

Lower-scaling advanced correlation methods with local Resolution of Identity strategies

A large-scale parallel and in-memory Laplace-Transformed MP2 for both molecular and periodic systems

im Fachbereich Physik der Freien Universität Berlin eingereichte

Dissertation zur Erlangung des akademischen Grades

DOCTOR RERUM NATURALIUM

vorgelegt von

Arvid Conrad Ihrig

Berlin, 2017



MAX-PLANCK-GESELLSCHAFT

Freie Universität  Berlin

Erstgutachter (Betreuer): Prof. Dr. Matthias Scheffler
Zweitgutachter: Prof. Dr. Piet Brouwer
Tag der Disputation: 06.11.2017

ABSTRACT

First-principles methods are a valuable tool to study the electronic structure of systems and explore the quantum-mechanical mechanisms driving them on the microscopic scale. Semi-local Density-Functional Approximations (DFAs) are still very popular due to the good balance between accuracy and computational efficiency. However, the instantaneous electronic interactions are described with an averaged potential in these methods, which is often insufficient to describe charge transfers or subtle many-body correlations. A promising way to improve the description of such interactions is the usage of advanced correlation methods which explicitly treat the electronic interactions in a many-body picture. Well known methods form this class are e.g. the Random-Phase Approximation (RPA), 2nd Order Møller-Plesset Perturbation Theory (MP2) and Coupled-Cluster Theory (CC). Unfortunately, the time and memory requirements for all these methods suffer from a steep scaling with system size in their canonical formulations.

Several lower-scaling reformulations exist for these methods, but the at least $\mathcal{O}(N^3)$ memory requirements and the increased algorithmic complexity prevented any efficient parallelization of these methods so far. In my thesis I developed a new lower-scaling technique for advanced correlation methods to address this challenge. The key components of this technique are the Laplace-Transformation (LT) and a new local Resolution of Identity (RI) strategy, named RI-LVL. Together, they significantly reduce the memory requirement without sacrificing accuracy and allow for an efficient parallelization by means of *a priori* known sparsity patterns. In combination with established integral screening techniques, the computational scaling can also be reduced considerably. The potential of this new approach is demonstrated at the example of an in-memory LT-MP2 with a MPI/OpenMP hybrid parallelization. It features an at worst cubic computational scaling, a quadratic memory scaling and a high parallel efficiency. Based on the Numeric Atom-Centered Orbital (NAO) framework used in the Fritz Haber Institute ab initio molecular simulations (FHI-aims) code, this implementation can handle both molecular and periodic systems with the same numerical footing. I will also discuss in detail how the combination of the LT and RI-LVL can be applied to RPA calculations.

The accuracy and performance of my LT-MP2 is demonstrated in detail for both molecular and periodic systems. In addition, the convergence of the MP2 level adsorption energy of water on a rutile titanium dioxide (TiO_2) surface in the low coverage limit is explored. This new LT-MP2 recovers more than 99.9% of the canonical MP2 correlation energy and can handle TiO_2 surfaces with up to 270 atoms in the unit cell.

ZUSAMMENFASSUNG

“First principles” Methoden wie Dichte-Funktional Theorie (DFT) sind ein wichtiges Werkzeug um die elektronische Struktur von Systemen und die ihr zugrunde liegenden quanten-mechanischen Effekte zu erforschen. Semi-lokale Approximationen für DFT sind aufgrund ihrer guten Balance zwischen Genauigkeit und Rechenaufwand immer noch sehr weit verbreitet. Allerdings wird die Coulomb-Wechselwirkung der Elektronen hierbei mit einem gemittelten Feld angenähert, welches oftmals nicht ausreicht um subtile Viel-Teilchen Wechselwirkungen akkurat zu modellieren. Um diese zu beschreiben ist es notwendig korrelierte Methoden zu verwenden, die diese explizit modellieren. Bekannte Methoden dieser Art sind beispielsweise die Random-Phase Approximation (RPA), Møller-Plesset Störungstheorie 2ter Ordnung (MP2) oder Coupled-Cluster Theorie. Der Rechenaufwand und RAM-Bedarf dieser Methoden skalieren allerdings ungünstig mit der Systemgröße und lassen diese schnell sehr kostspielig werden.

Es existieren mehrere Varianten, die die Skalierung der Rechenzeit reduzieren, aber der weiterhin hohe RAM-Bedarf und die zusätzliche Komplexität haben bisher eine effiziente Parallelisierung ebendieser verhindert. In meiner Doktorarbeit habe ich einen neuen Ansatz zur Reduktion des Skalierungsverhaltens korrelierter Methoden entwickelt, der eine effiziente Parallelisierung erlaubt. Die Hauptkomponenten sind eine Laplace-Transformation (LT) und RI-LVL, eine neue Variante des “Resolution of Identity” Ansatzes unter Ausnutzung von Lokalität. Daraus ergibt sich ein signifikant reduzierter RAM-Bedarf und eine im Voraus bekannte Struktur der dünnbesetzten Matrizen, wodurch eine effiziente Parallelisierung ermöglicht wird. Der nötige Rechenaufwand kann dramatisch reduziert werden wenn dieser Ansatz mit etablierten Techniken zur Vorauswahl der signifikanten Integrale kombiniert wird. Das Potential dieses neuen Ansatzes wird am Beispiel einer LT-MP2 mit effizienter MPI/OpenMP Hybrid-Parallelisierung demonstriert. Der Rechenaufwand skaliert maximal kubisch mit der Systemgröße und der RAM-Bedarf wächst nur quadratisch an. Die Implementierung im FHI-aims Software-Paket basiert auf Numerischen Atom-zentrierten Orbitalen (NAOs), welche es erlauben molekulare und periodische Systeme mit denselben numerischen Werkzeugen zu handhaben. Außerdem werde ich im Detail darlegen wie dieses neue Konzept auf die RPA übertragen werden kann.

Die Leistungsfähigkeit und Genauigkeit meiner neuen LT-MP2 wird detailliert für molekulare und periodische Systeme dargelegt. Außerdem wird die Konvergenz der MP2 Adsorptionsenergie eines isolierten Wasser-Moleküls auf einer Titandioxid-Oberfläche (Rutil) untersucht. Meine LT-MP2 Implementierung reproduziert mehr als 99.9% der kanonischen MP2 Korrelations-Energie und kann bis zu 270 Atome in der TiO_2 Oberfläche handhaben.

CONTENTS

1	INTRODUCTION	1
2	ELECTRONIC STRUCTURE THEORY	5
2.1	The Born-Oppenheimer Approximation	5
2.2	The world of DFT and WFT	6
2.3	Hartree-Fock	9
2.4	Canonical MP2	12
2.5	Variants of the MP2	14
2.6	Density Functional Theory	17
2.7	The Random Phase Approximation	21
2.8	Periodic Boundary conditions	23
2.9	Embedded Cluster Models for surface calculations . .	28
2.10	Resolution of Identity	30
2.11	Numeric Atom-Centered Orbitals	32
3	LOCALIZED RESOLUTION OF IDENTITY	35
3.1	Concept of the RI-LVL	35
3.2	The definition of the RI-LVL coefficients	36
3.3	The construction of the Auxiliary Basis Set	38
3.4	Formal Error Analysis	40
3.5	RI-LVL in a RI-V framework	41
3.6	RI-LVL in periodic systems	42
3.7	Accuracy of the RI-LVL	43
3.7.1	Gaussian Basis Set Comparison	44
3.7.2	HF calculations	46
3.7.3	MP2 calculations	47
3.8	PBE0 and RPA calculations	49
3.8.1	RI-LVL and heavier elements	50
3.9	Scaling analysis	51
3.10	Conclusions	54
4	LT-MP2 FOR MOLECULAR AND PERIODIC SYSTEMS	55
4.1	LT-MP2 in a NAO framework	56
4.2	The LT-MP2 for periodic systems	58
4.3	The integration points and weights	60
4.4	The generalized RI-LVL for the transformed ERIs . . .	61
4.5	Selecting the local ABS for the transformed ERIs	64
4.6	Formal scaling of the integral transformation	66
4.7	The integral screening criterion	67
4.8	Parallelization for the LT-MP2	70
5	ACCURACY AND PERFORMANCE OF THE LT-MP2	73
5.1	Accuracy	74
5.2	Scalability with system size and computational resources	74
5.3	LT-MP2 in the ISOL22 test set	80
6	LT-MP2 APPLIED TO WATER ADSORPTION ON TITANIUM DIOXIDE	85
6.1	Computational Setup	86
6.2	Results	88

6.3	Performance	93
7	OUTLOOK: APPLICATION TO THE RANDOM-PHASE AP- PROXIMATION	97
7.1	RPA in an RI-framework	97
7.2	Laplace-Transformed RPA for materials with band gap	98
7.3	Laplace-Transformed RPA with fractional occupation .	102
8	CONCLUSIONS	105
A	SUPPLEMENTARY DATA FOR RI-LVL ACCURACY ANALYSIS	109
B	PARAMETERS FOR THE LT-MP2	113
B.1	Quadrature and Integral Screening	113
B.2	The generalized RI-LVL	114
C	NOTATION CONVENTIONS AND ACRONYMS	121
C.1	Scaling parameters	121
C.2	Acronyms	121
D	SELBSTSTÄNDIGKEITSERKLÄRUNG	125
E	CURRICULUM VITAE	127
F	PUBLICATIONS	129
G	ACKNOWLEDGMENTS	131
H	BIBLIOGRAPHY	133

INTRODUCTION

“First-principles” electronic structure theory methods are a class of numerical approximations to the true Schrödinger equation without any system-dependent empirical parameters. Over the last decades these methods have become a powerful theoretical tool to explain experimental findings and more recently also to predict materials with desired properties. [1–5] Density-Functional Theory (DFT) contributed a lot to the success of first-principles methods because a multitude of Density-Functional Approximations (DFAs) exist which provide a reasonable accuracy and high computational efficiency. However, the approximate descriptions of the electronic correlations in local and hybrid DFAs are often not sufficiently accurate to describe some important physical and chemical properties driven by significant charge transfers or subtle many-body (or even strong) correlations.

It is often necessary to use advanced correlation methods which treat the correlation effects explicitly in a many-body picture to describe such electronic interactions accurately. One group of advanced correlation methods are high-level DFAs like the Random-Phase Approximation (RPA) [6, 7] and renormalized 2nd Order Perturbation Theory (rPT2) [7], which include information about unoccupied states in the exchange-correlation functional. Another group of advanced correlation methods, the Wave Function Theory (WFT), originates from the field of quantum chemistry. Popular methods from WFT include for example the 2nd Order Møller-Plesset Perturbation Theory (MP2) and Coupled-Cluster Theory with Singles, Doubles and perturbative Triples (CCSD(T)). The MP2 itself is known to be not sufficiently accurate for many systems of interest, [8, 9] but in recent years a new class of DFAs based on it has emerged, the double-hybrids. Following the idea of established hybrid functionals, these functionals add a weighted MP2-like contribution to correlation energy. Since the MP2 calculation is the computationally most complex part of the calculation, the computational performance of double-hybrids is similar to MP2, but it has been demonstrated that their accuracy can surpass the canonical MP2 by far. [10–12] In their canonical formulations, all advanced correlation methods suffer from a steep computational scaling with system size. RPA scales as $\mathcal{O}(N^6)$ in its original formulation and variants using a Resolution of Identity (RI) technique still scale with $\mathcal{O}(N^4)$. MP2 scales $\mathcal{O}(N^5)$ and CCSD(T), the “gold standard” of WFT, even has $\mathcal{O}(N^7)$ scaling. An alternative are Quantum Monte Carlo (QMC) techniques which rely on a statistical evaluation of the exact Schrödinger equation. QMC methods have been proven to be very accurate and feature a lower scaling with system size than the previously mentioned methods, but the effort re-

quired to convergence the statistical errors adds a large prefactor to these calculations. [13, 14]

A variety of lower-scaling reformulations exploiting the locality of the electronic interactions has been proposed in the literature. Methods based on different techniques like orbital transformations, system fragmentation or integral transformations have been demonstrated to reach a considerably reduced, sometimes even almost linear, scaling for e.g. MP2, RPA and Coupled-Cluster Theory (CC). [15–17] The reduced computational scaling however comes at the price of significantly more complex algorithms which are difficult to parallelize efficiently. Most of the lower-scaling implementations published so far feature only a very limited parallelization, e.g. by making use of thread-parallelized linear algebra libraries. Another challenge for the large-scale parallelization of first-principles methods is the usually at least a $\mathcal{O}(N^3)$ memory requirement. Sophisticated data distribution and communication patterns are necessary to obtain a high parallel efficiency and the *irregular* parallelism introduced by many lower-scaling algorithms further complicates this challenge.

To further extend the realm of systems that are accessible to advanced correlation methods it is therefore imperative to develop new algorithms and techniques that combine the benefits of lower-scaling techniques and large-scale parallelism. A promising approach to reach this aim are locality exploits for the Resolution of Identity (RI), an established technique from quantum chemistry. The RI approach reduces the computational effort for the evaluation of Electron Repulsion Integrals (ERIs), which is the most costly step in many advanced correlation methods. In the first part of my project I explored the benefits of locality exploits in the RI strategy for Hartree-Fock (HF) theory and hybrid functionals. This gave rise to the RI-LVL variant, which features a strongly reduced $\mathcal{O}(N^2)$ memory footprint without significant loss of accuracy. In the second part of the project, the RI-LVL is generalized for the use in advanced correlation methods and combined with the Laplace-Transformation (LT) technique. The resulting approach transforms methods like the MP2 and RPA into an atomic basis function framework that is more suitable for locality exploits than the canonical formulations in terms of molecular orbitals. In addition to the favorable $\mathcal{O}(N^2)$ memory consumption of the RI-LVL, these formulations also allow the efficient usage of screening techniques to reduce the computational scaling. At the example of the MP2 I will demonstrate the practical uses of my new technique. The resulting implementation exhibits an at worst cubic computational scaling and a quadratic memory requirement. The significantly reduced memory consumption enables the algorithm to work completely in memory with an efficient MPI/OpenMP hybrid parallelization utilizing more than 1000 cores. This new algorithm is using the Numeric Atom-Centered Orbital (NAO) framework in the Fritz Haber Institute ab initio molecular simulations (FHI-aims) code and can treat molecular and periodic systems on the same numerical ground. I also describe in detail how the same strategy can be

used to obtain an at most $\mathcal{O}(N^3)$ RPA method with $\mathcal{O}(N^2)$ memory consumption.

This thesis is organized as follows: An overview over the various available first-principles methods is given in chapter 2, where I also discuss those relevant for my project in more detail. In chapter 3 I introduce the RI-LVL, the localized RI strategy which is a vital component of my lower-scaling approach. Its accuracy and performance are demonstrated using various test systems and different first-principles methods. The concepts behind the combination of the RI-LVL and LT into an efficiently parallelized Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP2) implementation are described in detail in chapter 4. The accuracy and superior performance of this new algorithm compared to the canonical MP2 is presented in chapter 5. In chapter 6 the new LT-MP2 is applied to molecular adsorption processes on titanium dioxide surfaces to demonstrate its practical use for periodic systems. In particular the convergence behavior of the adsorption energy for water in the low coverage limit will be investigated. titanium dioxide surfaces are used in many catalytic reactions and have been studied extensively in both experiment and theory. (see e.g. [18–22]) As mentioned before, this new technique is not specific to the MP2 and its application to the RPA methods seems very promising. The future extension of this new approach to the RPA is discussed in detail in chapter 7. Finally, I summarize my findings in the conclusions.

The notation conventions and acronyms used throughout my thesis are listed in Appendix C, including the symbols used for analyzing time and memory scaling. Atomic units are used throughout the thesis. [23]

ELECTRONIC STRUCTURE THEORY

A central problem of electronic structure theory is the solution of the time-independent Schrödinger equation to obtain the wave-function Ψ for a many-particle system.

$$\hat{\mathcal{H}}\Psi(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}) = \mathcal{E}\Psi(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}) \quad (2.1)$$

$\hat{\mathcal{H}}$ denotes the full non-relativistic Hamiltonian containing the electronic and nuclear kinetic energy as well as the Coulomb terms for electron-electron, electron-nucleus and nucleus-nucleus interactions.

$$\begin{aligned} \hat{\mathcal{H}} = & - \sum_i \frac{1}{2} \nabla_i^2 - \sum_A \frac{1}{2M_A} \nabla_A^2 - \sum_i \sum_A \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} \\ & + \sum_{i<j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_{A<B} \frac{Z_B Z_A}{|\mathbf{R}_A - \mathbf{R}_B|} \end{aligned} \quad (2.2)$$

Sums over i and j involve all electrons in the system, those over A and B the nuclei. Apparently the solution of this equation is a very complex multi-dimensional problem which explicitly depends on the $3N_{elec}$ electronic coordinates and the $3N_{nuc}$ atomic coordinates. Solving this equation directly for anything but the simplest model systems is practically impossible and approximations are necessary.

In this chapter I will at first provide a short overview of both Wave Function Theory (WFT) and Density-Functional Theory (DFT). Afterwards, Hartree-Fock (HF) theory, 2nd Order Møller-Plesset Perturbation Theory (MP2) and Random-Phase Approximation (RPA) will be discussed in more detail for both molecular and periodic systems. Special attention will be given to the Resolution of Identity (RI) technique, which is an important tool to reduce the computational effort in the aforementioned methods and methods based on them. At the end of this chapter, I will briefly discuss Numeric Atom-Centered Orbitals (NAOs) which are employed by the Fritz Haber Institute ab initio molecular simulations (FHI-aims) code I use throughout this project.

2.1 THE BORN-OPPENHEIMER APPROXIMATION

A fundamental approximation shared by most electronic structure theory methods is the Born-Oppenheimer approximation. [23] It decouples the nuclear and electronic motion based on the fact that the nuclei move several orders of magnitude slower than the electrons. It allows to approximately factorize the total wave-function into an electronic and a nuclear part.

$$\Psi(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}) \approx \Psi_{elec}(\{\mathbf{r}_i\}; \{\mathbf{R}_A\}) \cdot \Psi_{nuc}(\{\mathbf{R}_A\}) \quad (2.3)$$

The electronic problem can then be solved for a fixed arrangement of the nuclei which do not change their position on the electronic time scale, i.e. Ψ_{elec} has a parametric dependence on the nuclear coordinates. This ansatz leads to the electronic Schrödinger equation with the electronic Hamiltonian $\hat{\mathcal{H}}_{elec}$.

$$\hat{\mathcal{H}}_{elec} \Psi_{elec}(\{\mathbf{r}_i\}) = \mathcal{E}_{elec} \Psi_{elec}(\{\mathbf{r}_i\}) \quad (2.4)$$

$$\begin{aligned} \hat{\mathcal{H}}_{elec} = & -\sum_i \frac{1}{2} \nabla_i^2 - \sum_i \sum_A \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} \\ & + \sum_{i < j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_{A < B} \frac{Z_B Z_A}{|\mathbf{R}_A - \mathbf{R}_B|} \end{aligned} \quad (2.5)$$

Finding proper approximate solutions to this equation is the goal of many electronic structure theory methods. Those named *ab initio* or “first-principles” methods do not contain any system-dependent empirical parameters.

The Born-Oppenheimer approximation is used throughout my thesis and thus only the electronic Hamiltonian is considered. As it can be seen from Equation 2.5, the nuclei are assumed to be classical point charges with fixed positions in this approximation. While the Born-Oppenheimer approximation is usually a good one for the ground-state calculations at zero temperature discussed here, it can introduce significant errors in finite temperature simulations and other scenarios where the assumption of fixed, classical positions for the nuclei is a bad approximation. In such cases additional corrections are necessary to include nuclear quantum effects. Due to the uncertainty principle, the nuclei will vibrate around their equilibrium positions even at zero temperature. This effect can have an important contribution to the total energy and is called the zero point energy. To lowest order, the phonon modes (or vibrational modes in molecules) can be approximated as quantum harmonic oscillators (quasi-harmonic approximation [24]). The zero point energy is then the contribution from each mode at zero phonon occupation. Furthermore the approximation of the nuclear probability density with a fixed, classical position becomes less and less valid with increasing temperature, in particular for light elements like hydrogen. The arising nuclear quantum effects can play an important role in Molecular Dynamics (MD) simulations already at room temperatures. (see e.g. Ref [25], in particular figure 13) Due to the neglect of phonon effects, the Born-Oppenheimer approximation without further corrections is generally not sufficient when phonon-driven mechanisms are investigated, e.g. phonon-mediated superconductivity. [26]

2.2 THE WORLD OF DFT AND WFT

A multitude of different first-principles methods exist to solve the electronic Schrödinger equation. Most first-principles approaches for the description of the electronic ground-state originate either from Density-Functional Theory (DFT) or Wave Function Theory (WFT). DFT uses the electronic ground-state density as the central quantity

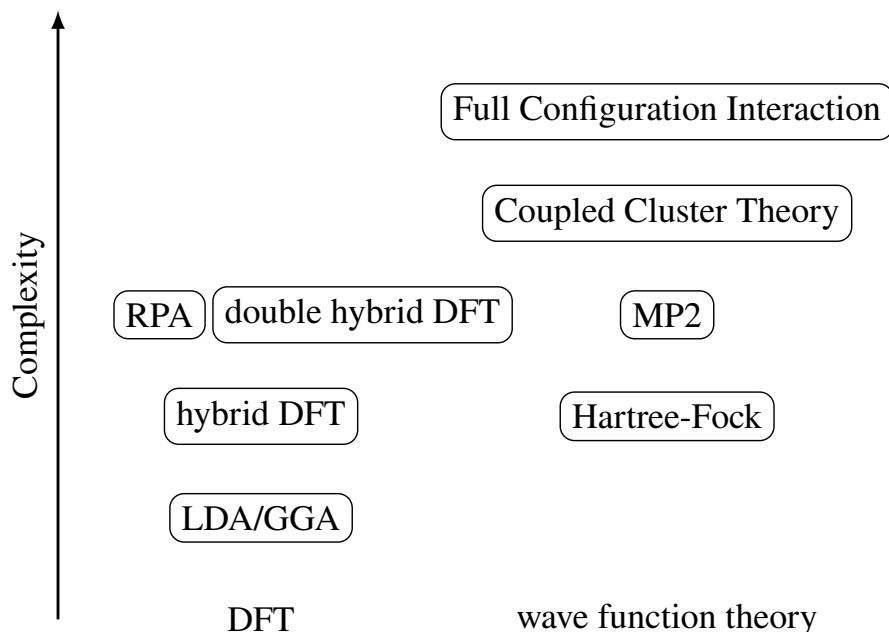


Figure 2.1: Overview of the common methods in DFT and quantum chemistry for the analysis of the electronic ground-state. A higher “complexity” implies an increased computational effort, but does not guarantee an improved accuracy. Many of these methods have been reformulated as lower-scaling variants, but these are generally a lot more complex in their implementation. Usually, only the canonical variants of these methods are available as efficiently parallelized algorithms that can utilize hundreds and thousands of cores on a computing cluster.

to describe the ground-state properties of the system, while wave-function methods like Hartree-Fock (HF) use approximations to the true wave-function as descriptor for the system of interest. Both feature a variety of different approximations that vary in computational cost and accuracy, as shown in Figure 2.1. The first-principles methods dedicated to the analysis of excited states like the *GW* approximation [27] and Time-Dependent Density Functional Theory [28] are beyond the scope of this thesis and will not be discussed.

Wave-function based methods follow a clear hierarchy of approximations. The simplest among them is the HF theory, which approximates the true wave-function as a single Slater determinant. It is the cheapest wave-function based method, but it completely misses electron correlation beyond the Pauli-Principle. A large fraction of the missing correlation energy in the HF solution can be recovered by means of perturbation theory, which yields the 2nd Order Møller-Plesset Perturbation Theory (MP2) and its higher order equivalents. MP2 is generally not considered adequate if very accurate energies are needed to determine e.g. chemical reaction barrier heights and rates. However, double hybrid functionals, which contain a weighted MP2 correlation contribution as their computationally most demanding term, have been demonstrated to reach much higher accuracies

than MP2 itself. [10, 12] HF and MP2 will be discussed in detail in the following sections.

Another method from WFT is the Full Configuration Interaction (FCI), which is a formally exact theory. FCI explicitly considers all excited states (“configurations”) that can be build from the HF solution by exciting one or more electrons into previously unoccupied orbitals. Unfortunately, the number of possible configurations grows exponentially with system size and the method is so expensive that it cannot be applied to any system with more than a few tens of electrons. Alavi combined FCI with Quantum Monte Carlo sampling techniques and was able to compute small periodic systems with about 60 electrons per unit cell. [29] To reduce the computational complexity, the Configuration Interaction (CI) expansion can be truncated to only include excited determinants up to a given order, resulting in e.g. Configuration Interaction with Single and Double Excitations (CISD). However, truncated CI variants cannot be applied to extended systems because they are not size-consistent [23], i.e. the energy does not scale linearly with the number of particles N when $N \rightarrow \infty$. A more robust and size-consistent alternative is the Coupled-Cluster Theory (CC). [30] In this approach the wave-function is written as an excitation operator $\exp(\hat{T})$ acting on the reference determinant, e.g. a HF solution. $\hat{T} = \hat{T}_1 + \hat{T}_2 + \hat{T}_3 + \dots$ contains all single, double, triple and so forth excitation operators. In contrast to the CI methods, one does not restrict the maximum excitation order of the considered determinants, but instead \hat{T} is truncated. CCSD for example only contains the single \hat{T}_1 and double \hat{T}_2 excitation operators. Since \hat{T} appears in an exponential operator, higher order determinants are still partially included, e.g. those quadruple excitations generated from two independent double-excitations. CCSD with a perturbative treatment of the triple excitations, known as CCSD(T), is widely considered as the “gold standard” of quantum chemistry. It is still a very demanding method with a high computational scaling and memory requirements, which is usually applied to small or medium-sized systems. In recent years some lower scaling variants have been developed too, e.g. the Domain-based local Pair Natural Orbital (DLPNO) based implementation from the Neese group. [17]

A different approach is the DFT which uses the ground-state density to describe the system instead of the full ground-state wave-function. The density is easier to handle because it has only 3 degrees of freedom instead of $3N$ like the ground-state wave-function and most properties of interest can be expressed in terms of the ground-state density too. A key challenge in DFT is the choice of an appropriate Density-Functional Approximation (DFA). A huge variety of DFAs exist, which differ in the approximations for the exchange correlation functional. Local Density Approximations (LDAs), Generalized Gradient Approximations (GGAs) and meta-Generalized Gradient Approximations (mGGAs) provide computationally very efficient DFAs which use only the electronic density and its derivatives as input. More sophisticated DFAs like the hybrid functionals, Random-Phase Approximation (RPA) and double hybrid functionals

also use the orbitals as input. DFT and the available DFA types will be discussed in more detail later in this chapter.

2.3 HARTREE-FOCK

The simplest wave-function based method is the HF theory. In this method an approximate solution to the electronic Schrödinger equation (Equation 2.4) is obtained by using a single Slater determinant as trial wave-function. A N -electron Slater determinant $|\Phi\rangle$ is built from N one-electron spin-orbitals $\chi_1 \dots \chi_N$

$$|\Phi\rangle = \frac{1}{\sqrt{N!}} \begin{vmatrix} \chi_1(\mathbf{x}_1) & \chi_2(\mathbf{x}_1) & \dots & \chi_N(\mathbf{x}_1) \\ \chi_1(\mathbf{x}_2) & \chi_2(\mathbf{x}_2) & \dots & \chi_N(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \chi_1(\mathbf{x}_N) & \chi_2(\mathbf{x}_N) & \dots & \chi_N(\mathbf{x}_N) \end{vmatrix} \quad (2.6)$$

where $\mathbf{x}_1 \dots \mathbf{x}_N$ denote the combined spatial and spin coordinates of the N electrons.

This choice of the wave-function ensures that the anti-symmetry principle and thus the Pauli-Exclusion is satisfied. The anti-symmetry principle states that the exchange of any two electron coordinates must yield the same wave-function but with inverse sign. As a consequence, two electrons cannot occupy the same spin-orbital χ_i as demanded by the Pauli-Exclusion principle. The energy of this wave-function is the expectation value $\langle \Phi | \hat{\mathcal{H}} | \Phi \rangle$ of the electronic Hamiltonian $\hat{\mathcal{H}}$ as defined in Equation 2.5. Using functional variation [23] one can obtain a set of equations for the optimal spin-orbitals which form the Slater determinant with the lowest energy, i.e. the ground-state within the HF approximation. After an unitary transformation, the canonical HF equations (Equation 2.7) are obtained. The sums include all occupied spin-orbitals.

$$\epsilon_i \chi_i(\mathbf{x}_1) = \left[\hat{h}(\mathbf{x}_1) + \sum_{j \neq i} \hat{\mathcal{J}}_j(\mathbf{x}_1) - \hat{\mathcal{K}}_j(\mathbf{x}_1) \right] \chi_i(\mathbf{x}_1) \quad (2.7)$$

$$\hat{h}(\mathbf{x}_1) = -\frac{1}{2} \nabla_{\mathbf{r}_1}^2 - \sum_A \frac{Z_A}{|\mathbf{r}_1 - \mathbf{R}_A|} \quad (2.8)$$

$$\hat{\mathcal{J}}_j(\mathbf{x}_1) \chi_i(\mathbf{x}_1) = \left[\int d\mathbf{x}_2 \frac{\chi_j^*(\mathbf{x}_2) \chi_j(\mathbf{x}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \right] \chi_i(\mathbf{x}_1) \quad (2.9)$$

$$\hat{\mathcal{K}}_j(\mathbf{x}_1) \chi_i(\mathbf{x}_1) = \left[\int d\mathbf{x}_2 \frac{\chi_j^*(\mathbf{x}_2) \chi_i(\mathbf{x}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \right] \chi_j(\mathbf{x}_1) \quad (2.10)$$

$\hat{h}(\mathbf{x}_1)$ is the single-particle Hamiltonian, which contains the kinetic energy of an electron in spin-orbital χ_i and its electrostatic interaction with the nuclei. The Coulomb operator $\hat{\mathcal{J}}_i(\mathbf{x}_1)$ describes a mean-field approximation to the instantaneous Coulomb interaction with all other electrons. The exchange operator $\hat{\mathcal{K}}_i(\mathbf{x}_1)$ does not have a classical interpretation and arises from the anti-symmetry requirement. In contrast to the Coulomb operator, the exchange operator is a non-local operator which depends on the value of χ_i at all points \mathbf{x}_1 . It

is important to note that the HF equations are non-linear because the Coulomb and exchange operators depend on the solution of the equations. They must therefore be solved iteratively.

Since $\hat{\mathcal{J}}_i(\mathbf{x}_1)\chi_i(\mathbf{x}_1) - \hat{\mathcal{K}}_i(\mathbf{x}_1)\chi_i(\mathbf{x}_1) = 0$, the sum in Equation 2.7 can be extended to include i as well.

$$\left[\hat{h}(\mathbf{x}_1) + \sum_j \hat{\mathcal{J}}_j(\mathbf{x}_1) - \hat{\mathcal{K}}_j(\mathbf{x}_1) \right] \chi_i(\mathbf{x}_1) = \epsilon_i \chi_i(\mathbf{x}_1) \quad (2.11)$$

$$\left[\hat{h}(\mathbf{x}_1) + \hat{v}^{HF}(\mathbf{x}_1) \right] \chi_i(\mathbf{x}_1) = \epsilon_i \chi_i(\mathbf{x}_1) \quad (2.12)$$

$$\hat{f}(\mathbf{x}_1)\chi_i(\mathbf{x}_1) = \epsilon_i \chi_i(\mathbf{x}_1) \quad (2.13)$$

$\hat{f}(\mathbf{x}_1)$ is the Fock operator which contains the single-particle Hamiltonian $\hat{h}(\mathbf{x}_1)$ and an effective one-electron potential operator $\hat{v}^{HF}(\mathbf{x}_1)$, the combined Coulomb and Exchange operators.

The energy \mathcal{E}_{HF} of the ground-state wave-function within the HF approximation is then given as the expectation value of the electronic Hamiltonian as shown in Equation 2.14.

$$\mathcal{E}_{HF} = \sum_i \langle \chi_i | \hat{h} | \chi_i \rangle + \frac{1}{2} \sum_{i,j} \langle \chi_i | \hat{\mathcal{J}}_j | \chi_i \rangle + \langle \chi_i | \hat{\mathcal{K}}_j | \chi_i \rangle \quad (2.14)$$

Since the HF theory does not contain any electron correlation besides the Pauli-Exclusion, it is commonly considered as “uncorrelated” theory. The difference to the true ground-state energy \mathcal{E}_0 is therefore defined as the “correlation energy” $\mathcal{E}_{corr} = \mathcal{E}_0 - \mathcal{E}_{HF}$. [23]

Once the spin-orbitals forming the ground-state are known, the Fock operator \hat{f} becomes a well-defined Hermitian operator. Equation 2.13 then yields an infinite set of spin-orbitals as eigenfunctions of the Fock operator. The N spin-orbitals used to construct the HF ground-state $|\Phi_0\rangle$ are the eigenstates with the lowest eigenvalues. By choosing a different set of spin-orbitals, one can construct additional N -electron determinants. The notation $|\Phi_i^a\rangle$ denotes a single excitation where the electron occupying χ_i in the ground-state has been excited to the state χ_a . In the same way, double excitations $|\Phi_{ij}^{ab}\rangle$, triple excitations $|\Phi_{ijk}^{abc}\rangle$ and higher order excitations are defined with respect to the HF ground-state. The well-known Brillouin’s theorem (Equation 2.15) follows from the orthogonality of the Fock operator eigenstates χ_i . It states that no singly excited determinant interacts directly with the HF ground-state wave-function. [23]

$$\langle \Phi_0 | \hat{\mathcal{H}} | \Phi_i^a \rangle = \langle \chi_i | \hat{f} | \chi_a \rangle = 0 \quad (2.15)$$

The HF equations can be simplified for spin-unpolarized (“closed-shell”) systems the by imposing the additional restriction that the spin-up and down orbitals share the same spatial orbitals. [23] The

resulting restricted HF theory is shown in Equation 2.16 where the sum now includes all occupied spatial orbitals $\psi_i(\mathbf{r}_1)$.

$$\epsilon_i \psi_i(\mathbf{r}_1) = \left[\hat{h}(\mathbf{r}_1) + \sum_j 2\hat{\mathcal{J}}_j(\mathbf{r}_1) - \hat{\mathcal{K}}_j(\mathbf{r}_1) \right] \psi_i(\mathbf{r}_1) \quad (2.16)$$

$$\hat{h}(\mathbf{r}_1) = -\frac{1}{2} \nabla_{\mathbf{r}_1}^2 - \sum_I \frac{Z_I}{|\mathbf{r}_1 - \mathbf{R}_I|} \quad (2.17)$$

$$\hat{\mathcal{J}}_j(\mathbf{r}_1) \psi_i(\mathbf{r}_1) = \left[\int d\mathbf{r}_2 \frac{\psi_j^*(\mathbf{r}_2) \psi_j(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \right] \psi_i(\mathbf{r}_1) \quad (2.18)$$

$$\hat{\mathcal{K}}_j(\mathbf{r}_1) \psi_i(\mathbf{r}_1) = \left[\int d\mathbf{r}_2 \frac{\psi_j^*(\mathbf{r}_2) \psi_i(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \right] \psi_j(\mathbf{r}_1) \quad (2.19)$$

It should be noted that the Coulomb operator has a factor of two in this formulation, while the exchange term does not. This is due to the fact that the exchange term is non-zero only for electron pairs with the same spins.

To solve the (restricted) HF equations numerically, it is necessary to expand each of the spatial orbitals $\psi_i(\mathbf{r}_1)$ in a finite set of basis functions $\varphi_t(\mathbf{r}_1)$. Throughout my thesis I will assume that the basis functions are real-valued, but the expansion coefficients $c_{i,t}$ can in general be complex.

$$\psi_i(\mathbf{r}_1) = \sum_t c_{i,t} \varphi_t(\mathbf{r}_1) \quad (2.20)$$

Inserting this expansion into Equation 2.16, we obtain:

$$\epsilon_i \sum_t c_{i,t} \varphi_t(\mathbf{r}_1) = \left[\hat{h}(\mathbf{r}_1) + \sum_j 2\hat{\mathcal{J}}_j(\mathbf{r}_1) - \hat{\mathcal{K}}_j(\mathbf{r}_1) \right] \sum_t c_{i,t} \varphi_t(\mathbf{r}_1) \quad (2.21)$$

After multiplying $\varphi_s(\mathbf{r}_1)$ to the left and integrating over \mathbf{r}_1 , the problem turns into a matrix equation.

$$\begin{aligned} \epsilon_i \sum_t c_{i,t} \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) &= \sum_t c_{i,t} \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \\ &\quad \left[\hat{h}(\mathbf{r}_1) + \sum_j 2\hat{\mathcal{J}}_j(\mathbf{r}_1) - \hat{\mathcal{K}}_j(\mathbf{r}_1) \right] \varphi_t(\mathbf{r}_1) \end{aligned} \quad (2.22)$$

The *Roothaan equations* shown in Equation 2.23 are obtained after introducing the overlap matrix $S_{s,t}$ and the Fock matrix $F_{s,t}$.

$$\sum_t c_{i,t} F_{s,t} = \epsilon_i \sum_t c_{i,t} S_{s,t} \quad (2.23)$$

$$F_{s,t} = \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \left[\hat{h}(\mathbf{r}_1) + \sum_j 2\hat{\mathcal{J}}_j(\mathbf{r}_1) - \hat{\mathcal{K}}_j(\mathbf{r}_1) \right] \varphi_t(\mathbf{r}_1) \quad (2.24)$$

$$S_{s,t} = \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \quad (2.25)$$

Upon inserting the expressions for the exchange and Coulomb operators (Equation 2.18 and 2.19) into the definition of the Fock matrix

(Equation 2.24) the central computational challenge of the HF theory becomes evident.

$$\begin{aligned}
F_{s,t} &= \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \hat{h}(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \\
&\quad + 2 \sum_j \int d\mathbf{r}_1 \int d\mathbf{r}_2 \frac{\varphi_s(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \psi_j^*(\mathbf{r}_2) \psi_j(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \\
&\quad - \sum_j \int d\mathbf{r}_1 \int d\mathbf{r}_2 \frac{\varphi_s(\mathbf{r}_1) \psi_j(\mathbf{r}_1) \psi_j^*(\mathbf{r}_2) \varphi_t(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \quad (2.26)
\end{aligned}$$

$$= \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \hat{h}(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) + 2 \sum_j (st|jj) - \sum_j (sj|jt) \quad (2.27)$$

To obtain the Fock matrix elements, it is necessary to evaluate the Electron Repulsion Integrals (ERIs) $(st|jj)$ which are six-dimensional spatial integrals. Since both the number of occupied orbitals and basis functions scales linear with system size, their number increases as $\mathcal{O}(N^4)$ with system size.

The ERIs in Equation 2.26 depend on both basis functions and spatial orbitals. Since the spatial orbitals will change in each step of an iterative solution, it is preferable to reformulate the ERIs only in terms of the fixed basis functions by inserting the basis expansion from Equation 2.20 into Equation 2.26.

$$\begin{aligned}
F_{s,t} &= \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \hat{h}(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \\
&\quad + 2 \sum_{u,v} \sum_j c_{j,u}^* c_{j,v} (st|uv) - \sum_{u,v} \sum_j c_{j,u}^* c_{j,v} (sv|ut) \quad (2.28)
\end{aligned}$$

$$= \int d\mathbf{r}_1 \varphi_s(\mathbf{r}_1) \hat{h}(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) + \sum_{u,v} D_{u,v} [2(st|uv) - (sv|ut)] \quad (2.29)$$

$$D_{u,v} = \sum_j c_{j,u}^* c_{j,v} \quad (2.30)$$

Using this formulation, only the density matrix $D_{u,v}$ has to be updated during the iteration to generate the Fock matrix for the next cycle. Starting from an initial guess for the density matrix, the HF equations can then be evaluated iteratively until the density matrix no longer changes. This procedure is known as Self-Consistent Field (SCF) approach. The convergence of the SCF and its acceleration are a huge field of study, but will not be discussed here. The focus of my project are methods for an efficient evaluation of the ERIs $(st|uv)$ which do not rely on the availability of analytic integral expressions.

2.4 CANONICAL MP2

In the previous section the HF theory was considered as an approximate solution for the eigenfunction with the lowest eigenvalue of the exact Hamiltonian. Alternatively, one can also view HF as the exact solution to the approximate Hamiltonian $\hat{\mathcal{F}} = \sum_i \hat{f}(i)$ which is a more convenient starting point for corrections based on perturbation theory. The eigenstates of $\hat{\mathcal{F}}$ are the HF ground-state $|\Phi_0\rangle$ and all excited Slater determinants $|\Phi_i^a\rangle$, $|\Phi_{ij}^{ab}\rangle$, $|\Phi_{ijk}^{abc}\rangle$ etc. The corresponding

eigenvalues are just the sum of the orbital eigenvalues given by Equation 2.13. The full electronic Hamiltonian can then be expressed as the Fock Hamiltonian $\hat{\mathcal{F}}$ plus a correction $\hat{\mathcal{H}}_1$.

$$\hat{\mathcal{H}} = \hat{\mathcal{F}} + \hat{\mathcal{H}}_1 \quad (2.31)$$

$$\hat{\mathcal{H}}_1 = \hat{\mathcal{H}} - \hat{\mathcal{F}} = \sum_{i<j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_i \hat{v}^{HF}(i) \quad (2.32)$$

If the correction $\hat{\mathcal{H}}_1$ is small, standard perturbation theory can be applied to this problem and leads to the expressions in equations 2.33 to 2.35. $\mathcal{E}_0^{(0)}$ represents the unperturbed ground-state energy of the non-interacting Fock system. $\mathcal{E}_0^{(1)}$ and $\mathcal{E}_0^{(2)}$ are the first and second order perturbation corrections towards the exact ground-state energy \mathcal{E}_0 of the full interacting Hamiltonian. [23]

$$\mathcal{E}_0^{(0)} = \langle \Phi_0 | \hat{\mathcal{F}} | \Phi_0 \rangle \quad (2.33)$$

$$\mathcal{E}_0^{(1)} = \langle \Phi_0 | \hat{\mathcal{H}}_1 | \Phi_0 \rangle = \langle \Phi_0 | \hat{\mathcal{H}} - \hat{\mathcal{F}} | \Phi_0 \rangle \quad (2.34)$$

$$\mathcal{E}_0^{(2)} = \sum_{n \neq 0} \frac{|\langle \Phi_0 | \hat{\mathcal{H}}_1 | \Phi_n \rangle|^2}{\mathcal{E}_0^{(0)} - \mathcal{E}_n^{(0)}} \quad (2.35)$$

From Equation 2.33 and 2.34 it is immediately obvious that these two contributions together yield the ground-state total energy at the HF level. The first contribution to the correlation energy \mathcal{E}_{corr} is thus the second order correction $\mathcal{E}_0^{(2)}$. The sum in Equation 2.35 includes all eigenstates of the Fock Hamiltonian except the ground-state $|\Phi_0\rangle$. However, from Brillouin's theorem (Equation 2.15) it follows that single excitations do not contribute to this sum. Furthermore, any term involving a triple or higher order excitation must vanish because $\hat{\mathcal{H}}_1$ only contains one- and two-electron operators. Thus only the double excitations remain and the second order energy correction becomes Equation 2.36.

$$\mathcal{E}_0^{(2)} = \sum_{i<j} \sum_{a<b} \frac{|\langle \Phi_0 | \hat{\mathcal{H}}_1 | \Phi_{ij}^{ab} \rangle|^2}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (2.36)$$

$$= 2 \sum_{i<j} \sum_{a<b} \frac{|\langle \chi_i \chi_a | \chi_j \chi_b \rangle|^2 - (\langle \chi_i \chi_a | \chi_j \chi_b \rangle)^* \langle \chi_i \chi_b | \chi_j \chi_a \rangle}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (2.37)$$

$$= \frac{1}{2} \sum_{i,j} \sum_{a,b} \frac{|\langle \chi_i \chi_a | \chi_j \chi_b \rangle|^2 - (\langle \chi_i \chi_a | \chi_j \chi_b \rangle)^* \langle \chi_i \chi_b | \chi_j \chi_a \rangle}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (2.38)$$

Equation 2.38 is the second order energy correction to the HF ground-state known as 2nd Order Møller-Plesset Perturbation Theory (MP2) and commonly applied as a post-SCF correction to a converged HF calculation. Higher order corrections can be derived in a similar fashion. For the closed-shell case, Equation 2.38 can be simplified to Equation 2.39 by integrating out the spins analytically.

$$\mathcal{E}_0^{(2)} = \sum_{i,j} \sum_{a,b} \frac{2 |\langle \psi_i \psi_a | \psi_j \psi_b \rangle|^2 - (\langle \psi_i \psi_a | \psi_j \psi_b \rangle)^* \langle \psi_i \psi_b | \psi_j \psi_a \rangle}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (2.39)$$

In a canonical MP2 implementation the computational effort scales as $\mathcal{O}(N^5)$. The dominating step is the transformation of the ERIs from the basis set representation $(\varphi_s\varphi_u|\varphi_t\varphi_v)$ to the molecular one $(\psi_i\psi_a|\psi_j\psi_b)$ used in Equation 2.39. As it can be seen from the pseudo-code steps below, this transformation has a $\mathcal{O}(N^5)$ scaling with respect to the system size.

$$\text{step 1: } \sum_s (\varphi_s\varphi_u|\varphi_t\varphi_v) c_{i,s} = (\psi_i\varphi_u|\varphi_t\varphi_v) \quad \sim N_{\text{bas}}^4 N_{\text{occ}} \quad (2.40)$$

$$\text{step 2: } \sum_t (\psi_i\varphi_u|\varphi_t\varphi_v) c_{j,t} = (\psi_i\varphi_u|\psi_j\varphi_v) \quad \sim N_{\text{bas}}^3 N_{\text{occ}}^2 \quad (2.41)$$

$$\text{step 3: } \sum_u (\psi_i\varphi_u|\psi_j\varphi_v) c_{a,u} = (\psi_i\psi_a|\psi_j\varphi_v) \quad \sim N_{\text{bas}}^2 N_{\text{occ}}^2 N_{\text{virt}} \quad (2.42)$$

$$\text{step 4: } \sum_v (\psi_i\psi_a|\psi_j\varphi_v) c_{b,v} = (\psi_i\psi_a|\psi_j\psi_b) \quad \sim N_{\text{bas}} N_{\text{occ}}^2 N_{\text{virt}}^2 \quad (2.43)$$

The MP2 is also demanding in terms of memory, since the number of molecular ERIs scales as $\mathcal{O}(N^4)$. In principle these elements can be evaluated one occupied orbital pair at a time in the canonical algorithm, but the computational overhead increases dramatically because one has to frequently recompute the ERIs in the basis set representation. A commonly used strategy to address this issue is the Resolution of Identity (RI), which will be discussed in detail in section 2.10.

