

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

**Cholesky Decomposed
Density Matrices
in Laplace Transform
Møller-Plesset Perturbation Theory**

von

Lucien Cyril Clin

aus

Tübingen

2012

Erklärung:

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

Eidstattliche Versicherung:

Diese Dissertation wurde eigenständig und ohne unerlaubte Hilfe erarbeitet.

München, den 26.04.2012

Dissertation eingereicht am:	26.04.2012
1. Gutachter:	Prof. Christian Ochsenfeld
2. Gutachterin:	Prof. Regina de Vivie-Riedle
Mündliche Prüfung am:	04.06.2012

”Was kann der Mensch im Leben mehr gewinnen,
Als dass sich Gott-Natur ihm offenbare?
Wie sie das Feste lässt zu Geist verrinnen,
Wie sie das Geisterzeugte fest bewahre.”

Johann Wolfgang von Goethe

Acknowledgements

I hereby wish to express my gratitude to all persons who have supported, stimulated or influenced me in some way along the path that led to this dissertation.

To begin, I'd like to thank my doctoral supervisor Prof. Christian Ochsenfeld for the opportunity of working in his research group, for his kindness and his support.

Further thanks are due to all present and former group members Dr. Matthias Beer, Dr. Tomasz Berezniak, Iris Blank, Dr. Bernd Doser, Denis Flaig, Dr. Matti Hanni, Jutta Kluge, Dr. Jörg Kußmann, Dr. Daniel Lambrecht, Dr. Martin Löffler, Arne Lünser, Dr. Boris Maryasin, Marina Maurer, Simon Maurer, Sven Roßbach, Dr. Keyarash Sadeghian, Johannes Schäffer, Dr. Benedikt Schmitt, Dr. Sabine Schweizer, Helga Senft, Philipp Snajberk, Dr. Chris Sumowski and Dr. Jan Zienau for their friendliness and the good atmosphere within the group.

I wish to express particular gratitude to Dr. Martin Löffler, Dr. Jan Zienau and Simon Maurer for proofreading and for valuable comments, and thank Dr. Martin Löffler in general for many stimulating and insightful discussions and reflections about and beyond science.

I also thank Prof. Regina de Vivie-Riedle for providing the co-referee report to this dissertation.

Among the many great people whom I thank for their friendship, though I shall not list them all explicitly in order not to omit anyone, I'd like to acknowledge particularly Dr. Karsten Holtin for numerous hours of intensive discussions on physical chemistry during the studies, as well as Dr. Max Wohland and Dr. Jochen Neumaier with whom I very much enjoyed exploring the practical benefits of a chemical education in various fermentation experiments.

Finally, I wish to express my deepest gratitude to my mother and my late father, as well as my brothers, for their everlasting support throughout my life.

Zusammenfassung

Die Berechnung von Korrelationsenergien mittels Møller-Plesset Störungstheorie zweiter Ordnung (MP2) [1] in ihrer kanonischen Formulierung ist aufgrund des hohen $\mathcal{O}(N^5)$ Skalenvhaltens nur für Systeme mit höchstens ca. 100 Atomen möglich. Die Entwicklung niedrigskalierender Reformulierungen ist daher unerlässlich, um den Anwendungsbereich dieser Methode auf größere Systeme zu erweitern. Einen möglichen Ansatz hierfür liefert die von Almlöf und Häser vorgeschlagene Laplace Transformation [2,3], mit deren Hilfe MP2 in der Basis der atomzentrierten Orbitale dargestellt werden kann (AO-MP2) [4]. Da diese Orbitale per Definition lokal sind, können Vorabschätzungsverfahren verwendet werden, um numerisch insignifikante Beiträge zur Energie zu vernachlässigen und somit das Skalenvverhalten zu reduzieren.

