

Dissertation zur Erlangung des Doktorgrades
der Fakultät für Chemie und Pharmazie
der Ludwig-Maximilians-Universität München

Development of Low-Scaling
Electronic Structure Methods
Using Rank Factorizations
and an Attenuated Coulomb Metric

Arne Lünser

aus

Hamburg

2017

Erklärung

Diese Dissertation wurde im Sinne von § 7 der Promotionsordnung vom 28. November 2011 von Herrn Prof. Dr. Christian Ochsenfeld betreut.

Eidesstattliche Versicherung

Diese Dissertation wurde eigenständig und ohne unerlaubte Hilfe erarbeitet.

München, 8. Juni 2017

Arne Lünser

Dissertation eingereicht am	2. Mai 2017
1. Gutachter	Prof. Dr. Christian Ochsenfeld
2. Gutachterin	Prof. Dr. Regina de Vivie-Riedle
Mündliche Prüfung am	1. Juni 2017

Mein Dank gilt zuerst Herrn Prof. Dr. Christian Ochsenfeld für die Möglichkeit, diese Dissertation unter seiner Anleitung anzufertigen; für seine Geduld sowie seine Großzügigkeit mit seiner eigenen Zeit.

Bei Frau Prof. Dr. Regina de Vivie-Riedle bedanke ich mich herzlich für die Anfertigung des Zweitgutachtens.

Dem gesamten Arbeitskreis Ochsenfeld danke ich für die unkomplizierte Kameradschaft und den unverzichtbaren wissenschaftlichen Austausch. Insbesondere gilt mein Dank Dr. Asbjörn M. Burow, Dr. Jörg Kußmann, Henry F. Schurkus und Matthias Beuerle.

Bei meiner Familie und meinen Freunden – insbesondere Matthias Trunk und Tim Gatzemeier – bedanke ich mich für die moralische Unterstützung ohne die diese Arbeit schlichtweg nicht existierte.

Abstract

Novel low-scaling techniques for molecular electronic structure and property calculations are introduced. Through the use of rank-revealing matrix factorizations, overheads compared to canonical molecular orbital-based formulations are virtually eliminated. Asymptotic computational complexity is linear or sub-linear (depending on the property) through the use of sparsity-preserving transformations throughout. For electron correlation energy calculations within the random phase approximation, these techniques are combined with an attenuated Coulomb metric in the resolution-of-the-identity to improve the accuracy over existing low-scaling methods and to reduce the scaling compared to existing canonical methods. For the resolution-of-the-identity itself, a novel method for the compression of auxiliary bases is introduced, powered by removal of the particle-hole-interaction nullspace through projection. Furthermore, efficient schemes for the calculation of molecular response properties at the Hartree–Fock and density functional theory levels are introduced: For the linear scaling calculation of vibrational frequencies, the exact cancellation of different long-range operator derivatives is employed in combination with Laplace-transformed and Cholesky-decomposed coupled-perturbed self-consistent field theories. Using related techniques, calculations of indirect nuclear spin-spin coupling constants with asymptotically constant time complexity are demonstrated and used to explore the dependence of spin-spin couplings in a peptide on the size of a surrounding solvent environment.

Contents

List of Publications	11
1 Introduction	13
2 Theory	15
2.1 Self-Consistent Field Methods	15
2.1.1 Hartree–Fock Theory	15
2.1.2 Kohn–Sham Theory	16
2.1.3 Roothaan–Hall Equations	17
2.1.4 Density Matrix Formulations	17
2.2 Molecular Properties	18
2.2.1 Hellmann–Feynman Theorem	18
2.2.2 Derivatives of the SCF Energy	19
2.2.3 Coupled-Perturbed SCF	20
2.2.4 Laplace-Transformed Density Matrix-Based CPSCF	21
2.3 Resolution-of-the-Identity	21
2.4 Correlation Energies from the Random Phase Approximation	23
2.4.1 Adiabatic Connection Correlation Energy	23
2.4.2 Frequency Integration of the Polarization Propagator	25
2.4.3 Dyson Equation and the Random Phase Approximation	26
3 Additional Results	27
3.1 The Pivoted Cholesky Factorization for DL-CPSCF	27
3.2 Rank Reduction Through Cholesky Factorization	28
3.3 Useable Matrix Sparsity	29
3.4 Sparsity of the Virtual Density Matrix	30
3.5 Sparsity of Cholesky-Decomposed Density Matrices	31
3.6 Local Perturbations	32
3.6.1 Nuclear Displacement	32
3.6.2 Magnetic Perturbations	33
3.7 Local Metrics in the Resolution-of-the-Identity	33
4 Summary	35
Bibliography	37

Publications	43
Article I	43
Article II	55
Article III	79
Article IV	91
Academic Teachers	105

List of Publications

Results of this dissertation have been published in several peer-reviewed articles. The following lists those articles and the author's contributions to each of them. Articles **I-IV** are considered the main part of this work, and are collected here in their entirety starting from page 43.

- I** "Vanishing-Overhead Linear-Scaling Random Phase Approximation by Cholesky Decomposition and an Attenuated Coulomb-Metric",
A. Luenser, H. F. Schurkus, and C. Ochsenfeld,
J. Chem. Theor. Comput. **13**, 1647 (2017).
Contributions by A. Luenser: *Conjoint development of the theory with H. F. Schurkus. All Calculations. Most of the implementation and writing.*
- II** "Almost error-free resolution-of-the-identity correlation methods by null space removal of the particle-hole interactions",
H. F. Schurkus, A. Luenser, C. Ochsenfeld,
J. Chem. Phys. **146**, 211106 (2017).
Contributions by A. Luenser: *Conjoint development of the theory with H. F. Schurkus. Major parts of the implementation. Parts of the calculations and writing.*
- III** "Computation of indirect nuclear spin–spin couplings with reduced complexity in pure and hybrid density functional approximations",
A. Luenser, J. Kussmann, and C. Ochsenfeld,
J. Chem. Phys. **145**, 124103 (2016).
Contributions by A. Luenser: *Parts of the theory. All calculations. Most of the writing and implementation.*
- IV** "A reduced-scaling density matrix-based method for the computation of the vibrational Hessian matrix at the self-consistent field level",
J. Kussmann, A. Luenser, M. Beer, and C. Ochsenfeld,
J. Chem. Phys. **142**, 094101 (2015).
Contributions by A. Luenser: *Parts of the implementation and derivation. All calculations. Most of the writing.*

- V** "Nuclear Magnetic Shieldings of Stacked Aromatic and Antiaromatic Molecules",
D. Sundholm, M. Rauhalhti, N. Özcan, R. Mera-Adasme, J. Kussmann, A. Luenser,
and C. Ochsenfeld,
J. Chem. Theory Comput. **13**, 1952 (2017).
Contributions by A. Luenser: *Parts of the computational results and writing.*
- VI** "Advances in molecular quantum chemistry contained in the Q-Chem 4 program
package",
Y. Shao *et al.*,
Mol. Phys. **113**, 184 (2015).
Contributions by A. Luenser: *Parts of the implementation and writing.*

1 Introduction

An ongoing quest in quantum chemistry is the calculation of the correlation energy. For small molecules, coupled-cluster theories truncated to double^[1] or (perturbative) triple^[2] excitations reliably provide excellent energetics and properties.^[3–6] However, this only holds when the Hartree–Fock reference wave function is a good approximation to the exact wave function. If the electronic ground state cannot be qualitatively captured in a single Slater determinant, truncated coupled-cluster theories fail. The same is true for Møller–Plesset^[7] many-body perturbation theory.

For such electronic situations, the standard approach is to resort to multi-configuration theories, in which the ground state wave function ansatz is a superposition of multiple Slater determinants.^[8] Multi-configurational self-consistent field theory is well established,^[9] as is the application of perturbation theory.^[10,11] In case of coupled-cluster theories, this generalization has not yet been fully realized, but several competing approaches have been put forward.^[6,12] All multi-configurational methods share some common limitations, though.

First, the computational demand increases extremely steeply with the number of configurations. The number of electrons which can be correlated by conventional multi-configurational approaches is of the order of ten to twenty, beyond which the necessary computational resources rise to intractable levels. This limits applications to molecular systems of small to moderate size.

The second limitation lies in the selection of which orbitals to include in the multi-configurational treatment. For small systems, chemical insight allows for an educated guess and often works well. This requires intricate knowledge of both the theoretical treatment and the molecular system under study, however.

For medium- to large-sized molecular systems, Kohn–Sham density functional theory (KS-DFT)^[13] is the *de facto* standard, because it yields useful accuracy at affordable computational cost. In fact, KS-DFT is often the *only* available choice, squarely due to the steep cost of other methods. The limitations of KS-DFT are relatively well understood.^[14,15] First, non-local correlation effects, such as dispersion forces, cannot be captured by local or semi-local density functionals. Approaches to remedy this include empirical^[16] or non-empirical^[17–20] corrections and non-local potentials.^[21–25] Second, most approximate exchange–correlation functionals suffer from self-interaction errors.^[26] Finally, static correlation is a problem in Kohn–Sham theories as much as it is in wave function theories. Several approaches have been proposed, including but not limited to, refs. 27–31.

Among the many approaches which have been put forward to overcome these limitations of KS-DFT, the random phase approximation (RPA) holds great promise. As a

correlation method, the RPA can be rigorously derived from the adiabatic connection (cf. sec. 2.4). It contains only well-defined approximations and no empirical parameters, is size-consistent^[32] and contains an *ab initio* description of dispersion. Additionally, it is non-perturbative, and convergent even for small-gap systems and metals. These desirable features have led to increased interest in the RPA over the last 15 years, even though the method itself can be traced to 1951.^[33]

In a Gaussian basis set, calculation of the RPA correlation energy scales as $O(N^6)$ with molecular size N if computed through diagonalization, or $O(N^5)$ when using iterative matrix sign-function methods.^[34] A breakthrough by Eshuis *et al.*^[35] reduced this to $O(N^4 \log N)$ using resolution-of-the-identity methods (cf. sec. 2.3) and a numerical integration. Since then, multiple authors have reported even lower-scaling algorithms.^[36–39]

This work includes article **I** (page 43), which takes the approach by Schurkus and Ochsenfeld^[38] and improves both accuracy and performance dramatically. Accuracy is enhanced by introducing a Coulomb metric attenuated by a complementary error function in the resolution-of-the-identity (cf. sec. 3.7), and performance is improved by pivoted Cholesky decomposition of density-like matrices (cf. sec. 3.2).

Article **II** (page 55) describes a novel method for the generation and compression of auxiliary bases in resolution-of-the-identity calculations, specifically correlation energy calculations which require the use of particle-hole pair interactions, such as second-order Møller–Plesset perturbation theory and the random phase approximation.

Aside from the molecular correlation energy, calculation of molecular properties (cf. sec. 2.2) is another important field of quantum chemistry. Articles **III** (page 79) and **IV** (page 91) are concerned with the calculation of second-order molecular properties at the self-consistent field level of theory. Because second- and higher-order must normally be calculated as derivatives of the electronic energy, computational cost is prohibitive in many cases. Explicitly using the locality of the perturbing operators (cf. sec. 3.6) to obtain sparse matrix representations has enabled the calculations of harmonic vibrational frequencies with $O(N)$ scaling (page 91), and the calculation of indirect nuclear spin–spin coupling constants with $O(1)$ scaling (page 79), both at the (hybrid-) density functional level of theory.

This work is a cumulative dissertation. As such, articles **I** through **IV** are considered the main part of this work and are collected in their entirety starting from page 43. Chapter 2 contains a succinct overview of the required theoretical background. Chapter 3 comprises additional results obtained during the work on this dissertation which have not been included in articles **I** through **IV**. Finally, chapter 4 summarizes the results.

2 Theory

2.1 Self-Consistent Field Methods

Hartree–Fock^[40–42] and Kohn–Sham^[13] theories are two of the most commonly employed methods for electronic structure calculations. Both provide approximate solutions to the Schrödinger equations by transforming the many-electron problem to a set of effective one-electron problems, where the electron–electron interaction is folded into an effective one-electron potential. This potential is itself a dependent on the one-electron wave functions. The solutions and potential are iteratively refined until they are self-consistent, i. e., the one-electron wave function solution set gives rise to the same potential (“field”) as was used to obtain the wave functions. Hence the term self-consistent field (SCF) theory.

I will review Hartree–Fock and Kohn–Sham theory only very briefly in the following, noting that countless textbooks are available on the topic (for example refs. 8,43–45). The aim here is to provide some context to the equations in the following sections and establish a consistent notation, without reiterating too much textbook knowledge.

2.1.1 Hartree–Fock Theory

Hartree–Fock theory starts with the exact, non-relativistic molecular Hamiltonian \hat{H}

$$\hat{H} = - \sum_A \frac{1}{2m_A} \nabla_A^2 - \sum_i \frac{1}{2} \nabla_i^2 + \sum_{B>A} \frac{Z_A Z_B}{r_{ab}} - \sum_{iA} \frac{Z_A}{r_{iA}} + \sum_{j>i} \frac{1}{r_{ij}}, \quad (2.1.1)$$

(in atomic units), where A, B index nuclei with mass m_A , charges Z_A, Z_B , and i, j index electrons. $r_{xy} = |\mathbf{r}_y - \mathbf{r}_x|$ denotes the distance between particles x, y . The Born–Oppenheimer approximation is assumed, and the ansatz for the molecular wave function is a single Slater determinant $|\Psi_{\text{SD}}\rangle$ built from an orthonormal set of molecular orbitals $\{|\varphi\rangle\}$ (MOs, one-particle wave functions). The Hartree–Fock energy expression is the expectation value of the exact Hamiltonian with this Slater determinant:

$$E_{\text{HF}} = \frac{\langle \Psi_{\text{SD}} | \hat{H} | \Psi_{\text{SD}} \rangle}{\langle \Psi_{\text{SD}} | \Psi_{\text{SD}} \rangle}. \quad (2.1.2)$$

This expression is variationally minimized by forming the functional derivative with respect to the molecular orbitals $|\varphi_i\rangle$ and finding the roots of the resulting expression, which yields the Hartree–Fock equations (canonical form)^[43]

$$\hat{F} |\varphi_i\rangle = \varepsilon_i |\varphi_i\rangle, \quad (2.1.3)$$

with

$$\hat{F} = \hat{h} + \hat{v}_{\text{HF}}, \quad (2.1.4)$$

$$\hat{h} = -\frac{1}{2}\nabla^2 + v_{\text{ext}}(\mathbf{r}) \quad (2.1.5)$$

$$\hat{v}_{\text{HF}} = \hat{J} - \hat{K}, \quad (2.1.6)$$

$$\hat{J} = \sum_j \int d\mathbf{r}' \varphi_j^*(\mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|} \varphi_j(\mathbf{r}') = \int d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}, \quad (2.1.7)$$

$$\hat{K} = \sum_j \int d\mathbf{r}' \varphi_j^*(\mathbf{r}') \frac{\mathcal{P}}{|\mathbf{r} - \mathbf{r}'|} \varphi_j(\mathbf{r}'). \quad (2.1.8)$$

The external potential v_{ext} includes the field generated by the nuclei, and possibly other external fields. The non-local Fock potential \hat{v}_{HF} comprises the classical electron-electron interaction in the Coulomb operator \hat{J} , and the non-classical exchange interaction in the exchange operator \hat{K} . The permutation operator \mathcal{P} in Eq. 2.1.8 exchanges the electronic coordinates \mathbf{r} and \mathbf{r}' .

Equation 2.1.2 leads to the definition of the correlation energy E_{C} , which is just the difference between the exact energy and the Hartree–Fock energy^[43]

$$E_{\text{C}} = E_{\text{exact}} - E_{\text{HF}}. \quad (2.1.9)$$

2.1.2 Kohn–Sham Theory

Kohn–Sham theory builds on the Hohenberg–Kohn theorems^[46] to construct an approximate one-electron potential analogous to the Fock potential \hat{v}_{HF} (Eq. 2.1.6), called the Kohn–Sham potential v_{KS} . Unlike the Fock potential, the Kohn–Sham potential is a functional of the electron density ρ rather than the molecular orbitals $|\varphi_i\rangle$. Analogously to the Hartree–Fock equations (2.1.3), the Kohn–Sham equations read^[44,45]

$$\hat{F} |\varphi_i\rangle = \varepsilon_i |\varphi_i\rangle, \quad (2.1.10)$$

$$\hat{F} = -\frac{1}{2}\nabla^2 + v_{\text{ext}}(\mathbf{r}) + v_{\text{KS}}(\mathbf{r}), \quad (2.1.11)$$

$$v_{\text{KS}}(\mathbf{r}) = \int d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + v_{\text{XC}}(\mathbf{r}), \quad (2.1.12)$$

$$v_{\text{XC}}(\mathbf{r}) = \frac{\delta E_{\text{XC}}[\rho]}{\delta \rho(\mathbf{r})}. \quad (2.1.13)$$

Compared to the Hartree–Fock equations, the exchange operator is replaced by the exchange–correlation (XC) potential $v_{\text{XC}}(\mathbf{r})$, which is the functional derivative of the exchange–correlation energy functional $E_{\text{XC}}[\rho]$ with respect to the density $\rho(\mathbf{r})$. While the exact form of $E_{\text{XC}}[\rho]$ is unknown for the general case, various constraints and physical insights have been used to obtain a plethora of approximations (collectively known as density functional approximations).^[45]

As the name suggests, exchange-correlation functionals attempt to fold exchange and correlation effects into the effective potential. This is in contrast to Hartree–Fock theory, which only incorporates exchange effects and—by definition—no electron correlation.

2.1.3 Roothaan–Hall Equations

The Hartree–Fock and Kohn–Sham equations are most often (but not exclusively) solved by expanding the molecular orbitals $\{|\varphi_i\rangle\}$ in some (generally non-orthogonal) one-electron basis set $\{|\mu\rangle\}$

$$|\varphi_i\rangle = \sum_{\mu} |\mu\rangle C_{\mu i}, \quad (2.1.14)$$

which when inserted into Eqs. 2.1.3 or 2.1.10 yields the Roothaan–Hall^[47,48] equations

$$\sum_{\nu} F_{\mu\nu} C_{\nu i} = \sum_{\nu} S_{\mu\nu} C_{\nu i} \varepsilon_i \quad (2.1.15)$$

$$F_{\mu\nu} = \langle \mu | \hat{F} | \nu \rangle, \quad (2.1.16)$$

$$S_{\mu\nu} = \langle \mu | \nu \rangle, \quad (2.1.17)$$

which is simply the matrix representation of the Hartree–Fock or Kohn–Sham equations in the basis $\{|\mu\rangle, |\nu\rangle, \dots\}$. Many different types of basis sets are in use, including but not limited to plane waves and augmented plane waves, wavelets, atom-centered Gaussian or Slater functions modeled after atomic orbitals (AOs). The results presented in this work, particularly in chapter 3 and articles **I–IV**, used Gaussian functions exclusively. However, many of the principles are directly applicable to any other type of local basis set. In the following, $\{i, j, \dots\}$ will be used for occupied MOs, $\{a, b, \dots\}$ will be virtual MOs, and $\{p, q, \dots\}$ will be general (occupied or virtual) MOs.

2.1.4 Density Matrix Formulations

It is very often convenient to work with reduced one-particle density matrices instead of molecular orbitals. For AO-based linear-scaling SCF theories, this is in fact required, because the canonical molecular orbitals offer no locality. Section 3.3 discusses this aspect in more detail.

In SCF-type calculations, the density operator is

$$\hat{\rho} = \sum_p |\varphi_p\rangle n_p \langle \varphi_p|, \quad (2.1.18)$$

where the occupation numbers n_p are zero (virtual orbitals) or unity (occupied orbitals) for pure states, such as all states considered in this work. The atomic orbital representation of the density operator in SCF theories reduces to

$$P_{\mu\nu} = \sum_{i \in \text{occ}} C_{\mu i} C_{i\nu}^{\dagger}, \quad (2.1.19)$$

which is the dyadic product of all occupied MO coefficient vectors $\underline{\mathbf{C}}$

$$\mathbf{P} = \underline{\mathbf{C}}\underline{\mathbf{C}}^\dagger. \quad (2.1.20)$$

For pure states, the density operator is a projection operator. Later, it will be useful to define a “virtual” density, which is the complement of the density (operator), defined by

$$\mathbf{Q} = \overline{\mathbf{C}}\overline{\mathbf{C}}^\dagger = \mathbf{S}^{-1} - \mathbf{P}, \quad (2.1.21)$$

where $\overline{\mathbf{C}}$ are the virtual MO coefficient vectors.

Using density matrices in the atomic orbital basis, the energy expression in Hartree–Fock and Kohn–Sham theories become^[49]

$$E_{\text{SCF}} = \sum_{\mu\nu} P_{\mu\nu} h_{\nu\mu} + \frac{1}{2} P_{\mu\nu} G_{\nu\mu}(\mathbf{P}) + V_{\text{NN}} \quad (2.1.22)$$

$$= \text{Tr} \left(\mathbf{P}\mathbf{h} + \frac{1}{2} \mathbf{P}\mathbf{G}(\mathbf{P}) \right) + V_{\text{NN}}, \quad (2.1.23)$$

where V_{NN} is the nuclear repulsion energy, and $\mathbf{G}(\mathbf{P})$ is the matrix representation of the Fock (Eq. 2.1.6) or Kohn–Sham (Eq. 2.1.12) potential, respectively. This notation shows the explicit dependence of the effective potential \mathbf{G} on the density matrix \mathbf{P} .

2.2 Molecular Properties

Apart from energies, prediction of molecular properties is of high relevance. Predictions can aid in interpretation of experimental data, substantiate or challenge hypotheses.

For virtually all molecular electronic structure methods, gradients of the energy with respect to nuclear coordinates are among the most important properties. They allow optimization of molecular geometries toward points of zero gradients, i. e. stationary points on the potential hyper surface. Other properties accessible include harmonic and higher-order vibrational frequencies including infrared and Raman intensities, nuclear magnetic shieldings tensors, indirect nuclear spin-spin coupling tensors, static and dynamic polarizabilities and hyper-polarizabilities, g -tensors, hyperfine coupling constants.^[50]

Some molecular properties can be calculated as simple expectation values of the ground-state wave function or (spin) densities, for example dipole moments, or hyperfine coupling constants. Higher-order properties, such as the gradient (first-order) or vibrational frequencies (second-order) are more involved.

2.2.1 Hellmann–Feynman Theorem

The Hellmann–Feynman theorem^[51,52] states that the first derivative of an expectation value requires only evaluation of an expectation value of the unperturbed wave function

with the operator derivative, but no wave function derivatives:

$$E = \langle \Psi | \hat{H} | \Psi \rangle \quad (2.2.1)$$

$$\frac{dE}{d\lambda} = \left\langle \frac{d}{d\lambda} \Psi | \hat{H} | \Psi \right\rangle + \langle \Psi | \frac{d}{d\lambda} \hat{H} | \Psi \rangle + \langle \Psi | \hat{H} | \frac{d}{d\lambda} \Psi \rangle \quad (2.2.2)$$

$$= E \left\langle \frac{d}{d\lambda} \Psi | \Psi \right\rangle + E \langle \Psi | \frac{d}{d\lambda} \Psi \rangle + \langle \Psi | \frac{d}{d\lambda} \hat{H} | \Psi \rangle \quad (2.2.3)$$

$$= E \frac{d}{d\lambda} \langle \Psi | \Psi \rangle + \langle \Psi | \frac{d}{d\lambda} \hat{H} | \Psi \rangle \quad (2.2.4)$$

$$= \langle \Psi | \frac{d}{d\lambda} \hat{H} | \Psi \rangle. \quad (2.2.5)$$

While the theorem holds only for complete bases and is therefore typically not applicable in quantum chemistry, it can still provide valuable insight. In the context of linear-scaling computation of vibrational frequencies as second-order energy derivatives with respect to nuclear coordinates, the decay behavior of the interaction can be inferred to be more rapid than cursory examination suggests. Details of this will be discussed in sec. 3.6. In most practical calculations, energy derivatives are not evaluated through the Hellmann–Feynman theorem, but by explicitly computing the derivatives of the energy expression of the method in question.

2.2.2 Derivatives of the SCF Energy

The first derivative of the Hartree–Fock or Kohn–Sham energy^[53–57] in a density matrix-based form (Eq. 2.1.23) with respect to a general perturbation λ (we use the shorthand $\frac{d}{d\lambda} \mathbf{Z} = \mathbf{Z}^\lambda$) yields

$$\frac{dE_{\text{SCF}}}{d\lambda} = \frac{d}{d\lambda} \text{Tr} \left(\mathbf{P} \mathbf{h} + \frac{1}{2} \mathbf{P} \mathbf{G}(\mathbf{P}) \right) + \frac{d}{d\lambda} V_{\text{NN}} \quad (2.2.6)$$

$$= \text{Tr} \left(\mathbf{P} \mathbf{h}^\lambda + \frac{1}{2} \mathbf{P} \mathbf{G}^\lambda(\mathbf{P}) + \mathbf{P}^\lambda (\mathbf{h} + \mathbf{G}(\mathbf{P})) \right) + V_{\text{NN}}^\lambda \quad (2.2.7)$$

$$= \text{Tr} \left(\mathbf{P} \mathbf{h}^\lambda + \frac{1}{2} \mathbf{P} \mathbf{G}^\lambda(\mathbf{P}) + \mathbf{P}^\lambda \mathbf{F} \right) + V_{\text{NN}}^\lambda \quad (2.2.8)$$

$$= \text{Tr} \left(\mathbf{P} \mathbf{h}^\lambda + \frac{1}{2} \mathbf{P} \mathbf{G}^\lambda(\mathbf{P}) - \mathbf{W} \mathbf{S}^\lambda \right) + V_{\text{NN}}^\lambda, \quad (2.2.9)$$

with $\mathbf{W} = \mathbf{P} \mathbf{F} \mathbf{P}$. The identity $\text{Tr}(\mathbf{P}^\lambda \mathbf{F}) = -\text{Tr}(\mathbf{W} \mathbf{S}^\lambda)$ can be understood from inspection of the subspace projections of \mathbf{P}^λ and \mathbf{F} :^[58]

$$\mathbf{P} \mathbf{S} \mathbf{P}^\lambda \mathbf{S} \mathbf{P} = \mathbf{P}_{\text{oo}}^\lambda = -\mathbf{P} \mathbf{S}^\lambda \mathbf{P}, \quad (2.2.10)$$

$$\mathbf{P} \mathbf{S} \mathbf{P}^\lambda (\mathbf{1} - \mathbf{S} \mathbf{P}) = \mathbf{P}_{\text{ov}}^\lambda, \quad (2.2.11)$$

$$(\mathbf{1} - \mathbf{P} \mathbf{S}) \mathbf{P}^\lambda \mathbf{P} = \mathbf{P}_{\text{vo}}^\lambda, \quad (2.2.12)$$

$$(\mathbf{1} - \mathbf{P} \mathbf{S}) \mathbf{P}^\lambda (\mathbf{1} - \mathbf{S} \mathbf{P}) = \mathbf{P}_{\text{vv}}^\lambda = \mathbf{0}. \quad (2.2.13)$$

Brillouin's theorem requires that the occupied-virtual and virtual-occupied subspace projections of \mathbf{F} vanish. Hence, the only subspace shared between \mathbf{F} and \mathbf{P}^λ is the occupied-occupied space. The first energy derivative may therefore be evaluated without having to resort to any explicit density or wave function derivatives. This is in accordance with the Hellmann–Feynman theorem (Eq. 2.2.5) and Wigner's $(2n + 1)$ rule.

For second derivatives, however, a density response is required. The general expression is^[53–57]

$$\frac{d^2 E_{\text{SCF}}}{d\xi d\lambda} = \frac{d}{d\xi} \text{Tr} \left(\mathbf{P} \mathbf{h}^\lambda + \frac{1}{2} \mathbf{P} \mathbf{G}^\lambda(\mathbf{P}) - \mathbf{W} \mathbf{S}^\lambda \right) + \frac{d}{d\xi} V_{\text{NN}}^\lambda \quad (2.2.14)$$

$$= \text{Tr} \left(\mathbf{P} \mathbf{h}^{\lambda\xi} + \frac{1}{2} \mathbf{P} \mathbf{G}^{\lambda\xi}(\mathbf{P}) - \mathbf{W} \mathbf{S}^{\lambda\xi} \right) \quad (2.2.15)$$

$$+ \text{Tr} \left(\mathbf{P}^\xi \mathbf{h}^\lambda + \mathbf{P}^\xi \mathbf{G}^\lambda(\mathbf{P}) - \mathbf{W}^\xi \mathbf{S}^\lambda \right) + V_{\text{NN}}^{\lambda\xi}. \quad (2.2.16)$$

Evidently, the density response (or perturbed density) \mathbf{P}^ξ is required for second derivatives, and may be obtained through solution of the coupled-perturbed SCF (CPSCF) equations, discussed in sec. 2.2.3.

2.2.3 Coupled-Perturbed SCF

Differentiation of the Roothaan–Hall equations (2.1.15) with respect to perturbation ξ yields

$$\sum_\nu F_{\mu\nu}^\xi C_{\nu i} + F_{\mu\nu} C_{\nu i}^\xi = \sum_\nu S_{\mu\nu}^\xi C_{\nu i} \varepsilon_i + S_{\mu\nu} C_{\nu i}^\xi \varepsilon_i + S_{\mu\nu} C_{\nu i} \varepsilon_i^\xi. \quad (2.2.17)$$

Inserting $C_{\nu i}^\xi = \sum_a C_{\nu a} U_{ai}^\xi$ and multiplying by $\sum_\mu C_{a\mu}^\dagger$ from the left gives

$$\sum_{a\mu\nu} C_{a\mu}^\dagger F_{\mu\nu}^\xi C_{\nu i} + C_{a\mu}^\dagger F_{\mu\nu} C_{\nu a} U_{ai}^\xi = \sum_{a\mu\nu} C_{a\mu}^\dagger S_{\mu\nu}^\xi C_{\nu i} \varepsilon_i + C_{a\mu}^\dagger S_{\mu\nu} C_{\nu a} U_{ai}^\xi \varepsilon_i, \quad (2.2.18)$$

or equivalently

$$F_{ai}^\xi + \varepsilon_a U_{ai}^\xi = S_{ai}^\xi \varepsilon_i + U_{ai}^\xi \varepsilon_i. \quad (2.2.19)$$

After rearranging terms, we obtain a compact form of the coupled-perturbed SCF (CPSCF) equations^[55,59–61]

$$U_{ai}^\xi = \frac{S_{ai}^\xi \varepsilon_i - F_{ai}^\xi}{\varepsilon_a - \varepsilon_i}, \quad (2.2.20)$$

$$P_{\mu\nu}^\xi = \sum_{ai} C_{\mu a} U_{ai}^\xi C_{i\nu}^\dagger + \sum_{ia} C_{\mu i} U_{ia}^{\xi\dagger} C_{a\nu}^\dagger - \sum_{\lambda\sigma} P_{\mu\lambda} S_{\lambda\sigma}^\xi P_{\sigma\nu}. \quad (2.2.21)$$

Because $F_{\mu\nu}^\xi$ depends on the perturbed density matrix \mathbf{P}^ξ , the CPSCF equations must be solved iteratively.

2.2.4 Laplace-Transformed Density Matrix-Based CPSCF

A transformation of Eq. 2.2.20 to the AO basis has been introduced by Beer and Ochsenfeld^[62] by means of the well-established Laplace transformation^[63]

$$\frac{1}{\varepsilon_a - \varepsilon_i} = \int_0^\infty dt e^{-(\varepsilon_a - \varepsilon_i)t} \quad (2.2.22)$$

$$\simeq \sum_{\alpha=1}^{\tau} w_\alpha e^{-(\varepsilon_a - \varepsilon_i)t_\alpha}, \quad (2.2.23)$$

which is used in Møller–Plesset perturbation theory,^[64–66] CPSCF,^[62,67–69] and in a different form in the random phase approximation.^[38] In the context of CPSCF, Eq. 2.2.20 is rewritten as

$$U_{ai}^\xi \simeq \sum_{\alpha=1}^{\tau} w_\alpha e^{-\varepsilon_a t_\alpha} \left(S_{ai}^\xi \varepsilon_i - F_{ai}^\xi \right) e^{\varepsilon_i t_\alpha} \quad (2.2.24)$$

$$= \sum_{\alpha=1}^{\tau} w_\alpha e^{-(\varepsilon_a - \varepsilon_F)t_\alpha} \left(S_{ai}^\xi \varepsilon_i - F_{ai}^\xi \right) e^{(\varepsilon_i - \varepsilon_F)t_\alpha}, \quad (2.2.25)$$

where ε_F is arbitrary. Ayala and Scuseria^[66] pointed out that setting ε_F to a value between $\varepsilon_{\text{HOMO}}$ and $\varepsilon_{\text{LUMO}}$ (the highest occupied and lowest unoccupied orbital energies, respectively) means the exponential factor is always smaller than unity. We will make use of this in sec. 3.2.

The density matrix-based Laplace-transformed CPSCF equations introduced by Beer and Ochsenfeld^[62] read

$$\mathbf{P}_{\text{vo}}^\xi = \sum_{\alpha=1}^{\tau} \mathbf{Q}^{(\alpha)} \left(-\mathbf{h}^\xi - \mathbf{G} \left[\mathbf{P}^\xi \right] \right) \mathbf{P}^{(\alpha)}, \quad (2.2.26)$$

$$P_{\mu\nu}^{(\alpha)} = \sqrt{w_\alpha} \sum_i C_{\mu i} \exp(t_\alpha \varepsilon_i) C_{i\nu}^\dagger, \quad (2.2.27)$$

$$Q_{\mu\nu}^{(\alpha)} = \sqrt{w_\alpha} \sum_a C_{\mu a} \exp(-t_\alpha \varepsilon_a) C_{a\nu}^\dagger. \quad (2.2.28)$$

We refer to $\mathbf{P}^{(\alpha)}$ and $\mathbf{Q}^{(\alpha)}$ loosely as “pseudo-density matrices”. Like the regular one-particle density matrices \mathbf{P} and $\mathbf{Q} = \mathbf{S}^{-1} - \mathbf{P}$, they become sparsely populated for molecular systems of sufficient size and non-zero band gap.

2.3 Resolution-of-the-Identity

In the context of quantum chemistry, the resolution-of-the-identity (RI)^[70–77] is used to approximate four-center two-electron repulsion integrals (ERIs) over atomic basis functions of the form

$$(\mu\nu|\lambda\sigma) = \iint d\mathbf{r} d\mathbf{r}' \mu(\mathbf{r}) \nu(\mathbf{r}) \frac{1}{|\mathbf{r} - \mathbf{r}'|} \lambda(\mathbf{r}') \sigma(\mathbf{r}') \quad (2.3.1)$$

as

$$(\mu\nu|\lambda\sigma) \approx \sum_{PQ} (\mu\nu|P) (P|Q)^{-1} (Q|\lambda\sigma), \quad (2.3.2)$$

$$(\mu\nu|P) = \iint d\mathbf{r} d\mathbf{r}' \mu(\mathbf{r}) \nu(\mathbf{r}) \frac{1}{|\mathbf{r} - \mathbf{r}'|} P(\mathbf{r}'), \quad (2.3.3)$$

where we introduced the shorthand notation

$$(P|Q)^{-1} \equiv (\mathbf{C}^{-1})_{PQ}, \quad (2.3.4)$$

$$C_{PQ} = (P|Q) = \iint d\mathbf{r} d\mathbf{r}' P(\mathbf{r}) \frac{1}{|\mathbf{r} - \mathbf{r}'|} Q(\mathbf{r}'), \quad (2.3.5)$$

and P, Q are elements of an auxiliary basis set, which is typically about three times as big as the atomic orbital basis set. RI is also known as “Density Fitting”, because derivation of Eq. 2.3.2^[74] proceeds by fitting an auxiliary density $\tilde{\rho}$ to the exact density ρ ,

$$\rho(\mathbf{r}) = \sum_{\mu\nu} \mu(\mathbf{r}) \nu(\mathbf{r}) P_{\nu\mu} \approx \tilde{\rho}(\mathbf{r}) = \sum_P \chi_P(\mathbf{r}) C_{\mu\nu}^P, \quad (2.3.6)$$

and minimizing the residual self-repulsion R

$$R = \iint d\mathbf{r} d\mathbf{r}' \frac{(\tilde{\rho} - \rho)(\mathbf{r}) (\tilde{\rho} - \rho)(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}. \quad (2.3.7)$$

Equation 2.3.2 has been the ubiquitous form of RI for some time. It is used to approximate ERIs in Kohn–Sham,^[74,78,79] Hartree–Fock,^[80] second-order Møller–Plesset,^[77,81,82] as well as coupled-cluster theories.^[83,84]

However, Eq. 2.3.2 is only uniquely defined through Eq. 2.3.7. Alternatively, minimizing the norm

$$|\delta_{\mu\nu}|^2 = \int d\mathbf{r} \left| (\mu\nu)(\mathbf{r}) - \sum_P \chi_P(\mathbf{r}) C_{\mu\nu}^P \right|^2 \quad (2.3.8)$$

leads to^[74]

$$(\mu\nu|\lambda\sigma) \approx \sum_{PQRS} (\mu\nu|P) (P|Q)^{-1} (Q|R) (RS)^{-1} (S|\lambda\sigma), \quad (2.3.9)$$

with

$$(P|Q)^{-1} \equiv (\mathbf{S}^{-1})_{PQ}, \quad (2.3.10)$$

$$S_{PQ} = \int d\mathbf{r} P(\mathbf{r}) Q(\mathbf{r}), \quad (2.3.11)$$

$$(\mu\nu|P) = \int d\mathbf{r} \mu(\mathbf{r}) \nu(\mathbf{r}) P(\mathbf{r}). \quad (2.3.12)$$

In terms of linear algebra, equations 2.3.2 and 2.3.9 differ in their definition of the norm (and hence the metric thereby induced). Equation 2.3.2 is known as Coulomb-metric RI (or

“RI-V”), while Eq. 2.3.9 corresponds to the overlap-metric RI (or “RI-SVS”). In the limit of a complete auxiliary basis, any metric will give the correct result. In finite bases, Eq. 2.3.2 has been found to yield superior accuracy,^[74,76] which is why it was universally adopted.

Coulomb-metric RI has a distinct disadvantage compared to overlap-metric RI, however. In the context of linear-scaling techniques, the three-center two-electron integrals (Eq. 2.3.3) offer only limited useable sparsity (cf. sec. 3.3) in their atomic orbital representations. In local bases (e. g., Gaussian functions), the product vector $|\mu\nu\rangle$ vanishes if $|\mu\rangle$ and $|\nu\rangle$ are centered on points sufficiently far apart. This criterion is universally used in all programs using local basis functions for virtually all integral evaluations, and it is also useful in the evaluation of Eq. 2.3.3. However, the $\frac{1}{r}$ coupling between bra and ket in $(\mu\nu|P)$ means whenever $|\mu\nu\rangle$ is significant, the integral $(\mu\nu|P)$ will also be significant. In other words, the number of significant Coulomb-metric three-center RI integrals grows as $O(N^2)$. In contrast, overlap-metric three-center integrals offer extremely high useable sparsity, since the integral decays exponentially in the distance of bra and ket.

Overlap-metric RI was reintroduced by Schurkus and Ochsenfeld^[38] for an effective linear-scaling RPA algorithm within the RI approximation. Techniques to combine the advantages of overlap- and Coulomb-metric RI is discussed in article I (page 43) and sec. 3.7.

2.4 Correlation Energies from the Random Phase Approximation

The Random Phase Approximation (RPA) originates in a series of publications by Bohm and Pines.^[33,85,86] Langreth and Perdew^[87] first defined the RPA correlation energy. In modern electronic structure literature, the RPA correlation energy is most often derived from the adiabatic connection^[87–90] and the fluctuation-dissipation theorem.^[91] Here, we follow Eshuis *et al.*^[92] for the first steps.