Another challenge all methods involving the virtual states have in common is the slow convergence with respect to the basis set size. Smooth one-particle basis functions are not well suited to represent the electronic cusp, thus one usually needs very large basis sets to reach convergence. [31, 32] One way to address this slow convergence is the introduction of explicitly correlated basis functions known as F12 methods, but this approach increases the computational complexity of the MP2 to $\mathcal{O}(N^6)$. [31] A popular alternative is the extrapolation towards the Complete Basis Set (CBS) limit with correlation-consistent basis sets. [32, 33]

2.5 VARIANTS OF THE MP2

A variety of techniques has been proposed to improve the accuracy or reduce the computational effort of the MP2. Most of the lower scaling correlation methods are motivated by the near-sightedness of the electronic correlation. [34–37] This feature can however not be exploited in the representation of canonical HF orbitals which have a delocalized nature. Local MP2 methods employ orbital transformations to recast the occupied orbitals into localized ones. After choosing an appropriate localized representation for the virtual orbital space, a system-size independent subset of it can be assigned to each occupied orbital pair. The previously system-size dependent sums over virtual orbitals then turn into locally restricted sums for each occupied orbital pair and significantly reduce the scaling. Schütz and coworkers used Projected Atomic Orbitals (PAOs) to represent the virtual space. PAOs are obtained by projecting out the occupied molecular orbitals from the space of the atom-centered basis functions. [34] A

local subset of the PAOs is then assigned to each localized occupied orbital pair according to its spatial extend. Another approach are the Pair Natural Orbitals (PNOs) originally used for CI by Ahlrichs and others. [38] PNOs are linear combinations of virtual orbitals and are generated individually for each pair of occupied orbitals. Closely connected to the PNOs are the Orbital Specific Virtuals (OSVs), which assign a subset of functions representing the virtual space to each individual occupied orbital instead of making assignments to pairs. [35] The local virtual space of a given occupied orbital pair is then just the union of their virtual orbitals. OSVs yield less compact virtual subspaces, but the effort to generate these virtual subspaces is considerably reduced. Furthermore, these approaches can be combined, e.g. Schmitz and coworkers combined OSVs and PNOs [36] while the Neese group combined PNOs and PAOs in their DLPNO schemes. [37]

These local restrictions of the virtual spaces can cause discontinuities of the potential energy surface if the virtual orbitals included in the local subspaces change radically as response to a small change in the geometry. [36, 39]

An alternative approach without partitioning of the virtual orbital space is the Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP2) originally introduced by Häser. [40] In this method the denominator with the molecular eigenvalues is transformed into the integral of an exponential and then approximated with a quadrature scheme.

$$\frac{1}{x} = \int_0^{\infty} e^{-xt} dt \quad \forall x > 0$$

$$\Rightarrow \frac{1}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} = - \int_0^{\infty} e^{-(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)q} dq \quad (2.44)$$

$$\approx - \sum_q w_q e^{-(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)t_q} \quad (2.45)$$

Inserting Equation 2.45 into the MP2 correlation energy for closed-shell systems (Equation 2.39) then yields:

$$\mathcal{E}_0^{(2)} \approx - \sum_q w_q \sum_{i,j} \sum_{a,b} e^{-(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)t_q} \left[2 (\psi_i \psi_a | \psi_j \psi_b)^* (\psi_i \psi_a | \psi_j \psi_b) - (\psi_i \psi_a | \psi_j \psi_b)^* (\psi_i \psi_b | \psi_j \psi_a) \right] \quad (2.46)$$

After inserting the basis set representation of the spatial molecular orbitals ψ_i (Equation 2.20), one obtains Equation 2.47.

$$\mathcal{E}_0^{(2)} = -2 \sum_q w_q \sum_{i,j} \sum_{a,b} \sum_{\substack{s,t,u,v \\ s',t',u',v'}} c_{i,s'} c_{j,t'} c_{a,u'}^* c_{b,v'}^* c_{i,s}^* c_{j,t}^* c_{a,u} c_{b,v}$$

$$\times e^{-(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)t_q} (s'u' | t'v') (su | tv)$$

$$+ \sum_q w_q \sum_{i,j} \sum_{a,b} \sum_{\substack{s,t,u,v \\ s',t',u',v'}} c_{i,s'} c_{j,t'} c_{a,u'}^* c_{b,v'}^* c_{i,s}^* c_{j,t}^* c_{a,v} c_{b,u}$$

$$\times e^{-(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)t_q} (s'u' | t'v') (sv | tu) \quad (2.47)$$

This formulation has the distinct advantage that the different tensors are no longer coupled via the denominator and one can thus now interchange the evaluation order of the sums. The expansion coefficients of the molecular orbitals and the exponentials with their eigenvalues can be reformulated as pseudo-density matrices, which reveals another advantage of the LT-MP2.

$$\mathcal{E}_0^{(2)} = - \sum_q w_q \sum_{s,t,u,v} \sum_{s',t',u',v'} (s'u'|t'v') X_{s',s}^q X_{t',t}^q Y_{u',u}^q Y_{v',v}^q \left[2(su|tv) - (sv|tu) \right] \quad (2.48)$$

$$X_{s',s}^q = \sum_i c_{i,s'} c_{i,s}^* e^{\epsilon_i t q} \quad (2.49)$$

$$Y_{u',u}^q = \sum_a c_{a,u'} c_{a,u}^* e^{-\epsilon_a t q} \quad (2.50)$$

$X_{s',s}^q$ denotes the pseudo-density matrix for the occupied orbitals and $Y_{t',t}^q$ is its counterpart for the virtual orbitals. These pseudo-density matrices feature a similar spatial decay behavior as the normal density matrix in HF (Equation 2.30) and thus allow the usage of integral screening techniques to dramatically reduce the number of significant terms in the summations.

In contrast to the local MP2 methods, the LT-MP2 does not explicitly enforce any fictitious locality on the orbitals. The accuracy is solely controlled by the quadrature scheme and the integral screening. By improving these, the LT-MP2 can be smoothly converged to the canonical MP2 result if no further approximation is made.

Ochsenfeld and his coworkers invested a lot of effort to improve the LT-MP2 technique by deriving superior screening criteria and other enhancement techniques like Cholesky-decomposed densities. [41, 42] The LT-MP2 will be discussed in more detail in chapter 4.

Most of these locality exploiting methods are implemented only with very limited parallelization (e.g. by using thread-parallel linear algebra libraries) and can only use a handful of cores at most. A notable exception to this statement is the recently published Divide-Expand-Consolidate 2nd Order Møller-Plesset Perturbation Theory (DEC-MP2) which is a massively-parallelized version of the fragmentation MP2, i.e. the global problem is divided into a multitude of small independent sub-problems. [43–45] This method scales linear with the number of fragments, but the computational scaling of the MP2 calculations within each fragment remains unchanged. Using most of the available resources on the “Titan” super computer (14952 out of 18688 nodes with 16 cores each and GPU support) they managed to compute a supra molecular wire system with 24400 basis functions within 19 hours. A massive-parallel canonical MP2 for molecular systems including analytical gradients also has been published recently. [46]

Another approach to reduce the computational complexity of the MP2 is the Spin-Opposite-Scaled 2nd Order Møller-Plesset Perturbation Theory (SOS-MP2). [47] This method is based on Grimme’s Spin-Component-Scaled 2nd Order Møller-Plesset Perturbation The-

ory (SCS-MP2) [8] which partitions the MP2 energy into same-spin and opposite-spin contributions and rescales them with empirical factors to improve the overall accuracy of the MP2. SOS-MP2 is based on the observation that the same-spin component with the computationally more complicated exchange term is rescaled with a factor $\frac{1}{3}$ in the SCS-MP2. Thus SOS-MP2 completely discards the same-spin term and only evaluates the opposite-spin term with an adjusted scaling factor. Avoiding the evaluation of the exchange term results in a MP2 algorithm that can be implemented with a fourth-order scaling.

Regardless of the chosen MP2 flavor, the evaluation of the ERIs always remains one of the biggest computational challenges in the algorithm. Since this is not a MP2 exclusive problem, but affects any other method using such ERIs as well, the established strategies to reduce the effort needed for these integrals will be discussed separately in section 2.10.

2.6 DENSITY FUNCTIONAL THEORY

An alternative approach to describe quantum mechanical systems on a first-principles level is Density-Functional Theory (DFT). While the central entity used to describe the system in HF and MP2 has been the wave-function, DFT uses the electronic density instead. Hohenberg and Kohn demonstrated in their important theorems [48], that there exists a bijective mapping from the external potential $v_{ext}(\mathbf{r}_1)$ (i.e. the potential formed by the nuclei in Equation 2.5) to the ground-state wave-function in non-degenerate systems and also from the ground-state wave-function to the electronic density of the ground-state. Therefore the external potential $v_{ext}(\mathbf{r}_1)$, the ground-state density $n_0(\mathbf{r}_1)$ and the ground-state wave-function Ψ_0 uniquely define each other. As a consequence, any ground-state observable can be expressed as a density functional, in particular the total energy of the ground-state. It follows that the ground-state density for a given $v_{ext}(\mathbf{r}_1)$ must minimize the total energy functional $E[n]$.

A detailed discussion of the Hohenberg-Kohn theorem including proofs, a generalization to systems with degenerate ground-states and additional proofs of existence needed for a detailed derivation of DFT can be found in Ref. [49].

DFT exploits this important theorem and the fact that most properties of interest can be expressed in terms of the electronic density to describe the ground-state of a system without having to explicitly make use of the ground-state wave-function. The practical use of DFT is based on the Kohn-Sham (KS) scheme, which maps the complicated interacting system with the Hamiltonian from Equation 2.5 to a *non-interacting* reference system with the Hamiltonian shown in Equation 2.51.

$$\hat{\mathcal{H}}_{KS} = - \sum_i \frac{1}{2} \nabla_i^2 + v_s(\mathbf{r}_1) \quad (2.51)$$

The Hohenberg-Kohn theorem also holds true for non-interacting systems, thus the multiplicative potential $v_s(\mathbf{r}_1)$ can be chosen such that

it gets infinitesimal close to the ground-state density of the interacting system. To obtain the ground-state density of the interacting system one thus “only” needs to determine the potential $v_s(\mathbf{r}_1)$ and calculate the density of the non-interacting reference system.

The exact ground-state wave-function for a non-interacting system is a Slater determinant built from the N eigenstates with the lowest eigenvalues of the single-particle Schrödinger equation in Equation 2.52. The ground-state density of the non-interacting system can be immediately calculated from the Slater determinant.

$$\left[-\frac{1}{2}\nabla^2 + v_s(\mathbf{r}_1) \right] \chi_i(\mathbf{x}_1) = \epsilon_i \chi_i(\mathbf{x}_1) \quad (2.52)$$

$$n_0(\mathbf{r}_1) = \sum_i |\chi_i(\mathbf{x}_1)|^2 \quad (2.53)$$

It should be noted that the single-particle orbitals $\chi_i(\mathbf{x}_1)$ are uniquely determined functionals of the external potential v_s too.

The total electronic energy of the ground-state in the interacting system can be reformulated as functional of the density and thus be evaluated with the density from Equation 2.53.

$$E[n] = T_s[n] + E_H[n] + E_{ext}[n] + E_{xc}[n] \quad (2.54)$$

$$T_s[n] = \sum_i \int d\mathbf{x}_1 \chi_i^*(\mathbf{x}_1) \nabla^2 \chi_i(\mathbf{x}_1) \quad (2.55)$$

$$E_H[n] = \frac{1}{2} \int d\mathbf{r}_1 d\mathbf{r}_2 \frac{n(\mathbf{r}_1)n(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \quad (2.56)$$

$$E_{ext}[n] = \int d\mathbf{r}_1 n(\mathbf{r}_1) v_{ext}(\mathbf{r}_1) \quad (2.57)$$

Here, $T_s[n]$ denotes the kinetic energy of the non-interacting particles and is an implicit density functional because it depends on the orbitals forming the density, not the density itself. It can however be proven that the orbitals χ_i are also uniquely determined by the external potential. [49] The next contribution is the Hartree-term $E_H[n]$ which is the classical Coulomb interaction. In contrast to the Coulomb operator in HF theory (Equation 2.9), this term includes the interaction of each orbital with itself. $E_{ext}[n]$ denotes the energy contribution from the external potential of the nuclei. Finally, $E_{xc}[n]$ denotes the exchange correlation functional which contains all parts of total energy not covered by the other terms. There exists no expression of $E_{xc}[n]$ which would be of any use in real calculations, thus suitable approximations are necessary for $E_{xc}[n]$. As a consequence, the widely used DFAs will usually not fully cancel the self-interaction in the Hartree-term, giving rise to the well-known “self-interaction error” of DFAs.

An explicit expression for the effective potential $v_s(\mathbf{r}_1)$ in the non-interacting reference system can be obtained by investigating how Equation 2.54 reacts to small variations of the ground-state density. [49]

$$\left[-\frac{1}{2}\nabla^2 + v_{ext}(\mathbf{r}_1) + v_H[n](\mathbf{r}_1) + v_{xc}[n](\mathbf{r}_1) \right] \chi_i(\mathbf{x}_1) = \epsilon_i \chi_i(\mathbf{x}_1) \quad (2.58)$$

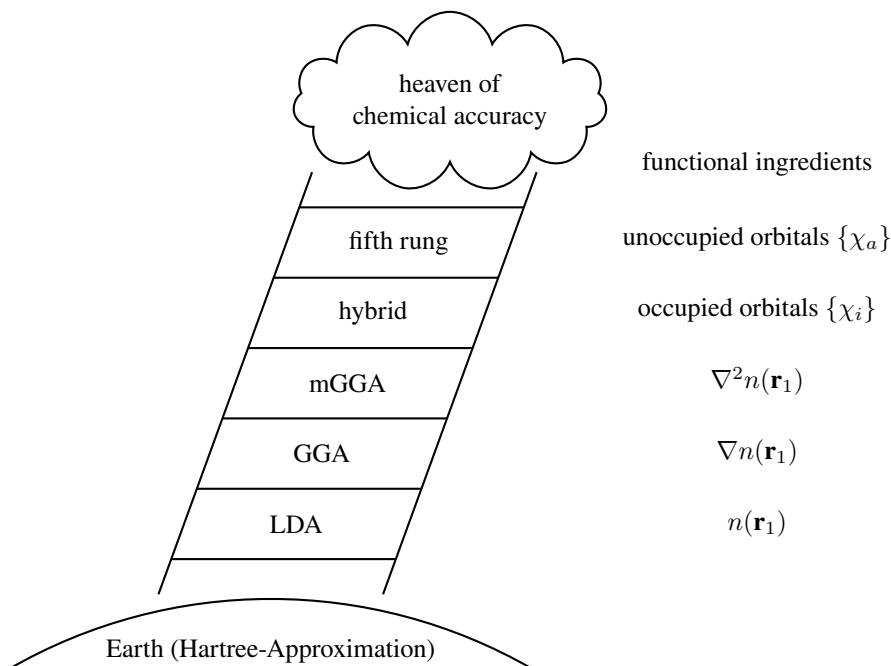


Figure 2.2: “Jacob’s ladder”, a classification of the various DFAs by John Perdew according to their input variables. Each rung also includes the ingredients from all lower rungs. A higher rung on the ladder implies a higher computational complexity, but does not guarantee a higher accuracy.

$$v_H[n](\mathbf{r}_1) = \int d\mathbf{r}_2 \frac{n(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \quad (2.59)$$

$$v_{xc}[n](\mathbf{r}_1) = \frac{\delta E_{xc}[n]}{\delta n} \quad (2.60)$$

In Equation 2.58 $v_{ext}(\mathbf{r}_1)$ is the external potential of the interacting system, $v_H[n](\mathbf{r}_1)$ is the Hartree potential and $v_{xc}[n](\mathbf{r}_1)$ is the exchange correlation potential, given as the functional derivative of the exchange correlation functional.

From an algorithmic point of view Equation 2.58 looks quite similar to the HF equations, because it is also a non-linear eigenvalue problem. Thus, the KS equations have to be solved iteratively just like the HF equations. The KS equation can be brought into a form suitable for numerics by expanding the single-particle orbitals in a finite set of basis functions. This procedure is similar to the derivation of the Roothan equations for HF shown in section 2.3.

As mentioned before, there exists no expression for the exact exchange correlation functional that can be used in real calculations. Over the years, a variety of different approximations has been suggested, which vary greatly in their accuracy and also their computational cost.

The various DFAs are often classified according to Jacob’s ladder (Figure 2.2) proposed by John Perdew. [50] On the lowest rung of this ladder there are the functionals belonging to the LDA which were the first functionals used in DFT. In these functionals, $E_{xc}[n]$ is solely a function of the electronic density itself. LDA functionals like VWN

[51] are parametrized fits of the Homogeneous Electron Gas (HEG) and designed such that the known density limits are reproduced. [52–55] Another well known LDA functional is the SIC-LSD [56] which contains a correction term that aims at minimizing the self-interaction error present in DFAs and its more recent extension FSIC-LSD. [57] The second rung are the GGA functionals where the exchange correlation functional also depends on the gradient of the density to take the local density changes better into account. Well known representatives of this functional class are B-LYP [58, 59] and PBE [60] as well as its modifications like the rPBE. [61] or PBEsol [62]. On the third level one finds the mGGAs which additionally depend on the second derivative of the density or alternatively the orbital kinetic energy density $\tau(\mathbf{r}_1) = \frac{1}{2} \sum_i |\nabla \chi_i(\mathbf{r}_1)|^2$. The TPSS [63] is a well established mGGA and more recently the SCAN functional [64] has been published, which is designed to fulfill a large number of exact constraints.

Up to this level, all exchange correlation functionals are local or semi-local. On the fourth rung the functionals become dependent on the occupied single-particle orbitals and can thus contain HF-like exact exchange terms. Most functionals of this kind are “hybrid functionals” which mix a certain fraction of exact exchange with the exchange part of a LDA or GGA functional, i.e.

$$E_{\text{hybrid}} = E_c^{\text{LDA/GGA}} + (1 - \alpha)E_x^{\text{LDA/GGA}} + \alpha E_x^{\text{HF}} \quad (2.61)$$

The PBE0 functional [65] is the popular extension of the PBE functional to this level and contains 25% of HF-like exact exchange. HSE06 [66] is another popular hybrid functional, which separates the HF-like exact exchange into a short and long-range contribution. Only the short-range part is used in the functional, the long-range contribution to exchange energy is treated completely on the PBE level. For the exchange term only, the Coulomb operator is split as $\frac{1}{r} = \frac{\text{erfc}(\omega r)}{r} + \frac{\text{erf}(\omega r)}{r}$ where the first term is the short-range (SR) part. The exchange-correlation energy in the HSE06 functional is then given by:

$$E_{\text{HSE06}} = E_c^{\text{PBE}} + (1 - \alpha)E_x^{\text{PBE,SR}} + \alpha E_x^{\text{HF,SR}} + E_x^{\text{PBE,LR}} \quad (2.62)$$

In particular for solids this screened Coulomb interaction reduces the effort for evaluating the exact exchange considerably.

On the last rung of Jacob’s ladder one finds those functionals which depend on both the occupied and unoccupied orbitals. The additional dependence on the unoccupied orbitals considerably increases the complexity of these functionals, which are therefore mostly applied as post-corrections to converged calculations of lower-rung functionals. One popular and well-known functional from this class is the RPA which is based on the correspondent high-density limit approximation of the HEG and discussed in more detail in the next section. [67] Other representatives of this functional class are the “double hybrids” that emerged in recent years. [10] Successful examples of such functionals are Grimme’s B2-PLYP [68] and Zhang’s XYG3 [69], which add a part of the MP2-like correlation as post-SCF correction.

“MP2-like” means that the MP2 correlation energy (see section 2.4) is computed with the KS orbitals instead of the HF orbitals. Functionals from this class generally provide more accurate than lower-rung functionals or pure MP2. [10, 12] The weighted MP2 contribution is the computationally most expensive term in their evaluation and an efficient MP2 implementation is therefore crucial for the overall computational efficiency of these functionals.

It should be noted that most functionals have problems with the description of weak dispersion interactions. A variety of correction schemes exist, which can be combined with most functionals, e.g. the Tkatchenko-Scheffler (TS) scheme [70] (and its extension to surfaces vdW^{surf} [71]), Grimme’s D3 correction [72] and the Many-Body-Dispersion (MBD). [73]

2.7 THE RANDOM PHASE APPROXIMATION

A formal expression for the exact exchange-correlation functional can be derived by means of the Adiabatic Connection. An in-depth derivation can be found in e.g. Ref. [49]. Many DFAs are based on computationally feasible approximations to this exact formalism. In the first step, the Hamiltonian of the interacting system, $\hat{\mathcal{H}}$ is partitioned into the Hamiltonian of the KS system $\hat{\mathcal{H}}_{KS}$ and a remainder $\hat{\mathcal{H}}_1$.

$$\hat{\mathcal{H}}_1 = v(\mathbf{r}_1 - \mathbf{r}_2) - \int d\mathbf{r}_1 n(\mathbf{r}_1) [v_H(\mathbf{r}_1) + v_{xc}(\mathbf{r}_1)] \quad (2.63)$$

A coupling constant λ can be introduced into this equation such that the scaled Hamiltonian $\hat{\mathcal{H}}_\lambda$ is equivalent to $\hat{\mathcal{H}}_{KS}$ for $\lambda = 0$ and becomes $\hat{\mathcal{H}}$ for $\lambda = 1$. Under the assumption of a non-changing density, i.e. $n_\lambda(\mathbf{r}_1) = n(\mathbf{r}_1) \forall \lambda \in [0, 1]$, the scaled Hamiltonian $\hat{\mathcal{H}}_\lambda$ becomes

$$\hat{\mathcal{H}}_\lambda = \hat{T} + \int d\mathbf{r}_1 v_\lambda n(\mathbf{r}_1) + \lambda v(\mathbf{r}_1 - \mathbf{r}_2) \quad (2.64)$$

where v_λ is only known explicitly for two special cases, $v_{\lambda=0} = v_s$ and $v_{\lambda=1} = v_{ext}$. After defining $\Phi_0(\lambda)$ as the ground state of the λ -scaled system and $E_0(\lambda)$ as its eigenvalue, one obtains:

$$\hat{\mathcal{H}}_\lambda |\Phi_0(\lambda)\rangle = E_0(\lambda) |\Phi_0(\lambda)\rangle \quad (2.65)$$

$$\Rightarrow E_0(\lambda) = \langle \Phi_0(\lambda) | \hat{\mathcal{H}}_\lambda | \Phi_0(\lambda) \rangle \quad (2.66)$$

At this point coupling constant integration can now be used, i.e. one first takes the derivative with respect to λ and then integrates over λ from 0 to 1. This results in

$$E_0(1) - E_0(0) = \int_0^1 d\lambda \left\langle \Phi_0(\lambda) \left| \int d\mathbf{r}_1 n(\mathbf{r}_1) \frac{dv_{\lambda(\mathbf{r}_1)}}{d\lambda} + v(\mathbf{r}_1 - \mathbf{r}_2) \right| \Phi_0(\lambda) \right\rangle \quad (2.67)$$

$$\begin{aligned} &= \int d\mathbf{r}_1 n(\mathbf{r}_1) [v_{ext}(\mathbf{r}_1) - v_s(\mathbf{r}_1)] \\ &+ \int_0^1 d\lambda \langle \Phi_0(\lambda) | v(\mathbf{r}_1 - \mathbf{r}_2) | \Phi_0(\lambda) \rangle \end{aligned} \quad (2.68)$$

From this formula one can extract an expression for the exchange-correlation functional by inserting the definition of the total electronic energy, Equation 2.54, for $E_0(1)$ and its equivalent for the non-interacting KS system for $E_0(0)$.

$$E_H + E_{xc} = \int_0^1 d\lambda \langle \Phi_0(\lambda) | v(\mathbf{r}_1 - \mathbf{r}_2) | \Phi_0(\lambda) \rangle \quad (2.69)$$

This expression can be rewritten in terms of the time-ordered response function $\chi^\lambda(\mathbf{r}_1 t, \mathbf{r}_2 t')$.

$$\chi^\lambda(\mathbf{r}_1 t, \mathbf{r}_2 t') = -i [\langle \Phi_0(\lambda) | T \hat{n}(\mathbf{r}_1 t) \hat{n}(\mathbf{r}_2 t') | \Phi_0(\lambda) \rangle - n_0(\mathbf{r}_1) n_0(\mathbf{r}_2)] \quad (2.70)$$

Here, T denotes the time-ordering operator and n_0 is the ground-state density which is λ independent by construction. The exchange-correlation energy can then be written as:

$$E_{xc} = \frac{1}{2} \iint d\mathbf{r}_1 d\mathbf{r}_2 \int_0^1 d\lambda v(\mathbf{r}_1 - \mathbf{r}_2) [i\chi^\lambda(\mathbf{r}_1 0, \mathbf{r}_2 0) - n(\mathbf{r}_1)\delta(\mathbf{r}_1 - \mathbf{r}_2)] \quad (2.71)$$

Since $\chi^\lambda(\mathbf{r}_1 t, \mathbf{r}_2 t')$ only depends on the difference $t - t'$, the equation can also be written in terms of its Fourier transformation and symmetry properties can be used to restrict the frequency integral to positive values.

$$E_{xc} = \frac{1}{2} \iint d\mathbf{r}_1 d\mathbf{r}_2 \int_0^1 d\lambda v(\mathbf{r}_1 - \mathbf{r}_2) \times \left[i \int_0^\infty \frac{d\omega}{\pi} \chi^\lambda(\mathbf{r}_1, \mathbf{r}_2, \omega) - n(\mathbf{r}_1)\delta(\mathbf{r}_1 - \mathbf{r}_2) \right] \quad (2.72)$$

Equation 2.72 can be split into an HF-like exact exchange contribution and the remaining correlation energy. [49]

$$E_x = \frac{1}{2} \iint d\mathbf{r}_1 d\mathbf{r}_2 v(\mathbf{r}_1 - \mathbf{r}_2) \left[i \int_0^\infty \frac{d\omega}{2\pi} \chi^0(\mathbf{r}_1, \mathbf{r}_2, \omega) - n(\mathbf{r}_1)\delta(\mathbf{r}_1 - \mathbf{r}_2) \right] \quad (2.73)$$

$$E_c = \frac{i}{2} \iint d\mathbf{r}_1 d\mathbf{r}_2 \int_0^1 d\lambda v(\mathbf{r}_1 - \mathbf{r}_2) \int_0^\infty \frac{d\omega}{\pi} [\chi^\lambda(\mathbf{r}_1, \mathbf{r}_2, \omega) - \chi^0(\mathbf{r}_1, \mathbf{r}_2, \omega)] \quad (2.74)$$

From the analytic structure of the response function it is known that it has no poles in the upper right quadrant of the complex frequency plane. Using contour integration, the above expression can thus be reformulated in terms of imaginary frequencies, which avoids any poles that are arbitrarily close to the real frequency axis. [49]

$$E_x = \frac{1}{2} \iint d\mathbf{r}_1 d\mathbf{r}_2 v(\mathbf{r}_1 - \mathbf{r}_2) \left[\int_0^\infty \frac{d\omega}{2\pi} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) - n(\mathbf{r}_1)\delta(\mathbf{r}_1 - \mathbf{r}_2) \right] \quad (2.75)$$

$$E_c = \frac{1}{2} \iint d\mathbf{r}_1 d\mathbf{r}_2 \int_0^1 d\lambda v(\mathbf{r}_1 - \mathbf{r}_2) \int_0^\infty \frac{d\omega}{\pi} [\chi^\lambda(\mathbf{r}_1, \mathbf{r}_2, i\omega) - \chi^0(\mathbf{r}_1, \mathbf{r}_2, \omega)] \quad (2.76)$$

$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega)$, the response function of the non-interacting KS system is known analytically. [67]

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \sum_{m,n} (f_m - f_n) \frac{\psi_m^*(\mathbf{r}_1)\psi_n(\mathbf{r}_1)\psi_n^*(\mathbf{r}_2)\psi_m(\mathbf{r}_2)}{\epsilon_m - \epsilon_n - i\omega} \quad (2.77)$$

Equation 2.75 is equivalent to HF-like exact exchange and thus can be evaluated explicitly with the same strategies as used for hybrid functionals. The correlation energy (Equation 2.76) on the other hand is of no direct use due to $\chi^\lambda(\mathbf{r}_1, \mathbf{r}_2, i\omega)$. To obtain a useful DFA from this expression, a good approximation must be found for $\chi^\lambda(\mathbf{r}_1, \mathbf{r}_2, i\omega)$. In the case of the RPA, it is approximated by a Dyson-type equation.

$$\begin{aligned} \chi^\lambda(\mathbf{r}_1, \mathbf{r}_2, i\omega) &= \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) \\ &+ \lambda \iint d\mathbf{r}_3 d\mathbf{r}_4 \chi^0(\mathbf{r}_1, \mathbf{r}_3, i\omega) v(\mathbf{r}_3 - \mathbf{r}_4) \chi^\lambda(\mathbf{r}_4, \mathbf{r}_2, i\omega) \end{aligned} \quad (2.78)$$

This expression can be used to obtain a power expansion series for the response function of the λ -scaled system and the λ integration can then be evaluated analytically for each order. Using the notation $\text{Tr}[AB] = \iint d\mathbf{r}_1 d\mathbf{r}_2 A(\mathbf{r}_1, \mathbf{r}_2)B(\mathbf{r}_2, \mathbf{r}_1)$, one finally obtains:

$$E_c = \frac{1}{2\pi} \int_0^\infty d\omega \sum_{n=2}^{\infty} \frac{1}{n} \text{Tr} \left[(\chi^0(i\omega)v)^n \right] \quad (2.79)$$

Equation 2.79 is the correlation energy in the RPA approximation. From the definition of the KS system response function in Equation 2.77 it becomes obvious that the evaluation of the RPA formula involves the calculation of molecular orbital ERIs as the most expensive step.

2.8 PERIODIC BOUNDARY CONDITIONS

To describe extended systems like bulk crystals or surfaces it is necessary to introduce periodic boundary conditions. A periodic system is build from an elementary unit cell, which is periodically repeated by shifting it by the lattice vectors $\mathbf{R}_i = n_1\mathbf{a}_1 + n_2\mathbf{a}_2 + n_3\mathbf{a}_3$ where the n are integers and the \mathbf{a}_n are the lattice vectors of the crystal. The electronic Hamiltonian $\hat{\mathcal{H}}$ (see Equation 2.5) for periodic systems can then be expressed as

$$\begin{aligned} \hat{\mathcal{H}} &= - \sum_i \frac{1}{2} \nabla_i^2 - \sum_i \sum_{\mathbf{R}_n} \sum_A \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A - \mathbf{R}_n|} \\ &+ \sum_{i < j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \frac{1}{2} \sum_{\mathbf{R}_n, \mathbf{R}_m} \sum'_{A,B} \frac{Z_B Z_A}{|\mathbf{R}_A + \mathbf{R}_n - \mathbf{R}_B - \mathbf{R}_m|} \end{aligned} \quad (2.80)$$

where the sums over \mathbf{R}_n and \mathbf{R}_m include all cells in the infinite crystal. The prime in the last sum indicates that terms with $\mathbf{R}_n = \mathbf{R}_m$ and $A = B$ are not included. The sums over A and B include all atoms within the unit cell.

The direct lattice also determines the crystal's reciprocal lattice with lattice vectors \mathbf{b}_1 , \mathbf{b}_2 and \mathbf{b}_3 . These are connected to the direct lattice by the conditions

$$\mathbf{a}_i \cdot \mathbf{b}_j = \delta_{i,j}. \quad (2.81)$$

The space spanned by the vectors $\mathbf{k}_i = k_1 \mathbf{b}_1 + k_2 \mathbf{b}_2 + k_3 \mathbf{b}_3$ with values of the k_i in the range of $[-\pi, \pi]$ is called the first Brillouin zone. Any point outside the first Brillouin zone can be mapped to exactly one point in the first Brillouin zone using translations with reciprocal lattice vectors. It is therefore sufficient to analyze k-dependent objects inside the first Brillouin zone only. [74]

From the periodicity of the system follows the Bloch theorem which states that the eigenstates of any single-particle Hamiltonian with translational symmetry can be written as Bloch functions $\psi_{i,\mathbf{k}_1}(\mathbf{r}_1)$ shown in Equation 2.82. [74]

$$\psi_{i,\mathbf{k}_1}(\mathbf{r}_1) = u(\mathbf{r}_1) e^{i\mathbf{k}_1 \mathbf{r}_1} \quad (2.82)$$

$u(\mathbf{r}_1)$ is a function that has the same periodicity as the lattice of the system, i.e. $u(\mathbf{r}_1) = u(\mathbf{r}_1 + \mathbf{R}_n)$. \mathbf{k}_1 is a continuous index which spans the entire first Brillouin zone. For numerical calculations, \mathbf{k}_1 has to be sampled with a finite set of "k-points".

Since the HF Hamiltonian (Equation 2.13) is a single-particle Hamiltonian, the eigenfunctions of its periodic counterpart are Bloch functions. If atom-centered basis functions like Gaussian-Type Orbitals (GTOs) or Numeric Atom-Centered Orbitals (NAOs) are used, it is therefore advantageous to form Bloch-like functions $\varphi_{s,\mathbf{k}_1}(\mathbf{r}_1)$ from the atomic basis functions. The crystalline HF eigenstates $\psi_{i,\mathbf{k}_1}(\mathbf{r}_1)$ can then be expressed as linear combinations of these Bloch-like functions.

$$\varphi_{s,\mathbf{k}_1}(\mathbf{r}_1) = \frac{1}{\sqrt{N}} \sum_{\mathbf{R}_n} e^{i\mathbf{k}_1 \mathbf{R}_n} \varphi_s(\mathbf{r}_1 - \mathbf{R}_n) \quad (2.83)$$

$$\psi_{i,\mathbf{k}_1}(\mathbf{r}_1) = \sum_s c_{i(\mathbf{k}_1),s} \varphi_{s,\mathbf{k}_1}(\mathbf{r}_1) \quad (2.84)$$

Here N denotes the number of unit cells included in the sum and is in principle infinite, but for numerical applications we need to restrict ourselves to a finite set of lattice cells.

The elements of the Fock and overlap matrices of the crystalline basis functions φ_{s,\mathbf{k}_1} between different k-points are zero due to the translational symmetry. [75] Therefore, the eigenvalue problem in Equation 2.85 can be solved for each k-point individually.

$$\sum_t c_{i(\mathbf{k}_1),t} F_{s,t}(\mathbf{k}_1) = \epsilon_i^{\mathbf{k}_1} \sum_t c_{i(\mathbf{k}_1),t} S_{s,t}(\mathbf{k}_1) \quad (2.85)$$

$$F_{s,t}(\mathbf{k}_1) = \sum_{\mathbf{R}_n} e^{i\mathbf{k}_1 \mathbf{R}_n} F_{s,t}(\mathbf{R}_n) \quad (2.86)$$

$$S_{s,t}(\mathbf{k}_1) = \sum_{\mathbf{R}_n} e^{i\mathbf{k}_1 \mathbf{R}_n} S_{s,t}(\mathbf{R}_n) \quad (2.87)$$

Here $F_{s,t}(\mathbf{R}_n)$ and $S_{s,t}(\mathbf{R}_n)$ denote the Fock and overlap matrix elements between the atom-centered basis functions with φ_s in the zeroth unit cell and φ_t shifted by \mathbf{R}_n . The computational effort for the

HF in periodic systems thus scales linearly with the number of \mathbf{k} -points. However, it should be noted that the different \mathbf{k} -points are still connected via the density matrix which is used to construct the Fock matrix. The periodic density matrix for closed-shell systems can be written as

$$D_{u,v}(\mathbf{k}_1) = 2 \sum_j c_{j(\mathbf{k}_1),u}^* c_{j(\mathbf{k}_1),v}. \quad (2.88)$$

A detailed discussion about the numerical challenges of the HF equations and other mean-field methods in periodic systems can be found in Ref. [75].

The Born von Karman (BvK) boundary conditions are a useful concept for numerical calculations of solids. In the BvK boundary conditions the infinite crystal is approximated by a macro-lattice where the unit cell (often referred to as “supercell”) contains $N_1 \times N_2 \times N_3$ primitive unit cells. All operators and functions are then required to obey periodic boundary conditions at the edges of the macro-lattice. [75] Special relations between the BvK macro-lattice and the primitive unit cell can be exploited if special \mathbf{k} -point sets are used. If a regular Monkhorst-Pack mesh[76] with $N_1 \times N_2 \times N_3$ \mathbf{k} -points is used, the translational symmetry can be exploited to switch between real and reciprocal space depending on computational needs. The CRYSTAL - a computational tool for solid state chemistry and physics (CRYSTAL) code for example evaluates the Fock-matrix in a real-space representation and then Fourier-transforms it into the reciprocal space to exploit the block-structure of the Fock-matrix in the diagonalization step. Afterwards the updated density matrix is computed in reciprocal space and then Fourier-transformed back into the real-space representation to evaluate the Fock-matrix for the next SCF iteration. [77]

As it can be seen from Equation 2.81, the volume of the reciprocal unit cell is inverse proportional to the volume of the real-space unit cell. The reciprocal unit cell corresponding to the supercell is therefore much smaller than the one belonging to the primitive unit cell as illustrated in Figure 2.3. If a regular Γ -centered \mathbf{k} -mesh is used, all \mathbf{k} -points of the primitive unit cell except Γ will be located outside the first Brillouin zone of the supercell. Since the first Brillouin zone contains all the necessary information, these additional \mathbf{k} -points can be obtained by translating the Γ point in the supercell by multiples of the supercell’s reciprocal lattice vectors. Therefore, the supercell calculation with the Γ -point and the primitive unit cell with the corresponding Γ -centered \mathbf{k} -grid provide the same information. However, since the system’s periodicity is not exploited in the supercell approach, it will be computationally less efficient because the scaling with the number of basis functions per unit cell is higher than the one with the number of \mathbf{k} -points. However, the supercell’s real-space representation exhibits a lot more sparsity than the reciprocal space. This offers the possibility to use integral screening and other sparsity exploits to improve the computational performance and also reduce the formal scaling. The implementation for periodic HF and hybrid functionals in the aims code uses such a strategy and is presented in

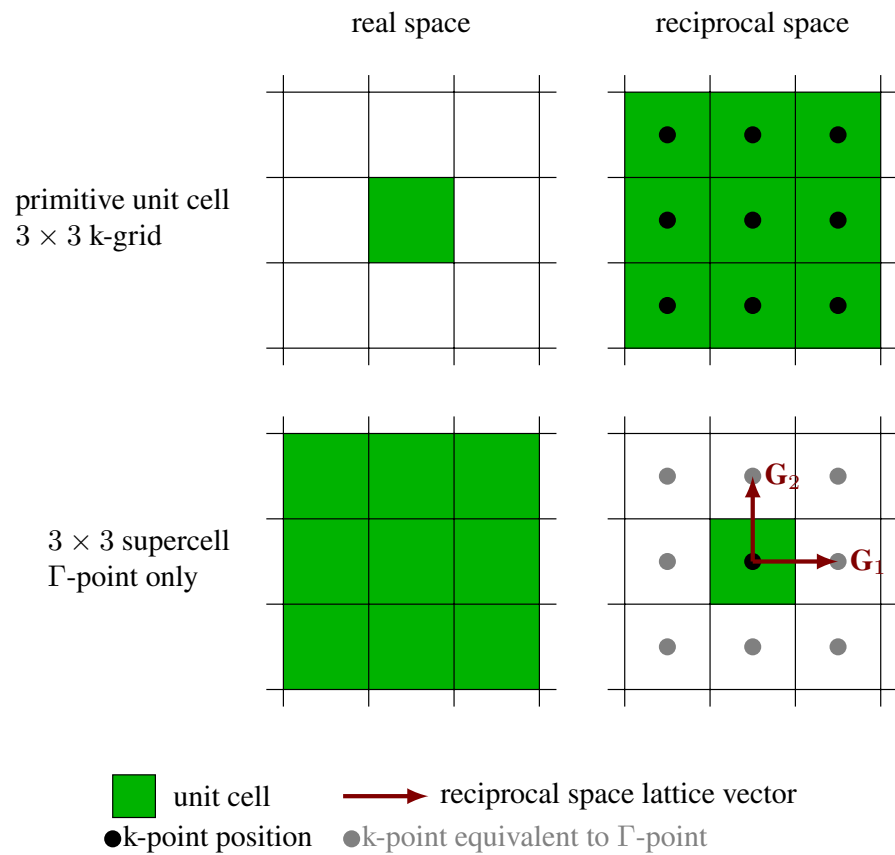


Figure 2.3: Comparison between a primitive unit-cell and the BvK supercell. The larger real-space cell of the BvK maps to a smaller cell in reciprocal space. If a regular Γ -centered k-point grid is used for the primitive unit cell, all k-points except Γ will be located outside the first Brillouin zone of the supercell's reciprocal cell. However, as indicated by the red arrows, they can be folded back onto the supercell Γ -point by translating them with reciprocal lattice vectors and are thus equivalent to it.

Ref. [78]. It should be noted that these techniques involving the BvK boundary conditions only depend on the basic properties of the crystal lattice and are thus not specific to the methods used as examples.

