonical orbitals, while U_{Ii} (U_{Aa}) represents a transformation from the occupied (virtual) pseudocanonical orbitals to the occupied (virtual) local orbitals. The cost of the fragment calculation is determined by the number of occupied MOs in the EOS (O_{EOS}), occupied MOs in the AOS (O_{AOS}), virtual MOs in the AOS (V_{AOS}), standard AOs ($N_{\mathrm{AO,AOS}}$), and auxiliary AOs ($N_{\mathrm{aux,AOS}}$). Note that these dimensions all refer to a fragment (EOS or AOS subscript), not the full molecular system. The determination of $N_{\mathrm{AO,AOS}}$ and $N_{\mathrm{aux,AOS}}$ is discussed below. The scaling of each step is shown in Algorithm 1.

ALGORITHM 1. DEC-RI-MP2 algorithm for calculating EOS amplitudes t_{ij}^{ab} and EOS integrals g_{aibj}. The scaling with the fragment size of the time-dominating step is given in parentheses for each term. Bold: Done by local master using OpenMP parallelization. Normal font: MPI-parallelized across n nodes, where each MPI process utilizes OpenMP parallelization. The notation is described in the text.

1 Calculate $(\alpha\|\mu\nu)$	$(N_{\mathrm{aux,AOS}} N_{\mathrm{AO,AOS}}^2)$
2 Calculate $(\alpha\|\beta)$	$\boldsymbol{(N_{\mathrm{aux,AOS}}^2)}$
3 Construct $(\alpha\|\beta)^{-\frac{1}{2}}$	$\boldsymbol{(N_{\mathrm{aux,AOS}}^3)}$
4 $(\beta\|AI) = \sum_{\mu} C_{\mu A}\left[\sum_{\nu} C_{\nu I}(\beta\|\mu\nu)\right]$	$(N_{\mathrm{aux,AOS}} N_{\mathrm{AO,AOS}}^2 O_{\mathrm{AOS}})$
5 $C_{AI}^{\alpha} = \sum_{\beta}(\alpha\|\beta)^{-\frac{1}{2}}(\beta\|AI)$	$(N_{\mathrm{aux,AOS}}^2 V_{\mathrm{AOS}} O_{\mathrm{AOS}})$
6 $t_{IJ}^{AB} = \sum_{\alpha} \frac{C_{AI}^{\alpha} C_{BJ}^{\alpha}}{\epsilon_I + \epsilon_J - \epsilon_A - \epsilon_B}$	$(N_{\mathrm{aux,AOS}} V_{\mathrm{AOS}}^2 O_{\mathrm{AOS}}^2)$
7 $t_{ij}^{ab} = \sum_{A} U_{aA}\left(\sum_{B} U_{bB}\left[\sum_{I} U_{iI}\left(\sum_{J} U_{jJ} t_{IJ}^{AB}\right)\right]\right)$	$(O_{\mathrm{EOS}} V_{\mathrm{AOS}}^2 O_{\mathrm{AOS}}^2)$
8 $C_{ai}^{\alpha} = \sum_{A} U_{aA}\left(\sum_{I} U_{iI} C_{AI}^{\alpha}\right)$	$(N_{\mathrm{aux,AOS}} V_{\mathrm{AOS}} O_{\mathrm{AOS}})$
9 $g_{aibj} = \sum_{\alpha} C_{ai}^{\alpha} C_{bj}^{\alpha}$	$(N_{\mathrm{aux,AOS}} V_{\mathrm{AOS}}^2)$

The memory requirements of the standard RI-MP2 method are very small, since the method does not require the storage of the full doubles amplitudes, but only requires the fitting coefficients ($N_{\text{aux}}^2 VO$) for the full molecular system. However, the DEC-RI-MP2 method additionally requires a transformation from the pseudocanonical basis to the local basis. In practice, pseudocanonical amplitudes t_{IJ}^{AB} are generated and immediately transformed to the local EOS (first transformation in step 7 in Algorithm 1), $t_{IJ}^{AB} \rightarrow t_{iJ}^{AB}$, and the t_{iJ}^{AB} amplitudes are stored before they are fully transformed to the local basis ($t_{iJ}^{AB} \rightarrow t_{ij}^{ab}$). The memory requirements are therefore $V_{\text{AOS}}^2 O_{\text{AOS}} O_{\text{EOS}}$ in addition to the $N_{\text{aux,AOS}}^2 V_{\text{AOS}} O_{\text{AOS}}$ requirements for the fitting coefficients. Thus, although the memory requirements of the DEC-RI-MP2 scheme are slightly more involved than for the conventional RI-MP2 method, the storage of doubles amplitudes with four AOS indices (t_{IJ}^{AB}) is avoided in the DEC-RI-MP2 scheme, which also does not use any I/O. We also note that, since the doubles amplitudes are stored in the local basis (t_{ij}^{ab}), it is also possible to construct the MP2 density, molecular gradient, electrostatic potential, etc., within the DEC framework.[49,51,69] Once the EOS amplitudes and EOS integrals have been determined, the actual evaluation of the fragment energy (Eq. (5) or (6)) is a minor task, which is not shown in Algorithm 1.

The parallelization of the DEC-RI-MP2 scheme is discussed in detail in Section IV. For now, we just note that each fragment calculation is parallelized over n compute nodes, where one of the nodes is assigned to be the local master. In Algorithm 1, the steps that are performed only by the local master are shown in bold, while the steps which are parallelized across the n nodes using the message passing interface (MPI) are shown in regular font.

The number of occupied (O_{AOS}) and virtual (V_{AOS}) orbitals assigned to a fragment is defined by the fragment optimization procedure.[48] To determine the number of AOs in a fragment, $N_{\text{AO,AOS}}$, we need to consider the expansion of a localized MO assigned to atomic site P, ϕ_r^P, in terms of atomic orbitals χ_μ and MO coefficients $c_{\mu r}^P$,

$$\phi_r^P = \sum_\mu \chi_\mu c_{\mu r}^P, \tag{9}$$

where r is a general MO label which may refer to either an occupied or a virtual orbital. Even though the bulk of the localized MO ϕ_r^P is confined to a small volume of space, the localized MO has small tail coefficients far from P. To reduce the number of 4-center AO integrals to be determined, we therefore introduce an approximate MO $\tilde{\phi}_r^P$,

$$\tilde{\phi}_r^P = \sum_{\mu \in \{P\}} \chi_\mu \tilde{c}_{\mu r}^P, \tag{10}$$

where $\{P\}$ is the so-called atomic extent (AE) which is a subset of atomic orbitals close to atomic site P, and the $\tilde{c}_{\mu r}^P$ coefficients are determined such that $\tilde{\phi}_r^P$ resembles ϕ_r^P as much as possible in a least squares sense. The practical determination of $\{P\}$ and $\tilde{c}_{\mu r}^P$ is detailed in Ref. 48. The number of AOs ($N_{\text{AO,AOS}}$) thus corresponds to the number of atomic basis functions included in $\{P\}$.