2.4.1 Adiabatic Connection Correlation Energy

A well-known result from the adiabatic connection formalism for the correlation energy E_C is

$$E_C = \int_0^1 d\alpha \left(\langle \Psi_0^\alpha | \hat{V}_{ee} | \Psi_0^\alpha \rangle - \langle \Psi_0^{\alpha=0} | \hat{V}_{ee} | \Psi_0^{\alpha=0} \rangle \right), \quad (2.4.1)$$

where α is the coupling strength, Ψ_0^α is the ground-state wave function at coupling strength α , and \hat{V}_{ee} is the exact electron-electron interaction. In second quantization, \hat{V}_{ee} may be rewritten as^[8]

$$\hat{V}_{ee} = \frac{1}{2} \sum_{pqrs} \langle pq|rs \rangle \hat{c}_p^\dagger \hat{c}_q^\dagger \hat{c}_s \hat{c}_r, \quad (2.4.2)$$

with the creation (\hat{c}_p^\dagger) and annihilation (\hat{c}_p) operators. In terms of the two-particle density operator $\hat{P}(x_1, x_2)$,

$$\hat{V}_{ee} = \iint dx_1 dx_2 \frac{\hat{P}(x_1, x_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} = \iint dx_1 dx_2 \hat{P}(x_1, x_2) v(x_1, x_2), \quad (2.4.3)$$

where $x_i = (\mathbf{r}_i, \sigma_i)$ is a combined spin-space coordinate and $v(x_1, x_2) = 1/|\mathbf{r}_1 - \mathbf{r}_2|$ is the bare Coulomb interaction. Using electron field operators^[93]

$$\hat{\psi}^\dagger(x) = \sum_p \phi_p^*(x) \hat{c}_p^\dagger, \quad (2.4.4)$$

the two-particle density operator reads

$$\hat{P}(x_1, x_2) = \frac{1}{2} \hat{\psi}^\dagger(x_1) \hat{\psi}^\dagger(x_2) \hat{\psi}(x_2) \hat{\psi}(x_1), \quad (2.4.5)$$

which may be factorized as

$$\hat{P}(x_1, x_2) = \frac{1}{2} [\hat{\rho}(x_1) \hat{\rho}(x_2) - \delta(x_1 - x_2) \hat{\rho}(x_1)], \quad (2.4.6)$$

using the fermionic anti-commutator relations^[93]

$$[\hat{\psi}(x_1), \hat{\psi}^\dagger(x_2)]_+ = \delta(x_1 - x_2), \quad (2.4.7)$$

$$[\hat{\psi}(x_1), \hat{\psi}(x_2)]_+ = [\hat{\psi}^\dagger(x_1), \hat{\psi}^\dagger(x_2)]_+ = 0. \quad (2.4.8)$$

Introducing the density-fluctuation operator^[93]

$$\Delta\hat{\rho}(x) = \hat{\rho}(x) - \rho(x), \quad (2.4.9)$$

equation 2.4.6 becomes

$$\hat{P}(x_1, x_2) = \frac{1}{2} [\Delta\hat{\rho}(x_1) \Delta\hat{\rho}(x_2) + \hat{\rho}(x_1) \rho(x_2) + \rho(x_1) \hat{\rho}(x_2) - \delta(x_1 - x_2) \hat{\rho}(x_1)]. \quad (2.4.10)$$

Inserting Eqs. 2.4.3 and 2.4.10 into the correlation energy expression from the adiabatic connection (Eq. 2.4.1) leaves just

$$E_C = \frac{1}{2} \int_0^1 d\alpha \iint dx_1 dx_2 \left[\frac{\langle \Psi_0^\alpha | \Delta\hat{\rho}(x_1) \Delta\hat{\rho}(x_2) | \Psi_0^\alpha \rangle}{|\mathbf{r}_1 - \mathbf{r}_2|} - \frac{\langle \Psi_0^{\alpha=0} | \Delta\hat{\rho}(x_1) \Delta\hat{\rho}(x_2) | \Psi_0^{\alpha=0} \rangle}{|\mathbf{r}_1 - \mathbf{r}_2|} \right], \quad (2.4.11)$$

because all other terms vanish identically for Ψ_0^α and $\Psi_0^{\alpha=0}$. Inserting the resolution-of-the-identity

$$1 = \sum_n |\Psi_n^\alpha\rangle \langle \Psi_n^\alpha| \quad (2.4.12)$$

between $\Delta\hat{\rho}(x_1)$ and $\Delta\hat{\rho}(x_2)$ yields

$$E_C = \frac{1}{2} \int_0^1 d\alpha \iint dx_1 dx_2 \sum_n \left[\frac{\langle \Psi_0^\alpha | \hat{\rho}(x_1) | \Psi_n^\alpha \rangle \langle \Psi_n^\alpha | \hat{\rho}(x_2) | \Psi_0^\alpha \rangle}{|\mathbf{r}_1 - \mathbf{r}_2|} - \frac{\langle \Psi_0^{\alpha=0} | \hat{\rho}(x_1) | \Psi_n^{\alpha=0} \rangle \langle \Psi_n^{\alpha=0} | \hat{\rho}(x_2) | \Psi_0^{\alpha=0} \rangle}{|\mathbf{r}_1 - \mathbf{r}_2|} \right]. \quad (2.4.13)$$

Written more compactly,

$$E_C = \frac{1}{2} \int_0^1 d\alpha \iint dx_1 dx_2 \sum_n \left[\rho_{0n}^\alpha(x_1) \rho_{0n}^\alpha(x_2) v(x_1, x_2) - \rho_{0n}^{\alpha=0}(x_1) \rho_{0n}^{\alpha=0}(x_2) v(x_1, x_2) \right], \quad (2.4.14)$$

with

$$\rho_{0n}^\alpha(x) = \langle \Psi_0^\alpha | \hat{\rho}(x) | \Psi_n^\alpha \rangle. \quad (2.4.15)$$

2.4.2 Frequency Integration of the Polarization Propagator

Next, we inspect the Lehmann representation of the (retarded) polarization propagator^[93,94] at coupling strength α

$$\chi_\alpha(x_1, x_2, \omega) = \lim_{\eta \rightarrow 0^+} \sum_n \left(\frac{\rho_{0n}^\alpha(x_1) \rho_{0n}^\alpha(x_2)}{\omega - \Omega_n^\alpha + i\eta} - \frac{\rho_{0n}^\alpha(x_1) \rho_{0n}^\alpha(x_2)}{\omega + \Omega_n^\alpha + i\eta} \right), \quad (2.4.16)$$

which has first-order poles at $z_0 = \Omega_n^\alpha - i\eta$ and $z_1 = -\Omega_n^\alpha - i\eta$ with residues

$$\text{Res } \chi_\alpha(x_1, x_2, z)|_{z=z_0} = \rho_{0n}^\alpha(x_1) \rho_{0n}^\alpha(x_2), \quad (2.4.17)$$

$$\text{Res } \chi_\alpha(x_1, x_2, z)|_{z=z_1} = -\rho_{0n}^\alpha(x_1) \rho_{0n}^\alpha(x_2). \quad (2.4.18)$$

Using Cauchy's residue theorem, integration along the imaginary axis can be performed by counter-clockwise integration along the contour Γ enclosing z_1 , yielding

$$\int_{-\infty}^{\infty} du \chi_\alpha(x_1, x_2, iu) = \oint_{\Gamma} dz \chi_\alpha(x_1, x_2, z) \quad (2.4.19)$$

$$= 2\pi i \text{Res } \chi_\alpha(x_1, x_2, z)|_{z=z_1} \quad (2.4.20)$$

$$= -2\pi i \rho_{0n}^\alpha(x_1) \rho_{0n}^\alpha(x_2). \quad (2.4.21)$$

Inserting this result in Eq. 2.4.14 gives a useful expression for the correlation energy

$$E_C = -\frac{1}{4\pi} \text{Im} \int_0^1 d\alpha \int_{-\infty}^{\infty} du \iint dx_1 dx_2 \left[\chi_\alpha(x_1, x_2, iu) - \chi_0(x_1, x_2, iu) \right] v(x_1, x_2), \quad (2.4.22)$$

known as the adiabatic-connection fluctuation-dissipation-theorem correlation energy.^[87,89,95–97]

2.4.3 Dyson Equation and the Random Phase Approximation

We switch to matrix notation, where multiplication is understood to mean

$$(\mathbf{ab})_{xx'} = \sum_{x''} \mathbf{a}_{xx''} \mathbf{b}_{x''x'} \quad (2.4.23)$$

$$= \int dx'' a(x, x'') b(x'', x'). \quad (2.4.24)$$

The polarization propagator χ_α at coupling strength α obeys a Dyson- or Bethe–Salpeter-type equation^[34,93,98]

$$\chi_\alpha(x_1, x_2, \omega) = \chi_0(x_1, x_2, \omega) + \iint dx'_1 dx'_2 \chi_0(x_1, x'_1, \omega) \left[\alpha v(x'_1, x'_2) + f_\alpha^{\text{XC}}(x'_1, x'_2, \omega) \right] \chi_\alpha(x'_2, x_2, \omega), \quad (2.4.25)$$

or

$$\chi_\alpha(\omega) = \chi_0(\omega) + \chi_0(\omega) \left[\alpha \mathbf{v} + f_\alpha^{\text{XC}}(\omega) \right] \chi_\alpha(\omega), \quad (2.4.26)$$

where f_α^{XC} is the (frequency-dependent) exchange–correlation kernel at coupling strength α . The (direct) random phase approximation simply omits f_α^{XC} so that Eq. 2.4.26 becomes

$$\chi_\alpha^{\text{RPA}}(\omega) = \chi_0(\omega) + \alpha \chi_0(\omega) \mathbf{v} \chi_\alpha^{\text{RPA}}(\omega), \quad (2.4.27)$$

which may be rearranged to

$$\chi_\alpha^{\text{RPA}}(\omega) = \left[\chi_0^{-1}(\omega) - \alpha \mathbf{v} \right]^{-1}, \quad (2.4.28)$$

where the coupling-strength integral may be evaluated analytically

$$\int_0^1 d\alpha \text{Tr} \left(\chi_\alpha^{\text{RPA}}(\omega) \mathbf{v} \right) = \int_0^1 d\alpha \text{Tr} \left(\left[\chi_0^{-1}(\omega) - \alpha \mathbf{v} \right]^{-1} \mathbf{v} \right) \quad (2.4.29)$$

$$= -\text{Tr} \ln (\mathbf{1} - \chi_0(\omega) \mathbf{v}). \quad (2.4.30)$$

Combining this result with 2.4.22 yields a definition of the RPA correlation energy

$$E_C^{\text{RPA}} = \frac{1}{4\pi} \text{Im} \int_{-\infty}^{\infty} du \text{Tr} \left[\ln (\mathbf{1} - \chi_0(iu) \mathbf{v}) - \chi_0(iu) \mathbf{v} \right]. \quad (2.4.31)$$

3 Additional Results

This chapter includes additional results and remarks not included in articles **I-IV**.

3.1 The Pivoted Cholesky Factorization for DL-CPSCF

Both the occupied ($\mathbf{P}^{(\alpha)}$) and virtual ($\mathbf{Q}^{(\alpha)}$) pseudo-densities are positive semi-definite and therefore amenable to a pivoted Cholesky factorization:

$$\mathbf{P}^{(\alpha)} = \sqrt{w_\alpha} \exp(t_a \mathbf{P}\mathbf{F}) \mathbf{P} = \tilde{\mathbf{P}}^{(\alpha)} \tilde{\mathbf{P}}^{(\alpha)\top}, \quad (3.1.1)$$

$$\mathbf{Q}^{(\alpha)} = \sqrt{w_\alpha} \exp(-t_a \mathbf{Q}\mathbf{F}) \mathbf{Q} = \tilde{\mathbf{Q}}^{(\alpha)} \tilde{\mathbf{Q}}^{(\alpha)\top}. \quad (3.1.2)$$

The occupied matrices have rank N_{occ} , and the virtual matrices have rank N_{virt} . The Cholesky factors $\tilde{\mathbf{P}}^{(\alpha)}$ and $\tilde{\mathbf{Q}}^{(\alpha)}$ consequently have dimension $N_{\text{basis}} \times N_{\text{occ}}$ and $N_{\text{basis}} \times N_{\text{virt}}$, respectively. The canonical DL-CPSCF equation (2.2.26) can be rewritten using the Cholesky factors as

$$\mathbf{P}_{\text{vo}}^\xi = \sum_{\alpha=1}^{\tau} \tilde{\mathbf{Q}}^{(\alpha)} \tilde{\mathbf{Q}}^{(\alpha)\top} \left(-\mathbf{h}^x - \mathbf{G} [\mathbf{P}^\xi] \right) \tilde{\mathbf{P}}^{(\alpha)} \tilde{\mathbf{P}}^{(\alpha)\top}. \quad (3.1.3)$$

This expression was first given by the author in articles **III** (page 79) and **IV** (page 91).

Because not all matrices here are square, the order of multiplications is important to performance. The optimum sequence depends on matrix sparsity, as well as the ratio of occupied to virtual orbitals. We will omit discussion of linear dependencies in the atomic orbital basis, because they are easily eliminated by projection, and reduce the total number of molecular orbitals only by a small fraction over the number of atomic orbitals.

In typical calculations with basis sets of useful size (i. e. non-minimal bases), the number of occupied orbitals is much smaller than the total number of orbitals. We will assume $N_{\text{occ}} = \frac{1}{10} N_{\text{basis}}$ (and therefore $N_{\text{virt}} = \frac{9}{10} N_{\text{basis}}$) in the following, which can be considered typical for basis sets of double- ζ quality.

In the limiting case where all matrices are densely filled, evaluation of Eq. 2.2.26 requires $2N_{\text{basis}}^3$ multiplications per Laplace point α . The optimal strategy for evaluation of Eq. 3.1.3 is to perform the Cholesky factorization only on $\mathbf{P}^{(\alpha)}$, and proceed by multiplying the integrand $-\mathbf{h}^x - \mathbf{G} [\mathbf{P}^x]$ by the left Cholesky factor $\tilde{\mathbf{P}}^{(\alpha)}$ from the right, followed by multiplication with $\mathbf{Q}^{(\alpha)}$ from the left, and finally multiplication by $\tilde{\mathbf{P}}^{(\alpha)\top}$ from the right. Assuming $N_{\text{occ}} = \frac{1}{10} N_{\text{basis}}$, this sequence requires $0.3N_{\text{basis}}^3$ multiplications, improving

computational time $6\times$. Performing the Cholesky factorization also on $\mathbf{Q}^{(\alpha)}$ increases the number of multiplications to $0.38N_{\text{basis}}^3$.

This discussion typically also holds for sparse matrices. In many cases, however, factorizing both $\mathbf{Q}^{(\alpha)}$ and $\mathbf{P}^{(\alpha)}$ does turn out to be superior after all. This is due to rank reductions beyond N_{occ} and N_{virt} , discussed in sec. 3.2.

3.2 Rank Reduction Through Cholesky Factorization

For a fixed number of integration nodes τ in Eq. 3.1.3, the sum of the numerical ranks of all $\mathbf{Q}^{(\alpha)}$ becomes smaller the smaller the HOMO-LUMO gap gets. This follows from the numerical quadrature of the Laplace-transformed term $\frac{1}{x}$. In the interval $x \in [1, R]$,

$$\frac{1}{x} \simeq \sum_{\alpha=1}^{\tau} \bar{w}_{\alpha} \exp(-\bar{t}_{\alpha}x). \quad (3.2.1)$$

The corresponding roots $\{w_{\alpha}\}$ and weights $\{t_{\alpha}\}$ for the interval $y = Ax \in [A, B]$ with $B = AR$ are^[99]

$$\frac{1}{y} \simeq \sum_{\alpha=1}^{\tau} w_{\alpha} \exp(-t_{\alpha}y), \quad (3.2.2)$$

$$w_{\alpha} = \bar{w}_{\alpha}/A, \quad (3.2.3)$$

$$t_{\alpha} = \bar{t}_{\alpha}/A. \quad (3.2.4)$$

In this particular application, the quadrature interval is

$$y \in [\varepsilon_{\text{LUMO}} - \varepsilon_{\text{HOMO}}, \varepsilon_{\text{high}} - \varepsilon_{\text{low}}], \quad (3.2.5)$$

where $\varepsilon_{\text{HOMO}}$ and $\varepsilon_{\text{LUMO}}$ are the HOMO and LUMO energies, respectively, and $\varepsilon_{\text{high}}$ and ε_{low} are the overall highest and lowest MO energies, respectively. For closing HOMO-LUMO gaps, therefore, t_{α} in Eq. 3.2.2 grows and the numerical rank of $\mathbf{Q}^{(\alpha)}$ in Eq. 2.2.28 shrinks.

This means, for instance, that the pivoted Cholesky decomposition provides larger rank reductions for non-hybrid KS calculations than for HF calculations (for a fixed number of nodes τ), since the former generally produces smaller gaps. In turn, the integration interval y in HF calculations is smaller than in KS calculations, meaning fewer quadrature points τ may be sufficient.

The numerical properties of the Cholesky factorization are discussed in Ref. 100. In LAPACK (routine `dpstrf`), accuracy is controlled by a tolerance parameter. We have found that the pseudo-density matrices are generally well-behaved with regard to the pivoted Cholesky factorization. The default tolerance in LAPACK introduces errors in the order of the machine precision. The numerical rank of the Cholesky factors can be artificially reduced at the cost of accuracy by using a less strict tolerance parameter. In our tests,

however, forcing the numerical rank to be just 10 % below the rank revealed by the default tolerance leads to unacceptable errors.

As a general rule, numerical results are improved when performing the orbital energy shift (Eq. 2.2.25) discussed in sec. 2.2.4. This is particularly true when using loose tolerances. In exact arithmetic, ε_F in Eq. 2.2.25 can be chosen arbitrarily. In double precision arithmetic, however, shifting the orbital energies by as little as ± 5 Hartree leads to severe numerical instabilities. Thus, we always choose $\varepsilon_F = (\varepsilon_{\text{HOMO}} + \varepsilon_{\text{LUMO}})/2$, which means all occupied MO energies will be negative, and all virtual MO energies will be positive. Consequently, the argument to the exponential function in Eqs. 2.2.27 and 2.2.28 is always negative, and numerical singularities are avoided.

In AO-RPA calculations, derivatives of pseudo-density matrices must be evaluated. In ω -CDD-RPA (article I, page 43), these matrices are Cholesky-factorized. Particularly in pure (non-hybrid) Kohn–Sham calculations, which form the basis of an RPA calculation, a few low-lying virtual orbitals with negative orbital energies are the norm rather than the exception. In such a case, the spectral representation of a virtual pseudo-density matrix derivative

$$-\frac{\partial}{\partial t_\alpha} \mathbf{Q}_{\mu\nu}^{(\alpha)} = \sqrt{w_\alpha} \sum_a C_{\mu a} \varepsilon_a \exp(-t_\alpha \varepsilon_a) C_{a\nu}^\dagger \quad (3.2.6)$$

reveals the matrix is manifestly not positive semi-definite and hence not amenable to Cholesky factorization, unless all virtual orbital energies are shifted to positive values.

Articles III (page 79) and IV (page 91) established that performing Cholesky-factorizing the virtual pseudo-density matrices in addition to their occupied counterparts can increase performance and improve numerical stability, particularly when using sparse algebra. This is despite the fact that the theoretical number of multiplications increases, as discussed in sec. 3.1.

3.3 Useable Matrix Sparsity

The key difficulty in the design and implementation of low-scaling variants of *ab initio* theories in the atomic orbitals lies in the matrix sparsity of relevant quantities. It is well known that the overlap, Fock, and density matrices become sparsely populated—in a local basis (such as Gaussian functions)—for molecules of sufficient size with non-vanishing band gaps.

However, unlike other problem domains—such as finite-element calculations—only certain specific forms of matrix sparsity are useful for AO-based low-scaling methods. Finite-element calculations often see matrices of dimensions of several hundred thousand rows and columns, with extremely high (> 99 %) sparsity. Large SCF calculations, on the other hand, might have matrix dimension of several ten thousand rows and columns, with just 50–90 % sparsity. The immediate consequence is that most sparse matrix storage schemes, such as the compressed-sparse-row (CSR) format, have extremely poor performance in this problem domain.

The most successful sparse storage format for *ab initio* calculations is the blocked-compressed sparse row (BCSR) format,^[101–105] which divides the source matrix into sub-matrices (typically 100×100 in size) and discards all blocks whose matrix norm is below a numerical threshold. This format allows for matrix multiplications with very low overhead, because the sub-matrix multiplications are delegated to high-performance libraries, often supplied by the processor vendor.

The caveat of the BCSR format, however, is that it only works well for banded sparse matrices (see figure 3.1 on the facing page for examples). If the significant elements of a matrix are dispersed all over the matrix, effectively all blocks must be allocated, and no compression is possible with the BCSR format.

This leads to a concept of useable matrix sparsity within AO-based quantum chemical calculations. The useable sparsity can be measured by, for example, the matrix bandwidth. An alternative criterion is the percentage of allocated BCSR blocks required for a certain accuracy.

Designing an AO-based low-scaling quantum chemical algorithm requires that the matrix operations preserve useable sparsity. This requirement is the key difficulty. It effectively eliminates the use of eigendecompositions and singular value decompositions in the algorithm, because the eigenvectors and singular vectors of the Fock and SCF density matrices have no useable sparsity. One highly useful sparsity-preserving matrix factorization is the Cholesky decomposition, extensively used in this work in articles **I** through **IV**. Two very important aspects of useable matrix sparsity will be discussed in sections 3.4 and 3.5.

3.4 Sparsity of the Virtual Density Matrix

One of the more common misleading statements in much of the literature concerning the sparsity of density matrices concerns the virtual density matrix (Eq. 2.1.21 on page 18). It is true that, for typical molecular calculations using atom-centered Gaussian basis sets, the *contravariant atomic orbital* representation of the virtual density matrix \mathbf{Q} is much less sparse than the corresponding occupied density matrix representation. However, Eq. 2.1.21 simplifies to

$$\mathbf{Q} = \mathbf{1} - \mathbf{P} \quad (3.4.1)$$

when the atomic orbital basis has been orthogonalized. In other words, the occupied and virtual density matrices have identical sparsity patterns in an orthogonalized basis. There is no physical reason whatsoever dictating the virtual density matrix be any less local (sparse) than the occupied density matrix. The confusion arises due to the presence of \mathbf{S}^{-1} in Eq. 2.1.21. In Gaussian orbitals, the overlap matrix \mathbf{S} is extremely sparse, because elements $S_{\mu\nu}$ decay as $\exp(-r^2)$ in the distance r of their centers. Matrix inversion will destroy a significant portion (but by no means all) of this sparsity. Hence, in the contravariant AO basis, the virtual density matrix \mathbf{Q} is much less sparse than its occupied counterpart \mathbf{P} .

In an orthogonalized basis, however, the sparsity of \mathbf{Q} and \mathbf{P} is identical. If the aim is to obtain a sparse matrix representation of the occupied *and* virtual density matrices, this

translates to finding an a sparsity-preserving orthogonalization of the atomic orbital basis.

Löwdin orthogonalization^[106] provides orthogonalized AOs which are most similar to the original AOs in a least-squares sense,^[107,108] and hence inherit most of their locality. Another possible choice is the Cholesky factorization of \mathbf{S} as $\mathbf{S} = \mathbf{L}\mathbf{L}^T$. Using \mathbf{L} to orthogonalize contravariant quantities, and \mathbf{L}^{-1} to orthogonalize covariant quantities. This procedure is equivalent to Gram–Schmidt orthogonalization, and is much faster to compute than Löwdin’s orthogonalization. Like Löwdin’s procedure, it preserves much of the existing sparsity. Sparsity patterns of occupied and virtual density matrices (computed for $n\text{-C}_{80}\text{H}_{162}$, PBE/def2-SVP) are given in figure 3.1.




Figure 3.1: Sparsity patterns of occupied and virtual density matrices in different bases. From left to right: \mathbf{P} and \mathbf{Q} (contravariant AO basis); \mathbf{P} and \mathbf{Q} (orthogonal Cholesky basis). In the orthogonalized basis, occupied and virtual density matrices have identical sparsity patterns, whereas the virtual density matrix is significantly less sparse in the contravariant AO basis.

3.5 Sparsity of Cholesky-Decomposed Density Matrices

An interesting and highly surprising result of the different orthogonalization techniques arose in the course of preparing article I (page 43). We used electronically sparse, highly symmetric, linear molecules as examples to demonstrate the computational complexity of our new method, ω -CDD-RPA (article I). The question of whether to factorize virtual density-like quantities in addition to factorizing occupied density-like quantities posed itself. The surprising result was that factorizing virtual density matrices in either the (non-orthogonal) contravariant AO basis, or the (orthogonal) Löwdin basis gave mostly local pseudo-molecular orbitals, but invariably showed a number of completely non-local MOs which spanned the whole AO space. When factorizing Cholesky-orthogonalized density matrices, however, sparsity and banded structure in the pseudo-MOs was perfectly preserved. Figure 3.2 shows typical examples of the phenomenon.

We attribute this behavior to an interaction between the molecular symmetry and the orthogonalization procedure. A possible (partial) explanation is as follows: Löwdin’s method, also known as *symmetric* orthogonalization,^[43] tries to orthogonalize symmetry-equivalent AO basis functions a symmetry-equivalent footing. The combined orthogonality *and* “symmetry” constraints forced the generation of several non-local, far-reaching orthogonalized orbitals. In contrast, the Cholesky orthogonalization procedure does not have any sym-




Figure 3.2: Sparsity patterns of the column-reordered Cholesky factors of $\mathbf{QFQF}^{(\alpha)}$. Left: intermediate Löwdin orthogonalization. A large number of pseudo-MOs span the whole AO space. Right: intermediate Cholesky orthogonalization. The useable matrix sparsity is fully retained.

metry constraints whatsoever. In fact, the transformation matrices, being triangular, are maximally asymmetric in some sense. This additional degree of freedom enables Cholesky-orthogonalized density matrices to yield very sparse pseudo-MOs when factorized.

3.6 Local Perturbations

Articles **III** (page 79) and **IV** (page 91) of this dissertation are concerned with the linear-scaling computation of second-order molecular properties at the SCF level. Specifically, **III** describes the evaluation of indirect nuclear spin–spin coupling constants with constant time-complexity. **IV** describes the evaluation of harmonic vibrational frequencies. For both of these properties, the perturbation (cf. sec. 2.2.2) is *local*, i. e., the effect of the perturbation is spatially confined to only parts of the molecule. This is in contrast to *global* perturbations, which affect the whole molecular electronic structure—an external magnetic field, for example. Local perturbations require considerations beyond those necessary for global perturbations.

3.6.1 Nuclear Displacement

The local perturbation in article **IV** (page 91) is the nuclear displacement, that is, the infinitesimal movement of an individual nucleus in all three spatial directions. The Hellmann–Feynman theorem (sec. 2.2.1) suggests the decay behavior of the perturbation is dictated only by the operator derivatives. In the case of nuclear displacement, the crucial operator would seem to be the derivative of the electron-nuclear interaction with respect to nuclear displacement. Whereas the electron-nuclear interaction decays as r^{-1} , its derivative decays as r^{-2} .

This line of argumentation holds in the limit of a complete basis. Typical basis sets commonly in use for SCF calculations are nowhere near complete, however. This gives rise to additional terms in the derivative of the SCF energy expression, which are long-range (r^{-1} decay), and effectively prohibit a linear-scaling behavior. Article **IV** showed that, even in incomplete bases, exploiting the r^{-2} decay is still possible. Consequently, evaluation of a density matrix perturbed by an individual nuclear displacement decays rapidly enough to enable constant-time evaluation. Evaluation of *all* nuclear displacements in a molecule, a prerequisite for the calculation of harmonic vibrational frequencies, is therefore possible in linear time.

3.6.2 Magnetic Perturbations

The perturbations in article **III** (page 79) are Fermi contact, spin-dipole, as well as paramagnetic and diamagnetic spin-orbit interaction operators. The Fermi contact operator decays exponentially in Gaussian basis set. The Spin-dipole operator decays as r^{-3} , but its atomic orbital representation increases strongly in sparsity after projection onto the virtual-occupied subspace, which is the only subspace required to solve the CPSCF equations. Finally, the combination of paramagnetic and diamagnetic spin-orbit operators also decays as r^{-3} .^[109] The individual parts, however, only decay as r^{-2} . Because they are calculated very differently, combined evaluation is not yet possible to the best of the author's knowledge. Whether or not this is possible in principle is subject to further research.

3.7 Local Metrics in the Resolution-of-the-Identity

Coulomb- and overlap-metric resolution-of-the-identity (RI) have been discussed in sec. 2.3. Article **I** (page 43) goes beyond overlap-metric RI. One of the techniques introduced by the author in that article is the usage of an attenuated Coulomb-metric RI within RPA calculations.

In fact, any inner product defined for the union of vector spaces spanned by the union of atomic orbital and auxiliary bases induces an eligible RI metric. The Coulomb metric minimizes the error in the Hartree energy (cf. sec. 21). In the limit of infinite bases, however, all metrics yield the exact result.

As discussed in sec. 2.3, the Coulomb metric affords no useable sparsity beyond that granted by the Gaussian product theorem (or its equivalent for other local bases). The overlap metric, on the other hand, affords very high useable sparsity at the cost of accuracy.

Attenuated Coulomb operators have been known for a long time. Gill *et. al.* published a series of articles^[110–116] in which the resolution

$$\frac{1}{r_{12}} = \frac{\operatorname{erf}(\omega r_{12})}{r_{12}} + \frac{\operatorname{erfc}(\omega r_{12})}{r_{12}} \quad (3.7.1)$$

was disseminated, where erf is the error function, and erfc is the complementary error function. This partitioning splits the long-range $\frac{1}{r}$ -operator into a short-range (erfc) and a




Figure 3.3: Sparsity patterns of three-center two-electron RI integrals in different metrics in the def2-SVP basis, for $n\text{-C}_{80}\text{H}_{162}$. Shown is the pattern for one matrix $(\mu\nu | \text{erfc}(\omega r_{12}) / r_{12} | P)$ at a fixed auxiliary function P . From left to right: $\omega = 0$ (Coulomb metric); $\omega = 0.1$; $\omega = 1$; $\omega = \infty$ (overlap metric).

long-range (erf) part, with the attenuation parameter ω .

The first use of the erfc-operator within the resolution-of-the-identity was, to the best of the author’s knowledge, given by Jung *et al.* [117]. In 2007, Jung *et al.* [118] presented a variant of scaled-opposite-spin second-order Møller–Plesset perturbation theory [119] using an erfc-attenuated Coulomb metric. Later, attenuated Coulomb operators were studied in the context of RI-HF and RI-KS calculations. [120,121]

Article **I** introduces the erfc-attenuated Coulomb metric for resolution-of-the-identity calculations of the random-phase-approximation correlation energy. It is demonstrated that the erfc-Coulomb metric can combine the advantages of the (long-range but accurate) Coulomb- and (short-range but less accurate) overlap-metrics. It was shown that, over a considerable interval of attenuation parameters ω , energies matched the accuracy of the (unattenuated) Coulomb metric, while sparsity (and consequently, performance) rivaled that of the overlap metric.

Figure 3.3 shows sparsity patterns of three-center two-electron RI integrals in different metrics in the def2-SVP basis, for $n\text{-C}_{80}\text{H}_{162}$. Concerning useable sparsity, variation of the attenuation parameter ω provides a smooth interpolation between the Coulomb ($\omega = 0$) and overlap ($\omega = \infty$) RI metrics. This aspect, as well as accuracy in RPA correlation energies is discussed in detail in article **I** (page 43).

4 Summary

This work introduces novel methods for the low-scaling calculation of molecular properties and correlation energies.

Using an attenuated Coulomb metric in the resolution-of-the-identity and Cholesky-decomposed pseudo-density matrices, correlation energies from the random phase approximation can be computed with $\mathcal{O}(N)$ scaling. The newly introduced method has virtually no overhead compared to the canonical molecular orbital-based formulation [Eshuis *et al.*^[35], $\mathcal{O}(N^4)$] for small molecules, and outperforms it by an order of magnitude for large molecules due to superior scaling.

A technique for the *ad-hoc* compression of auxiliary basis sets for resolution-of-the-identity (RI) correlation energy calculations (e. g. MP2, RPA) has been presented. Projection onto the particle-hole space followed by construction of a rank-reducing auxiliary basis transformation enables calculation of correlation energies to approach RI-free reference accuracy, while keeping the performance and memory benefits of RI intact.

Molecular property calculations at the Hartree–Fock and Kohn–Sham density functional theories with linear and sub-linear scaling behavior have been developed. Linear-scaling calculation of harmonic vibrational frequencies for molecules with a non-vanishing electronic gap have been demonstrated, made possible by explicit consideration of cancellation in long-range derivatives of the Coulomb interaction. No real-space cutoffs or thresholds are required and no further approximations are introduced in our novel technique. Only physically justifiable arguments—cancellation of long-range operators and distance-decay in the density matrix—are employed in the derivation and implementation.

Similar methods were used to enable $\mathcal{O}(1)$ calculation of nuclear spin–spin couplings in molecular systems, i. e., the computational effort to compute one coupling constant is asymptotically independent of system size. Here too, no additional approximations over existing response theory are necessary. The new method was used to study the convergence behavior of spin–spin couplings in an aminopyrazole peptide with respect to increases in the solvation environment size. The study revealed that large solvent environments (500 atoms and more) are required to converge the observables, underlining the necessity of low-scaling methods for molecular property calculations.

Bibliography

- [1] G. D. Purvis and R. J. Bartlett, *J. Chem. Phys.* **76**, 1910 (1982).
- [2] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon, *Chem. Phys. Lett.* **157**, 479 (1989).
- [3] K. L. Bak, J. Gauss, T. Helgaker, P. Jørgensen, and J. Olsen, *Chem. Phys. Lett.* **319**, 563 (2000).
- [4] K. L. Bak, P. Jørgensen, J. Olsen, T. Helgaker, and W. Klopper, *J. Chem. Phys.* **112**, 9229 (2000).
- [5] S. Coriani, D. Marchesan, J. Gauss, C. Hättig, T. Helgaker, and P. Jørgensen, *J. Chem. Phys.* **123**, 184107 (2005).
- [6] R. J. Bartlett and M. Musiał, *Rev. Mod. Phys.* **79**, 291 (2007).
- [7] C. Møller and M. S. Plesset, *Phys. Rev.* **46**, 618 (1934).
- [8] T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley, Chichester, 2000).
- [9] M. W. Schmidt and M. S. Gordon, *Annu. Rev. Phys. Chem.* **49**, 233 (1998).
- [10] K. Andersson, P. A. Malmqvist, B. O. Roos, A. J. Sadlej, and K. Wolinski, *J. Phys. Chem.* **94**, 5483 (1990).
- [11] K. Andersson, P. Malmqvist, and B. O. Roos, *J. Chem. Phys.* **96**, 1218 (1992).
- [12] A. Köhn, M. Hanauer, L. A. Mück, T.-C. Jagau, and J. Gauss, *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **3**, 176 (2013).
- [13] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
- [14] A. J. Cohen, P. Mori-Sanchez, and W. Yang, *Science* **321**, 792 (2008).
- [15] A. J. Cohen, P. Mori-Sánchez, and W. Yang, *Chem. Rev.* **112**, 289 (2012).
- [16] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, *J. Chem. Phys.* **132**, 154104 (2010).
- [17] A. D. Becke and E. R. Johnson, *J. Chem. Phys.* **122**, 154104 (2005).

- [18] A. D. Becke and E. R. Johnson, *J. Chem. Phys.* **123**, 154101 (2005).
- [19] A. D. Becke and E. R. Johnson, *J. Chem. Phys.* **124**, 014104 (2006).
- [20] A. D. Becke and E. R. Johnson, *J. Chem. Phys.* **127**, 154108 (2007).
- [21] C. F. Richardson and N. W. Ashcroft, *Phys. Rev. B* **50**, 8170 (1994).
- [22] Y. Andersson, D. C. Langreth, and B. I. Lundqvist, *Phys. Rev. Lett.* **76**, 102 (1996).
- [23] M. Dion, H. Rydberg, E. Schröder, D. C. Langreth, and B. I. Lundqvist, *Phys. Rev. Lett.* **92**, 246401 (2004).
- [24] O. A. Vydrov and T. Van Voorhis, *J. Chem. Phys.* **133**, 244103 (2010).
- [25] K. Lee, É. D. Murray, L. Kong, B. I. Lundqvist, and D. C. Langreth, *Phys. Rev. B* **82**, 081101 (2010).
- [26] J. P. Perdew and A. Zunger, *Phys. Rev. B* **23**, 5048 (1981).
- [27] P. Mori-Sánchez, A. J. Cohen, and W. Yang, *J. Chem. Phys.* **124**, 091102 (2006).
- [28] J. P. Perdew, V. N. Staroverov, J. Tao, and G. E. Scuseria, *Phys. Rev. A* **78**, 052513 (2008).
- [29] A. D. Becke, *J. Chem. Phys.* **138**, 074109 (2013).
- [30] A. D. Becke, *J. Chem. Phys.* **138**, 161101 (2013).
- [31] A. D. Becke, *J. Chem. Phys.* **139**, 021104 (2013).
- [32] M. Fuchs, Y.-M. Niquet, X. Gonze, and K. Burke, *J. Chem. Phys.* **122**, 094116 (2005).
- [33] D. Bohm and D. Pines, *Phys. Rev.* **82**, 625 (1951).
- [34] F. Furche, *J. Chem. Phys.* **129**, 114105 (2008).
- [35] H. Eshuis, J. Yarkony, and F. Furche, *J. Chem. Phys.* **132**, 234114 (2010).
- [36] M. Kaltak, J. Klimeš, and G. Kresse, *J. Chem. Theory Comput.* **10**, 2498 (2014).
- [37] M. Kállay, *J. Chem. Phys.* **142**, 204105 (2015).
- [38] H. F. Schurkus and C. Ochsenfeld, *J. Chem. Phys.* **144**, 031101 (2016).
- [39] J. Wilhelm, P. Seewald, M. Del Ben, and J. Hutter, *J. Chem. Theory Comput.* **12**, 5851 (2016).
- [40] D. R. Hartree, *Math. Proc. Cambridge Philos. Soc.* **24**, 89 (1928).

- [41] D. R. Hartree, *Math. Proc. Cambridge Philos. Soc.* **24**, 111 (1928).
- [42] V. Fock, *Zeitschrift für Phys.* **61**, 126 (1930).
- [43] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry* (Dover, Mineola, New York, 1996).
- [44] R. G. Parr and W. Yang, *Density-Functional Theory of Atoms and Molecules* (Oxford University Press, New York, 1989).
- [45] E. Engel and R. M. Dreizler, *Density Functional Theory* (Springer, Berlin, 2011).
- [46] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).
- [47] C. C. J. Roothaan, *Rev. Mod. Phys.* **23**, 69 (1951).
- [48] G. G. Hall, *Proc. R. Soc. A Math. Phys. Eng. Sci.* **205**, 541 (1951).
- [49] R. McWeeny, *Rev. Mod. Phys.* **32**, 335 (1960).
- [50] J. Gauss, in *Mod. Methods Algorithms Quantum Chem.*, NIC Series, Vol. 3, edited by J. Grotendorst (John von Neumann Institute for Computing, Jülich, 2000) pp. 541–592.
- [51] H. Hellmann, *Einführung in die Quantenchemie* (Franz Deuticke, Leipzig, 1937).
- [52] R. P. Feynman, *Phys. Rev.* **56**, 340 (1939).
- [53] D. M. Bishop and M. Randić, *J. Chem. Phys.* **44**, 2480 (1966).
- [54] R. Moccia, *Chem. Phys. Lett.* **5**, 260 (1970).
- [55] J. Gerratt and I. M. Mills, *J. Chem. Phys.* **49**, 1719 (1968).
- [56] J. Gerratt and I. M. Mills, *J. Chem. Phys.* **49**, 1730 (1968).
- [57] P. Pulay, *Mol. Phys.* **17**, 197 (1969).
- [58] J. L. Dodds, R. McWeeny, and A. J. Sadlej, *Mol. Phys.* **34**, 1779 (1977).
- [59] J. A. Pople, R. Krishnan, H. B. Schlegel, and J. S. Binkley, *Int. J. Quantum Chem.* **16**, 225 (2009).
- [60] Y. Osamura, Y. Yamaguchi, P. Saxe, D. Fox, M. Vincent, and H. Schaefer, *J. Mol. Struct. THEOCHEM* **103**, 183 (1983).
- [61] M. Frisch, M. Head-Gordon, and J. Pople, *Chem. Phys.* **141**, 189 (1990).
- [62] M. Beer and C. Ochsenfeld, *J. Chem. Phys.* **128**, 221102 (2008).