Das Ziel dieser Dissertation bestand in der Untersuchung Cholesky-zerlegter pseudo-Dichtematrizen (CDD) innerhalb dieses AO-MP2 Ansatzes. Aus technischen Gründen war die von Doser *et al.* vorgestellte AO-MP2 Implementierung [5–7] auf die MP2-Energiekomponente antiparallelen Spins (opposite spin) beschränkt und somit auf ein empirisches Skalierungsverfahren angewiesen (Scaled Opposite Spin—SOS-MP2) [8]. Die aus der Cholesky-Zerlegung [9] der auftretenden pseudo-Dichtematrizen resultierende CDD-MP2 Methode [10,11] enthält im Gegensatz dazu auch die Komponente parallelen Spins, wodurch die *ab initio* Eigenschaft wieder hergestellt wird. Da für Systeme mit lokalisierter Elektronenstruktur die über die Cholesky-Zerlegung erzeugten Orbitale ebenfalls lokal sind, lassen sich innerhalb dieses Ansatzes ebenso Vorabschätzungsverfahren [12] zur Reduktion des Skalenvhaltens anwenden. Für kleine Systeme konnte nachgewiesen werden, dass die implementierte CDD-MP2 Methode im Prinzip lineares Skalenvverhalten ermöglicht. Bei mittelgroßen Systemen mit vernünftigen Basissätzen treten jedoch schwerwiegende Leistungs- und Genauigkeitsprobleme auf, welche im Detail erörtert werden. Da diese Probleme dem integral-direkten CDD Ansatz leider innewohnen, und somit nicht behoben werden konnten, erscheint die Anwendbarkeit dieser Methode auf realistische Fragestellungen nicht möglich.

Mit diesem Projekt eng verwoben war die Erweiterung der RI-CDD-MP2 Methode von Zienau *et al.* [10,11,13], welche sich aus der Kombination des CDD-Ansatzes mit der "Resolution of the Identity" (RI) [14–18] zur Zerlegung und näherungsweise Berechnung der Integrale ergibt. Diese Methode hatte sich zwar bereits als sehr effizient erwiesen, war in ihrer bisheri-

gen Implementierung jedoch ebenfalls auf die MP2-Komponente antiparallelen Spins beschränkt gewesen. Daher wurde sie im Rahmen dieser Arbeit zur vollständigen Berechnung der MP2-Energie ergänzt. Hierbei wurde zudem die bisher verwendete Integralvorabschätzung aktualisiert, und eine Parallelisierung verschiedener Schritte des Algorithmus durchgeführt. Aufgrund ihres quadratischen Skalenverhaltens stellt diese Methode in ihrer vervollständigten Form nun eine ernsthafte und effiziente Alternative zu kanonischem RI-MP2 für die Berechnung von MP2-Energien mittlerer bis großer Moleküle dar.

Abschließend wurde untersucht, inwiefern sich die von Boman *et al.* [19] vorgeschlagene methodenspezifische Cholesky-Zerlegung (Method Specific Cholesky Decomposition – MSCD) zusammen mit CDD anwenden lässt, um eine effiziente Zerlegung der Zwei-Elektronen-Integralmatrix durchzuführen. Dies wird durch die Eigenschaft motiviert, dass die Cholesky-Zerlegung eine beliebige Genauigkeit der zerlegten Größe zulässt, wohingegen die RI-Näherung aufgrund der Unvollständigkeit der verwendeten Hilfsbasen mit einem systematischen Fehler behaftet ist. Obwohl diese vorläufige Untersuchung auf sehr kleine Systeme beschränkt werden musste, konnten bereits bedeutende Verringerungen des Zerlegungsranges verglichen mit konventioneller Cholesky-Zerlegung beobachtet werden, wodurch das Potential der MSCD für diese Fragestellung nachgewiesen werden konnte.

Abstract

The evaluation of correlation energies in the canonical formulation of second order Møller-Plesset Perturbation Theory (MP2) [1] is limited to systems of about 100 atoms, due to the method's steep $\mathcal{O}(N^5)$ scaling. In order to extend the method's applicability to larger systems, it is therefore imperative to develop alternative formulations that allow for efficient scaling reduction. One such approach is the Laplace transform formalism introduced by Almlöf and Häser [2, 3], with which MP2 can be expressed in the basis of atom-centered orbitals (AO-MP2) [4], whose local character allows to take advantage of the short range of correlation effects. The overall scaling can thus be reduced through the application of integral pre-selection schemes to discard all numerically irrelevant contributions to the energy.

This dissertation is concerned with the study of Cholesky decomposed pseudo-density (CDD) matrices within this AO-MP2 scheme. For technical reasons, namely, the AO-MP2 implementation of Doser *et al.* [5–7] is restricted to the evaluation of the opposite spin component of MP2, and is thus bound to the empirical scaled opposite spin parametrization procedure (SOS-MP2) [8]. Applying a Cholesky decomposition [9] to the occurring pseudo-density matrices, the same spin component required for full MP2 energies is naturally included in the resulting CDD-MP2 method [10, 11], whereby the *ab initio* character is restored. The investigation of the CDD-approach was further motivated by the fact that the orbitals generated by the decomposition are localized (for electronically non-delocalized systems), and thus allow for the pre-selection of only numerically significant integrals [12]. However, although it could be shown on simple systems that the method does in principle scale linearly, its application to even moderately sized systems with large basis sets is yet hampered by severe technical and numerical difficulties, which are analysed and discussed in detail.