Concerning the set of auxiliary AOs used in the RI approximation, we investigate two choices:

A: Include the full auxiliary AO basis set in all fragment calculations.

B: Include auxiliary AO basis functions on atoms of the atomic extent (e.g., for atomic fragment P, include auxiliary AOs on atoms in the $\{P\}$ space, see Eq. (10)).

Option **A** is only interesting for analysis purposes, since it would destroy the linear-scaling of the DEC-RI-MP2 algorithm ($N_{\text{aux,AOS}}$ in Algorithm 1 would be the full molecular system). Option **B** is the practical choice and is appealing due to its simplicity and since it ensures a linear-scaling DEC-RI-MP2 algorithm. Naturally, both options will lead to the full RI-MP2 result in the limit where the FOT approaches zero. The approximation associated with option **B** is expected to be small, since, for an atomic fragment P, only ERIs g_{aibj} ($i, j \in \underline{P}$), where the local virtual orbitals a and b are far from P, will be poorly described using the RI approximation, and such integrals have very small energy contributions. In essence, option **B** corresponds to a local domain fitting procedure, and local domain fitting has previously been used successfully for MP2.[12,70,71] Options **A** and **B** are compared in Section III C.

C. DEC-RI-MP2 approximations

In order to investigate the efficiency of the DEC-RI-MP2 model, it is convenient to summarize the approximations introduced so far:

Approx AOS: The amplitude equations are solved in the restricted AOS, and the virtual orbital summations are restricted in Eqs. (5) and (6).

Approx AE: The orbitals in the AOS are spanned by the restricted set of AOs in the atomic extent.

Approx PAIR: Distant pair interaction energies are neglected.

These approximations are present for all wave function models within the DEC framework, while the following two approximations arise from the introduction of the resolution of the identity:

Approx AAE: The auxiliary AOs are restricted to the atomic extent (option **B**).

Approx RI: The RI approximation.

Approx AE, **Approx AOS**, **Approx PAIR**, and **Approx AAE** are DEC specific, while **Approx RI** represents the intrinsic RI error. When the FOT is tightened, the errors associated with **Approx AE**, **Approx AOS**, **Approx PAIR**, and **Approx AAE** all decrease. However, the RI error persists. Thus, as the FOT approaches zero, the DEC-RI-MP2 energy approaches the RI-MP2 energy, not the canonical MP2 energy. In Section III B, we investigate the convergence of the DEC-RI-MP2 with the FOT and compare the DEC specific errors to the intrinsic RI error.

III. NUMERICAL RESULTS

A. Computational details and molecular systems

In order to investigate the performance of the DEC-RI-MP2 algorithm both in terms of errors in the total correlation energy and in terms of timings, we have selected a set of medium to large molecules (see supplementary material for the molecular geometries[72]). All calculations have been performed using Dunning's correlation consistent cc-pVTZ basis set[73] and the cc-pVTZ-RI auxiliary basis set.[23,74,75] The considered molecular systems are as follows:

(a) Coronene ($C_{24}H_{12}$): 78 occupied orbitals, 810 virtual orbitals, 888 atomic basis functions, 2304 auxiliary basis functions.
(b) Tetrahexacontanoic acid ($C_{64}O_2H_{128}$): 264 occupied orbitals, 3508 virtual orbitals, 3773 atomic basis functions, 9186 auxiliary basis functions.
(c) Heptapeptide (Asn-Phe-Gly-Ala-Ile-Leu-Ser): 208 occupied orbitals, 2240 virtual orbitals, 2448 atomic basis functions, 6155 auxiliary basis functions.
(d) Valinomycin ($C_{54}H_{90}N_6O_{18}$): 300 occupied orbitals, 3300 virtual orbitals, 3600 atomic basis functions, 9018 auxiliary basis functions.

These molecules have rather different chemical structures, including a delocalized aromatic electronic structure (coronene), a highly one-dimensional molecule (tetrahexacontanoic acid), and a more compact structure (valinomycin). Furthermore, all considered molecules are rather large, since the DEC-RI-MP2 scheme is only useful for large molecular systems, as we discuss will in detail in Section III D.

All calculations have been performed using a local version of the LSDalton program, and the methods used in this paper are part of the kernel of the Dalton2016 suite.[76,77]

B. The DEC correlation energy

In this section, we want to study the effect of the different approximations (RI and DEC) on the canonical MP2 correlation energy. We therefore introduce the following notations to quantify energy errors:

$$\delta\text{RI} = \left|E_{\text{corr}}^{\text{DEC-RI-MP2}} - E_{\text{corr}}^{\text{DEC-MP2}}\right|, \tag{11}$$

$$\delta\text{DEC} = \left|E_{\text{corr}}^{\text{DEC-RI-MP2}} - E_{\text{corr}}^{\text{RI-MP2}}\right|, \tag{12}$$

$$\delta\text{DEC-RI} = \left|E_{\text{corr}}^{\text{DEC-RI-MP2}} - E_{\text{corr}}^{\text{MP2}}\right|, \tag{13}$$

where δRI denote the error associated with the RI approximation, δDEC the error associated with the DEC scheme, and δDEC-RI the error associated with both RI and DEC. We also introduce the DEC recovery of the canonical RI-MP2 correlation energy,

$$\Delta\text{DEC} = \frac{E_{\text{corr}}^{\text{DEC-RI-MP2}}}{E_{\text{corr}}^{\text{RI-MP2}}}. \tag{14}$$

The error associated with both RI and DEC (δDEC-RI) are only reported for the coronene system as the other systems were too large to perform a canonical MP2 calculation of triple zeta quality. Furthermore, the RI errors (δRI) are not reported for the heptapeptide and the valinomycin systems (the DEC-MP2 calculations can in principle be performed, but for these systems they are very expensive for tight FOT values).

In Tables I–IV, we examine the convergence of the total correlation energy as the FOT is tightened. As for the DEC-MP2 scheme, the error δDEC in the total energy decreases by roughly an order of magnitude when decreasing the FOT by an order of magnitude, demonstrating that the combined DEC error associated with **Approx AE**, **Approx AOS**, **Approx PAIR**, and **Approx AAE** systematically decreases with the FOT. Thus, although in principle there are several approximations in DEC, they are all controlled by one threshold.