- [63] J. Almlöf, Chem. Phys. Lett. **181**, 319 (1991).
- [64] M. Häser and J. Almlöf, J. Chem. Phys. **96**, 489 (1992).
- [65] M. Häser, Theor. Chim. Acta **87**, 147 (1993).
- [66] P. Y. Ayala and G. E. Scuseria, J. Chem. Phys. **110**, 3660 (1999).
- [67] M. Beer, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **134**, 074102 (2011).
- [68] J. Kussmann, A. Luenser, M. Beer, and C. Ochsenfeld, J. Chem. Phys. **142**, 094101 (2015).
- [69] A. Luenser, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **145**, 124103 (2016).
- [70] E. Baerends, D. Ellis, and P. Ros, Chem. Phys. **2**, 41 (1973).
- [71] J. L. Whitten, J. Chem. Phys. **58**, 4496 (1973).
- [72] B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, J. Chem. Phys. **71**, 4993 (1979).
- [73] B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, J. Chem. Phys. **71**, 3396 (1979).
- [74] O. Vahtras, J. Almlöf, and M. Feyereisen, Chem. Phys. Lett. **213**, 514 (1993).
- [75] B. Dunlap, J. Mol. Struct. THEOCHEM **529**, 37 (2000).
- [76] B. Dunlap, J. Mol. Struct. THEOCHEM **501-502**, 221 (2000).
- [77] B. I. Dunlap, Phys. Chem. Chem. Phys. **2**, 2113 (2000).
- [78] K. Eichkorn, O. Treutler, H. Öhm, M. Häser, and R. Ahlrichs, Chem. Phys. Lett. **240**, 283 (1995).
- [79] M. Sierka, A. Hogeckamp, and R. Ahlrichs, J. Chem. Phys. **118**, 9136 (2003).
- [80] F. Weigend, Phys. Chem. Chem. Phys. **4**, 4285 (2002).
- [81] M. Feyereisen, G. Fitzgerald, and A. Komornicki, Chem. Phys. Lett. **208**, 359 (1993).
- [82] F. Weigend and M. Häser, Theor. Chem. Accounts Theory, Comput. Model. (Theoretica Chim. Acta) **97**, 331 (1997).
- [83] C. Hättig and F. Weigend, J. Chem. Phys. **113**, 5154 (2000).
- [84] E. Epifanovsky, D. Zuev, X. Feng, K. Khistyayev, Y. Shao, and A. I. Krylov, J. Chem. Phys. **139**, 134105 (2013).
- [85] D. Pines and D. Bohm, Phys. Rev. **85**, 338 (1952).

- [86] D. Bohm and D. Pines, Phys. Rev. **92**, 609 (1953).
- [87] D. Langreth and J. Perdew, Solid State Commun. **17**, 1425 (1975).
- [88] O. Gunnarsson and B. I. Lundqvist, Phys. Rev. B **13**, 4274 (1976).
- [89] D. C. Langreth and J. P. Perdew, Phys. Rev. B **15**, 2884 (1977).
- [90] J. Harris, Phys. Rev. A **29**, 1648 (1984).
- [91] H. B. Callen and T. A. Welton, Phys. Rev. **83**, 34 (1951).
- [92] H. Eshuis, J. E. Bates, and F. Furche, Theor. Chem. Acc. **131**, 1084 (2012).
- [93] A. L. Fetter and J. D. Walecka, *Quantum Theory of Many-Particle Systems* (McGraw-Hill, New York, 1971).
- [94] J. Oddershede, P. Jørgensen, and D. L. Yeager, Comput. Phys. Reports **2**, 33 (1984).
- [95] Y. M. Niquet, M. Fuchs, and X. Gonze, Phys. Rev. A **68**, 032507 (2003).
- [96] F. Furche and T. Van Voorhis, J. Chem. Phys. **122**, 164106 (2005).
- [97] J. G. Ángyán, R.-F. Liu, J. Toulouse, and G. Jansen, J. Chem. Theory Comput. **7**, 3116 (2011).
- [98] M. Petersilka, U. J. Gossmann, and E. K. U. Gross, Phys. Rev. Lett. **76**, 1212 (1996).
- [99] A. Takatsuka, S. Ten-no, and W. Hackbusch, J. Chem. Phys. **129**, 044112 (2008).
- [100] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia, 2002).
- [101] S. Itoh, P. Ordejón, and R. M. Martin, Comput. Phys. Commun. **88**, 173 (1995).
- [102] A. Canning, G. Galli, F. Mauri, A. De Vita, and R. Car, Comput. Phys. Commun. **94**, 89 (1996).
- [103] M. Challacombe, J. Chem. Phys. **110**, 2332 (1999).
- [104] M. Challacombe, Comput. Phys. Commun. **128**, 93 (2000).
- [105] C. Saravanan, Y. Shao, R. Baer, P. N. Ross, and M. Head-Gordon, J. Comput. Chem. **24**, 618 (2003).
- [106] P. Löwdin, J. Chem. Phys. **18**, 365 (1950).
- [107] B. C. Carlson and J. M. Keller, Phys. Rev. **105**, 102 (1957).


- [108] L. Piela, in *Ideas Quantum Chem.*, Vol. 1 (Elsevier, Waltham, Massachusetts, 2014) 2nd ed., pp. e99–e103.
- [109] M. A. Watson, P. Salek, P. Macak, M. Jaszunski, and T. Helgaker, *Chem. - A Eur. J.* **10**, 4627 (2004).
- [110] S. A. Varganov, A. T. B. Gilbert, E. Deplazes, and P. M. W. Gill, *J. Chem. Phys.* **128**, 201104 (2008).
- [111] P. M. Gill and A. T. Gilbert, *Chem. Phys.* **356**, 86 (2009).
- [112] T. Limpanuparb and P. M. W. Gill, *Phys. Chem. Chem. Phys.* **11**, 9176 (2009).
- [113] T. Limpanuparb, A. T. B. Gilbert, and P. M. W. Gill, *J. Chem. Theory Comput.* **7**, 830 (2011).
- [114] T. Limpanuparb and P. M. W. Gill, *J. Chem. Theory Comput.* **7**, 2353 (2011).
- [115] T. Limpanuparb, J. W. Hollett, and P. M. W. Gill, *J. Chem. Phys.* **136**, 104102 (2012).
- [116] T. Limpanuparb, J. Milthorpe, A. P. Rendell, and P. M. W. Gill, *J. Chem. Theory Comput.* **9**, 863 (2013).
- [117] Y. Jung, A. Sodt, P. M. W. Gill, and M. Head-Gordon, *Proc. Natl. Acad. Sci.* **102**, 6692 (2005).
- [118] Y. Jung, Y. Shao, and M. Head-Gordon, *J. Comput. Chem.* **28**, 1953 (2007).
- [119] Y. Jung, R. C. Lochan, A. D. Dutoi, and M. Head-Gordon, *J. Chem. Phys.* **121**, 9793 (2004).
- [120] S. Reine, E. Tellgren, A. Krapp, T. Kjærgaard, T. Helgaker, B. Jansik, S. Høst, and P. Salek, *J. Chem. Phys.* **129**, 104101 (2008).
- [121] P. Merlot, T. Kjaergaard, T. Helgaker, R. Lindh, F. Aquilante, S. Reine, and T. B. Pedersen, *J. Comput. Chem.* **34**, 1486 (2013).

Publications

Article I

”Vanishing-Overhead Linear-Scaling Random Phase Approximation by Cholesky Decomposition and an Attenuated Coulomb-Metric”,
A. Luenser, H. F. Schurkus, and C. Ochsenfeld,
J. Chem. Theor. Comput. **13**, 1647 (2017).


Vanishing-Overhead Linear-Scaling Random Phase Approximation by Cholesky Decomposition and an Attenuated Coulomb-Metric

Arne Luenser, Henry F. Schurkus, and Christian Ochsenfeld*

Chair of Theoretical Chemistry, Department of Chemistry, University of Munich (LMU), Butenandtstrasse 5-13, D-81377 Munich, Germany

Center for Integrated Protein Science Munich (CIPSM), Butenandtstrasse 5-13, D-81377 Munich, Germany

ABSTRACT: A reformulation of the random phase approximation within the resolution-of-the-identity (RI) scheme is presented, that is competitive to canonical molecular orbital RI-RPA already for small- to medium-sized molecules. For electronically sparse systems drastic speedups due to the reduced scaling behavior compared to the molecular orbital formulation are demonstrated. Our reformulation is based on two ideas, which are independently useful: First, a Cholesky decomposition of density matrices that reduces the scaling with basis set size for a fixed-size molecule by one order, leading to massive performance improvements. Second, replacement of the overlap RI metric used in the original AO-RPA by an attenuated Coulomb metric. Accuracy is significantly improved compared to the overlap metric, while locality and sparsity of the integrals are retained, as is the effective linear scaling behavior.



1. INTRODUCTION

The random phase approximation (RPA) correlation method has for its excellent description of long-ranged correlation effects piqued the interest of many researchers in the field of quantum chemistry, especially since its advancement from a computationally prohibitively expensive method in the original $O(M^6)$ formulation^{1,2} to a method computationally tractable enough for use within rung five³ DFT schemes.⁴ Many authors have contributed to this algorithmic improvement: Furche and co-workers^{5–8} put forward an effective $O(M^4)$ scaling formulation. While their approach employs the resolution-of-the-identity (RI) approximation,⁹ Yang and co-workers¹⁰ suggested to employ the tensor hypercontraction (THC) developed by Martinez and co-workers^{11–13} instead. Kresse and co-workers¹⁴ were able to reduce the scaling to $O(M^3)$ for periodic systems. Kallay¹⁵ and also ourselves¹⁶ have shown how to further reduce the effective computational complexity to linear for molecules with nonvanishing band gaps. Most recently, Hutter and co-workers¹⁷ have made significant steps toward highly parallelized evaluation of these formulations. Extensive analysis of the results obtainable within these new formulations has since shown the many strengths of the RPA.^{5,18–22} Analytical gradients for the RPA energy are given in refs 23 and 24. Benchmark studies have shown the importance of evaluating the RPA with large basis sets of up to quadruple- ζ size.

Our previously presented atomic orbital (AO) formulation shows linear scaling with the molecule size M and thus extends the applicability of the RPA to molecules comprising thousands of atoms. To do so, we moved from the Coulomb RI metric,

which gives the most accurate results for the commonly employed auxiliary basis sets, to the potentially less accurate short-range overlap metric. Also, the scaling with basis set size N_{basis} and auxiliary basis set size N_{aux} for any fixed molecular size ($N_{\text{occ}} = \text{const}$), which is important for the prefactor of the evaluation time, increases from $O(N_{\text{aux}}^2 N_{\text{basis}} N_{\text{occ}})$ to $O(N_{\text{aux}}^2 N_{\text{basis}}^2)$, thereby limiting the applicability to basis sets much smaller than quadruple- ζ sizes.

The present work supplies remedies for these drawbacks by introducing two concepts from related fields to the context of RPA. First, by pivoted Cholesky decomposition^{25–27} of the density and pseudodensity matrices (CDD^{28,29}) within the AO formulation the scaling with basis set size is brought back down to $O(N_{\text{aux}}^2 N_{\text{basis}} N_{\text{occ}})$ (section 3.1). We attain a prefactor which allows—even in the case of small molecules—for evaluation times competitive with the canonical MO-based formulation as will be shown in section 4.1. Second, by exchanging the overlap metric RI for a formulation using the Coulomb metric attenuated by the complementary error function^{30–32} in section 3.2 we arrive at a formulation that can smoothly interpolate between the overlap metric and exact Coulomb metric result. We show in section 4.2 that choices for the attenuation parameter ω are available which combine the best of both worlds—rendering results numerically equivalent to the Coulomb metric results, but local enough to be evaluated in linear time.

Neither of these concepts change the complexity with molecular size, thus either one or both can be applied while still allowing

Received: December 21, 2016

Published: March 6, 2017

for linear scaling evaluation as will be verified in section 4.3. By combining the two, we finally arrive at a formulation we term ω -CDD-RPA, which is competitive with the canonical $O(M^4)$ formulation regardless of molecule and basis set size but has the additional advantage of $O(M)$ scaling, allowing for the application to extended systems.

2. REVIEW

2.1. RI-RPA in Molecular and Atomic Orbital Bases. In the following, we will use μ, ν for atomic orbitals (AOs), i, j for occupied molecular orbitals (MOs), a, b for virtual molecular orbitals, and M, N for auxiliary RI functions. The RI-RPA correlation energy in its standard Coulomb metric form is⁶

$$E_C^{\text{RPA}} = \frac{1}{2\pi} \int_0^\infty \text{Tr}(\ln[\mathbf{1} + \mathbf{Q}(u)] - \mathbf{Q}(u)) du \quad (1)$$

where

$$\mathbf{Q}_{MN}(u) = 2\mathbf{S}_M^T \mathbf{G}(u) \mathbf{S}_N \quad (2)$$

$$G_{ai,bj}(u) = \frac{(\epsilon_a - \epsilon_i)\delta_{ai,bj}}{(\epsilon_a - \epsilon_i)^2 + u^2} \quad (3)$$

$$S_{ai}^M = \sum_N (ai|N) L_{NM} \quad (4)$$

$$\mathbf{L}\mathbf{L}^T = \mathbf{C}^{-1} \quad (5)$$

$$C_{MN} = \left(M \left| \frac{1}{r_{12}} \right| N \right) \quad (6)$$

using the Mulliken notation for two- and three-center integrals. The frequency integration in eq 1 is performed by Clenshaw-Curtis quadrature.⁶ Computational complexity is $O(N_{\text{aux}}^2 N_{\text{occ}} N_{\text{virt}})$ per frequency point u , where N_{aux} is the number of auxiliary functions, and N_{occ} and N_{virt} are the numbers of occupied and virtual MOs, respectively.

The reformulation in the atomic orbital basis presented by two of us, Schurkus and Ochsenfeld,¹⁶ used the integral transform

$$\frac{1}{(\epsilon_a - \epsilon_i)^2 + u^2} = \int_0^\infty \frac{\sin(ut)}{u} e^{-(\epsilon_a - \epsilon_i)t} dt \quad (7)$$

which together with the overlap RI metric

$$(ailbj) \approx \sum_{MN} B_{ai}^M \tilde{C}_{MN} B_{bj}^N \quad (8)$$

$$\tilde{\mathbf{C}} = \mathbf{S}^{-1} \mathbf{C} \mathbf{S}^{-1} \quad (9)$$

$$B_{\mu\nu}^M = (\mu\nu|M) = \int \chi_\mu(\mathbf{r}) \chi_\nu(\mathbf{r}) \chi_M(\mathbf{r}) d\mathbf{r} \quad (10)$$

and

$$\mathbf{Q}(u) = \begin{cases} u < u^*: & 2 \sum_{\alpha=1}^{\tau} w_\alpha \cos(ut_\alpha) \mathbf{F}_{\text{INT}}^{(t_\alpha)} \\ u > u^*: & \frac{2}{u^2} \mathbf{F}_0 + 2 \sum_{\alpha=1}^{\tau} w_\alpha \frac{\cos(ut_\alpha)}{u^2} \mathbf{F}_{\text{D}}^{(t_\alpha)} \end{cases} \quad (11)$$

enabled effective linear scaling. u^* specifies the frequency where we switch from one formula to the other, which was chosen as

0.5 au. w_α and t_α are the weights and roots for the quadrature of the integral in eq 7. Equation 1 is modified to

$$E_C^{\text{AO-RPA}} = \frac{1}{2\pi} \int_0^\infty \text{Tr}(\ln[\mathbf{1} + \mathbf{Q}(u)\tilde{\mathbf{C}}] - \mathbf{Q}(u)\tilde{\mathbf{C}}) du \quad (12)$$

that is, multiplication with the long-range $\frac{1}{r_{12}}$ -operator matrix $\tilde{\mathbf{C}}$ is deferred until the last step of the algorithm. All other matrices are local and sparse for large molecules. The three types of \mathbf{F} matrices are given by

$$(\mathbf{F}_0)_{MN} = \text{Tr}(\mathbf{P}\mathbf{F}\mathbf{P}_M \bar{\mathbf{P}}_N) - \text{Tr}(\mathbf{P}\mathbf{B}_M \bar{\mathbf{P}}\mathbf{F}\bar{\mathbf{P}}_N) \quad (13)$$

$$(\mathbf{F}_{\text{INT}}^{(t_\alpha)})_{MN} = \text{Tr}(\underline{\mathbf{P}}^{(\alpha)} \mathbf{B}_M \bar{\mathbf{P}}^{(\alpha)} \mathbf{B}_N) \quad (14)$$

$$(\mathbf{F}_{\text{D}}^{(t_\alpha)})_{MN} = 2\text{Tr}(\underline{\underline{\mathbf{P}}}^{(\alpha)} \mathbf{B}_M \bar{\bar{\mathbf{P}}}^{(\alpha)} \mathbf{B}_N) - \text{Tr}(\underline{\underline{\mathbf{P}}}^{(\alpha)} \mathbf{B}_M \bar{\mathbf{P}}^{(\alpha)} \mathbf{B}_N) - \text{Tr}(\underline{\mathbf{P}}^{(\alpha)} \mathbf{B}_M \bar{\bar{\mathbf{P}}}^{(\alpha)} \mathbf{B}_N) \quad (15)$$

\mathbf{P} , $\bar{\mathbf{P}}$, and \mathbf{F} are occupied and virtual one-particle density matrices, and the Kohn–Sham matrix, respectively. $\underline{\mathbf{P}}^{(\alpha)}$ and $\bar{\mathbf{P}}^{(\alpha)}$, as well as their first ($\underline{\underline{\mathbf{P}}}^{(\alpha)}$, $\bar{\bar{\mathbf{P}}}^{(\alpha)}$) and second ($\underline{\underline{\underline{\mathbf{P}}}}^{(\alpha)}$, $\bar{\bar{\bar{\mathbf{P}}}}^{(\alpha)}$) derivatives with respect to t_α are the usual Laplace pseudodensity matrices defined in ref 16.

2.2. Notes on the Frequency Quadrature. The trace of the matrix logarithm in eq 12 can be more efficiently evaluated by Cholesky-factorizing $\tilde{\mathbf{C}} = \mathbf{L}\mathbf{L}^T$, and using the Mercator series for $\ln(1+x)$ to rewrite

$$\text{Tr} \ln(\mathbf{1} + \mathbf{Q}(u)\tilde{\mathbf{C}}) = \text{Tr} \ln(\mathbf{1} + \mathbf{L}^T \mathbf{Q}(u) \mathbf{L}) \quad (16)$$

$$= \ln(\det \mathbf{L}' \det \mathbf{L}'^T) \quad (17)$$

$$= 2 \ln \prod_i L'_{ii} \quad (18)$$

where $\mathbf{1} + \mathbf{L}^T \mathbf{Q}(u) \mathbf{L}$ itself was Cholesky-factorized as $\mathbf{L}'\mathbf{L}'^T$.¹⁴ This strategy is about $10 \times$ faster than diagonalization of $\mathbf{1} + \mathbf{Q}(u)\tilde{\mathbf{C}}$ (which is not symmetric). The point at which the $O(N_{\text{aux}}^3)$ evaluation of eq 16 becomes the time-determining step is thus pushed further outward.

An important and possibly under-appreciated advantage of computing $\mathbf{Q}(u)$ by eq 11 (AO¹⁶) versus eq 2 (MO⁶) lies in the decoupling of the u -integration from the time-determining step. In eq 2, $\mathbf{G}(u)$ must be recalculated from scratch for each u -frequency point and contracted with three-center integrals S_{ai}^M from left and right, which is the dominant step at $O(N_u N_{\text{aux}}^2 N_{\text{occ}} N_{\text{virt}})$, where N_u is the number of u -quadrature nodes. In other words, the computational time required for conventional RI-RPA is directly proportional to the number of u -quadrature nodes. In contrast, the integral transform in eq 7 enables us to precompute the matrices \mathbf{F}_0 , $\mathbf{F}_{\text{INT}}^{(t_\alpha)}$, and $\mathbf{F}_{\text{D}}^{(t_\alpha)}$ at a formal cost of $O(\tau N_{\text{aux}}^2 N_{\text{basis}}^2)$, where τ is the number of Laplace quadrature nodes, because they do not depend on u . $\mathbf{Q}(u)$ is then constructed via eq 11 at a negligible cost of $O(\tau N_{\text{aux}}^2)$ per u -quadrature node, making it possible to employ hundreds of quadrature nodes essentially for free. When the electronic structure is sparse and a local RI metric is used, eq 11 (AO) asymptotically scales linearly, whereas eq 2 (MO) always scales quartically.

3. THEORY AND RESULTS

3.1. Complexity Reduction through Cholesky Factorization. In this work, we report dramatic performance improvements to the RI-RPA energy evaluation by pivoted Cholesky factorization. Formal complexity is reduced from $O(N_{\text{aux}}^2 N_{\text{basis}}^2)$ to $O(N_{\text{aux}}^2 N_{\text{basis}} N_{\text{occ}})$, resulting in drastic runtime reductions especially for the large basis sets typically used in RPA calculations.

We emphasize that we use the term “pivoted Cholesky factorization”²⁶ in the following sense: a positive semidefinite matrix \mathbf{A} may be decomposed as $\mathbf{P}^{-1}\mathbf{A}\mathbf{P} = \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is lower triangular, and \mathbf{P} is a permutation matrix, or equivalently $\mathbf{A} = \mathbf{P}\mathbf{L}\mathbf{L}^T\mathbf{P}^{-1}$. A permutation matrix \mathbf{P} is a unitary matrix which permutes the rows of a matrix \mathbf{A} when premultiplied ($\mathbf{P}\mathbf{A}$), or the columns when postmultiplied ($\mathbf{A}\mathbf{P}$). $\mathbf{P}\mathbf{L}$ has dimension $\dim \mathbf{A} \times \text{rank } \mathbf{A}$ and is *not* lower triangular. After resorting the columns of $\mathbf{P}\mathbf{L}$ by another permutation matrix \mathbf{U} to enhance useable sparsity (see ref 33 for details on this reordering procedure), the complete factorization is $\mathbf{A} = \mathbf{P}\mathbf{L}\mathbf{U}\mathbf{U}^{-1}\mathbf{L}^T\mathbf{P}^{-1}$, which we abbreviate as “ $\mathbf{A} = \mathbf{L}\mathbf{L}^T$ ”.

We proceed by performing pivoted Cholesky decompositions of all occupied-type density matrices (\mathbf{P} , $\mathbf{P}\mathbf{F}\mathbf{P}$, $\underline{\mathbf{P}}^{(\alpha)}$, $\underline{\mathbf{P}}^{(\alpha)}$, and $\underline{\mathbf{P}}^{(\alpha)}$) in eqs 13–15. This reveals the matrices’ ranks, which is N_{occ} or smaller. We then reorder the columns of each left Cholesky factor for sparsity, as described in ref 33. The resulting Cholesky factor is a transformation matrix to a local pseudo-MO basis.^{28,29} We transform one AO index of the three-center integrals \mathbf{B}_M to this local basis, which reduces the dimension of each \mathbf{B}_M from $N_{\text{basis}} \times N_{\text{basis}}$ to $N_{\text{basis}} \times N_{\text{occ}}$ while at the same time preserving any sparsity of the original matrix (see remarks below). The resulting expressions for \mathbf{F}_0 , $\mathbf{F}_{\text{INT}}^{(t_a)}$, and $\mathbf{F}_D^{(t_a)}$ are

$$(\mathbf{F}_0)_{MN} = -\text{Tr}(\mathbf{V}_M^T \mathbf{P} \mathbf{V}_N) - \text{Tr}(\mathbf{W}_M^T \mathbf{P} \mathbf{F} \mathbf{P} \mathbf{W}_N) \quad (19)$$

$$(\mathbf{F}_{\text{INT}}^{(t_a)})_{MN} = \text{Tr}(\mathbf{Z}_M^{(\alpha)T} \underline{\mathbf{P}}^{(\alpha)} \mathbf{Z}_N^{(\alpha)}) \quad (20)$$

$$(\mathbf{F}_D^{(t_a)})_{MN} = -2\text{Tr}(\mathbf{X}_M^{(\alpha)T} \underline{\mathbf{P}}^{(\alpha)} \mathbf{X}_N^{(\alpha)}) - \text{Tr}(\mathbf{Y}_M^{(\alpha)T} \underline{\mathbf{P}}^{(\alpha)} \mathbf{Y}_N^{(\alpha)}) - \text{Tr}(\mathbf{Z}_M^{(\alpha)T} \underline{\mathbf{P}}^{(\alpha)} \mathbf{Z}_N^{(\alpha)}) \quad (21)$$

where the occupied density matrices have been absorbed:

$$\mathbf{V}_M = \mathbf{B}_M \mathbf{L}' \text{ with } -\mathbf{P}\mathbf{F}\mathbf{P} = \mathbf{L}'\mathbf{L}'^T \quad (22)$$

$$\mathbf{W}_M = \mathbf{B}_M \mathbf{L} \text{ with } \mathbf{P} = \mathbf{L}\mathbf{L}^T \quad (23)$$

$$\mathbf{X}_M^{(\alpha)} = \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ with } -\underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)T} \quad (24)$$

$$\mathbf{Y}_M^{(\alpha)} = \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ with } \underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)T} \quad (25)$$

$$\mathbf{Z}_M^{(\alpha)} = \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ with } \underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)T} \quad (26)$$

Note that $\mathbf{P}\mathbf{F}\mathbf{P}$ and $\underline{\mathbf{P}}^{(\alpha)}$ are negative semidefinite and are therefore multiplied by -1 before factorization, which leads to sign changes in eqs 19 and 21 compared to eqs 13 and 15. Shifting the MO energies downward uniformly by $\frac{1}{2}(\epsilon_{\text{HOMO}} + \epsilon_{\text{LUMO}})$ (as described in ref 34; HOMO is highest occupied MO and LUMO is the lowest unoccupied MO) before constructing the pseudodensity matrices gives the numerically most stable results.

The question of whether to also factorize the virtual pseudodensities arises naturally. In the context of coupled-perturbed self-consistent field calculations, factorization of the virtual

subspace has performance benefits,³⁵ mostly due to numerical rank deficiency of some pseudodensities far beyond their analytical rank of N_{virt} . In AO-RPA, doing the same results in simpler and more symmetric equations, which are given in Appendix A. When it comes to performance, however, we have found this approach disappointing. The reason is twofold. First, correlation methods such as RPA require large bases for good results, in which case $N_{\text{basis}} \approx N_{\text{virt}}$ and the rank reduction from factorization is negligible. Second, the Cholesky factorization of the virtual subspace preserves sparsity less well than in the occupied subspace. Consequently, the rank reduction is often counteracted by loss of sparsity, and performance degrades. For small systems, when no sparse algebra is used, however, we do use the equations given in Appendix A.

Finally, we point out another surprising pitfall. The pivoted Cholesky decomposition will sometimes yield pseudo-MOs of which a small number are not local, but extend over the whole AO space, which thwarts attempts to benefit from blocked sparse algebra. We have found this problem can be solved by orthogonalizing the density matrices in the contravariant AO basis before Cholesky factorization, and reverting the orthogonalization afterward. However, the orthogonalization cannot be Löwdin’s symmetric orthogonalization, but must be some sparsity-preserving but nonsymmetric orthogonalization. We use the Cholesky factors of the AO overlap matrix.

3.2. RI-RPA Using the erfc-Attenuated Coulomb Metric.

Key to reduce the computational effort and scaling in extended molecules is the choice of the RI metric. As mentioned above, ref 16 uses the overlap metric³⁶ (eq 10), which is local in the sense that auxiliary functions M do not overlap with AO basis function pairs $\mu\nu$ if there is enough distance between their centers, which leads to sparsity in the matrix representations. In contrast, the Coulomb metric (eq 4) couples auxiliary functions and AO basis function pairs over effectively infinite distances ($\frac{1}{r_{12}}$ decay), and no sparsity can be gained.

The Coulomb metric has been shown to be optimal in fitting density-like repulsions,³⁷ and has subsequently been the ubiquitous choice. This metric has no disadvantage when transforming to the canonical MO basis where all sparsity is lost anyway. In local bases, such as the atomic orbitals and Cholesky pseudo-MOs (see section 3.1), local RI metrics have important advantages.

The overlap metric is very local, because it decays as $\exp(-r_{12}^2)$ in Gaussian basis sets. The downside is decreased accuracy.³⁷ Other metrics have been described in the literature, especially Coulomb metrics attenuated by a complementary error function (erfc)^{30,31} or by a Gaussian function.³² For correlation energy calculations, the erfc-attenuated Coulomb metric was first used by Jung et al.³¹ for scaled-opposite-spin MP2 calculations.

Here, we employ the erfc-attenuated Coulomb metric within RI-RPA for the first time, and show that it combines the advantages of the Coulomb and overlap metrics.

The two- and three-center integrals necessary are given by

$$(M|\mu\nu)_\omega = \left(M \left| \frac{\text{erfc}(\omega r_{12})}{r_{12}} \right| \mu\nu \right) \quad (27)$$

$$(\mathbf{S}_\omega)_{MN} = \left(M \left| \frac{\text{erfc}(\omega r_{12})}{r_{12}} \right| N \right) \quad (28)$$

$$\tilde{\mathbf{C}}_\omega = \mathbf{S}_\omega^{-1} \mathbf{C} \mathbf{S}_\omega^{-1} \quad (29)$$

such that the resolution-of-the-identity is formally

$$(ailbj) = \sum_{MN} (ailM)_\omega (\tilde{C}_\omega)_{MN} (Nl bj)_\omega \quad (30)$$

Adamson et al.³⁸ have described how to efficiently calculate these integrals. \tilde{C}_ω is most reliably computed as $(S_\omega C^{-1} S_\omega)^{-1}$. As with the overlap metric, deferring the multiplication with the two-center integrals to the last step of the algorithm (eq 12) is necessary to retain locality through the time-determining steps. The attenuation strength is controlled by the parameter ω . erfc-attenuation has the pleasant property of allowing continuous variation between the $\frac{1}{r_{12}}$ and overlap operators

$$\lim_{\omega \rightarrow 0} \frac{\text{erfc}(\omega r_{12})}{r_{12}} = \frac{1}{r_{12}} \quad (31)$$

$$\lim_{\omega \rightarrow \infty} \frac{\text{erfc}(\omega r_{12})}{r_{12}} = \delta(r_{12}) \quad (32)$$

where δ is Dirac's function.

Figure 1 shows the effect of different values for ω over several orders of magnitude. Plotted are the RPA correlation energy




Figure 1. Effect of the attenuation parameter ω on the accuracy and performance of ω -CDD-RPA in the erfc-Coulomb metric for linear n -C₈₀H₁₆₂ (def2-SVP). The dashed horizontal lines are reference energy values calculated with the Coulomb ($\omega = 0$) and overlap ($\omega = \infty$) RI metrics. For a wide range $\omega \in [0, 0.1]$, the erfc-Coulomb metric reproduces the (optimal) Coulomb metric results to sub-mHartree accuracy in absolute energies at a fraction of the cost. Our recommended value of $\omega = 0.1$ is indicated by the dashed vertical line. Divergence (presumably because of limited floating point precision) is observed after $\omega \gtrsim 500$ (shaded graph area).

(top) and the wall time required (bottom, 12 threads on an Intel Xeon E5-2620 machine) as a function of ω (logarithmic scale). The limits of $\omega \rightarrow 0$ and $\omega \rightarrow \infty$ are correctly recovered. The overlap metric as well as high- ω attenuation violate the variational upper bound property⁶ of the Coulomb metric. For relative energies, this bound does not hold for the Coulomb metric either, though. For $\omega \gtrsim 500$, our implementation begins to diverge, presumably due to the limits of double precision floating point arithmetic. For ω -values between approximately 0.5 and 1.1, the erfc-attenuated absolute energies are lower than for the Coulomb metric. The computational effort

rises dramatically as the attenuation approaches zero, exceeding 28 h, and bottoms out at around 30 min for maximum attenuation. The most interesting region is between $\omega \simeq 0.1$ (where the Coulomb metric result is recovered to within 100 μ Hartree in the absolute energy) and $\omega \simeq 1$ (where the correlation energy starts to move toward the overlap metric result). The wall times required for calculations in this ω -interval are around 2 h, but yield *absolute* energies within mHartree of the Coulomb metric result, which takes over 10 \times as long to compute. We will present more data in sections 4.1 and 4.2, including from calculations with larger basis sets, which corroborate these findings. From the combined results, we recommend $\omega = 0.1$ as a starting point when using the erfc-attenuated Coulomb metric in RI-RPA.

4. CALCULATIONS

We used the Perdew–Burke–Ernzerhof (PBE) functional^{39,40} to obtain Kohn–Sham orbitals, using def2-SVP, def2-TZVP, and def2-QZVP basis sets.^{41,42} The RI approximation uses the corresponding RI counterparts.^{43,44} We use RI only for $(ailbj)$ -type integrals in the correlation part of the RPA energy, not for the Hamiltonian expectation value with the Kohn–Sham orbitals or during the preceding SCF calculation. All calculations use 100 quadrature nodes for the frequency integration, and a fixed integration interval of $u \in [0; 400]$ a.u., using the coordinate mapping described in ref 6. We emphasize that 100 quadrature nodes is more than typically used in RI-RPA calculations. We chose a generous number of nodes and a large integration interval to eliminate any errors from the quadrature, and used that same number of quadrature nodes for all calculations in this work in order to facilitate easy comparisons. We stress that fewer quadrature nodes may be sufficient (depending on the system under study), which would improve the efficiency of the MO-based formulation.

Where applicable, 15 Laplace quadrature points using the weights and nodes of ref 45. were employed. Core orbitals were frozen in all calculations. None of the results presented include corrections for possible basis set superposition errors. While this may render the statistical results slightly worse, the error is expected to be systematic across all methods presented. Some caution is advised when comparing the results below to counterpoise-corrected results.

Our ω -CDD-RPA method as well as the MO (ref 6) and AO (ref 16) formulations of RI-RPA, as well as RI-free RPA were implemented in the FermiONs++ program.^{46,47} We checked our MO-RI implementation against the canonical implementation in Turbomole7.0⁴⁸ to verify correctness and comparable performance. All runtimes given are wall times, not CPU times.

4.1. Performance: S66 Set. The full S66 set⁴⁹ of small-molecule interaction energies (mean reference interaction energy -5.5 kcal mol⁻¹, maximum reference interaction energy -19.8 kcal mol⁻¹) in double-, triple-, and quadruple- ζ bases serves as a performance and accuracy benchmark for small molecules. All calculations were performed using 12 threads on a dual-processor Intel Xeon E5–2620 machine. The small size of the molecules in the test set means no sparsity in the RI integrals or density matrices can be expected and no computational efficiency can be gained from a local metric, that is, all metrics yield identical performance.

The following results, therefore, measure (1) the RI-independent performance of the AO, ω -CDD, and MO formulations of RI-RPA for small systems without electronic sparsity and (2) their accuracy using different RI metrics, also including RI-free RPA results (computed via eq 9 in ref 6). The accuracy of all methods and

Table 1. S66: Root Mean Square (RMSD), Mean Absolute (MAD), and Maximum Absolute (MAX) Deviations from the Updated CCSD(T)/CBS Reference Interaction Energies⁵⁰ in kcal mol⁻¹ for Coulomb-, erfc-Attenuated Coulomb- ($\omega = 0.1$, $\omega = 0.5$), and Overlap-RI Metric RPA in the AO, ω -CDD, and MO Formulations, as well as for RI-Free RPA

basis	RI metric	RMSD/kcal mol ⁻¹			MAD/kcal mol ⁻¹			MAX/kcal mol ⁻¹		
		AO	ω -CDD	MO	AO	ω -CDD	MO	AO	ω -CDD	MO
def2-SVP	no RI			0.95			0.75			3.43
	Coulomb	0.98	0.98	0.97	0.77	0.77	0.77	3.45	3.45	3.46
	$\omega = 0.1$	0.98	0.98	0.97	0.77	0.77	0.77	3.44	3.44	3.46
	$\omega = 0.5$	1.00	1.00	1.00	0.81	0.81	0.81	3.44	3.44	3.45
	overlap	1.04	1.04	1.04	0.87	0.87	0.87	3.57	3.57	3.58
def2-TZVP	no RI			0.88			0.56			3.83
	Coulomb	0.86	0.86	0.88	0.56	0.56	0.56	3.81	3.81	3.83
	$\omega = 0.1$	0.87	0.87	0.88	0.56	0.56	0.56	3.81	3.81	3.83
	$\omega = 0.5$	0.90	0.90	0.91	0.59	0.59	0.59	3.88	3.88	3.89
	overlap	0.95	0.95	0.96	0.64	0.64	0.65	3.99	3.99	4.01
def2-QZVP	no RI			0.62			0.47			3.11
	Coulomb	0.62	0.62	0.62	0.47	0.47	0.46	3.11	3.11	3.11
	$\omega = 0.1$	0.62	0.62	0.62	0.47	0.47	0.46	3.11	3.11	3.11
	$\omega = 0.5$	0.63	0.63	0.63	0.48	0.48	0.47	3.11	3.11	3.11
	overlap	0.64	0.64	0.64	0.51	0.51	0.50	3.11	3.11	3.11

RI metrics for the S66 set can be summarized briefly as being very similar. Detailed values for the root mean square (RMSD), mean absolute (MAD), and maximum absolute (MAX) deviations from the updated⁵⁰ CCSD(T)/CBS reference interaction energies can be found in Table 1. We note the Laplace transform of eq 7 and its quadrature introduces no significant error in the S66 set. Likewise, the Cholesky decompositions of section 3.1 do not introduce any error over the AO formulation.

Regarding the RI metric, the Coulomb metric is generally best. However, the moderately erfc-attenuated ($\omega = 0.1$) Coulomb metric gives practically identical results. Increasing the attenuation to $\omega = 0.5$ leads to slightly inferior results, which are, however, magnitudes below the error inherent to RPA. Overlap metric results are slightly worse, but still an order of magnitude below the RPA error itself. Increasing the attenuation systematically increases the error for a given RI basis set, but only very slightly. Generally, the differences between RI metrics decrease for larger AO bases (and correspondingly larger RI bases). For quadruple- ζ bases, the differences between different RI metrics almost vanish for the S66 test set.

We now move to performance characteristics. Cumulative wall times for the complete S66 set ($3 \times 66 = 198$ calculations) are given in Table 2 for AO-RPA and ω -CDD-RPA. Exploiting the

Table 2. S66: Accumulated Wall Times and Speedups for AO- and ω -CDD-RPA Correlation Energy Calculations ($\omega = 0.1$)

basis	time (s)		speedup
	AO	ω -CDD	ω -CDD v AO
def2-SVP	2453	373	7×
def2-TZVP	16408	1476	11×
def2-QZVP	432405	15879	27×

rank deficiency of occupied (pseudo-) density matrices leads to large performance improvements of ω -CDD. The speedups are larger for larger bases, reaching 27× for a quadruple- ζ basis. As was discussed in section 3.1, ω -CDD reduces the computational complexity with respect to basis set size for a fixed-size molecule from N^4 to N^3 , the same as MO-RI-RPA, while the linear scaling of AO-RPA with molecular size remains unchanged.

ω -CDD-RPA is also competitive with MO-RI-RPA, slightly outperforming it for double- and triple- ζ bases, and showing equal performance for a quadruple- ζ basis. ω -CDD-RPA has extremely low overhead and the same scaling behavior as MO-RI-RPA in the limit of dense matrices (i.e., no usable matrix sparsity, such as in the S66 set). We emphasize that runtimes of ω -CDD-RPA are proportional to the number τ of Laplace quadrature points, while runtimes of MO-RI-RPA are proportional to the number N_u of u -frequency quadrature points. Theoretical parity is reached when $4\tau + 2 = N_u$. Which of the two formulations will outperform the other depends on which quadrature can be carried out with fewer points.

4.2. Accuracy: L7 Set. The L7 benchmark set⁵¹ comprises dispersion-dominated molecular systems of much larger size (up to 112 atoms, mean reference interaction energy -18.2 kcal mol⁻¹, maximum reference interaction energy -31.3 kcal mol⁻¹) than those in the S66 set. Using the def2-TZVP basis, we calculated RMSD, MAD, and MAX deviations from the QCISD(T)/CBS reference interaction energies⁵¹ for different attenuation parameters ω for both the MO and ω -CDD formulations of RI-RPA. The results are shown in Figure 2.