Another closely related project has been to extend the RI-CDD-MP2 algorithm of Zienau *et al.* [10, 11, 13], an alternative to the integral-direct CDD-MP2 method based on the resolution of the identity (RI) [14–18] for integral approximation. This implementation, which had already been established as very efficient, was still limited to the opposite spin part and has here been completed for calculation of total MP2 energies, along with an adaptation of the previously employed integral pre-selection protocol, and a parallelization of various steps of the algorithm. Due to its quadratic scaling, it now represents a serious alternative to canonical RI-MP2 for

calculations on medium to large systems.

As a final project, it has been studied whether the method specific Cholesky decomposition (MSCD) proposed in general terms by Boman *et al.* [19] can be efficiently used to obtain a low-rank factorization of the two-electron integral matrix in the specific context of CDD-MP2. This is motivated by the arbitrary numerical accuracy provided by the Cholesky decomposition, whereas the finiteness of pre-optimized RI auxiliary basis sets invariably introduces a systematic error in the approximated entities. Despite the fact that these preliminary investigations were restricted to very small systems, the observed rank reduction with respect to conventional Cholesky decomposition were significant enough to demonstrate the potential of the MSCD for this particular purpose.

Contents

1	Introduction	1
2	Fundamentals	5
2.1	The Schrödinger equation	6
2.2	The Born-Oppenheimer approximation	6
2.3	The Hartree-Fock method	7
2.3.1	Hartree-Fock theory	8
2.3.2	LCAO and the Roothaan-Hall equation	10
2.3.3	Symmetric orthonormalization	11
2.3.4	The Fock matrix and the Hartree-Fock energy	12
2.3.5	The Self-Consistent Field method	13
2.4	The density operator and matrix	14
2.5	Linear scaling Hartree-Fock	16
2.5.1	Integral screening and direct SCF	17
2.5.2	Density-matrix based Hartree-Fock	17
2.5.3	Linear scaling Fock matrix construction	18
3	Electron correlation	21
3.1	Configuration Interaction and Coupled Cluster	22
3.2	The Random Phase Approximation	23
3.3	Many-Body Perturbation Theory	24
3.3.1	Møller-Plesset Perturbation Theory	25
3.3.2	Second order Møller-Plesset theory (MP2)	26
3.3.3	Parametrized MP2	27
3.4	Local correlation methods	27
3.4.1	Orbital localization methods	28
3.4.2	Local MP2	30
3.5	Explicitly correlated methods	31

4	Laplace-Transform Møller-Plesset Perturbation Theory	33
4.1	Laplace-Transform MP2	33
4.1.1	Orbital invariance and relation to local MP2	36
4.1.2	Relation to Lyapunov equations	37
4.2	Atomic-orbital MP theory (AO-MP)	38
4.2.1	Pseudo-density matrices and operators	39
4.2.2	The MP n energy as a functional of the density matrix	40
5	Cholesky Decomposed Densities in MP2	41
5.1	Cholesky-decomposed pseudo-density matrices (CDD-MP2)	41
5.1.1	Cholesky decomposition	43
5.1.2	CDD as localized pseudo-MO	43
5.1.3	Integral screening	44
5.2	Implementation of CDD-MP2	47
5.2.1	The Cholesky decomposition algorithm	48
5.2.2	Screening prerequisites	50
5.2.3	External screening for FTI	51
5.2.4	Internal screening for HTI	52
5.2.5	Formation of the HTI	54
5.2.6	Last transformations to FTI	54
5.2.7	Contraction to MP2 energy	56
5.3	Results and discussion	56
5.3.1	Matrix sparsities	57
5.3.2	Integral numbers and accuracy	58
5.3.3	Timings	61
5.3.4	Conclusion	65
6	Resolution of the Identity CDD-MP2	81
6.1	Resolution of the Identity and Density Fitting	81
6.2	The RI-CDD-MP2 algorithm	83
6.3	OMP Parallelization	84
6.4	Results and discussion	85
6.4.1	Timings	85
6.4.2	Accuracy	86
6.4.3	Parallel RI-CDD-MP2	88
6.4.4	Conclusion	90