Sun and Bartlett [79] showed that the periodic closed-shell MP2 can be derived from perturbation theory just the same way as in the cluster case.

$$\mathcal{E}_0^{(2)} = \sum_{i,j} \sum_{\mathbf{k}_1, \mathbf{k}_2} \sum_{a,b} \sum_{\mathbf{q}_1, \mathbf{q}_2} \left[\frac{2 \left| \langle \psi_{i, \mathbf{k}_1} \psi_{a, \mathbf{q}_1} | \psi_{j, \mathbf{k}_2} \psi_{b, \mathbf{q}_2} \rangle \right|^2}{\epsilon_i^{\mathbf{k}_1} + \epsilon_j^{\mathbf{k}_2} - \epsilon_a^{\mathbf{q}_1} - \epsilon_b^{\mathbf{q}_2}} - \frac{\Re \left[\langle \psi_{i, \mathbf{k}_1} \psi_{a, \mathbf{q}_1} | \psi_{j, \mathbf{k}_2} \psi_{b, \mathbf{q}_2} \rangle^* \langle \psi_{i, \mathbf{k}_1} \psi_{b, \mathbf{q}_2} | \psi_{j, \mathbf{k}_2} \psi_{a, \mathbf{q}_1} \rangle \right]}{\epsilon_i^{\mathbf{k}_1} + \epsilon_j^{\mathbf{k}_2} - \epsilon_a^{\mathbf{q}_1} - \epsilon_b^{\mathbf{q}_2}} \right] \quad (2.89)$$

Here the sums over i and a include all occupied/virtual Bloch orbitals and the sums over \mathbf{k} -points include all points in the first Brillouin zone. They also showed that Equation 2.89 can be further simplified by exploiting the translational symmetry. If \mathbf{k}_1 , \mathbf{q}_1 and \mathbf{q}_2 are given, there is only one choice for \mathbf{k}_2 in the first Brillouin zone which gives a non-zero contribution. After defining $T(\mathbf{k}_2)$ as an operator which folds \mathbf{k}_2 into the first Brillouin zone by translating it by reciprocal lattice vectors, the MP2 correlation energy of the periodic system can be expressed as

$$\mathcal{E}_0^{(2)} = \sum_{i,j} \sum_{\mathbf{k}_1, \mathbf{q}_1, \mathbf{q}_2} \sum_{a,b} \left[\frac{2 \left| \langle \psi_{i, \mathbf{k}_1} \psi_{a, \mathbf{q}_1} | \psi_{j, T(\mathbf{q}_1 + \mathbf{q}_2 - \mathbf{k}_1)} \psi_{b, \mathbf{q}_2} \rangle \right|^2}{\epsilon_i^{\mathbf{k}_1} + \epsilon_j^{\mathbf{k}_2} - \epsilon_a^{\mathbf{q}_1} - \epsilon_b^{\mathbf{q}_2}} - \frac{\Re \left[\langle \psi_{i, \mathbf{k}_1} \psi_{a, \mathbf{q}_1} | \psi_{j, T(\mathbf{q}_1 + \mathbf{q}_2 - \mathbf{k}_1)} \psi_{b, \mathbf{q}_2} \rangle^* \langle \psi_{i, \mathbf{k}_1} \psi_{b, \mathbf{q}_2} | \psi_{j, T(\mathbf{q}_1 + \mathbf{q}_2 - \mathbf{k}_1)} \psi_{a, \mathbf{q}_1} \rangle \right]}{\epsilon_i^{\mathbf{k}_1} + \epsilon_j^{\mathbf{k}_2} - \epsilon_a^{\mathbf{q}_1} - \epsilon_b^{\mathbf{q}_2}} \right] \quad (2.90)$$

Compared to the canonical MP2 in cluster systems, the scaling now has an additional cubic dependence on the number of \mathbf{k} -points.

As it was the case for molecules, a variety of lower-scaling methods exists for periodic systems as well. Pisani and coworkers extended the local MP2 methods to the periodic case. [80] Scuseria and coworkers adopted the LT-MP2 to the periodic case and demonstrated that this formulation allows one to fold the \mathbf{k} -points into the pseudo-density matrices at negligible cost. [81]

A massive-parallel implementation of the canonical MP2 for periodic systems including forces has been published by Del Ben and coworkers. [82] This implementation has been successfully applied to Molecular Dynamics (MD) simulations of bulk water at the MP2 level. [83]

If the RPA is applied periodic systems, the KS system response function (Equation 2.77) includes the summations over all sampled \mathbf{k} -points. [84]

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \sum_{\substack{m,n \\ \mathbf{k}_1, \mathbf{k}_2}} (f_{m, \mathbf{k}_1} - f_{n, \mathbf{k}_2}) \frac{\psi_{m, \mathbf{k}_1}^*(\mathbf{r}_1) \psi_{n, \mathbf{k}_2}(\mathbf{r}_1) \psi_{n, \mathbf{k}_2}^*(\mathbf{r}_2) \psi_{m, \mathbf{k}_1}(\mathbf{r}_2)}{\epsilon_{m, \mathbf{k}_1} - \epsilon_{n, \mathbf{k}_2} - i\omega} \quad (2.91)$$

and the total energy is evaluated with the same formula as before. In a more recent publication, the Kresse group also presented a periodic RPA implementation that uses the aforementioned supercell approach to obtain a cubic scaling with respect to the number of basis functions in the unit cell and a linear scaling with the number of k-points. [85]

2.9 EMBEDDED CLUSTER MODELS FOR SURFACE CALCULATIONS

An important special case of periodic systems are surfaces, which are needed for modeling e.g. catalytic processes on substrates or Hybrid Inorganic-Organic Systems (HIOS). [21, 86] Such systems can be computed with the various methods described in the previous sections if periodic boundary conditions are applied to the directions parallel to the surface. However, to obtain converged results and model the desired coverage of the surface with the adsorbed molecules it is often necessary to use very large unit cells. Especially for more advanced methods which are often required to obtain an accurate description of the interactions between molecules and the substrates, the computational effort quickly becomes very demanding or even prohibitive.

One possible approach to limit the computational effort of such calculations is the use of hybrid models which combine two or more calculation techniques into one calculation. In these techniques, the system under investigation is divided into two or more parts. The central region contains the site of interest like an active center of a protein or a reaction site which is modeled with a first-principles method. The classical region contains the remaining environment and is usually modeled with an inexpensive force field, which only interacts with the quantum mechanical region via electrostatic interactions. Some methods also allow adding a third layer in between these two, which is then usually computed with a lower-level first-principles method. Popular hybrid methods are the ONIOM embedding model (ONIOM) [87, 88] and Quantum Mechanics / Molecular mechanics embedding techniques (QM/MM). ONIOM on the other hand performs several independent calculations and then extrapolates the results. For a two-layer calculation for example, the classical force-field energy of the entire system and the model system of the quantum mechanical region are calculated as well as the first-principles calculation for the model system.

$$E^{ONIOM}(QM : MM) = E_{MM}^{system} + E_{QM}^{model} - E_{MM}^{model} \quad (2.92)$$

Interactions between the two fragments of the system are included via the not canceling terms in the force-field calculations and the usage of artificial link atoms in the model system when covalent bonds cross the borders between the fragments.

In QM/MM approaches like the Embedded Cluster Model (ECM) the regions are coupled directly and computed in a single calculation. In the ECM, the quantum mechanical region of the system is the adsorbed molecule and an appropriately sized chunk of the substrate




Figure 2.4: Schematic representation of the Embedded Cluster Model by Berger and coworkers. The system is divided into a quantum region where the adsorption takes place and the surrounding parts of the surface which are modeled with cheaper inter-atomic force fields. Pseudo-potentials are used at the boundary region between the quantum and classical region to minimize the errors introduced by the embedding. The outer point charges are chosen such that the correct long-range electrostatic potential of the system is reproduced.

surface around the adsorption site (usually taken as a half sphere), while the remainder of the substrate is modeled with point charges.

In the model used by Berger and coworkers (see Figure 2.4), the classical region of the surface is modeled as a larger half-sphere surrounding the quantum mechanical region and confined by a set of outer point charges which are modeled such that they reproduce the long-range electrostatic potential of the infinite surface. The principal challenge arises at the boundary between the two regions. To minimize the impact of the interface on the electronic structure inside the quantum mechanical region it is necessary to introduce artificial capping potentials which mimic the not explicitly considered electrons as good as possible. The form and parameters of these potentials depend a lot on the system of interest and also the properties one wants to investigate. [89, 90]

The ECM has the distinct advantage that it allows to compute the electronic structure in the quantum mechanical region with high-level methods that are too expensive to be applied to a periodic slab calculation. Recently, the Neese and Reuter groups demonstrated how such embedding techniques can be combined with DLPNO based CCSD(T) calculations to compute the adsorption energies of small molecules on rutile titanium dioxide (TiO_2) surfaces. [22] It also allows for a straight-forward treatment of charged surface defects or adsorbed molecules. [89] On the other hand, the neglect of any explicit periodicity in the model is also a limitation, because coverage-dependent effects cannot be explored. For example, the study of NO adsorption on TiO_2 surfaces by Hättig and coworkers [21] compared calculated infrared spectroscopy data with experimental data and

data from other theoretical investigations. They found several differences in their binding energies and equilibrium structures compared to the other studies of the same system which they attributed to the higher coverages considered in the calculations with periodic boundary conditions. Also it is necessary to compare to calculations with periodic boundary conditions to verify that the artificial boundary region does not alter the physical behavior. And finally, it is difficult to determine the potentially fractional charge contained in the quantum mechanical region when studying systems with delocalized states. [91]

2.10 RESOLUTION OF IDENTITY

The central challenge the methods presented in the previous sections have in common is the evaluation of ERIs. ERIs of basis functions ($su|tv$) are needed for HF and DFT with hybrid functionals, while the canonical MP2 and methods based on it also require the evaluation of molecular orbital ERIs ($ia|jb$). Computing them once and keeping them in memory afterwards is not an option for any but the smallest systems since the number of these integrals scales with $\mathcal{O}(N_{\text{bas}}^4)$ or $\mathcal{O}(N_{\text{occ}}^2 N_{\text{virt}}^2)$. If GTOs or plane waves are used as basis set, the ERIs of the basis functions can be calculated analytically. [92] However, even with these over decades highly-optimized methods the evaluation of ERIs still remains a computationally very demanding task because the ERIs have to be recomputed in each SCF step or for each pair (or batch) of occupied orbitals in a MP2 calculation.

The standard strategy to address this problem is the Resolution of Identity (RI), also known as ‘‘Density Fitting’’. [93, 94] The central idea of the RI approach is to express products of functions in an auxiliary basis set (ABS) and then evaluate the Coulomb interaction only in this auxiliary basis.

$$q_{s,t}(\mathbf{r}_1) = \varphi_s(\mathbf{r}_1)\varphi_t(\mathbf{r}_1) \approx \sum_{\mu}^{ABS} C_{s,t}^{\mu} P_{\mu}(\mathbf{r}_1) = \tilde{q}_{s,t}(\mathbf{r}_1) \quad (2.93)$$

$$(st|uv) \approx \sum_{\mu}^{ABS} \sum_{\lambda}^{ABS} C_{s,t}^{\mu} (\mu|\lambda) C_{u,v}^{\lambda} \quad (2.94)$$

If the number of auxiliary functions P_{μ} is considerably smaller than the square of the basis functions, Equation 2.94 grants a significant reduction of the required memory. The usually employed ABS is a factor 3-10 larger than the original basis set. [95, 96] The largest object in this representation are the RI expansion coefficients $C_{s,t}^{\mu}$ which have a memory requirement that scales as $\mathcal{O}(N_{\text{bas}}^2 N_{\text{aux}})$, i.e. the scaling with respect to system size has been reduced to cubic. The RI coefficients can either be kept in memory as it is usually done in massively parallelized codes (e.g. [82, 96]) or stored on disk and fetched when needed as in many lower-scaling implementations in quantum chemistry codes. (e.g. [36, 42])

The obvious choice for the RI coefficients would be to minimize the differences between the true and the approximated density. This ap-

proach is known as RI-SVS [94], but is rarely used because the RI-V variant is clearly superior to it. RI-V [93, 94] determines the coefficients by minimizing the Coulomb self-interaction of the residual of the fitted density $\delta q_{s,t}$. As Whitten has shown in Ref [93] this optimization criterion yields quadratic errors with respect to the density residual $\delta q_{s,t}$. In RI-SVS on the other hand the errors scale linearly with respect to $\delta q_{s,t}$.

$$\delta q_{s,t} = q_{s,t} - \tilde{q}_{s,t} \quad (2.95)$$

$$(\delta q_{s,t} | \delta q_{s,t}) = (st|st) - 2 \sum_{\mu} (st|\mu) C_{s,t}^{\mu} + \sum_{\mu,\lambda} C_{s,t}^{\mu} (\mu|\lambda) C_{s,t}^{\lambda} \quad (2.96)$$

This entity has to be minimized with respect to the expansion coefficients $C_{s,t}^{\mu}$ to obtain the RI-V coefficients.

$$\frac{\partial (\delta q_{s,t} | \delta q_{s,t})}{\partial C_{s,t}^{\mu}} \stackrel{!}{=} 0 \quad \forall \mu \quad (2.97)$$

$$\Rightarrow (st|\mu) = \sum_{\lambda} C_{s,t}^{\lambda} (\lambda|\mu) \quad \forall \mu \quad (2.98)$$

By introducing the Coulomb matrix of the auxiliary basis functions $V_{\mu,\lambda} = (\mu|\lambda)$, the coefficients in RI-V can be written as Equation 2.99.

$$C_{s,t}^{\mu} = \sum_{\lambda} (st|\lambda) V_{\lambda,\mu}^{-1} \quad (2.99)$$

$$\Rightarrow (st|uv) \approx \sum_{\mu} \sum_{\lambda}^{ABS} (st|\mu) V_{\mu,\lambda}^{-1} (\lambda|uv) \quad (2.100)$$

In practical calculations $V_{\mu,\lambda}^{-1}$ is usually split into two square roots $\sum_{\kappa} V_{\mu,\kappa}^{-0.5} V_{\kappa,\lambda}^{-0.5}$ which are then absorbed into the RI coefficients to eliminate one sum over auxiliary functions. If molecular orbitals are required, the RI coefficients of the atomic orbitals can be contracted with the molecular orbital expansion coefficients to yield a new set of molecular RI coefficients, which reduces the fifth-order scaling of the transformation step.

$$\text{step 1: } \sum_s C_{s,t}^{\mu} c_{i,s} = \hat{C}_{i,t}^{\mu} \quad \sim N_{\text{aux}} N_{\text{bas}}^2 N_{\text{occ}} \quad (2.101)$$

$$\text{step 2: } \sum_t \hat{C}_{i,t}^{\mu} c_{a,t} = \tilde{C}_{i,a}^{\mu} \quad \sim N_{\text{aux}} N_{\text{bas}} N_{\text{occ}} N_{\text{virt}} \quad (2.102)$$

However, the evaluation of the molecular integrals in the MP2 then scales with $\mathcal{O}(N_{\text{occ}}^2 N_{\text{virt}}^2 N_{\text{aux}})$. Despite the overall scaling of the MP2 is thus not reduced by RI-V, the memory requirement and the prefactor can be reduced dramatically. The RI-MP2 is a standard choice to perform MP2 calculations nowadays and available in most electronic structure codes with atom-centered basis functions.

A variety of local RI schemes have been proposed to further reduce the computational effort and also the memory requirements of RI based calculations. Among them is the RI-LVL, which is a part of my thesis project and will be presented in detail in chapter 3 together with other local RI schemes.

Mathematically RI can also be viewed as a tensor factorization, i.e. the rank 4 tensor of ERIs is represented as a product of rank 2 (Coulomb matrix) and rank 3 tensors, the RI coefficients. In recent years the Tensor-Hyper-Contraction (THC) has become more popular, which factorizes the ERI tensor only in terms of rank 2 tensors:

$$(st|uv) \approx \sum_{P,Q} X_s^P X_t^P Z^{P,Q} X_u^Q X_v^Q \quad (2.103)$$

where the indices P, Q denote auxiliary indices, which can e.g. be a grid. As in the RI, the factorization is not uniquely defined. One choice is the least squares THC which imposes a form for the X matrices and then tries to find the optimal factorization. [97] In contrast to the RI, this method can reduce the formal scaling of the exchange term, MP2 in the THC framework thus scales with the fourth power of the system size. Local variants of the THC also have been recently developed. [98]

Another strategy that has been proposed for the evaluation of the ERIs is the pseudo-spectral method. [99, 100] In this method the ERIs are evaluated with a combination of analytic integrations for Gaussian basis functions and real-space grid integrals. Talman used Fourier space techniques to obtain spherical expansions of the product of two NAOs on a point between their centers and then evaluated the Coulomb interaction of these expansions. [101–105]

2.11 NUMERIC ATOM-CENTERED ORBITALS

For my project I used the Fritz Haber Institute ab initio molecular simulations code, which is an electronic structure code for all-electron calculations using Numeric Atom-Centered Orbitals (NAOs) as basis set. NAOs are atom-centered basis functions like GTOs but they do not have any fixed analytic form. [106–109] NAOs are usually separated into an angular and radial part. The former is usually a real-valued spherical harmonic Y_{l_s, m_s} and the latter can be any smooth and normalizable function $R_s(|\mathbf{r}_1|)$ of the distance to the center.

$$\varphi_s(\mathbf{r}_1) = R_s(|\mathbf{r}_1|) Y_{l_s, m_s} \left(\frac{\mathbf{r}_1}{|\mathbf{r}_1|} \right) \quad (2.104)$$


The full flexibility of the radial part can be used to represent the correct near-nuclear behavior in a single basis function instead of having to approximate it by a superposition of several basis functions. Thus NAO basis sets are usually more compact than other basis sets and require less basis functions per atom to reach a similar accuracy. Furthermore, one can apply radial constraints to ensure that the basis functions are strictly zero beyond a given radius and do not have any Gaussian-like tails. This strictly limited radial confinement can be used to exploit sparsity in larger systems without invoking approximations like distance-based cut-off criteria. On the other hand, NAOs also have a significant drawback: there are no analytical solutions available for the ERIs and a direct evaluation of the six-dimensional

integrals is usually prohibitively expensive. Therefore, the RI or related techniques are mandatory in a NAO framework to evaluate these integrals with reasonable effort.

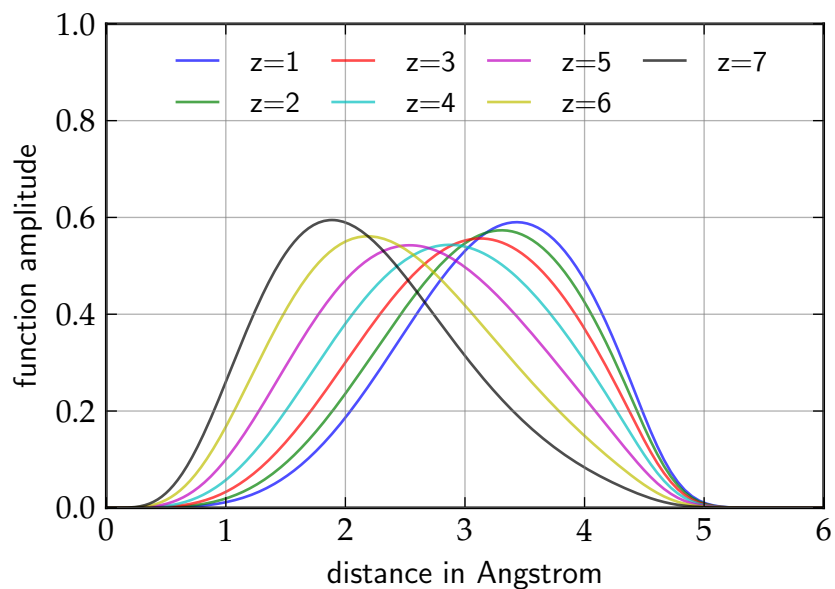
The Fritz Haber Institute ab initio molecular simulations (FHI-aims) code supports several types of NAOs which are determined as solutions of the radial Schrödinger equation shown in Equation 2.105. Atomic and ionic functions are obtained by using the LDA solutions of the free atom or ion to determine the effective potentials v_s . Hydrogen-like functions are the solutions for the hydrogen atom with an effective core charge z .

$$\left[-\frac{1}{2} \frac{d^2}{dr^2} + \frac{l_s(l_s + 1)}{2r^2} + v_s(r) + v_{cut}(r) \right] rR_s(r) = \epsilon_s rR_s(r) \quad (2.105)$$

All radial functions are subject to a confinement potential v_{cut} which forces the functions to smoothly decay to zero before reaching the chosen threshold radius. A selection of hydrogen-like functions is depicted in Figure 2.5. GTOs can of course also be represented as NAOs and are also supported in the FHI-aims code, but rarely used. A detailed discussion of the NAOs used in the FHI-aims code can be found in Refs. [110] and [32].



(a) shapes for confined functions with different angular momenta (The effective charge z is equal to 1 in all cases.)



(b) effect of different effective charges on the shape of a 5g hydrogen-like function.

Figure 2.5: Hydrogen-like radial functions in the FHI-aims code. The confinement potential starts at 4\AA and has a width of 2\AA .

As discussed in the previous chapter, a central challenge of many electronic structure theory methods is the evaluation of Electron Repulsion Integrals (ERIs). The Resolution of Identity (RI) strategy can be used to reduce the computational effort and required memory to a level that makes these methods routinely applicable for small and medium-sized systems. For large systems however, the cubic scaling memory requirement of the RI-V coefficients still poses a significant challenge. To overcome this limitation it is necessary to exploit the locality of atom-centered basis functions in the RI scheme.

In this chapter I will present the RI-LVL, a localized variant of the RI. In the first part of the chapter the construction of the RI-LVL and the underlying sparsity assumption are discussed. Afterwards it will be demonstrated that RI-LVL can provide the same level of accuracy as the standard RI-V when combined with an appropriate enhanced auxiliary basis set, while providing a superior computational performance and requiring considerably less memory. The material presented in this chapter has already been published in the *New Journal of Physics*. [111] The first version of the RI-LVL has been implemented into the Fritz Haber Institute ab initio molecular simulations (FHI-aims) code by Jürgen Wieferink. I finalized the implementation and performed the extensive accuracy and performance benchmarks presented in this chapter.

3.1 CONCEPT OF THE RI-LVL

The standard RI-V (discussed in section 2.10) is based on the idea to expand the products of basis function pairs $\varphi_s\varphi_t$ appearing in the ERIs in an auxiliary basis set (ABS) and then compute the Coulomb interactions between these auxiliary functions P_μ only.

$$\varrho_{s,t}(\mathbf{r}_1) = \varphi_s(\mathbf{r}_1)\varphi_t(\mathbf{r}_1) \approx \sum_{\mu}^{ABS} C_{s,t}^{\mu} P_{\mu}(\mathbf{r}_1) = \tilde{\varrho}_{s,t}(\mathbf{r}_1) \quad (3.1)$$

$$(st|uv) \approx \sum_{\mu}^{ABS} \sum_{\lambda}^{ABS} C_{s,t}^{\mu} (\mu|\lambda) C_{u,v}^{\lambda} \quad (3.2)$$

For the RI-V shown in Equation 3.2 it is necessary to store N_{aux} coefficients for each function pair, resulting in a total memory requirement of $\mathcal{O}(N_{\text{aux}}N_{\text{bas pairs}})$. Formally, the computational effort per ERI would be $\mathcal{O}(N_{\text{aux}}^2)$, but in practice the Coulomb matrix $V_{\mu,\lambda} = (\mu|\lambda)$ is split into a product of its square-root and absorbed into the RI coefficients.

This eliminates one of the sums over auxiliary basis functions and yields a $\mathcal{O}(N_{\text{aux}})$ effort per ERI.

$$(st|uv) \approx \sum_{\mu}^{\text{ABS}} \sum_{\lambda}^{\text{ABS}} \sum_{\kappa}^{\text{ABS}} [C_{s,t}^{\mu} V_{\mu,\kappa}^{0.5}] [V_{\kappa,\lambda}^{0.5} C_{u,v}^{\lambda}] \quad (3.3)$$

$$= \sum_{\mu}^{\text{ABS}} \tilde{C}_{s,t}^{\mu} \tilde{C}_{u,v}^{\lambda} \quad (3.4)$$

However, the product of the two Numeric Atom-Centered Orbitals (NAOs) $q_{s,t} = \varphi_s \varphi_t$ has a spatially confined extent and an auxiliary basis function on a distant atom should not be required to accurately represent this local quantity. One way to exploit this locality is the *à priori* restriction of the ABS used for the expansion to contain only those functions centered on the same atoms as φ_s or φ_t . (see Figure 3.1) This local ABS is denoted as $\mathcal{P}(S, T)$ and the resulting local RI method has been named RI-LVL.

$$q_{s,t}(\mathbf{r}_1) = \varphi_s(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \approx \sum_{\mu}^{\mathcal{P}(S,T)} C_{s,t}^{\mu} P_{\mu}(\mathbf{r}_1) = \tilde{q}_{s,t}(\mathbf{r}_1) \quad (3.5)$$

$$(st|uv) \approx \sum_{\mu}^{\mathcal{P}(S,T)} \sum_{\lambda}^{\mathcal{P}(U,V)} C_{s,t}^{\mu} (\mu|\lambda) C_{u,v}^{\lambda} \quad (3.6)$$

The expression in Equation 3.6 requires only a sparse subset of the RI coefficients, which yields a $\mathcal{O}(N_{\text{local aux}} N_{\text{bas pairs}})$ memory requirement. The cubic memory requirement of RI-V is thus reduced to a quadratic one. Additionally, both sums are now restricted to system-size independent subsets of the ABS, which yields a constant effort $\mathcal{O}(N_{\text{local aux}}^2)$ per ERI.

It should be noted that other research groups have also investigated the potential of locality exploits in the RI. Reine *et al* replaced the global Coulomb metric with a localized one in their work. [112] Sodt *et al* followed an approach more similar to RI-LVL and restricted the RI expansion to those functions within a sphere around the atom hosting the first basis function. [113, 114] Pisani *et al* followed a similar approach, but used more general fitting domains in periodic local-MP2 methods. [80, 115] Merlot *et al* tried to use only the auxiliary basis functions on the two atoms hosting the basis functions, but had to resort to adding additional auxiliary functions along the connection line of the atoms to reach the desired accuracy. [116] More recently the Head-Gordon group also published a more restricted approach using only the auxiliary basis functions on the two atoms, but only applied it to the exchange-type integrals in Hartree-Fock (HF). [117] This last publication is similar to my own work and has been published while the RI-LVL manuscript was in preparation.

3.2 THE DEFINITION OF THE RI-LVL COEFFICIENTS

As in the standard RI, there is no unique definition of the expansion coefficients in this localized approach. Following the strategy of the




Figure 3.1: A schematic representation of the RI-LVL concept. Only the auxiliary basis functions centered on the atoms S and T are used to expand the products of basis functions centered on them.

standard RI-V, I choose the coefficients such that they minimize the Coulomb self-interaction $(\delta q_{s,t} | \delta q_{s,t})$ of the residual fitting error $\delta q_{s,t}$.

$$\delta q_{s,t}(\mathbf{r}_1) = \sum_{\mu}^{\mathcal{P}(S,T)} C_{s,t}^{\mu} P_{\mu}(\mathbf{r}_1) - \varphi_s(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \quad (3.7)$$

$$\begin{aligned} (\delta q_{s,t} | \delta q_{s,t}) &= \iint \left(\sum_{\mu}^{\mathcal{P}(S,T)} C_{s,t}^{\mu} P_{\mu}(\mathbf{r}_1) - \varphi_s(\mathbf{r}_1) \varphi_t(\mathbf{r}_1) \right) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \\ &\quad \times \left(\sum_{\lambda}^{\mathcal{P}(S,T)} C_{s,t}^{\lambda} P_{\lambda}(\mathbf{r}_2) - \varphi_s(\mathbf{r}_2) \varphi_t(\mathbf{r}_2) \right) d\mathbf{r}_1 d\mathbf{r}_2 \quad (3.8) \end{aligned}$$

$$= \sum_{\mu,\lambda}^{\mathcal{P}(S,T)} C_{s,t}^{\mu} C_{s,t}^{\lambda} (\mu|\lambda) - 2 \sum_{\mu}^{\mathcal{P}(S,T)} C_{s,t}^{\mu} (st|\mu) + (st|st) \quad (3.9)$$

To minimize this term, I take the derivative with respect to the coefficients for each auxiliary basis function P_{μ} and solve the resulting system of equations.

$$\frac{\partial (\delta q_{s,t} | \delta q_{s,t})}{\partial C_{s,t}^{\mu}} \stackrel{!}{=} 0 \quad \forall \mu \in \mathcal{P}(S, T) \quad (3.10)$$

$$\Rightarrow 0 = \sum_{\lambda}^{\mathcal{P}(S,T)} C_{s,t}^{\lambda} (\mu|\lambda) - (st|\mu) \quad (3.11)$$

$$\Rightarrow \sum_{\lambda}^{\mathcal{P}(S,T)} C_{s,t}^{\lambda} (\mu|\lambda) = (st|\mu) \quad \forall \mu \in \mathcal{P}(S, T) \quad (3.12)$$

Introducing the local Coulomb matrix $L_{\mu,\lambda}^{S,T} = (\mu|\lambda) \forall \mu, \lambda \in \mathcal{P}(S, T)$, this set of equations for the P_μ can be rewritten as a matrix equation.

$$\sum_{\lambda}^{\mathcal{P}(S,T)} L_{\mu,\lambda}^{S,T} C_{s,t}^\lambda = (st|\mu) \quad \forall \mu \in \mathcal{P}(S, T) \quad (3.13)$$

$$\Rightarrow C_{s,t}^\lambda = \begin{cases} \sum_{\mu}^{\mathcal{P}(S,T)} (L^{S,T})_{\mu,\lambda}^{-1} (st|\mu) & \lambda \in \mathcal{P}(S, T) \\ 0 & \text{else} \end{cases} \quad (3.14)$$

Inserting these expansion coefficients into Equation 3.6, one obtains

$$(st|uv) \approx \sum_{\mu,\lambda}^{\mathcal{P}(S,T)} \sum_{\kappa,\nu}^{\mathcal{P}(U,V)} (st|\mu) \left(L^{S,T}\right)_{\mu,\lambda}^{-1} V_{\lambda,\kappa} \left(L^{U,V}\right)_{\kappa,\nu}^{-1} (v|uv) \quad (3.15)$$

as final expression for the ERIs in this approximation. The matrix product appearing in the last equation gave rise to the method's name "RI-LVL".

3.3 THE CONSTRUCTION OF THE AUXILIARY BASIS SET

In the FHI-aims code [110], the ABS is not a predefined set of functions like the orbital basis set (OBS), but is instead generated at runtime based on the used orbital basis functions. As their counterparts in the OBS, the auxiliary basis functions P_μ are NAOs consisting of a radial part R_μ and a spherical harmonic $Y_{l,m}$.

$$P_\mu(\mathbf{r}_1) = R_\mu(\mathbf{r}_1) Y_{l_\mu, m_\mu}(\mathbf{r}_1) \quad (3.16)$$

The ERIs $(st|uv)$ with the most significant contributions to the total energy are those where the involved basis functions are centered on the same atom. Thus the ABS is constructed such that it contains all the onsite products $\varphi_s \varphi_t$ exactly to ensure a high accuracy of these integrals. The procedure for generating the ABS is outlined in Figure 3.2 and is the same as the one used for RI-V based calculations in FHI-aims. [96]

At first all possible products of the radial functions used in the OBS are computed for each species in the system. A Gram-Schmidt-Orthonormalization is then applied to remove any linearly dependent products from the resulting set. The angular part of the function products is represented as a spherical harmonics expansion, which is restricted by the Clebsch-Gordon coefficients to those angular momenta l satisfying $|l_1 - l_2| \leq l \leq |l_1 + l_2|$. [118] l_1 and l_2 are the angular momenta associated with the radial functions forming the radial part of the product. Each obtained radial function product is thus multiplied with all spherical harmonics that satisfy this condition and the resulting NAOs are then distributed to all atoms of the species in the system to form the ABS.

This general procedure for the construction of the ABS has been demonstrated to be very accurate in the standard RI-V approach. [96] In addition, it also provides a systematic way to further improve the ABS according to higher requirements. As indicated by the gray

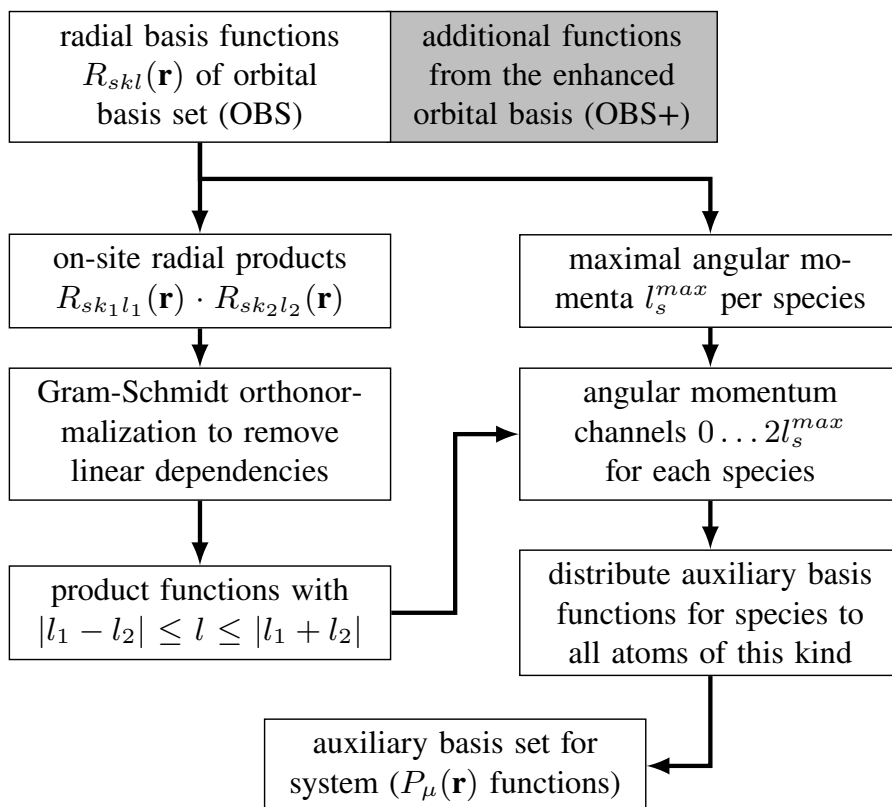


Figure 3.2: The construction scheme for the ABS from the OBS. (The radial function subscripts s, k, l denote species, radial function index and angular momentum.) A detailed explanation of the steps in this scheme is provided in the text. One can also construct the ABS from an OBS+, which is a superset of the OBS and contains additional functions which are not used in the expansion of the KS or HF one-electron orbitals, but only for the ABS construction. (gray box)

	OBS	OBS+
number of φ_s in OBS/OBS+	39	39
radial functions types in OBS/OBS+	4s, 3p, 2d, 1f, 1g	4s, 3p, 2d, 1f, 2g
number of onsite radial function products	150	198
number of linearly independent onsite radial products	50	70
radial functions types in the ABS/ABS+	10s, 9p, 8d, 7f, 6g, 4h, 3i, 2j, 1k	11s, 10p, 9d, 10f, 9g, 7h, 6i, 5j, 3k
number of P_μ in ABS/ABS+	310	518

Table 3.1: The size of the ABS for a carbon atom with a standard and augmented auxiliary basis for the tier2 basis set in FHI-aims, corresponding to the steps outlined in Figure 3.2. The notation $2d$ here indicates the inclusion of two different radial d -type functions, which are then combined with all $Y_{l,m}$ of the given angular momentum.

box in Figure 3.2, a superset of the OBS, the enhanced orbital basis set (OBS+), can be used as input to generate a larger ABS labeled as enhanced auxiliary basis set (ABS+). It is important to note that these additional functions are not used in the expansion of the HF or Kohn-Sham (KS) orbitals, all other parts of the calculation are thus not affected by these functions. This new ABS+ provides the same accuracy for on-site integrals, but yields a larger number of auxiliary basis functions per atom, which allows a more accurate representation of off-site integrals. Table 3.1 shows a comparison of the ABS generated for the carbon atom from the standard tier2 NAO basis set [110] and the same basis set enhanced by an additional g -type function.

For simplicity of notation, I will only use the term ABS hereafter for both the ABS constructed from the standard OBS and the ABS+.

3.4 FORMAL ERROR ANALYSIS

For any ERI represented in a RI approximation, Equation 3.17 holds exactly. $\tilde{q}_{s,t}$ is the RI approximation to the exact pair product $q_{s,t}$, and $\delta q_{s,t} = \tilde{q}_{s,t} - q_{s,t}$ is the fitting residual.

$$(st|uv) = (\tilde{q}_{s,t}|\tilde{q}_{u,v}) + (\tilde{q}_{s,t}|\delta q_{u,v}) + (\delta q_{s,t}|\tilde{q}_{u,v}) + (\delta q_{s,t}|\delta q_{u,v}) \quad (3.17)$$

The minimization criterion for RI-V (Equation 2.98) can be rewritten as $0 = (\mu|\delta q_{s,t})$. This implies that the terms in Equation 3.17 which are linear in the fitting residual must vanish since $\tilde{q}_{s,t}$ is a superposition of the $\{P_\mu\}$. The RI-V method therefore has a quadratic error with respect to the fitting residual.

RI-LVL does not inherit this feature due to the restricted optimization criterion in Equation 3.10. This potential problem of localized RI

methods can be addressed by the Dunlap correction. [119] By including the linear error terms explicitly in the approximation of the ERI, one obtains Equation 3.18.

$$(\widetilde{st|uv}) = (\widetilde{q}_{s,t}|q_{u,v}) + (q_{s,t}|\widetilde{q}_{u,v}) - (\widetilde{q}_{s,t}|\widetilde{q}_{u,v}) \quad (3.18)$$

This expression however has a few significant drawbacks. To compute these correction terms, one has to evaluate and store all three-center integrals $(\mu|st)$ as in RI-V instead of just those where $\mu \in \mathcal{P}(ST)$. This would nullify all memory savings of RI-LVL and also a considerable amount of the computational time savings. Furthermore, it has been reported [116] that this correction can lead to severe convergence issues in rare cases because the matrix of the corrected ERIs is not guaranteed to be positive semi-definite. A negative eigenvalue in this matrix corresponds to an attractive electron-electron interaction and will lead to a high ratio of rejected Self-Consistent Field (SCF) steps or even a total failure to converge. The Dunlap correction is not used in RI-LVL because of these drawbacks. In the next sections I will show that very accurate RI-LVL results can be obtained if an appropriate ABS is employed.

3.5 RI-LVL IN A RI-V FRAMEWORK

An efficient RI-LVL implementation which fully exploits the imposed sparsity is available in the FHI-aims code for HF and hybrid functionals. Its linear scaling exchange evaluation with dramatically lowered memory consumption will be discussed in section 3.9. However, for advanced correlation methods like 2nd Order Møller-Plesset Perturbation Theory (MP2) and the Random-Phase Approximation (RPA) it is still a challenge to design a local RI implementation from both theoretical and technical points of view and will be covered in the following chapters.

To nevertheless judge the accuracy of this new RI at those levels of theory, an additional variant, RI-LVL_{full}, has been implemented. This method uses the expansion coefficients of the RI-LVL (Equation 3.14), but transforms their sparse representation back into a dense tensor structure by explicitly storing all the zero entries. These bloated-up coefficients are then multiplied with the square-root of the Coulomb matrix as in the RI-V (see Equation 3.3) and afterwards used in the existing RI-V routines. With RI-LVL_{full} it is possible to judge the accuracy of RI-LVL itself for beyond Density-Functional Theory (DFT) methods if no further approximation is used. However, RI-LVL_{full} combines the disadvantages from both RI-V and RI-LVL, because it requires the larger ABS from the RI-LVL and due to the back-transformation all memory savings are lost. RI-LVL_{full} is therefore primarily a benchmark tool, but not a recommended choice for production calculations.

3.6 RI-LVL IN PERIODIC SYSTEMS

The periodic version of the RI-LVL has been tested extensively by my colleague Sergey Levchenko in Ref [78]. This implementation uses the equivalence between a primitive unit cell with k-grid and a Born von Karman (BvK) supercell as discussed in section 2.8. The exchange contribution to the Fock-matrix is evaluated in the real-space supercell where the locality from RI-LVL can be combined with integral screening techniques to improve the performance. I will not present his results here because no noteworthy differences have been found for the behavior of RI-LVL in cluster and periodic systems. Instead, I will only discuss how the RI-LVL affects the evaluation of ERIs formed by crystalline basis functions. The general RI expression for periodic systems is obtained by expanding the crystalline basis function pairs in a set of Bloch-like auxiliary basis functions $P_{\mu,\mathbf{k}_1}(\mathbf{r}_1)$ with not yet determined coefficients $C_{s,t}^{\mu}(\mathbf{k}_1, \mathbf{q}_1)$.

$$\varphi_{s,\mathbf{k}_1}^*(\mathbf{r}_1)\varphi_{t,\mathbf{q}_1}(\mathbf{r}_1) = \sum_{\mu} C_{s,t}^{\mu}(\mathbf{k}_1, \mathbf{q}_1)P_{\mu,\mathbf{q}_1-\mathbf{k}_1}(\mathbf{r}_1) \quad (3.19)$$

$$P_{\mu,\mathbf{k}_1}(\mathbf{r}_1) = \frac{1}{\sqrt{N}} \sum_{\mathbf{R}_1} e^{i\mathbf{k}_1\mathbf{R}_1} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_1) \quad (3.20)$$

The set of auxiliary basis functions in Equation 3.19 can be restricted by the fact that the overlap $\int \varphi_{s,\mathbf{k}_1}^*(\mathbf{r}_1)\varphi_{t,\mathbf{q}_1}(\mathbf{r}_1)P_{\mu,\mathbf{k}_2}(\mathbf{r}_1)d\mathbf{r}_1$ is non-zero only if $\mathbf{k}_2 + \mathbf{k}_1 - \mathbf{q}_1 = 0$.