From Table I, we note that the RI error (δRI) of the DEC-RI-MP2 calculation converges to the error associated with the RI approximation for the canonical calculations on coronene ($5.55 \cdot 10^{-4}$ a.u.) when tightening the FOT. The same observation can be made for the carbon chain in Table II even though the exact number could not be calculated. It is important to realize that at approximately FOT = 10^{-5} a.u., the DEC-RI-MP2 calculation is dominated by the RI error (**Approx RI**) and further tightening of the FOT is futile. We also note that the DEC-RI-MP2 model provides results of similar accuracy for the different systems under consideration in Tables I–IV, i.e., independently of the spatial structure (1, 2, or 3-dimensional) and the chemical structure (conjugated or not conjugated system).

C. Selection of the auxiliary basis functions (Approx AAE)

Having demonstrated the general convergence of the DEC-RI-MP2 energy with the FOT, we now examine **Approx AAE** in more detail. In Table V, we compare option **A** (include

TABLE I. Total errors (δDEC, δRI, and δDEC-RI, a.u.) and recovery (ΔDEC, %) in the total correlation energy of the coronene molecule for different values of the FOT (a.u.). Speed up of DEC-RI-MP2 with respect to canonical RI-MP2 is also reported.

FOT	δDEC	δRI	δDEC-RI	ΔDEC	Speed up
$1.0 \cdot 10^{-3}$	$7.576 \cdot 10^{-2}$	$5.255 \cdot 10^{-4}$	$7.632 \cdot 10^{-2}$	98.153	0.127
$1.0 \cdot 10^{-4}$	$6.475 \cdot 10^{-3}$	$5.598 \cdot 10^{-4}$	$7.032 \cdot 10^{-3}$	99.842	0.045
$1.0 \cdot 10^{-5}$	$4.430 \cdot 10^{-4}$	$5.703 \cdot 10^{-4}$	$1.000 \cdot 10^{-3}$	99.989	0.029
$1.0 \cdot 10^{-6}$	$2.447 \cdot 10^{-5}$	$5.555 \cdot 10^{-4}$	$5.816 \cdot 10^{-4}$	99.999	0.023
$1.0 \cdot 10^{-7}$	$8.749 \cdot 10^{-7}$	$5.555 \cdot 10^{-4}$	$5.580 \cdot 10^{-4}$	100.00	0.022

TABLE II. Total errors (δDEC and δRI, a.u.) and recovery (ΔDEC, %) in the total correlation energy of the tetrahexacontanoic acid for different values of the FOT (a.u.). Speed up of DEC-RI-MP2 with respect to canonical RI-MP2 is also reported.

FOT	δDEC	δRI	ΔDEC	Speed up
$1.0 \cdot 10^{-3}$	$1.285 \cdot 10^{-1}$	$1.823 \cdot 10^{-3}$	98.986	25.8
$5.0 \cdot 10^{-4}$	$6.492 \cdot 10^{-2}$	$2.416 \cdot 10^{-3}$	99.488	24.8
$1.0 \cdot 10^{-4}$	$1.414 \cdot 10^{-2}$	$2.066 \cdot 10^{-3}$	99.888	17.2
$5.0 \cdot 10^{-5}$	$8.393 \cdot 10^{-3}$	$1.918 \cdot 10^{-3}$	99.934	14.6
$1.0 \cdot 10^{-5}$	$1.821 \cdot 10^{-3}$	$2.050 \cdot 10^{-3}$	99.986	6.7
$5.0 \cdot 10^{-6}$	$8.520 \cdot 10^{-4}$	$2.051 \cdot 10^{-3}$	99.993	3.7

all auxiliary functions) and option **B** (include auxiliary functions associated with atomic centers in the atomic extent $\{P\}$) for tetrahexacontanoic acid. The difference between options **A** and **B** (δ_{B-A} in Table V) should be compared with the DEC error for this molecule in Table II. It is clear that the error related to option **B** is minuscule compared to the error associated with the DEC partitioning scheme, and option **B** is therefore well-justified. Note that option **B** is the practical implementation used for all other DEC-RI-MP2 calculations reported in this paper.

The numbers in Table V indicate that it is in fact possible to further reduce the size of the auxiliary space. For the FOT = 10^{-5} a.u. calculation, the largest and smallest fitting domains contained 3197 and 1147 auxiliary functions, respectively. These fitting domains are quite large compared to what has typically been used in the literature,[12,14,70,71,78,79] and the δ_{B-A} numbers in Table V strongly support the conclusion that smaller auxiliary domains can be used without affecting the precision. In the present implementation, we have used option **B** due to the simplicity of the approach, but we are currently investigating alternative choices leading to smaller auxiliary domains. Possible options include the use of Natural Auxiliary Functions (NAFs),[80] or 3 center integral screening techniques such as the SQVℓ[81] or Löwdin charges.[12]

D. Performance analysis

The overall goal of DEC is to provide a framework to calculate energies and properties at the CC level of theory for molecular systems where a conventional implementation hits a scaling wall. Furthermore, the DEC scheme is designed to utilize the thousands of computing cores available on modern supercomputers. One of the consequences of the massively parallel character of DEC is a large amount of recalculations, which makes it vastly inefficient for small molecular systems compared to a conventional implementation. In this section we analyze the general performance of the DEC-RI-MP2 algorithm and try to give simple rule of thumb for the general user to choose between DEC or conventional implementations to treat a given problem. We have therefore reported speed ups of the DEC-RI-MP2 simulations with respect to the canonical RI-MP2 calculations for both coronene and the tetrahexacontanoic acid (see Tables I and II).

The calculations on coronene were performed on the Eos cluster at Oak Ridge National Laboratory (ORNL).[82] The speed ups given in Table I show that the coronene molecule is clearly too small to demonstrate the usefulness of the DEC algorithm, i.e., the crossover between a conventional implementation and the DEC scheme occurs for a system that is larger than coronene. Indeed, in the case of a small molecular system and a tight FOT value, the DEC algorithm is basically repeating a conventional calculation using (almost) the full orbital space for each fragment, and consequently no speed ups are observed in Table I. For coronene, there are 300 fragments if all pair fragments are included. In the limit where the FOT approaches zero, a DEC calculation would therefore be 300 times slower than a canonical RI-MP2 calculation. However, because the fragment calculations are completely independent, it is always possible to bring down the time to solution of a DEC calculation by using many nodes.