ω -CDD-RPA yields lower errors throughout, which is clearly fortuitous. The more significant result is between the different attenuation parameters ω . Unlike the smaller molecules of the S66 test set, the larger molecules of the L7 set are sensitive to the choice of the RI metric. For all three measures (RMSD, MAD, MAX), the unattenuated Coulomb metric consistently yields the best results. Similarly to the results of section 4.1, lower attenuation yields superior results. However, a moderate attenuation of $\omega = 0.1$ degrades the accuracy by only 0.02 kcal mol⁻¹ in RMSD and MAD, and 0.05 kcal mol⁻¹ in MAX, which we consider numerical noise. Increasing the attenuation to $\omega = 0.2$ or $\omega = 0.5$ increases the error measures very slightly. The overlap metric ($\omega = \infty$), however, yields an RMSD almost one kcal mol⁻¹ higher than the Coulomb metric, an MAD which is 0.6 kcal mol⁻¹ higher, and a maximum absolute error 2.5 kcal mol⁻¹ higher. Combined with the results presented in sections 3.2 and 4.1, this strongly supports the erfc-attenuated Coulomb metric in place of the overlap metric in RI-RPA calculations.

Finally, we note that while the systems comprising this benchmark set are considerably larger than those of the S66 set




Figure 2. L7: Root mean square (RMSD), mean absolute (MAD), and maximum absolute (MAX) deviations in kcal mol⁻¹ from the QCISD(T)/CBS reference interaction energies⁵¹ for different attenuation parameters ω (def2-TZVP). Left column: Canonical MO-RI-RPA. Right column: ω -CDD-RPA.

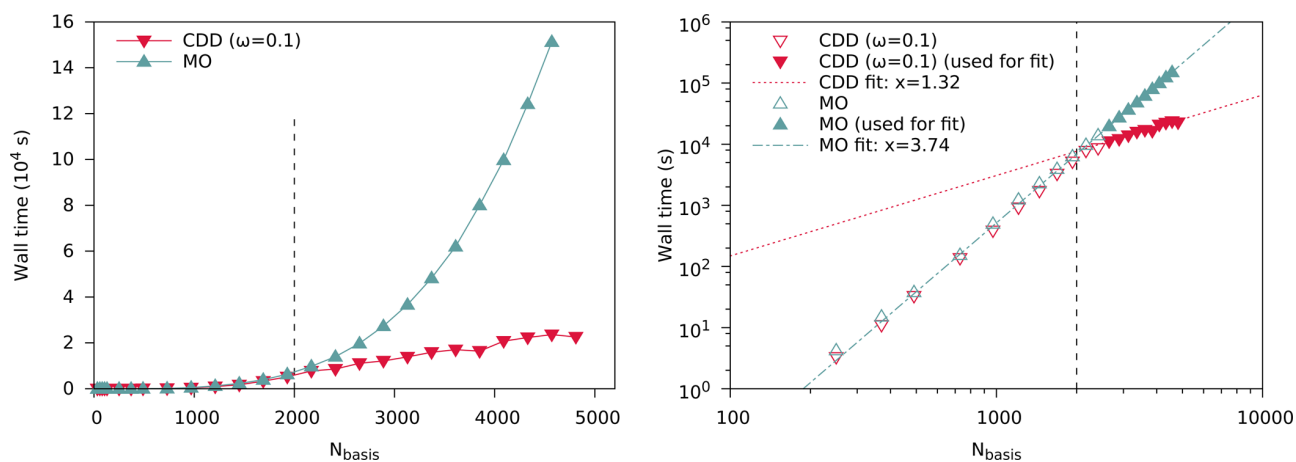


Figure 3. Plots of wall time against number of basis functions showing the computational complexity of the MO and ω -CDD formulations for linear n -alkanes (def2-SVP, erfc-attenuated Coulomb metric, $\omega = 0.1$). Left: Linear plot. Right: Log–log plot, with linear fits for MO and ω -CDD. The vertical dashed line indicates the point after which sparse algebra is used.

discussed in section 4.1, they are not as large as to provide usable matrix sparsity. Consequently, dense matrix algebra is used for all calculations, and performance is similar for the ω -CDD and MO variants of RI-RPA.

4.3. Scaling Behavior. Utilizing the rank deficiency of density matrices (cf. section 3.1) reduces the computational pre-factor of AO-RPA, as well as its formal scaling from $O(N_{\text{aux}}^2 N_{\text{basis}}^2)$

to $O(N_{\text{aux}}^2 N_{\text{basis}} N_{\text{occ}})$ in the limit of densely populated matrices. In the limit of large molecules with sparse electronic structures, however, we must ensure that the Cholesky factorization does not destroy usable matrix sparsity. Similarly, the introduction of an erfc-attenuated Coulomb metric instead of the overlap metric has to preserve the asymptotic $O(N)$ scaling of AO-RPA. In the following, we focus on these two aspects. Calculations in this

section used 16 threads on a dual-processor Intel Xeon E5–2667 machine. We switched the ω -CDD formulation from dense to sparse algebra at 2000 basis functions. Ideally, this would be automatically decided based on usable sparsity in the density matrix and three-center integrals (eq 27). For the purpose of this work, we switched to sparse algebra as soon as it was faster, which for both test sets in sections 4.3.1 and 4.3.2 happened to be at 2000 basis functions. Structure files for the test systems can be found online.⁵²

4.3.1. Alkanes. Our first test set are linear n -alkanes of increasing length. We calculated the RI-RPA correlation energy with canonical MO-RPA and ω -CDD-RPA ($\omega = 0.1$), which includes construction of $\mathbf{Q}(u)$ and evaluation of the matrix logarithm via eq 16. The data is plotted in Figure 3 in linear and log–log plots. From the linear fit in the log–log plots, MO-RPA scales as $O(N^{3.7})$, which is slightly better than the theoretical $O(N^4)$. Our own ω -CDD-RPA scales as $O(N^{1.3})$ for large systems.

Selected data points are given in Table 3. ω -CDD-RPA gives speedups over MO-RI-RPA of around 20–40% for systems with around 1000 basis functions. As discussed in section 4.1, whether or not a speedup can be attained for these system sizes depends critically on the number of u -frequency quadrature nodes used.

For large molecular sizes, however, the scaling behavior of ω -CDD will always give large speedups over MO-RPA. Table 3

Table 3. Wall Times and Computational Complexity of the ω -CDD and MO Formulations for Linear n -Alkanes (def2-SVP, erfc-Attenuated Coulomb Metric, $\omega = 0.1$)^a

	N_{basis}	ω -CDD		MO		total speedup
		time (s)	$O(N^x)$	time (s)	$O(N^x)$	
$\text{C}_{20}\text{H}_{42}$	490	33		38		1.2×
$\text{C}_{50}\text{H}_{102}$	1210	919	3.7	1244	3.9	1.4×
$\text{C}_{100}\text{H}_{202}$	2410	8585	3.2	13 995	3.5	1.6×
$\text{C}_{120}\text{H}_{242}$	2890	12100	1.9	27 359	3.7	2.3×
$\text{C}_{150}\text{H}_{302}$	3610	17022	1.5	62 025	3.7	3.6×
$\text{C}_{200}\text{H}_{402}$	4810	22593	1.0	*182 875	(3.7)	*8.1×

^aThe column headed $O(N^x)$ contains the scaling exponents relative to the row above. Values marked with an asterisk (*) are extrapolated conservatively.

shows the scaling exponents for ω -CDD and MO for progressively larger molecules. While ω -CDD-RPA reaches perfect $O(N)$ scaling for the largest system, MO-RPA scales as $O(N^{3.7})$. ω -CDD outperforms MO 8-fold for the largest system.

4.3.2. Glycine Chains. As a second test set for the scaling behavior, we used glycine chains of increasing length. The corresponding linear and log–log plots are given in Figure 4. Similarly to the results of section 4.3.1, linear fits in the log–log plots reveal $O(N^{3.6})$ scaling for MO, and $O(N^{1.4})$ scaling for ω -CDD. These timing data contain both the construction of $\mathbf{Q}(u)$ via eq 2 (MO) or eqs 11 and 33–35 (ω -CDD), as well as evaluation of the matrix logarithm (eq 16).

Selected data points are given in Table 4, where we separately resolved the construction of $\mathbf{Q}(u)$ and evaluation of the matrix logarithm (eq 16). The complexity of the time-determining steps in ω -CDD-RPA (eqs 33–35) reaches $O(N^{1.3})$ very quickly, whereas the time-determining step in MO-RPA scales as $O(N^4)$. Evaluation of the matrix logarithm is identical for both formulations, and predictably has $O(N^3)$ complexity with a small prefactor. For the largest system, the time required for this step is less than 8% of the total time for ω -CDD-RPA, and will become dominant only for extremely large molecules. The total speedup (including matrix logarithm evaluation) of ω -CDD over MO steadily rises from 1.3× up to over 13× for the largest test system.

5. SUMMARY

Cholesky decomposition of densities (CDD) in atomic orbital RI-RPA¹⁶ calculations reduces the scaling with respect to basis set size by one order, removing the main drawback of the AO formulation over the canonical MO formulation. CDD can accelerate AO-RPA calculations by over 5× for double- ζ , over 10× for triple- ζ , and over 25× for quadruple- ζ basis sets, without any decrease in accuracy.

Depending on the number of necessary nodes for the frequency or Laplace quadrature, even for small molecular systems our method is competitive with the canonical MO-RI-RPA formulation of ref 6, due to the decoupling of the numerical integration in the frequency domain from the time-determining step. For sparse electronic systems, our ω -CDD-RPA method outperforms the canonical formulation over 10×.

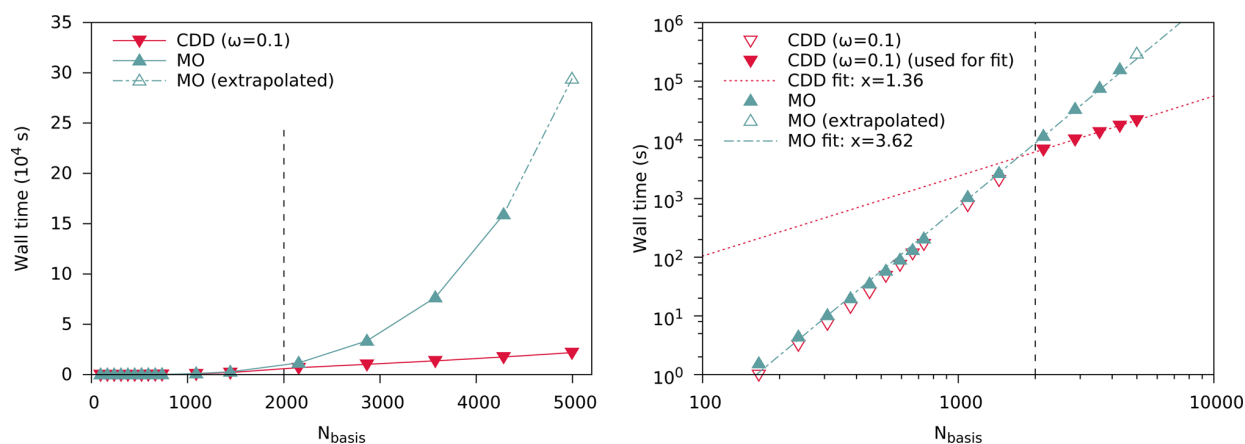


Figure 4. Plots of wall time against number of basis functions showing the computational complexity of the MO and ω -CDD formulations for glycine chains (def2-SVP, erfc-attenuated Coulomb metric, $\omega = 0.1$). Left: Linear plot. Right: Log–log plot, with linear fits for MO and ω -CDD. The vertical dashed line indicates the point after which sparse algebra is used.

Table 4. Wall Times and Computational Complexity of the ω -CDD and MO Formulations for Glycine Chains (def2-SVP, erfc-Attenuated Coulomb Metric, $\omega = 0.1$)^a

	N_{basis}	ω -CDD		MO		Tr ln (eq16)		total
		time (s)	$O(N^x)$	time (s)	$O(N^x)$	time (s)	$O(N^x)$	speedup
Gly ₂₀	1444	3690		2649		53 (3%)		1.3×
Gly ₃₀	2154	6719	3.0	11456	3.7	163 (2%)	2.8	1.7×
Gly ₄₀	2864	9850	1.3	33150	3.7	338 (3%)	2.6	3.3×
Gly ₅₀	3574	12970	1.2	76074	3.8	624 (5%)	2.8	5.6×
Gly ₆₀	4284	16402	1.3	158148	4.0	1063 (6%)	2.9	9.1×
Gly ₇₀	4994	20144	1.3	*292347	(4)	1667 (8%)	2.9	*13.5×

^aThe column headed $O(N^x)$ contains the scaling exponents relative to the row above. Values marked with an asterisk (*) are extrapolated. See text for detailed descriptions of the columns.

Furthermore, we introduced the use of the erfc-attenuated Coulomb metric for the resolution-of-the-identity in RPA calculations, which reproduces Coulomb-metric results to very high accuracy even in absolute energies while preserving the locality and performance benefits of the overlap metric. We demonstrated that the effective $O(N)$ scaling of AO-RPA in the overlap metric is fully retained, while improving accuracy at the same time.

Further research is required to obtain more efficient quadratures for the Laplace transform of eq 7, to increase accuracy and lower the number of necessary nodes. Similarly, the behavior of the transform for small-gap systems warrants further study.

■ APPENDIX A: RI-RPA EQUATIONS WITH CHOLESKY DECOMPOSITION OF OCCUPIED AND VIRTUAL SUBSPACES

When performing Cholesky factorizations also on all virtual pseudodensity matrices, eqs 19–21 simplify to

$$(\mathbf{F}_0)_{MN} = -\text{Tr}(\mathbf{U}_M^T \mathbf{U}_N) - \text{Tr}(\mathbf{V}_M^T \mathbf{V}_N) \quad (33)$$

$$(\mathbf{F}_{\text{INT}}^{(t_a)})_{MN} = \text{Tr}(\mathbf{W}_M^{(\alpha)T} \mathbf{W}_N^{(\alpha)}) \quad (34)$$

$$(\mathbf{F}_D^{(t_a)})_{MN} = -2\text{Tr}(\mathbf{X}_M^{(\alpha)T} \mathbf{X}_N^{(\alpha)}) - \text{Tr}(\mathbf{Y}_M^{(\alpha)T} \mathbf{Y}_N^{(\alpha)}) - \text{Tr}(\mathbf{Z}_M^{(\alpha)T} \mathbf{Z}_N^{(\alpha)}) \quad (35)$$

with

$$\mathbf{U}_M = \bar{\mathbf{L}}^T \mathbf{B}_M \underline{\mathbf{L}} \text{ and } \mathbf{P} = \underline{\mathbf{L}} \underline{\mathbf{L}}^T, \bar{\mathbf{P}} \bar{\mathbf{P}} = \bar{\mathbf{L}}' \bar{\mathbf{L}}'^T \quad (36)$$

$$\mathbf{V}_M = \bar{\mathbf{L}}^T \mathbf{B}_M \underline{\mathbf{L}}' \text{ and } -\mathbf{P} \bar{\mathbf{P}} = \underline{\mathbf{L}}' \underline{\mathbf{L}}'^T, \bar{\mathbf{P}} = \bar{\mathbf{L}} \bar{\mathbf{L}}^T \quad (37)$$

$$\mathbf{W}_M^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)T} \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ and } \underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)T}, \bar{\mathbf{P}}^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)} \bar{\mathbf{L}}^{(\alpha)T} \quad (38)$$

$$\mathbf{X}_M^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)T} \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ and } -\underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)T}, \bar{\mathbf{P}}^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)} \bar{\mathbf{L}}^{(\alpha)T} \quad (39)$$

$$\mathbf{Y}_M^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)T} \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ and } \underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)T} \quad (40)$$

$$\mathbf{Z}_M^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)T} \mathbf{B}_M \underline{\mathbf{L}}^{(\alpha)} \text{ and } \bar{\mathbf{P}}^{(\alpha)} = \bar{\mathbf{L}}^{(\alpha)} \bar{\mathbf{L}}^{(\alpha)T} \quad (41)$$

■ AUTHOR INFORMATION

Corresponding Author

*E-mail: christian.ochsenfeld@uni-muenchen.de.

ORCID

Christian Ochsenfeld: 0000-0002-4189-6558

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

The authors thank Dr. A. M. Burow and M. Beuerle (LMU Munich) for helpful discussions. Financial support was provided by the Deutsche Forschungsgemeinschaft (DFG) for the project Oc35/4-1, and the cluster of excellence EXC 114 “Center for Integrated Protein Science Munich” (CIPSM).

■ REFERENCES

- (1) Langreth, D. C. D.; Perdew, J. P. J. The Exchange-Correlation Energy of a Metallic Surface. *Solid State Commun.* **1975**, *17*, 1425.
- (2) Langreth, D. C.; Perdew, J. P. Exchange-correlation energy of a metallic surface: Wave-vector analysis. *Phys. Rev. B* **1977**, *15*, 2884.
- (3) Perdew, J. P. Jacob's ladder of density functional approximations for the exchange-correlation energy. *AIP Conf. Proc.* **2000**, 1–20.
- (4) Mezei, P. D.; Csonka, G. L.; Ruzsinszky, A.; Kállay, M. Construction and Application of a New Dual-Hybrid Random Phase Approximation. *J. Chem. Theory Comput.* **2015**, *11*, 4615–4626.
- (5) Furche, F. Molecular tests of the random phase approximation to the exchange-correlation energy functional. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2001**, *64*, 195120.
- (6) Eshuis, H.; Yarkony, J.; Furche, F. Fast computation of molecular random phase approximation correlation energies using resolution of the identity and imaginary frequency integration. *J. Chem. Phys.* **2010**, *132*, 234114.
- (7) Eshuis, H.; Bates, J. E.; Furche, F. Electron correlation methods based on the random phase approximation. *Theor. Chem. Acc.* **2012**, *131*, 1–18.
- (8) Eshuis, H.; Furche, F. Basis set convergence of molecular correlation energy differences. *J. Chem. Phys.* **2012**, *136*, 084105.
- (9) Dunlap, B. I.; Rösch, N. On the gaussian-type orbitals approach to local density functional theory. *J. Chim. Phys. Physico-Chimie Biol.* **1989**, *86*, 671.
- (10) Shenvi, N.; van Aggelen, H.; Yang, Y.; Yang, W. Tensor hypercontracted ppRPA: Reducing the cost of the particle-particle random phase approximation from $O(r^6)$ to $O(r^4)$. *J. Chem. Phys.* **2014**, *141*, 024119.
- (11) Hohenstein, E. G.; Parrish, R. M.; Martínez, T. J. Tensor hypercontraction density fitting. I. Quartic scaling second- and third-order Møller-Plesset perturbation theory. *J. Chem. Phys.* **2012**, *137*, 044103.
- (12) Parrish, R. M.; Hohenstein, E. G.; Martínez, T. J.; Sherrill, C. D. Tensor hypercontraction. II. Least-squares renormalization. *J. Chem. Phys.* **2012**, *137*, 224106.
- (13) Hohenstein, E. G.; Parrish, R. M.; Sherrill, C. D.; Martínez, T. J. Communication: Tensor hypercontraction. III. Least-squares tensor hypercontraction for the determination of correlated wavefunctions. *J. Chem. Phys.* **2012**, *137*, 221101.
- (14) Kaltak, M.; Klimeš, J.; Kresse, G. Low Scaling Algorithms for the Random Phase Approximation: Imaginary Time and Laplace Transformations. *J. Chem. Theory Comput.* **2014**, *10*, 2498–2507.
- (15) Kállay, M. Linear-scaling implementation of the direct random-phase approximation. *J. Chem. Phys.* **2015**, *142*, 204105.

- (16) Schurkus, H. F.; Ochsenfeld, C. Communication: An effective linear-scaling atomic-orbital reformulation of the random-phase approximation using a contracted double-Laplace transformation. *J. Chem. Phys.* **2016**, *144*, 031101.
- (17) Wilhelm, J.; Seewald, P.; Del Ben, M.; Hutter, J. Large-Scale Cubic-Scaling Random Phase Approximation Correlation Energy Calculations Using a Gaussian Basis. *J. Chem. Theory Comput.* **2016**, *12*, 5851–5859.
- (18) Jiang, H.; Engel, E. Random-phase-approximation-based correlation energy functionals: Benchmark results for atoms. *J. Chem. Phys.* **2007**, *127*, 184108.
- (19) Eshuis, H.; Furche, F. A Parameter-Free Density Functional That Works for Noncovalent Interactions. *J. Phys. Chem. Lett.* **2011**, *2*, 983–989.
- (20) Paier, J.; Ren, X.; Rinke, P.; Scuseria, G. E.; Grüneis, A.; Kresse, G.; Scheffler, M. Assessment of correlation energies based on the random-phase approximation. *New J. Phys.* **2012**, *14*, 043002.
- (21) Ren, X.; Rinke, P.; Joas, C.; Scheffler, M. Random-phase approximation and its applications in computational chemistry and materials science. *J. Mater. Sci.* **2012**, *47*, 7447–7471.
- (22) Mussard, B.; Reinhardt, P.; Ángyán, J. G.; Toulouse, J. Spin-unrestricted random-phase approximation with range separation: Benchmark on atomization energies and reaction barrier heights. *J. Chem. Phys.* **2015**, *142*, 154123.
- (23) Rekkedal, J.; Coriani, S.; Iozzi, M. F.; Teale, A. M.; Helgaker, T.; Pedersen, T. B. Communication: Analytic gradients in the random-phase approximation. *J. Chem. Phys.* **2013**, *139*, 081101.
- (24) Burrow, A. M.; Bates, J. E.; Furche, F.; Eshuis, H. Analytical First-Order Molecular Properties and Forces within the Adiabatic Connection Random Phase Approximation. *J. Chem. Theory Comput.* **2014**, *10*, 180–194.
- (25) Koch, H.; Sánchez De Merás, A.; Pedersen, T. B. Reduced scaling in electronic structure calculations using Cholesky decompositions. *J. Chem. Phys.* **2003**, *118*, 9481–9484.
- (26) Higham, N. J. Cholesky factorization. *Wiley Interdiscip. Rev. Comput. Stat.* **2009**, *1*, 251–254.
- (27) Harbrecht, H.; Peters, M.; Schneider, R. On the low-rank approximation by the pivoted Cholesky decomposition. *Appl. Numer. Math.* **2012**, *62*, 428–440.
- (28) Zienau, J.; Clin, L.; Doser, B.; Ochsenfeld, C. Cholesky-decomposed densities in Laplace-based second-order Møller-Plesset perturbation theory. *J. Chem. Phys.* **2009**, *130*, 204112.
- (29) Maurer, S. A.; Clin, L.; Ochsenfeld, C. Cholesky-decomposed density MP2 with density fitting: Accurate MP2 and double-hybrid DFT energies for large systems. *J. Chem. Phys.* **2014**, *140*, 224112.
- (30) Jung, Y.; Sodt, A.; Gill, P. M. W.; Head-Gordon, M. Auxiliary basis expansions for large-scale electronic structure calculations. *Proc. Natl. Acad. Sci. U. S. A.* **2005**, *102*, 6692–6697.
- (31) Jung, Y.; Shao, Y.; Head-Gordon, M. Fast evaluation of scaled opposite spin second-order Møller-Plesset correlation energies using auxiliary basis expansions and exploiting sparsity. *J. Comput. Chem.* **2007**, *28*, 1953–1964.
- (32) Reine, S.; Tellgren, E.; Krapp, A.; Kjærgaard, T.; Helgaker, T.; Jansik, B.; Host, S.; Salek, P. Variational and robust density fitting of four-center two-electron integrals in local metrics. *J. Chem. Phys.* **2008**, *129*, 104101.
- (33) Kussmann, J.; Luenser, A.; Beer, M.; Ochsenfeld, C. A reduced-scaling density matrix-based method for the computation of the vibrational Hessian matrix at the self-consistent field level. *J. Chem. Phys.* **2015**, *142*, 094101.
- (34) Ayala, P. Y.; Scuseria, G. E. Linear scaling second-order Møller-Plesset theory in the atomic orbital basis for large molecular systems. *J. Chem. Phys.* **1999**, *110*, 3660.
- (35) Luenser, A.; Kussmann, J.; Ochsenfeld, C. Computation of indirect nuclear spin-spin couplings with reduced complexity in pure and hybrid density functional approximations. *J. Chem. Phys.* **2016**, *145*, 124103.
- (36) Baerends, E.; Ellis, D.; Ros, P. Self-consistent molecular Hartree-Fock-Slater calculations I. The computational procedure. *Chem. Phys.* **1973**, *2*, 41–51.
- (37) Vahtras, O.; Almlöf, J.; Feyereisen, M. W. Integral approximations for LCAO-SCF calculations. *Chem. Phys. Lett.* **1993**, *213*, 514.
- (38) Adamson, R. D.; Dombroski, J. P.; Gill, P. M. W. Efficient calculation of short-range Coulomb energies. *J. Comput. Chem.* **1999**, *20*, 921–927.
- (39) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **1996**, *77*, 3865–3868.
- (40) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple [Phys. Rev. Lett. *77*, 3865 (1996)]. *Phys. Rev. Lett.* **1997**, *78*, 1396–1396.
- (41) Weigend, F.; Furche, F.; Ahlrichs, R. Gaussian basis sets of quadruple zeta valence quality for atoms H-Kr. *J. Chem. Phys.* **2003**, *119*, 12753.
- (42) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Phys. Chem. Chem. Phys.* **2005**, *7*, 3297.
- (43) Weigend, F.; Häser, M.; Patzelt, H.; Ahlrichs, R. RI-MP2: optimized auxiliary basis sets and demonstration of efficiency. *Chem. Phys. Lett.* **1998**, *294*, 143–152.
- (44) Hättig, C. Optimization of auxiliary basis sets for RI-MP2 and RI-CC2 calculations: Core-valence and quintuple- ζ basis sets for H to Ar and QZVPP basis sets for Li to Kr. *Phys. Chem. Chem. Phys.* **2005**, *7*, 59–66.
- (45) Takatsuka, A.; Ten-no, S.; Hackbusch, W. Minimax approximation for the decomposition of energy denominators in Laplace-transformed Møller-Plesset perturbation theories. *J. Chem. Phys.* **2008**, *129*, 044112.
- (46) Kussmann, J.; Ochsenfeld, C. Pre-selective screening for matrix elements in linear-scaling exact exchange calculations. *J. Chem. Phys.* **2013**, *138*, 134114.
- (47) Kussmann, J.; Ochsenfeld, C. Preselective Screening for Linear-Scaling Exact Exchange-Gradient Calculations for Graphics Processing Units and General Strong-Scaling Massively Parallel Calculations. *J. Chem. Theory Comput.* **2015**, *11*, 918–922.
- (48) TURBOMOLE, version 7.0; TURBOMOLE GmbH, 2015; <http://www.turbomole.com>. TURBOMOLE is a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH, 1989–2007, and TURBOMOLE GmbH, since 2007.
- (49) Rezac, J.; Riley, K. E.; Hobza, P. S66: A Well-balanced Database of Benchmark Interaction Energies Relevant to Biomolecular Structures. *J. Chem. Theory Comput.* **2011**, *7*, 2427–2438.
- (50) Rezac, J.; Riley, K. E.; Hobza, P. Extensions of the S66 Data Set: More Accurate Interaction Energies and Angular-Displaced Non-equilibrium Geometries. *J. Chem. Theory Comput.* **2011**, *7*, 3466–3470.
- (51) Sedlak, R.; Janowski, T.; Pitoňák, M.; Řezáč, J.; Pulay, P.; Hobza, P. Accuracy of Quantum Chemical Methods for Large Noncovalent Complexes. *J. Chem. Theory Comput.* **2013**, *9*, 3364–3374.
- (52) Structures are available online at www.cup.lmu.de/pc/ochsenfeld/.

Article II

”Almost error-free resolution-of-the-identity correlation methods by null space removal of the particle-hole interactions”,

H. F. Schurkus, A. Luenser, C. Ochsenfeld,

J. Chem. Phys. **146**, 211106 (2017).

Communication: Almost error-free resolution-of-the-identity correlation methods by null space removal of the particle-hole interactions

Henry F. Schurkus, Arne Luenser, and Christian Ochsenfeld

Chair of Theoretical Chemistry and Center for Integrated Protein Science Munich (CIPSM),
 Department of Chemistry, University of Munich (LMU), Butenandtstr. 7, D-81377 Munich, Germany

(Received 20 April 2017; accepted 24 May 2017; published online 7 June 2017)

We present a method to improve upon the resolution-of-the-identity (RI) for correlation methods. While RI is known to allow for drastic speedups, it relies on a cancellation of errors. Our method eliminates the errors introduced by RI which are known to be problematic for absolute energies. In this way, independence of the error compensation assumption for relative energies is also achieved. The proposed method is based on the idea of starting with an oversized RI basis and projecting out all of its unphysical parts. The approach can be easily implemented into existing RI codes and results in an overhead of about 30%, while effectively removing the RI error. In passing, this process alleviates the problem that for many frequently employed basis sets no optimized RI basis sets have been constructed. In this paper, the theory is presented and results are discussed exemplarily for the random phase approximation and Møller-Plesset perturbation theory. *Published by AIP Publishing.* [<http://dx.doi.org/10.1063/1.4985085>]

I. INTRODUCTION

The resolution-of-the-identity (RI) approximation^{1–6} is a common approach throughout the electronic structure theories of quantum chemistry in which an identity is resolved into a product of terms dependent on a preoptimized (auxiliary) basis. This allows to reformulate theories into a computationally more tractable form. For example, within explicitly correlated methods^{7–9} (F12) several three-electron integrals occur which are commonly split into products of two-electron integrals via RI. However, the by far most frequently encountered use of RI is the case of 4-center-2-electron repulsion integrals (ERIs) occurring in almost any correlation method, which are split into products of at most 3-center-2-electron integrals.^{10,11} The final formulations obtained in this case are identical to those that can be derived from another perspective, often called density-fitting¹² (DF). For this reason, in theories containing only this special type of RI the terms RI and DF are often used interchangeably.¹³ While we will focus on this major case of RI, i.e., DF, in the following, we expect the transfer to other types of RI to be straightforward. As an example, we sketch the application to F12 in the [supplementary material](#).

Preoptimized incomplete auxiliary bases have to balance accuracy against computational performance. Several attempts and discussions have therefore revolved around improving the RI approach.^{14–17} While a lot of effort has been put into advancing upon the approximation formula itself^{18–23} as well as the auxiliary basis sets,^{24–26} the null space structure of the physical models remains more or less opaque.

In this paper, we present an approach which explicitly recognizes the physical model of the correlation method. By starting with an oversized auxiliary basis, we eliminate the RI error and then project out any contribution spanned by the auxiliary basis which is not used by the physical model. While the theory extends to any precomputable null space structure

of the physical model, here we focus on the concrete examples of second order Møller-Plesset perturbation theory²⁷ (MP2; as a post-Hartree-Fock (HF) theory) and the direct random phase approximation²⁸ (RPA; as a post-Kohn-Sham (KS) theory) because the supports of their physical models are both encompassed by particle-hole space. This allows us to treat both of these theories simultaneously.

In formulating our theory, we employ the Cholesky decomposition (CD) and the singular value decomposition (SVD) which have been studied in related works by Aquilante, Lindh, and Pedersen^{21–23} and Kállay,²⁹ respectively.

II. THEORY

Consider a generic correlation method (Einstein's sum convention is used throughout)

$$E_C = Q_{\mu\nu\lambda\sigma}(\mu\nu|\lambda\sigma) \quad (1)$$

which contains a single set of ERIs

$$(\mu\nu|\lambda\sigma) \equiv (\mu\nu|\frac{1}{r_{12}}|\lambda\sigma) := \iint \chi_\mu(r_1)\chi_\nu(r_1)\frac{1}{|r_1-r_2|} \times \chi_\lambda(r_2)\chi_\sigma(r_2) dr_1 dr_2 \quad (2)$$

over basis functions $\{\chi\}$, where Q is given by the physical model $(\mu\nu|\frac{1}{r_{12}}|I)$ denotes a corresponding three-center integral. Note that if Q again depends on ERIs, the following considerations can be applied to each of them in turn. Even if the actual computation may be done in basis functions $\chi_\mu, \chi_\nu, \dots$, most physical models will not describe the interaction of all basis functions, but instead of a physically relevant subspace. Writing a theory like Eq. (1), however, hides this information in the physical model Q .

For the ERIs [Eq. (2)], the RI results in the following equality (see the [supplementary material](#) for details):

$$(\mu\nu|\lambda\sigma) = B_{\mu\nu,I}C_{I,J}B_{J,\lambda\sigma}^T, \quad (3)$$

where I and J are indices over an auxiliary basis set. The B and C matrices are computed in any chosen metric m_{12} as

$$B_{\mu\nu,I} = (\mu\nu|m_{12}|I), \quad (4)$$

$$C_{I,J} = (I|m_{12}|K)^{-1}(K|L)(L|m_{12}|J)^{-1}, \quad (5)$$

with the most common choice being the Coulomb metric $m_{12} = 1/r_{12}$. Note that $(K|L) = (K|\frac{1}{r_{12}}|L)$ is a Coulomb integral independent of the chosen metric. In typical approaches, a preoptimized auxiliary basis set is employed which is only 3-4 times larger in cardinality than the set $\{\chi\}$ because by the reasoning of DF one finds that Eq. (3) is the best approximation in a metric-specific sense also when the auxiliary basis is incomplete. While drastically improving performance, this introduces errors of several meVs in absolute energies. In contrast, our considerations start with a saturated auxiliary basis, such that Eq. (3) in fact effectively constitutes an identity.

By employing the RI, the correlation energy [Eq. (1)] becomes

$$E_C = \text{Tr}\{\tilde{Q}C\} \quad \text{with} \quad \tilde{Q}_{I,J} = B_{\mu\nu,I}Q_{\mu\nu\lambda\sigma}B_{\lambda\sigma,J}. \quad (6)$$

A. Example cases: RPA and MP2

As two exemplary cases we consider the RPA and MP2 methods. As post-KS/post-HF methods, they can be formulated in molecular orbitals as linear combinations of atomic orbitals. $c_{\mu i}$, $c_{\nu a}$, etc. denote the coefficient matrix elements, where i, j, k, \dots index the occupied orbitals and a, b, c, \dots the virtual orbitals. Within the RI-RPA^{11,30-33} in the time-determining step, one has to compute a quantity of the form

$$B_{I,ia}^T \frac{\varepsilon_a - \varepsilon_i}{(\varepsilon_a - \varepsilon_i)^2 + \omega^2} B_{ia,J}, \quad (7)$$

where ω is a variable to be integrated over in a later step of the algorithm and ε_i , etc. denote orbital energies. Therefore, we can associate

$$Q_{\mu\nu\lambda\sigma} \sim c_{\mu i}c_{\nu a} \frac{\varepsilon_a - \varepsilon_i}{(\varepsilon_a - \varepsilon_i)^2 + \omega^2} c_{\lambda i}c_{\sigma a}. \quad (8)$$

Within RI-MP2³⁴ on the other hand, the time-determining step is the recombination of the ERIs according to Eq. (3) in computing

$$t_{ij}^{ab} B_{ia,I}C_{I,J}B_{J,jb}^T \quad \text{with} \quad t_{ij}^{ab} = -\frac{(ia|jb) - (ib|ja)}{\varepsilon_a + \varepsilon_b - \varepsilon_i - \varepsilon_j}, \quad (9)$$

so we can associate

$$Q_{\mu\nu\lambda\sigma} \sim c_{\mu i}c_{\nu a}t_{ij}^{ab}c_{\lambda j}c_{\sigma b}. \quad (10)$$

Both of these methods share a commonality: Their physical models are built solely around particle-hole (ph) interactions as

$$Q_{\mu\nu\lambda\sigma} = c_{\mu i}c_{\nu a} \cdots c_{\lambda i}c_{\sigma a}. \quad (11)$$

To treat both methods within the same derivation, we will only focus on this common property and restrict ourselves to treat the inner part as a black box.

B. Null space projection

We aim to project out all auxiliary basis functions and linear combinations in the null space of the physical model ($\text{Null}(\tilde{Q}) = \{x \mid \tilde{Q}x = 0\}$) because these will not contribute to the final energy, or equivalently project onto its complement, called support.

We begin with a thought experiment: If \tilde{Q} was known beforehand, one could construct a minimal-rank projector P onto the support of \tilde{Q} such that

$$\tilde{Q}P = \tilde{Q}. \quad (12)$$

To do so, one would first apply SVD to \tilde{Q}

$$\tilde{Q} = U\Sigma V^T, \quad (13)$$

where Σ is a diagonal matrix containing only $\text{rank}(\tilde{Q})$ non-zero singular values and the columns of the unitary matrices U and V are the corresponding left and right singular vectors. We now discard all singular vectors with a corresponding zero singular value. In numerical implementations, it is useful to further discard those corresponding to very small singular values. This step is justified by the Eckart-Young-Mirsky theorem.³⁵ Precisely which numerical values qualify as very small in this context will be further analyzed in Sec. III A.

As the left and right singular vectors form an orthogonal system each, recombining the remaining singular vectors according to

$$P^L = UU^T, \quad P^R = VV^T \quad (14)$$

defines a left and a right projector. The right projector fulfils Eq. (12) and the left projector can similarly be introduced to the left of \tilde{Q} at any point in the derivation.

Therefore one can redefine Eqs. (4) and (5) without changing the final correlation energy, Eq. (6),

$$B_{\mu\nu,I} \rightarrow B_{\mu\nu,J}V_{J,I}, \quad (15)$$

$$C \rightarrow V^T C V. \quad (16)$$

This reduces the effective size of the auxiliary basis in all the time determining steps to the width of the singular vector matrices, which *per constructionem* is exactly limited by the rank of the physical model.

For the case at hand, there is no need for additional steps in the construction. Nonetheless, we want to propose one that may prove useful in other cases. Once the projector matrix is known, it can be decomposed by pivoted CD as

$$P = LL^T, \quad (17)$$

where the number of columns in the Cholesky factor L again equals the rank of P . Therefore, Eqs. (15) and (16) can be carried out with L instead of V . This may prove useful, for example, when P is sparse because unlike the singular decomposition, pivoted CD on a sparse matrix can return sparse factors. Although not necessary, we use the Cholesky factors throughout, since their use does not deteriorate the results compared to the use of V .

C. Auxiliary matrix construction

Computing the exact physical model is almost always part of the time determining step. To exploit the previous

discussion, it is therefore necessary to find an auxiliary matrix H , such that

$$\text{Null}(H) \subseteq \text{Null}(\tilde{Q}), \quad (18)$$

and construct the projectors from H instead of \tilde{Q} . The central idea to finding H is to recognize that because in a k -linear map $N : W \rightarrow X$, e.g., a matrix or tensor product, the zero element in W is mapped to the zero element in X , in a succession of maps $M : V \rightarrow W, N : W \rightarrow X, \dots$ the null space of the first map is contained within the null space of the succession of all maps.

Returning to the generic example, Eq. (11), this means

$$\text{Null}(c_{\mu i} c_{\nu a} B_{\mu\nu, I}) \subseteq \text{Null}(\dots c_{\mu i} c_{\nu a} B_{\mu\nu, I}) \quad (19)$$

allowing the definition of H for this case as

$$H_{ia, I} = c_{\mu i} c_{\nu a} B_{\mu\nu, I}. \quad (20)$$

By the SVD decomposition of H , it becomes obvious that $H^T H$ in fact has the same span as H (although in numerical implementations a tighter threshold for the singular values is needed) so we may equally define

$$H_{I, J} = B_{I, \mu\nu}^T c_{\mu i} c_{\nu a} c_{\lambda i} c_{\sigma a} B_{\lambda\sigma, J}, \quad (21)$$

which is favorable as it allows for a cheaper SVD. Applying the discussion of Sec. II B to H instead of \tilde{Q} thus delivers the new projection. Because it does not make use of all the internal structure of \tilde{Q} , it may eliminate less than—however due to Eq. (18) never more than—the projector constructed directly from \tilde{Q} . Therefore no additional error is introduced.

While the argument extends naturally to any other operator with a non-trivial null space structure, e.g., $\tilde{Q}C$, by replacing \tilde{Q} with that operator in Eq. (18), we will only evaluate the straightforward example, which we have presented here, in Sec. III.

III. NUMERICAL EVALUATION

All calculations have been done in a development version of the program package FermiONs++^{36,37} employing the Coulomb metric. All RPA calculations have been based on self-consistent orbitals obtained with the Perdew-Burke-Ernzerhof (PBE) functional.

A. Accuracy

Two practical questions have yet to be addressed: The auxiliary basis to start the projection from and the numerical threshold to choose for the SVD.

In principle, any auxiliary basis large enough to fulfill Eq. (3) as an identity can be taken as the starting point. While a naïve overly large auxiliary basis can be used, it is advisable to make use of specialized auxiliary basis sets whenever available, to lessen the overhead of the projector construction.