CONTENTS

7	Method Specific Cholesky Decomposition	103
7.1	Decompositions of the ERI matrix	103
7.1.1	Resolution of the Identity	104
7.1.2	Cholesky decomposition	104
7.1.3	Tensor approximation	105
7.2	Method specific Cholesky decomposition (MSCD)	106
7.2.1	Technical details	108
7.2.2	Results and discussion	108
8	Conclusion and outlook	111
A	Exponentially weighted energy differences in MP2	117
B	Kronecker Product Approximation of ERI	121
C	List of abbreviations	123
D	List of academic teachers	133

CONTENTS

Chapter 1

Introduction

With the advent of the quantum theory at the beginning of the twentieth century, our world view and understanding of Nature have experienced a revolution—both practical and philosophical—similar to none other in the history of science. Although the hitherto prevailing *classical* theory, consisting in essence of Newtonian mechanics and Maxwellian field theory, was very successfully applied to a wide range of fields ranging from mechanics, hydrodynamics, electrodynamics to astronomy, the puzzling and seemingly paradoxical behaviour of Nature at its microscopic level—stability of atoms, wave-particle duality—has required the development of a fundamentally different theory of quantum phenomena. While the classical theory yet retains validity at the macroscopic level of everyday experience, quantum theory has required a profound reshaping of concepts pertaining to the small scale, and caused a partial abandon of principles that were previously taken for granted. Within the classical world view, the evolution of physical systems, which were conceived of as consisting of palpable, *material* objects possessing well-defined physical quantities such as position and momentum, was mostly described in the relatively pictorial and intuitive language of differential equations in Euclidian space. In quantum mechanics, physically observable entities lose their concreteness and mental graspability in favour of a highly abstract mathematical formalism, in which they become associated with operators acting on functions in configuration space, and the uncertainty principle sets a limit to the simultaneous definition of conjugate quantities. Also, in stark contrast to the entirely *deterministic* classical paradigm, the wavefunction of quantum theory merely represents *potentialities*, i.e. it allows assertions of a statistical nature, which only become *factual* through an act of measurement. Therefore, at least according to the Copenhagen interpretation of quantum theory, the Cartesian split, one of the basic tenets of classical physics which ascribed an independent, *objective* reality to physical phenomena and viewed us as separate, non-disturbing observers, cannot be rigorously upheld anymore. Namely, the *interaction* of a measurement apparatus with the quantum system that is under inquiry, which is required in order to make information accessible to an experimenter, blurs the clear-cut separation of subject and object. Further, the inevitable

disturbance that the measuring process creates, according to the uncertainty principle, relativises the validity of classical concepts such as, for instance, particles moving on well-defined trajectories, to the extent that they may be considered as useful visualisations which, however, possess no intrinsic *objective* value. Due to these implications, and along with the fact that for communication and interpretation, the outcome of measurements can only be expressed in *language*, whose evolution and constituent notions have been shaped by experience at the macroscopic level, and is therefore only of limited applicability to speak of sub-atomic phenomena, the quantum theory, more than any realization before, forces upon us the fundamental distinction between Reality and models of it, between how Nature actually *is*, and how we *describe* it. In the words of Heisenberg [20]: "*Natural science does not simply describe and explain Nature as it is 'in itself'. It is rather part of the interplay between Nature and ourselves. It describes Nature as exposed to our questioning and our methods.*"

On a more practical footing, chemistry has also been profoundly influenced by the new insights brought about by the quantum theory. Hitherto, chemistry had been based on many rather vaguely defined—though of course still useful—concepts and notions about the structure and behaviour of the constituent parts of the substances it dealt with, such that in many respects, it yet had more traits of a *knack* or *craft* than a *science* on a similar level as physics. The accurate description of atomic structure and behaviour provided by quantum mechanics, however, put chemistry on a firmer theoretical basis and at the same time dissolved its until then sharply marked border with physics. Not only did quantum theory explain the *systematic* character of the empirically established "periodic system of elements", the foundation of chemistry, but also provided explanations for phenomena ranging from the chemical bond to the observed spectra of atoms and molecules. This better understanding of molecular structure and behaviour has in turn led to formidable advances in synthetic and especially analytical techniques, so that today spectroscopic methods belong to the basic tools of any chemist.