To exploit locality in Equation 3.19, the RI-LVL concept can be applied to the pairs of atomic basis functions in different unit cells which form the crystalline basis functions.

$$\begin{aligned} \varphi_s(\mathbf{r}_1 - \mathbf{R}_1)\varphi_t(\mathbf{r}_1 - \mathbf{R}_2) &= \sum_{\mu}^{\mathcal{P}(S)} C_{s(\mathbf{R}_1),t(\mathbf{R}_2)}^{\mu(\mathbf{R}_1)} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_1) \\ &+ \sum_{\mu}^{\mathcal{P}(T)} C_{s(\mathbf{R}_1),t(\mathbf{R}_2)}^{\mu(\mathbf{R}_2)} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_2) \end{aligned} \quad (3.21)$$

Upon inserting these RI-LVL expansions into the crystalline basis functions, one obtains the RI-LVL approximation for crystalline basis function ERIs.

$$\varphi_{s,\mathbf{k}_1}^*(\mathbf{r}_1)\varphi_{t,\mathbf{q}_1}(\mathbf{r}_1) = \frac{1}{N} \sum_{\mathbf{R}_1, \mathbf{R}_2} e^{-i\mathbf{k}_1\mathbf{R}_1} e^{i\mathbf{q}_1\mathbf{R}_2} \varphi_s(\mathbf{r}_1 - \mathbf{R}_1)\varphi_t(\mathbf{r}_1 - \mathbf{R}_2) \quad (3.22)$$

$$\begin{aligned} &= \frac{1}{N} \sum_{\mathbf{R}_1, \mathbf{R}_2} e^{-i\mathbf{k}_1\mathbf{R}_1} e^{i\mathbf{q}_1\mathbf{R}_2} \left[\sum_{\mu}^{\mathcal{P}(S)} C_{s(\mathbf{R}_1),t(\mathbf{R}_2)}^{\mu(\mathbf{R}_1)} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_1) \right. \\ &\quad \left. + \sum_{\mu}^{\mathcal{P}(T)} C_{s(\mathbf{R}_1),t(\mathbf{R}_2)}^{\mu(\mathbf{R}_2)} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_2) \right] \end{aligned} \quad (3.23)$$

Here the sums over \mathbf{R}_1 and \mathbf{R}_2 include all cell pairs for which $\varphi_s(\mathbf{r}_1 - \mathbf{R}_1)\varphi_t(\mathbf{r}_1 - \mathbf{R}_2)$ has a finite overlap. The RI coefficients feature a trans-

lational invariance, i.e. $C_{s(\mathbf{R}_1),t(\mathbf{R}_2)}^{\mu(\mathbf{R}_1)} = C_{s(\mathbf{R}_0),t(\mathbf{R}_2-\mathbf{R}_1)}^{\mu(\mathbf{R}_0)}$ where \mathbf{R}_0 denotes the central unit cell.

$$\begin{aligned} \varphi_{s,\mathbf{k}_1}^*(\mathbf{r}_1)\varphi_{t,\mathbf{q}_1}(\mathbf{r}_1) &= \frac{1}{N} \sum_{\mathbf{R}_1,\mathbf{R}_2} e^{-i\mathbf{k}_1\mathbf{R}_1} e^{i\mathbf{q}_1\mathbf{R}_2} \left[\sum_{\mu}^{\mathcal{P}(S)} C_{s(\mathbf{R}_0),t(\mathbf{R}_2-\mathbf{R}_1)}^{\mu(\mathbf{R}_0)} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_1) \right. \\ &\quad \left. + \sum_{\mu}^{\mathcal{P}(T)} C_{s(\mathbf{R}_1-\mathbf{R}_2),t(\mathbf{R}_0)}^{\mu(\mathbf{R}_0)} P_{\mu}(\mathbf{r}_1 - \mathbf{R}_2) \right] \end{aligned} \quad (3.24)$$

$$= \sum_{\mu} [C_{t,s}^{\mu}(-\mathbf{k}_1, \mathbf{R}_0) + C_{s,t}^{\mu}(\mathbf{q}_1, \mathbf{R}_0)] P_{\mu,\mathbf{q}_1-\mathbf{k}_1}(\mathbf{r}_1) \quad (3.25)$$

$$C_{s,t}^{\mu}(\mathbf{q}_1, \mathbf{R}_0) = \begin{cases} \frac{1}{\sqrt{N}} \sum_{\mathbf{R}_1} e^{i\mathbf{q}_1\mathbf{R}_1} C_{s(\mathbf{R}_1),t(\mathbf{R}_0)}^{\mu(\mathbf{R}_0)} & \mu \in \mathcal{P}(S) \\ 0 & \text{else} \end{cases} \quad (3.26)$$

The definition of the reciprocal space RI coefficients within the RI-LVL strategy can now be obtained by comparing Equation 3.25 and 3.19.

$$C_{s,t}^{\mu}(\mathbf{k}_1, \mathbf{q}_1) = C_{t,s}^{\mu}(-\mathbf{k}_1, \mathbf{R}_0) + C_{s,t}^{\mu}(\mathbf{q}_1, \mathbf{R}_0) \quad (3.27)$$

RI-LVL is thus able to represent the RI coefficients in a form that depends on a single k-point only. Therefore the persistent memory requirement scales linearly with the number of k-points. Finally, the ERIs of the crystalline basis functions are given as

$$\begin{aligned} (\varphi_{s,\mathbf{k}_1}\varphi_{t,\mathbf{q}_1} | \varphi_{u,\mathbf{k}_2}\varphi_{v,\mathbf{q}_2}) &= \sum_{\mu}^{\mathcal{P}(ST)} \sum_{\lambda}^{\mathcal{P}(UV)} C_{s,t}^{\mu}(\mathbf{k}_1, \mathbf{q}_1) \\ &\quad \times V_{\mu,\lambda}(\mathbf{q}_1 - \mathbf{k}_1, \mathbf{q}_2 - \mathbf{k}_2) C_{u,v}^{\lambda}(\mathbf{k}_2, \mathbf{q}_2) \end{aligned} \quad (3.28)$$

where $V_{\mu,\lambda}(\mathbf{k}_1, \mathbf{q}_1)$ is the Coulomb matrix element of two auxiliary Bloch basis functions P_{μ,\mathbf{q}_1} and P_{λ,\mathbf{k}_1} .

3.7 ACCURACY OF THE RI-LVL

In the following section I will at first demonstrate the accuracy of RI-LVL with a suitable ABS in a direct comparison to NorthWest computational Chemistry (NWChem) [120], a Gaussian code which can evaluate the ERIs analytically. This overview section is followed by a detailed analysis of the impact the ABS choice has on the accuracy for HF, MP2, PBE0 and RPA calculations.

As explained in section 3.3 all on-site products of basis function pairs are represented exactly by the employed ABS. The change from RI-V to RI-LVL therefore only affects the off-site integrals. The OBS+ must be designed such that the auxiliary basis functions on the two host atoms in off-site integrals can make up for the lost flexibility in the restricted expansion. The additions from the OBS+ should also be chosen as general as possible to ensure a broad applicability to many systems. Following this rationale, only hydrogen-like functions with the lowest principal quantum number for the given angular momentum have been added to the OBS+ throughout this project. The potential in Equation 2.105 is therefore $V(r) = -\frac{z}{r}$ where z is the effective charge assigned to the hydrogen-like function. Hydrogen-like

functions have the distinct feature that their shape is not depending on the element of the host atom. The products of OBS+ function pairs thus provide a common set of new auxiliary basis functions that is shared by all atom kinds in the system. At the same time, the products between radial OBS and OBS+ functions will also give rise to additional, atom-kind specific auxiliary basis functions. The OBS+ in the following is thus the set of hydrogen-like functions which are added to all atomic species in the system. As shown later, the results do not show strong dependencies on the choice of the effective charge z , thus all functions in the OBS+ will use the same effective charge.

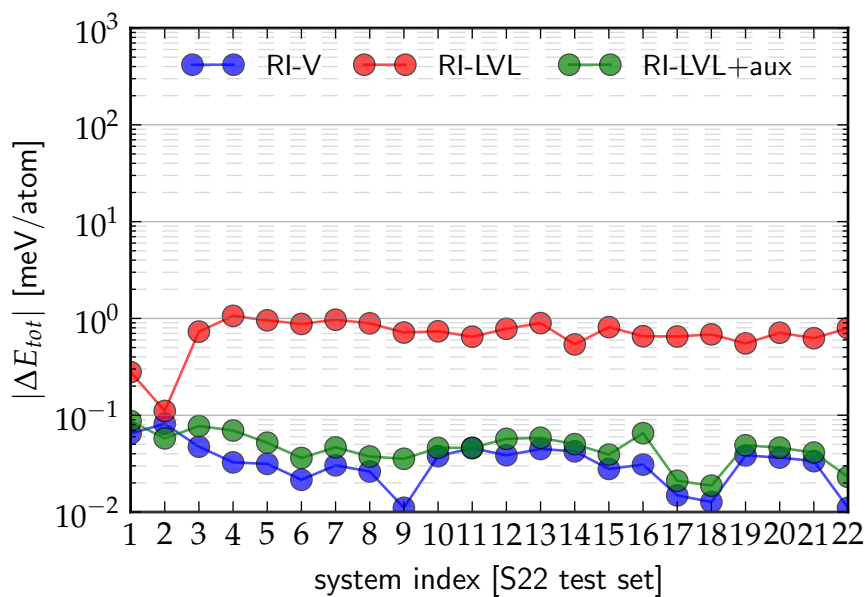
The notation $spd(z=x)$ denotes the functions that have been added to the OBS+, where spd represent the angular momenta of the added functions (s -type, p -type function, etc.) and $z=x$ shows the effective charge used for the potential in the radial function generation.

Unless mentioned otherwise, "tight" integration settings have been used for all calculations. All MP2 and RPA calculations for RI-LVL have been done with the RI-LVL_{full} algorithm and are therefore only suitable to demonstrate the reachable accuracy, but not the computational performance.

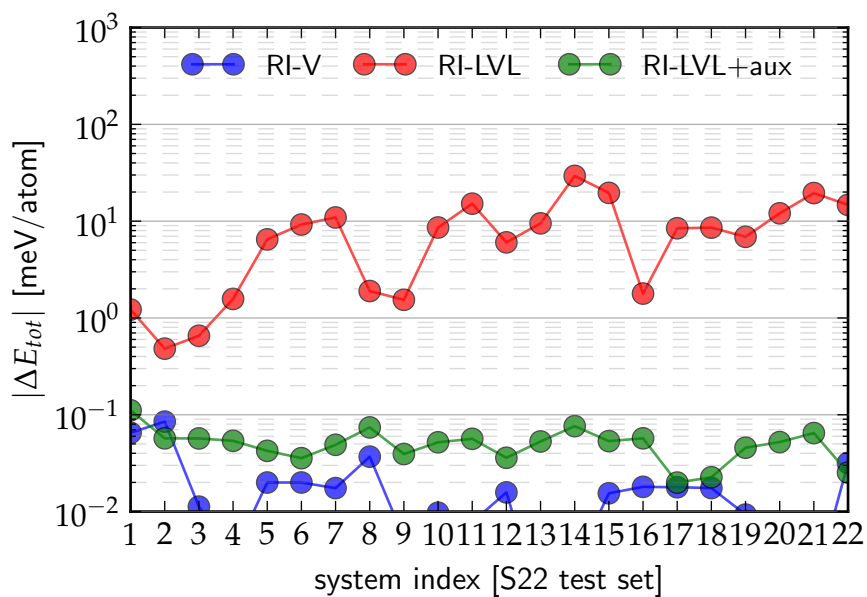
In the following, the RI-LVL calculations with different ABS are benchmarked against the RI-V calculations with the standard ABS constructed from the pure OBS. Adding additional auxiliary basis functions to these RI-V calculations has no significant impact on the results. (Data can be found in Appendix A.) "Errors per atom" denote the error in the total energy divided by the number of non-hydrogen atoms.

3.7.1 Gaussian Basis Set Comparison

ERIs cannot be evaluated directly with reasonable cost in a NAO framework, thus the reference point for benchmarking the RI-LVL must be the standard RI-V implementation, which is an approximation itself. To confirm that the RI-V calculations are a well-chosen reference, I used Dunning's valence correlation-consistent cc-pVTZ Gaussian basis set [33] and compared the HF and MP2 total energies to the NWChem code. The geometries from the S22 test set [9] are chosen as the benchmark data set, which contain 22 weakly bonded molecular dimers with 6 to 30 atoms per structure. The included systems range from water dimers to adenine-thymine complexes. The results of this comparison are shown in Figure 3.3, where I also included the results of RI-LVL calculations using either the standard ABS or one constructed from an OBS+ with an additional g -type function. The RI-V implementation in FHI-aims provides very accurate results with total energy errors below 0.1 meV/atom for both HF and MP2 calculations. RI-LVL with the ABS constructed from the standard OBS however shows considerably larger errors, about 1 meV/atom for HF and up to 30 meV/atom for MP2 calculations. If the ABS is enhanced by adding a g -type function to the OBS used to



(a) HF total energy errors per atom



(b) MP2 total energy errors per atom

Figure 3.3: Absolute total energy errors in the S22 test set for HF and MP2 calculations with the cc-pVTZ basis set for different RI methods. The reference values are the NWChem calculations with analytic ERIs. In the data series denoted as "RI-LVL+aux" an OBS+ with an additional g-type function has been used to construct the ABS.




Figure 3.4: Absolute total energy per atom errors for RI-LVL based HF calculations in the S22 test set with the tier2 basis set and different choices for the ABS. All OBS+ functions are using $z = 1.0$. The reference values are the RI-V calculations without OBS+.

generate it, RI-LVL provides the same level of accuracy as RI-V with maximum errors on the scale of 0.1 meV/atom for both HF and MP2.

Since the RI-V implementation in FHI-aims has proven to be very accurate, RI-V results will be used from here on as reference to benchmark the accuracy of RI-LVL with NAO basis sets.

3.7.2 HF calculations

To assess the accuracy of RI-LVL for HF theory, I analyzed the impact of different ABS choices on the resulting total energy. The performance of RI-LVL is shown in more detail in Figure 3.4 for all systems in the S22 test set. RI-LVL with the ABS generated from the standard OBS gives acceptable, but notable errors with a Root-Mean Square Deviation (RMSD) of 0.6 meV/atom and a maximum error of 1.9 meV/atom. This plot suggests that RI-LVL can be systematically converged very close to the RI-V results by adding functions of increasing angular momentum to the OBS+. If one follows this hierarchical approach up to the g -type function, the errors can be reduced to negligible values of less than 0.13 meV/atom with a RMSD below 0.06 meV/atom. For practical applications, the ABS should of course be chosen as small as possible to optimize the performance and memory footprint. From the plot it is also evident that low angular momentum functions provide almost no improvement. Probably the low angular momentum channels of the ABS are already well-saturated by the radial function products provided by the OBS.

<i>HF RI-LVL</i>	$ \Delta E_{tot} $ (meV/atom)	
	RMSD	MAX
<i>OBS+</i>		
<i>tier2</i>	0.594	1.891
<i>tier2 + spdfg(z=1)</i>	0.052	0.124
<i>tier2 + g(z=1)</i>	0.066	0.153
<i>tier2 + g(z=2)</i>	0.059	0.137
<i>tier2 + g(z=3)</i>	0.054	0.127
<i>tier2 + g(z=4)</i>	0.053	0.125
<i>tier2 + g(z=5)</i>	0.053	0.123
<i>tier2 + g(z=6)</i>	0.047	0.107
<i>tier2 + g(z=7)</i>	0.046	0.098

Table 3.2: Total energy per atom errors for RI-LVL based HF calculations with tier2 basis sets and different choices of the effective charge assigned to the g -type function in the *OBS+*.

The f -type and g -type functions on the other hand provide a substantial improvement to the accuracy and the question arises if one of these functions alone would provide a similar accuracy as the hierarchical approach. Table 3.2 shows how different effective charges for a single g -type function in the *OBS+* affect the accuracy of the RI-LVL. Regardless of the chosen effective charge, the single g -type function provides the same level of accuracy as the hierarchical approach. To understand the independence on the chosen effective charge, it is necessary to remember the construction of the radial basis functions in FHI-aims. As explained in section 2.11, a confinement potential is employed for the construction of the radial functions to limit their spatial extent. As a consequence, the hydrogen-like functions with different effective charges are confined to the same region in space. (see Figure 2.5) If no confinement potential would be applied, the g -type functions with different effective charges would differ considerably in the outer region and the results would probably show a notable dependence on the effective charge.

3.7.3 *MP2 calculations*

As aforementioned in subsection 3.7.1, the errors introduced by the RI-LVL are much more significant for MP2 than for HF. I therefore repeated the systematic accuracy analysis shown in the previous section for the MP2. The MP2 calculations use the NAO-VCC-3Z valence-correlation consistent NAO basis set [32] which has been designed for MP2 and other beyond DFT methods.

Figure 3.5 shows the same systematic convergence for RI-LVL with increasing angular momentum of the functions in the *OBS+* that was already observed for HF theory. The RI-LVL errors when using just the standard ABS are considerably larger and have a RMSD of 14 meV/atom. (largest errors up to 33 meV/atom) By including up to g -type functions in the *OBS+* the results can be converged within 0.1 meV/atom. As it was the case in the HF analysis, the higher angular




Figure 3.5: Absolute total energy per atom errors for RI-LVL based MP2@HF calculations in the S22 test set with the NAO-VCC-3Z basis set and different choices for the ABS. All OBS+ functions are using $z = 1.0$. The reference values are the RI-V calculations without OBS+.

<i>MP2@HF RI-LVL</i>	$ \Delta E_{tot} $ (meV/atom)	
	RMSD	MAX
<i>OBS+</i>		
NAO-VCC-3Z	13.919	32.970
NAO-VCC-3Z + <i>spdfg</i> ($z=1$)	0.047	0.104
NAO-VCC-3Z + <i>g</i> ($z=1$)	0.063	0.111
NAO-VCC-3Z + <i>g</i> ($z=2$)	0.064	0.114
NAO-VCC-3Z + <i>g</i> ($z=3$)	0.059	0.104
NAO-VCC-3Z + <i>g</i> ($z=4$)	0.054	0.097
NAO-VCC-3Z + <i>g</i> ($z=5$)	0.036	0.070
NAO-VCC-3Z + <i>g</i> ($z=6$)	0.028	0.074
NAO-VCC-3Z + <i>g</i> ($z=7$)	0.046	0.082

Table 3.3: Total energy errors for MP2@HF calculations using RI-LVL and NAO-VCC-3Z basis sets with different choices for the ABS. The reference is the RI-V calculations using the standard ABS.

<i>PBE0 RI-LVL</i>	ΔE_{tot} (meV/atom)	
	RMSD	MAX
<i>OBS+</i>		
<i>tier2</i>	0.068	0.185
<i>tier2 + s(z=1)</i>	0.070	0.188
<i>tier2 + sp(z=1)</i>	0.060	0.159
<i>tier2 + spd(z=1)</i>	0.039	0.091
<i>tier2 + spdf(z=1)</i>	0.012	0.028
<i>tier2 + spdfg(z=1)</i>	0.007	0.017

Table 3.4: Total energy errors per atom for PBE0 calculations using RI-LVL and tier2 basis sets with different choices for the ABS. RI-V calculations using the standard ABS are used as reference point.

momentum functions play a key role in improving the accuracy of the results. Table 3.3 shows that adding a single g -type function to the OBS+ suffices to reduce the errors to the scale of 0.1 meV/atom. Therefore, the OBS+ containing a single g -type function in addition to the functions from the OBS is sufficient to provide very accurate results with the RI-LVL method in both HF and MP2 calculations.

3.8 PBE0 AND RPA CALCULATIONS

To demonstrate RI-LVL’s broad range of applicability, I also analyzed its performance for PBE0 and RPA. The accuracy analysis follows the same strategy as for HF and MP2, which was shown in detail in the previous sections. The PBE0 calculations are performed with the tier2 basis set and the NAO-VCC-3Z basis set is used for the RPA calculations.

Table 3.4 shows the impact of the different OBS+ choices for RI-LVL based PBE0 calculations in the S22 test set. In contrast to HF theory, the ABS constructed from the standard OBS is already sufficient to provide accurate results with errors below 0.2 meV/atom. It can be further reduced by an order of magnitude if a g -type function is added to the OBS+. The different behavior compared to RI-LVL based HF calculations is not surprising because the PBE0 functional only contains 25% exact exchange. (see Equation 2.61 and context) The errors made in the ERIs thus have a smaller impact on the total energy computed in each SCF step and the final converged SCF solution.

The RPA@PBE0 calculations with different choices for the OBS+ are shown in Figure 3.6. The calculations exhibit the same systematic convergence with increasing angular momentum of the added functions in the OBS+ as in the MP2 calculations. Adding a single g -type function to the OBS+ also gives very accurate results with errors below 0.4 meV/atom in all cases and average errors below 0.1 meV/atom. The RPA results obtained from the OBS+ with a single g -type function are also very accurate and do not show any significant dependence on the chosen effective charge of the added g -type function. (see Table A.5)

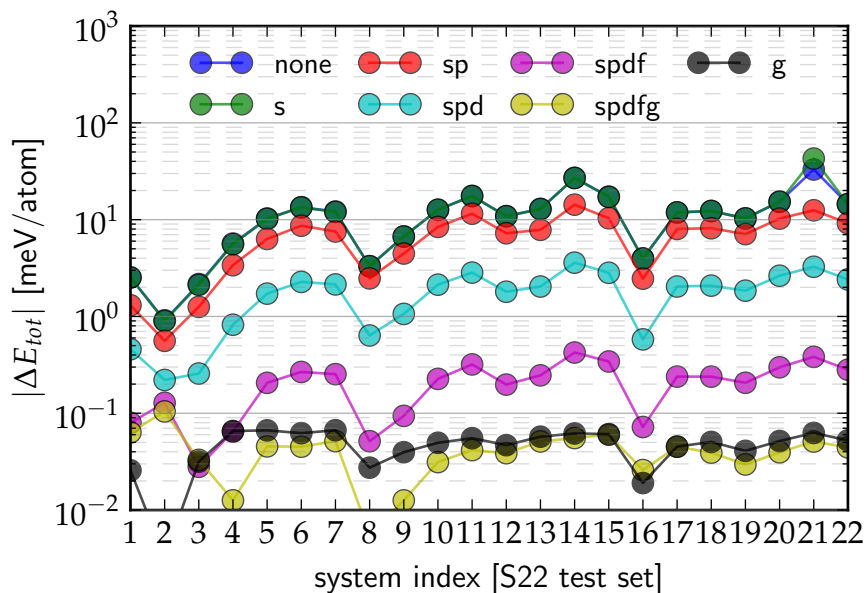


Figure 3.6: Absolute total energy per atom errors for RI-LVL based RPA@PBE0 calculations in the S22 test set with the NAO-VCC-3Z basis set and different choices for the ABS. All OBS+ functions are using $z = 1.0$. The reference values are the RI-V calculations without OBS+.




Figure 3.7: Clusters used to benchmark the accuracy of RI-LVL. The geometries are taken from references [121–123].

3.8.1 RI-LVL and heavier elements

In the previous sections, the S22 test set was used to benchmark the accuracy of RI-LVL, but it contains no element beyond the second row of the periodic table. To verify that RI-LVL also provides reliable results for heavier elements, I complemented the PBE0 and RPA results with calculations for copper, gold and titanium dioxide clusters. Figure 3.7 shows the geometries of all calculated clusters. Since there is no valence-correlation consistent NAO basis set available for the elements in these clusters, the tier4 (copper) and tier3 (all other elements) basis sets have been used instead. These are the largest NAO basis sets available in FHI-aims for these elements and thus also serve as a test if RI-LVL can handle very large basis sets. The OBS+ of the oxygen atoms in titanium dioxide was enhanced with a g -type func-

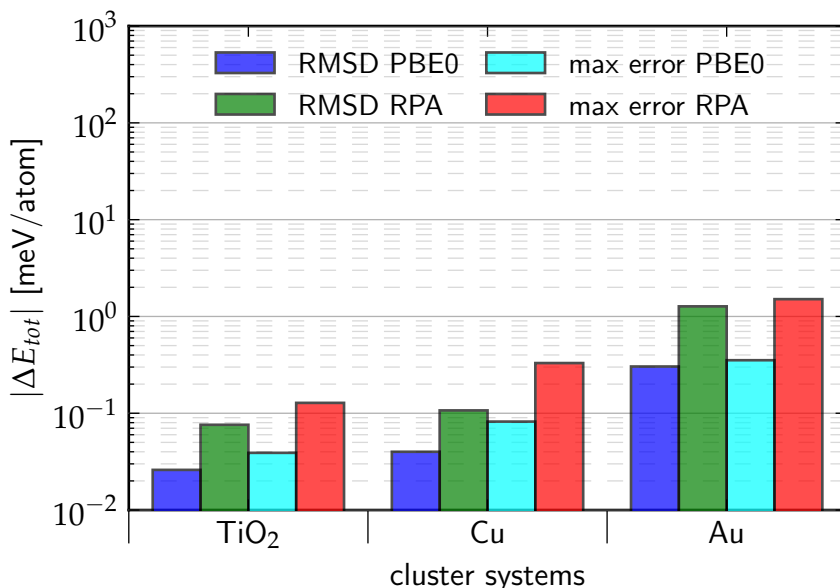


Figure 3.8: Total energy errors per atom of the RI-LVL RPA@PBE0 calculations for TiO₂, Cu and Au clusters. The RI-V calculations with the standard ABS are the reference values. Except for the oxygen atoms in TiO₂, no OBS+ functions were used in either RI-V or RI-LVL calculations. Details about the used basis sets can be found in the text.

tion to ensure that only the impact of the ABS on the titanium atoms is reflected in the results.

For these elements “very tight” integration settings are needed for the RI-V approach to ensure a sufficient accuracy for the Coulomb matrix of the auxiliary basis functions. It should be noted that RI-LVL does not need these integration settings because the Coulomb matrix is evaluated with the logarithmic spherical Bessel transforms (logSBT) algorithm, [102, 103] while the RI-V relies on grid-based integrations. For the sake of consistency, the same integration settings are used for RI-LVL in these calculations.

Figure 3.8 shows the results for both PBE0 and RPA calculations in this setup. In all systems, the PBE0 errors are below 0.5 meV/atom and the RPA are errors below 1.6 meV/atom. RI-LVL using the standard ABS is already well converged for heavier elements. Adding additional functions to the OBS+ to construct the ABS does not yield any notable change in the results. This observation also holds for smaller basis sets. (see Table A.6 for details) These results demonstrate the applicability of RI-LVL to systems with d- and f-electrons.

3.9 SCALING ANALYSIS

In addition to testing the accuracy of the RI-LVL, I also investigated the scaling with system size in terms of both memory and computational time in HF calculations. For this purpose fully extended oligo-Alanine chains have been chosen as benchmark systems. One of them,


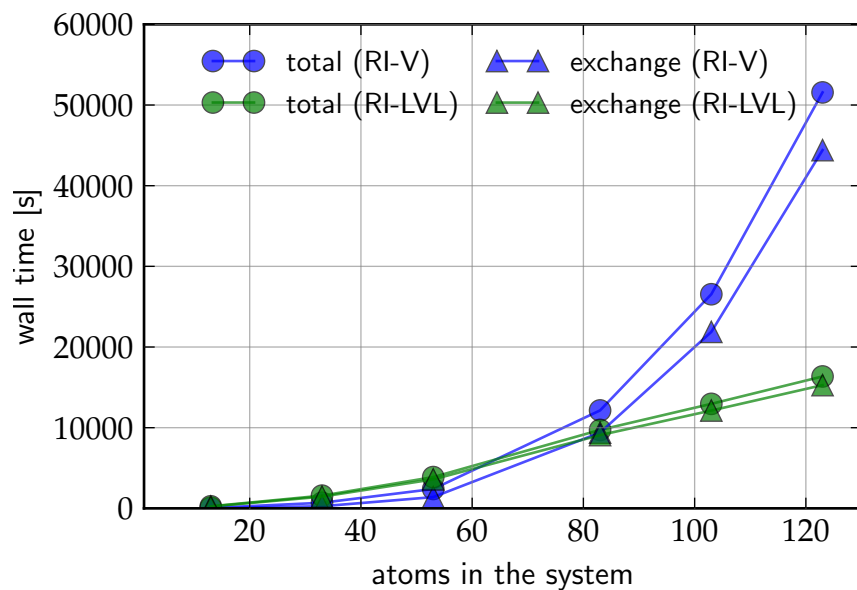
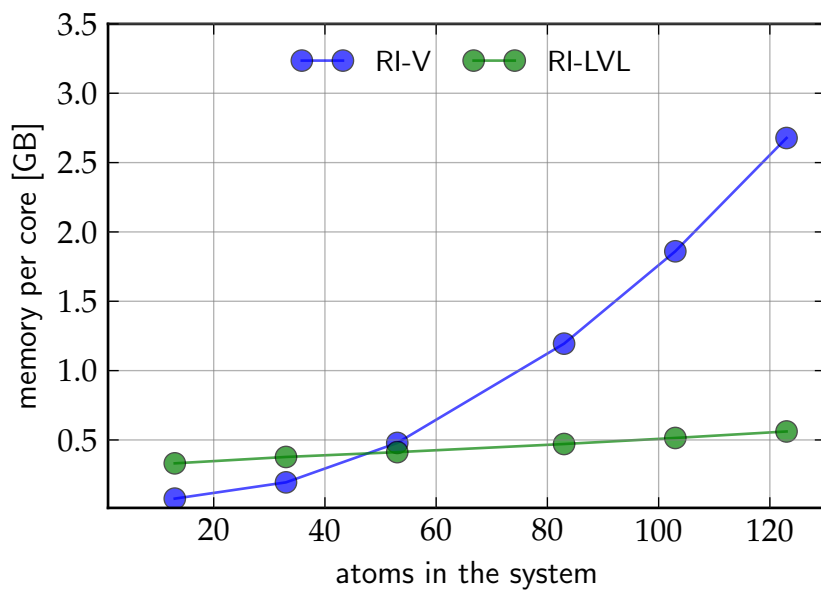


Figure 3.9: Alanine-8 system used for the computational scaling analysis.

Alanine-8, is depicted in Figure 3.9. For the comparison between RI-V and RI-LVL tier2 basis sets with “tight” integration settings have been used. In the case of RI-LVL an extra g -type function with $z = 6$ was added to the OBS+ to accommodate the results from the accuracy study. All calculations used 180 cores of an Infiniband-connected Intel cluster. The RI-LVL implementation for HF in FHI-aims is documented in detail in Ref. [78] and makes use of linear-scaling density-matrix based screening techniques from quantum chemistry. [124, 125] The total computational time requirements for the self-consistent HF calculations are shown in figure 3.10a as well as the evaluation of the exact exchange term, which is the rate-determining step in these calculations. RI-LVL is more expensive in small systems compared to the dense matrix algebra used in the RI-V. This is a consequence of both the larger auxiliary basis set required for RI-LVL and the overhead caused by the more complex implementation to exploit the sparsity. However, RI-LVL exhibits a lower scaling with system size than RI-V and breaks even at about 70 atoms. The RI-V implementation exhibits a cubic scaling for the exact exchange term, while the sparsity and screening used in RI-LVL give it an almost linear scaling with system size. The remaining computational time not spent for the exact exchange evaluation is dominated by the generation of the RI coefficients. As it can be seen in figure 3.10a, this step requires a notable amount of time in the RI-V calculations. In RI-LVL however, this expense remains negligible even for large systems due to the *a priori* enforced locality constraints in the RI coefficients. This difference is also visible in figure 3.10b which shows the required memory as a function of the system size. RI-V again shows the superior performance for small systems, but the memory consumption quickly rises and RI-LVL becomes more efficient already at about 50 atoms. It should be noted that NAOs in FHI-aims have a strictly limited radial extent and thus any basis function pair product $\varphi_s \varphi_t$ with host atoms further away than the sum of their radial extents is exactly zero. This locality can be used to effectively reduce the memory requirements in RI-V calculations, especially for linear systems. Taking Alanine-8 as



(a) total computational time and time spent for the exchange matrix evaluation



(b) memory requirement per core

Figure 3.10: Computational scaling and memory consumption for self-consistent HF calculations using the RI-V and RI-LVL algorithms for fully extended oligo-alanine chains. All calculations have been performed with 180 cores and tier2 basis sets. RI-LVL additionally used a g -type function for the OBS+ to provide accurate results.

example, about 50% of the basis function pairs have no overlap. However, the benefit is smaller in compact three-dimensional compounds with similar system size. The memory savings of RI-LVL are therefore even more pronounced than in this scaling example. The total energy differences between the RI-V and RI-LVL calculations are below 0.02 meV/atom and do not increase with system-size.

3.10 CONCLUSIONS

In this chapter I presented the RI-LVL, a localized RI scheme. I demonstrated that the accuracy of the RI-LVL method can be systematically controlled and the level of accuracy provided by the commonly used RI-V can be obtained for elements all across the periodic table. In combination with algorithms that make use of the sparsity in the coefficients, RI-LVL can dramatically reduce the memory requirements of calculations that require the evaluation of ERIs. And combined with integral screening algorithms, RI-LVL can also provide a superior computational performance with a reduced scaling.

The applicability of RI-LVL to HF-based calculations in periodic systems has been demonstrated by Sergey Levchenko in Ref [78]. The reported accuracy and performance match well with the results for cluster systems reported here. RI-LVL can thus be applied to molecular and periodic systems alike.

LT-MP₂ FOR MOLECULAR AND PERIODIC SYSTEMS

In the previous chapter I introduced the concept of the RI-LVL and demonstrated how it can be applied to Hartree-Fock (HF) theory and methods based on it to significantly reduce both the memory and computational costs without compromising the accuracy. I also showed that the atomic Electron Repulsion Integrals (ERIs) expanded with RI-LVL are accurate enough for advanced correlation methods like the 2nd Order Møller-Plesset Perturbation Theory (MP2) and the Random-Phase Approximation (RPA), which depend on the unoccupied states as well as the occupied ones.

A variety of reduced scaling approaches has been proposed for advanced correlation methods. (see .e.g. section 2.5 for an overview of the available lower-scaling MP2 variants) A key concept in these approaches is to take the “near sightedness” of the electronic correlation explicitly into account and rewrite the formalism in terms of spatially localized molecular orbitals or atomic orbitals. Combined with criteria for pre-selecting the significant contributions, present local correlation methods are tremendously successful in reducing the computational scaling. However, the manipulation of the atomic and transformed ERIs remains a big challenge, even when Resolution of Identity (RI) strategies are used. A disk-storage approach is often used by quantum chemistry codes (e.g. [36, 42]) because it is easy to implement and one does not need to make additional approximations beyond the well established RI-V. However, this approach prevents any efficient parallelization because the file-system quickly becomes the bottleneck. On the other hand, keeping the RI coefficients in memory while exploiting locality gives rise to much more complicated algorithms. These are usually a lot more difficult to parallelize efficiently than the canonical formulations. Therefore most of the demonstrations of local correlation methods feature only a limited parallelization, e.g. by using thread-parallel linear algebra libraries in serial code.

In this second part of my project I combined a generalized version of the RI-LVL from the previous chapter with the Laplace-Transformation (LT) technique. The LT can be used to recast the eigenstates appearing in the correlation energy expressions of e.g. MP2 and RPA into a set of localized functions, which is the prerequisite for using locality exploits such as RI-LVL. Since these new functions are not as strongly localized as the atomic basis functions discussed in the previous chapter, a generalization of the RI-LVL is needed to provide the desired level of accuracy. The new lower-scaling technique that arises from the combination of RI-LVL and the LT is a powerful tool which can be applied to molecular and periodic systems alike. Although I will introduce these concepts explicitly at the example of the MP2, it

is important to realize that this work is not restricted to the MP2. In chapter 7 I demonstrate in detail how the concepts presented here can be applied to the RPA as well. Applying this new lower-scaling technique to the MP2, I obtained a lower-scaling and well-parallelized MP2 which reduces the required memory in an *a priori* predictable way.

4.1 LT-MP2 IN A NAO FRAMEWORK

As already discussed in detail in section 2.5, the Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP2) is based on the idea to rewrite the denominator in the correlation energy formula to avoid the spatially non-decaying coupling between the molecular orbitals. After this transformation step, Equation 4.1 is the starting point for all LT-MP2 implementations. $c_{i,s}$ denotes the expansion coefficients of the molecular orbital ψ_i into the set of orbital basis functions φ_s .

$$\mathcal{E}_{MP2} = - \sum_q w_q \sum_{s,t,u,v} \sum_{s',t',u',v'} (s't'|u'v') X_{s',s}^q X_{u',u}^q Y_{t',t}^q Y_{v',v}^q \left[2(st|uv) - (sv|ut) \right] \quad (4.1)$$

$$X_{s',s}^q = \sum_i c_{i,s'} c_{i,s} e^{(\epsilon_i - \epsilon_F) t_q} \quad (4.2)$$

$$Y_{t',t}^q = \sum_a c_{a,t'} c_{a,t} e^{-(\epsilon_a - \epsilon_F) t_q} \quad (4.3)$$

$$\epsilon_F = \frac{\epsilon_{\text{HOMO}} + \epsilon_{\text{LUMO}}}{2} \quad (4.4)$$

A key feature of the LT-MP2 is the use of the pseudo-density matrices $X_{s',s}^q$ and $Y_{t',t}^q$. Both feature a similar decay behavior as the real density matrix. As an illustration, Figure 4.1 shows the rapid decay of the pseudo-density matrix in a water cluster with 50 molecules. Note that for $t_q = 0$, $X_{s',s}^q$ is equivalent to the common density matrix. ϵ_F is a “constructive unity” added to ensure that the exponents will always have negative arguments and thus provide a smoother decay of the pseudo-density matrices. It does not arise from physical considerations (see the derivation in section 2.5 without it), thus the quality of the HOMO-LUMO gap provided by the preceding HF is of no concern.

Due to this decay behavior many terms in the summation will yield insignificant contributions and their evaluation can be avoided with a suitable screening criterion. (discussed in detail in section 4.7) To evaluate this expression numerically, we need to rewrite the equation in terms of transformed basis functions $\varphi_{\bar{s}}$ and $\varphi_{\underline{t}}$ as defined in Equation 4.5.

$$\varphi_{\bar{s}} = \sum_{s'} X_{s,s'}^q \varphi_{s'} \quad \varphi_{\underline{t}} = \sum_{t'} Y_{t,t'}^q \varphi_{t'} \quad (4.5)$$

There are several ways how Equation 4.1 can be reformulated with these transformed functions. One possible choice are the half trans-




Figure 4.1: Decay feature of the occupied orbital pseudo-density matrix $X_{s't',s}^q$ in a cluster with 50 water molecules. Shown are the elements of the oxygen's d -type function with magnetic quantum number $m = 0$ as a function of the inter-atomic distance.

formed ERIs from Equation 4.6 which have been used in the past by the Ochsenfeld group. [41]

$$E = - \sum_q w_q \sum_{s,t,u,v} 2 (\bar{s}t|uv)^q (st|\bar{u}v)^q - (\bar{s}t|\bar{u}v)^q (sv|ut)^q \quad (4.6)$$

In this formulation no ERI needs more than two transformations, which makes it suitable for an on-the-fly evaluation of the ERIs without RI strategies. Due to the decay properties of the pseudo-density matrices, the number of atomic integrals needed to form a half transformed integral in Equation 4.6 quickly becomes independent of system size. Combined with the screening criterion which preselects the significant function tuples, this formulation allows for a LT-MP2 implementation with a low computational scaling. On the other hand, half transformed integrals are not suited for an implementation which uses a RI to evaluate the ERIs. Computing the transformed integrals on the fly will become very inefficient due to the communication overhead as soon as we need to distribute the RI coefficients among several processes. To avoid this communication bottleneck, the transformed ERIs need to be evaluated with a RI as well. As a consequence, we need a full set of RI coefficients for every transformation pattern appearing in the correlation energy formula. Equation 4.6 would thus require us to store three sets of transformed RI coefficients.

An alternative is the usage of the fully transformed integrals shown in Equation 4.7, as they have been used in the original publications. [15, 81]

$$E = - \sum_q w_q \sum_{s,t,u,v} (\bar{s}t|\bar{u}v)^q \left[2(st|uv) - (sv|ut) \right] \quad (4.7)$$

This formulation has the advantage that only two sets of RI coefficients are needed: one for the atomic ERIs and one to represent the $\varphi_{\bar{s}}\varphi_{\underline{t}}$ pairs. The atomic integrals are exactly the same as the ones already encountered in the HF theory, which has been discussed extensively in the previous chapter. RI-LVL is thus the optimal choice to handle these integrals.

Unfortunately, the auxiliary basis set (ABS) used in RI-LVL is too restrictive to accurately represent the transformed function pairs $\varphi_{\bar{s}}\varphi_{\underline{t}}$, which contain significant contributions from nearby atoms via the pseudo-density matrices $X_{s',s}^q$ and $Y_{t',t}^q$. (see Equation 4.5) To address this issue, I generalized the concept of the RI-LVL expansion to provide the necessary accuracy for the transformed integrals $(\bar{s}\underline{t}|\underline{u}\underline{v})$, which will be presented in section 4.4.

4.2 THE LT-MP2 FOR PERIODIC SYSTEMS

Starting from the closed-shell expression for the MP2 correlation energy in periodic systems (Equation 2.89), the LT-MP2 in terms of crystalline orbitals can be derived just the same way as the molecular case. (see derivation in section 2.5) Indexes for the integration point q are omitted from here on for the simplicity of notation.

$$\mathcal{E}_{MP2} = - \sum_{i,j} \sum_{\mathbf{k}_1,\mathbf{k}_2} \sum_{a,b} \sum_{\mathbf{q}_1,\mathbf{q}_2} \left[\frac{2 (\psi_{i,\mathbf{k}_1} \psi_{a,\mathbf{q}_1} | \psi_{j,\mathbf{k}_2} \psi_{b,\mathbf{q}_2})^* (\psi_{i,\mathbf{k}_1} \psi_{a,\mathbf{q}_1} | \psi_{j,\mathbf{k}_2} \psi_{b,\mathbf{q}_2})}{\epsilon_i^{\mathbf{k}_1} + \epsilon_j^{\mathbf{k}_2} - \epsilon_a^{\mathbf{q}_1} - \epsilon_b^{\mathbf{q}_2}} - \frac{\Re \left[(\psi_{i,\mathbf{k}_1} \psi_{a,\mathbf{q}_1} | \psi_{j,\mathbf{k}_2} \psi_{b,\mathbf{q}_2})^* (\psi_{i,\mathbf{k}_1} \psi_{b,\mathbf{q}_2} | \psi_{j,\mathbf{k}_2} \psi_{a,\mathbf{q}_1}) \right]}{\epsilon_i^{\mathbf{k}_1} + \epsilon_j^{\mathbf{k}_2} - \epsilon_a^{\mathbf{q}_1} - \epsilon_b^{\mathbf{q}_2}} \right] \quad (4.8)$$

$$= - \sum_q w_q \sum_{\mathbf{k}_1,\mathbf{k}_2} \sum_{\mathbf{q}_1,\mathbf{q}_2} \sum_{s,t,u,v} \sum_{s',t',u',v'} \left(s'_{\mathbf{k}_1} t'_{\mathbf{q}_1} | u'_{\mathbf{k}_2} v'_{\mathbf{q}_2} \right) X_{s',s}^{\mathbf{k}_1} X_{u',u}^{\mathbf{k}_2} Y_{t',t}^{\mathbf{q}_1} Y_{v',v}^{\mathbf{q}_2} \left[2 (s_{\mathbf{k}_1} t_{\mathbf{q}_1} | u_{\mathbf{k}_2} v_{\mathbf{q}_2}) - (s_{\mathbf{k}_1} v_{\mathbf{q}_2} | u_{\mathbf{k}_2} t_{\mathbf{q}_1}) \right] \quad (4.9)$$

$$X_{s',s}^{\mathbf{k}_1} = \sum_i c_{i(\mathbf{k}_1),s'}^* c_{i(\mathbf{k}_1),s} e^{(\epsilon_i^{\mathbf{k}_1} - \epsilon_F) t_q} \quad (4.10)$$

$$Y_{t',t}^{\mathbf{q}_1} = \sum_a c_{a(\mathbf{q}_1),t'}^* c_{a(\mathbf{q}_1),t} e^{-(\epsilon_a^{\mathbf{q}_1} - \epsilon_F) t_q} \quad (4.11)$$

ϵ_F is now obtained from the valence band maximum and conduction band minimum of the band structure at the sampled k-points. As before, ϵ_F does not have a physical meaning in this equation, but is added for numerical reasons only. The translational symmetry can be used to eliminate one of the sums over k-points as in the canonical formulation. The evaluation of this equation for each k-point tuple is then similar to the cluster case which has been already discussed in the previous section.

As explained in section 2.8, one can also sample the Born von Karman (BvK) supercell at the Γ -point to obtain the same information as when sampling the primitive unit cell with the corresponding k-grid. As illustrated in Figure 4.2 the primitive unit cell sampled with a k-

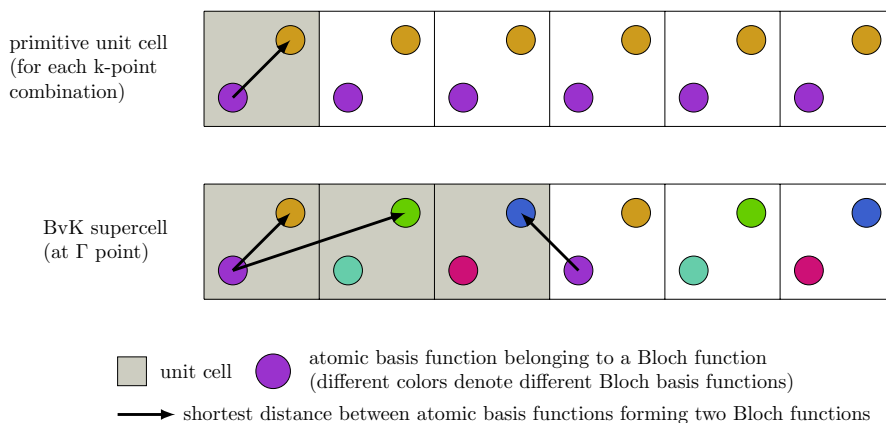


Figure 4.2: Comparison between a primitive unit-cell and the BvK supercell for the screening in the LT-MP2. In case of the primitive cell, the distances between the atomic basis functions used to assemble a Bloch basis function are nearby in real-space for any k -point combination. Since the k -points only add phase-factors, all ERIs of Bloch basis functions will have significant contributions. In the BvK supercell on the other hand, the atomic functions contained in a Bloch function are further apart. Therefore the ERIs in this representation offer a lot more sparsity that can be exploited by integral screening techniques.

grid offers little or no sparsity in the to be evaluated ERIs. Due to the limited size of the unit cells, any two Bloch functions will contain atomic basis functions that are located nearby in real space. Since the k -points only add phase factors to these atomic basis functions, all ERIs of Bloch functions will contain some significant ERIs of atomic basis functions and can thus not be neglected in the evaluation. The ERIs tensor is thus not sparse at all in this representation and the primitive unit cell with a k -grid is thus not an appealing choice for methods like the LT-MP2 which rely heavily on screening.