For the tetrahexacontanoic acid (Table II), the calculations were performed using our local cluster.[83] While using DEC on the coronene system is clearly unfavorable, the situation is very different for the highly linear tetrahexacontanoic acid. It is seen that even without using the massively parallel feature of DEC (all independent calculations are done one after an other), the DEC-RI-MP2 simulation provides shorter time to solution than the standard method for all FOTs in Table II. We also note that at approximately FOT = 10^{-5} a.u., the DEC-RI-MP2 calculation starts to be dominated by the RI error, and the DEC-RI-MP2 calculation is still 6.7 times faster than the conventional RI-MP2 calculation.

TABLE III. Total errors (δDEC, a.u.) and recovery (ΔDEC, %) in the total correlation energy of the heptapeptide for different values of the FOT (a.u.).

FOT	δDEC	ΔDEC
$1.0 \cdot 10^{-3}$	$1.162 \cdot 10^{-1}$	98.944
$1.0 \cdot 10^{-4}$	$1.455 \cdot 10^{-2}$	99.868
$1.0 \cdot 10^{-5}$	$1.575 \cdot 10^{-3}$	99.986
$1.0 \cdot 10^{-6}$	$1.606 \cdot 10^{-4}$	99.999

TABLE IV. Total errors (δDEC, a.u.) and recovery (ΔDEC, %) in the total correlation energy of the valinomycin molecule for different values of the FOT (a.u.).

FOT	δDEC	ΔDEC
$1.0 \cdot 10^{-3}$	$2.349 \cdot 10^{-1}$	98.512
$5.0 \cdot 10^{-4}$	$1.205 \cdot 10^{-1}$	99.237
$1.0 \cdot 10^{-4}$	$3.064 \cdot 10^{-2}$	99.806
$5.0 \cdot 10^{-5}$	$1.774 \cdot 10^{-2}$	99.888

TABLE V. Comparison of options **A** and **B** concerning the choice of auxiliary basis functions for tetrahexacontanoic acid using a cc-pVTZ(cc-pVTZ-RI) basis.

FOT	δ_{B-A}[a]
$1.0 \cdot 10^{-3}$	$4.1 \cdot 10^{-6}$
$1.0 \cdot 10^{-4}$	$6.1 \cdot 10^{-7}$
$1.0 \cdot 10^{-5}$	$5.3 \cdot 10^{-8}$

[a]Difference in total correlation energy between options **A** and **B**.

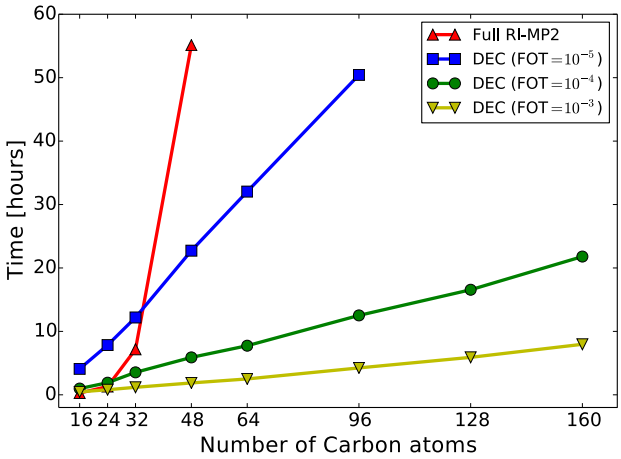

FIG. 1. Comparison of the linear-scaling DEC-RI-MP2 method and the N^5 scaling standard canonical RI-MP2 method for a set of alkane chains of increasing length. The cc-pVTZ(cc-pVTZ-RI) basis was employed.[83]

In order to further illustrate the crossover between conventional and DEC implementations as well as to demonstrate the linear-scaling behavior of the DEC-RI-MP2 code, we have performed DEC-RI-MP2/cc-pVTZ(cc-pVTZ-RI) calculations on alkane chains of increasing sizes, from $C_{16}H_{34}$ (956 basis functions) to $C_{160}H_{322}$ (9308 basis functions). The results are given in Figure 1 for calculations performed using our local cluster.[83] It is seen that the DEC-RI-MP2 method indeed is scaling linearly with the system size (N), while the standard RI-MP2 method scales as N^5. Obviously, tightening the FOT increases the prefactor, but it does not change the overall linear-scaling behavior of the method. Of course, the linear-scaling behavior emerges quite early for these highly linear systems, while 3-dimensional systems would enter the linear-scaling regime for larger system sizes.

The timings reported in Tables I and II and Figure 1 should guide the general user to choose between the conventional and DEC methods for RI-MP2 calculations. For highly linear systems, the DEC algorithm in general outperforms the conventional implementation. For more complicated structures, the conventional algorithm is to be preferred if it is at all feasible. However, when the canonical implementation hits the scaling wall, the DEC scheme is still feasible and becomes the method of choice. In particular, if many compute nodes are available, the DEC algorithm always provides a very short time to solution.

IV. PARALLELIZATION OF THE DEC SCHEME

In this section, we summarize the DEC parallelization strategy[65] and provide numerical results to investigate the parallel performance of the DEC-RI-MP2 scheme. As a test system, we use a cluster of 200 water molecules, $(H_2O)_{200}$, and the standard AO and auxiliary basis sets, cc-pVTZ and cc-pVTZ-RI, respectively. The calculations have been performed using a FOT = 10^{-4} a.u and the Titan supercomputing system at ORNL.[84] In DEC, the different fragment calculations may be run in parallel because they are independent of each other. We refer to this as *coarse grained parallelization*, while the parallelization of the individual fragment calculations is referred to as *medium grained parallelization* (Algorithm 1). The coarse grained parallelization is common to all wave function models (RI-MP2, MP2, CCSD, CCSD(T), . . .), while the medium grained parallelization is dependent on the model of choice. At the coarse grained level of parallelization, one MPI process called the global master dynamically distributes the fragment jobs to the local masters according to a list where the computational demands of each fragment is prioritized (largest fragments are calculated first). At the medium grained level, each local master has a set of associated local slaves, which all together define a fragment slot. Each slot carries out the fragment calculation and sends back the fragment energy to the global master who adds up the fragment energy contributions. Since the fragment calculations are of very different size, we have chosen an approach where the slots divide dynamically to ensure a good medium grained parallelization. The medium grained parallelization is investigated in Section IV A, while the coarse grained parallel performance is discussed in Section IV B.