To find a suitable choice for both the starting basis and the threshold, we evaluated all 198 monomers and dimers within the S66 test set of small to medium sized molecules.³⁸ In Fig. 1 we compare the deviation in energy from the RI-free result to the total number of auxiliary basis functions left

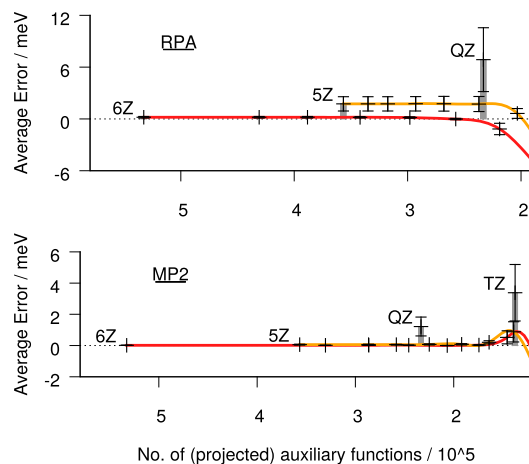


FIG. 1. Mean errors (left axis) and standard deviations (error bars) against the number of auxiliary functions using ph-projection. Solid lines: Projection from cc-pV6Z-RI (red) and cc-pV5Z-RI (orange). Larger thresholds result in fewer auxiliary basis functions. Gray bars labelled XZ = 6Z, 5Z, QZ, TZ represent the corresponding unprojected cc-pVXZ-RI bases. Mean error and standard deviation nearly vanish for cc-pV6Z-RI. Null space removal grants sub-meV accuracy with an auxiliary basis set size only slightly larger than the canonical RI basis.

after the projection along the continuous range of thresholds. The variation over the different molecules is expressed by the empirical standard deviation of the energy differences (at threshold values 0 and $10^{-10}, 10^{-9}, \dots, 10^{-5}$, respectively). The unprojected RI results with the canonical discrete set of preoptimized auxiliary basis sets (cc-pVXZ-RI) are also given for comparison.

RPA/cc-pVQZ results are effectively converged with an auxiliary basis two cardinality numbers larger than the AO basis, cc-pV6Z-RI (Fig. 1, top). Kállay's very different reasoning²⁹ had focused on reducing the size of the typical auxiliary basis (here cc-pVQZ-RI). In contrast, our understanding as presented above, has led us to consider saturated auxiliary basis sets. We find that even when aiming for the same auxiliary size after reduction, projections based on larger basis sets lead to significantly better results. SVD thresholds tighter than 10^{-6} hardly change the results, so we recommend the combination cc-pV6Z-RI $\rightarrow 10^{-6}$ and use it for the following analysis.

For MP2 (Fig. 1, bottom), where the typical basis set size is triple- ζ , here cc-pVTZ, already the pentuple- ζ auxiliary basis can be considered converged. Because projecting from a smaller basis means that less auxiliary functions are left after

TABLE I. Comparison of accuracy gain and computational overhead of cc-pV6Z-RI $\rightarrow 10^{-6}$ ph-projection against canonical cc-pVQZ-RI for molecules in Fig. 2 (geometries and more detailed results are given in the supplementary material). All calculations were carried out at the RPA@PBE/cc-pVQZ level with 12 concurrent evaluation threads on an Intel E5-2620@2.0GHz architecture.

Molecule	A	B	C
Canonical RI error (meV)	30.2	40.9	56.5
ph-projection error (meV)	0.7	0.6	0.8
Total walltime overhead	+27%	+35%	+14%




FIG. 2. Selection of molecules employed for testing (see Table I; $C_{26}H_{22}O_2S$,³⁹ adenosine-thymine base pair, and C_{60}).

projection with the same threshold, we recommend tightening the threshold to 10^{-7} in this case.

B. Overhead

Our recommended choices for the projector construction (see Sec. III A) in eliminating the RI error result in slightly more auxiliary functions entering the energy evaluation than canonical RI (e.g., cc-pVQZ-RI for cc-pVQZ). Additional compute time overhead is caused by the construction of the projector itself. While the most expensive step scales formally as $\mathcal{O}(N^4)$ [Eq. (21); asymptotic limit with a local metric: $\mathcal{O}(N)$] the same holds true for the RI correlation methods [MP2: $\mathcal{O}(N^5)$, RPA: $\mathcal{O}(N^4 \log(N))$], so the overhead is approximately given by a constant fraction of the total runtime. In Table I we compare the overall overhead of our method to the gains in accuracy for some larger molecules (Fig. 2). Timings comprise the complete correlation energy calculation after the generation of the Kohn–Sham orbitals. This includes projector construction, integral evaluation and transformation, and evaluation of the Hamiltonian expectation value as part of the full RPA energy. Our recommended choices eliminate the RI error to an insignificant residue of less than one meV in absolute energy also for these larger molecules, causing a total overhead of only about 30% compared to the canonical RI approach, which shows two orders of magnitude larger errors.

C. Potential energy surfaces

As the projection is essentially error-free, so are potential energy surfaces evaluated with it. We computed the dissociation curve of the C_{60} dimer in Fig. 3. The equilibrium distance can be deduced from experiment to be 9.93 Å,⁴⁰ in good agreement with our results. Obtaining a good experimental estimate for the binding energy is more intricate.^{41,42}




FIG. 3. cc-pV6Z-RI \rightarrow 10^{-6} ph-projected potential energy dissociation curve of the C_{60} dimer. Geometries were taken from Ref. 43. The experimental equilibrium distance^{40,44} is given as a black line for reference. All calculations have been counter-poise corrected.⁴⁵

IV. CONCLUSION

We have introduced the concept of null space projection of the physical model. The theory was presented in a general manner and is not limited to the discussed applications at the MP2 and RPA levels. For these specifically we have shown that the RI error can be eliminated to a residue of less than one meV while overall runtime increases by only about 30%, within a few lines of code. Other correlation methods may similarly benefit from our presented null space projection idea.

SUPPLEMENTARY MATERIAL

See [supplementary material](#) for further results and discussions as indicated in the text.

ACKNOWLEDGMENTS

We acknowledge funding by DFG: Oc35/4-1, SFB749, and EXC114.

- ¹S. F. Boys and I. Shavitt, “A fundamental calculation of the energy surface for the system of three hydrogen atoms,” Technical Report, University of Wisconsin, 1959.
- ²J. L. Whitten, *J. Chem. Phys.* **58**, 4496 (1973).
- ³E. Baerends, D. Ellis, and P. Ros, *Chem. Phys.* **2**, 41 (1973).
- ⁴B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, *J. Chem. Phys.* **71**, 4993 (1979).
- ⁵C. Van Alsenoy, *J. Comput. Chem.* **9**, 620 (1988).
- ⁶O. Vahtras, J. Almlöf, and M. W. Feyereisen, *Chem. Phys. Lett.* **213**, 514 (1993).
- ⁷W. Kutzelnigg, *Theor. Chim. Acta* **68**, 445 (1985).
- ⁸C. Hättig, W. Klopper, A. Köhn, and D. P. Tew, *Chem. Rev.* **112**, 4 (2012).
- ⁹L. Kong, F. A. Bischoff, and E. F. Valeev, *Chem. Rev.* **112**, 75 (2012).
- ¹⁰M. Feyereisen, G. Fitzgerald, and A. Komornicki, *Chem. Phys. Lett.* **208**, 359 (1993).
- ¹¹H. Eshuis, J. Yarkony, and F. Furche, *J. Chem. Phys.* **132**, 234114 (2010).
- ¹²H.-J. Werner, F. R. Manby, and P. J. Knowles, *J. Chem. Phys.* **118**, 8149 (2003).
- ¹³In Ref. 12 the term DF was introduced as a synonym to RI, highlighting the difference in perspective. Some authors have adapted this nomenclature [cf. J. G. Hill, *Int. J. Quantum Chem.* **113**, 21 (2013)], others have not [cf. T. Kjærgaard, *J. Chem. Phys.* **146**, 044103 (2017)].
- ¹⁴I. Duchemin, J. Li, and X. Blase, *J. Chem. Theory Comput.* **13**, 1199 (2017).
- ¹⁵F. Weigend, M. Kattannek, and R. Ahlrichs, *J. Chem. Phys.* **130**, 164106 (2009).
- ¹⁶L. Grajciar, *J. Comput. Chem.* **36**, 1521 (2015).
- ¹⁷S. Manzer, P. R. Horn, N. Mardirossian, and M. Head-Gordon, *J. Chem. Phys.* **143**, 024113 (2015).
- ¹⁸E. G. Hohenstein, R. M. Parrish, and T. J. Martínez, *J. Chem. Phys.* **137**, 044103 (2012).
- ¹⁹R. M. Parrish, E. G. Hohenstein, T. J. Martínez, and C. D. Sherrill, *J. Chem. Phys.* **137**, 224106 (2012).
- ²⁰E. G. Hohenstein, R. M. Parrish, C. D. Sherrill, and T. J. Martínez, *J. Chem. Phys.* **137**, 221101 (2012).
- ²¹F. Aquilante, T. B. Pedersen, A. Sánchez de Merás, and H. Koch, *J. Chem. Phys.* **125**, 174101 (2006).
- ²²T. B. Pedersen, F. Aquilante, and R. Lindh, *Theor. Chem. Acc.* **124**, 1 (2009).
- ²³F. Aquilante, L. Boman, J. Boström, H. Koch, R. Lindh, A. Sánchez de Merás, and T. B. Pedersen, *Linear-Scaling Techniques in Computational Chemistry and Physics* (Springer, Dordrecht, The Netherlands, 2011), Vol. 13, pp. 301–343.
- ²⁴F. Weigend, A. Köhn, and C. Hättig, *J. Chem. Phys.* **116**, 3175 (2002).
- ²⁵C. Hättig, *Phys. Chem. Chem. Phys.* **7**, 59 (2005).
- ²⁶G. L. Stoychev, A. A. Auer, and F. Neese, *J. Chem. Theory Comput.* **13**, 554 (2017).
- ²⁷C. Møller and M. S. Plesset, *Phys. Rev.* **46**, 618 (1934).
- ²⁸D. Bohm and D. Pines, *Phys. Rev.* **92**, 609 (1953).
- ²⁹M. Kállay, *J. Chem. Phys.* **141**, 244113 (2014).

- ³⁰F. Furche, *J. Chem. Phys.* **129**, 114105 (2008).
- ³¹H. Eshuis and F. Furche, *J. Chem. Phys.* **136**, 084105 (2012).
- ³²H. F. Schurkus and C. Ochsenfeld, *J. Chem. Phys.* **144**, 031101 (2016).
- ³³A. Luenser, H. F. Schurkus, and C. Ochsenfeld, *J. Chem. Theory Comput.* **13**, 1647 (2017).
- ³⁴F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs, *Chem. Phys. Lett.* **294**, 143 (1998).
- ³⁵C. Eckart and G. Young, *Psychometrika* **1**, 211 (1936).
- ³⁶J. Kussmann, M. Beer, and C. Ochsenfeld, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **3**, 614 (2013).
- ³⁷J. Kussmann, A. Luenser, M. Beer, and C. Ochsenfeld, *J. Chem. Phys.* **142**, 094101 (2015).
- ³⁸J. Řezáč, K. E. Riley, and P. Hobza, *J. Chem. Theory Comput.* **7**, 2427 (2011).
- ³⁹T. Gatzemeier, M. van Gemmeren, Y. Xie, D. Hoffer, M. Leutzsch, and B. List, *Science* **351**, 949 (2016).
- ⁴⁰P. A. Heiney, J. E. Fischer, A. R. McGhie, W. J. Romanow, A. M. Denenstein, J. P. McCauley, Jr., A. B. Smith, and D. E. Cox, *Phys. Rev. Lett.* **66**, 2911 (1991).
- ⁴¹J. Abrefah, D. R. Olander, M. Balooch, and W. J. Siekhaus, *Appl. Phys. Lett.* **60**, 1313 (1992).
- ⁴²F. Tournus, J.-C. Charlier, and P. Mélinon, *J. Chem. Phys.* **122**, 094315 (2005).
- ⁴³D. I. Sharapa, J. T. Margraf, A. Hesselmann, and T. Clark, *J. Chem. Theory Comput.* **13**, 274 (2017).
- ⁴⁴Y.-y. Ohnishi, K. Ishimura, and S. Ten-no, *J. Chem. Theory Comput.* **10**, 4857 (2014).
- ⁴⁵S. Boys and F. Bernardi, *Mol. Phys.* **19**, 553 (1970).

SUPPLEMENTARY MATERIAL

Almost Error-Free Resolution-of-the-Identity Correlation Methods By Null Space Removal Of The Particle-Hole Interactions

Henry F. Schurkus, Arne Luenser, and Christian Ochsenfeld

*Chair of Theoretical Chemistry and Center for Integrated
Protein Science Munich (CIPSM), Department of Chemistry,
University of Munich (LMU), Butenandtstr. 7, D-81377 Munich,
Germany*

I. GEOMETRIES

A. $C_{26}H_{22}O_2S$

Element	x/Bohr	y/Bohr	z/Bohr
C	2.30641074	10.10228689	5.48436316
H	2.28581272	11.98539897	5.63780892
C	-0.36792968	9.06841773	5.35888534
C	-2.28354505	9.20485595	7.15166852
H	-2.03939243	10.10247586	8.68744894
C	-4.57710565	7.99467534	6.64711164
H	-5.89367784	8.04569795	7.86560705
C	-5.11528075	6.52700065	4.02723314
H	-6.51955513	5.88101667	4.07065904
C	-3.06135632	6.63067103	2.53393376
H	-3.32931949	5.79125468	0.96924053
C	-0.76892956	7.81704109	3.05700995
C	1.55845713	7.95933746	1.54881953
C	2.16109080	6.93718460	-0.80370052
H	0.92747758	6.05184791	-1.76368139
C	4.61301044	7.24275332	-1.71700516
H	5.05766300	6.54771205	-3.31174503
C	6.42242321	8.55649092	-0.32276522
H	8.08689398	8.73034572	-0.96924053
C	5.81430934	9.62002878	2.01218038
H	7.04017468	10.53824671	2.94853967
C	3.37353908	9.30879089	2.94230358
C	3.94159075	8.93141258	7.55191251
H	4.05421843	7.08099276	7.30605914
H	5.67031221	9.64100474	7.47518963
C	4.16609021	8.67611058	11.98332027

C	2.87597419	9.24321739	14.45961739
H	3.42796319	10.95625413	15.03503899
C	-0.07917952	9.19786397	14.38100478
H	-0.90971416	10.64917363	13.50077035
C	-0.81069251	6.57908150	13.46278686
H	-1.54636289	6.17619189	11.87636178
C	-0.22922378	4.94333457	15.30130140
H	-0.45977037	3.16377948	15.24801113
C	0.86814018	6.44812348	17.49262782
H	0.79595264	5.61437632	19.18771215
C	3.55646457	7.28735085	16.57459887
H	4.41515612	5.80467174	15.77770136
C	-0.62984572	8.93991635	17.23071178
H	-2.45191965	8.74073922	17.60695625
H	0.09391939	10.33755779	18.24303806
C	5.23850979	8.32821200	18.61739281
O	2.81418014	9.48113391	9.98285620
O	6.22211224	7.68343745	11.77601732
S	8.33066865	8.89116142	18.08014367
C	8.82861148	9.85529968	21.15340527
H	10.42013882	10.40350923	21.78136126
C	6.70304754	9.74588454	22.57750287
H	6.61933267	10.20130854	24.31227146
C	4.60753024	8.85903607	21.13167342
H	2.95685447	8.64927647	21.80328208

B. DNA (AT-base pair)

Element	x/Bohr	y/Bohr	z/Bohr
---------	--------	--------	--------

H	-50.07585257	-13.74775756	7.89017349
O	-48.92330861	-12.77341477	6.89031940
H	-48.44426304	-15.75974896	4.43953358
H	-44.47621612	-12.74658066	5.94829092
C	-48.15059960	-13.72167934	4.49792612
H	-51.43437669	-11.83818931	3.00617632
H	-49.22812144	-13.14285622	0.61718455
C	-45.34586809	-13.14134444	4.12319343
H	-42.25333128	-14.26648738	1.76519317
H	-43.08348797	-16.49844290	4.12394932
C	-49.49079337	-12.20857563	2.43415622
O	-45.33793124	-10.73288850	2.79036960
H	-48.37264242	-8.80480093	4.35033851
C	-44.10394008	-15.15106818	2.75219713
C	-47.86241637	-9.84037085	2.64429377
H	-48.53270222	-10.54901815	-2.66470281
O	-45.40917391	-16.45441228	0.96583902
N	-48.05819199	-8.13111357	0.41498386
C	-48.38416975	-8.76284901	-2.06603757
C	-47.94197383	-5.54540131	0.46562852
N	-48.48375831	-6.83343864	-3.57441696
H	-48.32558824	0.74946538	-4.94503532
H	-48.60999202	-2.23970340	-6.24800149
N	-47.63168081	-4.05629713	2.51144602
H	-47.38412668	-0.43463701	3.28528887
C	-48.20370090	-4.76891285	-1.98099990
N	-48.38057927	-1.12797752	-4.73716545
C	-47.60768128	-1.65785673	1.86213612
C	-48.14644220	-2.15485470	-2.43207752
N	-47.83898376	-0.62682216	-0.42216482

H	-44.91066416	10.68961377	-6.04882435
H	-46.52997047	8.12374364	-7.49994504
H	-48.23507036	10.85609864	-6.53108246
O	-47.05455845	4.26643467	-5.21602205
H	-47.76056013	2.90658775	-0.81995217
C	-46.66225130	9.60642275	-6.07622538
C	-47.24542079	5.71528769	-3.40112908
N	-47.64075149	4.78421963	-0.99739745
O	-48.23431447	5.19277842	3.23388832
C	-47.09027427	8.43649331	-3.53019737
H	-47.20989393	11.73141978	-1.56903960
C	-47.88736075	6.19093176	1.18693698
C	-47.32346647	9.84981948	-1.43505802
H	-50.98764543	17.48355713	0.16742973
N	-47.71426184	8.76341593	0.88212415
H	-47.50506915	9.22054068	4.85546230
O	-50.38312204	17.04589656	1.81678270
H	-53.55748399	14.81979919	2.41110156
H	-52.79195593	16.91871799	4.90969744
C	-47.97069767	10.34814026	3.19571585
H	-44.41102057	12.32063639	3.76131088
H	-46.55378102	13.75569441	1.38932665
C	-51.73446519	15.64787717	3.49391463
O	-50.49858431	11.23328797	3.32308339
H	-51.44552607	13.07161355	6.56226294
C	-46.33891916	12.72333702	3.15886619
C	-50.52787506	13.56521001	4.78535347
H	-47.44100744	16.15999295	5.32694897
C	-47.73467088	14.12211230	5.26591082
O	-47.02526769	13.04383457	7.62239930

H -47.88811664 13.90422688 8.96637252
H -44.33354180 -17.71674934 0.23035761

C. C60

Element	x/Bohr	y/Bohr	z/Bohr
C	-1.36154741	-4.88927519	4.34643598
C	1.36154741	4.88927519	-4.34643598
C	-1.36154741	-4.88927519	-4.34643598
C	1.36154741	-4.88927519	-4.34643598
C	-1.36154741	4.88927519	-4.34643598
C	1.36154741	-4.88927519	4.34643598
C	-1.36154741	4.88927519	4.34643598
C	1.36154741	4.88927519	4.34643598
C	-2.68624518	-5.70798339	2.20303021
C	2.68624518	5.70798339	-2.20303021
C	-2.68624518	-5.70798339	-2.20303021
C	2.68624518	-5.70798339	-2.20303021
C	-2.68624518	5.70798339	-2.20303021
C	2.68624518	-5.70798339	2.20303021
C	-2.68624518	5.70798339	2.20303021
C	2.68624518	5.70798339	2.20303021
C	-2.20303021	-2.68624518	5.70798339
C	2.20303021	2.68624518	-5.70798339
C	-2.20303021	-2.68624518	-5.70798339
C	2.20303021	-2.68624518	-5.70798339
C	-2.20303021	2.68624518	-5.70798339
C	2.20303021	-2.68624518	5.70798339
C	-2.20303021	2.68624518	5.70798339

C 2.20303021 2.68624518 5.70798339
 C -4.88927519 -4.34643598 1.36154741
 C 4.88927519 4.34643598 -1.36154741
 C -4.88927519 -4.34643598 -1.36154741
 C 4.88927519 -4.34643598 -1.36154741
 C -4.88927519 4.34643598 -1.36154741
 C 4.88927519 -4.34643598 1.36154741
 C -4.88927519 4.34643598 1.36154741
 C 4.88927519 4.34643598 1.36154741
 C -1.32469777 -6.54946600 0.00000000
 C 1.32469777 -6.54946600 0.00000000
 C -1.32469777 6.54946600 0.00000000
 C 1.32469777 6.54946600 0.00000000
 C -6.54946600 0.00000000 1.32469777
 C 6.54946600 0.00000000 -1.32469777
 C -6.54946600 0.00000000 -1.32469777
 C 6.54946600 0.00000000 1.32469777
 C 0.00000000 -1.32469777 6.54946600
 C 0.00000000 1.32469777 -6.54946600
 C 0.00000000 -1.32469777 -6.54946600
 C 0.00000000 1.32469777 6.54946600
 C -4.34643598 -1.36154741 4.88927519
 C 4.34643598 1.36154741 -4.88927519
 C -4.34643598 -1.36154741 -4.88927519
 C 4.34643598 -1.36154741 -4.88927519
 C -4.34643598 1.36154741 -4.88927519
 C 4.34643598 -1.36154741 4.88927519
 C -4.34643598 1.36154741 4.88927519
 C 4.34643598 1.36154741 4.88927519
 C -5.70798339 -2.20303021 2.68624518

C	5.70798339	2.20303021	-2.68624518
C	-5.70798339	-2.20303021	-2.68624518
C	5.70798339	-2.20303021	-2.68624518
C	-5.70798339	2.20303021	-2.68624518
C	5.70798339	-2.20303021	2.68624518
C	-5.70798339	2.20303021	2.68624518
C	5.70798339	2.20303021	2.68624518

II. EXTENDED TABLES

Here, we present more detailed results for the three example molecules of the main text. All calculations are RPA energy evaluations and have been carried out in a cc-pVQZ atomic orbital basis and based on PBE orbitals.

N_{atoms} : Number of atoms in the molecule.

N_{aux} : Number of (projected) pure auxiliary functions

6Z-RI: Evaluation with cc-pV6Z-RI auxiliary basis.

QZ-RI: Evaluation with cc-pVQZ-RI auxiliary basis.

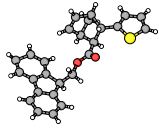
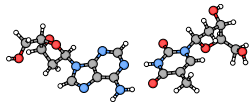

proj: Evaluation with cc-pV6Z-RI auxiliary basis projected to threshold 10^{-6} .

Speedup vs 6Z-RI: comparing “proj” with 6Z-RI

Overhead vs QZ-RI: comparing “proj” with QZ-RI

E_{total} : Absolute correlated energy.

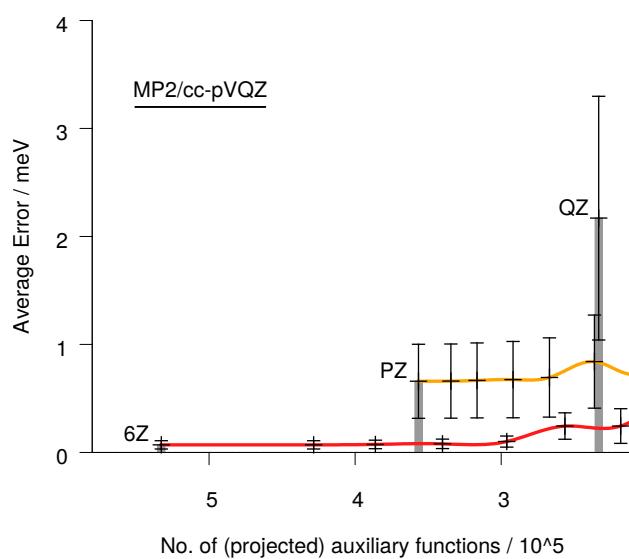
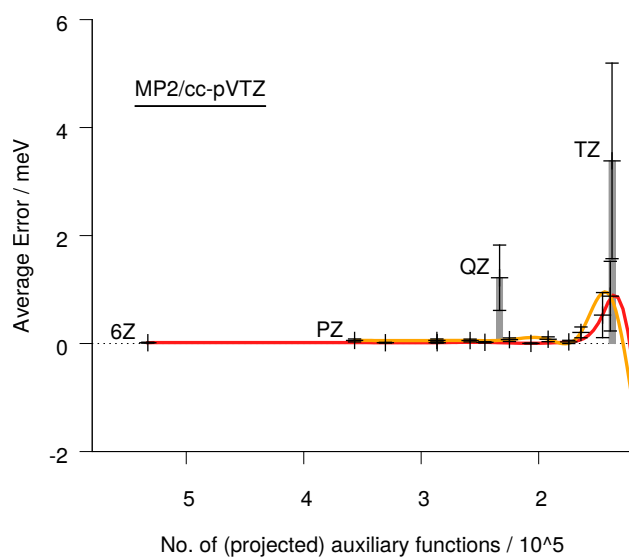
E_C^{RPA} : RPA correlation energy.

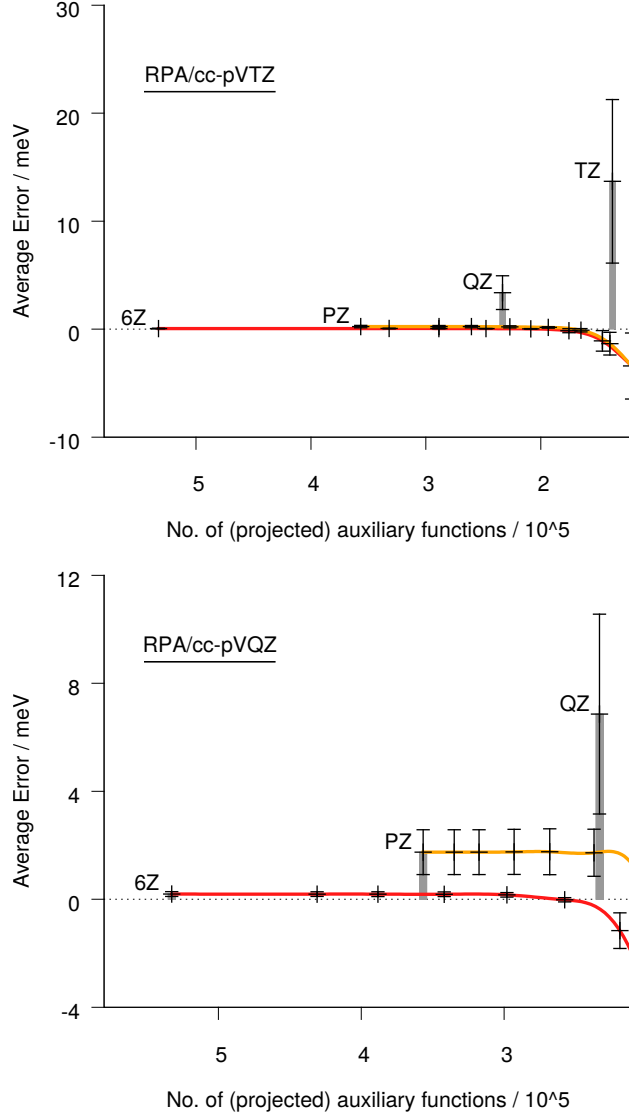
			
N_{atoms}	51	62	60
Speedup vs 6Z-RI	2.5x	2.8x	1.9x
Naux-Reduction	-50.0%	-49.3%	-47.7%
Error of proj.	2.6e-5 hartree	2.4e-5 hartree	2.8e-5 hartree
Error of QZ-RI	1.1e-3 hartree	1.5e-3 hartree	2.1e-3 hartree
Overhead vs QZ-RI	+27.2%	+35.4%	+14.4%
cc-pV6Z-RI:			
walltime[s]	27890	56531	172617
N_{aux}	11346	13685	16980
Etotal[hartree]	-1552.5715171846	-1764.3369534545	-2272.0007428025
E_C^{RPA} [hartree]	-8.09901141777614	-10.90586404074700	-14.45599212934066
cc-pVQZ-RI:			
walltime[s]	8611	14874	80667
N_{aux}	5064	6105	7920
Etotal[hartree]	-1552.5704047828	-1764.3354520156	-2272.0007428025
E_C^{RPA} [hartree]	-8.09789901599577	-10.90436260185441	-14.45391667695902
cc-pV6Z-RI→1e-6:			
walltime[s]	10954	20141	92273
N_{aux}	5650	6935	8873
Etotal[hartree]	-1552.5715434597	-1764.3369772871	-2272.0007428025
E_C^{RPA} [hartree]	-8.09903769282070	-10.90588787332585	-14.45601973832719

III. EXTENSION ON S66 ACCURACY

In the following we present more results as in fig. 1 of the main text (see caption of fig. 1 for details). While in fig. 1 we show only the typical combinations of atomic orbital

basis sets and correlation methods, here we present all combinations of RPA/MP2 with cc-pVTZ/cc-pVQZ.





IV. DENSITY FITTING AS A SPECIAL CASE OF RESOLUTION-OF-THE-IDENTITY

In the literature both terms, “density fitting” (DF) and “resolution-of-the-identity” (RI), are used to denote the factorization of the ERIs as

$$(\mu\nu|\lambda\sigma) = (\mu\nu|m_{12}|I)(I|m_{12}|J)^{-1}(J|K)(K|m_{12}|L)^{-1}(L|m_{12}|\lambda\sigma) . \quad (1)$$

However, they differ in perspective: While RI considers an exact identity to be introduced, DF considers how to minimize the error of fitting a product of functions by a linear combination over single functions in some metric. We now show that both perspectives yield

eq. 1 – RI as an exact identity over a complete set, DF as an optimal approximation over an incomplete set.

In DF one is interested in fitting the “density” $|\mu\nu\rangle$ with an in general incomplete auxiliary function set

$$|\mu\nu\rangle \approx |\tilde{\mu}\tilde{\nu}\rangle = \alpha_I^{\mu\nu}|I\rangle \quad (2)$$

where the coefficients α are to be determined such as to minimize the error within a metric m_{12}

$$0 \stackrel{!}{=} \frac{\partial}{\partial \alpha_K^{\mu\nu}} (\mu\nu - \tilde{\mu}\tilde{\nu}|m_{12}|\lambda\sigma - \tilde{\lambda}\sigma) \quad (3)$$

$$= -(K|m_{12}|\lambda\sigma) + (K|m_{12}|J)\alpha_J^{\lambda\sigma} \quad (4)$$

So by solving for α

$$\alpha_J^{\lambda\sigma} = (J|m|K)^{-1}(K|m|\lambda\sigma) \quad (5)$$

one obtains eq. 1 as

$$(\mu\nu|\lambda\sigma) \approx \alpha_I^{\mu\nu}(I|J)\alpha_J^{\lambda\sigma} \quad (6)$$

We are now going to show eq. 1 again, but not from a DF perspective but by insertion and resolution of an identity. Consider the ERI $(\mu\nu|\lambda\sigma)$ as a chain of operations. One can insert two identities into it, obtaining

$$(\mu\nu|\mathbf{I}|\frac{1}{r}|\mathbf{I}|\lambda\sigma) . \quad (7)$$

Again, we will omit the $\frac{1}{r_{12}}$ operator and replace it with a vertical bar whenever an integral contains no other operators. A three center overlap integral is therefore denoted by $(\mu\nu\lambda)$.

Let us first only consider the identity to the right of $\frac{1}{r}$. Given a complete and orthogonal set $\{|\tilde{R}\rangle\}$ one can “resolve” this identity as

$$\mathbf{I} = |\tilde{R}\rangle\langle\tilde{R}| \quad (8)$$

so that

$$(\mu\nu|\lambda\sigma) = (\mu\nu|\tilde{R}\rangle\langle\tilde{R}|\lambda\sigma) \quad (9)$$

Given a complete, however non-orthogonal set $\{|I\rangle\}$ one can similarly “resolve” the identity. To find the explicit form we first note that

$$|I\rangle = |\tilde{R}\rangle\langle\tilde{R}I| \quad (10)$$

and so by completeness of the new set $\{|I\rangle\}$

$$|\tilde{R}\rangle = |I\rangle(I\tilde{R})^{-1} \quad (11)$$

where we have introduced the notational shorthand of $(I\tilde{R})^{-1}$ meaning the (I, \tilde{R}) -th element of the inverse of the matrix $(\tilde{R}I)$ for easier readability.

If $\{|J\rangle\}$ is a complete set of functions $J(x)$ and $m_{12} = m(x_1, x_2)$ is a two-point function (which we call metric), in many cases

$$m_J(x) = \int J(x')m(x', x) dx' \quad (12)$$

forms again a complete set $\{|m_J\rangle\}$. Specifically this is the case for the choices $m(x_1, x_2) = \delta(x_1 - x_2)$ (overlap metric), $m(x_1, x_2) = \frac{1}{|x_1 - x_2|}$ (Coulomb metric), and $m(x_1, x_2) = \frac{\text{erfc}(w|x_1 - x_2|)}{|x_1 - x_2|}$ (attenuated Coulomb metric). With $\{|m_J\rangle\}$ being a complete set one can write in analogy to eq. 10

$$(m_J| = (J|m_{12}|\tilde{R})(\tilde{R}| \quad (13)$$

as

$$(\tilde{R}| = (\tilde{R}|m_{12}|J)^{-1}(m_J| \quad (14)$$

Expanding $|\tilde{R}\rangle$ in eq. 8 according to eqs. 11 and 14 respectively one obtains

$$\mathbf{I} = |I\rangle(I\tilde{R})^{-1}(\tilde{R}|m_{12}|J)^{-1}(m_J| \quad (15)$$

where for the inner matrix product the elements can be found as

$$(I\tilde{R})^{-1}(\tilde{R}|m_{12}|J)^{-1} = ((I|m_{12}|\tilde{R})(\tilde{R}J))^{-1} = (I|m_{12}|J)^{-1} \quad (16)$$

due to eq. 8. So concluding one finds

$$\mathbf{I} = |I\rangle(I|m_{12}|J)^{-1}(J|m_{12}| \quad (17)$$

or by exchanging $|I\rangle$ with $(I|$, etc. throughout the argument

$$\mathbf{I} = |m|I\rangle(I|m|J)^{-1}(J| \quad (18)$$

Inserted into eq. 7 this yields the final formulation, eq. 1, which we take as the starting point in the main text (eq. 3 and following).

V. OUTLOOK: KERNEL PROJECTION FOR EXPLICITLY CORRELATED METHODS

While the main text is focused solely on RI in the context of ERIs, for which it is equivalent to DF, in the following we briefly outline some related ideas for use in another type of RI, as it applies to explicitly correlated (F12) methods (c.f. [W. Klopper, C.C.M. Samson: J. Chem. Phys. 116 (2002), 6397] and references therein). There, integrals of the form

$$\langle xy|F_{12}o_1g_{12}|ij\rangle \quad (19)$$

occur where F_{12} and g_{12} are operators acting on both electrons, while o_1 acts on the first electron only, projecting it onto the occupied space, while the second electron is effectively acted on by an identity operation

$$o_1 = |k\rangle \langle k| \otimes \mathbf{I} \quad (20)$$

Writing eq. 19 in explicit integral form makes it obvious that this is in fact a three-electron integral

$$\langle xy|F_{12}o_1g_{12}|ij\rangle = \quad (21)$$

$$= \int \langle x(r)y(r')|F_{12}(r,r')|k(r)\rangle \langle k(r'')|g_{12}(r'',r')|i(r'')j(r')\rangle dr' \quad (22)$$

$$= \iiint x(r)y(r')k(r'')F_{12}g_{32}k(r)j(r')i(r'')drdr'dr'' \quad (23)$$

$$= \langle xyk|F_{12}g_{32}|kji\rangle \quad (24)$$

However, within the RI the identity in eq. 20 is resolved according to eq. 8. For this outlook we disregard, that commonly the complementary auxiliary basis set (CABS, [E.F. Valeev: Chem. Phys. Lett. 395 (2004), 190]) approach is used, by which certain contributions are already projected out of the otherwise formally complete auxiliary set. Within RI eq. 20 then becomes

$$o_1 = |kI\rangle \langle kI| \quad (25)$$

and therefore

$$\langle xy|F_{12}o_1g_{12}|ij\rangle = \langle xy|F_{12}|kI\rangle \langle kI|g_{12}|ij\rangle \quad (26)$$

which is a sum of products of two-electron integrals.

With $\langle xy|F_{12}|kI\rangle$ being linear in each index (specifically I) one could consider it as a linear operator $F_{xyk,I}$ and build an auxiliary matrix as

$$H_{IJ} = F_{I,xyk}^\dagger F_{xyk,J} \quad (27)$$

This would suggest that the projector P^R constructed from H (following the scheme of the main text) would not change the result when introduced as

$$\langle xy|F_{12}|kI\rangle \langle kI|g_{12}|ij\rangle = \langle xy|F_{12}|kI\rangle P_{IJ}^R \langle kJ|g_{12}|ij\rangle \quad (28)$$

How well this approach would perform in terms of auxiliary reduction compared to overhead costs caused by the projector construction remains to be seen and will have to be studied in future work.

Article III

“Computation of indirect nuclear spin–spin couplings with reduced complexity in pure and hybrid density functional approximations”,

A. Luenser, J. Kussmann, C. Ochsenfeld,

J. Chem. Phys. **145**, 124103 (2016).

Computation of indirect nuclear spin–spin couplings with reduced complexity in pure and hybrid density functional approximations

Arne Luenser, Jörg Kussmann, and Christian Ochsenfeld^{a)}

Chair of Theoretical Chemistry, Department of Chemistry, University of Munich (LMU Munich) and Center for Integrated Protein Science Munich (CIPSM, LMU Munich), Butenandtstr. 5–13, D-81377 Munich, Germany

(Received 8 May 2016; accepted 18 August 2016; published online 22 September 2016)

We present a (sub)linear-scaling algorithm to determine indirect nuclear spin–spin coupling constants at the Hartree–Fock and Kohn–Sham density functional levels of theory. Employing efficient integral algorithms and sparse algebra routines, an overall (sub)linear scaling behavior can be obtained for systems with a non-vanishing HOMO–LUMO gap. Calculations on systems with over 1000 atoms and 20 000 basis functions illustrate the performance and accuracy of our reference implementation. Specifically, we demonstrate that linear algebra dominates the runtime of conventional algorithms for 10 000 basis functions and above. Attainable speedups of our method exceed 6× in total runtime and 10× in the linear algebra steps for the tested systems. Furthermore, a convergence study of spin–spin couplings of an aminopyrazole peptide upon inclusion of the water environment is presented: using the new method it is shown that large solvent spheres are necessary to converge spin–spin coupling values. *Published by AIP Publishing.* [<http://dx.doi.org/10.1063/1.4962260>]

I. INTRODUCTION

Nuclear magnetic resonance spectroscopy (NMR) is an important tool to determine chemical structures in liquid or solid state. In recent years, a particular focus has been on the structural investigation of large biochemical systems like proteins.^{1–4} While theoretical prediction of molecular magnetic properties like NMR shieldings, spin-spin couplings, hyperfine couplings, or *g*-tensors using *ab initio* methods (see, e.g., Refs. 5–8 for overviews) is routine nowadays, the unfavorable scaling behavior of conventional algorithms prevents the application to large molecular systems. For self-consistent field (SCF) methods, computational complexity is $O(N^3 N_{\text{pert}})$, where N is a measure for system size, e.g., number of atoms, and N_{pert} is the number of perturbations.

In the last decade, linear-scaling methods for Hartree–Fock (HF) and Kohn–Sham density functional (KS–DFT) theories have been introduced which allow the theoretical investigation of large chemical systems with more than 1000 atoms (see, e.g., Ref. 9 for a recent review). Considering magnetic properties, linear-scaling methods^{10–14} allow the investigation of large chemical and biochemical systems.^{10,15,16}

The goal of the present work is to reduce the computational scaling for calculating indirect nuclear spin–spin coupling constants (ISSCs) at the density functional level of theory (DFT). Conventionally, the scaling is $O(N^3)$ per perturbing nucleus. Most ISSCs in larger molecules are vanishingly small and thus inaccessible experimentally.¹⁷ Ideally, computation of a single ISSC would scale as $O(1)$, i.e., be asymptotically independent of system size. See also our related work on nuclei-selected NMR with $O(1)$ scaling.¹²

We will review the required theory in Sections II A and II B. The obstacles to be overcome and our approach are

described in Sections II C through II D. Sections III and IV show exemplary calculations demonstrating the computational scaling. Finally, Section V summarizes our results.