The BvK supercell on the other hand has much more potential for sparsity exploits. As the figure illustrates, the periodic repetitions of the atomic basis functions in the primitive cells constituting the BvK supercell are now associated with different Bloch functions. The distances between the atomic basis functions forming the Bloch basis functions thus increase and with increasing size of the supercell more and more Bloch function tuples will yield ERIs with negligible contributions.

Since the performance of the LT-MP2 is strongly dependent on the efficiency of the integral screening, I decided to use the supercell approach for the periodic version of my LT-MP2 implementation. For a calculation sampling only the Γ point ($\mathbf{k}_1 = 0$) Equation 4.9 becomes very similar to the molecular formula given in Equation 4.1. The only difference is that the atomic basis functions have been replaced by their crystalline counterparts at $\mathbf{k}_1 = 0$. The Bloch functions are real-valued for the Γ point and thus the same implementation can be used for both molecular and periodic systems. The only difference are the

RI coefficients, eigenvectors and eigenvalues used as input. In the FHI-aims code, the Coulomb matrix for beyond DFT methods in periodic systems is usually generated with the truncated Coulomb operator proposed by Spencer and Alavi [126] for the Γ point to treat the singularity, but with the normal Coulomb operator for all other k-points. [127]

It should be noted that multiple k-points can also be used explicitly in a LT-MP2 framework as demonstrated by Scuseria and coworkers. [81] To exploit sparsity and avoid the cubic scaling with k-points, it is however necessary to expand the Bloch basis functions explicitly in terms of the atomic basis functions. (see Equation 2.82) After this expansion, all k-dependencies can be adsorbed into the pseudo-density matrices at little cost. However, instead of summing over the finite number of Bloch Basis functions in the unit cell, one now has to handle infinite summations over unit cells. These must be truncated carefully in such a way that the long-range Coulomb interactions of the solid are not perturbed.

4.3 THE INTEGRATION POINTS AND WEIGHTS

The first choice one has to make for a LT-MP2 implementation is the set of the integration points and weights used to approximate the denominator of the Kohn-Sham (KS) or HF eigenvalues.

$$\frac{1}{\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j} = \int_0^\infty e^{-s(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)} ds \quad (4.12)$$

$$\approx \sum_q^{N_q} w_q e^{-t_q(\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j)} \quad (4.13)$$

Standard quadrature schemes like Gauss-Legendre are not sufficient to provide accurate results with a small number of integration points. (see e.g. the figures in [128]) Häser already realized this in his early works on the LT-MP2 and suggested an alternative scheme which chooses the weights and grid points such that they minimize the summed squares of differences between the true denominator and its approximation for all orbital-tuples.[40]

A better alternative is the minimax approximation. [128] In this method the integration points and weights are determined by minimizing the Chebyshev-norm of the fitting error $\eta(x)$ of the function $\frac{1}{x}$ in a given interval.

$$\frac{1}{x} = \int_0^\infty e^{-xs} ds \approx \sum_q^{N_q} w_q e^{-t_q x} \quad (4.14)$$

$$\eta(x) = \frac{1}{x} - \sum_q^{N_q} w_q e^{-t_q x} \quad (4.15)$$

The Chebyshev-norm $\|\cdot\|_{C,\mathcal{I}}$ is the maximum absolute value of a function in its definition interval \mathcal{I} .

$$\|\eta(x)\|_{C,\mathcal{I}} = \max_{x \in \mathcal{I}} |\eta(x)| \quad (4.16)$$

$$= \max_{x \in \mathcal{I}} \left| \frac{1}{x} - \sum_q^{N_q} w_q e^{-t_q x} \right| \quad (4.17)$$

The accuracy of the fit is only relevant in the range spanned by the energy denominators appearing in the MP2, i.e. it suffices to minimize the Chebyshev-norm in the interval spanned by the eigenvalue tuples.

$$\mathcal{I} = [2(\epsilon_{\text{LUMO}} - \epsilon_{\text{HOMO}}), 2(\epsilon_{\text{max}} - \epsilon_{\text{min}})] \quad (4.18)$$

A small Chebyshev-norm of the fitting error $\eta(x)$ in this interval therefore guarantees that the energy denominator will be accurately represented for all function tuples without having to consider them explicitly in the optimization procedure. It has been shown that the minimax approximation provides a considerably higher accuracy than other quadrature schemes for the same number of grid points and that already five grid points are usually sufficient to reach micro-Hartree accuracies for the valence-shell correlated MP2. [128]

For my project I decided to use the minimax scheme. The source-code for this method has been kindly provided by Prof. Seiichiro Tenno.[128]

4.4 THE GENERALIZED RI-LVL FOR THE TRANSFORMED ERIS

As outlined in the previous sections, the key challenge for a RI based LT-MP2 implementation is the choice of the RI strategy for the evaluation of the fully transformed integrals. A general definition of the RI expansion of the transformed ERIs can be written as

$$(\bar{s}\underline{t}|\bar{u}\underline{v}) = \sum_{s',t',u',v'} X_{s',s} Y_{t',t} (s't'|u'v') X_{u',u} Y_{v',v} \quad (4.19)$$

$$\approx \sum_{\mu,\lambda} M_{\bar{s},\underline{t}}^\mu V_{\mu,\lambda} M_{\bar{u},\underline{v}}^\lambda \quad (4.20)$$

where the $M_{\bar{s},\underline{t}}^\mu$ are the RI coefficients which approximate the transformed densities as shown in Equation 4.21.

$$q_{\bar{s},\underline{t}}(\mathbf{r}_1) = \sum_{s',t'} X_{s',s} Y_{t',t} \varphi_{s'}(\mathbf{r}_1) \varphi_{t'}(\mathbf{r}_1) \approx \sum_{\mu}^{ABS} M_{\bar{s},\underline{t}}^\mu P_{\mu}(\mathbf{r}_1) = \tilde{q}_{\bar{s},\underline{t}}(\mathbf{r}_1) \quad (4.21)$$

Starting from RI-LVL based HF or KS calculations, the simplest choice for the $M_{\bar{s},\underline{t}}^\mu$ is the RI-LVL_{full} introduced in section 3.5. In this RI technique we simply recast our sparse coefficients into a dense tensor with all auxiliary basis functions in the system. Inserting our

RI coefficients $C_{s',t'}^\mu$ for the atomic integrals into Equation 4.21, we obtain:

$$Q_{\bar{s},\underline{t}}(\mathbf{r}_1) \approx \sum_{s',t'} \sum_{\mu} X_{s',s} Y_{t',t} C_{s',t'}^\mu P_{\mu}(\mathbf{r}_1) \quad (4.22)$$

$$\Rightarrow M_{\bar{s},\underline{t}}^\mu = \sum_{s',t'} X_{s',s} Y_{t',t} C_{s',t'}^\mu \quad (4.23)$$

As one can see from Equation 4.23, the RI-LVL_{full}-based RI coefficients for the transformed integrals follow directly from the atomic ones after two tensor-matrix multiplications. This formula holds for atomic RI coefficients calculated from both RI-V and RI-LVL alike. By multiplying the $M_{\bar{s},\underline{t}}^\mu$ with the square-root of the Coulomb matrix, we can then reduce the the cost of a single ERI evaluation to a simple dot-product over auxiliary basis functions as it is commonly done in RI-V.

A big disadvantage of this approach is however the memory requirement. The memory footprint of the RI coefficient tensor scales with $\mathcal{O}(N_{\text{bas}}^2 N_{\text{aux}})$, while the RI coefficients for the molecular ERIs in the canonical MP2 require $\mathcal{O}(N_{\text{occ}} N_{\text{virt}} N_{\text{aux}})$. Although the formal scaling is cubic in both cases, the prefactor is much smaller in the canonical implementation. Furthermore, the multiplication with the Coulomb matrix square-root corresponds to a transformation into a new set of delocalized, but orthogonal auxiliary basis functions, which destroys any locality one could hope to exploit.

A LT-MP2 implementation using RI-LVL_{full} to represent the transformed ERIs thus has no practical value in everyday applications, but it is a useful benchmark tool to clearly separate the errors introduced by the Laplace-Transformation and the integral screening from those caused by the generalized RI-LVL used for the transformed ERIs. (introduced in the following)

For practical applications it is important to reduce the required memory considerably and to do so in predictable way. Both goals can be achieved by restricting the ABS used to expand a transformed function pair *a priori* as it was done in RI-LVL for the atomic ERIs.

Given the successful application of the RI-V optimization criterion to the RI-LVL in the previous chapter, I decided to use it as the starting point for the derivation of the transformed local RI coefficients. Similar to the procedure in section 3.2, I choose the coefficients such that they minimize the Coulomb self-interaction of the fitting residual $\delta Q_{\bar{s},\underline{t}} = Q_{\bar{s},\underline{t}} - \tilde{Q}_{\bar{s},\underline{t}}$. At this point I do not make any assumptions about the choice of the local ABS used for the transformed pair $\varphi_{\bar{s}}\varphi_{\underline{t}}$ and take the notation $\mathbb{P}(\bar{s},\underline{t})$ as a general local ABS for the given function pair.

$$\delta q_{\bar{s},\underline{t}}(\mathbf{r}_1) = \sum_{s',t'} X_{s',s} Y_{t',t} \varphi_{s'}(\mathbf{r}_1) \varphi_{t'}(\mathbf{r}_1) - \sum_{\mu}^{\mathbb{P}(\bar{s},\underline{t})} M_{\bar{s},\underline{t}}^{\mu} P_{\mu} \quad (4.24)$$

$$\begin{aligned} (\delta q_{\bar{s},\underline{t}} | \delta q_{\bar{s},\underline{t}}) &= \sum_{\mu,\lambda}^{\mathbb{P}(\bar{s},\underline{t})} M_{\bar{s},\underline{t}}^{\mu} M_{\bar{s},\underline{t}}^{\lambda} (\mu | \lambda) - 2 \sum_{\mu}^{\mathbb{P}(\bar{s},\underline{t})} \sum_{s',t'} X_{s',s} Y_{t',t} M_{\bar{s},\underline{t}}^{\mu} (s' t' | \mu) \\ &+ \sum_{s',t',u',v'} X_{s',s} Y_{t',t} (s' t' | u' v') X_{u',s} Y_{v',t} \end{aligned} \quad (4.25)$$

To minimize this term, I take its derivative with respect to the transformed RI coefficients $M_{\bar{s},\underline{t}}^{\mu}$.

$$\begin{aligned} \frac{\partial (\delta q_{\bar{s},\underline{t}} | \delta q_{\bar{s},\underline{t}})}{\partial M_{\bar{s},\underline{t}}^{\mu}} &\stackrel{!}{=} 0 \quad \forall \mu \in \mathbb{P}(\bar{s},\underline{t}) \\ \Rightarrow 0 &= \sum_{\lambda}^{\mathbb{P}(\bar{s},\underline{t})} M_{\bar{s},\underline{t}}^{\lambda} (\mu | \lambda) - \sum_{s',t'} X_{s',s} Y_{t',t} (s' t' | \mu) \end{aligned} \quad (4.26)$$

$$\Rightarrow \sum_{\lambda}^{\mathbb{P}(\bar{s},\underline{t})} M_{\bar{s},\underline{t}}^{\lambda} (\mu | \lambda) = \sum_{s',t'} X_{s',s} Y_{t',t} (s' t' | \mu) \quad \forall \mu \in \mathbb{P}(\bar{s},\underline{t}) \quad (4.27)$$

Equation 4.27 looks very similar to the case of the atomic ERIs (see Equation 3.12) at a first glance, but has the very important difference that each transformed function pair $\varphi_{\bar{s}}\varphi_{\underline{t}}$ requires all atomic function pairs for the generation of the right-hand side of the equation system. Computing all the three-center integrals $(s' t' | \mu)$ is however a significant effort, in particular if they would need to be recomputed for each integration point of the LT. Storing them on the other hand would require a cubic scaling amount of memory just as in a RI-V-based canonical MP2.

It can be seen from Equation 4.25 that the problematic three-center integrals are introduced by the cross-term between the exact transformed density and its RI approximation. The exact transformed function pair (Equation 4.5) is however just a weighted sum of orbital basis pairs. As known from the results shown in chapter 3, the RI implementations in FHI-aims are capable of providing very high accuracies for the atomic ERIs and thus also the orbital basis pair densities forming them. If one replaces the true orbital pairs in Equation 4.25 with an accurately converged RI-LVL representation, one obtains:

$$\begin{aligned} \delta q_{\bar{s},\underline{t}}(\mathbf{r}_1) &\approx \sum_{s',t'}^{\mathcal{P}(S',T')} \sum_{\kappa} X_{s',s} Y_{t',t} C_{s',t'}^{\kappa} P_{\kappa}(\mathbf{r}_1) - \sum_{\mu}^{\mathbb{P}(\bar{s},\underline{t})} M_{\bar{s},\underline{t}}^{\mu} P_{\mu}(\mathbf{r}_1) \quad (4.28) \\ (\delta q_{\bar{s},\underline{t}} | \delta q_{\bar{s},\underline{t}}) &= \sum_{\mu,\lambda}^{\mathbb{P}(\bar{s},\underline{t})} M_{\bar{s},\underline{t}}^{\mu} M_{\bar{s},\underline{t}}^{\lambda} (\mu | \lambda) \\ &- 2 \sum_{\mu}^{\mathbb{P}(\bar{s},\underline{t})} \sum_{s',t'} \sum_{\kappa}^{\mathcal{P}(S',T')} X_{s',s} Y_{t',t} M_{\bar{s},\underline{t}}^{\mu} C_{s',t'}^{\kappa} (\kappa | \mu) \\ &+ \sum_{s',t',u',v'} \sum_{\mu}^{\mathcal{P}(S',T')} \sum_{\kappa}^{\mathbb{P}(u',v')} X_{s',s} Y_{t',t} C_{s',t'}^{\mu} (\mu | \kappa) C_{u',v'}^{\kappa} X_{u',s} Y_{v',t} \end{aligned} \quad (4.29)$$

Here $\mathcal{P}(S', T')$ denotes the ABS for RI-LVL as introduced in the previous chapter, i.e. it only contains the auxiliary functions centered at the same two atoms as $\varphi_{s'}$ and $\varphi_{t'}$. Taking the derivative of Equation 4.29 then yields:

$$\frac{\partial (\delta Q_{\bar{s}, \underline{t}} | \delta Q_{\bar{s}, \underline{t}})}{\partial M_{\bar{s}, \underline{t}}^\mu} \stackrel{!}{=} 0 \quad \forall \mu \in \mathbb{P}(\bar{s}, \underline{t}) \quad (4.30)$$

$$\Rightarrow 0 = \sum_{\lambda}^{\mathbb{P}(\bar{s}, \underline{t})} M_{\bar{s}, \underline{t}}^\lambda (\mu | \lambda) - \sum_{s', t'}^{\mathcal{P}(S', T')} \sum_{\kappa} X_{s', s} Y_{t', t} C_{s', t'}^\kappa (\kappa | \mu) \quad (4.31)$$

$$\Rightarrow \sum_{\lambda}^{\mathbb{P}(\bar{s}, \underline{t})} M_{\bar{s}, \underline{t}}^\lambda (\mu | \lambda) = \sum_{s', t'}^{\mathcal{P}(S', T')} \sum_{\kappa} X_{s', s} Y_{t', t} C_{s', t'}^\kappa (\kappa | \mu) \quad \forall \mu \in \mathbb{P}(\bar{s}, \underline{t}) \quad (4.32)$$

Equation 4.32 no longer requires the evaluation of any additional three-center integrals. Only an accurate RI-LVL representation of the basis function pair densities and the Coulomb matrix are needed, which are already provided by the RI-LVL needed to evaluate the atomic ERIs appearing in the LT-MP2.

For the supercell approach used in periodic systems, the atomic basis and auxiliary basis functions in the preceding discussion need to be replaced by their Bloch counterparts. The derivation then follows the same steps as in the molecular case and is thus not shown here explicitly.

4.5 SELECTING THE LOCAL ABS FOR THE TRANSFORMED ERIS

To use Equation 4.32 in practice, it is still necessary to choose a local set of auxiliary basis functions $\mathbb{P}(\bar{s}, \underline{t})$ for the transformed function pairs. The self-evident choice is to use same ABS as in the RI-LVL, i.e. define $\varphi_{\bar{s}}$ and $\varphi_{\underline{t}}$ to be centered at the same atoms as φ_s and φ_t and then use only the auxiliary basis functions on these two atoms. Equation 4.5 however reveals that each transformed function is a superposition of basis functions centered at different atoms, and thus a more delocalized quantity than a Numeric Atom-Centered Orbital (NAO) basis function. It is therefore not surprising that the RI-LVL is not sufficient to model the transformed integrals accurately, even if a very large enhanced orbital basis set (OBS+) is used to generate the ABS.

To accurately model the transformed integrals it is necessary to include auxiliary basis functions from other atoms as well. The natural extension would be to include not just the auxiliary basis functions on atoms S and T to model the pair $\varphi_{\bar{s}}\varphi_{\underline{t}}$, but also those on the neighboring atoms within a chosen distance. Inspecting Equation 4.32 however reveals that this approach becomes computationally quite inefficient. In general, this approach would yield a different local ABS for each atom pair in the system, i.e. there is one linear equation system (LES) to solve per atom pair. Solvers for this problem class usually scale with the third power of the size of the matrix in the system. Since additional functions from nearby atoms are now included, the ma-

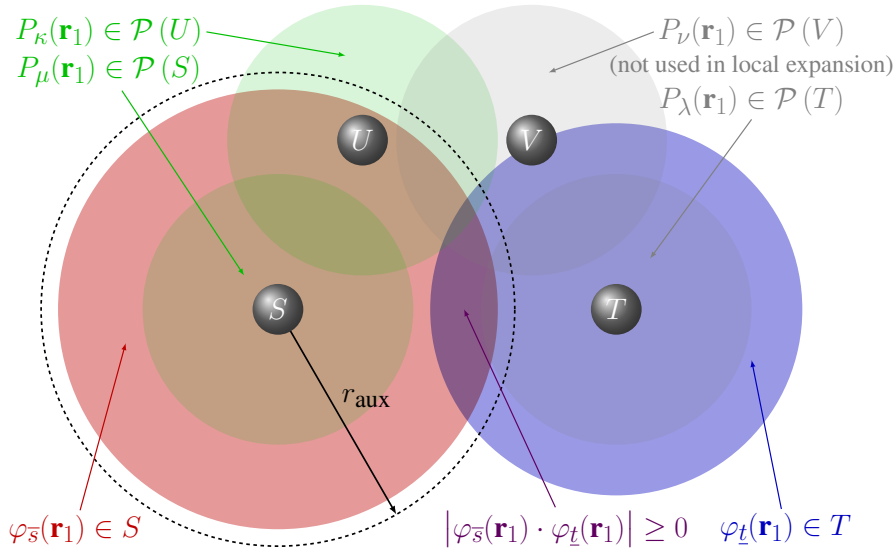


Figure 4.3: A schematic representation of the local auxiliary basis set used for the transformed function pairs. Only those auxiliary basis functions hosted on atoms within a distance r_{aux} from the host atom of the “occupied” transformed function $\varphi_{\bar{s}}$ are included in the local ABS of the pair $\varphi_{\bar{s}}\varphi_{\underline{t}}$. The on-site auxiliary basis functions on the atom hosting the “virtual” transformed function $\varphi_{\underline{t}}$ do not receive any special treatment.

trices are much larger than their counterparts in the RI-LVL case for the atomic ERIs. This results in a considerably larger computational effort per pair. Furthermore, the global ABS in FHI-aims is not guaranteed to be free of linear dependencies, orthogonality is only ensured among functions centered at the same atom. Using efficient solvers for symmetric LES is therefore prone to numerical instabilities. Instead, the equations must be solved using a Singular Value Decomposition (SVD) or Rank-Revealing QR-Decomposition (RR-QR), which are able to handle such possibly rank-deficient problems. Alternatively, one can also use an (symmetric) Eigenvalue Decomposition (EVD) and remove all almost zero eigenvalues to solve the LES. These methods do not change the formal scaling, but they have a considerably larger prefactor than the symmetric solvers for full-rank LES.

To reduce the computational effort, it is therefore necessary to choose a different local ABS $\mathbb{P}(\bar{s}, \underline{t})$ for the transformed RI coefficients which reduces the scaling of the $\mathcal{O}(N_{\text{atom pairs}})$ distinct matrices we need to factorize. Instead of choosing all atoms within a distance r_{aux} of either atom, I decided to only include those in the vicinity of the atom belonging to the “occupied” transformed function $\varphi_{\bar{s}}$ in the local ABS as shown in Figure 4.3. In periodic calculations, the distance criterion also checks the mirror atoms in neighboring unit cells to determine the smallest distance between two atoms. This choice of the local ABS reduces the number of distinct matrices we need to factorize to only $\mathcal{O}(N_{\text{atoms}})$. These factorizations can either be recomputed at every in-

tegration grid point or precomputed once because the local Coulomb matrices in Equation 4.32 do not depend on the integration grid point. The accuracy of this choice will be demonstrated in chapter 5.

To summarize, the transformed RI coefficients $M_{\bar{s},\underline{t}}^\lambda$ are determined by solving the LES with $\mathbb{P}(\bar{s},\underline{t})$ including all auxiliary basis functions with host atoms at most r_{aux} away from atom S .

$$\sum_{\lambda} \mathbb{P}(\bar{s},\underline{t}) M_{\bar{s},\underline{t}}^\lambda(\mu|\lambda) = \sum_{\kappa} \mathbb{P}(s,t) \sum_{s',t'} X_{s',s} Y_{t',t} C_{s',t'}^\kappa(\kappa|\mu) \quad \forall \mu \in \mathbb{P}(\bar{s},\underline{t})$$

It should be noted that this choice of the local auxiliary basis set is similar to an older work by Head-Gordon for exchange integrals in HF theory. [113, 114]

4.6 FORMAL SCALING OF THE INTEGRAL TRANSFORMATION

Another very important aspect is the scaling of the transformation step with system size. A straight-forward implementation would scale with the fourth power of the system size and thus quickly become the dominating step of the calculation and nullify the gains from the integral screening in the subsequent evaluation step.

The previously discussed choice for the local ABS however can be used to limit the scaling to $\mathcal{O}(N^3)$ at most. During the transformation step, we need to contract the RI coefficients tensor with three different matrices, as shown in Equation 4.33.

$$\sum_{\mu} \sum_{s',t'} C_{s',t'}^\mu X_{s',s} Y_{t',t} V_{\mu,\lambda} = \hat{C}_{\bar{s},\underline{t}}^\lambda \quad (4.33)$$

If these contractions are executed in a specific order, the *a priori* locality constraints in the chosen RI scheme can be used to reduce the formal scaling.

$$\text{step 1: } \sum_{\mu} C_{s',t'}^\mu V_{\mu,\lambda} = \bar{C}_{s',t'}^\lambda \quad \sim N_{\text{bas pairs}} N_{\text{local aux}} N_{\text{aux}} \quad (4.34)$$

$$\text{step 2: } \sum_{s'} \bar{C}_{s',t'}^\lambda X_{s',s} = \tilde{C}_{s,t'}^\lambda \quad \sim N_{\text{bas pairs}} N_{\text{aux}} N_{\text{local bas}} \quad (4.35)$$

$$\text{step 3: } \sum_{t'} \tilde{C}_{s,t'}^\lambda Y_{t',t} = \hat{C}_{s,t}^\lambda \quad \sim N_{\text{bas}}^3 N_{\text{local aux}} \quad (4.36)$$

In the first step we use the fact that the RI-LVL for atomic ERIs only uses the auxiliary basis functions centered on the same atoms. Thus the μ axis has at most $2N_{\text{local aux}}$ significant entries, where $N_{\text{local aux}}$ is the maximum number of auxiliary basis functions per atom. Furthermore, the strictly limited radial extend of the NAOs will limit the number of basis function pairs $N_{\text{bas pairs}}$ with non-zero overlap in the large system limit, giving the first contraction step an asymptotically quadratic scaling. The choice of the local ABS for the transformed ERIs can be exploited in the second step. By construction the ABS for each occupied transformed function $\varphi_{\bar{s}}$ is limited to a subset $\{P_\lambda\}$ of the global ABS. This relation can be inverted to obtain the set of occupied transformed basis functions $\{\varphi_{\bar{s}}\}$ which belong to a given

auxiliary function P_λ . With this information the sum in Equation 4.35 only needs to be evaluated for those $\varphi_{\bar{s}}$ which belong to a given P_λ . Since the radius r_{aux} to determine the local ABS will be small in practice (see the discussion in chapter 5), the number of $\varphi_{\bar{s}}$ belonging to an auxiliary basis function is effectively constant. The second step therefore has a cubic scaling with system size and also features the same asymptotically quadratic scaling as the previous step. The last contraction only needs to be done for those auxiliary basis functions that belong to the local ABS of the already transformed index. Therefore, this step also has a cubic scaling, but since the previous step destroyed the strict vanishing overlap criterion of the atomic basis functions, its asymptotic scaling is not reduced.

Thus, the integral transformation using this local ABS has an overall cubic scaling with system size. A further reduction of the scaling could be achieved if the sparsity of the pseudo-density matrices is taken into account to skip insignificant contributions.

4.7 THE INTEGRAL SCREENING CRITERION

The performance of the LT-MP2 is highly dependent on the integral screening criterion that should skip as many basis function tuples with small contributions to the MP2 correlation energy as possible in Equation 4.7. The original LT-MP2 implementations by Häser [15] and Scuseria [81] employed the so-called QQZZ-screening [129] which is based on the Cauchy-Schwarz inequality shown in Equation 4.37. [93]

$$(st|uv) \leq (st|st)^{\frac{1}{2}} (uv|uv)^{\frac{1}{2}} \quad (4.37)$$

This inequality holds for any kind of orbitals forming the ERI. Applying this inequality to the ERIs appearing in the LT-MP2, one obtains the strict upper bounds for both the atomic and the transformed ERIs shown in Equation 4.38 to 4.41.

$$|(st|uv)| \leq Q_{s,t} Q_{u,v} \quad (4.38)$$

$$|(\bar{s}t|\bar{u}v)| \leq Z_{s,t} Z_{u,v} \quad (4.39)$$

$$Q_{s,t} = (st|st)^{\frac{1}{2}} \quad (4.40)$$

$$Z_{s,t} = (\bar{s}t|\bar{s}t)^{\frac{1}{2}} \quad (4.41)$$

The $Q_{s,t}$ matrix is straight forward to compute. If a RI is used, $Z_{s,t}$ can also be obtained easily because the effort to evaluate the matrix is small compared to the generation of the transformed coefficients. In an implementation without RI however, $Z_{s,t}$ cannot be obtained directly because the required transformations would negate a significant part of the expected time savings. Instead one has to rely on upper bound estimates for the elements of this matrix which only require the evaluation of half-transformed ERIs. [15]

Using QQZZ, an upper bound estimate for the magnitude of each summand in Equation 4.7 can be made.

$$|(\bar{s}t|\bar{u}v) [2(st|uv) - (sv|ut)]| \leq Z_{s,t} Z_{u,v} [2Q_{s,t} Q_{u,v} - Q_{s,v} Q_{u,t}] \quad (4.42)$$

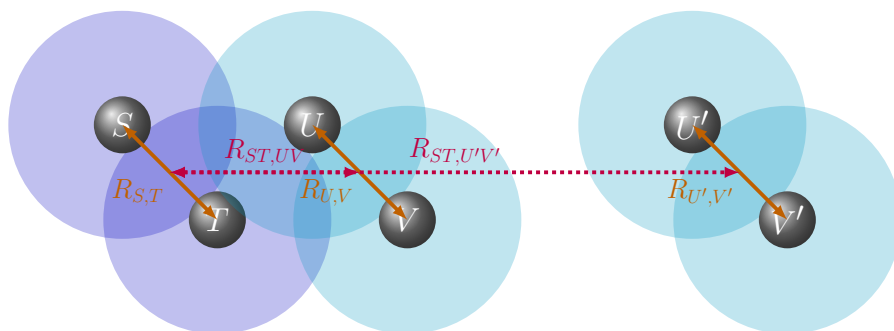


Figure 4.4: A schematic representation of the shortcoming of the QQZZ screening approach. If the ERI ($st|uv$) is estimated with this method, only the distances $R_{S,T}$ and $R_{U,V}$ between the function pairs forming the left and right side of the integral are taken into account. The distance between the two charge distributions $R_{ST,UV}$ on the other hand is not accounted for in the estimate. As a consequence, the ERI ($st|u'v'$) where the right side functions from ($st|uv$) have been shifted further away, will yield exactly the same estimate, although the integral will be considerably smaller in reality.

The number of significant integrals can be reduced significantly by discarding all terms below a chosen threshold. [15, 81]

However, QQZZ has a significant limitation. It only takes the distance between the functions in the two pairs into account, but not the distance between the pairs as shown in Figure 4.4. It is evident that the latter has an important impact on the sparsity of both the atomic and transformed ERIs. Therefore excluding this distance decay in the screening criterion will lead to a significant underestimation of the locality in the MP2 correlations especially for large systems.

The Ochsenfeld group tried to overcome this shortcoming by introducing Multipole Based Integral Estimates (MBIE). [41, 130] This method allows to include the distance decay of the Coulomb operator in those ERIs which are formed by non-overlapping pairs. At first the ERI is rewritten as a multipole series, which is known to converge for non-overlapping charge distributions. Lambrecht and his coworkers then replaced the multipole moments with absolute multipole moments $\mathcal{M}_{s,t}^{(n)} = \int r^n |\varphi_s(r)\varphi_t(r)| dr$ which are upper bounds to the absolute values of the real multipole moments. They then used the convergence properties of the multipole series to demonstrate that the N th order term in their estimate is an upper bound for all higher order terms (including N th order), if the distance R between the charge centers is replaced by a rescaled distance R' . The rescaled distance $R' = R - R_{st} - R_{uv}$ is the distance between the two charge centers in the multipole expansion minus the radial extends of the expanded charges. In the resulting formula, the sum over the higher order terms can be carried out analytically, which then leads to the

final MBIE equations at N th order. The second order MBIE is shown in Equation 4.43. [130]

$$|(st|uv)| \leq \left| \frac{\mathcal{M}_{s,t}^{(0)} \mathcal{M}_{u,v}^{(0)}}{R} \right| + \left| \frac{\mathcal{M}_{s,t}^{(1)} \mathcal{M}_{u,v}^{(0)} + \mathcal{M}_{s,t}^{(0)} \mathcal{M}_{u,v}^{(1)}}{R^2} \right| + \left| \frac{\mathcal{M}_{s,t}^{(2)} \mathcal{M}_{u,v}^{(0)} + 2\mathcal{M}_{s,t}^{(1)} \mathcal{M}_{u,v}^{(1)} + \mathcal{M}_{s,t}^{(0)} \mathcal{M}_{u,v}^{(2)}}{R^3 - R'^2} \right| \quad (4.43)$$

As shown in a follow-up paper [41], a similar expression can be obtained for the transformed integrals by applying the multipole expansion to each term in Equation 4.19. They argued that the transformation with the pseudo-density matrices makes the zeroth-order multipole moments vanish, thus the transformed absolute multipoles $\tilde{\mathcal{M}}_{s,t}^{(n)}$ can be set to zero for the zeroth order. As a consequence, the terms appearing in the second order MBIE are the leading terms in the integral bounds of the transformed ERIs.

$$(\bar{s}t|\bar{u}v) = \sum_{s',t',u',v'} X_{s',s} Y_{t',t} (s't'|u'v') X_{u',u} Y_{v',v} \quad (4.44)$$

$$|(\bar{s}t|\bar{u}v)| \leq \left| \frac{2\tilde{\mathcal{M}}_{s,t}^{(1)} \tilde{\mathcal{M}}_{u,v}^{(1)}}{R'^3 - R'^2} \right| \quad (4.45)$$

While MBIE provides strict upper bounds to the ERIs, it still overestimates many integrals significantly because the absolute multipole moments are used. To overcome this shortcoming, the Ochsenfeld group later developed the QQZZR4-screening, which is an empirical combination of QQZZ with the distance-based corrections from the MBIE. [131, 132]

$$|(st|uv)| \approx \frac{Q_{s,t} Q_{u,v}}{R'} \quad (4.46)$$

$$|(\bar{s}t|\bar{u}v)| \approx \frac{Z_{s,t} Z_{u,v}}{R'^3} \quad (4.47)$$

These integral estimates are no longer true upper bounds, but the estimates are much closer to the real magnitudes of ERIs with well-separated charge extends than either QQZZ or MBIE and have been proven to yield accurate results. In both MBIE and QQZZR4 the conventional QQZZ screening is used for ERIs with overlapping charge distributions.

In my implementation both the QQZZ and QQZZR4 methods are available. The charge extends are determined by means of the “outer radii” (this is the outermost radius where the absolute function value exceeds a chosen threshold) of the functions forming the pair. For simplicity, the same charge centers and extends are used for the atomic and transformed ERIs in my LT-MP2 project.

As already stated before, the QQZZ-screening based on the Cauchy-Schwarz inequality is applicable to any functions forming the charge densities. Thus it can be directly applied to the ERIs formed by crystalline orbitals as well. The generalization of the QQZZR4-screening screening to periodic systems is however more complicated. Since

the original derivation is based on multipole expansions of localized charges, it cannot be applied to the ERIs of crystalline basis functions directly. Expanding the crystalline basis functions in the atomic basis functions and applying the QQZZR4 screening to each term in the resulting sums leads to a considerably increased computational effort for the integral screening and is of no practical use. Additional approximations like using an upper bound for the distance between each unit cell pair also did not yield any noteworthy improvements. Thus the periodic implementation with crystalline orbitals only uses QQZZ screening at present. The distance corrections from the QQZZR4 screening can be incorporated more easily if a real-space representation is used instead of crystalline functions.

4.8 PARALLELIZATION FOR THE LT-MP2

Another important challenge for implementing the LT-MP2 in a RI framework is the parallelization strategy. Due to the significance of integral screening in the LT-MP2, the method does not expose *regular parallelism*. One cannot predict with reasonable effort which chunk of data will be needed on a process at a given time. Thus the algorithm cannot be formulated as a sequence of communication and computation phases. Instead, the LT-MP2 features *irregular parallelism*. The individual chunks of work are still independent of each other and can be processed in parallel, but the real computational effort for each chunk is not known in advance. Thus, a dynamic load balancing and a flexible communication network are needed.

To address this problem, I decided to use a hybrid approach which combines the Open Multi-Processing (OpenMP) [133] and Message Passing Interface (MPI) [134] frameworks. OpenMP is a thread-based framework, where a single process spawns several threads which operate in shared memory. Threads spawned by the same process can access the same memory regions directly and thus no data transfer overhead occurs. The usage of shared memory however implies that OpenMP is bounded by the physical hardware, i.e. all threads spawned by one process must run on the same computer or cluster node. MPI on the other hand is a communication framework to exchange data between processes by means of explicitly invoked communication functions. As such, using MPI causes communication overhead, but data can be transmitted across a network infrastructure and thus be shared among processes on different machines.

OpenMP is particularly appealing for irregular parallelism problems, because it has a set of powerful features aimed explicitly at such problems and offers the advantage of shared memory. Thus, the LT-MP2 implementation is optimized for a moderate number of processes with many threads each. MPI is used for the high-level workload distribution and data exchange between the processes. Once a chunk of work has been assigned to a process, it fetches all required data from other processes. From there on, the assigned work is split into smaller subtasks, which are processed in parallel using OpenMP

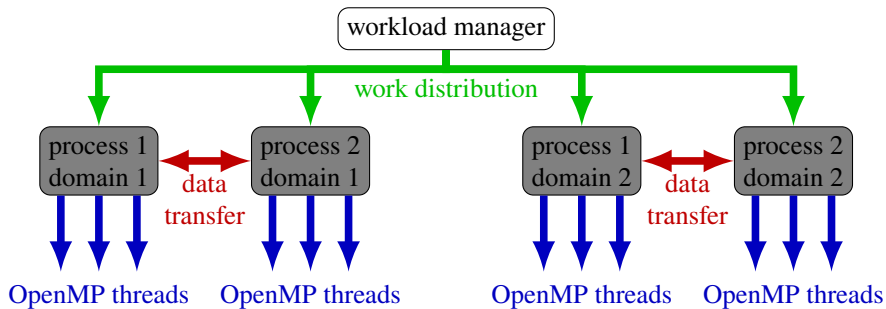
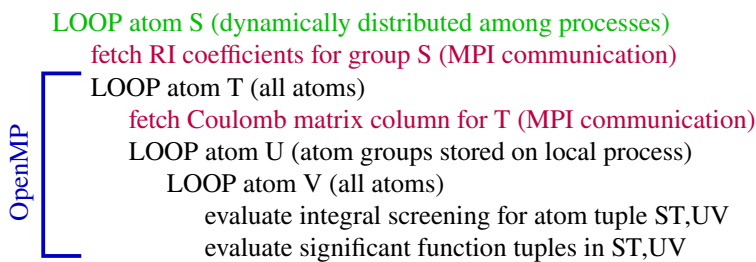
Pseudo code for the integral evaluation

Figure 4.5: A schematic representation and pseudo-code of the parallelization structure used for the integral evaluation. The color coding relates steps in the pseudo-code to the corresponding elements in the schematic. Available processes are divided into one or more computational domains, which all have a complete copy of the required data. Thus global data transfers can be avoided and communication only happens inside a computational domain. The processes then use OpenMP threads to handle the assigned workload and use shared memory as much as possible to minimize the required communication. The workload chunks are distributed dynamically at runtime by a central controller to take the workload reduction due to the screening into account.

and without any cross-process communication. The parallelization structure is illustrated in Figure 4.5 with a schematic representation and a pseudo-code.

To minimize the communication overhead, the RI coefficients $C_{s,t}^\mu$ and $M_{s,t}^\mu$ are divided into atomic blocks $C_{S,T}^{\forall\mu}$ and $M_{S,T}^{\forall\mu}$. All coefficient blocks belonging to the same atom S on the first basis function index are assigned to the same process. Such a set of coefficients is hereafter referred to as “atom group” $C_{S,\forall T}^{\forall\mu}$. The atom groups are distributed such among the processes that the required memory per process is as balanced as possible.

Consider again the energy formula of the LT-MP2:

$$\mathcal{E}_{MP2} = - \sum_q w_q \sum_{s,t,u,v} (\bar{s}t|\bar{u}v)^q \left[2(st|uv) - (sv|ut) \right] \quad (4.48)$$

This formula can be divided into a set of atomic contributions.

$$\mathcal{E}_{MP2} = - \sum_q w_q \sum_{S,T,U,V}^{\text{atoms}} \sum_{\substack{s \in S, t \in T, \\ u \in U, v \in V}} (\bar{s}t|\bar{u}v)^q \left[2(st|uv) - (sv|ut) \right] \quad (4.49)$$

Each work assignment to a process then consists of all the locally stored atom groups U and one atom group S which can be stored on another process. Since each atom group includes all atoms on the second basis function index, the process only needs to fetch the atomic and transformed coefficients for the atom group S . Afterwards the process can evaluate all ERIs belonging to this chunk of work without the need to fetch any more RI coefficients later. This feature is particularly important to avoid communication overhead in the evaluation of the exchange term. Furthermore, there is an important symmetry we can exploit.

$$(\bar{s}t|\bar{u}v) \left[2(st|uv) - (sv|ut) \right] = (\bar{u}v|\bar{s}t) \left[2(uv|st) - (ut|sv) \right] \quad (4.50)$$

This symmetry implies that the contributions from the atomic group pairs S, U and U, S will be equivalent and we thus need to evaluate only one. The decision which pair will be evaluated is made at runtime and is used to manage the load balancing between the processes.

If the provided computational resources have enough memory available to hold two or more copies of the RI coefficients, the nodes can be split into multiple computational “domains”. Each computational domain then holds a complete copy of the required data, which allows the restriction of data transfers to only occur within the bounds of these domains. Only negligible communication for the workload distribution still happens in the global scope between all participating processes.

In the previous chapter I demonstrated how the RI-LVL can be generalized and combined with the Laplace-Transformation (LT) to obtain a new technique for the reduction of the computational and memory scaling in advanced correlation methods. At the example of the 2nd Order Møller-Plesset Perturbation Theory (MP2), I discussed in detail how this new technique give rise to a new Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP2) implementation with efficient large-scale parallelization and a $\mathcal{O}(N^2)$ memory consumption.

In this chapter I demonstrate the practical applicability of this new LT-MP2. Using water clusters and titanium dioxide surfaces as examples, I will at first prove the accuracy of the implementation. To be of practical use, the algorithm must have a reduced scaling with system size, break even with the canonical MP2 at reasonable system sizes and feature a considerably reduced memory footprint. Furthermore, it should also exhibit a good parallel efficiency despite being an irregular parallel problem. I will demonstrate that my implementation fulfills all these requirements with water clusters containing up to 150 molecules. To prove that this LT-MP2 implementation can also handle covalently bonded systems with a similar accuracy, the results for the water clusters are complemented with a study of the ISOL22 test set, which contains isomerization reactions for molecules with up to 51 atoms. [11]

All calculations shown in this chapter are frozen-core MP2 calculations, catering to the well established valence-correlation consistent Numeric Atom-Centered Orbital (NAO) basis sets. [32] All-electron calculations can be handled by the LT-MP2 as well, but studies from other groups (see e.g. [128]) demonstrated that they will usually require a larger number of quadrature points to reach a similar accuracy. All LT-MP2 calculations use the eigenvectors and eigenvalues of converged Hartree-Fock (HF) calculations with RI-LVL as starting point. The NAO-VCC-2Z basis set has been used throughout this chapter with an additional f -type function in the enhanced orbital basis set (OBS+). For titanium the tier1 basis without OBS+ has been used since no NAO-VCC-2Z basis set is available for this element.

The accuracy of the canonical MP2 implementation in FHI-aims has been proven by comparisons to other codes for both periodic and molecular systems, thus the accuracy analysis in this chapter will only compare to the canonical MP2 in FHI-aims. (see e.g. the comparison to NWChem in chapter 3 and Refs [96, 127].) A performance comparison to the recently published DEC-MP2 from Kjærgaard and co-workers will be presented at the end of the next chapter.

5.1 ACCURACY

I have carefully analyzed the accuracy controlling parameters of the LT-MP2, in particular the number of quadrature points N_q , the integral screening threshold and the radius of the local auxiliary basis set (ABS) r_{aux} for the generalized RI-LVL. A short overview of these results is presented here, a more detailed analysis for the interested reader can be found in Appendix B.