A. Medium grained parallelization

The *medium grained parallelization* of the DEC-RI-MP2 method is very similar to the parallelization of the standard canonical RI-MP2 calculation and as such several options are possible. Katouda and Nakajima[10] performed the MPI work distribution based on the set of virtual orbitals, while in a previous paper,[9] the MPI parallelization was based on the set of occupied orbitals. Hättig *et al.*[85] concluded that the time-determining steps of RI-MP2 was most efficiently parallelized over the pairs of occupied orbital indices, because a parallelization over auxiliary basis functions would require the communication of 4-index MO integrals and thus require transfer rates which can only be reached with high performance networks. However, in connection with the DEC-RI-MP2 code, the 4-index MO integrals and amplitudes that must be communicated are the ones residing in the small EOS space. The *medium grained parallelization* of the DEC-RI-MP2 model was therefore chosen to be based on a parallelization over auxiliary basis functions. Each node in the slot is thus simply assigned a subset of the auxiliary basis functions in each of the MPI-parallelized steps of Algorithm 1.

The individual fragment calculations can have significantly different sizes depending on the chemical environment. For example, the smallest pair fragment in the water cluster calculation (FOT = 10^{-4} a.u.), contained 17 occupied MOs, 222 virtual MOs, and 416 AOs, while the biggest fragment contained 73 occupied MOs, 902 virtual MOs, and 1928 AOs. The computational cost of the individual fragment calculations is the standard $O(N_{\text{frag}}^5)$ scaling, where N_{frag} is a measure of the fragment size (see Algorithm 1). These differences in fragment sizes thus lead to huge differences in terms of computational requirements.

To minimize the time-to-solution for the total calculation, it is important to know how many nodes can be used efficiently for a given fragment at the medium grained level of parallelization. In order to investigate the medium grained

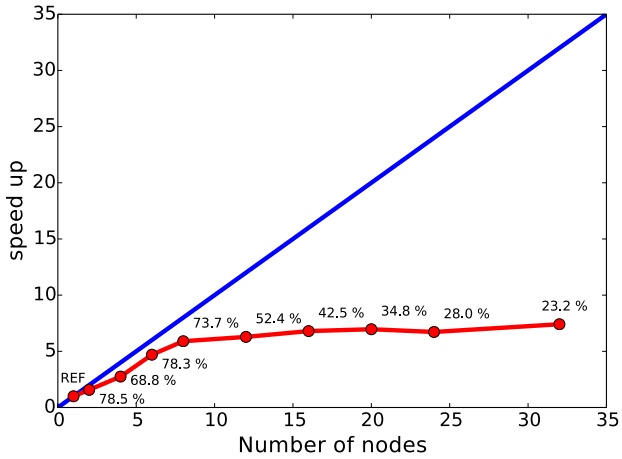

FIG. 2. Medium grained scaling: Speed up for a pair fragment consisting of 38 occupied orbitals, 350 virtual orbitals, 690 basis functions, and 1923 auxiliary basis functions in a DEC-RI-MP2/cc-pVTZ(cc-pVTZ-RI) energy calculation on a $(H_2O)_{200}$ cluster using FOT = 10^{-4} a.u. (The blue line represents ideal scaling and the relative speed up (in %) compared to ideal behavior are given for each point.)

parallel performance, we present relative timings for two fragments of very different sizes. In Figure 2, the relative timings are given for a small pair fragment (with 38 occupied orbitals, 350 virtual orbitals, 690 basis functions, and 1923 auxiliary basis functions) of the $(H_2O)_{200}$ calculation, while Figure 3 shows the relative timings for a large pair fragment with 69 occupied orbitals, 723 virtual orbitals, 1397 basis functions, and 3999 auxiliary basis functions.

For the small fragment, the performance is quickly saturated and for six nodes, we obtain 78.3% of the idealized speed up. For the large fragment, the parallel performance is superlinear with the number of nodes. The reason for this superlinear behaviour is that the fitting coefficients C_{bj}^{γ} have been distributed among the nodes, and this reduces the datasize and improves cache memory performance. The

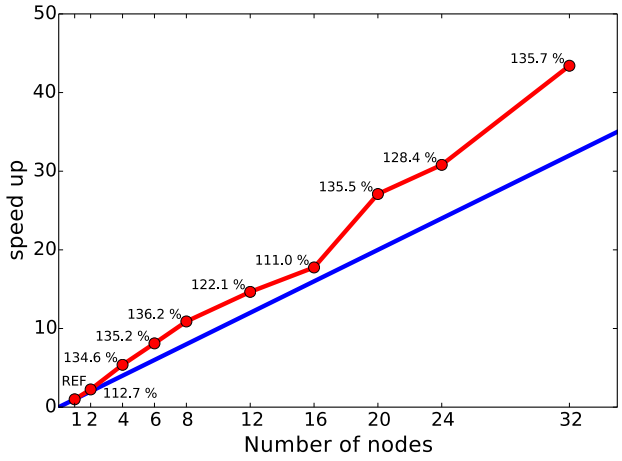

FIG. 3. Medium grained scaling: Speed up for a pair fragment consisting of 69 occupied orbitals, 723 virtual orbitals, 1397 basis functions, and 3999 auxiliary basis functions in a DEC-RI-MP2/cc-pVTZ(cc-pVTZ-RI) calculation on a $(H_2O)_{200}$ cluster using FOT = 10^{-4} a.u. (The blue line represents ideal scaling and the relative times (in %) compared to ideal behavior are given for each point.)

smaller memory requirements on the individual node also allow for more efficient AO to MO transformations.

The fragment size, which can be quantified in terms of $N_{\text{aux,AOS}}$, O_{AOS}, and V_{AOS}, thus determines the efficiency of the parallelization in Algorithm 1. The most expensive computational step (step 6 of the Algorithm 1) of a fragment calculation scales as $N_{\text{aux,AOS}}V_{\text{AOS}}^2O_{\text{AOS}}^2$, while, in the most expensive communication step, the 3 center ERIs $(\alpha|AI)$ of size $N_{\text{aux,AOS}}V_{\text{AOS}}O_{\text{AOS}}$ are communicated among n nodes. We therefore require that the ratio between the computationally most expensive step and the most expensive communication step is large. Thus, roughly speaking, the larger $N_{\text{aux,AOS}}V_{\text{AOS}}^2O_{\text{AOS}}^2/(N_{\text{aux,AOS}}V_{\text{AOS}}O_{\text{AOS}}) = O_{\text{AOS}}V_{\text{AOS}}$ is, the more nodes can efficiently be used in Algorithm 1. In practice, we use the following condition to determine whether a slot containing n nodes should divide or not:

$$\text{Divide slot if: } n > O_{\text{AOS}}V_{\text{AOS}}/X. \tag{15}$$

The large number of fragments that must be calculated in a DEC calculation allows us to use a conservative number of nodes for each fragment, and empirical investigations have established that $X = 8000$ is a reasonable value. For example, if a slot of 4 nodes receives a fragment job where Eq. (15) is satisfied, the slot would divide into two new slots each containing two nodes. One of these slots would calculate the fragment job in question, while the other slot would receive another small fragment job from the global master.