The field of quantum chemistry, whose primary concern is the solution of Schrödinger's equation for atoms and molecules in order to provide theoretical insights on energetics and properties, is nowadays a well-established, though still actively developing part of chemistry. In many situations, the accurate reproduction of difficultly interpretable experimental data by theoretical means is indispensable to a correct understanding, and theoretical predictions are often helpful for the guidance of experimental work. Although the analytical solution of Schrödinger's equation is not possible for systems of chemical interest, there yet exist many methods which allow for systematic numerical approximation of the solution. In essence, the methodical approach hereby consists of transforming Schrödinger's differential equation into a matrix mechanics problem by the introduction of an appropriate basis set, in order to make it accessible to treatment by the methods of algebra. This algebraization, on the other hand, is accompanied by a high computational complexity of the resulting algorithms, and as computer power does not evolve

as rapidly as the interest in investigating ever larger systems, it is imperative to develop more efficient, low-complexity alternatives while retaining control over the numerical precision.

The method due to Hartree and Fock (HF) [21,22], which constitutes the foundation of *ab initio* quantum chemistry, i.e. the branch that endeavours to obtain a solution to Schrödinger’s equation from first principles only, without having recourse to empirical parameters—apart from the inevitably parametrized basis sets, of course—is one of the methods with lowest computational cost. In spite of a formal $\mathcal{O}(M^4)$ complexity— M is some suitable measure of the system size—in its most naive implementation, algorithmic improvements have reduced the method’s scaling to linear such that systems of the order of thousands of atoms are nowadays routinely treated. Chapter 2 of this dissertation is concerned with a review of the HF method and its scaling reduction.

Due to its simplistic mean-field description of electronic interactions, however, the Hartree-Fock method lacks the so-called electron correlation and thus fails to fully recover the energy and properties derived from it. Yet, for the accurate treatment of most situations in chemistry, it is of paramount importance to take electron correlation effects into account. For that end, a variety of post-HF methods exists, which are usually grouped into a hierarchy in order of increasing accuracy, and the correlation energy can—in principle—be exactly evaluated, though of course with a computational effort that rapidly becomes intractable. Among this hierarchy of methods, second order Møller-Plesset perturbation theory (MP2) [1] has assumed a key position ever since its inception in 1934, due to its good balance of accuracy and complexity. However, despite being the formally simplest correlation method, the computational scaling of already $\mathcal{O}(M^5)$ in its canonical formulation quickly leads to insurmountable difficulties which preclude the treatment of molecules exceeding approximately 100 atoms with reasonably sized basis sets. In the last decades, significant progress has been made in the search for efficient low-scaling reformulations of MP2, but there still remains the need to explore further possibilities in order to extend the perturbational treatment of electron correlation to yet larger systems. The common denominator of all these low-complexity approaches is the introduction of a basis of local orbitals, which enables the exploitation of the short-range character of electronic correlation effects to reduce the number of numerically significant interaction terms, and hence of the algorithms’ overall scaling. An overview of correlation methods and approaches for the scaling reduction of MP2 is given in chapter 3.

Along the aforementioned line of thought, the subject of the present thesis has been the study of one such alternative approach to MP2, based on Cholesky decomposed density matrices within the framework of Laplace transform MP2 theory [2,3]. The application of the Laplace transformation to the energy denominator allows to reformulate MP2 in terms of atom-centered orbitals (AO-MP2) [4], which are local by definition and hence an ideal basis for scaling reduction. However, for technical reasons, the previously completed linearly scaling AO-MP2 implemen-

tation of Doser *et al.* [5–7], upon which the work presented here is based, was restricted to the computation of merely the opposite spin component of the MP2 energy, and was therefore bound to the empirical scaled opposite spin (SOS-MP2) parametrization scheme [8]. It was then realized that through the application of a Cholesky decomposition [9] to the occurring pseudo-density matrices (CDD—Cholesky Decomposed Densities), the exchange-type term required for the same spin component of MP2 is straightforwardly included in the resulting CDD-MP2 method [10, 11, 23], hereby restoring the method’s *ab initio* character. Since for systems having a localized electronic structure, the pseudo-densities possess a sparse character that is preserved by the Cholesky decomposition, their Cholesky factors represent localized orbitals, and as such establish CDD-MP2 as a potentially linearly scaling MP2 reformulation. Due to these promising features, the implementation and study of the CDD-MP2 method has constituted the main topic of this dissertation. The Laplace transform formalism of MP2 is described in chapter 4, where the relationship of this formalism to Lyapunov equations, which seems to have so far been overlooked, is pointed out along with potential applications. The implementation of CDD-MP2 and the obtained results are then described and discussed in chapter 5.

Parallely to the CDD-MP2 implementation that has been in our focus, the CDD procedure has also been used by Zienau *et al.* in conjunction with the resolution of the identity (RI) for the approximation of two-electron repulsion integrals in a RI-CDD-MP2 implementation [10, 11, 13]. This algorithm, however, had only been implemented for calculation of the opposite spin component, and as a side project of the present thesis, it has been completed to include the exchange-type term as well, for computation of total MP2 energies. In addition, a parallelization of the algorithm has been undertaken, and results on both these contributions to RI-CDD-MP2 are given in chapter 6.