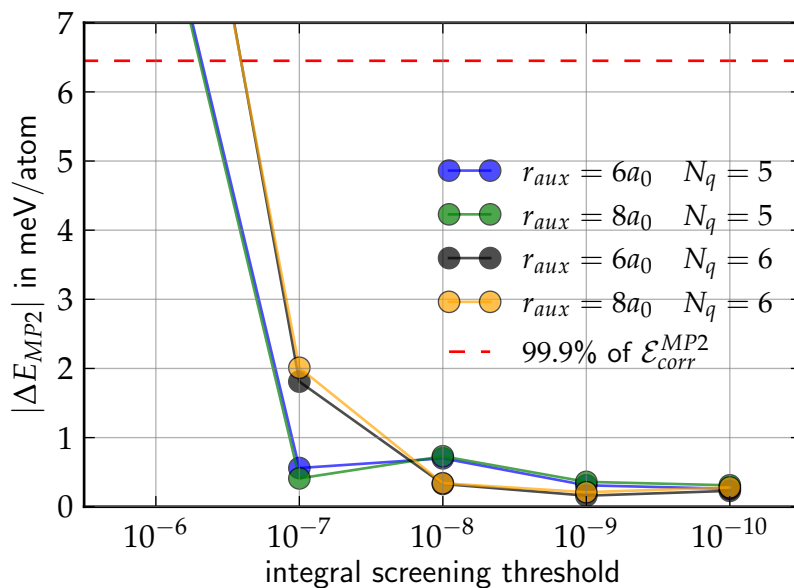
Taking a cluster with 20 water molecules and a titanium dioxide surface with 48 atoms as examples, Figure 5.1 shows the errors per atom in the MP2 correlation energy for various choices of the LT-MP2 parameters. Using 5 or 6 quadrature points is sufficient to reach an accuracy of about 1 meV/atom. This is in agreement with the reports from the authors of the minimax approximation [128] who suggest to use 5 to 6 quadrature points for frozen-core MP2 calculations. The generalized RI-LVL used for the evaluation of the transformed Electron Repulsion Integrals (ERIs) converges quickly with the radius r_{aux} of the local ABS and a value of 6 Bohr suffices to yield very accurate results. For the integral screening threshold a value of 10^{-8} or 10^{-9} is needed to ensure a good convergence. It should be noted that tighter integral thresholds are required for covalently bonded systems. 10^{-7} yields acceptable errors for the water clusters, but errors larger than 12 meV/atom for the titanium dioxide surface. Additional convergence tests including the larger NAO-VCC-3Z basis set can be found in the appendix.

Unless mentioned otherwise the calculations shown in this chapter use 5 quadrature points, an integral threshold of 10^{-9} and $r_{\text{aux}} = 6a_0$.


5.2 SCALABILITY WITH SYSTEM SIZE AND COMPUTATIONAL RESOURCES

To demonstrate the scalability of my LT-MP2 implementation I use the water cluster geometries from Ref. [97] which range from 10 to 150 molecules and are snapshots from a Molecular Dynamics (MD) trajectory using classical force-fields.

The accuracy of the LT-MP2 does not degrade with system size as shown in Figure 5.2. The correlation energy errors are below 1 meV/atom compared to the canonical MP2 using both RI-LVL_{full} and RI-V in all cases. This good agreement indicates that the hybrid local Resolution of Identity (RI) strategy combining the standard RI-LVL for the atomic ERIs with the generalized RI-LVL (section 4.5) for the transformed ERIs is a numerically stable and accurate solution to address the memory problem. Regardless of the chosen reference point, the LT-MP2 recovers more than 99.9% of the total MP2 correlation energy, which the Neese group often sets as required target accuracy for their local correlation methods. [37] This target accuracy translates to an upper limit for the acceptable errors of about 6.5 meV/atom in this test case.



(a) Accuracy of the LT-MP2 for a water cluster with 20 molecules



(b) Accuracy of the LT-MP2 for a titanium dioxide slab with 2 layers (24 atoms per layer)

Figure 5.1: Absolute errors per atom in the MP2 correlation energy of the LT-MP2 compared to the canonical MP2 calculations. N_q denotes the number of quadrature points and r_{aux} is the radius of the generalized RI-LVL for the transformed integrals. Errors below the dashed red lines recover more than 99.9% of the total MP2 correlation energy. For hydrogen and oxygen the NAO-VCC-2Z basis sets with a f -type function in the OBS+ are used, for titanium the tier1 basis without OBS+.




Figure 5.2: Absolute errors per atom in the MP2 correlation energy of the LT-MP2 compared to the canonical MP2 calculation with different RI strategies. The dashed red line denotes the maximum allowed error to recover at least 99.9% of the total correlation energy.

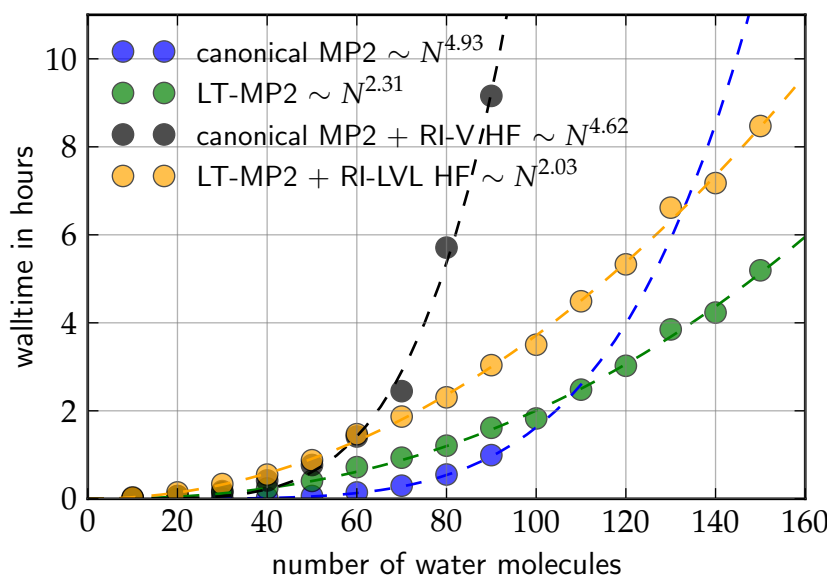


Figure 5.3: Total time used for the calculation of the MP2 correlation energy with both the canonical and LT-MP2 implementations. In addition the total time including the overhead for the generation of the atomic RI coefficients and the HF SCF cycle is shown for both variants. All calculations used 4 cluster nodes with 32 cores each.

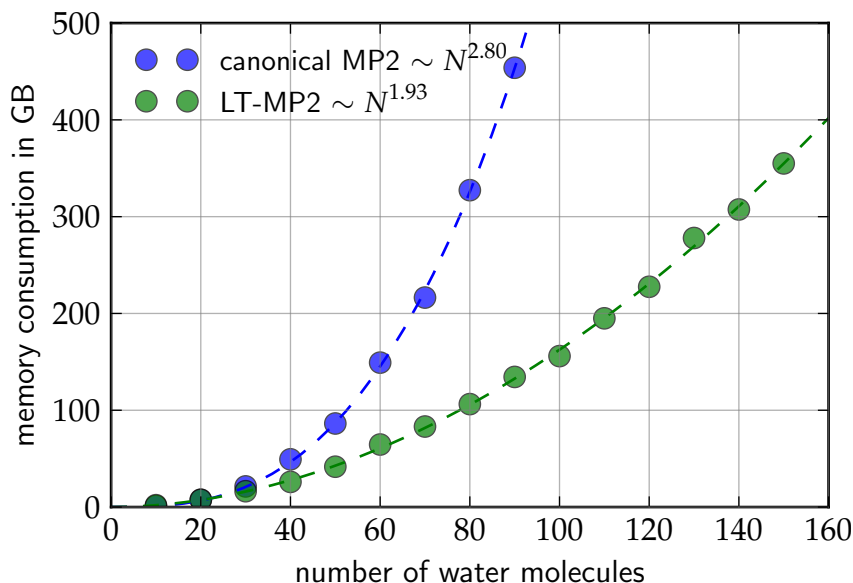


Figure 5.4: Memory consumption for the dominant objects in both canonical MP2 (blue) and LT-MP2. (green) These are the RI coefficients for both atomic and transformed/molecular orbitals. The LT-MP2 curve also includes the Coulomb matrix of the auxiliary basis functions and the precomputed solution matrices for the equation systems we need to solve for the transformed RI coefficients.

The computational scaling of the LT-MP2 compared to its canonical counterpart is demonstrated in Figure 5.3. If we only consider the calculation of the MP2 correlation energy itself, the fifth-order scaling of the canonical formulation can be clearly seen, while the LT-MP2 features an overall less than cubic scaling $\mathcal{O}(N^{2.3})$. Due to a considerably larger prefactor, the break-even happens at about 110 molecules. However, for the canonical MP2 with RI-V the dominating step at these system sizes is not the MP2 correlation itself, but the preceding generation of the RI coefficients. These are needed for the HF Self-Consistent Field (SCF) cycle as well, but they usually dominate both HF and MP2. If the time required to generate the atomic RI coefficients and evaluate the HF SCF cycle is included in the comparison, the overall picture changes dramatically in favor of the LT-MP2. The use of the RI-LVL dramatically reduces the effort for the RI coefficient generation and the break-even already happens at about 60 molecules. In principle, one could also use RI-LVL for the canonical MP2 by means of the RI-LVL_{full} algorithm, but this approach requires the usage of an OBS+ to maintain the accuracy. (see section 3.5) This would significantly increase the memory requirements for the molecular RI coefficients and is thus of no practical use.

The scaling of the memory usage is another important aspect of the algorithm. Figure 5.4 compares the memory consumption of the LT-MP2 with the canonical MP2 using the standard RI-V. The shown data only includes the persistent objects which dominate the total memory requirement. Some additional memory is needed in both

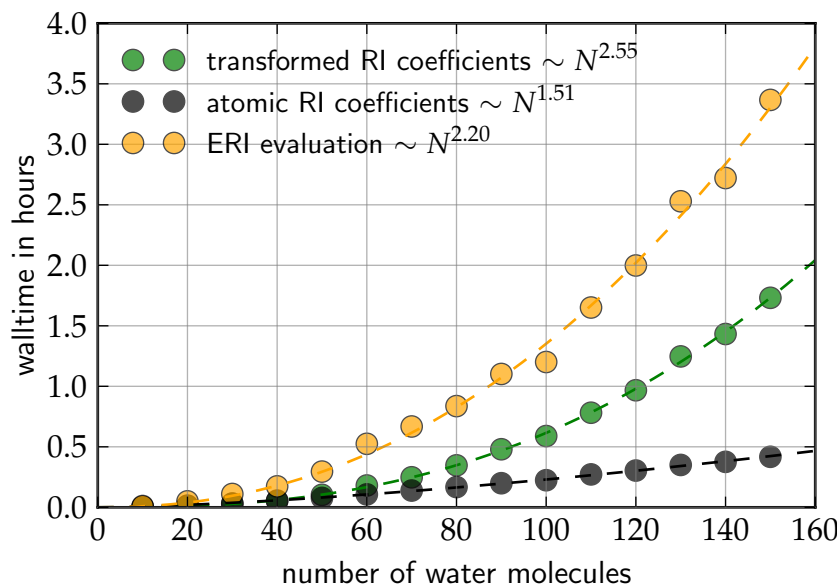


Figure 5.5: Time spent for the main steps in the LT-MP2 and their scaling with system size. The timings are summed over all quadrature points. All calculations used 4 cluster nodes with 32 cores each.




Figure 5.6: Scalability of the two main steps in the LT-MP2 with the number of available cores. For this analysis, two processes with 16 threads each have been used per computational domain. (see section 4.8 for details about the parallelization scheme) The low parallelization efficiency in the less time-consuming (see Figure 5.5) transformation step is caused by the use of an easier to implement global synchronization. The parallel efficiency of this step can be significantly improved with a more sophisticated implementation relying on process-to-process communication.

implementations for intermediate results and communication buffers. The LT-MP2 features a quadratic scaling with system size, while the canonical variant shows a near cubic scaling. Both curves have a scaling exponent which is below the expected theoretical value. This is caused by sparsity exploits in the atomic RI coefficients. Due to the strict radial confinement of NAOs in the FHI-aims code, any basis function pair with no overlap does not need to be stored. The asymptotic scaling of the atomic RI coefficients thus reduces to quadratic in RI-V and linear in RI-LVL. The molecular RI coefficients in the canonical MP2 do not have a similar sparsity and thus need to be stored completely, yielding a cubic scaling. The transformed RI coefficients in the LT-MP2 do not have a strict zero overlap criterion either, but they decay with increasing distance between the transformed basis functions. If one would take advantage of this property, the memory footprint of the dominant object in my LT-MP2 implementation could be reduced to scale asymptotically linear with system size as well.

Figure 5.5 shows a more detailed analysis of the computational scaling of the main steps in the LT-MP2. These are the preparation of the atomic RI coefficients, the generation of the transformed RI coefficients and the evaluation of the significant ERIs. All of these steps feature a less than cubic scaling behavior, which indicates that the overall scaling shown in Figure 5.3 will not degrade if the system size is increased further. It should also be noted that the generation of the atomic RI coefficients at present does not make use of Open Multi-Processing (OpenMP) parallelization due to implementation constraints. With OpenMP support, this step could be sped up by a factor 10 to 15, thus reducing it to a negligible part of the overall calculation.

The scalability of my LT-MP2 implementation with respect to the available computational resources is presented in Figure 5.6. The evaluation of the significant ERIs is the dominant step in the calculation. (see Figure 5.5) More importantly, it is the technically more involved one due to the irregular parallelism. As shown in Figure 5.6 a near perfect parallelization scalability can be achieved for this step if the domain-based memory distribution (see section 4.8) is used to avoid global communication and the unnecessary re-evaluation of intermediates. The generation of the transformed RI coefficients on the other hand has a considerably lower parallelization efficiency. Since this part of the method has not been the overall bottleneck so far, an easier to implement global synchronization has been chosen for this step. If this synchronization is replaced with direct process-to-process communication without global synchronization barriers, the parallel efficiency of this step can be significantly improved.

As discussed in the previous chapter, the generation of the transformed RI coefficients can be limited to a cubic scaling with system size if the contractions are executed in a specific order. (see discussion in section 4.6) The numerical confirmation is shown in Figure 5.7 where the computational scaling for the individual steps of the transformation is analyzed.

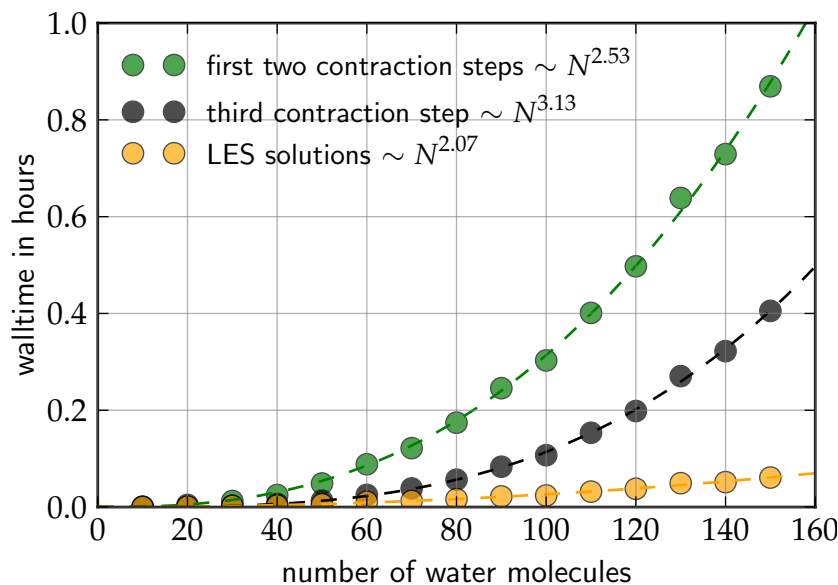


Figure 5.7: Total time spent for different stages in the generation of the transformed RI coefficients. The shown times are per calculation, not per quadrature point. See section 4.5 and equations 4.34 to 4.36 for a detailed explanation of the steps in the RI coefficient transformation. All calculations used 4 cluster nodes with 32 cores each.

5.3 LT-MP2 IN THE ISOL22 TEST SET

The LT-MP2 is relying heavily on integral screening and it is important to verify that the quality of the integral screening does not degrade with increasing system size. The analysis of the water clusters in the previous section demonstrated that the chosen parameters are well-converged for non-covalently bonded systems. To verify the applicability of my LT-MP2 implementation to covalently bonded systems, I also investigated the ISOL22 test set [12] which is a subset of the ISOL test set [11]. It contains 22 isomerization reactions with system sizes ranging from 24 to 51 atoms. The test set contains the start and end point geometries for each reaction and provides theoretical reference values for the energy difference between them. The accuracy of MP2 itself for this test set has been investigated in previous literature [11] where it yielded overall good results, but a rather high error in one test case.

Here I focus on the accuracy of the numerical approximations made in my LT-MP2 implementation, thus the NAO-VCC-2Z basis set is used and the canonical MP2 with RI-V is used as reference. The results of this comparison are shown in Figure 5.8. The error per atom is below 1.9 meV/atom in all cases and yields a good Root-Mean Square Deviation (RMSD) of about 1.0 meV/atom. Errors below 3.8 meV/atom are sufficient to recover more than 99.9% of the total MP2 correlation energy in all of these systems.

In addition to the analysis of the correlation energy itself, the MP2 corrections to the isomerization energies are listed in Table 5.1. These

reaction index	atoms (non-hydrogen)	MP2 contribution to isomerization energy in meV	
		canonical MP2	LT-MP2
2	41 (17)	875	870
3	24 (12)	-176	-175
5	32 (16)	490	488
6	48 (28)	43	49
7	51 (21)	493	485
8	43 (22)	-237	-242
9	32 (12)	155	151
10	35 (15)	391	393
11	30 (20)	-746	-752
12	40 (24)	23	14
13	26 (14)	284	283
14	26 (14)	224	220
15	42 (18)	35	28
16	51 (28)	291	279
17	44 (24)	250	230
18	39 (23)	276	283
19	36 (23)	-74	-70
20	28 (16)	60	54
21	36 (20)	131	139
22	44 (29)	412	408
23	39 (18)	446	469
24	52 (19)	142	143

Table 5.1: MP2 isomerization energy corrections for the ISOL test set using NAO-VCC-2Z basis sets.




Figure 5.8: Absolute MP2 correlation errors per atom in different systems from the ISOL22 test set with NAO-VCC-2Z basis sets. The LT-MP2 results are compared to the RI-LVL_{full} based canonical MP2 calculations. Errors smaller than the dotted red line recover more than 99.9% of the MP2 correlation energy. The reactions 1 and 4 are excluded in the ISOL22 test set.

values are not to be confused with the isomerization energy itself, because the HF contribution is not considered here. Comparing the corrections to the isomerization energies, it can be concluded that the errors introduced by the LT-MP2 are systematic because the differences in the isomerization energy contributions are consistently smaller than the errors of the individual systems. In reaction 16 (51 atoms, shown in Figure 5.9) for example, the isomerization energy contribution differs by 12 meV, while the starting and end geometries have total errors of 45 and 34 meV.

A part of the small errors in the isomerization energies is potentially caused by the choice of the local ABS, in particular for those




Figure 5.9: Initial and final geometry of the isomerization reaction 16 in the ISOL22 test set.

reactions like 16 where the geometry does not change too drastically. As explained in section 4.5, the local ABS includes all auxiliary basis functions within a certain distance r_{aux} to the host atom of the “occupied” transformed basis function. This implies that a small change of the geometry might give rise to a notable change of the local ABS, which can cause discontinuities of the total energy. This can happen for example for the atoms in the different rings of this system which are closer to each other once the initially planar ring structure has been folded.

To avoid this potential problem, the construction of the local ABS in the generalized RI-LVL should be using a bump-function like it was suggested by the Head-Gordon group in their atomic RI. [113, 114] This bump function adds a small transition region where the contributions of auxiliary basis functions to the local RI coefficients are gradually screened out. At the outer radius of this transition area they are filtered out completely and can be removed from the local ABS without introducing any discontinuities in the potential energy surface.

In conclusion, the application to the ISOL22 test set demonstrates that the control parameters of the LT-MP2 are transferable and do not have a strong dependence on the structure of the system they are applied to.

LT-MP₂ APPLIED TO WATER ADSORPTION ON TITANIUM DIOXIDE

In the previous chapters I introduced the RI-LVL, my new localized Resolution of Identity (RI) strategy for Hartree-Fock (HF) and demonstrated how it can be generalized and combined with a Laplace-Transformation (LT) to obtain a local RI strategy that can be used for advanced correlation methods. This new technique can be used to considerably reduce the memory requirement of such methods and combined with integral screening, it can also reduce the computational scaling significantly. As I have demonstrated at the example of the MP2 in the previous chapter, this new generalized RI-LVL framework reaches a very high accuracy for both molecular and periodic systems. I also demonstrated the considerable computational time and memory savings achieved by this new lower-scaling algorithm for a set of water clusters both at the HF and MP2 level. To prove the practical applicability of my method for periodic systems as well, I also investigated the problem of water adsorption on a rutile titanium dioxide (TiO₂) surface in the low coverage limit. The adsorption energy \mathcal{E}_{ads} of a molecule onto a substrate is defined as the total energy of the composite system minus the total energies of the substrate and molecule in isolation.

$$\mathcal{E}_{ads} = \mathcal{E}_{composite} - \mathcal{E}_{substrate} - \mathcal{E}_{molecule} \quad (6.1)$$

Interactions at surfaces are a particular interesting topic due to their many possible applications, but their theoretical modeling is very challenging. (see e.g. [18–22]) It is often desirable to use advanced correlation methods to obtain an accurate description of the interaction between molecule and substrate. Embedded Cluster Models (ECMs) (see section 2.9) can be used instead of periodic boundary conditions to keep the computational effort for such methods tractable. A drawback of the ECM is however that periodicity is considered only in the force-field model of the ideal surface. Coverage-dependent effects therefore cannot be studied within this framework. Periodic boundary conditions on the other hand are well-suited to investigate coverage-dependent effects, but a tremendous effort is required to investigate the low-coverage limit since large unit cells are required to avoid interactions between the adsorbed molecule and its periodic images. Lower-scaling variants of advanced correlation methods are particularly suited for such tasks because the workload reduction achieved by sparsity exploits and integral screening does not rely on the presence of symmetries in the investigated cell.

In this chapter I analyze the convergence of the adsorption energy for a water molecule on a rutile TiO₂ surface with respect to system size using periodic boundary conditions. In addition to the results obtained with my Laplace-Transformed 2nd Order Møller-Plesset Per-

turbation Theory (LT-MP2) implementation, I will also present rPBE, HF and HSE06 results. Van-der-Waals corrections should in principle be included for these functionals, but unfortunately the surface-specific functionals in the FHI-aims code (Many-Body-Dispersion [73] and vdW^{surf} [71]) cannot be used with TiO_2 surfaces at present.

My LT-MP2 implementation clearly outperforms the canonical MP2 in FHI-aims both with respect to computational time and memory consumption. More importantly, its reduced scaling features considerably increase the accessible system size. Furthermore, it uses the same numerical setup for both molecular and periodic systems, which has the advantage that it is not necessary to reach the Complete Basis Set (CBS) limit to make comparisons between periodic boundary conditions and the ECM. My new LT-MP2 implementation thus paves the way towards a direct comparison of periodic boundary conditions and the ECM for advanced correlation methods.

6.1 COMPUTATIONAL SETUP

The surface geometries used in this investigation are cut out along the (110) plane from a bulk TiO_2 crystal with experimental lattice constants taken from Ref. [135]. The water molecule is placed at a distance of 2.5\AA above a top-layer titanium atom as shown in Figure 6.1. The TiO_2 slab is build from O- Ti_2O_2 -O tri-layers as depicted in the figure. (Hereafter simply referred to as “layers”.) The periodic images of the slab along the z-axis have been separated by a 100\AA vacuum layer and decoupled electrostatically via a dipole correction. [136]

The same geometries are used for all methods to allow a direct comparison of the convergence behavior of the adsorption energy without having to account for effects from different geometries. In principle the use of experimental geometries (or geometries obtained from other methods) can cause artificial strain on the atoms if the lattice parameters differ from those predicted by the chosen method. Lee and coworkers investigated how different lattice constants affect the phonon modes and thus the dielectric constants of TiO_2 . [137] In their study they demonstrated that the HSE06 functional overestimates the static dielectric constants by up to 150% with its predicted equilibrium structures, but a much better description is obtained when the Local Density Approximation (LDA) geometry is used instead. The convergence behavior of the adsorption energy studied in this work is however not a function of the nuclear derivatives and thus should not be strongly affected by the choice of lattice parameters. For other applications, like surface relaxations in the presence of a vacancy in the bridging oxygen row it is however crucial to avoid artificial stress in the system. [89]

In all calculations I used the NAO-VCC-2Z basis set for the oxygen and hydrogen atoms. No NAO-VCC-XZ valence-correlation consistent NAO basis sets are available at present for titanium, therefore the tier1 basis set was used instead. Whenever Electron Repulsion Integrals (ERIs) are needed, RI-LVL was used with an enhanced orbital




Figure 6.1: Structure of the investigated TiO_2 systems. Along the z -axis (perpendicular to the surface) the system is composed of stacking tri-layers. The primitive surface cell (green) has an extent of 6.5\AA along the x -axis and 2.96\AA along the y -axis. The water molecule is placed at a distance of 2.5\AA to the titanium atom in the top surface layer.

basis set (OBS+) containing an additional f -type function ($z = 3.0$) for the oxygen and hydrogen atoms. In accordance with the results from subsection 3.8.1, no OBS+ has been used for the titanium atoms. “Tight” integration settings have been used in all calculations. [110] Larger basis sets have not been tested in this study because the effort is too high for the canonical MP2 with its fifth order scaling. The LT-MP2 should be able to handle such system sizes better, but as can be seen from the following results, additional code optimizations will be necessary at first to further improve its performance. For my LT-MP2 implementation I used the parameters tested in the previous chapter: 5 quadrature points, a screening threshold of 10^{-9} and a radius of $6a_0$ for the generalized RI-LVL. The LT-MP2 uses the supercell approach explained in section 2.8, i.e. all calculations use only the Γ point from a technical point of view, but the clean surface supercell containing a given number of primitive unit cells is exactly equivalent to sampling the primitive unit cell with a Γ -centered k -grid with the same number of k -points. In the following the notation $X \times Y$ will denote a surface unit cell which contains X primitive cells along the x -axis and Y many primitive cells along the y -axis. When necessary, the notation $X \times Y \times Z$ will additionally indicate the number of layers Z .





Figure 6.2: Convergence of the adsorption energy as a function of the number of layers with a 1×2 surface unit cell for both HF and the LT-MP2. The LT-MP2 graph only shows the contribution of the MP2 correlation to the adsorption energy, it does not contain the HF contribution.

6.2 RESULTS


Figure 6.2 shows the convergence of the adsorption energy with the number of layers in the system. The adsorption energy converges quickly with the number of layers and is essentially converged for both HF and MP2 when four layers are used. Three layers are already converged within a window of 5 – 6 meV. I decided to use three layers in the following calculations since it is to be expected that large unit cells will be needed to reach the low coverage limit. (A $3 \times 5 \times 3$ system already contains 270 atoms.)

The convergence of the adsorption energy with system size for rPBE, HF, HSE06 and LT-MP2 is shown in Figure 6.3. Each panel shows the convergence with increasing number of primitive cells along the y-axis for a fixed extent along the x-axis. As it can be seen from these results, all mean-field methods (rPBE, HF and HSE06) converge smoothly with increasing cell size. However, the LT-MP2 exhibits a notably different convergence behavior which requires further investigation. While the convergence with system size is similar to the mean-field methods if only one primitive cell is used along the x-axis, it becomes much slower when two or three cells are used along this axis.


To identify if this convergence behavior is the natural convergence with respect to cell size (or the corresponding k-grid) or an issue in the LT-MP2 implementation, I investigated how the individual contributions to the adsorption energy converge. The energy of the water molecule in the periodic cell did not exhibit any unexpected behavior and thus will not be discussed. The bare surface on the other hand



(a) 1 primitive cell along the x-axis



(b) 2 primitive cells along the x-axis



(c) 3 primitive cells along the x-axis

Figure 6.3: Convergence of the adsorption energy of $\text{H}_2\text{O}@Ti\text{O}_2$ as a function of the number of primitive cells along y-axis. (see Figure 6.1) All calculations use three layers along the z-direction. The MP2 energies do not include the HF contribution.

number of cells along y-axis	HF energy per primitive surface cell (eV)		
	1 x-cell	2 x-cells	3 x-cells
2	-164150.921	-164151.688	-164151.710
3	-164152.613	-164154.749	-164154.765
4	-164152.664	-164155.019	-164155.035
5	-164152.677	-164155.061	-164155.087
6	-164152.679	-164155.060	
7	-164152.679	-164155.061	

Table 6.1: Energy per TiO_2 primitive surface cell for different system sizes at the HF level of theory. All calculations contain three layers.




Figure 6.4: LT-MP2 correlation energy per primitive TiO_2 surface cell. All calculations use three layers along the z-direction.

reveals an interesting behavior which seems to be the cause of the observed convergence patterns. Table 6.1 contains the energies per primitive surface cell (see Figure 6.1) at the HF level for all investigated cell sizes. Due to the supercell approach, each calculation of the clean surface is exactly equivalent to a calculation of the primitive unit cell with as many k-points as we have cells in the supercell. At this level of theory, a smooth convergence is observed. Combined with the quickly converging adsorption energies at the HF level it is therefore safe to assume that all LT-MP2 calculations start from well-defined HF solutions. As illustrated by the results in Figure 6.4, the LT-MP2 calculations based upon these HF solutions exhibit a dramatically different convergence behavior. Similar to the adsorption energies (Figure 6.3) the LT-MP2 correlation converges rapidly if the surface cell contains only a single primitive cell along the x-axis. If the surface cell is however enlarged to two or three primitive cells along the x-axis the same slow convergence behavior as for the adsorption energy is observed.

Furthermore, the LT-MP2 correlation energy per primitive surface cell appears to be converged with respect to the x-axis extent (or k-point sampling) if only few primitive cells along the y-axis are considered: 2×2 and 3×2 differ by 30 meV and 2×3 and 3×3 only by 3 meV. However, this observation breaks down for the 2×4 and 3×4 system pair where the energy difference suddenly increases dramatically and reaches the scale of 500 meV. Since no similar observation was made on the HF level, these results require further attention to verify that they are not caused by the approximations made in the LT-MP2.

The first possible source of errors is the quadrature scheme. As described in section 4.3 the interval spanned by the energy denominators in the MP2 governs the accuracy of the minimax scheme for a fixed number of grid points. The HF eigenvalues determining this interval are listed in Table 6.2. These eigenvalues do not exhibit any notable change with system size. As a consequence it is safe to assume that the quadrature scheme exhibits an equal performance in all system sizes.

The other two parameters in the LT-MP2 affecting the accuracy are the integral screening threshold and the radius of the local auxiliary basis set (ABS) for the generalized RI-LVL. (see section 4.5) As a test case for further investigation I picked the 2×4 and 3×4 cell pair, because the 210 atoms in the 3×4 cell are about the largest system the canonical MP2 can handle with reasonable effort. For these two systems I obtained reference results with the canonical MP2 and repeated the LT-MP2 calculations with a tighter screening threshold 10^{-10} . The results of this comparison are listed in Table 6.3. As it is evident from these results, the same jump in the energy per primitive surface cell also occurs when the canonical MP2 implementation is used. The observed discrepancy in the energy per primitive surface cell is thus not an artifact caused by the approximations made in the LT-MP2 implementation. Instead, this comparison proves the reliability of the chosen parameters for calculations with large unit cells. The

cell size	selected eigenvalues ϵ in eV		
	lowest frozen-core ϵ	highest ϵ	HOMO-LUMO gap
1×2	-38.329	1420.480	9.429
1×3	-38.355	1420.387	9.633
1×4	-38.359	1420.376	9.582
1×5	-38.358	1420.376	9.477
1×6	-38.358	1420.375	9.443
1×7	-38.358	1420.375	9.436
2×2	-38.273	1420.499	9.398
2×3	-38.437	1420.318	9.936
2×4	-38.433	1420.290	9.969
2×5	-38.430	1420.289	9.902
2×6	-38.431	1420.287	9.879
2×7	-38.430	1420.288	9.878
3×2	-38.273	1420.498	9.404
3×3	-38.437	1420.317	9.939
3×4	-38.433	1420.289	9.978
3×5	-38.435	1420.287	9.928

Table 6.2: Special HF eigenvalues which determine the range spanned by the energy denominators in the frozen-core MP2.

MP2 correlation per primitive surface cell in eV	surface cell dimensions	
	2×4	3×4
LT-MP2 with 10^{-9} screening	-104.238	-103.764
LT-MP2 with 10^{-10} screening	-104.215	-103.773
canonical MP2	-104.210	-103.782

Table 6.3: Energy per primitive surface cell in TiO_2 for different surface cell sizes. All calculations contained three layers.

implementation	surface cell size	# of cores	walltime in hours	persistent memory (GB)
canonical MP2	3×4	1280	14.70	1800.7
LT-MP2 (10^{-9})	3×4	1024	10.25	953.4
LT-MP2 (10^{-10})	3×4	1024	14.21	953.4
LT-MP2 (10^{-9})	3×5	1024	16.34	1483.0

Table 6.4: Performance comparison between canonical MP2 and LT-MP2 for a TiO_2 surface with three layers. Values in parentheses denote the screening thresholds used for the LT-MP2. The persistent memory requirement is the memory required to store the persistent data of the calculation like the RI coefficients and the Coulomb matrix. These are the dominating contributions to the total memory requirement. Additional memory needed for communication buffers and intermediate results is not included. The calculations have been performed on different clusters with different processor models, thus the timings cannot be compared exactly with each other.

calculation with 10^{-9} as screening threshold yields errors on the scale of 1.5 meV/atom and the one with 10^{-10} reduces the error to about 0.5 meV/atom. A considerably larger error of 5.75 meV/atom would have been sufficient to recover 99.9% of the MP2 correlation energy in these systems. My LT-MP2 implementation has thus no problem to provide the desired level of accuracy even for condensed matter systems with more than 200 non-hydrogen atoms. It can also be concluded that the investigated supercell sizes (and their corresponding k-grids) are not yet large enough to obtain converged adsorption energies in the low coverage limit.

6.3 PERFORMANCE

The calculation for the $3 \times 4 \times 3$ system contains 216 atoms and 5616 basis functions. The canonical MP2 implementation in FHI-aims requires about 15 hours after finishing the HF Self-Consistent Field (SCF) cycle to calculate the MP2 correlation energy for this system with 1280 cores. As shown in chapter 2, the canonical MP2 has a $\mathcal{O}(N^5)$ scaling with respect to system size. With the optimistic assumption of a perfect parallel scalability, it would be necessary to at least triple the computational resources to finish the $3 \times 5 \times 3$ system in a comparable amount of time. This steep increase of the required computational resources again emphasizes the need for lower-scaling techniques with efficient parallelization, like the combination of RI-LVL and the LT. Table 6.4 shows the memory and computational time consumed by both implementations for this surface size. The timing results cannot be compared exactly because the calculations were performed on different computing clusters, but they nevertheless give a good estimate of the performance differences between the




Figure 6.5: Computational scaling with system size for the dominant steps in the LT-MP2 implementation. All calculations used one cell along the x-axis and three layers.

two implementations. In addition, the table also presents the LT-MP2 results for the 3×5 system, which the canonical MP2 cannot handle with reasonable effort. As it is evident from the results, the LT-MP2 reduces the computational time considerably for these large systems. More importantly, the required memory is significantly reduced in the LT-MP2. For the LT-MP2 the values displayed in the table are the minimum amount of memory needed to store a complete copy of the persistent data. Since the computing nodes on the used cluster provided considerably more memory than needed, the calculations employed two computational domains (see section 4.8) to use the surplus memory to reduce the communication effort.

Finally, Figure 6.5 demonstrates that the superior computational scaling of the LT-MP2 observed for molecular systems also applies to periodic systems. The quasi linear-scaling in the generation of the bare RI coefficients occurs because the generation of the crystalline RI coefficients involves the generation of atomic RI coefficient blocks with atoms outside of the unit cell. Thus the number of required atomic RI coefficient blocks does not increase quadratically with system size.

It is also interesting to compare to the DEC-MP2 by Kjærsgaard and co-workers, which is another lower-scaling MP2 with massive parallelization support. [44] In their paper, they used 1-aza-adamantane-trione (AAT) supramolecular wires to demonstrate the scalability and power of their method. The system with ten monomers (no periodic boundary conditions) contains 6110 basis functions and is thus on a similar scale as the previously discussed TiO₂ surfaces. The 3×4 system has 5616 basis functions and 3×5 contains 7020 basis functions. A direct quantitative comparison between these calculations is of course not feasible, but a few interesting qualitative conclusions

implementation	system (basis functions)	# of cores	walltime in hours
DEC-MP2 [44]	AAT ₁₀ (6110)	59808	10.94
LT-MP2 (10 ⁻⁹)	TiO ₂ 3 × 4 (5616)	1024	10.25
LT-MP2 (10 ⁻⁹)	TiO ₂ 3 × 5 (7020)	1024	16.34

Table 6.5: Performance comparison between DEC-MP2 of Kjærgaard and co-workers and my LT-MP2 implementation. Details about the comparison and the points to consider when comparing the numbers from these very different systems can be found in the text.

can nevertheless be drawn. The calculations presented in this chapter have been performed on the Draco and EOS clusters of the Max-Planck society, while the DEC-MP2 calculations were performed on the Titan supercomputer. The CPUs used on the Draco cluster (Intel Xeon E5-2698 v3 with 16 cores) provide a considerably higher computational performance per core than those on Titan (AMD Opteron 6274 with 16 cores). Titan is however a GPU-focused supercomputer and each CPU is complemented by a NVIDIA Tesla K20X accelerator card. Each K20X card provides a theoretical peak performance of 1.3 Tflops [138], while Draco’s Xeon E5-2698 v3 CPUs (16 cores) are reported to have a theoretical peak performance of 26.5 Gflops per core. [139] For the sake of a rough comparison it can thus safely be assumed that the calculations on Titan have a computational power per core available that is at least equal (probably higher) to the calculations presented in this chapter. Table 6.5 lists the timing comparisons for these two methods. The finite, one-dimensional wire-system composed of non-covalently bonded molecules certainly features a lot more sparsity than the two-dimensional TiO₂ surface with periodic boundary conditions investigated in this study. The AAT wire system is thus a more favorable system for sparsity exploits and lower-scaling methods relying on them. Given these observations, it is remarkable that my LT-MP2 can reach comparable wall-times in systems that are equally sized (measured in the number of basis functions) with a dramatically lower number of cores. It should be noted that a recent follow-up publication [45] suggested additional modifications to the DEC-MP2 to obtain a speed-up of up to 9.5 which would bring the two implementations closer in terms of performance. However, this modification was only demonstrated numerically for smaller systems with less than 2000 basis functions and an alkane chain with 2812 basis functions and it was estimated that about 99.8% of the MP2 correlation energy was recovered.

This comparison to a state-of-the-art lower-scaling MP2 with massive parallelization support clearly shows that my LT-MP2 implementation is a very promising approach.

OUTLOOK: APPLICATION TO THE RANDOM-PHASE APPROXIMATION

Beyond the common semi-local Density-Functional Approximations (DFAs), the Random-Phase Approximation (RPA) is one of the most frequently used methods for the evaluation of electronic interactions. The exchange contribution of the RPA is the same as the Hartree-Fock (HF) exact exchange, it thus can be directly evaluated efficiently with the RI-LVL method presented in chapter 3. The correlation energy \mathcal{E}_c in the RPA (Equation 7.1, see section 2.7 for a derivation) on the other hand is computationally more involved. [67]

$$E_c = \frac{1}{2\pi} \int_0^\infty d\omega \sum_{n=2}^{\infty} \frac{1}{n} \text{Tr} \left[(\chi^0(i\omega)v)^n \right] \quad (7.1)$$

$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega)$ denotes the independent-particle response function of the Kohn-Sham (KS) system and is known analytically.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \sum_{m,n} \frac{(f_m - f_n) \psi_m^*(\mathbf{r}_1) \psi_n(\mathbf{r}_1) \psi_n^*(\mathbf{r}_2) \psi_m(\mathbf{r}_2)}{\epsilon_m - \epsilon_n - i\omega} \quad (7.2)$$

Here, f_m denotes the occupation number of the state. In this chapter I will demonstrate how the generalized RI-LVL concept and the Laplace-Transformation (LT) presented in the previous chapters can be applied to RPA and the potential benefits of such a combination. I will at first analyze the simpler case of a system with finite band-gap and integer occupation and then outline how the results can be generalized to systems with zero band-gap and fractional occupation numbers. Schurkus and Ochsenfeld already demonstrated the accuracy and benefits of combining the RPA with a LT for molecules with a finite HOMO-LUMO gap in a recent publication. [16] I will present a different derivation here based on imaginary frequencies, which gives rise to a different working equation and avoids the need to use two different evaluation techniques for the high- and low-frequency limit.

7.1 RPA IN AN RI-FRAMEWORK

The pairs of KS states appearing in the response function can be represented with a Resolution of Identity (RI) expansion $\psi_m^*(\mathbf{r}_1) \psi_n(\mathbf{r}_1) = \sum_{\mu} C_{m,n}^{\mu} P_{\mu}(\mathbf{r}_1)$. [67] As usual, this expansion is not unique and no specific choice of the coefficients will be imposed at this time. Using this

expansion, the response function can be represented in the basis of the auxiliary basis functions $\{P_\mu\}$.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \sum_{\mu, \lambda} \chi_{\mu, \lambda}^0(i\omega) P_\mu(\mathbf{r}_1) P_\lambda(\mathbf{r}_2) \quad (7.3)$$

$$\chi_{\mu, \lambda}^0(i\omega) = \sum_{m, n} \frac{(f_m - f_n) C_{m, n}^{P_\mu} C_{n, m}^{P_\lambda}}{\epsilon_m - \epsilon_n - i\omega} \quad (7.4)$$

If this expansion is inserted into Equation 7.1, the implicit spatial integrals in it (see Equation 2.78 and text referencing it) can be evaluated explicitly and yield the Coulomb matrix of the auxiliary basis functions $V_{\mu, \lambda}$. The RPA correlation energy in the RI framework can then be evaluated with Equation 7.5.

$$E_c = \frac{1}{2\pi} \int_0^\infty d\omega \sum_{n=2}^\infty \frac{1}{n} \text{Tr} \left[\left(\sum_{\kappa} \chi_{\mu, \kappa}^0(i\omega) V_{\kappa, \lambda} \right)^n \right] \quad (7.5)$$

$$V_{\mu, \lambda} = \iint \frac{P_\mu(\mathbf{r}_1) P_\lambda(\mathbf{r}_1)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2 \quad (7.6)$$

In this RI variant of the RPA, the computationally most demanding step is the generation of the response function for each point on the frequency grid. [67] From Equation 7.4 it can be seen that $\mathcal{O}(N_{\text{states}}^2)$ contributions to a $\mathcal{O}(N_{\text{aux}}^2)$ matrix need to be computed. The formation of the response matrix at a given frequency thus scales as $\mathcal{O}(N^4)$ and has a $\mathcal{O}(N^3)$ memory requirement to store the RI coefficients. In the following derivation, I will therefore solely focus on the calculation of the frequency-dependent response matrix.

The expression for the response matrix can be manipulated by means of a LT, but it is not as straight-forward as for 2nd Order Møller-Plesset Perturbation Theory (MP2). Formally, the LT $F(s)$ of a function $f(t)$ is defined as

$$F(s) = \int_0^\infty f(t) e^{-st} dt \quad (7.7)$$

where s is a complex variable, and $f(t)$ is a complex-valued function of the real variable t . [140] The function $f(t)$ must be defined for $t \geq 0$ and be integrable over $(0, \infty)$. Furthermore, it must be subject to an exponential growth limitation:

$$|f(t)| \leq K e^{ct} \quad (7.8)$$

with a real constant c . If these conditions hold, the integral in Equation 7.7 is absolutely convergent for $\Re(s) > c$. In the case of $F(s) = \frac{1}{s}$ the original function is given as $f(t) = 1$, thus we have $c = 0$ in this case. [140] The real part of the denominator in Equation 7.2 can be both positive and negative since both sums include all states. This concept therefore cannot be applied directly to the RPA.