B. Loss and coarse grained scaling

Concerning the parallel performance of the DEC-RI-MP2 scheme as a whole, we distinguish between local and global loss of efficiency.[65] The local loss occurs in the individual fragment calculation at the medium grained level of parallelization and was analyzed in Section IV A. Local loss is present due to nonideal load balancing, communication, and the non-parallelized steps in Algorithm 1. Global loss refers to the coarse grained parallelization level and occurs when all jobs have been distributed by the global master and some of the nodes wait for the remaining jobs to finish. It may in principle also occur if many local masters are trying to send/receive fragment job information to/from the global master at the same time, but we have not observed this to be a practical problem. In Figure 4, we investigate the coarse grained parallel performance for a calculation on the $(H_2O)_{200}$ cluster using between 400 and 1600 nodes. The calculation using 400 nodes is taken as reference and corresponds to a time-to-solution of 8 h and 14 min. The coarse grained scaling behavior for the DEC-RI-MP2 code is close to ideal in the considered range, although it slowly starts to deteriorate for 1600 nodes for reasons detailed below.

Equation (15) was used to define the fragment slot sizes for the calculations in Figure 4, and some local loss is therefore expected. Global loss is present for any DEC calculation using more than one fragment slot. The local and global losses for the calculations in Figure 4 are compared in Table VI. The global loss is increasing with the number of nodes, reflecting the decrease in relative speed up in Figure 4, where the calculation with 400 nodes was used as reference. This

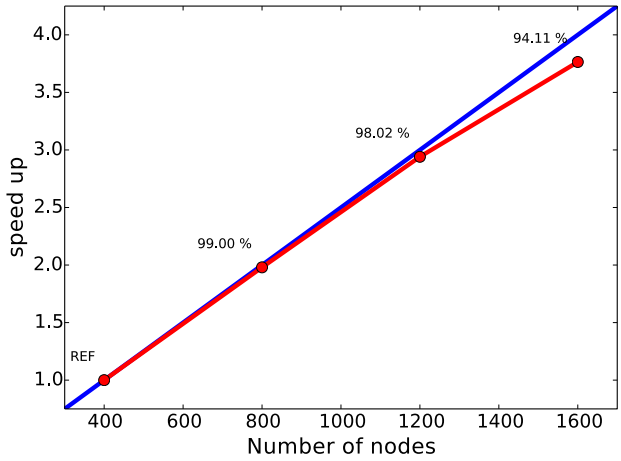

FIG. 4. Coarse grained scaling: Speed up for DEC-RI-MP2/cc-pVTZ(cc-pVTZ-RI) calculations on $(H_2O)_{200}$ cluster using FOT = 10^{-4} a.u. compared to the reference calculation using 400 nodes. The blue line represents ideal scaling and the relative times (in %) compared to ideal behavior are given for each point (relative to reference calculation).

TABLE VI. Local and global loss for DEC-RI-MP2/cc-pVTZ(cc-pVTZ-RI) calculations on $(H_2O)_{200}$ cluster using FOT = 10^{-4} a.u.

Number of nodes:	400	800	1200	1600
Local loss (%)	11.9	11.3	10.3	9.66
Global loss (%)	2.14	4.13	6.91	10.5

increase in global loss happens because, as the number of nodes is increased, more nodes are waiting for the last few fragment jobs to finish at the end of the calculation. The local loss from Table VI is roughly constant with the number of nodes (≃10%), which means that the dividing procedure relying on Eq. (15) is working properly.

The current DEC-RI-MP2 parallelization scheme can be further improved, and work is being done in that direction. At present, Eq. (15) has been defined in order to use as many nodes as possible for a given fragment calculation without having a non-beneficial scaling behaviour (see Fig. 2) and keeping the local loss as low as possible. However, if a huge number of nodes is available, it may be beneficial to use more nodes for each fragment, which would reduce the global loss at the expense of increasing the local loss. An optimal balance between local and global losses to minimize the total loss would require a more advanced criterion for determining the slot sizes than the one in Eq. (15), or, at least X in Eq. (15) should be chosen in a more sophisticated manner than simply using a fixed predefined value. Nevertheless, the current status of the code allows us to efficiently exploit large supercomputer architectures and perform large DEC-RI-MP2 calculations with a short time-to-solution and low computational loss.

V. SUMMARY AND OUTLOOK

We have presented the linear-scaling and massively parallel DEC-RI-MP2 method, which shows substantial speed up compared to the DEC-MP2 algorithm. The method can be applied to systems that are much larger than the ones that can be treated with the RI-MP2 method with small and controllable errors. The massively parallel character of the algorithm makes it particularly well suited for very large computer architectures for which several thousand nodes can be used efficiently, resulting in a very short time-to-solution. The DEC error control ensures that the standard RI-MP2 energy can be obtained to the desired precision.

The results indicate that the size of the auxiliary basis set can be further reduced, and work in this direction is ongoing. We also currently investigate the possibility of introducing a Laplace transformation of the orbital energy denominators[16,17,31,32,86] to further reduce the computational cost of the DEC-RI-MP2 method. Current and future computer architectures utilize graphical processing units (GPUs), and the DEC-RI-MP2 scheme is ideally suited to exploit such hardware as has already been done for standard RI-MP2.[15,87]

It is clear from the last two points that the DEC-RI-MP2 performance may be further improved — both concerning the algorithm itself and adaptation to modern computer hardware. It is therefore our intention to further develop DEC-RI-MP2 algorithm such that it will be able to offer a tractable alternative to density functional theory (DFT) calculations of molecular energies. Finally, we note that molecular gradient has already been implemented for the DEC-MP2 model.[51] These developments are currently being adapted to the DEC-RI-MP2 model with the goal of extending the DEC-RI-MP2 model to be able to calculate molecular properties.

ACKNOWLEDGMENTS

This research used resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract No. DE-AC05-00OR22725.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (No. FP/2007-2013)/ERC Grant Agreement No. 291371. This work was further supported by the Research Council of Norway (RCN) through CoE Grant No. 179568/V30.

We acknowledge the PRACE Research Infrastructure resource Curie at the Très Grand Centre de Calcul (TGCC) operated by CEA near Paris, France.