II. THEORY

We will briefly review the underlying physics for the *ab initio* calculation of indirect spin-spin coupling constants and discuss the strategies to reduce the scaling behavior of the occurring terms. For a more detailed description of the theory we refer the reader to several excellent reviews on the theory of magnetic interactions in molecules.^{5–8,18}

A. Intramolecular magnetic interactions

Indirect nuclear spin–spin coupling refers to the coupling mechanism of the nuclear magnetic moments \mathbf{m} in a molecular system, which is *not* described by the classical through-space magnetic dipole–dipole interaction.¹⁹ Instead, the indirect coupling is conveyed by three mechanisms. They are the spin–orbit (SO), spin–dipole (SD), and Fermi contact (FC) interactions. In Ramsay’s theory of indirect spin–spin coupling,^{20,21} four operators describe these interactions (the spin–orbit coupling has a diamagnetic [DSO] and a paramagnetic [PSO] part)^{5,6}

$$\hat{h}_A^{\text{FC}} = -\frac{8\pi\alpha^2}{3} \sum_j \delta(\mathbf{r}_{jA}) \mathbf{m}_j, \quad (1)$$

$$\hat{h}_A^{\text{SD}} = \alpha^2 \sum_j \frac{r_{jA}^2 \mathbf{m}_j - 3(\mathbf{m}_j \cdot \mathbf{r}_{jA}) \mathbf{r}_{jA}}{r_{jA}^5}, \quad (2)$$

$$\hat{h}_A^{\text{PSO}} = \alpha^2 \sum_j \frac{-i\mathbf{r}_{jA} \times \nabla_j}{r_{jA}^3}, \quad (3)$$

^{a)}christian.ochsenfeld@uni-muenchen.de

$$\hat{h}_{AB}^{\text{DSO}} = \frac{\alpha^4}{2} \sum_j \frac{(\mathbf{r}_{jA} \cdot \mathbf{r}_{jB}) \mathbf{1} - \mathbf{r}_{jA} \mathbf{r}_{jB}^T}{r_{jA}^3 r_{jB}^3}, \quad (4)$$

where $i = \sqrt{-1}$ is the imaginary unit, j indexes electrons, A, B index nuclei, and α is the fine structure constant. In practical calculations, ISSCs are obtained by differentiation of the electronic energy with respect to the magnetic moments of the coupling nuclei

$$\mathcal{J}_{AB} = h \frac{\gamma_A \gamma_B}{4\pi^2} \mathcal{K}_{AB} = h \frac{\gamma_A \gamma_B}{4\pi^2} \frac{d^2 E}{d\mathbf{m}_A d\mathbf{m}_B}, \quad (5)$$

where h is Planck's constant and γ_A, γ_B are the gyromagnetic ratios of nuclei A, B . See Ref. 22 for a discussion of the units and the reduced coupling tensor \mathcal{K} . Most often, only the isotropic average $\mathcal{J}_{AB} = \frac{1}{3} \text{Tr}(\mathcal{J}_{AB})$ of the coupling tensor is reported.

B. Molecular properties in Hartree–Fock and Kohn–Sham DFT

In a density matrix form, the SCF energy for a restricted wave function is given by

$$E_{\text{SCF}} = \text{Tr} \left(\mathbf{P} \mathbf{h} + \frac{1}{2} \mathbf{P} \mathbf{G}[\mathbf{P}] \right), \quad (6)$$

$$G_{\mu\nu}[\mathbf{P}] = J_{\mu\nu}[\mathbf{P}] - \frac{1}{2} \gamma K_{\mu\nu}[\mathbf{P}] + V_{\mu\nu}^{\text{XC}}[\mathbf{P}], \quad (7)$$

$$J_{\mu\nu}[\mathbf{P}] = \sum_{\lambda\sigma} (\mu\nu|\lambda\sigma) P_{\lambda\sigma}, \quad (8)$$

$$K_{\mu\nu}[\mathbf{P}] = \sum_{\lambda\sigma} (\mu\lambda|\nu\sigma) P_{\lambda\sigma}, \quad (9)$$

where \mathbf{h} is the one-electron matrix, \mathbf{P} is the one-particle density matrix, and \mathbf{V}^{XC} is the Kohn–Sham exchange–correlation (XC) matrix, and $\mu, \nu, \lambda, \sigma$ are AO basis functions. For Hartree–Fock, \mathbf{V}^{XC} vanishes and $\gamma = 1$. For pure DFT, $\gamma = 0$. For hybrid DFT, $0 < \gamma \leq 1$.

The second derivative of the electronic energy with respect to the magnetic moments $\mathbf{m}_A, \mathbf{m}_B$ of nuclei A, B yields the indirect nuclear spin–spin coupling tensor between those nuclei,

$$\frac{d^2 E_{\text{SCF}}}{d\mathbf{m}_A d\mathbf{m}_B} = \text{Tr}(\mathbf{P}^{\mathbf{m}_A} \mathbf{h}^{\mathbf{m}_B} + \mathbf{P} \mathbf{h}^{\mathbf{m}_A \mathbf{m}_B}), \quad (10)$$

where \mathbf{A}^x is short for $\frac{d}{dx} \mathbf{A}$. We also used the fact that the basis functions do not depend on the magnetic moments. The perturbed one-electron matrices are given by

$$h_{\mu\nu}^{\mathbf{m}_B} = \langle \mu | \hat{h}_B^{\text{FC}} + \hat{h}_B^{\text{SD}} + \hat{h}_B^{\text{PSO}} | \nu \rangle, \quad (11)$$

$$h_{\mu\nu}^{\mathbf{m}_A \mathbf{m}_B} = \langle \mu | \hat{h}_{AB}^{\text{DSO}} | \nu \rangle. \quad (12)$$

The total coupling tensor comprises the following contributions:

$$\mathcal{J}_{AB} = \mathcal{J}_{AB}^{\text{FC}} + \mathcal{J}_{AB}^{\text{SD}} + \mathcal{J}_{AB}^{\text{SDFC}} + \mathcal{J}_{AB}^{\text{PSO}} + \mathcal{J}_{AB}^{\text{DSO}}, \quad (13)$$

$$\mathcal{J}_{AB}^{\text{FC}} = h \frac{\gamma_A \gamma_B}{4\pi^2} \text{Tr}(\mathbf{P}_A^{\text{FC}} \mathbf{h}_B^{\text{FC}}), \quad (14)$$

$$\mathcal{J}_{AB}^{\text{SD}} = h \frac{\gamma_A \gamma_B}{4\pi^2} \text{Tr}(\mathbf{P}_A^{\text{SD}} \mathbf{h}_B^{\text{SD}}), \quad (15)$$

$$\mathcal{J}_{AB}^{\text{SDFC}} = h \frac{\gamma_A \gamma_B}{4\pi^2} \text{Tr}(\mathbf{P}_A^{\text{SD}} \mathbf{h}_B^{\text{FC}} + \mathbf{P}_A^{\text{FC}} \mathbf{h}_B^{\text{SD}}), \quad (16)$$

$$\mathcal{J}_{AB}^{\text{PSO}} = h \frac{\gamma_A \gamma_B}{4\pi^2} \text{Tr}(\mathbf{P}_A^{\text{PSO}} \mathbf{h}_B^{\text{PSO}}), \quad (17)$$

$$\mathcal{J}_{AB}^{\text{DSO}} = h \frac{\gamma_A \gamma_B}{4\pi^2} \text{Tr}(\mathbf{P} \mathbf{h}_{AB}^{\text{DSO}}). \quad (18)$$

The SD and FC operators are real triplet operators, while the PSO operator is a purely imaginary singlet operator. Because singlet and triplet states do not mix, there are no PSO-FC or PSO-SD contributions. The mixed SD-FC contribution is traceless, and does not contribute to the isotropic coupling. The calculation of ISSCs within DFT was first discussed by Fukui,²³ later by Malkin *et al.*²⁴ The first complete implementation (no terms omitted) within KS-DFT was given in Refs. 25–28. The decay behavior of the individual contributions has been studied in Ref. 17.

C. Coupled perturbed SCF

The perturbed density matrices $\mathbf{P}^{\mathbf{m}_A}$ are the solutions to the coupled perturbed SCF (CPSCF) equations.^{29–31} Integral-direct molecular orbital (MO)-CPSCF³² computes two-electron integral contractions on-the-fly in the atomic orbital (AO) basis, and transforms to the canonical molecular orbital (MO) basis for solving the CPSCF equations

$$U_{ai}^x = \frac{S_{ai}^x \varepsilon_i - F_{ai}^x}{\varepsilon_a - \varepsilon_i}, \quad (19)$$

where i indexes occupied orbitals, a indexes virtuals, and x is a general perturbation. The perturbed density matrix \mathbf{P}^x is given by

$$P_{\mu\nu}^x = \sum_{ai} C_{\mu a} U_{ai}^x C_{i\nu}^\dagger + \sum_{ia} C_{\mu i} U_{ia}^x C_{a\nu}^\dagger - \sum_{\lambda\sigma} P_{\mu\lambda} S_{\lambda\sigma}^x P_{\sigma\nu}. \quad (20)$$

Since the canonical MO basis is delocal, the matrix \mathbf{C} of MO coefficients is densely populated. The number of significant elements therefore grows as $\mathcal{O}(N^2)$ with the system size N . The matrix multiplications of AO-MO transformations in Eqs. (19) and (20) thus scale as $\mathcal{O}(N^3)$, regardless of whether the perturbation is local or not.

To reduce scaling to linear or sub-linear, we use the density matrix-based Laplace-transformed CPSCF (DL-CPSCF) equations,^{12,14} which are formulated entirely in the AO basis

$$\mathbf{P}_{\text{vo}}^x = \sum_{\alpha=1}^{\tau} \mathbf{Q}^{(\alpha)} (-\mathbf{h}^x - \mathbf{G}[\mathbf{P}^x]) \mathbf{P}^{(\alpha)} \quad (21)$$

$$= \sum_{\alpha=1}^{\tau} \tilde{\mathbf{Q}}^{(\alpha)} \tilde{\mathbf{Q}}^{(\alpha)\dagger} (-\mathbf{h}^x - \mathbf{G}[\mathbf{P}^x]) \tilde{\mathbf{P}}^{(\alpha)} \tilde{\mathbf{P}}^{(\alpha)\dagger}, \quad (22)$$

$$\mathbf{P}^{(\alpha)} = \sqrt{w_\alpha} \exp(t_\alpha \mathbf{P} \mathbf{F}) \mathbf{P} = \tilde{\mathbf{P}}^{(\alpha)} \tilde{\mathbf{P}}^{(\alpha)\dagger}, \quad (23)$$

$$\mathbf{Q}^{(\alpha)} = \sqrt{w_\alpha} \exp(-t_\alpha \mathbf{Q} \mathbf{F}) \mathbf{Q} = \tilde{\mathbf{Q}}^{(\alpha)} \tilde{\mathbf{Q}}^{(\alpha)\dagger}. \quad (24)$$

In contrast to the MO coefficient matrix \mathbf{C} , the occupied ($\mathbf{P}^{(\alpha)}$) and virtual ($\mathbf{Q}^{(\alpha)}$) pseudo-density matrices are sparse for large systems. $\tilde{\mathbf{P}}^{(\alpha)}$ and $\tilde{\mathbf{Q}}^{(\alpha)}$ are the factors from a pivoted Cholesky decomposition of $\mathbf{P}^{(\alpha)}$ and $\mathbf{Q}^{(\alpha)}$, respectively.³³ The required number of roots τ is typically 5–10. We have omitted all terms involving basis function derivatives with respect to nuclear magnetic moments from the full DL-CPSCF

equations because they vanish for the perturbations discussed in this work. The full perturbed density matrix is thus simply $\mathbf{P}^x = \mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x$ with $\mathbf{P}_{\text{ov}}^x = (\mathbf{P}_{\text{vo}}^x)^\dagger$.

The efficient evaluation of the occurring quantities as well as the (sub)linear-scaling solution of the linear equation system in Eq. (20) is discussed in Section II D. Furthermore, it is also possible to reduce the prefactor of the DL-CPSCF algorithm as described in Sec. II C 1.

1. Reducing the computational prefactor by Cholesky decomposition

As discussed in Ref. 33, performing a Cholesky decomposition with complete pivoting on the occupied (Eq. (23)) and/or virtual (Eq. (24)) pseudo-density matrices can reduce the computational prefactor of the DL-CPSCF method. The time savings are more pronounced the larger the basis set is, because the complexity of evaluating Eq. (21) can be reduced from $O(N_{\text{basis}}^3)$ to $O(N_{\text{basis}}^2 N_{\text{occ}})$ in the limiting case where all matrices are densely filled.

Performing the decomposition on the occupied pseudo-densities reveals their numerical rank, which is N_{occ} or smaller, and reduces computational complexity hugely when $N_{\text{occ}} \ll N_{\text{basis}}$. Decomposition of the virtual counterparts is often very beneficial as well. Beyond rank reduction to N_{virt} , the numerical rank of the matrices can be much smaller. Pseudo-densities corresponding to Laplace points with large roots t_α often have a numerical rank one order of magnitude smaller than N_{virt} . This is obvious from Eq. (24), where $-t_\alpha$ is an argument to the exponential function. Consequently, Eq. (22) is most economical in our experience. We first multiply by the left Cholesky factor of the occupied pseudo-density matrix, then by the two virtual Cholesky factors, and finally by the right occupied Cholesky factor. For the largest molecule studied in Sec. III B 2, the numerical rank of the most sparse virtual pseudo-density matrix $\mathbf{Q}^{(\alpha)}$ is 2455 while the number of virtual orbitals is 19 160, an 87% reduction in rank.

D. (Sub)linear-scaling algorithms

We restrict ourselves to closed-shell, spin-unpolarized systems in this work. Because the operator matrix representations (Eq. (11)) involve only one-electron integrals, their computation is extremely rapid and consumes negligible overall time. This leaves three time-determining steps in the computation of the FC, PSO, and SD contributions. These are the construction of $\mathbf{K}[\mathbf{P}^x]$ (only when using hybrid density functional approximations), construction of $\mathbf{V}^{\text{XC}}[\mathbf{P}^x]$, and solution of Eq. (21), i.e., matrix multiplications.

1. (Sub)linear-scaling sparse algebra

If the matrices \mathbf{h}^x , $\mathbf{G}[\mathbf{P}^x]$, and \mathbf{P}^x in Eq. (21) are local, i.e., the number of significant elements is asymptotically constant with respect to system size, the matrix multiplications can be performed with $O(1)$ scaling. Our implementation (also detailed in Refs. 12 and 33) uses a modified blocked compressed sparse matrix format (BCSR).^{11,34–36} Matrices are

blocked into submatrices of $\sim 100 \times 100$ elements, which are only allocated in memory if their Frobenius norm exceeds a threshold, typically 10^{-6} . Obtaining compact matrix representations for sparse matrices in BCSR format requires the atoms be sorted by spatial proximity. We do this by applying the Cuthill–McKee algorithm to the atom connectivity matrix.^{37,38}

2. The exact-exchange matrix $\mathbf{K}[\mathbf{P}^x]$

For hybrid density functional approximations, we calculate the exact-exchange contribution

$$K_{\mu\nu}[\mathbf{P}^x] = \sum_{\lambda\sigma} (\mu\lambda|\nu\sigma) P_{\lambda\sigma}^x \quad (25)$$

using the LinK algorithm.^{39,40} Due to the coupling of bra- and ket-charge distributions via the local response density \mathbf{P}^x , the computational effort is asymptotically independent of system size for a single perturbation. Furthermore, our implementation directly scatters the resulting exchange integrals into sparse matrix format, thus also ensuring $O(1)$ scaling with respect to memory usage.

Increasing the fraction of HF exchange in the density functional approximation promotes the development of triplet instabilities in the wave function.⁴¹ This can strongly impede the quality of predictions, which is why pure density functionals may be preferable to hybrid functionals for increased robustness.

3. The exchange-correlation potential matrix $\mathbf{V}^{\text{XC}}[\mathbf{P}^x]$

The central step within the evaluation of the exchange-correlation (XC) potential and its derivatives is the calculation of the XC-energy and its derivatives based on the ground-state density $\rho(\mathbf{r})$. Using local basis functions, linear scaling behavior is possible.^{42,43}

In order to reduce the scaling to $O(1)$ when using local response densities \mathbf{P}^x , a preselection of those basis function pairs where \mathbf{P}^x has significant values is necessary. Since only a constant number of local basis functions remains, the functional values as well as the real-space representations of the densities have to be evaluated for a constant number of grid points in order to obtain the local response potential $\mathbf{V}^{\text{XC}}[\mathbf{P}^x]$.

4. The FC and SD contributions

The FC and SD contributions require evaluation of both $\mathbf{V}^{\text{XC}}[\mathbf{P}^x]$ and $\mathbf{K}[\mathbf{P}^x]$ (hybrid functionals only). The Coulomb-type two-electron integral contraction $\mathbf{J}[\mathbf{P}^x]$ vanishes because the response densities of α - and β -spin cancel for closed-shell triplet operators.

The FC operator (Eq. (1)) is extremely local with an $\exp(-r^2)$ decay in a Gaussian basis. The operator is isotropic and requires only a single response density matrix per perturbing nucleus. Figure 1 shows the sparsity pattern of the resulting converged response density matrix in the AO basis. The number of significant elements in the matrix is




FIG. 1. Sparsity patterns for perturbed density matrices obtained from DL-CPSCF for ISSCs in amylose₁₆ (PBE/pcJ-1). Darker areas mean higher absolute values of the matrix elements. White areas are numerically insignificant ($<10^{-5}$) and not allocated in memory. Shown are the perturbed density matrices for the FC (left), SD (center), and PSO (right) perturbations.

independent of the molecular size. Evaluation of the matrix in constant time will be demonstrated in Sec. III.

The SD operator (Eq. (2)) decays as r^{-3} , and requires six independent perturbations to be evaluated per perturbing nucleus (cf. Ref. 25). The sparsity of the SD operator's AO representation can be significantly enhanced by projection onto the virtual-occupied subspace before performing CPSCF. An exemplary sparsity pattern of a corresponding response density is given in Fig. 1. As with the FC operator, the electronic response to the SD perturbation is spatially local and does not further increase with system size beyond a certain point.

5. The PSO contribution

The PSO (Eq. (3)) operator is longer ranged at r^{-2} than both FC and SD. Three independent perturbations must be computed per perturbing nucleus. The XC potential matrix $\mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{PSO}}]$, as well as the Coulomb matrix $\mathbf{J}[\mathbf{P}^{\text{PSO}}]$, vanishes due to the purely imaginary Hermitian (i.e., skew-symmetric) nature of the operator. For pure density functional approximations (no exact-exchange contribution), this perturbation can therefore be evaluated non-iteratively using Eq. (22).

While the PSO operator is the longest-ranged of the three (see Fig. 1 for a sample sparsity pattern of the resulting response density), the evaluation is still extremely rapid in pure DFT calculations. This operator has the worst scaling exponent of the three, allowing for just below linear computational complexity for the systems considered in Sec. III. The absolute computational time, however, is very small. In hybrid DFT calculations, where $\mathbf{K}[\mathbf{P}^{\text{PSO}}]$ must be computed, this perturbation becomes the dominating contribution.

Like the SD operator, we project the PSO operator's AO representation to the virtual-occupied subspace before CPSCF. This enhances sparsity slightly and is beneficial for performance of large molecules.

6. The DSO contribution

The DSO contribution is calculated directly as an expectation value with the ground state density. In our

implementation, like in most others,^{25,26} it is computed numerically using existing grid-based numerical quadrature routines for KS-DFT. The expectation value $\mathcal{J}_{AB}^{\text{DSO}}$ is obtained by contracting the real-space representations of the DSO operator (Eq. (4)) and the density on the grid

$$\mathcal{J}_{AB}^{\text{DSO}} = \int \hat{h}_{AB}^{\text{DSO}}(\mathbf{r}) \rho(\mathbf{r}) d^3\mathbf{r} \approx \sum_g w_g \hat{h}_{AB}^{\text{DSO}}(\mathbf{r}_g) \rho(\mathbf{r}_g), \quad (26)$$

with grid points \mathbf{r}_g and weights w_g .

Thus, this step shows a very small prefactor as compared to the other contributions of the total ISSC tensor and will not dominate the overall computation time even for very large systems. Because its evaluation is so rapid, we will forego further discussion of this term.

III. ILLUSTRATIVE CALCULATIONS

Calculations were performed using the PBE functional^{44,45} (as recommended by Lutnæs⁴¹) and HF theory, using Jensen's pcJ-1 basis set designed specifically for spin-spin coupling calculations,^{46,47} as deposited in the EMSL basis set exchange.⁴⁸⁻⁵⁰ While the HF approximation is not recommended for the calculation of spin-spin coupling constants, recent work by Cheng *et al.*⁵¹ shows impressive improvements by employing the Tamm-Dancoff approximation. In the present context we employ the HF calculations for demonstrating the scaling behavior of hybrid density functional approximations containing exact exchange. We implemented our method in a development version of Q-Chem.⁵² Computations were performed on an Intel Xeon E5-2620 machine, using no symmetry and a single thread of execution. Structure files for the molecules can be found online.⁵³

A. Accuracy

To find the most economical thresholds and to assess the accuracy of our methods, we performed a series of calculations with varying sparsity thresholds (Frobenius norm) and integral thresholds. Our reference molecule was the adenine-thymine DNA base pair (62 atoms, 1215 basis functions). In each

case, we calculated the SD, FC, and PSO contributions to the spin–spin coupling values between one hydrogen nucleus and all other nuclei in the molecule, as well as between one carbon nucleus and all other nuclei. We consider settings to be accurate if the maximum absolute deviation in the final expectation value of all those ^1H – ^1H and ^1H – ^{13}C couplings differs from the reference value by less than 1% or less than 0.1 Hz.

First, the number of Laplace points (we use the minimax fitting of Takatsuka *et al.*⁵⁴) was fixed to seven in all calculations. Reference values were computed with MO-CPSCF using an Euler–Maclaurin/Lebedev grid with 99 radial and 590 spherical points, an integral threshold of 10^{-12} , and dense linear algebra. Second, we studied the influence of the DFT grid on accuracy. In our case, the SG-0 standard grid⁵⁵ gave sufficiently accurate (according to the criterion above) results compared to a 99-590 Euler–Maclaurin–Lebedev grid, using an integral threshold of 10^{-12} in both cases.

Finally, we searched for the most economical combination of sparsity thresholds and integral thresholds for DL-CPSCF. The reference in this case was an MO-CPSCF calculation using dense linear algebra and an integral threshold of 10^{-12} . We obtained sufficiently accurate (*vide supra*) results with an integral threshold of 10^{-8} and a sparsity threshold of 10^{-5} (Frobenius norm). The $\mathbf{V}^{\text{XC}}[\mathbf{P}^x]$ screening described in Sec. II D 3 is based on shell pairs. If no element in \mathbf{P}^x belonging to a particular shell pair exceeds 10^{-6} (absolute value), that shell pair is eliminated from the evaluation of \mathbf{V}^{XC} . Otherwise, the shell pair is considered significant, and all constituent basis functions are used in the potential matrix evaluation. We confirmed the accuracy of our final thresholds on the DNA₈ molecule (524 atoms, 10 770 basis functions).

To facilitate easy comparison, the number of CPSCF iterations was fixed to seven in all cases below, which is easily sufficient for the given accuracy. All times given below are

wall times, and correspond to the cumulative times of all seven iterations.

B. Pure density functionals

1. Amylose chains

We computed individual ISSCs on amylose chains of increasing length using the PBE density functional approximation. The perturbing nucleus is located near the end of the molecule in each case. Detailed wall times and scaling exponents are given in Table I.

The data show the most expensive steps of the calculation. They are (in decreasing order of cost for the largest system) evaluation of the XC response potential and solution of Eq. (22) for the SD perturbation (corresponding to 58% and 19% of the total time, respectively). Evaluation of $\mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{FC}}]$ accounts for 13%, and solution of Eq. (22) for the PSO contribution accounts for 6% of the overall time.

The FC contribution is extremely local because of its exponential decay. Correspondingly, it shows sub-linear scaling behavior even between the smaller systems. Because of its high sparsity and compact matrix representation, the DL-CPSCF iteration step (Eq. (22)) is extremely rapid for the FC perturbation even for the largest systems.

The scaling exponent for the computation of $\mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{FC}}]$ increases from 0.2 to 1.0 for the larger systems. This is a consequence of our implementation using the pre-existing DFT algorithms in Q-Chem, which can also be observed for the DNA systems studied in Sec. III B 2. Specifically, the converged perturbed densities evaluated on the real-space DFT grid have virtually the same number of significant grid points, basis functions, and basis function pairs from amylose₄ through amylose₃₂. The extent of the perturbative effect has decayed below numerical accuracy. Our screening procedure

TABLE I. Wall times (in seconds) for the calculation of an individual ISSC in amylose chains (PBE/pcJ-1). The columns headed $O(N^x)$ contain the scaling exponents between the systems to the left and right of that column. Exponents and sums calculated from unrounded numbers.

		Amylose ₄	$O(N^x)$	Amylose ₈	$O(N^x)$	Amylose ₁₆	$O(N^x)$	Amylose ₃₂
N_{atoms}		87		171		339		675
$N_{\text{electrons}}$		354		698		1386		2762
N_{basis}		1635		3223		6399		12751
DL-CPSCF	SD $\left\{ \begin{array}{l} \mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{SD}}] \\ \text{Eq. (22)} \end{array} \right.$	238	0.2	267	0.2	305	0.7 ^a	508
	FC $\left\{ \begin{array}{l} \mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{FC}}] \\ \text{Eq. (22)} \end{array} \right.$	39	0.2	46	0.4	60	1.0 ^a	119
	PSO Eq. (22)	5	1.0	10	0.2	12	0.6	19
	Σ	5	1.7	15	1.1	32	0.8	56
	Σ	337		424		524		868
	MO-CPSCF	SD $\left\{ \begin{array}{l} \mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{SD}}] \\ \text{Eqs. (19) and (20)} \end{array} \right.$	222	0.4	286	0.3	347	0.7 ^a
FC $\left\{ \begin{array}{l} \mathbf{V}^{\text{XC}}[\mathbf{P}^{\text{FC}}] \\ \text{Eqs. (19) and (20)} \end{array} \right.$		9	2.9	67	3.0	508	3.0	3879
PSO Eq. (22)		41	0.3	50	0.4	63	0.9 ^a	119
Σ		2	2.9	11	2.9	84	3.0	650
Σ		1	3.0	5	2.9	36	3.0	280
Σ		274		419		1038		5474

^aSee text for a detailed discussion.




FIG. 2. Sparsity patterns of several matrices for DNA₁₆ (PBE/pcJ-1). Darker areas mean higher absolute values of the matrix elements. White areas are numerically insignificant ($<10^{-5}$) and not allocated in memory. Shown (from left to right) are the ground state density matrix, an occupied pseudo-density Cholesky factor, as well as the least and most sparse virtual pseudo-density Cholesky factors.

within the evaluation of the XC matrix correctly recognizes this, and only evaluates those basis functions on the current grid batch which have significant elements in \mathbf{P}^x . However, the DFT numerical integration procedure currently available in Q-Chem performs several steps adversely affecting the scaling, which cannot be circumvented without rewriting a large part of the implementation. These steps do not surface in the evaluation of the regular Kohn–Sham potential matrix during SCF, but do in fact become dominant steps during the evaluation of local perturbations in systems with many basis functions as is the case here. Although these steps can be avoided, we have opted not to rewrite the corresponding code, but to discuss this shortcoming by showing the actual wall times in all cases, even though the resulting scaling exponents arguably understate the power of the DL-CPSCF method in the present context.

The SD contribution is the most expensive one for two reasons. First, six independent density matrices must be computed for this perturbation (cf. Sec. II D 4), each

of which is less sparse than the FC-perturbed density. The computational time grows sub-linearly for both the evaluation of the XC matrix and the solution of the DL-CPSCF equations.

Finally, the computational complexity of the PSO contribution decreases with larger systems. While scaling is nearly quadratic between the two small systems, it has reduced to below linear at $O(N^{0.7})$ for the largest systems. This is consistent with its slower decay behavior of r^{-2} , being the longest-ranged of the three contributions. The interaction has not decayed below the matrix sparsity threshold for these systems to enable truly constant scaling, but steadily reduces in complexity with molecular size. Nevertheless, the computational prefactor of the PSO evaluation is very small due to the vanishing response potential, i.e., one only has to solve the uncoupled perturbed KS-equations in a non-iterative manner.

Table I also contains timing data for regular MO-CPSCF as a reference. The routines for evaluation of the XC response potential are the exact same ones as used in the DL-CPSCF

TABLE II. Wall times (in seconds) for the calculation of an individual ISSC in DNA double strands (PBE/pcJ-1). The columns headed $O(N^x)$ contain the scaling exponents between the systems to the left and right of that column. Exponents and sums calculated from unrounded numbers.

		DNA ₁	DNA ₂	DNA ₄	$O(N^x)$	DNA ₈	$O(N^x)$	DNA ₁₆
	N_{atoms}	62	128	260		524		1 052
	$N_{\text{electrons}}$	260	580	1220		2 500		5 060
	N_{basis}	1215	2580	5310		10 770		21 690
DL-CPSCF	SD { $\mathbf{v}^{\text{XC}}[\mathbf{P}^{\text{SD}}]$	115	390	813	0.2	958	0.7 ^a	1 602
	Eq. (22)	22	148	485	1.0	986	0.4	1 312
	FC { $\mathbf{v}^{\text{XC}}[\mathbf{P}^{\text{FC}}]$	24	69	101	0.3	125	1.1 ^a	279
	Eq. (22)	3	16	49	1.0	99	0.6	148
	PSO Eq. (22)	2	16	90	2.0	362	1.0	748
	Σ	165	638	1538		2 531		4 090
MO-CPSCF	SD { $\mathbf{v}^{\text{XC}}[\mathbf{P}^{\text{SD}}]$	104	337	664	0.7	1 059	0.7 ^a	1 752
	Eqs. (19) and (20)	4	36	302	3.0	2 477	2.8	17 771
	FC { $\mathbf{v}^{\text{XC}}[\mathbf{P}^{\text{FC}}]$	24	70	100	0.5	142	1.0 ^a	287
	Eqs. (19) and (20)	1	6	50	3.0	414	3.0	3 355
	PSO Eqs. (19) and (20)	0	3	22	3.0	178	3.0	1 494
	Σ	132	451	1136		4 269		24 659

^aSee text for a detailed discussion.




FIG. 3. Real-space representation of the Fermi contact-perturbed density on the numerical DFT grid for DNA₈ and DNA₁₆. The points colored in orange are employed for both calculations; those in purple are used only for DNA₁₆.

calculations. The matrix multiplications required in MO-CPSCF (Eqs. (19) and (20)) show the expected $O(N^3)$ complexity with a low prefactor. The crossover point after which DL-CPSCF outperforms MO-CPSCF lies at eight amylose units, beyond which DL-CPSCF is much faster. The overall speedup is 2× for amylose₁₆, and exceeds 6× for amylose₃₂.

2. DNA strands

We also performed calculations on adenine-thymine DNA double strands with up to 16 base pairs (1052 atoms) as a more representative example of a space-filling biochemical structure that exhibits a more delocalized electronic structure. The density matrices (ground-state as well as occupied and virtual pseudo-density matrices) for these systems will become sparse only for large molecular sizes. Figure 2 shows exemplary sparsity patterns for the

largest molecular system. Detailed timing data are given in Table II.

As for amylose chains, computational time for the FC contribution is the smallest. Scaling of the XC matrix construction is below linear between four and eight base pairs. The scaling behavior deteriorates when moving to sixteen base pairs, which is an artefact of our implementation using pre-existing DFT algorithms as discussed in more detail in Sec. III B 1. Figure 3 is a graphical representation of the DNA₈ and DNA₁₆ molecules (the translucent atoms in the lower half of the image exist only in the larger molecule), and the perturbed density from the Fermi contact operator. The density is shown on the real-space DFT grid used during numerical integration. All grid points used during the computation are shown. The perturbation is nearly identical for the two different molecules. The perturbed density is localized around the perturbing nucleus, and does not reach the far end of the molecule.

The SD contribution once again dominates computational time, constituting 78% of the total in DNA₁₆. The XC matrix construction scales as $O(N^{0.2})$ between four and eight base pairs, whereas the linear algebra step exhibits sub-linear scaling only between eight and sixteen base pairs. The PSO contribution scales linearly between the two largest systems, due to its long range. The computational prefactor is small, however, accounting for just 18% of the total time.

DL-CPSCF has significant performance benefits over MO-CPSCF. It is competitive in performance for small molecules, and significantly faster for large systems. The total speedup is 6× for DNA₁₆. Again, this is mostly due to the computational savings from using linear-scaling sparse algebra routines, which become dominant at around 10 000 basis functions.

C. Approximations including exact exchange

To study the scaling behavior of the exact-exchange contribution $\mathbf{K}[\mathbf{P}^x]$ using LinK,^{39,40} we employed HF theory to the same amylose systems as in Section III B 1. Detailed timings are given in Table III, indicating that considering exact exchange strongly increases the computational effort as compared to pure DFT calculations (cf. Sec. III B).

TABLE III. Wall times (in seconds) for the calculation of an individual ISSC in amylose chains (HF/pcJ-1). The columns headed $O(N^x)$ contain the scaling exponents between the systems to the left and right of that column. Exponents calculated from unrounded numbers.

	Amylose ₄	$O(N^x)$	Amylose ₈	$O(N^x)$	Amylose ₁₆
N_{atoms}	87		171		339
N_{basis}	1635		3 223		6 399
$\mathbf{K}[\mathbf{P}^{\text{PSO}}]$	3641	0.8	6 313	0.7	10 089
$\mathbf{K}[\mathbf{P}^{\text{SD}}]$	4070	0.2	4 496	-0.0	4 438
$\mathbf{K}[\mathbf{P}^{\text{FC}}]$	542	0.1	579	0.2	650
Eq. (22) (PSO)	43	1.8	148	1.2	336
Eq. (22) (SD)	60	1.1	124	0.2	147
Eq. (22) (FC)	7	1.1	15	0.3	18
Σ	8363		11 674		15 678

The PSO perturbation dominates the computational effort for the largest systems, being the longest-ranged. The complexity of the evaluation of $\mathbf{K}[\mathbf{P}^{\text{PSO}}]$ using LinK is slightly below linear with a large prefactor as compared to the other contributions, while solution of Eq. (22) scales slightly above linear with a small prefactor. The FC and SD perturbations behave very favorably, with near constant time complexity between each of the systems in the evaluation of $\mathbf{K}[\mathbf{P}^x]$, and near constant complexity in the solution of Eq. (22) between the two larger molecules. Computation of $\mathbf{K}[\mathbf{P}^{\text{SD}}]$ is in fact slightly faster for amylose₁₆ compared to amylose₈ which is due to the different orientations of the two molecules.

IV. EXAMPLE APPLICATION: CONVERGENCE OF ISSCS WITH SYSTEM SIZE IN A WATER ENVIRONMENT

In order to illustrate the usefulness of the new possibilities for studying large molecular systems, we studied the convergence of ISSCs with the size of an exemplary model system. We chose an aminopyrazole peptide from Ref. 15, embedded in spheres of water of increasing radius. We calculated all ^1H - ^1H and ^1H - ^{13}C couplings from the naked solvate (89 atoms in total, 2134 unique ^1H - ^1H and ^1H - ^{13}C couplings, PBE/pcJ-1). Of those, a subset of couplings with magnitudes above 0.1 Hz was selected as our test set, resulting in 145 ^1H - ^1H and 76 ^1H - ^{13}C couplings. These ISSCs were calculated with the solvate embedded in increasingly large shells of solvent water, using the pcJ-1 basis for the solvate and def2-SV(P) for the solvent. Our reference system contains the solvate, plus 246 water molecules, for a total of 827 atoms and 6083 basis functions.

Statistical data given in Table IV show that system sizes of over 500 atoms are required to converge ^1H - ^{13}C couplings to ~ 1 Hz. ^1H - ^1H couplings require above 350 atoms to converge below 0.5 Hz. This convergence behavior is similar to that of nuclear magnetic shieldings.¹⁵

Figure 4 shows the size convergence of one ^1H - ^1H and one ^1H - ^{13}C coupling (arbitrarily selected) as a function of the total number of atoms in the calculation. While the convergence behavior of other couplings may vary, it shows the systematic convergence. It has to be noted that the results for individual couplings typically do not improve monotonically, whereas the whole set of coupling constants statistically continuously improves the overall result upon increasing the solvent radius, as can be seen from Table IV.




FIG. 4. Convergence of one ^1H - ^1H and one ^1H - ^{13}C coupling of an aminopyrazole peptide with increasing solvent radius.

The convergence behavior in other systems may of course vary, however, the data illustrate the general importance to systematically study the convergence of observables with respect to increasing the solvent spheres or the considered QM spheres in general. Because of its superior scaling properties, the DL-CPSCF method allows to perform such studies and to describe large molecular systems.

V. SUMMARY

We have presented a method to calculate indirect nuclear spin-spin coupling tensors at the DFT level with reduced time complexity. For pure density functional approximations, the observed scaling is sub-linear per perturbation as compared to the cubic scaling of conventional algorithms. When exact exchange is incorporated, overall observed scaling is also below linear per perturbation, albeit with a later onset and higher prefactor. We exploit the locality of the interaction (r^{-2} , r^{-3} , or exponential decay) without any real-space cutoffs and do not neglect any of the operators. Our method has near-constant time complexity (i.e., computational time is asymptotically independent of system size) in the evaluation of the Fermi contact and spin-dipole contributions for both pure and hybrid density functional approximations where speedups for the largest test systems exceed 6 times. This includes evaluation of the exchange-correlation and exact-exchange response potential matrices as well as the solution of the DL-CPSCF equations. Although the computation of the paramagnetic spin-orbit contribution scales only slightly better than linear for our largest test systems, its prefactor

TABLE IV. Root mean square (RMS), mean absolute (MAD), and maximum absolute (max) deviations for ISSCs (in Hz) of an aminopyrazole peptide with different solvent environments.

	No. atoms	Naked	98	137	239	350	476	665
^1H	RMS	0.49	0.41	0.36	0.23	0.15	0.10	0.06
	MAD	0.20	0.18	0.15	0.10	0.06	0.04	0.02
	Max	1.66	1.21	1.33	1.20	0.78	0.38	0.36
^{13}C	RMS	2.16	1.94	1.77	1.03	0.59	0.41	0.29
	MAD	0.34	0.30	0.27	0.15	0.08	0.06	0.05
	Max	8.43	7.97	6.85	3.95	3.27	1.64	1.12

is very small if pure KS-DFT is employed as compared to the other steps of the computation. A further reduction of the scaling behavior of this step is subject to further research.

In addition, a convergence study of spin–spin coupling values of a peptide in a water environment revealed that large system sizes (>500 atoms) are necessary to converge results to sub-Hertz accuracy.

ACKNOWLEDGMENTS

The authors thank M. Ludwig (LMU Munich) for useful discussions. Financial support was provided by the Deutsche Forschungsgemeinschaft (DFG) for Project No. Oc35/4-1, and the Cluster of Excellence EXC 114 “Center for Integrated Protein Science Munich” (CIPSM).