7.2 LAPLACE-TRANSFORMED RPA FOR MATERIALS WITH BAND GAP

Let us at first consider the simpler case of a material with a finite band gap and integer occupation numbers, i.e. all orbitals are either fully

occupied or empty. In this scenario we can split the state summations into sums over occupied and virtual states.

$$\sum_m = \sum_i + \sum_a \quad (7.9)$$

Due to the occupation factors, only the cross-terms are non-zero and we obtain:

$$\begin{aligned} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) &= 2 \sum_i \sum_a \frac{\psi_i^*(\mathbf{r}_1) \psi_a(\mathbf{r}_1) \psi_a^*(\mathbf{r}_2) \psi_i(\mathbf{r}_2)}{\epsilon_i - \epsilon_a - i\omega} \\ &\quad - 2 \sum_b \sum_j \frac{\psi_b^*(\mathbf{r}_1) \psi_j(\mathbf{r}_1) \psi_j^*(\mathbf{r}_2) \psi_b(\mathbf{r}_2)}{\epsilon_b - \epsilon_j - i\omega} \end{aligned} \quad (7.10)$$

Under the assumption of the finite band gap $\epsilon_a > \epsilon_i$ holds true for all state combinations. Thus the denominators in each double sum now have a well-defined sign and we can easily bring them into a form suitable for the transformation.

$$\begin{aligned} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) &= -2 \sum_i \sum_a \frac{\psi_i^*(\mathbf{r}_1) \psi_a(\mathbf{r}_1) \psi_a^*(\mathbf{r}_2) \psi_i(\mathbf{r}_2)}{\epsilon_a - \epsilon_i + i\omega} \\ &\quad - 2 \sum_j \sum_b \frac{\psi_j(\mathbf{r}_1) \psi_b^*(\mathbf{r}_1) \psi_b(\mathbf{r}_2) \psi_j^*(\mathbf{r}_2)}{\epsilon_b - \epsilon_j - i\omega} \end{aligned} \quad (7.11)$$

Noting that the second term is just the complex conjugate of the first one, the expression can be simplified.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = -4\Re \left[\sum_i \sum_a \frac{\psi_i^*(\mathbf{r}_1) \psi_a(\mathbf{r}_1) \psi_a^*(\mathbf{r}_2) \psi_i(\mathbf{r}_2)}{\epsilon_a - \epsilon_i + i\omega} \right] \quad (7.12)$$

Having the equation in this form, the same LT as for the MP2 can be applied, except for the fact that we have complex functions this time. As discussed in the previous section, this extension is covered by the definition and convergence requirements of the LT.

$$\frac{1}{\epsilon_a - \epsilon_i + i\omega} = \int_0^\infty e^{-i\omega q} e^{-(\epsilon_a - \epsilon_i)q} dq \quad (7.13)$$

Inserting this expression into Equation 7.12, one obtains

$$\begin{aligned} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) &= -4\Re \left[\int_0^\infty dq \sum_i \sum_a e^{-i\omega q} e^{-(\epsilon_a - \epsilon_i)q} \right. \\ &\quad \left. \times \psi_i^*(\mathbf{r}_1) \psi_a(\mathbf{r}_1) \psi_a^*(\mathbf{r}_2) \psi_i(\mathbf{r}_2) \right] \end{aligned} \quad (7.14)$$

If one now inserts the basis set expansions, the same pseudo-density matrices as in the Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP2) are obtained and the generalized RI-LVL could be used straight away to reduce the formation of the response function to a quadratic step with a preceding cubic coefficient generation step.

However, this formulation has a serious drawback that disqualifies it for practical usage, namely the type of the dependency on ω . The

frequency only occurs in an oscillatory prefactor and thus this expression does not vanish at large frequencies. As it has been shown in a recent publication by the Ochsenfeld group [16] such terms cause rapid oscillations in the high-frequency limit and make it very difficult to obtain converged results with a quadrature scheme for the LT.

A computationally more stable form can be obtained by modifying Equation 7.12 a bit before the transformation is applied.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = -4\Re \left[\sum_i \sum_a \frac{(1-i)\psi_i^*(\mathbf{r}_1)\psi_a(\mathbf{r}_1)\psi_a^*(\mathbf{r}_2)\psi_i(\mathbf{r}_2)}{(1-i)(\epsilon_a - \epsilon_i + i\omega)} \right] \quad (7.15)$$

$$= -4\Re \left[\sum_i \sum_a \frac{(1-i)\psi_i^*(\mathbf{r}_1)\psi_a(\mathbf{r}_1)\psi_a^*(\mathbf{r}_2)\psi_i(\mathbf{r}_2)}{\epsilon_a - \epsilon_i + \omega - i(\epsilon_a - \epsilon_i - \omega)} \right] \quad (7.16)$$

As can be seen from Equation 7.1, the response function is only evaluated for $\omega \geq 0$, thus the frequency component in the real part does not affect the validity of the LT. Applying the same transformation as before, we obtain

$$\begin{aligned} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = -4\Re \left[\int_0^\infty dq \sum_i \sum_a e^{-\omega q} e^{-i\omega q} e^{-\epsilon_a q} e^{\epsilon_i q} e^{i\epsilon_a q} e^{-i\epsilon_i q} \right. \\ \left. \times (1-i)\psi_i^*(\mathbf{r}_1)\psi_a(\mathbf{r}_1)\psi_a^*(\mathbf{r}_2)\psi_i(\mathbf{r}_2) \right] \quad (7.17) \end{aligned}$$

This formulation has the advantage that it features an exponential decay with the frequency and thus avoids any problems caused by the oscillatory factor at high frequencies. Since the expression is numerically well-behaved for both the high- and low-frequency limit, it is not necessary to adopt different evaluation schemes for these limits as it has been done in the work by Schurkus and Ochsenfeld. [16] After inserting the basis set expansion $\psi_m(\mathbf{r}_1) = \sum_s c_{m,s}\varphi_s(\mathbf{r}_1)$ into the formula, we can obtain new pseudo-density matrices. In the formulas below, the Fermi-energy ϵ_F has been added into the equations for numerical reasons discussed later.

$$\begin{aligned} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = -4\Re \left[\int_0^\infty dq \sum_{s,t,u,v} \sum_i \sum_a e^{-\omega q} e^{-i\omega q} e^{-\epsilon_a q} e^{\epsilon_i q} e^{i\epsilon_a q} e^{-i\epsilon_i q} \right. \\ \left. \times (1-i)c_{i,s}^*c_{a,t}c_{a,u}^*c_{i,v}\varphi_s(\mathbf{r}_1)\varphi_t(\mathbf{r}_1)\varphi_u(\mathbf{r}_2)\varphi_v(\mathbf{r}_2) \right] \quad (7.18) \end{aligned}$$

$$\begin{aligned} = -4\Re \left[\int_0^\infty dq \sum_{s,t,u,v} \sum_i \sum_a e^{-\omega q} e^{-i\omega q} \right. \\ \left. \times e^{-(\epsilon_a - \epsilon_F)q} e^{(\epsilon_i - \epsilon_F)q} e^{i(\epsilon_a - \epsilon_F)q} e^{-i(\epsilon_i - \epsilon_F)q} \right. \\ \left. \times (1-i)c_{i,s}^*c_{a,t}c_{a,u}^*c_{i,v}\varphi_s(\mathbf{r}_1)\varphi_t(\mathbf{r}_1)\varphi_u(\mathbf{r}_2)\varphi_v(\mathbf{r}_2) \right] \quad (7.19) \end{aligned}$$

$$\begin{aligned} = -4\Re \left[\int_0^\infty dq \sum_{s,t,u,v} e^{-\omega q} e^{-i\omega q} \right. \\ \left. \times (1-i)X_{s,v}^q Y_{t,u}^q \varphi_s(\mathbf{r}_1)\varphi_t(\mathbf{r}_1)\varphi_u(\mathbf{r}_2)\varphi_v(\mathbf{r}_2) \right] \quad (7.20) \end{aligned}$$

where the pseudo-densities $X_{s,v}$ and $Y_{t,\mu}$ are now defined as:

$$X_{s,v}^q = \sum_i c_{i,s}^* e^{(\epsilon_i - \epsilon_F)q} e^{-i(\epsilon_i - \epsilon_F)q} c_{i,v} \quad (7.21)$$

$$Y_{t,\mu}^q = \sum_a c_{a,t} e^{-(\epsilon_a - \epsilon_F)q} e^{i(\epsilon_a - \epsilon_F)q} c_{a,\mu}^* \quad (7.22)$$

Compared to the expressions used in the LT-MP2, the pseudo-density matrices now contain an additional phase factor. For eigenvalues far away from the Fermi-level this factor causes rapid oscillations in the Laplace quadrature. Since the terms also decay exponentially with the distance to the Fermi-level, this is however not a problem for the numerical quadrature.

Finally, two sums over basis functions can be contracted with the pseudo-densities to obtain an expression that bears some similarities to the LT-MP2 formulas discussed in previous chapters.

$$\begin{aligned} \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = & -4\Re \left[\int_0^\infty dq e^{-\omega q} e^{-i\omega q} (1-i) \sum_{u,v} \right. \\ & \left. \times \varphi_{\bar{v}}(\mathbf{r}_1) \varphi_{\underline{u}}(\mathbf{r}_1) \varphi_u(\mathbf{r}_2) \varphi_v(\mathbf{r}_2) \right] \quad (7.23) \end{aligned}$$

The function products $\varphi_{\bar{v}}(\mathbf{r}_1) \varphi_{\underline{u}}(\mathbf{r}_1)$ and $\varphi_u(\mathbf{r}_2) \varphi_v(\mathbf{r}_2)$ can now be represented with an RI approximation to obtain a computationally convenient form.

$$\varphi_{\bar{v}}(\mathbf{r}_1) \varphi_{\underline{u}}(\mathbf{r}_1) \approx \sum_{\mu} M_{\bar{v},\underline{u}}^{\mu} P_{\mu}(\mathbf{r}_1) \quad (7.24)$$

$$\varphi_u(\mathbf{r}_2) \varphi_v(\mathbf{r}_2) \approx \sum_{\lambda} C_{t,v}^{\lambda} P_{\lambda}(\mathbf{r}_2) \quad (7.25)$$

$$\begin{aligned} \Rightarrow \chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = & -4\Re \left[\sum_{\mu,\lambda} \int_0^\infty dq e^{-\omega q} e^{-i\omega q} (1-i) \sum_{u,v} \right. \\ & \left. \times M_{\bar{v},\underline{u}}^{\mu} C_{u,v}^{\lambda} P_{\mu}(\mathbf{r}_1) P_{\lambda}(\mathbf{r}_2) \right] \quad (7.26) \end{aligned}$$

Comparing to Equation 7.3, the RI representation of the independent-particle response function in the LT approach is now found.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = -4\Re \left[\sum_{\mu,\lambda} \int_0^\infty dq e^{-\omega q} e^{-i\omega q} \chi_{\mu,\lambda}^0(q) P_{\mu}(\mathbf{r}_1) P_{\lambda}(\mathbf{r}_2) \right] \quad (7.27)$$

$$\chi_{\mu,\lambda}^0(q) = (1-i) \sum_{u,v} M_{\bar{v},\underline{u}}^{\mu} C_{u,v}^{\lambda} \quad (7.28)$$

The $C_{u,v}^{\lambda}$ coefficients are for pure atomic basis functions, thus the RI-LVL as presented in chapter 3 can be used to accurately represent them. The transformed coefficients $M_{\bar{v},\underline{u}}^{\mu}$ only differ by the phase factors in the pseudo-density matrices from the ones encountered in the LT-MP2. The generalized RI-LVL used in the LT-MP2 should thus have no problems representing these entities.

If the integral over q is finally replaced by a numerical quadrature, another important feature besides the possibility to use local RI techniques becomes evident.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = -4\Re \left[\sum_{\mu, \lambda} \int_0^\infty dq e^{-\omega q} e^{-i\omega q} \chi_{\mu, \lambda}^0(q) P_\mu(\mathbf{r}_1) P_\lambda(\mathbf{r}_2) \right] \quad (7.29)$$

$$\approx -4\Re \left[\sum_q w_q e^{-\omega q} e^{-i\omega q} \sum_{\mu, \lambda} \chi_{\mu, \lambda}^0(q) P_\mu(\mathbf{r}_1) P_\lambda(\mathbf{r}_2) \right] \quad (7.30)$$

For each frequency, the RI representation of $\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega)$ can be build from a set of frequency-independent building blocks $\chi_{\mu, \lambda}^0(q)$ by means of simple matrix additions. If enough memory is available to store an $\mathcal{O}(N_{\text{aux}}^2)$ matrix for each point of the quadrature used to represent the LT, the effort for each ω point can be reduced considerably by pre-computing these blocks once.

If no further approximations are used, this would give rise to a formally cubic scaling RPA. The generation for each $\chi_{\mu, \lambda}^0(q)$ in Equation 7.28 scales quadratically since both auxiliary basis sets feature *a priori* enforced locality, i.e. each basis function pair only updates a system-size independent sub-block of the response matrix. The generation of the transformed RI coefficients on the other hand is scaling cubically with system size as discussed in section 4.6. It should however be noted that the transformed RI coefficients are only needed in Equation 7.28. Since the atomic RI coefficients $C_{u, v}^\lambda$ are strictly zero if the involved functions have no overlap, the transformed RI coefficients $M_{\bar{v}, \bar{u}}^\mu$ thus only need to be computed if the corresponding atomic coefficients are non-zero. Exploiting this feature leads to a asymptotically $\mathcal{O}(N^2)$ scaling with system size for the last contraction step and thus an overall asymptotic $\mathcal{O}(N^2)$ scaling for the contraction step. Additionally, the sparsity in the pseudo-density matrices can be used during the coefficient generation to further reduce the computational effort of this step.

The memory requirements of such an RPA implementation would also scale quadratically with system size since all involved matrices are $\mathcal{O}(N_{\text{aux}}^2)$ and the involved RI coefficients have been demonstrated to be quadratic scaling in the previous chapters. The combination of the generalized RI-LVL with Laplace-Transformation techniques is therefore also a very promising approach for other electronic structure theory methods besides the MP2.

7.3 LAPLACE-TRANSFORMED RPA WITH FRACTIONAL OCCUPATION

In the previous section the application of the LT to the RPA was outlined under the constraint of a finite band gap and integer occupation. If these constraints are lifted, the evaluation in the LT framework be-

comes a lot more involved. The sums over states can now be splitted into three contributions:

$$\sum_m = \sum_i + \sum_o^{\text{partial}} + \sum_a \quad (7.31)$$

where the middle term includes all states with fractional occupation. This results in nine different subsets we have to sum over. As in the previous case, the occupied-occupied and virtual-virtual blocks are zero. The occupied-virtual sums can be evaluated in the same way as for the band-gap case. Mixed terms that involve one sum over partially occupied states can be evaluated in a similar way since we know from the Fermi-Dirac Distribution [118] that all partially filled states will have eigenvalues that are higher than those of fully occupied orbitals and smaller than those of empty orbitals. Thus the denominators in each of these terms have a consistent sign for all terms of the sums. Since eigenstates with the same eigenvalue have the same occupation number according to the Fermi-Dirac distribution, all terms with equal eigenvalues are exactly zero due to the occupation number difference in Equation 7.2. Thus only terms with a non-zero eigenvalue difference must be taken into account for the LT. Apart from an fractional occupation number as additional factor, each of these terms follows a similar derivation as shown in the previous section.

The most challenging term is the one where both sums iterate over partially occupied states.

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \dots + \sum_o^{\text{partial}} \sum_p^{\text{partial}} (f_o - f_p) \frac{\psi_o^*(\mathbf{r}_1) \psi_p(\mathbf{r}_1) \psi_p^*(\mathbf{r}_2) \psi_o(\mathbf{r}_2)}{\epsilon_o - \epsilon_p - i\omega} \quad (7.32)$$

In contrast to the other contributions, the denominators in this summation have an indefinite sign, which prevents a direct application of the LT. A manipulation of the equation however can help to address this problem:

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \dots + \sum_o^{\text{partial}} \sum_p^{\text{partial}} i(f_o - f_p) \frac{\psi_o^*(\mathbf{r}_1) \psi_p(\mathbf{r}_1) \psi_p^*(\mathbf{r}_2) \psi_o(\mathbf{r}_2)}{i(\epsilon_o - \epsilon_p - i\omega)} \quad (7.33)$$

$$= \dots + \sum_o^{\text{partial}} \sum_p^{\text{partial}} i(f_o - f_p) \frac{\psi_o^*(\mathbf{r}_1) \psi_p(\mathbf{r}_1) \psi_p^*(\mathbf{r}_2) \psi_o(\mathbf{r}_2)}{\omega + i(\epsilon_o - \epsilon_p)} \quad (7.34)$$

The indefinite sign of the eigenvalues has now been moved into the imaginary part of the denominator, where it does not affect the validity of the LT. After applying the usual transformation, one obtains:

$$\chi^0(\mathbf{r}_1, \mathbf{r}_2, i\omega) = \dots + \sum_o^{\text{partial}} \sum_p^{\text{partial}} i(f_o - f_p) \int_0^\infty dq e^{-\omega q} e^{-i(\epsilon_o - \epsilon_p)q} \psi_o^*(\mathbf{r}_1) \psi_p(\mathbf{r}_1) \psi_p^*(\mathbf{r}_2) \psi_o(\mathbf{r}_2) \quad (7.35)$$

This expression decays for large frequencies, but it has the problem that the coupling term between eigenvalues and q is purely oscillatory, which could cause quadrature convergence problems if the eigenvalues become large. Fortunately, the summations under consideration only include partially occupied states near the Fermi level. By adding a constructive zero into the exponent, we can ensure that all terms in this summation will oscillate slowly and thus not cause any problems for the convergence. (See also the discussion in Schurkus' work about the high- and low-frequency limit integrations in Ref [16] for an example of the impact rapid oscillatory terms can have on the LT.)

$$\begin{aligned}
 &= \dots + \sum_o^{\text{partial}} \sum_p^{\text{partial}} i(f_o - f_p) \\
 &\int_0^\infty dq e^{-\omega q} e^{-i([\epsilon_o - \epsilon_F] - [\epsilon_p - \epsilon_F])q} \psi_o^*(\mathbf{r}_1) \psi_p(\mathbf{r}_1) \psi_p^*(\mathbf{r}_2) \psi_o(\mathbf{r}_2) \quad (7.36)
 \end{aligned}$$

This expression can now be manipulated further in the same spirit as shown above to obtain an RI expression for it. Compared to the case with a finite band-gap two potential problems arise: the magnitude of the denominators and the validity of the LT for $\omega = 0$ in the double-summation over partially occupied states. As discussed before, the Fermi-Dirac distribution ensures that only terms with non-zero eigenvalue difference must be considered explicitly, thus the denominator is always non-zero, but the closer the magnitude is to zero, the more challenging it becomes to find an accurate quadrature representation. It is therefore likely necessary to use tighter quadrature grids for the LT to retain the accuracy. The second issue is region of convergence of the LT in the double-summation over partially occupied states. As discussed before, the convergence of the integral is only guaranteed if the real part, i.e. ω , is larger than zero. At this stage, it is not clear if the LT can nevertheless handle this particular term or if a different evaluation strategy is needed. The application of the LT and localized RI strategies to systems with fractional occupation is therefore an open question that requires further investigation.

CONCLUSIONS

Surfaces are an often studied class of systems due to their manifold applications. For an accurate modeling of the electronic structure in these systems it is often necessary to go beyond Density-Functional Theory (DFT) with semi-local approximations. Among the promising approaches for the treatment of such systems are advanced correlation methods which treat the electronic interactions explicitly in a many-body picture. Popular representatives from this class of methods are e.g. 2nd Order Møller-Plesset Perturbation Theory (MP2), Random-Phase Approximation (RPA) and Coupled-Cluster Theory (CC). All these methods share the severe limitation that they have a high-order scaling with system size for both the computational time and memory requirement in their canonical formulations. A variety of lower-scaling methods has been proposed in the literature to address this problem. Until now most published lower-scaling algorithms only provide a very limited level parallelization, e.g. by using threaded linear algebra libraries.

To further extend the realm of application for advanced correlation methods it is imperative to develop lower-scaling techniques that support large-scale parallelism. In my thesis I developed a new technique to reach this goal for MP2 and RPA. Its key ingredients are a new localized Resolution of Identity (RI) strategy, named RI-LVL, and the Laplace-Transformation (LT). The LT removes the coupling of different orbitals through the eigenvalue differences in the denominators of the canonical formulations. After this transformation, the delocalized canonical molecular orbitals can be recast into a set of localized functions, which is the prerequisite for using localized RI strategies. The RI-LVL significantly reduces the required memory to a $\mathcal{O}(N^2)$ scaling with system size and its sparsity patterns are known *à priori*. This *à priori* known sparsity is an important feature which allows the design of an efficient (completely in-memory) data distribution and communication pattern in a large-scale parallelized code. This new technique can be complemented with established integral screening techniques from quantum chemistry to preselect the significant tuples of the localized functions and thus dramatically reduce the computational scaling. This scaling reduction however comes at the price of having *irregular* parallelism, i.e. the algorithm can no longer be formulated as a predictable sequence of global computation and communication phases. To obtain a good parallel performance one instead needs a dynamic load balancing and a flexible data communication framework. Taking advantage of the memory savings from the RI-LVL strategy, efficient hybrid OpenMP/MPI parallelization schemes can be developed to address this problem and maintain a high parallel efficiency while enjoying the benefits of the integral screening. In addition to discussing these concepts in detail at the example of the

MP2, I also explained how the same technique can be applied to the RPA.

To prove the practical applicability of my new technique, I implemented a Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP2) based on the Numeric Atom-Centered Orbital (NAO) framework of the Fritz Haber Institute ab initio molecular simulations (FHI-aims) code. For periodic systems the combination of RI-LVL and the LT has been paired with a supercell approach to take the maximum benefit from the integral screening. I demonstrated the accuracy of the resulting LT-MP2 algorithm for water clusters with up to 70 molecules, isomerization energies in the ISOL22 test set and also for titanium dioxide (TiO_2) surfaces with up to 210 atoms. In all cases, a very high accuracy was reached and more than 99.9% of the canonical MP2 correlation energy was recovered. The scaling of my new implementation has been investigated in detail with a series of water clusters ranging from 10 to 150 molecules. My LT-MP2 features an at most cubic scaling in all computational steps and the memory requirement scales quadratically with system size. The new algorithm breaks even with the canonical implementation both with respect to memory and time at reasonable system sizes in periodic and molecular systems alike. Furthermore I demonstrated the near perfect parallel scalability of the new algorithm with the number of cores for the irregular parallel problem.

I also used my new LT-MP2 implementation to investigate the convergence of the adsorption energy of water on rutile TiO_2 in the low coverage limit. It was found that the MP2 correlation converges very slowly with the size of the surface unit cell. In contrast, all investigated mean-field methods (HF, HSE06 and rPBE) showed a considerably faster and smoother convergence with system size. Within the system sizes accessible with the canonical MP2 implementation, it was also confirmed that the reference implementation yields the same results. The lower scaling features of the LT-MP2 considerably extend the range of system sizes that can be accessed with a reasonable commitment of computational resources. Unfortunately however, no final converged value for the MP2 correlation contribution to the adsorption energy could be obtained with the surface sizes investigated in this project. This slow convergence of the adsorption energy clearly emphasizes the need for large-scale parallel and lower scaling advanced correlation methods for the study of interactions at surfaces. The qualitative comparison to a state-of-the-art massively parallel lower-scaling MP2 method, the Divide-Expand-Consolidate 2nd Order Møller-Plesset Perturbation Theory (DEC-MP2), reveals that my LT-MP2 is capable of handling much more condensed systems with a similar number of basis functions with a considerably lower resource commitment and similar wall-times.

Despite its already impressive performance, the LT-MP2 implementation presented in this thesis still has a lot of potential for future improvements. The sparsity introduced by the integral screening is not yet fully exploited in all parts of the implementation, for example in the data communication. Apart from general numerical optimiza-

tions, in particular the implementation for periodic systems can still be improved on the conceptual level. In the present implementation, the underlying symmetry in the supercell is not yet used explicitly and the usage of not Γ -centered k-grids might also speed up the convergence. Other very exciting prospects are the usage of the presented LT-MP2 implementation in double-hybrid functionals and the application of this new lower-scaling technique to the RPA. As I showed in the previous chapter, applying the LT and RI-LVL to the RPA leads to expressions which are very similar to the corresponding expressions in the MP2. It is thus to be expected that the RPA will yield the same high accuracy as the LT-MP2 and the computational scaling is predicted to become cubically at worst.

SUPPLEMENTARY DATA FOR RI-LVL ACCURACY ANALYSIS

In chapter 3 the RI-V calculations with the standard auxiliary basis set (ABS) are used as reference to benchmark the accuracy of the RI-LVL. In this appendix, the effect (or lack thereof) of the additional functions in the enhanced orbital basis set (OBS+) in RI-V calculations is demonstrated for the S22 test set.

As in the RI-LVL presentation, the OBS+ is hierarchically improved by adding additional functions of increasing angular momentum. The following tables show how the Root-Mean Square Deviation (RMSD) and maximum absolute error per atom change with respect to the calculation without OBS+ for Hartree-Fock (HF), MP2, PBE0 and the RPA.

For all methods, even the largest OBS+ does not cause any significant changes in the total energy. The RI-V is thus well converged with respect to the ABS and the standard ABS can be used as reference to judge the accuracy of RI-LVL.

<i>HF RI-V</i>	$ \Delta E_{tot} $ (meV/atom)	
	RMSD	MAX
OBS+		
<i>tier2 + s(z=1)</i>	0.005	0.016
<i>tier2 + sp(z=1)</i>	0.009	0.024
<i>tier2 + spd(z=1)</i>	0.011	0.020
<i>tier2 + spdf(z=1)</i>	0.023	0.064
<i>tier2 + spdfg(z=1)</i>	0.154	0.367

Table A.1: RMSD and maximum absolute errors for HF calculations in the S22 test set using the RI-V method with tier2 basis sets. The reference point is the RI-V calculation with no additional functions used to construct the ABS.

<i>MP2@HF RI-V</i>	$ \Delta E_{tot} $ (meV/atom)	
	RMSD	MAX
OBS+		
<i>NAO-VCC-3Z + s(z=1)</i>	0.003	0.008
<i>NAO-VCC-3Z + sp(z=1)</i>	0.013	0.045
<i>NAO-VCC-3Z + spd(z=1)</i>	0.012	0.036
<i>NAO-VCC-3Z + spdf(z=1)</i>	0.011	0.027
<i>NAO-VCC-3Z + spdfg(z=1)</i>	0.031	0.098

Table A.2: Total energy per atoms errors for MP2@HF calculations using NAO-VCC-3Z basis sets with RI-V and different choices for the ABS. The reference point is the calculation with no extra functions for the ABS construction.

In subsection 3.8.1 it was also stated that the ABS enhancements do not have any significant impact on the accuracy for the heavier ele-

<i>PBE0 RI-V</i>	$ \Delta E_{tot} $ (meV/atom)	
	RMSD	MAX
<i>OBS+</i>		
<i>tier2 + s(z=1)</i>	0.001	0.002
<i>tier2 + sp(z=1)</i>	0.002	0.005
<i>tier2 + spd(z=1)</i>	0.003	0.009
<i>tier2 + spdf(z=1)</i>	0.005	0.011
<i>tier2 + spdfg(z=1)</i>	0.031	0.083

Table A.3: RMSD and maximum absolute errors for PBE0 calculations in the S22 test set using the RI-V method with tier2 basis sets. The reference point is the RI-V calculation with no additional functions used to construct the ABS.

<i>RPA@PBE0 RI-V</i>	$ \Delta E_{tot} $ (meV/atom)	
	RMSD	MAX
<i>OBS+</i>		
<i>NAO-VCC-3Z + s(z=1)</i>	0.006	0.018
<i>NAO-VCC-3Z + sp(z=1)</i>	0.025	0.096
<i>NAO-VCC-3Z + spd(z=1)</i>	0.018	0.063
<i>NAO-VCC-3Z + spdf(z=1)</i>	0.023	0.086
<i>NAO-VCC-3Z + spdfg(z=1)</i>	0.032	0.098

Table A.4: Total energy per atoms errors for RPA@PBE0 calculations using NAO-VCC-3Z basis sets with RI-V and different choices for the ABS. The reference point is the calculation with no extra functions for the ABS construction.

ments copper, titanium and gold. Table A.6 shows the the RMSD values for the RI-LVL RPA@PBE0 calculations in all three cluster groups for a tier1 and the largest employed basis set. We also found that it is quite challenging for the RI-V to generate a sufficiently accurate auxiliary basis Coulomb matrix to avoid numerical noise in the inversion step for very large basis sets. Thus the slight deterioration of the RMSD in the tier3 calculations for gold clusters is likely caused by the reference, not the RI-LVL calculations.

RPA@PBE0 RI-LVL	$ \Delta E_{tot} $ (meV/atom)	
	OBS+	
	RMSD	MAX
NAO-VCC-3Z	9.119	17.611
NAO-VCC-3Z + <i>spdfg</i> (<i>z=1</i>)	0.041	0.130
NAO-VCC-3Z + <i>g</i> (<i>z=1</i>)	0.059	0.385
NAO-VCC-3Z + <i>g</i> (<i>z=2</i>)	0.060	0.383
NAO-VCC-3Z + <i>g</i> (<i>z=3</i>)	0.059	0.387
NAO-VCC-3Z + <i>g</i> (<i>z=4</i>)	0.058	0.390
NAO-VCC-3Z + <i>g</i> (<i>z=5</i>)	0.059	0.406
NAO-VCC-3Z + <i>g</i> (<i>z=6</i>)	0.058	0.417
NAO-VCC-3Z + <i>g</i> (<i>z=7</i>)	0.057	0.402

Table A.5: Total energy errors for RPA@PBE0 calculations using RI-LVL and NAO-VCC-3Z basis sets with different choices for the ABS. The reference is the RI-V calculations using the standard ABS.

cluster	RPA@PBE0 RI-LVL	$ \Delta E_{tot} $ (meV/atom)
	OBS+	
Au	<i>tier1</i>	0.353
	<i>tier1</i> + <i>spdfg</i> (<i>z=1</i>)	0.401
	<i>tier3</i>	1.271
	<i>tier3</i> + <i>spdfg</i> (<i>z=1</i>)	1.387
Cu	<i>tier1</i>	0.010
	<i>tier1</i> + <i>spdfg</i> (<i>z=1</i>)	0.092
	<i>tier4</i>	0.107
	<i>tier4</i> + <i>spdfg</i> (<i>z=1</i>)	0.087
TiO ₂	<i>tier1</i>	0.096
	<i>tier1</i> + <i>spdfg</i> (<i>z=1</i>)	0.038
	<i>tier3</i>	0.076
	<i>tier3</i> + <i>spdfg</i> (<i>z=1</i>)	0.089

Table A.6: Influence of the ABS set on the accuracy of RPA@PBE0 calculations using the RI-LVL strategy. The reference points are the RI-V calculations with the standard ABS and very tight integration settings. For the titanium dioxide test-cases an OBS+ with an additional *g* function on oxygen was used. Only the OBS+ on the titanium atoms was changed.

PARAMETERS FOR THE LT-MP₂

This appendix provides more details about the impact the parameters of the Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory (LT-MP₂) have on the final accuracy. A compact summary of these results is presented in section 5.1. This chapter additionally contains results for the NAO-VCC-3Z basis set to prove that the LT-MP₂ implementation also works for larger basis sets. The test systems considered throughout this appendix are a small cluster of 10 water molecules as example for sparse three-dimensional systems and a fully extended oligo-alanine chain with 3 segments as covalently bonded example system.

B.1 QUADRATURE AND INTEGRAL SCREENING

At first the impact of the Laplace-Transformation (LT) and the integral screening needs to be assessed. Figure B.1 and B.2 show how the number of quadrature points and the integral screening threshold influence the error per atom in the MP₂ correlation energy for NAO-VCC-2Z and NAO-VCC-3Z basis sets.

These calculations use only the QQZZ screening, i.e. a true upper bound estimate is obtained for each ERI. Furthermore a RI-LVL_{full} based implementation was used for these tests as described in section 4.4. Thus the transformed RI coefficients are not restricted to local subsets of the global auxiliary basis set (ABS) and the approximations made in the generalized RI-LVL are thus not applied here. The errors shown in these plots are thus only originating from the quadrature and the QQZZ integral screening.

As it can be seen from these results, a screening threshold of 10^{-9} is necessary to ensure well converged results with NAO-VCC-2Z basis sets. To obtain a similar accuracy for NAO-VCC-3Z, the screening threshold must be tightened to 10^{-10} . This basis set dependence is in accordance with the findings of previous research from other groups. [141] 5 or 6 quadrature points are sufficient to yield very accurate results for frozen-core MP₂, as it was also found by the authors of the minimax quadrature scheme. [128]


Figure B.3 and B.4 show the additional errors introduced by the QQZZR4 screening. In contrast to the original QQZZ screening this criterion does not guarantee to provide an upper bound for the integral estimates, but its estimates are more accurate than the ones provided by QQZZ. The two parameters that control the accuracy in my implementation are the threshold for the outer radius used to determine the extent of the basis functions and the minimum extent that will be assigned to charge distributions. As it can be seen from the presented data, the implementation is rather insensitive to the threshold for the outer radius unless very large thresholds are chosen. An

outer radius threshold of 10^{-2} is thus used in all production calculations. The minimum extent of the charge distributions on the other hand has a larger impact on the overall accuracy. This is not surprising since the Numeric Atom-Centered Orbitals (NAOs) in FHI-aims have a strictly limited radial extent, but the same is not true for the transformed basis functions. As a consequence, the transformed basis function pair may have a non-zero overlap even when the corresponding atomic basis function pair has no more overlap. The same charge extents are used for both the atomic and transformed ERIs for implementation efficiency and simplicity reasons. Thus an appropriate choice of the minimum charge extent is crucial to ensure that the transformed ERIs are approximated correctly. Therefore, a minimum charge extent of $1a_0$ is used for all production calculations in molecular systems. The reciprocal space implementation for periodic systems does not use QQZZR4 screening.


B.2 THE GENERALIZED RI-LVL

The data in Figure B.5 demonstrates the accuracy of the generalized RI-LVL LT-MP2 as introduced in section 4.5. The shown errors are relative to the results obtained with a RI-LVL_{full} implementation, i.e. they only represent the additional approximation introduced by the generalized RI-LVL. The generalized RI-LVL converges quickly with the size of the local ABS and a sphere with $6a_0$ radius around the atom hosting the occupied transformed function is already well converged for both NAO-VCC-2Z and NAO-VCC-3Z basis sets.

The overall accuracy achieved by combining the approximations discussed in this appendix is discussed in chapter 5.

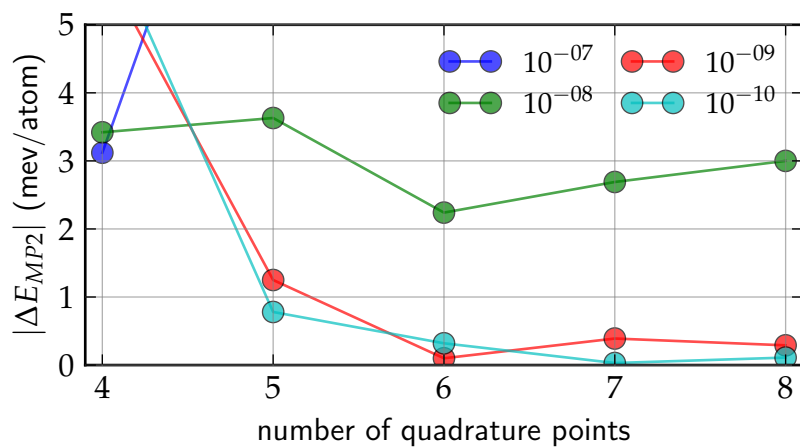


(a) small water cluster with 10 molecules




(b) fully extended polyaniline with 3 segments

Figure B.1: Absolute correlation energy errors per atom as function of the screening threshold and the number of quadrature points with NAO-VCC-2Z basis sets. The reference energies are the RI-V calculations with the canonical MP2. The transformed ERIs have been treated with the LT-MP2 based on RI-LVL_{full}, i.e. the errors shown here only originate from the quadrature and integral screening. (see section 4.4)




(a) HF total energy errors per atom




(b) MP2 total energy errors per atom

Figure B.2: Absolute correlation energy errors per atom as function of the screening threshold and the number of quadrature points with NAO-VCC-3Z basis sets. The reference energies are the RI-V calculations with the canonical MP2. The transformed ERIs have been treated with the LT-MP2 based on RI-LVL_{full}, i.e. the errors shown here only originate from the quadrature and integral screening. (see section 4.4)




(a) small water cluster with 10 molecules




(b) fully extended polyaniline with 3 segments

Figure B.3: Absolute correlation energy errors per atom caused by the QQZZR4 screening criterion with different threshold for the outer radius and minimum charge extent. (see section 4.7) The reference are the LT-MP2 calculations with RI-LVL_{full}. All calculations used 5 quadrature points and 10^{-9} as screening threshold.

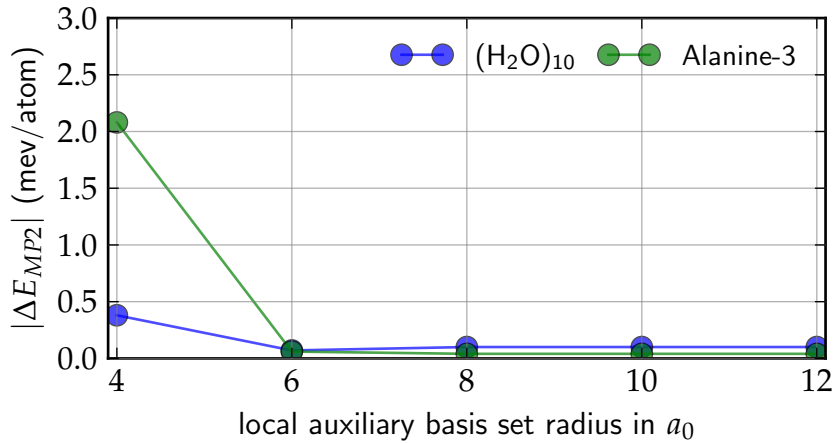


(a) HF total energy errors per atom

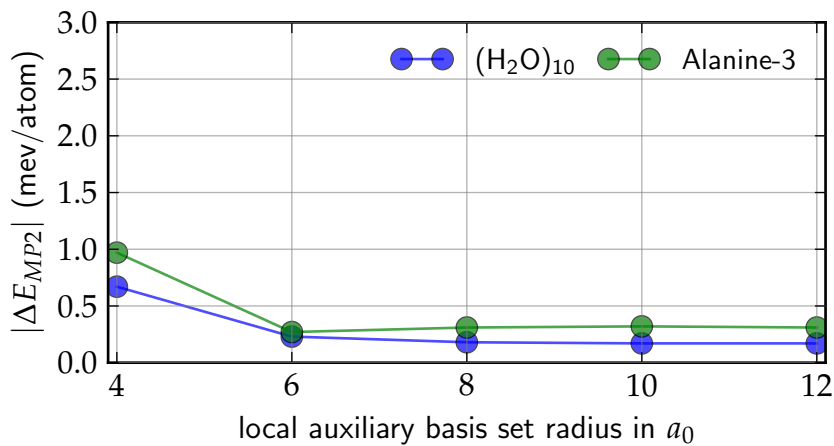


(b) MP2 total energy errors per atom

Figure B.4: Absolute correlation energy errors per atom caused by the QQZZR4 screening criterion with different threshold for the outer radius and minimum charge extent. (see section 4.7) The reference are the LT-MP2 calculations with RI-LVL_{full}. All calculations used 5 quadrature points and 10^{-9} as screening threshold.



(a) Accuracy of the generalized RI-LVL for NAO-VCC-2Z basis sets



(b) Accuracy of the generalized RI-LVL for NAO-VCC-3Z basis sets

Figure B.5: Absolute correlation energy errors per atom introduced by the generalized RI-LVL scheme for the transformed ERIs as a function of the r_{aux} used to determine the local auxiliary basis set. (see section 4.5) The reference are the LT-MP2 calculations with RI-LVL_{full}. All calculations used 5 quadrature points and 10^{-9} as screening threshold.

NOTATION CONVENTIONS AND ACRONYMS

Atomic units are used throughout the thesis.

$\psi_m, \psi_n, \psi_o, \psi_p, \dots$	any molecular orbital
$\psi_i, \psi_j, \psi_k, \psi_l, \dots$	occupied molecular orbital
$\psi_a, \psi_b, \psi_c, \psi_d, \dots$	unoccupied (virtual) molecular orbital
$\varphi_s, \varphi_t, \varphi_u, \varphi_v, \dots$	orbital basis set functions
$P_\mu, P_\lambda, P_\kappa, P_\nu, \dots$	auxiliary basis set functions

Throughout the thesis, the "error per atom" is the total energy error divided by the number of non-hydrogen atoms in the system.

C.1 SCALING PARAMETERS

N_{atoms}	number of atoms in the system
$N_{\text{atom pairs}}$	number of significant atom pairs (have at least one basis function pair with finite overlap)
N_{bas}	number orbital basis functions in the system
$N_{\text{bas pairs}}$	number of significant (non-zero overlap) orbital basis function pairs
$N_{\text{local bas}}$	(maximum) number of orbital basis functions per atom
N_{aux}	number of auxiliary basis functions in the system
$N_{\text{local aux}}$	(maximum) number of auxiliary basis functions per atom
N_{occ}	number of occupied molecular orbitals
N_{virt}	number of unoccupied (virtual) molecular orbitals

C.2 ACRONYMS

ABS	auxiliary basis set
ABS+	enhanced auxiliary basis set
BvK	Born von Karman periodic boundary conditions
CBS	Complete Basis Set
CC	Coupled-Cluster Theory
CCSD	Coupled-Cluster Theory with Singles and Doubles
CCSD(T)	Coupled-Cluster Theory with Singles, Doubles and perturbative Triples
CI	Configuration Interaction
CISD	Configuration Interaction with Single and Double Excitations

DEC-MP2	Divide-Expand-Consolidate 2nd Order Møller-Plesset Perturbation Theory
DFA	Density-Functional Approximation
DFT	Density-Functional Theory
DLPNO	Domain-based local Pair Natural Orbital
ECM	Embedded Cluster Model
ERI	Electron Repulsion Integral
EVD	Eigenvalue Decomposition [142, 143]
FCI	Full Configuration Interaction
GGA	Generalized Gradient Approximation
GTO	Gaussian-Type Orbital
GW	GW approximation
HEG	Homogeneous Electron Gas
HF	Hartree-Fock
HIOS	Hybrid Inorganic-Organic System
HOMO	Highest Occupied Molecular Orbital
KS	Kohn-Sham
LDA	Local Density Approximation
LES	linear equation system
logSBT	logarithmic spherical Bessel transforms [102, 103]
LT	Laplace-Transformation
LT-MP2	Laplace-Transformed 2nd Order Møller-Plesset Perturbation Theory
LUMO	Lowest Unoccupied Molecular Orbital
MBIE	Multipole Based Integral Estimates [41, 130]
MD	Molecular Dynamics
mGGA	meta-Generalized Gradient Approximation
minimax	minimax approximation [128]
MP2	2nd Order Møller-Plesset Perturbation Theory
NAO	Numeric Atom-Centered Orbital
OBS	orbital basis set
OBS+	enhanced orbital basis set

ONIOM	ONIOM embedding model (“our own n-layered integrated molecular orbital and molecular mechanics”) [144]
OSV	Orbital Specific Virtual
PAO	Projected Atomic Orbital
PNO	Pair Natural Orbital
QM/MM	Quantum Mechanics / Molecular Mechanics embedding technique
QQZZ	QQZZ-screening [129]
QQZZR4	QQZZR4-screening [131, 132]
RR-QR	Rank-Revealing QR-Decomposition [142, 143]
RI	Resolution of Identity
RI-LVL	localized Resolution of Identity
RI-LVL _{full}	localized Resolution of Identity (using a RI-V implementation)
RI-SVS	overlap-metric based RI
RI-V	coulomb-metric based RI
RMSD	Root-Mean Square Deviation
SCF	Self-Consistent Field
SCS-MP2	Spin-Component-Scaled 2nd Order Møller-Plesset Perturbation Theory
SOS-MP2	Spin-Opposite-Scaled 2nd Order Møller-Plesset Perturbation Theory
SVD	Singular Value Decomposition [142, 143]
TD-DFT	Time-Dependent Density Functional Theory
THC	Tensor-Hyper-Contraction
QMC	Quantum Monte Carlo
WFT	Wave Function Theory

DFT FUNCTIONALS

B2-PLYP	B2-PLYP (double hybrid) [68]
B-LYP	Becke-Lee-Yang-Parr (GGA) [58, 59]
FSIC-LSD	Fermi-Orbital Self-Interaction Corrected Local Spin Density (LDA) [57]

NOTATION CONVENTIONS AND ACRONYMS

HSE06	Heyd-Scuseria-Ernzerhof (range-separated hybrid) [66]
MBD	Many-Body-Dispersion (Van-der-Waals correction [73])
PBE	Perdew-Ernzerhof-Burke (GGA) [60]
PBEsol	Perdew-Ernzerhof-Burke for solids (GGA) [62]
PBE0	Perdew-Ernzerhof-Burke (hybrid) [65]
RPA	Random-Phase Approximation [6, 7]
rPBE	revised Perdew-Ernzerhof-Burke (GGA) [61]
rPT2	renormalized 2nd Order Perturbation Theory [7]
SCAN	Strongly Constrained and Appropriately Normed (mGGA) [64]
SIC-LSD	Self-Interaction Corrected Local Spin Density (LDA) [56]
TPSS	Tao-Perdew-Staroverov-Scuseria (mGGA) [63]
TS	Tkatchenko-Scheffler (Van-der-Waals correction [70])
vdW ^{surf}	Tkatchenko-Scheffler for surfaces (Van-der-Waals correction [71])
VWN	Vosko-Wilk-Nusair (LDA) [51]
XYG3	XYG3 (double hybrid) [69]

COMPUTATIONAL PACKAGES AND LIBRARIES

CRYSTAL	CRYSTAL - a computational tool for solid state chemistry and physics [145]
FHI-aims	Fritz Haber Institute ab initio molecular simulations [110]
MPI	Message Passing Interface [134]
NWChem	NorthWest computational Chemistry [120]
OpenMP	Open Multi-Processing [133]

SELBSTSTÄNDIGKEITSERKLÄRUNG

Hiermit versichere ich, alle Hilfsmittel und Hilfen angegeben zu haben und auf dieser Grundlage die Arbeit selbstständig verfasst zu haben. Die Arbeit wurde nicht schon einmal in einem früheren Promotionsverfahren angenommen oder als ungenügend beurteilt.