[1]J. L. Whitten, J. Chem. Phys. **58**, 4496 (1973).
[2]B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, J. Chem. Phys. **71**, 3396 (1979).
[3]B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, J. Chem. Phys. **71**, 4993 (1979).
[4]C. Van Alsenoy, J. Comput. Chem. **9**, 620 (1988).
[5]O. Vahtras, J. Almlöf, and M. W. Feyereisen, Chem. Phys. Lett. **213**, 514 (1993).
[6]C. Møller and M. S. Plesset, Phys. Rev. **46**, 618 (1934).
[7]M. Feyereisen, G. Fitzgerald, and A. Komornicki, Chem. Phys. Lett. **208**, 359 (1993).
[8]D. E. Bernholdt and R. J. Harrison, Chem. Phys. Lett. **250**, 477 (1996).
[9]M. Katouda and S. Nagase, Int. J. Quantum Chem. **109**, 2121 (2009).
[10]M. Katouda and T. Nakajima, J. Chem. Theory Comput. **9**, 5373 (2013).
[11]L. Maschio, D. Usvyat, F. R. Manby, S. Casassa, C. Pisani, and M. Schütz, Phys. Rev. B **76**, 075101 (2007).
[12]H.-J. Werner, F. R. Manby, and P. J. Knowles, J. Chem. Phys. **118**, 8149 (2003).
[13]F. R. Manby, J. Chem. Phys. **119**, 4607 (2003).