- ¹E. M. Bradbury and C. Nicolini, *Nato Science Series A: NMR in the Life Sciences* (Springer, 1986).
- ²A. Rahman, *One and Two Dimensional NMR Spectroscopy* (Elsevier, 2010).
- ³G. E. Roberts and L.-Y. C.-E. Lian, *Protein NMR Spectroscopy: Practical Techniques and Applications* (Wiley-VCH, Weinheim, 2011).
- ⁴K. Wüthrich, *Angew. Chem. Int. Ed.* **42**, 3340 (2003).
- ⁵T. Helgaker, M. Jaszuński, and K. Ruud, *Chem. Rev.* **99**, 293 (1999).
- ⁶J. Gauss, in *Modern Methods and Algorithms of Quantum Chemistry*, NIC Series Vol. 3, edited by J. Grotendorst (John von Neumann Institute for Computing, Jülich, 2000), pp. 541–592.
- ⁷*Calculation of NMR and EPR Parameters: Theory and Applications*, edited by M. Kaupp, M. Bühl, and V. G. Malkin (Wiley-VCH, Weinheim, 2004).
- ⁸T. Helgaker, S. Coriani, P. Jørgensen, K. Kristensen, J. Olsen, and K. Ruud, *Chem. Rev.* **112**, 543 (2012).
- ⁹J. Kussmann, M. Beer, and C. Ochsenfeld, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **3**, 614 (2013).
- ¹⁰C. Ochsenfeld, J. Kussmann, and F. Koziol, *Angew. Chem., Int. Ed.* **43**, 4485 (2004).
- ¹¹J. Kussmann and C. Ochsenfeld, *J. Chem. Phys.* **127**, 054103 (2007).
- ¹²M. Beer, J. Kussmann, and C. Ochsenfeld, *J. Chem. Phys.* **134**, 074102 (2011).
- ¹³H. Larsen, T. Helgaker, J. Olsen, and P. Jørgensen, *J. Chem. Phys.* **115**, 10344 (2001).
- ¹⁴M. Beer and C. Ochsenfeld, *J. Chem. Phys.* **128**, 221102 (2008).
- ¹⁵D. Flaig, M. Beer, and C. Ochsenfeld, *J. Chem. Theory Comput.* **8**, 2260 (2012).
- ¹⁶W. Pisula, Z. Tomovic, M. D. Watson, K. Müllen, J. Kussmann, C. Ochsenfeld, T. Metzroth, and J. Gauss, *J. Phys. Chem. B* **111**, 7481 (2007).
- ¹⁷M. A. Watson, P. Salek, P. Macak, M. Jaszuński, and T. Helgaker, *Chem. Eur. J.* **10**, 4627 (2004).
- ¹⁸D. Cremer and J. Gräfenstein, *Phys. Chem. Chem. Phys.* **9**, 2791 (2007).
- ¹⁹E. L. Hahn and D. E. Maxwell, *Phys. Rev.* **84**, 1246 (1951).
- ²⁰N. F. Ramsey and E. M. Purcell, *Phys. Rev.* **85**, 143 (1952).
- ²¹N. F. Ramsey, *Phys. Rev.* **91**, 303 (1953).
- ²²W. T. Raynes, *Magn. Reson. Chem.* **30**, 686 (1992).
- ²³H. Fukui, *J. Chem. Phys.* **65**, 844 (1976).
- ²⁴V. G. Malkin, O. L. Malkina, and D. R. Salahub, *Chem. Phys. Lett.* **221**, 91 (1994).
- ²⁵V. Sychrovský, J. Gräfenstein, and D. Cremer, *J. Chem. Phys.* **113**, 3530 (2000).
- ²⁶T. Helgaker, M. Watson, and N. C. Handy, *J. Chem. Phys.* **113**, 9402 (2000).
- ²⁷J. Autschbach and T. Ziegler, *J. Chem. Phys.* **113**, 936 (2000).
- ²⁸J. Autschbach and T. Ziegler, *J. Chem. Phys.* **113**, 9410 (2000).
- ²⁹J. Gerratt and I. Mills, *J. Chem. Phys.* **49**, 1719 (1968).
- ³⁰J. A. Pople, R. Krishnan, H. B. Schlegel, and J. S. Binkley, *Int. J. Quantum Chem.* **16**, 225 (1979).
- ³¹Y. Osamura, Y. Yamaguchi, P. Saxe, D. Fox, M. Vincent, and H. Schaefer, *J. Mol. Struct.: THEOCHEM* **103**, 183 (1983).
- ³²M. Frisch, M. Head-Gordon, and J. Pople, *Chem. Phys.* **141**, 189 (1990).
- ³³J. Kussmann, A. Luenser, M. Beer, and C. Ochsenfeld, *J. Chem. Phys.* **142**, 094101 (2015).
- ³⁴M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).
- ³⁵M. Challacombe, *Comput. Phys. Commun.* **128**, 93 (2000).
- ³⁶C. Saravanan, Y. Shao, R. Baer, P. N. Ross, and M. Head-Gordon, *J. Comput. Chem.* **24**, 618 (2003).
- ³⁷Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia, 2003).
- ³⁸L. Colombo and W. Sawyer, *Mater. Sci. Eng., B* **37**, 228 (1996).
- ³⁹C. Ochsenfeld, C. A. White, and M. Head-Gordon, *J. Chem. Phys.* **109**, 1663 (1998).
- ⁴⁰C. Ochsenfeld, *Chem. Phys. Lett.* **327**, 216 (2000).
- ⁴¹O. B. Lutnæs, T. Helgaker, and M. Jaszuński, *Mol. Phys.* **108**, 2579 (2010).
- ⁴²J. M. Pérez-Jordá and W. Yang, *Chem. Phys. Lett.* **241**, 469 (1995).
- ⁴³R. E. Stratmann, G. E. Scuseria, and M. J. Frisch, *Chem. Phys. Lett.* **257**, 213 (1996).
- ⁴⁴J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- ⁴⁵J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **78**, 1396 (1997).
- ⁴⁶F. Jensen, *J. Chem. Theory Comput.* **2**, 1360 (2006).
- ⁴⁷F. Jensen, *Theor. Chem. Acc.* **126**, 371 (2009).
- ⁴⁸See <https://bse.pnl.gov/bse/portal> for basis set definition.
- ⁴⁹D. Feller, *J. Comput. Chem.* **17**, 1571 (1996).
- ⁵⁰K. 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).
- ⁵¹C. Y. Cheng, M. S. Ryley, M. J. G. Peach, D. J. Tozer, T. Helgaker, and A. M. Teale, *Mol. Phys.* **113**, 1937 (2015).
- ⁵²See <http://www.q-chem.com> for development version of the Q-Chem program package.
- ⁵³See www.cup.lmu.de/pc/ochsenfeld/ for structures.
- ⁵⁴A. Takatsuka, S. Ten-no, and W. Hackbusch, *J. Chem. Phys.* **129**, 044112 (2008).
- ⁵⁵P. M. W. Gill, B. G. Johnson, and J. A. Pople, *Chem. Phys. Lett.* **209**, 506 (1993).

Article IV

“A reduced-scaling density matrix-based method for the computation of the vibrational Hessian matrix at the self-consistent field level”,
J. Kussmann, A. Luenser, M. Beer, and C. Ochsenfeld,
J. Chem. Phys. **142**, 094101 (2015).

A reduced-scaling density matrix-based method for the computation of the vibrational Hessian matrix at the self-consistent field level

Jörg Kussmann, Arne Luenser, Matthias Beer, and Christian Ochsenfeld^{a)}

Chair of Theoretical Chemistry, Department of Chemistry, University of Munich (LMU), Butenandtstr. 7, D-81377 München, Germany

(Received 6 August 2014; accepted 3 February 2015; published online 2 March 2015)

An analytical method to calculate the molecular vibrational Hessian matrix at the self-consistent field level is presented. By analysis of the multipole expansions of the relevant derivatives of Coulomb-type two-electron integral contractions, we show that the effect of the perturbation on the electronic structure due to the displacement of nuclei decays at least as r^{-2} instead of r^{-1} . The perturbation is asymptotically local, and the computation of the Hessian matrix can, in principle, be performed with $O(N)$ complexity. Our implementation exhibits linear scaling in all time-determining steps, with some rapid but quadratic-complexity steps remaining. Sample calculations illustrate linear or near-linear scaling in the construction of the complete nuclear Hessian matrix for sparse systems. For more demanding systems, scaling is still considerably sub-quadratic to quadratic, depending on the density of the underlying electronic structure. © 2015 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4908131>]

I. INTRODUCTION

Together with the evaluation of the forces acting on the nuclei at a specified molecular geometry, the prediction of vibrational frequencies in the harmonic approximation is among the most fundamental properties in quantum chemistry.^{1,2} Not only are the predictions of the vibrational frequencies and related thermochemical properties themselves quite accurate (especially when using empirical scaling³), but perhaps even more importantly, the sign of the force constants at zero-gradient molecular geometries reveals important information about the potential energy hypersurface at that point. It allows for characterization of true energetic minima, transition states, and higher-order saddle points. Prediction of vibrational frequencies is a routine procedure nowadays, albeit only for small to medium-sized molecules.

To obtain the harmonic force constants, the matrix of all second derivatives of the molecular energy (Hessian matrix) with respect to displacements of the nuclei must be evaluated. This requires solving the coupled perturbed self-consistent field (CPSCF) equations⁴ in order to find the response of the electronic structure to the perturbation, i.e., the perturbed (response) density matrices. Traditionally, this has been done completely in the molecular orbital (MO) basis (MO-CPSCF),⁴⁻⁶ necessitating a transformation of the two-electron integrals from the atomic orbital (AO) to the molecular orbital basis. The computational time for this step scales as $O(N^3)$, where N is some measure of system size such as the number of atoms. Alternatively, an integral-direct variant⁷ can be used, in which the two-electron integrals are evaluated and directly contracted with the response densities, while the linear equations are still solved in the MO basis. This approach scales as $O(N^4)$ for the case of vibrational frequencies.

Reduction of the computational complexity of self-consistent field (SCF) calculations, specifically Hartree-Fock (HF) and density functional theory (DFT), has spawned many different approaches. We are concerned with molecular calculations using a Gaussian basis only. Computation of the Coulomb matrix with $O(N)$ or $O(N \log N)$ scaling is possible with, for example, the Continuous Fast Multipole Method (CFMM),^{8,9} the Quantum Chemical Tree Code (QCTC),^{10,11} and the Fourier Transform Coulomb (FTC) method.¹²⁻¹⁶ For insulating systems, the exchange matrix can be obtained with $O(N)$ complexity through the NFX,¹⁷ ONX,^{18,19} LinK,^{20,21} and COSX²² methods. Numerical quadrature of the exchange-correlation energy functional is described in Refs. 23 and 24. The final bottleneck is the $O(N^3)$ diagonalization of the Fock matrix, which can be circumvented by density matrix minimization²⁵⁻³² or purification³³⁻³⁶ algorithms.

Likewise, reduced-scaling computation of response properties has received attention in the form of density matrix-based CPSCF,³⁷⁻⁴⁰ exponential parametrization of the density matrix,⁴¹⁻⁴³ density matrix perturbation theory,⁴⁴⁻⁴⁷ and density matrix-based Laplace-transformed CPSCF (DL-CPSCF),^{48,49} which is discussed in Sec. II. Reduced-scaling approaches for second nuclear derivatives are briefly reviewed in Sec. III B.

In this work, we describe the theory to evaluate the full nuclear-displacement Hessian matrix at the SCF level with effective linear scaling behavior. We show that the long-range Coulomb interaction due to the nuclear displacement decays more rapidly than is obvious from rudimentary inspection (r^{-2} instead of r^{-1}). The perturbation of the nuclear displacement on the electronic structure is spatially local and asymptotically does not increase in size for larger molecules. Our implementation at HF level was tested on several molecular systems of increasing size, showcasing near-linear scaling behavior in the time-determining steps. Evaluation of the far-field Coulomb contributions is asymptotically $O(N^2)$ with a very small prefactor, but can, in principle, be reduced to $O(N)$. The limitations

^{a)}Electronic mail: christian.ochsenfeld@uni-muenchen.de

of the method are discussed and an overview of possible future work is given.

II. THEORY

In the Born–Oppenheimer approximation, electronic and nuclear motion are uncoupled. Apart from the electronic Schrödinger equation, this ansatz also results in a nuclear Schrödinger equation whose Hamiltonian is

$$\hat{H}_{\text{nuc}} = \hat{T}_{\text{nuc}} + E_{\text{el}}(\{\mathbf{R}_A\}),$$

where the potential E_{el} is the electronic energy including the nuclear-nuclear repulsion. E_{el} is a function of the complete set of nuclear coordinates $\{\mathbf{R}_A\}$. For small nuclear displacements, a harmonic potential is often a good approximation

$$V_{\text{harmonic}} = \frac{1}{2}kr^2.$$

This model potential corresponds to the harmonic oscillator model for molecular vibrations. To obtain the vibrational eigenmodes and eigenfrequencies of a molecule in the harmonic approximation, the $3N_{\text{atoms}} \times 3N_{\text{atoms}}$ matrix of second derivatives of the electronic energy with respect to the cartesian nuclear coordinates (Hessian matrix for short) must be calculated. The Hessian matrix is then weighted with nuclear masses, and translational and rotational degrees of freedom are projected out. The eigenmodes and eigenfrequencies are obtained from the resulting matrix by diagonalization.

In a density matrix-based form, the HF or DFT expression for the electronic energy E_{el} is

$$E_{\text{SCF}} = \text{Tr} \left[\mathbf{P}\mathbf{h} + \frac{1}{2}\mathbf{P}\mathbf{G}[\mathbf{P}] \right] + V_{\text{NN}},$$

with the ground-state one-particle density matrix \mathbf{P} , the core Hamiltonian matrix $\mathbf{h} = \mathbf{T} + \mathbf{V}_{\text{eN}}$, the two-electron terms $\mathbf{G}[\mathbf{P}] = \mathbf{J}[\mathbf{P}] + \alpha\mathbf{K}[\mathbf{P}] + (1 - \alpha)\mathbf{V}_{\text{XC}}[\mathbf{P}]$ ($\alpha = 1$: HF, $0 < \alpha < 1$: hybrid DFT, $\alpha = 0$: pure DFT), and the nuclear-nuclear repulsion energy V_{NN} .

The first derivative with respect to a perturbation x gives

$$\frac{\partial E}{\partial x} = E^x = \text{Tr} \left[\mathbf{P}\mathbf{h}^x + \frac{1}{2}\mathbf{P}\mathbf{G}^x[\mathbf{P}] - \mathbf{W}\mathbf{S}^x \right] + V_{\text{NN}}^x,$$

with the energy-weighted density matrix $\mathbf{W} = \mathbf{P}\mathbf{F}\mathbf{P}$. The second derivative is given by

$$\begin{aligned} \frac{\partial^2 E}{\partial y \partial x} = E^{xy} = \text{Tr} \left[\mathbf{P}\mathbf{h}^{xy} + \frac{1}{2}\mathbf{P}\mathbf{G}^{xy}[\mathbf{P}] - \mathbf{W}\mathbf{S}^{xy} \right] \\ + \text{Tr} [\mathbf{P}^y\mathbf{h}^x + \mathbf{P}^y\mathbf{G}^x[\mathbf{P}] - \mathbf{W}^y\mathbf{S}^x] + V_{\text{NN}}^{xy}. \end{aligned} \quad (1)$$

The second matrix trace in this expression contains the first-order response density with respect to perturbation y , obtained from CPSCF.

In this work, we adapt the DL-CPSCF algorithm.⁴⁸ In what follows, $\{i, j, \dots\}$ refer to occupied MOs, $\{a, b, \dots\}$ refer to virtual MOs, $\{p, q, \dots\}$ refer to any (occupied or virtual) MO, and $\{\mu, \nu, \dots\}$ refer to AOs. We start from the Roothaan-Hall equations^{50,51}

$$\mathbf{F}\mathbf{C} = \mathbf{S}\mathbf{C}\boldsymbol{\varepsilon}$$

and form the derivative with respect to x

$$\mathbf{F}^x\mathbf{C} + \mathbf{F}\mathbf{C}^x = \mathbf{S}^x\mathbf{C}\boldsymbol{\varepsilon} + \mathbf{S}\mathbf{C}^x\boldsymbol{\varepsilon} + \mathbf{S}\mathbf{C}\boldsymbol{\varepsilon}^x.$$

Inserting $\mathbf{C}_{\mu p}^x = \sum_q C_{\mu q} U_{qp}^x$ and multiplying with \mathbf{C}^\dagger from the left gives

$$\mathbf{F}_{\text{MO}}^x + \boldsymbol{\varepsilon}\mathbf{U}^x = \mathbf{S}_{\text{MO}}^x\boldsymbol{\varepsilon} + \mathbf{U}^x\boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}^x.$$

Since the matrix $\boldsymbol{\varepsilon}^x$ is block-diagonal, i.e., occupied and virtual spaces are orthogonal, it suffices to inspect the virtual-occupied subspace. One obtains

$$U_{ai}^x = \frac{S_{ai}^x \boldsymbol{\varepsilon}_i - F_{ai}^x}{\boldsymbol{\varepsilon}_a - \boldsymbol{\varepsilon}_i}, \quad (2)$$

which is a short form of the canonical MO-CPSCF equations. The response densities, whose subspace projections satisfy⁵²

$$\mathbf{P}_{\text{oo}}^x = -\mathbf{P}\mathbf{S}^x\mathbf{P}, \quad \mathbf{P}_{\text{ov}}^x = \mathbf{P}_{\text{vo}}^{x\dagger}, \quad \mathbf{P}_{\text{vv}}^x = \mathbf{0},$$

are obtained from the matrix \mathbf{U}^x as

$$(\mathbf{P}_{\text{vo}}^x)_{\mu\nu} = \sum_{ai} C_{\mu a} U_{ai}^x C_{i\nu}^\dagger.$$

However, Eq. (2) cannot be readily transformed to the AO basis because of the orbital energy denominator. A Laplace transformation and numerical integration^{53–55} with τ roots t_α and weights w_α ,

$$\frac{1}{\boldsymbol{\varepsilon}_a - \boldsymbol{\varepsilon}_i} = \int_0^\infty e^{(\boldsymbol{\varepsilon}_i - \boldsymbol{\varepsilon}_a)t} dt \approx \sum_{\alpha=1}^\tau w_\alpha e^{(\boldsymbol{\varepsilon}_i - \boldsymbol{\varepsilon}_a)t_\alpha},$$

allow us to write Eq. (2) in the AO basis

$$(\mathbf{P}_{\text{vo}}^x)_{\mu\nu} = \sum_{\alpha=1}^\tau \sum_{\lambda\sigma} \bar{Q}_{\mu\lambda}^{(\alpha)} (\tilde{b}_{\lambda\sigma}^x - G_{\lambda\sigma}^x[\mathbf{P}^x]) \underline{P}_{\sigma\nu}^{(\alpha)}, \quad (3a)$$

$$\tilde{b}_{\lambda\sigma}^x = (\mathbf{S}^x\mathbf{P}\mathbf{F})_{\lambda\sigma} - h_{\lambda\sigma}^x - G_{\lambda\sigma}^x[\mathbf{P}], \quad (3b)$$

with

$$\begin{aligned} \underline{P}_{\sigma\nu}^{(\alpha)} &= \sqrt{w_\alpha} \sum_{i \in \text{occ}} C_{\sigma i} e^{t_\alpha \boldsymbol{\varepsilon}_i} C_{i\nu}^\dagger, \\ \bar{Q}_{\mu\lambda}^{(\alpha)} &= \sqrt{w_\alpha} \sum_{a \in \text{virt}} C_{\mu a} e^{-t_\alpha \boldsymbol{\varepsilon}_a} C_{a\lambda}^\dagger. \end{aligned}$$

As will be shown in Sec. III, it is crucial to move the occupied-occupied projection of the perturbed density (\mathbf{P}_{oo}^x) within the integrand from $\mathbf{G}[\mathbf{P}^x]$ to $\tilde{\mathbf{b}}^x$ to form the working DL-CPSCF equations

$$\mathbf{P}_{\text{vo}}^x = \sum_{\alpha=1}^\tau \bar{\mathbf{Q}}^{(\alpha)} (\mathbf{b}^x - \mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]) \underline{\mathbf{P}}^{(\alpha)}, \quad (4a)$$

$$\mathbf{b}^x = \mathbf{S}^x\mathbf{P}\mathbf{F} - \mathbf{h}^x - \mathbf{G}^x[\mathbf{P}] - \mathbf{G}[\mathbf{P}_{\text{oo}}^x]. \quad (4b)$$

\mathbf{b}^x is computed just once for each perturbation, as all components are directly accessible. Equation (4a) is then solved iteratively for each perturbation with the help of Pulay's *direct inversion of the iterative subspace* (DIIS).⁵⁶ The pseudo-densities of the occupied ($\underline{\mathbf{P}}^{(\alpha)}$) and virtual ($\bar{\mathbf{Q}}^{(\alpha)}$) subspace are computed only once, in a linear-scaling fashion as described in Refs. 48, 55, and 57

$$\begin{aligned} \underline{\mathbf{P}}^{(\alpha)} &= \sqrt{w_\alpha} \exp(t_\alpha \mathbf{P}\mathbf{F}) \mathbf{P}, \\ \bar{\mathbf{Q}}^{(\alpha)} &= \sqrt{w_\alpha} \exp(-t_\alpha \mathbf{Q}\mathbf{F}) \mathbf{Q}. \end{aligned}$$

III. SCALING BEHAVIOR

A. General considerations

For linear-scaling behavior, our method requires the system to have a local electronic structure, i.e., a non-zero gap between the highest occupied and lowest virtual orbital energies (HOMO-LUMO gap). The ground state Fock and density matrices, \mathbf{F} and \mathbf{P} , respectively, will then contain a number of significant elements, which grows as $\mathcal{O}(N)$ with system size.

We do not tackle the issue of extracting useful information from the Hessian matrix in this work (i.e., circumventing the $\mathcal{O}(N_{\text{atoms}}^3)$ diagonalization), because the computational prefactor of that step is very small. Hessian matrices for systems with several thousand atoms can be rapidly diagonalized on workstation computers.

The time-determining steps in the construction of the Hessian matrix are the determination of the electronic response density \mathbf{P}^x via CPSCF and the formation of two-electron integrals. Since three response densities (x, y, z directions) are required for each atom in the system, the computational time for the determination of each of them must be asymptotically constant to allow for overall linear scaling. Equation (4) can be solved in constant time if \mathbf{b}^x and $\mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]$ have an asymptotically constant number of significant elements (see Sec. III C). Technical details of the DL-CPSCF regarding the sparse-matrix implementation are discussed in Sec. IV B.

B. State of the art

The computationally expensive contributions to the second derivatives of the SCF energy with respect to nuclear positions are the derivatives of the Coulomb and exchange integral contractions, the derivatives of the exchange-correlation matrix (DFT only), and the solution of the CPSCF equations.

The derivatives required for the Coulomb contribution in the CFMM formulation were first given by Burant, Strain, Scuseria, and Frisch.^{58,59} Like CFMM itself, the gradient also scales linearly with system size because the integral derivatives are directly contracted with the ground state density matrices

$$\begin{aligned} E_J^x &= \text{Tr}[\mathbf{P}\mathbf{J}^x[\mathbf{P}]] = \sum_{\mu\nu} P_{\mu\nu}(\mu\nu|\lambda\sigma)^x P_{\lambda\sigma} \\ &= \sum_{\mu\nu\lambda\sigma} P_{\mu\nu}([\mu\nu]^x|\lambda\sigma)P_{\lambda\sigma} + \sum_{\mu\nu\lambda\sigma} P_{\mu\nu}(\mu\nu|\lambda\sigma)^x P_{\lambda\sigma} \\ &= 2 \sum_{\mu\nu\lambda\sigma} P_{\mu\nu}([\mu\nu]^x|\lambda\sigma)P_{\lambda\sigma}, \end{aligned}$$

where the far-field $|\lambda\sigma\rangle P_{\lambda\sigma}$ is generated only once for all derivatives.

For second derivatives, however, the explicit matrix $\mathbf{J}^x[\mathbf{P}]$ is required for each perturbation (cf. Sec. III C), which to date scales linearly per perturbation, leading to overall quadratic scaling in the Coulomb contribution. Similarly, the $\mathbf{J}[\mathbf{P}^x]$ and $\mathbf{J}^x[\mathbf{P}^y]$ terms scale quadratically in existing implementations. We show in the following that an overall linear-scaling is possible in principle.

Linear-scaling methods to compute the first-order exchange contributions $\mathbf{K}^x[\mathbf{P}]$ and $\mathbf{K}[\mathbf{P}^x]$ can be found in Refs. 20, 21, 42, and 60. Second-order derivatives of this term are briefly

discussed in Sec. IV D 2 and demonstrated to show linear scaling in Sec. V.

Solution of the CPSCF equations for nuclear derivatives with linear complexity has been proposed several times (cf. Sec. I) but is demonstrated here for the first time, using sparse matrix algebra in all steps.

C. Integrals and decay behavior

We distinguish between two types of terms in Eq. (1). First, quantities that are only calculated once and scale as $\mathcal{O}(N)$. Second, quantities that are calculated $3N$ times and, therefore, must scale as $\mathcal{O}(1)$ with respect to system size for overall linear scaling.

Quantities in the first category are the second-order integral derivatives contracted with the ground-state density. Linear scaling can, in principle, be achieved by adapting CFMM^{8,9} and LinK^{20,21}-based integral routines combined with sparse linear algebra. Details of their implementation are given in Sec. IV D 2.

The second category warrants a more detailed analysis. Here, the $\mathcal{O}(1)$ scaling of the CPSCF for a single perturbation is crucial in order to achieve an overall linear-scaling behavior. Apart from the integral evaluation, DL-CPSCF equation (4) has to be computed with $\mathcal{O}(1)$ complexity. This is possible if \mathbf{P}^x , \mathbf{b}^x , and $\mathbf{G}[\mathbf{P}^x]$ contain only an asymptotically constant number of significant elements for a single perturbation. We will show this to be the case in the following.

When using Eq. (3) instead of Eq. (4) (that is, the term $\mathbf{G}[\mathbf{P}_{\text{oo}}^x]$ is computed as part of $\mathbf{G}[\mathbf{P}^x]$ instead of \mathbf{b}^x), both \mathbf{b}^x and $\mathbf{G}[\mathbf{P}^x]$ contain a linear-scaling number of significant elements. On first glance, the problematic terms in \mathbf{b}^x appear to be \mathbf{V}_{en}^x and $\mathbf{J}^x[\mathbf{P}]$, due to their long-range r^{-1} coupling. Inspecting the explicit derivative formula for \mathbf{V}_{en}^x , however, reveals the term is in fact short-ranged

$$\begin{aligned} \frac{\partial}{\partial x} \int \chi_\mu \frac{Z_A}{r_A} \chi_\nu dr &= \int \frac{\partial \chi_\mu}{\partial x} \frac{Z_A}{r_A} \chi_\nu dr + \int \chi_\mu \frac{Z_A}{r_A} \frac{\partial \chi_\nu}{\partial x} dr \\ &\quad - \int \chi_\mu \frac{Z_A^x}{r_A^3} \chi_\nu dr. \end{aligned} \quad (5)$$

The first two terms in Eq. (5) are only non-zero for basis functions centered on the displaced nucleus. The last term represents an electric field integral which decays as r^{-2} and, therefore, results in only a constant number of significant matrix elements. A similar term occurs in the context of nuclei-selected magnetic shielding calculations.⁴⁹

The term $\mathbf{J}^x[\mathbf{P}]$ shows a linear-scaling number of significant elements due to the r^{-1} coupling of the locally perturbed charge-distribution $[[\lambda\sigma]^x]P_{\lambda\sigma}$ with the bra distributions $(\mu\nu|$

$$J_{\mu\nu}^x[\mathbf{P}] = \sum_{\lambda\sigma} ([\mu\nu]^x|\lambda\sigma)P_{\lambda\sigma} + \sum_{\lambda\sigma} (\mu\nu|[\lambda\sigma]^x)P_{\lambda\sigma}. \quad (6)$$

The first term on the right-hand side of Eq. (6) results in a constant number of matrix elements. The unperturbed far-field $|\lambda\sigma\rangle P_{\lambda\sigma}$ is the same for all perturbations and can be reused from the final SCF iteration. The second term, however, leads to linear growth in the number of significant matrix elements. Similarly, the number of significant elements in the matrix

$\mathbf{J}[\mathbf{P}^x]$,

$$J_{\mu\nu}[\mathbf{P}^x] = \sum_{\lambda\sigma} (\mu\nu|\lambda\sigma) P_{\lambda\sigma}^x,$$

grows linearly for each perturbation. For deeper analysis, we split $\mathbf{J}[\mathbf{P}^x]$ into separate contributions from the subspace

projections of \mathbf{P}^x

$$\mathbf{J}[\mathbf{P}^x] = \mathbf{J}[\mathbf{P}_{00}^x] + \mathbf{J}[\mathbf{P}_{0v}^x] + \mathbf{J}[\mathbf{P}_{v0}^x],$$

and inspect the multipole expansion^{61,62} of the long-range (i.e., far-field) component of $\mathbf{J}[\mathbf{P}_{00}^x]$,

$$\begin{aligned} J_{\mu\nu}^{\text{FF}}[\mathbf{P}_{00}^x] &= \sum_{\lambda\sigma\backslash\mu\nu} (\mu\nu|\lambda\sigma) (\mathbf{P}_{00}^x)_{\lambda\sigma} \\ &= \sum_{\lambda\sigma\backslash\mu\nu} \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{j=0}^{\infty} \sum_{k=-j}^j q_{lm}^{\mu\nu}(\mathbf{r}_P) T_{lm,jk}(\mathbf{r}_{PQ}) q_{jk}^{\lambda\sigma}(\mathbf{r}_Q) (\mathbf{P}_{00}^x)_{\lambda\sigma} \\ &= - \sum_{\substack{\lambda\sigma\backslash\mu\nu \\ \kappa\theta\backslash\mu\nu}} \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{j=0}^{\infty} \sum_{k=-j}^j q_{lm}^{\mu\nu}(\mathbf{r}_P) T_{lm,jk}(\mathbf{r}_{PQ}) q_{jk}^{\lambda\sigma}(\mathbf{r}_Q) P_{\lambda\kappa} S_{\kappa\theta}^x P_{\theta\sigma} \\ &= - \sum_{\substack{\lambda\sigma\backslash\mu\nu \\ \kappa\theta\backslash\mu\nu}} \left(\frac{q_{00}^{\mu\nu} q_{00}^{\lambda\sigma}}{r_{PQ}} + \frac{q_{00}^{\mu\nu} \left[\sum_{j=-1}^1 T'_{00,1j} q_{1j}^{\lambda\sigma} \right]}{r_{PQ}^2} + \frac{\left[\sum_{j=-1}^1 T'_{1j,00} q_{1j}^{\mu\nu} \right] q_{00}^{\lambda\sigma}}{r_{PQ}^2} + \dots \right) P_{\lambda\kappa} S_{\kappa\theta}^x P_{\theta\sigma}, \end{aligned} \quad (7)$$

with

$$\begin{aligned} q_{lm}^{\mu\nu}(\mathbf{r}_P) &= \int \chi_{\mu}(\mathbf{r}) \chi_{\nu}(\mathbf{r}) R_{lm}(\mathbf{r}_P) d\mathbf{r}, \\ T_{lm,jk}(\mathbf{r}_{PQ}) &= T'_{lm,jk}(\mathbf{r}_{PQ}) r_{PQ}^{-(l+j+1)} \\ &= (-1)^j I_{l+j,m+k}^*(\mathbf{r}_{PQ}), \end{aligned}$$

where R_{lm} and $I_{l+j,m+k}$ designate regular and irregular solid harmonics, respectively. We use the notation $\lambda\sigma\backslash\mu\nu$ to emphasize that the charge distribution $\chi_{\lambda}\chi_{\sigma}(\mathbf{P}_{00}^x)_{\lambda\sigma}$ must be fully disjoint from $\chi_{\mu}\chi_{\nu}$, i.e., the idempotency condition $\mathbf{PSP} = \mathbf{P}$ of the density matrix must hold for the set of all $\chi_{\lambda}\chi_{\sigma}$. Intuitively, this will be the case for two distinct molecular systems at large separations, where the density matrix is individually idempotent for each system.

Analogously, the multipole expansion for the far-field component of $\sum_{\lambda\sigma} (\mu\nu|[\lambda\sigma]^x) P_{\lambda\sigma}$ is

$$\sum_{\lambda\sigma\backslash\mu\nu} (\mu\nu|[\lambda\sigma]^x) P_{\lambda\sigma} = \sum_{\lambda\sigma\backslash\mu\nu} \left(\frac{q_{00}^{\mu\nu} (q_{00}^{\lambda\sigma})^x}{r_{PQ}} + \frac{q_{00}^{\mu\nu} \left[\sum_{j=-1}^1 T'_{00,1j} (q_{1j}^{\lambda\sigma})^x \right]}{r_{PQ}^2} + \frac{\left[\sum_{j=-1}^1 T'_{1j,00} q_{1j}^{\mu\nu} \right] (q_{00}^{\lambda\sigma})^x}{r_{PQ}^2} + \dots \right) P_{\lambda\sigma}, \quad (8)$$

where $(q_{jk}^{\lambda\sigma})^x = \partial q_{jk}^{\lambda\sigma} / \partial x$. Since $q_{00}^{\lambda\sigma} = S_{\lambda\sigma}$ and $(q_{00}^{\lambda\sigma})^x = S_{\lambda\sigma}^x$, the long-range, slowly decaying r^{-1} components in the sum of Eqs. (7) and (8) cancel

$$\begin{aligned} &\sum_{\substack{\lambda\sigma\backslash\mu\nu \\ \kappa\theta\backslash\mu\nu}} \left(\frac{q_{00}^{\mu\nu} (q_{00}^{\lambda\sigma})^x}{r_{PQ}} P_{\lambda\sigma} - \frac{q_{00}^{\mu\nu} q_{00}^{\lambda\sigma}}{r_{PQ}} P_{\lambda\kappa} S_{\kappa\theta}^x P_{\theta\sigma} \right) \\ &= \sum_{\lambda\sigma\backslash\mu\nu} \frac{S_{\mu\nu} S_{\lambda\sigma}^x P_{\lambda\sigma}}{r_{PQ}} - \sum_{\substack{\lambda\sigma\backslash\mu\nu \\ \kappa\theta\backslash\mu\nu}} \frac{S_{\mu\nu} S_{\lambda\sigma} P_{\lambda\kappa} S_{\kappa\theta}^x P_{\theta\sigma}}{r_{PQ}} \\ &= \sum_{\lambda\sigma\backslash\mu\nu} \frac{S_{\mu\nu} S_{\lambda\sigma}^x P_{\lambda\sigma}}{r_{PQ}} - \sum_{\kappa\theta\backslash\mu\nu} \frac{S_{\mu\nu} S_{\kappa\theta}^x P_{\theta\kappa}}{r_{PQ}} = 0, \end{aligned} \quad (9)$$

where we used the idempotency condition $\mathbf{PSP} = \mathbf{P}$. Building \mathbf{b}^x as given in Eq. (4b) by explicitly moving $\mathbf{J}[\mathbf{P}_{00}^x]$ from $\mathbf{G}[\mathbf{P}^x]$ to \mathbf{b}^x gives a local matrix where all components decay at least as r^{-2} . We use the conventional $O(N)$ CFMM for this term in our implementation (Sec. IV) because of the small computational prefactor of this step. One approach to consider the cancellation of far-field terms during the integral evaluation would be the introduction of a second well-separated criterion based on the extent of the density-coupled distribution $S_{\mu\nu} P_{\nu\lambda} S_{\lambda\sigma}$, i.e., ensuring that the distribution $\mu\nu$ is not coupled to $\lambda\sigma$ via the density matrix \mathbf{P} , so that Eq. (9) remains valid. However, the implementational effort is hardly justified by the benefit due to the small prefactor of the long-range contribution to the Coulomb matrices.

Finally, the far-field component of $\mathbf{J}[\mathbf{P}_{\text{ov}}^x]$ must be shown to decay faster than r^{-1} . Again, we invoke the multipole expansion

$$\begin{aligned}
 J_{\mu\nu}^{\text{FF}}[\mathbf{P}_{\text{ov}}^x] &= \sum_{\lambda\sigma \setminus \mu\nu} (\mu\nu|\lambda\sigma) P_{\text{ov},\lambda\sigma}^x = \sum_{\substack{\lambda\sigma \setminus \mu\nu \\ \kappa\theta \setminus \mu\nu}} (\mu\nu|\lambda\sigma)(\mathbf{P}\mathbf{S})_{\lambda\kappa} P_{\kappa\theta}^x (\mathbf{1} - \mathbf{S}\mathbf{P})_{\theta\sigma} \\
 &= - \sum_{\substack{\lambda\sigma \setminus \mu\nu \\ \kappa\theta \setminus \mu\nu}} \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{j=0}^{\infty} \sum_{k=-j}^j q_{lm}^{\mu\nu}(\mathbf{r}_P) T_{lm,jk}(\mathbf{r}_{PQ}) q_{jk}^{\lambda\sigma}(\mathbf{r}_Q) (\mathbf{P}\mathbf{S})_{\lambda\kappa} P_{\kappa\theta}^x (\mathbf{1} - \mathbf{S}\mathbf{P})_{\theta\sigma} \\
 &= - \sum_{\substack{\lambda\sigma \setminus \mu\nu \\ \kappa\theta \setminus \mu\nu}} \left(\frac{q_{00}^{\mu\nu} q_{00}^{\lambda\sigma}}{r_{PQ}} + \frac{q_{00}^{\mu\nu} \left[\sum_{j=-1}^1 T'_{00,1j} q_{1j}^{\lambda\sigma} \right]}{r_{PQ}^2} + \frac{\left[\sum_{j=-1}^1 T'_{1j,00} q_{1j}^{\mu\nu} \right] q_{00}^{\lambda\sigma}}{r_{PQ}^2} + \dots \right) (\mathbf{P}\mathbf{S})_{\lambda\kappa} P_{\kappa\theta}^x (\mathbf{1} - \mathbf{S}\mathbf{P})_{\theta\sigma} \\
 &= - \sum_{\substack{\lambda\sigma \setminus \mu\nu \\ \kappa\theta \setminus \mu\nu}} \left(0 + \frac{q_{00}^{\mu\nu} \left[\sum_{j=-1}^1 T'_{00,1j} q_{1j}^{\lambda\sigma} \right]}{r_{PQ}^2} + \frac{\left[\sum_{j=-1}^1 T'_{1j,00} q_{1j}^{\mu\nu} \right] q_{00}^{\lambda\sigma}}{r_{PQ}^2} + \dots \right) (\mathbf{P}\mathbf{S})_{\lambda\kappa} P_{\kappa\theta}^x (\mathbf{1} - \mathbf{S}\mathbf{P})_{\theta\sigma},
 \end{aligned}$$

where we used the cyclic permutation property of the trace, and $(\mathbf{1} - \mathbf{S}\mathbf{P})\mathbf{S}\mathbf{P} = \mathbf{0}$. The long-range r^{-1} component thus vanishes in all Coulomb-type matrices, leading to an r^{-2} coupling. The two matrices \mathbf{b}^x and $\mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]$ in Eq. (4) as well as the perturbed density matrix \mathbf{P}^x itself are local and contain an asymptotically constant number of significant elements for large system sizes. In principle, computing the electronic response due to a displacement of a single nucleus is, therefore, possible with $\mathcal{O}(1)$ time complexity.

IV. IMPLEMENTATIONAL DETAILS

A. Sparse storage format

All matrices in our implementation are stored in a variant of the blocked compressed sparse row format (BCSR).^{26,40,63,64} In this storage scheme, the matrix is divided into blocks of submatrices of dimension $\sim 100 \times 100$. Only blocks whose matrix norm (typically the Frobenius norm) is above a numerical threshold are allocated in memory. In contrast to the original BCSR format, the non-zero blocks are not necessarily located in a single contiguous chunk of memory but are individually allocated. This provides increased flexibility during the implementation of matrix operations, because allocating and freeing individual blocks during integral contraction procedures is possible on-the-fly and independently of all other blocks. The memory overhead is negligible. At the same time, matrix-matrix multiplications, matrix traces, or similar matrix operations are trivial to implement and cause virtually zero overhead compared to the corresponding dense matrix operations. The matrix operations on the individual blocks are delegated to highly tuned linear algebra (BLAS) libraries for maximum performance.

B. DL-CPSCF

For large molecular systems, the DL-CPSCF equations must be solved in several consecutive batches. In order to avoid unnecessary recalculation of a large number of two-electron integrals, we group the atoms by relative proximity and create

batches of spatially close atoms. The response densities \mathbf{P}^x in a batch thus share most of the significant two-electron integrals, reducing overhead. To obtain a suitable atom ordering, we apply the reverse Cuthill–McKee algorithm⁶⁵ to the connectivity matrix.^{40,66} This or a similar reordering is also required to obtain compact matrix representations for all quantities in DL-CPSCF equation (4). The matrix multiplications can be performed in a way that is both fast and has the necessary $\mathcal{O}(1)$ complexity only if the significant elements of each matrix are grouped together. We explicitly reproject matrices onto the virtual-occupied subspace ($\mathbf{A} \rightarrow (\mathbf{1} - \mathbf{S}\mathbf{P})\mathbf{A}\mathbf{S}$) to increase sparsity and obtain more local matrices. In our implementation, this is performed for \mathbf{b}^x and $\mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]$ (cf. Eq. (4) and Fig. 2).