As a final project, it has been investigated whether the method specific Cholesky decomposition (MSCD) proposed by Boman *et al.* [19], which represents a promising alternative to the RI-approximation of two-electron integrals, may be beneficially applied within the framework of CDD-MP2. Results thereof are presented in chapter 7.

To conclude this dissertation, two propositions of potentially interesting approaches are included as appendices. Although they could not be further investigated, they are deemed worthy of mention and are thus presented here without results. On a different footing than the usual approaches which are grounded in exploiting the spatial separation of electron pairs, it is proposed in appendix A that a reduction of the number of significant interaction terms for MP2 may possibly be achieved with an exponential weighting of the orbital energy differences occurring in the canonical expression. Appendix B is an illustration of how a matrix Kronecker Product Approximation (KPA) of the matrix of two-electron integrals may be interesting to study for application with MP2, as an alternative related to the canonical product (CP) format decomposition proposed by Benedikt *et al.* [24].

Chapter 2

Fundamentals

The first chapter of this dissertation shall be dedicated to reviewing some fundamentals of quantum chemistry in order to set the general context as well as establish the notation for our own work exposed in later chapters. The presentation naturally begins with the Schrödinger equation, along with the Born-Oppenheimer approximation that underlies nearly all non-relativistic quantum chemistry.

Then follows a thorough, though condensed, presentation of the Hartree-Fock (HF) method, whose importance to *ab initio* quantum chemistry cannot be overemphasized, of course as an electronic structure method in its own right, but particularly as the basis of the hierarchy of wavefunction based methods accounting for electron correlation, which allows systematic approximation to the solution of the Schrödinger equation. The essential concepts of the HF-method are recapitulated, followed by the algebraization that makes the HF equations accessible to numerical treatment, and a description of the SCF-cycle required for the solution of the HF-problem.

The theory of density matrices is then described, which has achieved such importance in modern quantum chemistry that a presentation of basic concepts is essential. Also, this formalism shall be encountered in slightly different form in chapter 4.

This chapter concludes with an overview of techniques that have been devised to reduce the computational scaling of the HF method, thus making it applicable to systems intractable with conventional HF-implementations. This further lays the foundation for the later chapters concerned with scaling reduction of correlation methods.

Though density functional theory (DFT) methods have long been established as viable, low-cost alternatives to wavefunction based methods, DFT will be entirely omitted, as the presented work has no connection with developments in that field.

2.1 The Schrödinger equation

The central equation of (non-relativistic) quantum mechanics is Schrödinger's wave equation [25]

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}, t) = \hat{\mathcal{H}} \Psi(\mathbf{r}, t), \quad (2.1)$$

where the wavefunction $\Psi(\mathbf{r}, t)$ describes the evolution of the physical system specified by the Hamilton operator $\hat{\mathcal{H}}$. The wavefunction explicitly depends on the time t and spatial coordinates \mathbf{r} only, whereas spin is accounted for by appending spin functions to the spatial wave functions. States having a well-defined, constant energy, i.e. not changing with time, are called stationary and are described by the eigenfunctions Ψ_k of the Hamiltonian. In that case, insertion of the eigenvalue E_k and integration of (2.1) yields [26]

$$\Psi_k(\mathbf{r}, t) = \psi_k(\mathbf{r}) e^{-\frac{i}{\hbar} E_k t}. \quad (2.2)$$

The energy eigenvalues and their corresponding eigenfunctions are thus obtained from the time-independent eigenvalue equation

$$\hat{\mathcal{H}} \psi_k(\mathbf{r}) = E_k \psi_k(\mathbf{r}), \quad (2.3)$$

whose solution is one of the central goals of quantum chemistry, in particular for the state having lowest energy, the ground state, which we shall exclusively be concerned with in this treatise. The Ψ_k provide a basis with which any wavefunction can be expanded according to

$$\Psi = \sum_k C_k \Psi_k. \quad (2.4)$$

The Hamiltonian of a molecular system (in the absence of external fields) is given in atomic units by