Berlin, den _____

Arvid Conrad Ihrig

CURRICULUM VITAE

Mein Lebenslauf wird aus Gründen des Datenschutzes in der elektronischen Fassung meiner Arbeit nicht veröffentlicht.

PUBLICATIONS

PUBLICATIONS RELATED TO THIS THESIS

- A. C. Ihrig, J. Wierink, I. Y. Zhang, M. Ropo, X. Ren, P. Rinke, M. Scheffler, and V. Blum. “Accurate localized resolution of identity approach for linear-scaling hybrid density functionals and for many-body perturbation theory.” In: *New J. Phys.* 17.9 (2015), p. 093020. DOI: 10.1088/1367-2630/17/9/093020

PUBLICATIONS CONCERNING OTHER TOPICS

- A. C. Ihrig, C. Schiffmann, and D. Sebastiani. “Specific quantum mechanical/molecular mechanical capping-potentials for biomolecular functional groups.” In: *J. Chem. Phys.* 135.21 (2011), p. 214107. DOI: 10.1063/1.3664300
- A. C. Ihrig, A. Scherrer, and D. Sebastiani. “Electronic density response to molecular geometric changes from explicit electronic susceptibility calculations.” In: *J. Chem. Phys.* 139.9 (2013), p. 094102. DOI: 10.1063/1.4819070

ACKNOWLEDGMENTS

First and foremost, I want to express my gratitude to my group leader Igor Ying Zhang. Your constant support throughout these four years has been really amazing!

I want to thank Prof. Dr. Matthias Scheffler for giving me the opportunity to be PhD student at the Fritz-Haber Institute and for refereeing my thesis. I also want to thank my second referee, Prof. Dr. Piet Brouwer.

I give thanks to my (former) office mates Franziska, Adriana, Mateusz and Dimitri for the always positive atmosphere and the other members of the Theory department and Igor's group for the enjoyable years. Special thanks go to Markus, Lydia, Björn, Oliver, Wael, Jürgen and Prof. Dr. Patrick Rinke for many motivating discussions. Thanks to Hanna, Birgit and Julia for always providing a helping hand when needed.

Finally, I want to thank my family and friends for their encouragement and support. I owe Annett a debt of gratitude for her unconditional support and always being there when I needed her.

BIBLIOGRAPHY

-
- [1] D. P. Chong. *Recent Advances in Density Functional Methods - Part I*. Singapore: World Scientific Publishing, 1995. DOI: 10.1142/9789812830586.
- [2] D. P. Chong. *Recent Advances in Density Functional Methods - Part II*. Singapore: World Scientific Publishing, 1997. DOI: 10.1142/9789812819468.
- [3] V. Barone, A. Bencini, and P. Fantucci. *Recent Advances in Density Functional Methods - Part III*. Singapore: World Scientific Publishing, 2002. DOI: 10.1142/9789812778161_fmatter.
- [4] W. Koch and M. C. Holthausen. *A Chemist's Guide to Density Functional Theory*. Weinheim: Wiley-VCH, 2001. DOI: 10.1002/3527600043.
- [5] A. Michaelides, T. J. Martinez, A. Alavi, G. Kresse, and F. R. Manby. "Preface: Special Topic Section on Advanced Electronic Structure Methods for Solids and Surfaces." In: *J. Chem. Phys.* 143.10 (2015), p. 102601. DOI: 10.1063/1.4930182.
- [6] F. Furche. "Molecular tests of the random phase approximation to the exchange-correlation energy functional." In: *Phys. Rev. B* 64.19 (2001), p. 195120. DOI: 10.1103/PhysRevB.64.195120.
- [7] X. Ren, P. Rinke, G. E. Scuseria, and M. Scheffler. "Renormalized second-order perturbation theory for the electron correlation energy: Concept, implementation, and benchmarks." In: *Phys. Rev. B* 88.3 (2013), p. 035120. DOI: 10.1103/PhysRevB.88.035120.
- [8] S. Grimme. "Improved second-order Møller-Plesset perturbation theory by separate scaling of parallel- and antiparallel-spin pair correlation energies." In: *J. Chem. Phys.* 118.20 (2003), pp. 9095–9102. DOI: 10.1063/1.1569242.
- [9] P. Jurečka, J. Šponer, J. Černý, and P. Hobza. "Benchmark database of accurate (MP2 and CCSD(T) complete basis set limit) interaction energies of small model complexes, DNA base pairs, and amino acid pairs." In: *Phys. Chem. Chem. Phys.* 8.17 (2006), pp. 1985–1993. DOI: 10.1039/B600027D.
- [10] I. Y. Zhang and X. Xu. *A New-Generation Density Functional*. Berlin: Springer, 2014. DOI: 10.1007/978-3-642-40421-4.
- [11] R. Huenerbein, B. Schirmer, J. Moellmann, and S. Grimme. "Effects of London dispersion on the isomerization reactions of large organic molecules: a density functional benchmark study." In: *Phys. Chem. Chem. Phys.* 12.26 (2010), p. 6940. DOI: 10.1039/c003951a.

- [12] L. Goerigk and S. Grimme. "Efficient and Accurate Double-Hybrid-Meta-GGA Density Functionals - Evaluation with the Extended GMTKN₃₀ Database for General Main Group Thermochemistry, Kinetics, and Noncovalent Interactions." In: *J. Chem. Theory Comput.* 7.2 (2011), pp. 291–309. DOI: 10.1021/ct100466k.
- [13] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. "Quantum Monte Carlo simulations of solids." In: *Rev. Mod. Phys.* 73.1 (2001), pp. 33–83. DOI: 10.1103/RevModPhys.73.33.
- [14] B. M. Austin, D. Y. Zubarev, and W. A. Lester. "Quantum monte carlo and related approaches." In: *Chem. Rev.* 112.1 (2012), pp. 263–288. DOI: 10.1021/cr2001564.
- [15] M. Häser. "Møller-Plesset (MP2) perturbation theory for large molecules." In: *Theor. Chim. Acta* 87.1-2 (1993), pp. 147–173. DOI: 10.1007/BF01113535.
- [16] H. F. Schurkus and C. Ochsenfeld. "Communication: An effective linear-scaling atomic-orbital reformulation of the random-phase approximation using a contracted double-Laplace transformation." In: *J. Chem. Phys.* 144.3 (2016). DOI: 10.1063/1.4939841.
- [17] C. Riplinger and F. Neese. "An efficient and near linear scaling pair natural orbital based local coupled cluster method." In: *J. Chem. Phys.* 138.3 (2013), p. 034106. DOI: 10.1063/1.4773581.
- [18] U. Diebold. "The surface science of titanium dioxide." In: *Surf. Sci. Rep.* 48.5-8 (2003), pp. 53–229. DOI: 10.1016/S0167-5729(02)00100-0.
- [19] A. Fujishima, X. Zhang, and D. A. Tryk. "TiO₂ photocatalysis and related surface phenomena." In: *Surf. Sci. Rep.* 63.12 (2008), pp. 515–582. DOI: 10.1016/j.surfrep.2008.10.001.
- [20] M. A. Henderson. "A surface science perspective on TiO₂ photocatalysis." In: *Surf. Sci. Rep.* 66.6-7 (2011), pp. 185–297. DOI: 10.1016/j.surfrep.2011.01.001.
- [21] D. Stodt, H. Noei, C. Hättig, and Y. Wang. "A combined experimental and computational study on the adsorption and reactions of NO on rutile TiO₂." In: *Phys. Chem. Chem. Phys.* 15.2 (2013), pp. 466–472. DOI: 10.1039/C2CP42653F.
- [22] A. Kubas, D. Berger, H. Oberhofer, D. Maganas, K. Reuter, and F. Neese. "Surface Adsorption Energetics Studied with "Gold Standard" Wave-Function-Based Ab Initio Methods: Small-Molecule Binding to TiO₂ (110)." In: *J. Phys. Chem. Lett.* 7.20 (2016), pp. 4207–4212. DOI: 10.1021/acs.jpcclett.6b01845.
- [23] A. Szabó and N. S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Mineola: Dover Publications, 1996.
- [24] A. Togo and I. Tanaka. "First principles phonon calculations in materials science." In: *Scr. Mater.* 108 (2015), pp. 1–5. DOI: 10.1016/j.scriptamat.2015.07.021.

- [25] C. Baldauf and M. Rossi. "Going clean: structure and dynamics of peptides in the gas phase and paths to solvation." In: *J. Phys. Condens. Matter* 27.49 (2015), p. 493002. DOI: 10.1088/0953-8984/27/49/493002.
- [26] S. Pisana, M. Lazzeri, C. Casiraghi, K. S. Novoselov, A. K. Geim, A. C. Ferrari, and F. Mauri. "Breakdown of the adiabatic Born-Oppenheimer approximation in graphene." In: *Nat. Mater.* 6.3 (2007), pp. 198–201. DOI: 10.1038/nmat1846.
- [27] L. Hedin. "New method for calculating the one-particle Green's function with application to the electron gas problem." In: *Phys. Rev.* 139.3A (1965), A796–A823.
- [28] M. A. L. Marques, C. A. Ullrich, F. Nogueira, A. Rubio, K. Burke, and E. K. U. Gross. *Time-Dependent Density Functional Theory*. Berlin: Springer, 2006. DOI: 10.1007/b11767107.
- [29] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi. "Towards an exact description of electronic wavefunctions in real solids." In: *Nature* 493.7432 (2012), pp. 365–370. DOI: 10.1038/nature11770.
- [30] R. J. Bartlett and M. Musiał. "Coupled-cluster theory in quantum chemistry." In: *Rev. Mod. Phys.* 79.1 (2007), pp. 291–352. DOI: 10.1103/RevModPhys.79.291.
- [31] C. Hättig, W. Klopper, A. Köhn, and D. P. Tew. "Explicitly Correlated Electrons in Molecules." In: *Chem. Rev.* 112.1 (2012), pp. 4–74. DOI: 10.1021/cr200168z.
- [32] I. Y. Zhang, X. Ren, P. Rinke, V. Blum, and M. Scheffler. "Numeric atom-centered-orbital basis sets with valence-correlation consistency from H to Ar." In: *New J. Phys.* 15.12 (2013), p. 123033. DOI: 10.1088/1367-2630/15/12/123033.
- [33] T. H. Dunning. "Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen." In: *J. Chem. Phys.* 90.2 (1989), pp. 1007–1023. DOI: 10.1063/1.456153.
- [34] M. Schütz, G. Hetzer, and H.-J. Werner. "Low-order scaling local electron correlation methods. I. Linear scaling local MP2." In: *J. Chem. Phys.* 111.13 (1999), pp. 5691–5705. DOI: 10.1063/1.479957.
- [35] J. Yang, Y. Kurashige, F. R. Manby, and G. K. L. Chan. "Tensor factorizations of local second-order Møller-Plesset theory." In: *J. Chem. Phys.* 134.4 (2011), p. 044123. DOI: 10.1063/1.3528935.
- [36] G. Schmitz, B. Helmich, and C. Hättig. "A scaling PNO-MP2 method using a hybrid OSV-PNO approach with an iterative direct generation of OSVs." In: *Mol. Phys.* 111.16-17 (2013), pp. 2463–2476. DOI: 10.1080/00268976.2013.794314.

- [37] P. Pinski, C. Riplinger, E. F. Valeev, and F. Neese. "Sparse maps - A systematic infrastructure for reduced-scaling electronic structure methods. I. An efficient and simple linear scaling local MP2 method that uses an intermediate basis of pair natural orbitals." In: *J. Chem. Phys.* 143.3 (2015), p. 034108. DOI: 10.1063/1.4926879.
- [38] R. Ahlrichs, H. Lischka, V. Staemmler, and W. Kutzelnigg. "PNO-CI (pair natural orbital configuration interaction) and CEPA-PNO (coupled electron pair approximation with pair natural orbitals) calculations of molecular systems. I. Outline of the method for closed-shell states." In: *J. Chem. Phys.* 62.4 (1975), pp. 1225–1234. DOI: 10.1063/1.430637.
- [39] F. Neese, F. Wennmohs, and A. Hansen. "Efficient and accurate local approximations to coupled-electron pair approaches: An attempt to revive the pair natural orbital method." In: *J. Chem. Phys.* 130.11 (2009), p. 114108. DOI: 10.1063/1.3086717.
- [40] M. Häser and J. Almlöf. "Laplace transform techniques in Møller-Plesset perturbation theory." In: *J. Chem. Phys.* 96.1 (1992), pp. 489–494. DOI: 10.1063/1.462485.
- [41] D. S. Lambrecht, B. Doser, and C. Ochsenfeld. "Rigorous integral screening for electron correlation methods." In: *J. Chem. Phys.* 123.18 (2005), p. 184102. DOI: 10.1063/1.2079987.
- [42] S. A. Maurer, L. Clin, and C. Ochsenfeld. "Cholesky-decomposed density MP2 with density fitting: Accurate MP2 and double-hybrid DFT energies for large systems." In: *J. Chem. Phys.* 140.22 (2014), p. 224112. DOI: 10.1063/1.4881144.
- [43] K. Kristensen, T. Kjærgaard, I.-M. Høyvik, P. Ettenhuber, P. Jørgensen, B. Jansik, S. Reine, and J. Jakowski. "The divide-expand-consolidate MP2 scheme goes massively parallel." In: *Mol. Phys.* 111.9-11 (2013), pp. 1196–1210. DOI: 10.1080/00268976.2013.783941.
- [44] T. Kjærgaard, P. Baudin, D. Bykov, J. J. Eriksen, P. Ettenhuber, K. Kristensen, J. Larkin, D. Liakh, F. Pawłowski, A. Vose, Y. M. Wang, and P. Jørgensen. "Massively parallel and linear-scaling algorithm for second-order Møller-Plesset perturbation theory applied to the study of supramolecular wires." In: *Comput. Phys. Commun.* 212 (2017), pp. 152–160. DOI: 10.1016/j.cpc.2016.11.002.
- [45] T. Kjærgaard. "The Laplace transformed divide-expand-consolidate resolution of the identity second-order Møller-Plesset perturbation (DEC-LT-RIMP2) theory method." In: *J. Chem. Phys.* 146.4 (2017), p. 044103. DOI: 10.1063/1.4973710.
- [46] M. Katouda and T. Nakajima. "MPI/OpenMP hybrid parallel algorithm for resolution of identity second-order Møller-Plesset perturbation calculation of analytical energy gradient for massively parallel multicore supercomputers." In: *J. Comput. Chem.* 38.8 (2017), pp. 489–507. DOI: 10.1002/jcc.24701.

- [47] Y. Jung, R. C. Lochan, A. D. Dutoi, and M. Head-Gordon. "Scaled opposite-spin second order Møller-Plesset correlation energy: An economical electronic structure method." In: *J. Chem. Phys.* 121.20 (2004), pp. 9793–9802. DOI: 10.1063/1.1809602.
- [48] P. Hohenberg and W. Kohn. "Inhomogeneous Electron Gas." In: *Phys. Rev.* 136.3B (1964), B864–B871. DOI: 10.1103/PhysRev.136.B864.
- [49] E. Engel and R. M. Dreizler. *Density Functional Theory*. Berlin: Springer, 2011. DOI: 10.1007/978-3-642-14090-7.
- [50] J. P. Perdew. "Jacob's ladder of density functional approximations for the exchange-correlation energy." In: *AIP Conf. Proc.* Vol. 577. Density Functional Theory and Its Application to Materials. AIP, 2001, pp. 1–20. DOI: 10.1063/1.1390175.
- [51] S. H. Vosko, L. Wilk, and M. Nusair. "Accurate spin-dependent electron liquid correlation energies for local spin density calculations: a critical analysis." In: *Can. J. Phys.* 58.8 (1980), pp. 1200–1211. DOI: 10.1139/p80-159.
- [52] E. Wigner. "On the interaction of electrons in metals." In: *Phys. Rev.* 46.11 (1934), pp. 1002–1011. DOI: 10.1103/PhysRev.46.1002.
- [53] E. Wigner. "Effects of the electron interaction on the energy levels of electrons in metals." In: *Trans. Faraday Soc.* 34 (1938), p. 678. DOI: 10.1039/tf9383400678.
- [54] M. Gell-Mann and K. Brueckner. "Correlation energy of an electron gas at high density." In: *Phys. Rev.* 106.2 (1957), pp. 364–368.
- [55] D. M. Ceperley and B. J. Alder. "Ground state of the electron gas by a stochastic method." In: *Phys. Rev. Lett.* 45.7 (1980), pp. 566–569. DOI: 10.1103/PhysRevLett.45.566.
- [56] J. P. Perdew and A. Zunger. "Self-interaction correction to density-functional approximations for many-electron systems." In: *Phys. Rev. B* 23.10 (1981), pp. 5048–5079. DOI: 10.1103/PhysRevB.23.5048.
- [57] M. R. Pederson, A. Ruzsinszky, and J. P. Perdew. "Communication: Self-interaction correction with unitary invariance in density functional theory." In: *J. Chem. Phys.* 140.12 (2014). DOI: 10.1063/1.4869581.
- [58] A. D. Becke. "Density-functional exchange-energy approximation with correct asymptotic behavior." In: *Phys. Rev. A* 38.6 (1988), pp. 3098–3100. DOI: 10.1103/PhysRevA.38.3098.
- [59] C. Lee, W. Yang, and R. G. Parr. "Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density." In: *Phys. Rev. B* 37.2 (1988), pp. 785–789. DOI: 10.1103/PhysRevB.37.785.
- [60] J. P. Perdew, K. Burke, and M. Ernzerhof. "Generalized Gradient Approximation Made Simple." In: *Phys. Rev. Lett.* 77.18 (1996), pp. 3865–3868. DOI: 10.1103/PhysRevLett.77.3865.

- [61] B. Hammer, L. B. Hansen, and J. K. Nørskov. "Improved adsorption energetics within density-functional theory using revised Perdew-Burke-Ernzerhof functionals." In: *Phys. Rev. B* 59.11 (1999), pp. 7413–7421. DOI: 10.1103/PhysRevB.59.7413.
- [62] J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou, and K. Burke. "Generalized gradient approximation for solids and their surfaces." In: 136406. April (2007), pp. 1–4. DOI: 10.1103/PhysRevLett.100.136406.
- [63] J. Tao, J. P. Perdew, V. N. Staroverov, and G. E. Scuseria. "Climbing the Density Functional Ladder: Nonempirical Meta-Generalized Gradient Approximation Designed for Molecules and Solids." In: *Phys. Rev. Lett.* 91.14 (2003), p. 146401. DOI: 10.1103/PhysRevLett.91.146401.
- [64] J. Sun, A. Ruzsinszky, and J. Perdew. "Strongly Constrained and Appropriately Normed Semilocal Density Functional." In: *Phys. Rev. Lett.* 115.3 (2015), p. 036402. DOI: 10.1103/PhysRevLett.115.036402.
- [65] J. P. Perdew, M. Ernzerhof, and K. Burke. "Rationale for mixing exact exchange with density functional approximations." In: *J. Chem. Phys.* 105.22 (1996), pp. 9982–9985. DOI: 10.1063/1.472933.
- [66] J. Heyd, G. E. Scuseria, and M. Ernzerhof. "Hybrid functionals based on a screened Coulomb potential." In: *J. Chem. Phys.* 118.18 (2003), pp. 8207–8215. DOI: 10.1063/1.1564060.
- [67] X. Ren, P. Rinke, C. Joas, and M. Scheffler. "Random-phase approximation and its applications in computational chemistry and materials science." In: *J. Mater. Sci.* 47.21 (2012), pp. 7447–7471. DOI: 10.1007/s10853-012-6570-4.
- [68] S. Grimme. "Semiempirical hybrid density functional with perturbative second-order correlation." In: *J. Chem. Phys.* 124.3 (2006), p. 034108. DOI: 10.1063/1.2148954.
- [69] Y. Zhang, X. Xu, and W. A. Goddard. "Doubly hybrid density functional for accurate descriptions of nonbond interactions, thermochemistry, and thermochemical kinetics." In: *Proc. Natl. Acad. Sci.* 106.13 (2009), pp. 4963–4968. DOI: 10.1073/pnas.0901093106.
- [70] A. Tkatchenko and M. Scheffler. "Accurate Molecular Van Der Waals Interactions from Ground-State Electron Density and Free-Atom Reference Data." In: *Phys. Rev. Lett.* 102.7 (2009), p. 073005. DOI: 10.1103/PhysRevLett.102.073005.
- [71] V. G. Ruiz, W. Liu, and A. Tkatchenko. "Density-functional theory with screened van der Waals interactions applied to atomic and molecular adsorbates on close-packed and non-close-packed surfaces." In: *Phys. Rev. B - Condens. Matter Mater. Phys.* 93.3 (2016), pp. 1–17. DOI: 10.1103/PhysRevB.93.035118.

- [72] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg. "A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu." In: *J. Chem. Phys.* 132.15 (2010), p. 154104. DOI: 10.1063/1.3382344.
- [73] A. Tkatchenko, R. A. DiStasio, R. Car, and M. Scheffler. "Accurate and Efficient Method for Many-Body van der Waals Interactions." In: *Phys. Rev. Lett.* 108.23 (2012), p. 236402. DOI: 10.1103/PhysRevLett.108.236402.
- [74] H. Ibach and H. Lüth. *Solid-State Physics (3rd Extensively Updated and Enlarged Edition)*. Berlin: Springer, 2003. DOI: 10.1007/978-3-662-05342-3.
- [75] R. A. Evarestov. *Quantum Chemistry of Solids (2nd Edition)*. Berlin: Springer, 2012. DOI: 10.1007/978-3-642-30356-2.
- [76] J. D. Pack and H. J. Monkhorst. "'Special points for Brillouin-zone integrations"—a reply." In: *Phys. Rev. B* 16.4 (1977), pp. 1748–1749. DOI: 10.1103/PhysRevB.16.1748.
- [77] P. Deák, T. Frauenheim, and M. R. Pederson, eds. *Computer Simulation of Materials at Atomic Level*. Weinheim: Wiley-VCH Verlag GmbH & Co. KGaA, 2000. DOI: 10.1002/3527603107.
- [78] S. V. Levchenko, X. Ren, J. Wieferink, R. Johanni, P. Rinke, V. Blum, and M. Scheffler. "Hybrid functionals for large periodic systems in an all-electron, numeric atom-centered basis framework." In: *Comput. Phys. Commun.* 192 (2015), pp. 60–69. DOI: 10.1016/j.cpc.2015.02.021.
- [79] J. Sun and R. J. Bartlett. "Second-order many-body perturbation-theory calculations in extended systems." In: *J. Chem. Phys.* 104.21 (1996), pp. 8553–8565. DOI: 10.1063/1.471545.
- [80] C. Pisani, L. Maschio, S. Casassa, M. Halo, M. Schütz, and D. Usvyat. "Periodic local MP2 method for the study of electronic correlation in crystals: Theory and preliminary applications." In: *J. Comput. Chem.* 29.13 (2008), pp. 2113–2124. DOI: 10.1002/jcc.20975.
- [81] P. Y. Ayala and G. E. Scuseria. "Linear scaling second-order Møller-Plesset theory in the atomic orbital basis for large molecular systems." In: *J. Chem. Phys.* 110.8 (1999), pp. 3660–3671. DOI: 10.1063/1.478256.
- [82] M. Del Ben, J. Hutter, and J. VandeVondele. "Forces and stress in second order Møller-Plesset perturbation theory for condensed phase systems within the resolution-of-identity Gaussian and plane waves approach." In: *J. Chem. Phys.* 143.10 (2015), p. 102803. DOI: 10.1063/1.4919238.
- [83] M. Del Ben, J. Hutter, and J. VandeVondele. "Probing the structural and dynamical properties of liquid water with models including non-local electron correlation." In: *J. Chem. Phys.* 143.5 (2015). DOI: 10.1063/1.4927325.

- [84] J. Harl and G. Kresse. "Cohesive energy curves for noble gas solids calculated by adiabatic connection fluctuation-dissipation theory." In: *Phys. Rev. B - Condens. Matter Mater. Phys.* 77.4 (2008), pp. 1–8. DOI: 10.1103/PhysRevB.77.045136.
- [85] M. Kaltak, J. Klimeš, and G. Kresse. "Cubic scaling algorithm for the random phase approximation: Self-interstitials and vacancies in Si." In: *Phys. Rev. B - Condens. Matter Mater. Phys.* 90.5 (2014), pp. 1–11. DOI: 10.1103/PhysRevB.90.054115.
- [86] O. T. Hofmann and P. Rinke. "Band Bending Engineering at Organic/Inorganic Interfaces Using Organic Self-Assembled Monolayers." In: *Adv. Electron. Mater.* (2017), p. 1600373. DOI: 10.1002/aelm.201600373.
- [87] T. Vreven, K. S. Byun, I. Komáromi, S. Dapprich, J. A. Montgomery, K. Morokuma, and M. J. Frisch. "Combining Quantum Mechanics Methods with Molecular Mechanics Methods in ONIOM." In: *J. Chem. Theory Comput.* 2.3 (2006), pp. 815–826. DOI: 10.1021/ct050289g.
- [88] L. W. Chung, W. M. Sameera, R. Ramozzi, A. J. Page, M. Hatanaka, G. P. Petrova, T. V. Harris, X. Li, Z. Ke, F. Liu, H. B. Li, L. Ding, and K. Morokuma. "The ONIOM Method and Its Applications." In: *Chem. Rev.* 115.12 (2015), pp. 5678–5796. DOI: 10.1021/cr5004419.
- [89] D. Berger, H. Oberhofer, and K. Reuter. "First-principles embedded-cluster calculations of the neutral and charged oxygen vacancy at the rutile TiO₂ (110) surface." In: *Phys. Rev. B* 92.7 (2015), p. 075308. DOI: 10.1103/PhysRevB.92.075308.
- [90] A. C. Ihrig, C. Schiffmann, and D. Sebastiani. "Specific quantum mechanical/molecular mechanical capping-potentials for biomolecular functional groups." In: *J. Chem. Phys.* 135.21 (2011), p. 214107. DOI: 10.1063/1.3664300.
- [91] D. Berger, A. J. Logsdail, H. Oberhofer, M. R. Farrow, C. R. A. Catlow, P. Sherwood, A. A. Sokol, V. Blum, and K. Reuter. "Embedded-cluster calculations in a numeric atomic orbital density-functional theory framework." In: *J. Chem. Phys.* 141.2 (2014), p. 024105. DOI: 10.1063/1.4885816.
- [92] P. M. W. Gill and J. a. Pople. "The prism algorithm for two-electron integrals." In: *Int. J. Quantum Chem.* 40.6 (1991), pp. 753–772. DOI: 10.1002/qua.560400605.
- [93] J. L. Whitten. "Coulombic potential energy integrals and approximations." In: *J. Chem. Phys.* 58.10 (1973), pp. 4496–4501. DOI: 10.1063/1.1679012.
- [94] O. Vahtras, J. Almlöf, and M. Feyereisen. "Integral approximations for LCAO-SCF calculations." In: *Chem. Phys. Lett.* 213.5 (1993), pp. 514–518.

- [95] K. Eichkorn, O. Treutler, H. Öhm, M. Häser, and R. Ahlrichs. "Auxiliary basis sets to approximate Coulomb potentials." In: *Chem. Phys. Lett.* 240.4 (1995), pp. 283–290. DOI: 10.1016/0009-2614(95)00621-A.
- [96] X. Ren, P. Rinke, V. Blum, J. Wieferink, A. Tkatchenko, A. Sanfilippo, K. Reuter, and M. Scheffler. "Resolution-of-identity approach to Hartree-Fock, hybrid density functionals, RPA, MP2 and GW with numeric atom-centered orbital basis functions." In: *New J. Phys.* 14.5 (2012), p. 053020. DOI: 10.1088/1367-2630/14/5/053020.
- [97] S. I. L. Kokkila Schumacher, E. G. Hohenstein, R. M. Parrish, L.-P. Wang, and T. J. Martínez. "Tensor Hypercontraction Second-Order Møller-Plesset Perturbation Theory: Grid Optimization and Reaction Energies." In: *J. Chem. Theory Comput.* 11.7 (2015), pp. 3042–3052. DOI: 10.1021/acs.jctc.5b00272.
- [98] C. Song and T. J. Martínez. "Atomic orbital-based SOS-MP2 with tensor hypercontraction. I. GPU-based tensor construction and exploiting sparsity." In: *J. Chem. Phys.* 144.17 (2016), p. 174111. DOI: 10.1063/1.4948438.
- [99] R. A. Friesner. "Solution of the Hartree-Fock equations for polyatomic molecules by a pseudospectral method." In: *J. Chem. Phys.* 86.6 (1987), pp. 3522–3531. DOI: 10.1063/1.451955.
- [100] T. J. Martinez and E. A. Carter. "Pseudospectral Møller-Plesset perturbation theory through third order." In: *J. Chem. Phys.* 100.5 (1994), pp. 3631–3638. DOI: 10.1063/1.466350.
- [101] J. D. J. Talman. "Numerical Fourier and Bessel transforms in logarithmic variables." In: *J. Comput. Phys.* 29.1 (1978), pp. 35–48. DOI: 10.1016/0021-9991(78)90107-9.
- [102] J. D. Talman. "Numerical calculation of four-center Coulomb integrals." In: *J. Chem. Phys.* 80.5 (1984), pp. 2000–2008. DOI: 10.1063/1.446963.
- [103] J. D. Talman. "Numerical methods for multicenter integrals for numerically defined basis functions applied in molecular calculations." In: *Int. J. Quantum Chem.* 93.2 (2003), pp. 72–90. DOI: 10.1002/qua.10538.
- [104] J. D. Talman. "Multipole expansions for numerical orbital products." In: *Int. J. Quantum Chem.* 107.7 (2007), pp. 1578–1584. DOI: 10.1002/qua.21308.
- [105] J. Talman. "NumSBT: A subroutine for calculating spherical Bessel transforms numerically." In: *Comput. Phys. Commun.* 180.2 (2009), pp. 332–338. DOI: 10.1016/j.cpc.2008.10.003.
- [106] B. Delley. "An all-electron numerical method for solving the local density functional for polyatomic molecules." In: *J. Chem. Phys.* 92.1 (1990), pp. 508–517. DOI: 10.1063/1.458452.

- [107] A. Zunger and A. J. Freeman. "Self-consistent numerical-basis-set linear-combination-of-atomic-orbitals model for the study of solids in the local density formalism." In: *Phys. Rev. B* 15.10 (1977), pp. 4716–4737. DOI: 10.1103/PhysRevB.15.4716.
- [108] J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal. "The SIESTA method for ab initio order- N materials simulation." In: *J. Phys. Condens. Matter* 14.11 (2002), pp. 2745–2779. DOI: 10.1088/0953-8984/14/11/302.
- [109] S. Kenny and A. Horsfield. "Plato: A localised orbital based density functional theory code." In: *Comput. Phys. Commun.* 180.12 (2009), pp. 2616–2621. DOI: 10.1016/j.cpc.2009.08.006.
- [110] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler. "Ab initio molecular simulations with numeric atom-centered orbitals." In: *Comput. Phys. Commun.* 180.11 (2009), pp. 2175–2196. DOI: 10.1016/j.cpc.2009.06.022.
- [111] A. C. Ihrig, J. Wieferink, I. Y. Zhang, M. Ropo, X. Ren, P. Rinke, M. Scheffler, and V. Blum. "Accurate localized resolution of identity approach for linear-scaling hybrid density functionals and for many-body perturbation theory." In: *New J. Phys.* 17.9 (2015), p. 093020. DOI: 10.1088/1367-2630/17/9/093020.
- [112] S. Reine, E. Tellgren, A. Krapp, T. Kjærgaard, T. Helgaker, B. Jansik, S. Høst, and P. Salek. "Variational and robust density fitting of four-center two-electron integrals in local metrics." In: *J. Chem. Phys.* 129.10 (2008), p. 104101. DOI: 10.1063/1.2956507.
- [113] A. Sodt, J. E. Subotnik, and M. Head-Gordon. "Linear scaling density fitting." In: *J. Chem. Phys.* 125.19 (2006), p. 194109. DOI: 10.1063/1.2370949.
- [114] A. Sodt and M. Head-Gordon. "Hartree-Fock exchange computed using the atomic resolution of the identity approximation." In: *J. Chem. Phys.* 128.10 (2008), p. 104106. DOI: 10.1063/1.2828533.
- [115] C. Pisani, M. Busso, G. Capecchi, S. Casassa, R. Dovesi, L. Maschio, C. Zicovich-Wilson, and M. Schütz. "Local-MP2 electron correlation method for nonconducting crystals." In: *J. Chem. Phys.* 122.9 (2005), p. 094113. DOI: 10.1063/1.1857479.
- [116] P. Merlot, T. Kjaergaard, T. Helgaker, R. Lindh, F. Aquilante, S. Reine, and T. B. Pedersen. "Attractive electron-electron interactions within robust local fitting approximations." In: *J. Comput. Chem.* 34.17 (2013), pp. 1486–1496. DOI: 10.1002/jcc.23284.

- [117] S. F. Manzer, E. Epifanovsky, and M. Head-Gordon. "Efficient Implementation of the Pair Atomic Resolution of the Identity Approximation for Exact Exchange for Hybrid and Range-Separated Density Functionals." In: *J. Chem. Theory Comput.* 11.2 (2015), pp. 518–527. DOI: 10.1021/ct5008586.
- [118] D. Bes. *Quantum Mechanics (Third Edition)*. Berlin: Springer, 2012. DOI: 10.1007/978-3-642-20556-9.
- [119] B. Dunlap. "Robust variational fitting: Gáspár's variational exchange can accurately be treated analytically." In: *J. Mol. Struct. THEOCHEM* 501-502 (2000), pp. 221–228. DOI: 10.1016/S0166-1280(99)00433-9.
- [120] M. Valiev, E. Bylaska, N. Govind, K. Kowalski, T. Straatsma, H. Van Dam, D. Wang, J. Nieplocha, E. Apra, T. Windus, and W. de Jong. "NWChem: A comprehensive and scalable open-source solution for large scale molecular simulations." In: *Comput. Phys. Commun.* 181.9 (2010), pp. 1477–1489. DOI: 10.1016/j.cpc.2010.04.018.
- [121] X. Ren, P. Rinke, and M. Scheffler. "Exploring the random phase approximation: Application to CO adsorbed on Cu(111)." In: *Phys. Rev. B* 80.4 (2009), p. 045402. DOI: 10.1103/PhysRevB.80.045402.
- [122] S. Bhattacharya, B. H. Sonin, C. J. Jumonville, L. M. Ghiringhelli, and N. Marom. "Computational design of nanoclusters by property-based genetic algorithms: Tuning the electronic properties of (TiO₂)_n clusters." In: *Phys. Rev. B* 91.24 (2015), p. 241115. DOI: 10.1103/PhysRevB.91.241115.
- [123] L. M. Ghiringhelli, P. Gruene, J. T. Lyon, D. M. Rayner, G. Meijer, A. Fielicke, and M. Scheffler. "Not so loosely bound rare gas atoms: finite-temperature vibrational fingerprints of neutral gold-cluster complexes." In: *New J. Phys.* 15.8 (2013), p. 083003. DOI: 10.1088/1367-2630/15/8/083003.
- [124] C. Ochsenfeld, C. A. White, and M. Head-Gordon. "Linear and sublinear scaling formation of Hartree–Fock-type exchange matrices." In: *J. Chem. Phys.* 109.5 (1998), pp. 1663–1669. DOI: 10.1063/1.476741.
- [125] E. Schwegler, M. Challacombe, and M. Head-Gordon. "Linear scaling computation of the Fock matrix. II. Rigorous bounds on exchange integrals and incremental Fock build." In: *J. Chem. Phys.* 106.23 (1997), pp. 9708–9717. DOI: 10.1063/1.473833.
- [126] J. Spencer and A. Alavi. "Efficient calculation of the exact exchange energy in periodic systems using a truncated Coulomb potential." In: *Phys. Rev. B* 77.19 (2008), p. 193110. DOI: 10.1103/PhysRevB.77.193110.

- [127] I. Y. Zhang, A. J. Logsdail, N. A. Richter, T. Morshedloo, X. Ren, S. Levchenko, L. Ghiringhelli, and M. Scheffler. "Test set for materials science and engineering: Accurate reference data for a hierarchy of electronic-structure approaches." In: *Manuscr. Prep.* (2017).
- [128] A. Takatsuka, S. Ten-no, and W. Hackbusch. "Minimax approximation for the decomposition of energy denominators in Laplace-transformed Møller-Plesset perturbation theories." In: *J. Chem. Phys.* 129.4 (2008), p. 044112. DOI: 10.1063/1.2958921.
- [129] M. Häser and R. Ahlrichs. "Improvements on the direct SCF method." In: *J. Comput. Chem.* 10.1 (1989), pp. 104–111. DOI: 10.1002/jcc.540100111.
- [130] D. S. Lambrecht and C. Ochsenfeld. "Multipole-based integral estimates for the rigorous description of distance dependence in two-electron integrals." In: *J. Chem. Phys.* 123.18 (2005), p. 184101. DOI: 10.1063/1.2079967.
- [131] S. A. Maurer, D. S. Lambrecht, D. Flaig, and C. Ochsenfeld. "Distance-dependent Schwarz-based integral estimates for two-electron integrals: Reliable tightness vs. rigorous upper bounds." In: *J. Chem. Phys.* 136.14 (2012), p. 144107. DOI: 10.1063/1.3693908.
- [132] D. S. Lambrecht, B. Doser, and C. Ochsenfeld. "Erratum: \Rigorous integral screening for electron correlation methods" [*J. Chem. Phys.* 123, 184102 (2005)]." In: *J. Chem. Phys.* 136.14 (2012), p. 149902. DOI: 10.1063/1.3693968.
- [133] OpenMP Architecture Review Board. *OpenMP Application Program Interface*. 2013. URL: <http://www.openmp.org/> (visited on 01/18/2017).
- [134] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard*. 2015. URL: <http://www.mpi-forum.org/docs/docs.html> (visited on 02/23/2016).
- [135] E. P. Meagher and G. A. Lager. "Polyhedral thermal expansion in the TiO₂ polymorphs: refinement of the crystal structures of rutile and brookite at high temperature." In: *Can. Mineral.* 17.23 (1979), pp. 77–85. DOI: 10.1103/PhysRevB.76.235307.
- [136] J. Neugebauer and M. Scheffler. "Adsorbate-substrate and adsorbate-adsorbate interactions of Na and K adlayers on Al(111)." In: *Phys. Rev. B* 46.24 (1992), pp. 16067–16080. DOI: 10.1103/PhysRevB.46.16067.
- [137] B. Lee, C.-k. Lee, C. S. Hwang, and S. Han. "Influence of exchange-correlation functionals on dielectric properties of rutile TiO₂." In: *Curr. Appl. Phys.* 11.1 (2011), S293–S296. DOI: 10.1016/j.cap.2010.11.104.
- [138] NVIDIA Corporation. *NVIDIA Tesla GPU Accelerators Datasheet*. URL: <https://www.nvidia.com/content/tesla/pdf/NVIDIA-Tesla-Kepler-Family-Datasheet.pdf> (visited on 08/09/2017).

- [139] Max Planck Computing & Data Facility. *Draco Cluster configuration*. URL: <http://www.mpcdf.mpg.de/services/computing/draco/configuration> (visited on 08/09/2017).
- [140] I. N. Bronstein and K. A. Semendjajew. *Taschenbuch der Mathematik*. Leipzig: BSB B. G. Teubner Verlagsgesellschaft, 1981.
- [141] B. Doser, J. Zienau, L. Clin, D. S. Lambrecht, and C. Ochsenfeld. "A Linear-Scaling MP2 Method for Large Molecules by Rigorous Integral-Screening Criteria." In: *Zeitschrift für Phys. Chemie* 224.3-4 (2010), pp. 397-412. DOI: 10.1524/zpch.2010.6113.
- [142] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. 1999. URL: http://www.netlib.org/lapack/lug/lapack%7B%5C_%7Dlug.html (visited on 02/21/2017).
- [143] Intel Corporation. *Intel Math Kernel Library (MKL)*. 2017. URL: <https://software.intel.com/en-us/mkl-reference-manual-for-fortran> (visited on 02/21/2017).
- [144] M. Svensson, S. Humbel, R. D. J. Froese, T. Matsubara, S. Sieber, and K. Morokuma. "ONIOM: A Multilayered Integrated MO + MM Method for Geometry Optimizations and Single Point Energy Predictions. A Test for Diels-Alder Reactions and Pt(P(*t*-Bu)₃)₂ + H₂ Oxidative Addition." In: *J. Phys. Chem.* 100.50 (1996), pp. 19357-19363. DOI: 10.1021/jp962071j.
- [145] R. Dovesi, R. Orlando, A. Erba, C. M. Zicovich-Wilson, B. Civalieri, S. Casassa, L. Maschio, M. Ferrabone, M. De La Pierre, P. D'Arco, Y. Noël, M. Causà, M. Rérat, and B. Kirtman. "CRYSTAL14: A program for the ab initio investigation of crystalline solids." In: *Int. J. Quantum Chem.* 114.19 (2014), pp. 1287-1317. DOI: 10.1002/qua.24658.
- [146] A. C. Ihrig, A. Scherrer, and D. Sebastiani. "Electronic density response to molecular geometric changes from explicit electronic susceptibility calculations." In: *J. Chem. Phys.* 139.9 (2013), p. 094102. DOI: 10.1063/1.4819070.