When building the pseudo-density matrices $\underline{\mathbf{P}}^{(\alpha)}$ and $\overline{\mathbf{Q}}^{(\alpha)}$ via matrix exponentials (see, e.g., Ref. 67), special care must be taken with respect to sparsity cutoffs, as too loose thresholds will lead to numerical noise and consequently to low sparsity. A recent discussion about matrix sparsity in electronic structure calculations can be found in Ref. 68. We also remark that the numerical quadrature fitting for the Laplace transformation we presently employ⁶⁹ requires knowledge of the system's HOMO-LUMO gap, which can, in principle, be obtained even within diagonalization-free SCF procedures through, e.g., the method put forward by Rubensson and Zahedi.⁷⁰

C. Cholesky decomposition of the pseudo-densities

Fill-in effects (creeping loss of matrix sparsity), particularly in the perturbed density matrices \mathbf{P}^x , adversely affect matrix multiplications in the main DL-CPSCF equation (4), as well as integral formation and contraction. Sparsity loss can be alleviated by performing the multiplications in Eq. (4a) stepwise. To this end, we perform a Cholesky decomposition with complete pivoting on $\underline{\mathbf{P}}^{(\alpha)}$ and $\overline{\mathbf{Q}}^{(\alpha)}$ ($\underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)\text{T}}$) and subsequently reorder the columns of the Cholesky factors

$$\underline{\mathbf{P}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)} \underline{\mathbf{L}}^{(\alpha)\text{T}} = \underline{\mathbf{L}}'^{(\alpha)} \underline{\mathbf{L}}'^{(\alpha)\text{T}},$$

with $\underline{\mathbf{L}}^{(\alpha)} = \underline{\mathbf{L}}^{(\alpha)}\mathbf{M}$. \mathbf{M} is a (unitary) permutation matrix which reorders the columns of $\underline{\mathbf{L}}^{(\alpha)}$ in ascending order of a weighted mean index defined by

$$w_j = \frac{\sum_i |\underline{L}_{ij}^{(\alpha)}| i}{\sum_i |\underline{L}_{ij}^{(\alpha)}|}.$$

$\overline{\mathbf{Q}}^{(\alpha)}$ is decomposed analogously as $\overline{\mathbf{Q}}^{(\alpha)} = \overline{\mathbf{L}}^{(\alpha)}\overline{\mathbf{L}}^{(\alpha)\top}$.

The resulting transformation matrices $\underline{\mathbf{L}}^{(\alpha)}$ and $\overline{\mathbf{L}}^{(\alpha)}$ have dimensions $N_{\text{basis}} \times N_{\text{occ}}$ and $N_{\text{basis}} \times N_{\text{virt}}$, respectively. They are sparse and define local pseudo-molecular orbitals, which are not orthogonal.⁷¹ The orthogonality of the occupied and virtual subspaces, however, is preserved. The intermediate quantities $\overline{\mathbf{L}}^{(\alpha)\top}[\mathbf{b}^x - \mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]]\underline{\mathbf{L}}^{(\alpha)}$ are non-zero only in their virtual-occupied blocks, which in turn are also sparse.

Performing the numerical integration in this stepwise manner doubles the number of matrix multiplications per iteration from 2τ to 4τ . However, because of their pronounced sparsity, multiplication with the Cholesky factors is so rapid that the numerical integration is typically faster nevertheless. Tests showed a 10%–15% reduction in computational time for the evaluation of Eq. (4a). We attribute the improved sparsity in \mathbf{P}^x to the lower absolute values in $\overline{\mathbf{L}}^{(\alpha)}$ compared to $\overline{\mathbf{Q}}^{(\alpha)}$. We have observed that sparse matrix multiplications accumulate less numerical noise when the multiplicands have absolute values of comparable magnitude (cf. Ref. 40). With the correct sequence of multiplications, the Cholesky decomposition reduces the complexity of Eq. (4a) from $O(N_{\text{basis}}^3)$ to $O(N_{\text{basis}}^2 N_{\text{occ}})$ for dense matrices. For local matrices, the scaling is unchanged at $O(1)$.

D. Integral routines

1. First-order derivatives

The first-order derivatives of various terms occurring in the Fock- or Kohn–Sham matrix must be evaluated such that only a constant amount of memory is required. The integrals and integral derivatives are scattered directly into sparse matrices.

The evaluation of the one-electron terms $\mathbf{S}^x = \frac{\partial \mathbf{S}}{\partial \mathbf{x}}$, $\mathbf{T}^x = \frac{\partial \mathbf{T}}{\partial \mathbf{x}}$, and $\mathbf{V}_{\text{en}}^x = \frac{\partial \mathbf{V}_{\text{en}}}{\partial \mathbf{x}}$ is straightforward. We stress that the evaluation of the electron-nuclear attraction term formally scales quadratically with system size due to the r^{-1} -coupling in the electron-nuclear attraction term. Although this could be reduced to linear by applying the fast multipole method, we currently use the conventional routine since the computational effort is negligible. Considering the sparse storage, the significant shell-pairs are directly identified by the displaced nuclei, so that direct scattering of the $O(1)$ significant elements into BCSR-format bears virtually no overhead compared to conventional methods.

Among the two-electron terms, the Coulomb term generally requires the most implementational effort when using the CFMM approach. Distinguishing between integral and density derivatives, the following two first-order Coulomb-derivatives must be evaluated during DL-CPSCF:

$$J_{\mu\nu}[\mathbf{X}] = \sum_{\lambda\sigma} (\mu\nu|\lambda\sigma)X_{\lambda\sigma}, \quad \mathbf{X} = (\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x) \text{ or } \mathbf{P}_{\text{oo}}^x, \quad (10)$$

$$J_{\mu\nu}^x[\mathbf{P}] = \sum_{\lambda\sigma} ([\mu\nu]^x|\lambda\sigma)P_{\lambda\sigma} + \sum_{\lambda\sigma} (\mu\nu|[\lambda\sigma]^x)P_{\lambda\sigma}. \quad (11)$$

The first matrix (Eq. (10)) is computed using standard CFMM routines where the far-field is built from the perturbed densities \mathbf{P}_{oo}^x or $\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x$. No special routines have been developed to reduce the loops over the constant number of significant boxes to $O(1)$ due to the small prefactor of the far-field part compared to the near-field part (see Sec. V A). However, the strategy suggested in Ref. 59 could be employed in order to reduce the computational effort. The near-field contribution is computed in $O(1)$ time by preselecting only those neighboring boxes which contain basis functions with significant response density elements.

For the Coulomb matrix $\mathbf{J}^x[\mathbf{P}]$ (Eq. (11)), the first term is obtained easily by a modification of the CFMM gradient routine.⁷² In our current implementation, the evaluation of the far-field contribution scales linearly with system size. Again, due to the small prefactor, this step is only a small part of the overall computational effort, but one could easily avoid this step by storing the final far-field from the preceding SCF calculation. However, for the evaluation of the second term of $\mathbf{J}^x[\mathbf{P}]$, a new CFMM routine for the generation of the perturbed far-field due to the displaced basis functions was implemented. Note that the conversion to and the handling of the Taylor expansion (steps 2–4 in Ref. 8) are not affected by nuclear displacement, only the multipole as it is generated and shifted to the center of the lowest-level box. Thus, the derivative of $|\lambda\sigma\rangle P_{\lambda\sigma}$ is formed as described for the bra-term in Ref. 72. The multipoles are then converted to Taylor expansions as in the original algorithm. Note that we have to perform steps 2–4 of the CFMM algorithm for each perturbation of the batch.

The exchange terms occurring within the CPSCF,

$$K_{\mu\nu}[\mathbf{X}] = \sum_{\lambda\sigma} (\mu\lambda|\nu\sigma)X_{\lambda\sigma}, \quad \mathbf{X} = (\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x) \text{ or } \mathbf{P}_{\text{oo}}^x,$$

$$K_{\mu\nu}^x[\mathbf{P}] = \sum_{\lambda\sigma} ([\mu\lambda]^x|\nu\sigma)P_{\lambda\sigma} + \sum_{\lambda\sigma} (\mu\lambda|[\nu\sigma]^x)P_{\lambda\sigma},$$

are evaluated with the LinK^{20,21} algorithm in $O(1)$ time, made possible by the \mathbf{P}/\mathbf{P}^x -coupling of the bra and ket distributions.

2. Second-order derivatives

The second-order derivatives in Eq. (1) are never formed explicitly but directly contracted with \mathbf{P} and \mathbf{P}^x to the corresponding contribution to the Hessian matrix elements. We use standard routines for the second-order derivatives of the one-electron terms \mathbf{T}^{xy} , $\mathbf{V}_{\text{en}}^{xy}$, and \mathbf{S}^{xy} . The mixed integral/density derivatives ($\text{Tr}[\mathbf{P}^y\mathbf{K}^x[\mathbf{P}]]$, for example) are obtained by a modification of the standard gradient routines in order to support density matrices in sparse format.

We have implemented linear-scaling integral routines to evaluate $\mathbf{J}^{xy}[\mathbf{P}]$ and $\mathbf{K}^{xy}[\mathbf{P}]$, respectively. We fell back on existing routines to evaluate $\mathbf{V}_{\text{xc}}^{xy}[\mathbf{P}]$.⁷³ For the exchange contribution $\mathbf{K}^{xy}[\mathbf{P}]$, an adaptation of the LinK gradient routine²¹ is straightforward. However, for the Coulomb term $\mathbf{J}^{xy}[\mathbf{P}]$,

$$\begin{aligned} \text{Tr}[\mathbf{P}\mathbf{J}^{xy}[\mathbf{P}]] = & 2 \sum_{\substack{\mu\nu \\ \lambda\sigma}} P_{\mu\nu}([\mu\nu]^{xy}|\lambda\sigma)P_{\lambda\sigma} \\ & + P_{\mu\nu}([\mu\nu]^x|[\lambda\sigma]^y)P_{\lambda\sigma}, \end{aligned} \quad (12)$$

two new CFMM routines were implemented along the lines of Ref. 59. For the first term in Eq. (12), the unperturbed far-field must contract with the second-order derivatives of the bra charge-distributions, where the derivative equations are obtained by a straightforward derivation of the gradient expressions in Ref. 72. The second term involves the interaction of the perturbed far-field and the perturbed bra charge-distributions, where the routines can be composed from the gradient algorithm⁷² and the far-field algorithm as it is used for the evaluation of $\mathbf{J}^x[\mathbf{P}]$ in Eq. (11). As with the first-order Coulomb derivatives, the far-field contribution scales as $O(N)$ per perturbation in our present implementation.

V. EXEMPLARY CALCULATIONS AND DISCUSSION

We have implemented our method in a development version of Q-Chem⁷⁴ to illustrate performance and scaling. Exemplary calculations showcasing the effective $O(N)$ construction (in the rate-determining steps) of the complete nuclear Hessian matrix were performed on a series of linear alkanes and amylose chains of increasing length. As a more demanding model system, we used stacked DNA base pairs and calculated the response density matrices due to the displacement of a single nucleus. In the following, scaling exponents reported for alkanes and amylose (Secs. V A and V B) are for the *complete* Hessian matrix, whereas scaling exponents for DNA fragments (Sec. V C) are for a *single nucleus* due to the high computational demand.

All calculations were performed at the Hartree–Fock level of theory, using the 6-31G(d) basis set and an integral cutoff threshold of 10^{-8} a.u. The sparsity criterion was set to 10^{-7} (Frobenius norm, modified BCSR format with an average 100×100 blocksize). Seven Laplace points using the min-max fitting⁶⁹ were employed for the numerical integration. For $\text{C}_{10}\text{H}_{22}$, $\text{C}_{20}\text{H}_{42}$, and $\text{C}_{40}\text{H}_{82}$, this provides a combined root mean square deviation of $\sim 0.2 \text{ cm}^{-1}$ in comparison to MO-based CPSCF without CFMM/LinK. All calculations were performed on a dual Intel Xeon E5-2620 machine using a single thread of execution. No molecular symmetry was exploited. All times given are actual wall times. Our method is not disk-bound, so processor and wall times are very close. Percentages

and scaling exponents are calculated from unrounded numbers. Coordinate files for all systems can be found online.⁷⁵

A. *n*-alkanes

The execution times for the creation of the complete molecular vibrational Hessian matrix for linear alkanes of increasing length are given in Table I. The onset of linear scaling occurs as early as 40 carbon atoms for this class of highly sparse systems. Near-linear scaling is observed when moving from $\text{C}_{40}\text{H}_{82}$ to $\text{C}_{80}\text{H}_{162}$ and beyond. Note that even for the largest system, the projection and diagonalization of the Hessian matrix take just 21 s.

The wall times for $\text{C}_{80}\text{H}_{162}$ and $\text{C}_{160}\text{H}_{322}$ broken down into detailed contributions are given in Table II. In both cases, the $\mathbf{J}[\mathbf{P}^x]$ and $\mathbf{K}[\mathbf{P}^x]$ contributions dominate, consuming around 70% of the total time. The three most computationally expensive terms ($\mathbf{K}[\mathbf{P}^x]$, the near-field part of $\mathbf{J}[\mathbf{P}^x]$, and $\mathbf{J}^x[\mathbf{P}^y]$ which contribute a combined 70% of the overall time) are all within the linear-scaling regime. The most time-consuming term with quadratic scaling behavior is the far-field part of $\mathbf{J}[\mathbf{P}^x]$ (which must be evaluated every CPSCF iteration for each perturbation) requires only about 4% of the total time for $\text{C}_{160}\text{H}_{322}$. This is in contrast to $\mathbf{J}^x[\mathbf{P}^y]$, which must be evaluated only once per perturbation, and consequently takes less total time. We note that all exchange-type contributions ($\mathbf{K}[\mathbf{P}^x]$, $\mathbf{K}^x[\mathbf{P}]$, $\mathbf{K}^x[\mathbf{P}^y]$, and $\mathbf{K}^{xy}[\mathbf{P}]$), as well as all near-field Coulomb-type contributions show linear or near-linear scaling, as does the DL-CPSCF iteration step (Eq. (4a)). Figure 1 shows the wall times and scaling exponents graphically, including results from conventional integral-direct MO-CPSCF without LinK or CFMM. In the case of MO-CPSCF, the timings for $\text{C}_{80}\text{H}_{162}$ and $\text{C}_{160}\text{H}_{322}$ were extrapolated using a conservative $O(N^3)$ scaling. The crossover of our method with the conventional algorithm is at around 150 atoms.

For visual reference, sparsity patterns of several matrices occurring during the calculation of $\text{C}_{80}\text{H}_{162}$ are shown in Fig. 2. The response density \mathbf{P}^x is local and changes very little compared to the corresponding matrix of $\text{C}_{40}\text{H}_{82}$. This confirms our observations concerning the range of the perturbative effect on the electronic structure. The interaction decays below the matrix sparsity threshold before reaching the far end of the molecule, i.e., is constant with respect to further system size increases. Also shown in Fig. 2 is the effect of explicit reprojection of $\mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]$, which increases matrix sparsity.

TABLE I. Wall times and scaling behavior (relative to previous entry) for the calculation of the complete nuclear Hessian matrix for linear alkanes ($\text{C}_n\text{H}_{2n+2}$, HF/6-31G(d)). The last two rows (labeled “conventional”) contain data from integral-direct MO-CPSCF without LinK/CFMM, where the data in parentheses were extrapolated.

Molecule	$\text{C}_{10}\text{H}_{22}$	$\text{C}_{20}\text{H}_{42}$	$\text{C}_{40}\text{H}_{82}$	$\text{C}_{80}\text{H}_{162}$	$\text{C}_{160}\text{H}_{322}$
N_{atoms}	32	62	122	242	482
N_{basis}	194	384	764	1524	3044
Wall time (h)	0.53	2.61	8.78	23.1	60.2
$O(N^x)$...	2.31	1.75	1.40	1.38
Wall time (h) (conventional)	0.11	0.84	7.20	(57.6)	898.5/(460.8)
$O(N^x)$ (conventional)	...	2.88	3.09	(3)	3.49 ^a /(3)

^aScaling exponent with respect to $\text{C}_{40}\text{H}_{82}$.

TABLE II. Detailed wall times in minutes (percent of total time in parentheses) and scaling behavior for the complete nuclear Hessian matrix of $C_{80}H_{162}$ and $C_{160}H_{322}$.

	$C_{80}H_{162}$		$C_{160}H_{322}$		$O(N^x)$
$\mathbf{K}[\mathbf{P}^x]$	668	(48%)	1525	(42%)	1.19
$\mathbf{J}[\mathbf{P}^x]$ (near-field)	341	(25%)	815	(23%)	1.26
$\mathbf{J}^x[\mathbf{P}^y]$ (near-field)	90	(7%)	197	(5%)	1.13
$\mathbf{J}[\mathbf{P}^x]$ (far-field)	44	(3%)	158	(4%)	1.85
Equation (4a)	66	(5%)	157	(4%)	1.26
$\mathbf{J}^x[\mathbf{P}^y]$ (far-field)	22	(2%)	86	(2%)	1.95
$\mathbf{K}^x[\mathbf{P}]$	24	(4%)	57	(2%)	1.27
$\mathbf{K}^x[\mathbf{P}^y]$	18	(1%)	48	(1%)	1.39
$\mathbf{J}^x[\mathbf{P}]$ (near-field)	7.8	(0.5%)	19	(0.5%)	1.28
$\mathbf{J}^x[\mathbf{P}]$ (far-field)	4.2	(0.3%)	16	(0.5%)	1.97
$\mathbf{J}^x[\mathbf{P}]$ (near-field)	6.2	(0.4%)	13	(0.3%)	1.03
$\mathbf{K}^x[\mathbf{P}]$	5.3	(0.4%)	11	(0.3%)	1.03
$\mathbf{J}^x[\mathbf{P}]$ (far-field)	2.3	(0.1%)	9.1	(0.3%)	1.97

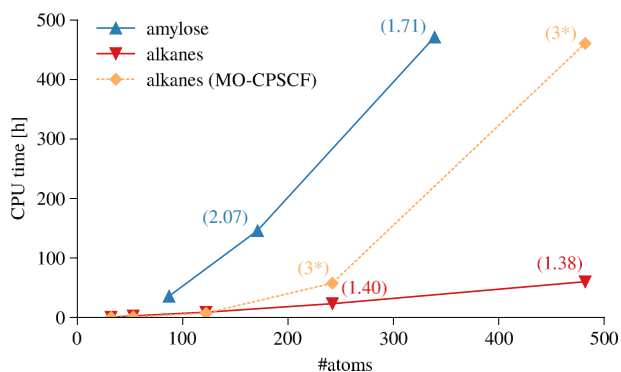


FIG. 1. Wall times for the calculation of the complete nuclear Hessian matrix for linear alkanes and amylose chains (HF/6-31G(d)). The numbers in parentheses are the observed scaling exponents $O(N^x)$ relative to the previous system. The asterisk indicates the timing data were extrapolated from the scaling exponent in parentheses. Refer to Tables I and III for numerical data.

B. Amylose chains

We computed the vibrational Hessian matrix for amylose chains of 4, 8, and 16 glucose units. Timing data and scaling exponents are given in Table III. The largest of these molecules has 339 atoms and 2979 basis functions. A scaling behavior which is considerably better than quadratic ($N_{\text{basis}}^{1.7}$) is observed when comparing the eight- and sixteen-unit amylose chains. The relative computational costs of the individual terms are listed in Table IV. They are essentially very similar to those of the alkanes. The exchange-type contractions as well

TABLE III. Wall times and scaling behavior (relative to previous entry) for the calculation of the complete nuclear Hessian matrix of amylose chains (HF/6-31G(d)).

Molecule	amylose ₄	amylose ₈	amylose ₁₆
N_{atoms}	87	171	339
N_{basis}	759	1499	2979
Wall time (h)	35.7	146.0	471.4
$O(N^x)$...	2.07	1.71

as the near-field contributions to the Coulomb-type contractions exhibit linear or near-linear scaling, while the far-field Coulomb-type contributions show, as discussed above, quadratic scaling with small prefactors. The DL-CPSCF iteration step does not show linear-scaling between 8 and 16 glucose units, because the matrices involved are not yet sufficiently sparse for these systems sizes. This is not a serious problem since this step constitutes only around 5% of the total computational time for the largest system. The scaling behavior of matrix multiplications improves for larger systems with more sparse matrices (cf. Sec. V A). Figure 1 shows the wall times and scaling exponents graphically.

C. DNA base pairs

We calculated the perturbed density matrices for all three cartesian directions of a single atom in stacked DNA base pairs of increasing size. The chosen perturbed nuclei are located on the periphery of each DNA molecule. Computation of these large, realistic, and very representative systems are challenging because of their demanding electronic structure. Full linear-scaling is only achieved for very large systems even in the underlying SCF calculation. The timing data are given in Table V. We observe linear scaling between increasing system sizes, leading to overall quadratic scaling for the complete Hessian matrix. Note that for the given fragment sizes, the number of significant elements in the first-order matrices still grows linearly with system size, i.e., our method provides quadratic scaling effort for the construction of the nuclear Hessian. This already substantially improves the conventional $O(N^4)$ scaling for the integral-direct CPSCF⁷ solved in the MO basis.

Detailed contributions are given in Table VI. The dominating contribution, the contraction of exchange integrals with the response density matrices ($\mathbf{K}[\mathbf{P}^x]$), shows sub-linear scaling between eight and sixteen DNA base pairs. The corresponding Coulomb-type contraction $\mathbf{J}[\mathbf{P}^x]$ scales linearly in the near-field contribution and somewhat worse than linear in




FIG. 2. Matrix sparsity patterns from $C_{80}H_{162}$ (HF/6-31G(d), modified BCSR format, sparsity threshold 10^{-7}). White areas are numerically insignificant and are not allocated in memory. Darker shades mean higher absolute matrix elements. (a) Response density \mathbf{P}^x . (b) Unprojected matrix of two-electron integral contractions $\mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x]$. (c) Subspace-projected matrix $(\mathbf{I} - \mathbf{SP})\mathbf{G}[\mathbf{P}_{\text{vo}}^x + \mathbf{P}_{\text{ov}}^x](\mathbf{PS})$. See text for a discussion of the long-range behavior.

TABLE IV. Detailed wall times in minutes (percent of total time in parentheses) and scaling behavior for the complete nuclear Hessian matrix of amylose chains (HF/6-31G(d)).

	amylose ₈		amylose ₁₆		$O(N^x)$
$\mathbf{K}[\mathbf{P}^x]$	3637.3	(42%)	10 814.9	(38%)	1.59
$\mathbf{J}[\mathbf{P}^x]$ (near-field)	3365.8	(38%)	10 759.3	(38%)	1.69
$\mathbf{J}^x[\mathbf{P}^y]$ (near-field)	743.5	(8.5%)	2313.1	(8.2%)	1.65
Equation (4a)	323.3	(3.7%)	1547.8	(5.5%)	2.28
$\mathbf{J}[\mathbf{P}^x]$ (far-field)	176.3	(2.0%)	865.7	(3.1%)	2.32
$\mathbf{K}^x[\mathbf{P}]$	69.1	(0.8%)	161.3	(0.6%)	1.24
$\mathbf{J}^x[\mathbf{P}^y]$ (far-field)	32.6	(0.4%)	149.8	(0.5%)	2.22
$\mathbf{K}^x[\mathbf{P}^y]$	58.4	(0.7%)	145.3	(0.5%)	1.33
$\mathbf{J}^x[\mathbf{P}]$ (near-field)	39.6	(0.5%)	86.7	(0.3%)	1.14
$\mathbf{J}^{xy}[\mathbf{P}]$ (near-field)	28.2	(0.4%)	56.5	(0.2%)	1.01
$\mathbf{K}^{xy}[\mathbf{P}]$	13.9	(0.2%)	28.8	(0.1%)	1.06
$\mathbf{J}^{xy}[\mathbf{P}]$ (far-field)	3.9	(<0.1 %)	14.5	(0.1%)	1.92
$\mathbf{J}^x[\mathbf{P}]$ (far-field)	3.1	(<0.1 %)	12.2	(<0.1 %)	2.00

TABLE V. Wall times and scaling behavior (relative to previous entry) for the calculation of the perturbed density matrices \mathbf{P}^x (three cartesian directions) for a single atom in stacked adenine–thymine DNA base pairs (HF/6-31G(d)).

Molecule	DNA ₄	DNA ₈	DNA ₁₆
N_{atoms}	260	524	1052
N_{basis}	2598	5290	10 674
Wall time (h)	2.09	4.46	9.16
$O(N^x)$...	1.07	1.02

the far-field contribution. This is consistent with the matrix sparsity of \mathbf{P}^x : for DNA₈, \mathbf{P}^x has around 90% significant elements (as measured by the number of significant blocks), while \mathbf{P}^x for DNA₁₆ has around 46%. This corresponds to an approximately linear growth in the number of significant elements in the matrix from DNA₈ to DNA₁₆. In other words, the perturbation range has not decayed below the sparsity threshold for these systems, and the best possible scaling would be linear per perturbation.

Importantly, these calculations illustrate the importance of utilizing sparse linear algebra and linear-scaling CPSCF algorithms such as DL-CPSCF for large-scale molecular property calculations. The largest of the calculations uses more than 10 000 basis functions. Consequently, matrix multiplications take considerable portions of the total wall time, reaching over 30% just for the DL-CPSCF iteration step (Eq. (4a)). Dense

matrix multiplications used in conventional CPSCF algorithms would dominate the computational expense very quickly for larger systems due to their $O(N^3)$ complexity. In contrast, sparse linear algebra shows decreasing complexity for larger systems, i.e., more sparse matrices, as shown for linear alkanes in Table II.

VI. CONCLUSION AND OUTLOOK

We have presented an AO-based CPSCF scheme for calculating the nuclear Hessian for molecular systems which opens the way to linear scaling in all time-determining steps. The key step making this possible is circumventing the long-range (r^{-1}) coupling through cancellation of far-field interactions between two perturbed Coulomb-type contractions of the two-electron integrals. Multipole expansions of the corresponding terms reveal a r^{-2} -decay, which is much more short ranged, so that all occurring first-order matrices in the DL-CPSCF equations contain an asymptotically constant number of significant elements. No approximations are introduced in the derivation.

We implemented the method, reducing the complexity of the rate-determining steps in the calculation of the complete nuclear displacement Hessian matrix to linear. Using DL-CPSCF theory combined with the CFMM and LinK integral evaluation schemes, obtaining the response of the electronic structure to a single nuclear displacement becomes possible with sub-linear complexity.

TABLE VI. Detailed wall times in minutes (percent of total time in parentheses) and scaling behavior for the calculation of the perturbed density matrices \mathbf{P}^x (three cartesian directions) for a single atom in stacked adenine–thymine DNA base pairs (HF/6-31G(d)).

	DNA ₈		DNA ₁₆		$O(N^x)$
$\mathbf{K}[\mathbf{P}^x]$	124.2	(45.8%)	194.4	(34.8%)	0.64
Equation (4a)	61.3	(22.6%)	169.8	(30.4%)	1.45
$\mathbf{J}[\mathbf{P}^x]$ (near-field)	66.8	(24.6%)	137.4	(24.6%)	1.03
$\mathbf{J}[\mathbf{P}^x]$ (far-field)	5.5	(2.0%)	16.1	(2.9%)	1.54
$\mathbf{K}^x[\mathbf{P}]$	0.9	(0.3%)	2.6	(0.5%)	1.45
$\mathbf{J}^x[\mathbf{P}]$ (near-field)	0.5	(0.2%)	1.2	(0.2%)	1.18
$\mathbf{J}^x[\mathbf{P}]$ (far-field)	0.3	(0.1%)	0.9	(<0.1 %)	1.45

Like all density matrix-based algorithms, our method requires a local electronic structure, i.e., a non-zero band gap (HOMO-LUMO gap) to achieve linear scaling. Furthermore, exploiting the long-range cancellation in the Coulombic contribution using the CFMM is difficult. In our implementation, evaluation of the far-field contributions to Coulomb-type matrices has quadratic scaling with a very small prefactor. For our realistic and difficult test system (stacked DNA base pairs with over 10 000 basis functions), evaluation of the far-field contribution takes less than 3% of the total CPSCF time. This step is expected to dominate the computational cost only for extremely large systems which are out of reach for the time being. The r^{-2} decay laid out in Sec. III C implies, however, that these contributions can, in principle, also be calculated with $O(N)$ scaling.

Finally, the eigendecomposition of the nuclear displacement Hessian matrix, which reveals the actual vibrational frequencies and scales as $O(N_{\text{atoms}}^3)$ with a very small prefactor, has not been studied in this work.

While in the present paper, we have focused on HF, it is expected that an $O(1)$ (overall $O(N)$) DFT implementation can be achieved along similar lines due to the local nature of current exchange-correlation functionals.

We note the evaluation of the nuclear Hessian matrix is an easily parallelizable problem. Evaluation of several batches of nuclei in parallel with shared- or distributed-memory approaches would yield large speedups on many-processor machines. Using the proposed method, fully analytical vibrational frequency calculations of large molecular systems are now accessible.

ACKNOWLEDGMENTS

We thank Dr. Simon A. Maurer and Professor M. Head-Gordon for helpful discussions, and a reviewer for bringing Ref. 70 to our attention. C.O. acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG) for the Project No. Oc35/4-1. Further funding was provided by the DFG via the initiatives SFB 749 “Dynamik und Intermediate molekularer Transformationen” (project C7) and the DFG cluster of excellence EXC 114 “Center for Integrated Protein Science Munich” (CIPSM).

- ¹P. Pulay, *Mol. Phys.* **17**, 197 (1969).
- ²P. Pulay, *WIREs Comput. Mol. Sci.* **4**, 169 (2014).
- ³J. P. Merrick, D. Moran, and L. Radom, *J. Phys. Chem. A* **111**, 11683 (2007).
- ⁴J. Gerratt and I. Mills, *J. Chem. Phys.* **49**, 1719 (1968).
- ⁵J. A. Pople, R. Krishnan, H. B. Schlegel, and J. S. Binkley, *Int. J. Quantum Chem.* **16**, 225 (1979).
- ⁶Y. Osamura, Y. Yamaguchi, P. Saxe, D. Fox, M. Vincent, and H. Schaefer, *J. Mol. Struct.: THEOCHEM* **103**, 183 (1983).
- ⁷M. Frisch, M. Head-Gordon, and J. Pople, *Chem. Phys.* **141**, 189 (1990).
- ⁸C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, *Chem. Phys. Lett.* **230**, 8 (1994).
- ⁹C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, *Chem. Phys. Lett.* **253**, 268 (1996).
- ¹⁰M. Challacombe, E. Schwegler, and J. Almlöf, *J. Chem. Phys.* **104**, 4685 (1996).
- ¹¹M. Challacombe and E. Schwegler, *J. Chem. Phys.* **106**, 5526 (1997).
- ¹²L. Füsti-Molnár and P. Pulay, *J. Chem. Phys.* **116**, 7795 (2002).
- ¹³L. Füsti-Molnár and P. Pulay, *J. Chem. Phys.* **117**, 7827 (2002).
- ¹⁴L. Füsti-Molnár and P. Pulay, *J. Mol. Struct.: THEOCHEM* **666-667**, 25 (2003).
- ¹⁵L. Füsti-Molnár, *J. Chem. Phys.* **119**, 11080 (2003).
- ¹⁶L. Füsti-Molnár and J. Kong, *J. Chem. Phys.* **122**, 074108 (2005).
- ¹⁷J. C. Burant, G. E. Scuseria, and M. J. Frisch, *J. Chem. Phys.* **105**, 8969 (1996).
- ¹⁸E. Schwegler and M. Challacombe, *J. Chem. Phys.* **105**, 2726 (1996).
- ¹⁹E. Schwegler, M. Challacombe, and M. Head-Gordon, *J. Chem. Phys.* **106**, 9708 (1997).
- ²⁰C. Ochsenfeld, C. A. White, and M. Head-Gordon, *J. Chem. Phys.* **109**, 1663 (1998).
- ²¹C. Ochsenfeld, *Chem. Phys. Lett.* **327**, 216 (2000).
- ²²F. Neese, F. Wennmohs, A. Hansen, and U. Becker, *Chem. Phys.* **356**, 98 (2009).
- ²³J. M. Pérez-Jordá and W. Yang, *Chem. Phys. Lett.* **241**, 469 (1995).
- ²⁴R. E. Stratmann, G. E. Scuseria, and M. J. Frisch, *Chem. Phys. Lett.* **257**, 213 (1996).
- ²⁵X.-P. Li, R. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10891 (1993).
- ²⁶M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).
- ²⁷R. Nunes and D. Vanderbilt, *Phys. Rev. B* **50**, 17611 (1994).
- ²⁸C. A. White, P. Maslen, M. S. Lee, and M. Head-Gordon, *Chem. Phys. Lett.* **276**, 133 (1997).
- ²⁹J. M. Millam and G. E. Scuseria, *J. Chem. Phys.* **106**, 5569 (1997).
- ³⁰A. D. Daniels, J. M. Millam, and G. E. Scuseria, *J. Chem. Phys.* **107**, 425 (1997).
- ³¹Y. Shao, C. Saravanan, M. Head-Gordon, and C. A. White, *J. Chem. Phys.* **118**, 6144 (2003).
- ³²H. Larsen, J. Olsen, P. Jørgensen, and T. Helgaker, *J. Chem. Phys.* **115**, 9685 (2001).
- ³³A. Palser and D. Manolopoulos, *Phys. Rev. B* **58**, 12704 (1998).
- ³⁴A. Niklasson, *Phys. Rev. B* **66**, 155115 (2002).
- ³⁵A. M. N. Niklasson, C. J. Tymczak, and M. Challacombe, *J. Chem. Phys.* **118**, 8611 (2003).
- ³⁶D. K. Jordan and D. A. Mazziotti, *J. Chem. Phys.* **122**, 84114 (2005).
- ³⁷C. Ochsenfeld and M. Head-Gordon, *Chem. Phys. Lett.* **270**, 399 (1997).
- ³⁸C. Ochsenfeld, J. Kussmann, and F. Koziol, *Angew. Chem., Int. Ed.* **43**, 4485 (2004).
- ³⁹J. Kussmann and C. Ochsenfeld, *J. Chem. Phys.* **127**, 204103 (2007).
- ⁴⁰J. Kussmann and C. Ochsenfeld, *J. Chem. Phys.* **127**, 054103 (2007).
- ⁴¹H. Larsen, P. Jørgensen, J. Olsen, and T. Helgaker, *J. Chem. Phys.* **113**, 8908 (2000).
- ⁴²H. Larsen, T. Helgaker, J. Olsen, and P. Jørgensen, *J. Chem. Phys.* **115**, 10344 (2001).
- ⁴³S. Coriani, S. Høst, B. Jansík, L. Thøgersen, J. Olsen, P. Jørgensen, S. Reine, F. Pawłowski, T. Helgaker, and P. Sałek, *J. Chem. Phys.* **126**, 154108 (2007).
- ⁴⁴A. Niklasson and M. Challacombe, *Phys. Rev. Lett.* **92**, 193001 (2004).
- ⁴⁵V. Weber, A. Niklasson, and M. Challacombe, *Phys. Rev. Lett.* **92**, 193002 (2004).
- ⁴⁶V. Weber, A. M. N. Niklasson, and M. Challacombe, *J. Chem. Phys.* **123**, 044106 (2005).
- ⁴⁷A. M. N. Niklasson, V. Weber, and M. Challacombe, *J. Chem. Phys.* **123**, 044107 (2005).
- ⁴⁸M. Beer and C. Ochsenfeld, *J. Chem. Phys.* **128**, 221102 (2008).
- ⁴⁹M. Beer, J. Kussmann, and C. Ochsenfeld, *J. Chem. Phys.* **134**, 074102 (2011).
- ⁵⁰C. C. J. Roothaan, *Rev. Mod. Phys.* **23**, 69 (1951).
- ⁵¹G. G. Hall, *Proc. R. Soc. London A* **205**, 541 (1951).
- ⁵²J. L. Dodds, R. McWeeny, and A. J. Sadlej, *Mol. Phys.* **34**, 1779 (1977).
- ⁵³M. Häser, *Theor. Chim. Acta* **87**, 147 (1993).
- ⁵⁴B. Doser, D. S. Lambrecht, J. Kussmann, and C. Ochsenfeld, *J. Chem. Phys.* **130**, 64107 (2009).
- ⁵⁵S. Schweizer, B. Doser, and C. Ochsenfeld, *J. Chem. Phys.* **128**, 154101 (2008).
- ⁵⁶P. Pulay, *J. Comput. Chem.* **3**, 556 (1982).
- ⁵⁷P. R. Surján, *Chem. Phys. Lett.* **406**, 318 (2005).
- ⁵⁸J. C. Burant, M. C. Strain, G. E. Scuseria, and M. J. Frisch, *Chem. Phys. Lett.* **248**, 43 (1996).
- ⁵⁹J. C. Burant, M. C. Strain, G. E. Scuseria, and M. J. Frisch, *Chem. Phys. Lett.* **258**, 45 (1996).
- ⁶⁰W. Liang, Y. Shao, C. Ochsenfeld, A. T. Bell, and M. Head-Gordon, *Chem. Phys. Lett.* **358**, 43 (2002).
- ⁶¹T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley-VCH, Weinheim, 2000).
- ⁶²S. A. Maurer, D. S. Lambrecht, D. Flaig, and C. Ochsenfeld, *J. Chem. Phys.* **136**, 144107 (2012).
- ⁶³M. Challacombe, *Comput. Phys. Commun.* **128**, 93 (2000).

- ⁶⁴C. Saravanan, Y. Shao, R. Baer, P. N. Ross, and M. Head-Gordon, *J. Comput. Chem.* **24**, 618 (2003).
- ⁶⁵Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia, 2003).
- ⁶⁶L. Colombo and W. Sawyer, *Mater. Sci. Eng.: B* **37**, 228 (1996).
- ⁶⁷C. Moler and C. Van Loan, *SIAM Rev.* **45**, 3 (2003).
- ⁶⁸E. H. Rubensson and E. Rudberg, *J. Comput. Chem.* **32**, 1411 (2011).
- ⁶⁹A. Takatsuka, S. Ten-no, and W. Hackbusch, *J. Chem. Phys.* **129**, 044112 (2008).
- ⁷⁰E. H. Rubensson and S. Zahedi, *J. Chem. Phys.* **128**, 176101 (2008).
- ⁷¹S. A. Maurer, L. Clin, and C. Ochsenfeld, *J. Chem. Phys.* **140**, 224112 (2014).
- ⁷²Y. Shao, C. A. White, and M. Head-Gordon, *J. Chem. Phys.* **114**, 6572 (2001).
- ⁷³B. G. Johnson and M. J. Fisch, *J. Chem. Phys.* **100**, 7429 (1994).
- ⁷⁴Development version of the Q-Chem program package, <http://www.q-chem.com>.
- ⁷⁵Structures are available online www.cup.lmu.de/pc/ochsenfeld/.

Academic Teachers

- Prof. Dr. Roland Beckmann
- Prof. Dr. Hans-Christian Böttcher
- Prof. Dr. Christoph Bräuchle
- Prof. Dr. Jeroen Buters
- Prof. Dr. Thomas Carell
- Prof. Dr. Patrick Cramer
- Prof. Dr. Hubert Ebert
- Prof. Dr. Klaus Förstemann
- Prof. Dr. Peter Gilch
- Prof. Dr. Achim Hartschuh
- Prof. Dr. Manfred Heuschmann
- Prof. Dr. Dirk Johrendt
- Prof. Dr. Konstantin Karaghiosoff
- Prof. Dr. Thomas M. Klapötke
- Prof. Dr. Peter Klüfers
- Prof. Dr. Paul Knochel
- Prof. Dr. Andreas Kornath
- Prof. Dr. Heinz Langhals
- Prof. Dr. Bettina Lotsch
- Prof. Dr. Herbert Mayr
- Prof. Dr. Dierk Niessing

- Prof. Dr. Christian Ochsenfeld
- Prof. Dr. Oliver Oeckler
- Prof. Dr. Hans-Rudolf Pfaendler
- Prof. Dr. Regina de Vivie-Riedle
- Prof. Dr. Joost Winterlin
- Prof. Dr. Martin Wirsing
- Prof. Dr. Hendrik Zipse
- PD Dr. Diemo Ködderitzsch
- PD Dr. Albert Lötz
- PD Dr. Dietmar E. Martin
- PD Dr. Gerhard Mestl
- PD Dr. Ján Minár
- Dr. Asbjörn M. Burow
- Dr. Thomas C. Jagau
- Dr. Steffen Jost
- Dr. Jörg Kussmann
- Dr. Werner Spahl
- Dr. David. S. Stephenson