$$\begin{aligned} \hat{\mathcal{H}} &= -\frac{1}{2} \sum_i^{N_{\text{el}}} \nabla_i^2 - \frac{1}{2} \sum_A^{N_{\text{atom}}} \frac{1}{M_A} \nabla_A^2 - \sum_i^{N_{\text{el}}} \sum_A^{N_{\text{atom}}} \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} + \sum_i^{N_{\text{el}}} \sum_{j < i} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_A^{N_{\text{atom}}} \sum_{B < A} \frac{Z_A Z_B}{|\mathbf{R}_A - \mathbf{R}_B|} \\ &= \hat{\mathcal{T}}_{\text{el}} + \hat{\mathcal{T}}_{\text{n}} + \hat{\mathcal{V}}_{\text{el,n}} + \hat{\mathcal{V}}_{\text{el,el}} + \hat{\mathcal{V}}_{\text{n,n}}. \end{aligned} \quad (2.5)$$

Indices i, j and A, B denote electrons and atomic nuclei, respectively, and $\hat{\mathcal{T}}$, $\hat{\mathcal{V}}$ represent operators of kinetic and potential energy. Due to the dependence of the wavefunction on the $3N$ spatial coordinates of the particles, analytical solution of the differential equation (2.3) is possible only for few model systems, and many-particle systems can only be solved approximately.

2.2 The Born-Oppenheimer approximation

The Born-Oppenheimer approximation, which underlies nearly all of quantum chemistry, is based on the assumption that the total wavefunction is factorizable into electronic and nuclear components, i.e.

$$\psi(\mathbf{r}_{\text{el}}, \mathbf{R}_{\text{n}}) = \psi_{\text{el}}(\mathbf{r}_{\text{el}}; \mathbf{R}_{\text{n}}) \psi_{\text{n}}(\mathbf{R}_{\text{n}}). \quad (2.6)$$

This is justified by the high difference in mass, but comparable momentum, of electrons and nuclei, from which electrons may be assumed to react instantaneously to nuclear displacements, such that their motion may be considered for static nuclear configurations. The nuclear repulsion potential $V_{n,n}$ thus becomes constant, and the electronic wavefunction, along with the electronic energy, depends only parametrically on the nuclear coordinates. It follows that after solution of the electronic Schrödinger equation

$$\hat{\mathcal{H}}_{\text{el}}\psi_{\text{el}} = E_{\text{el}}\psi_{\text{el}}, \quad (2.7)$$

with the electronic Hamiltonian operator

$$\hat{\mathcal{H}}_{\text{el}} = \hat{\mathcal{T}}_{\text{el}} + \hat{\mathcal{V}}_{\text{el},n} + \hat{\mathcal{V}}_{\text{el,el}}, \quad (2.8)$$

the total energy is obtained as the sum $E = E_{\text{el}} + V_{n,n}$ of the electronic energy and nuclear repulsion potential. The manifold of electronic energies for all nuclear conformations is called the Born-Oppenheimer *hypersurface*. Within this conceptual framework, geometry optimizations can be performed by searching stationary points on the potential energy surface according to

$$\text{grad}_A E = \frac{\partial E}{\partial \mathbf{R}_A} = 0. \quad (2.9)$$

Further, by expanding the energy in a Taylor series around a stationary point on the potential surface, molecular force constants, and from these harmonic vibrational frequencies, can be computed as the mixed second derivatives of the energy with respect to the nuclear coordinates:

$$k_{AB} = \frac{\partial^2 E}{\partial \mathbf{R}_A \partial \mathbf{R}_B}. \quad (2.10)$$

On the same footing, a whole variety of molecular properties are calculated as (mixed) derivatives of the energy with respect to particular quantities, e.g. nuclear magnetic resonance (NMR) shieldings as derivatives with respect to the external magnetic field and nuclear magnetic momenta.

2.3 The Hartree-Fock method

The Hartree-Fock (HF) method [21, 22, 27] is one of the most widely applied methods for the solution of the electronic Schrödinger equation. Despite its simplistic description of the inter-electronic potential as a mean-field, it provides fairly accurate results at relatively low computational cost. It is further the basis for the hierarchy of *ab initio* methods that take electron correlation into account, as will be discussed in chapter 3.