[14]H.-J. Werner and F. R. Manby, J. Chem. Phys. **124**, 054114 (2006).
[15]L. Vogt, R. Olivares-Amaya, S. Kermes, Y. Shao, C. Amador-Bedolla, and A. Aspuru-Guzik, J. Phys. Chem. A **112**, 2049 (2008).
[16]A. F. Izmaylov and G. E. Scuseria, Phys. Chem. Chem. Phys. **10**, 3421 (2008).
[17]T. Nakajima and K. Hirao, Chem. Phys. Lett. **427**, 225 (2006).
[18]L. Maschio, J. Chem. Theory Comput. **7**, 2818 (2011).
[19]R. A. Kendall and H. A. Früchtl, Theor. Chem. Acc. **97**, 158 (1997).
[20]D. Cremer, WIREs Comput. Mol. Sci. **1**, 509 (2011).
[21]D. E. Bernholdt and R. J. Harrison, J. Chem. Phys. **109**, 1593 (1998).
[22]F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs, Chem. Phys. Lett. **294**, 143 (1998).
[23]F. Weigend, A. Köhn, and C. Hättig, J. Chem. Phys. **116**, 3175 (2002).
[24]C. Hättig, Phys. Chem. Chem. Phys. **7**, 59 (2005).
[25]A. Hellweg and D. Rappoport, Phys. Chem. Chem. Phys. **17**, 1010 (2015).
[26]F. Aquilante, R. Lindh, and T. Bondo Pedersen, J. Chem. Phys. **127**, 114107 (2007).
[27]F. Aquilante, L. Gagliardi, T. Bondo Pedersen, and R. Lindh, J. Chem. Phys. **130**, 154107 (2009).
[28]H. Koch, A. Sánchez De Merás, and T. Bondo Pedersen, J. Chem. Phys. **118**, 9481 (2003).
[29]P. Pinski, C. Riplinger, E. F. Valeev, and F. Neese, J. Chem. Phys. **143**, 034108 (2015).
[30]Y. Kurashige, J. Yang, G. K.-L. Chan, and F. R. Manby, J. Chem. Phys. **136**, 124106 (2012).
[31]J. Almlöf, Chem. Phys. Lett. **181**, 319 (1991).
[32]M. Häser, Theor. Chim. Acta **87**, 147 (1993).
[33]B. Doser, D. S. Lambrecht, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **130**, 064107 (2009).
[34]S. A. Maurer, D. S. Lambrecht, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **138**, 014101 (2013).
[35]B. Doser, J. Zienau, L. Clin, D. S. Lambrecht, and C. Ochsenfeld, Z. Phys. Chem. **224**, 397 (2010).
[36]P. Y. Ayala and G. E. Scuseria, J. Chem. Phys. **110**, 3660 (1999).
[37]M. Kobayashi, Y. Imamura, and H. Nakai, J. Chem. Phys. **127**, 074103 (2007).
[38]Y. Mochizuki, K. Yamashita, T. Murase, T. Nakano, K. Fukuzawa, K. Takematsu, H. Watanabe, and S. Tanaka, Chem. Phys. Lett. **457**, 396 (2008).
[39]T. Ishikawa and K. Kuwata, Chem. Phys. Lett. **474**, 195 (2009).
[40]T. Ishikawa and K. Kuwata, J. Phys. Chem. Lett. **3**, 375 (2012).
[41]M. Katouda, Theor. Chem. Acc. **130**, 449 (2011).
[42]A. P. Rahalkar, M. Katouda, S. R. Gadre, and S. Nagase, J. Comput. Chem. **31**, 2405 (2010).
[43]V. Deev and M. A. Collins, J. Chem. Phys. **122**, 154102 (2005).
[44]Y. Guo, W. Li, and S. Li, J. Phys. Chem. A **118**, 8996 (2014).
[45]Y. Guo, W. Li, D. Yuan, and S. Li, Sci. China: Chem. **57**, 1393 (2014).
[46]J. Friedrich and M. Dolg, J. Chem. Theory Comput. **5**, 287 (2009).
[47]M. Ziółkowski, B. Jansík, T. Kjærgaard, and P. Jørgensen, J. Chem. Phys. **133**, 014107 (2010).
[48]K. Kristensen, M. Ziółkowski, B. Jansík, T. Kjærgaard, and P. Jørgensen, J. Chem. Theory Comput. **7**, 1677 (2011).
[49]K. Kristensen, I.-M. Høyvik, B. Jansík, P. Jørgensen, T. Kjærgaard, S. Reine, and J. Jakowski, Phys. Chem. Chem. Phys. **14**, 15706 (2012).
[50]I.-M. Høyvik, K. Kristensen, B. Jansík, and P. Jørgensen, J. Chem. Phys. **136**, 014105 (2012).
[51]K. Kristensen, P. Jørgensen, B. Jansík, T. Kjærgaard, and S. Reine, J. Chem. Phys. **137**, 114102 (2012).
[52]J. J. Eriksen, P. Baudin, P. Ettenhuber, K. Kristensen, T. Kjærgaard, and P. Jørgensen, J. Chem. Theory Comput. **11**, 2984 (2015).
[53]M. Ziółkowski, B. Jansík, P. Jørgensen, and J. Olsen, J. Chem. Phys. **131**, 124112 (2009).
[54]B. Jansík, S. Høst, K. Kristensen, and P. Jørgensen, J. Chem. Phys. **134**, 194104 (2011).
[55]I.-M. Høyvik, B. Jansík, and P. Jørgensen, J. Chem. Phys. **137**, 224114 (2012).
[56]I.-M. Høyvik, B. Jansík, and P. Jørgensen, J. Chem. Theory Comput. **8**, 3137 (2012).
[57]I.-M. Høyvik, B. Jansík, and P. Jørgensen, J. Comput. Chem. **34**, 1456 (2013).
[58]I.-M. Høyvik, K. Kristensen, T. Kjærgaard, and P. Jørgensen, Theor. Chem. Acc. **133**, 1417 (2013).
[59]S. F. Boys, Rev. Mod. Phys. **32**, 296 (1960).
[60]J. M. Foster and S. F. Boys, Rev. Mod. Phys. **32**, 300 (1960).
[61]C. Edmiston and K. Ruedenberg, Rev. Mod. Phys. **35**, 457 (1963).
[62]C. Edmiston and K. Ruedenberg, J. Chem. Phys. **43**, S97 (1965).
[63]J. Pipek and P. G. Mezey, J. Chem. Phys. **90**, 4916 (1989).
[64]C. Zhang and S. Li, J. Chem. Phys. **141**, 244106 (2014).
[65]K. Kristensen, T. Kjærgaard, I.-M. Høyvik, P. Ettenhuber, P. Jørgensen, B. Jansík, S. Reine, and J. Jakowski, Mol. Phys. **111**, 1196 (2013).
[66]T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic Structure Theory*, 1st ed. (Wiley, Chichester, England, 2000).
[67]The term "fragment canonical basis" was previously used to denote this basis.
[68]F. Weigend and M. Häser, Theor. Chem. Acc. **97**, 331 (1997).
[69]S. Jakobsen, K. Kristensen, and F. Jensen, J. Chem. Theory Comput. **9**, 3978 (2013).
[70]M. Schütz, H.-J. Werner, R. Lindh, and F. R. Manby, J. Chem. Phys. **121**, 737 (2004).
[71]D. Usvyat, L. Maschio, F. R. Manby, S. Casassa, M. Schütz, and C. Pisani, Phys. Rev. B **76**, 075102 (2007).
[72]See supplementary material at http://dx.doi.org/10.1063/1.4940732 for molecular geometries and the correlation energy of the systems presented in the result section.
[73]A. K. Wilson, D. E. Woon, K. A. Peterson, and T. H. Dunning, J. Chem. Phys. **110**, 7667 (1999).
[74]D. Feller, J. Comput. Chem. **17**, 1571 (1996).
[75]K. L. Schuchardt, B. T. Didier, T. Elsethagen, L. Sun, V. Gurumoorthi, J. Chase, J. Li, and T. L. Windus, J. Chem. Inf. Model. **47**, 1045 (2007).
[76]K. Aidas, C. Angeli, K. L. Bak, V. Bakken, R. Bast, L. Boman, O. Christiansen, R. Cimiraglia, S. Coriani, P. Dahle, E. K. Dalskov, U. Ekström, T. Enevoldsen, J. J. Eriksen, P. Ettenhuber, B. Fernández, L. Ferrighi, H. Fliegl, L. Frediani, K. Hald, A. Halkier, C. Hättig, H. Heiberg, T. Helgaker, A. C. Hennum, H. Hettema, E. Hjertenæs, S. Høst, I.-M. Høyvik, M. F. Iozzi, B. Jansík, H. J. A. Jensen, D. Jonsson, P. Jørgensen, J. Kauczor, S. Kirpekar, T. Kjærgaard, W. Klopper, S. Knecht, R. Kobayashi, H. Koch, J. Kongsted, A. Krapp, K. Kristensen, A. Ligabue, O. B. Lutnæs, J. I. Melo, K. V. Mikkelsen, R. H. Myhre, C. Neiss, C. B. Nielsen, P. Norman, J. Olsen, J. M. H. Olsen, A. Osted, M. J. Packer, F. Pawlowski, T. B. Pedersen, P. F. Provasi, S. Reine, Z. Rinkevicius, T. A. Ruden, K. Ruud, V. Rybkin, P. Salek, C. C. M. Samson, A. S. de Merás, T. Saue, S. P. A. Sauer, B. Schimmelpfennig, K. Sneskov, A. H. Steindal, K. O. Sylvester-Hvid, P. R. Taylor, A. M. Teale, E. I. Tellgren, D. P. Tew, A. J. Thorvaldsen, L. Thøgersen, O. Vahtras, M. A. Watson, D. J. D. Wilson, M. Ziółkowski, and H. Ågren, WIREs Comput. Mol. Sci. **4**, 269 (2013).
[77]LSDalton, a linear-scaling molecular electronic structure program, Release Dalton2016, 2015, http://daltonprogram.org.
[78]D. Kats, T. Korona, and M. Schütz, J. Chem. Phys. **125**, 104106 (2006).
[79]E. Goll, T. Leininger, F. R. Manby, A. Mitrushchenkov, H.-J. Werner, and H. Stoll, Phys. Chem. Chem. Phys. **10**, 3353 (2008).
[80]M. Kállay, J. Chem. Phys. **141**, 244113 (2014).
[81]D. S. Hollman, H. F. Schaefer, and E. F. Valeev, J. Chem. Phys. **142**, 154106 (2015).
[82]The calculations on Eos were performed using 20 XC30 compute-nodes, each with two sockets of 8 physical cores (Intel® Xeon® Processor E5-2670–2.60 GHz). Eos uses Cray's Aries interconnect, and each node has 64 GB of memory available.
[83]The calculations on Grendel cluster at the center for scientific computing aarhus (CCSCA) were performed using 21 SUN X2200 compute-nodes, each with 2 quadcore 2.3 GHz AMD Opteron CPU and 16GB of memory. The nodes are interconnected using a Gigabit Ethernet (GigE) network.
[84]The calculations on Titan were performed using up to 1600 compute nodes of the Titan supercomputing system at ORNL. Each node consists of a 16-core 2.2 GHz AMD Opteron CPU and 32 GB of memory. Two nodes share a Gemini high-speed interconnect router.
[85]C. Hättig, A. Hellweg, and A. Köhn, Phys. Chem. Chem. Phys. **8**, 1159 (2006).
[86]D. Kats, D. Usvyat, and M. Schütz, Phys. Chem. Chem. Phys. **10**, 3430 (2008).
[87]S. A. Maurer, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **141**, 051106 (2014).