2.3.1 Hartree-Fock theory

The Hartree-Fock method is based on an independent particle model (IPM), where it is assumed that the total wavefunction can be represented as an antisymmetrized product of mutually orthogonal one-particle spin-orbitals formed as products of spatial orbitals and spin functions

$$\varphi_i(\mathbf{x}) = \begin{cases} \phi_{\frac{i+1}{2}}(\mathbf{r})\alpha(\omega), & i \text{ odd} \\ \phi_{\frac{i}{2}}(\mathbf{r})\beta(\omega), & i \text{ even} \end{cases} \quad i \in N_{\text{el}} \quad (2.11)$$

The variational method [28] is then applied to determine the spin-orbitals whose Slater-determinant

$$\Phi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_{\text{el}}}) = \frac{1}{\sqrt{N_{\text{el}}!}} \det \left(\varphi_i(\mathbf{x}_j) \right)_{i,j \in N_{\text{el}}} \quad (2.12)$$

represents the best approximation to the exact ground state wavefunction of a system described by the electronic Hamiltonian $\hat{\mathcal{H}}$ (subscripts denoting electronic quantities shall henceforth be omitted). According to the variational theorem, the obtained energy represents an upper bound to the exact electronic energy.

The expectation value of the Hamiltonian for a Slater-determinant is given by the Slater-Condon rules as

$$\begin{aligned} E_{\text{HF}} &= \langle \Phi | \hat{\mathcal{H}} | \Phi \rangle \\ &= \sum_i^{N_{\text{el}}} \langle \varphi_i | \hat{H} | \varphi_i \rangle + \frac{1}{2} \sum_{i,j}^{N_{\text{el}}} \langle \varphi_i \varphi_j | (1 - \mathcal{P}_{(ij)}) \hat{C} | \varphi_i \varphi_j \rangle \\ &= \sum_i^{N_{\text{el}}} \langle i | \hat{H} | i \rangle + \frac{1}{2} \sum_{i,j}^{N_{\text{el}}} (\langle ij | ij \rangle - \langle ij | ji \rangle), \end{aligned} \quad (2.13)$$

where the Coulombic repulsion operator $\hat{C} = 1/|\mathbf{r}_1 - \mathbf{r}_2|$ has been introduced, and \mathcal{P} denotes the permutation arising from the alternating property of the determinant, the so-called *exchange* term. The *core* Hamiltonian operator

$$\hat{H}_{\text{core}}(1) = -\frac{1}{2} \nabla_1^2 - \sum_A^{N_{\text{atom}}} \frac{Z_A}{|\mathbf{r}_1 - \mathbf{R}_A|} \quad (2.14)$$

represents the kinetic energy and the nuclear attraction potential of a given electron, and the two-electron repulsion integrals (ERI)

$$\langle ij | ij \rangle \equiv \langle \varphi_i \varphi_j | \hat{C} | \varphi_i \varphi_j \rangle = \int \int \varphi_i^*(1) \varphi_j^*(2) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \varphi_i(1) \varphi_j(2) d\mathbf{x}_1 d\mathbf{x}_2 \quad (2.15)$$

describe the inter-electronic interaction.

Minimization of the expectation value (2.13) through the use of Lagrange's method of underdetermined multipliers to impose the constraints that the orbitals remain orthonormal, i.e. that

$\langle \varphi_i | \varphi_j \rangle = \delta_{ij}$, leads to the *Hartree-Fock* equations (sometimes also *Fock-Dirac* equations [29])

$$\hat{F}\varphi_i = \sum_j^{N_{\text{el}}} \epsilon_{ij} \varphi_j, \quad (2.16)$$

with the Lagrange-multipliers ϵ_{ij} . The Fock-operator is an effective one-electron Hamiltonian given by

$$\hat{F}(1) = \hat{H}(1) + \sum_j^{N_{\text{el}}} \langle \varphi_j | (1 - \mathcal{P}_{(ij)}) \hat{C} | \varphi_j \rangle = \hat{H}(1) + \sum_j^{N_{\text{el}}} \left(\hat{J}_j(1) - \hat{K}_j(1) \right), \quad (2.17)$$

where the one-electron Coulomb-operator

$$\hat{J}_j(1)\varphi_i(1) = \int d\mathbf{x}_2 \varphi_j^*(2) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \varphi_j(2) \varphi_i(1) \quad (2.18)$$

describes the repulsive Coulomb potential between an electron and the field of a second one in the orbital φ_j . The so-called exchange operator \hat{K} has no classical interpretation, but is defined by its action on an orbital as

$$\hat{K}_j(1)\varphi_i(1) = \langle \varphi_j | \mathcal{P}_{(ij)} \hat{C} | \varphi_j \rangle | \varphi_i \rangle = \int d\mathbf{x}_2 \varphi_j^*(2) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \varphi_i(2) \varphi_j(1) \quad (2.19)$$

The sum of Coulomb and exchange operators is termed Hartree-Fock-potential

$$\hat{V}^{\text{HF}}(1) = \sum_j^{N_{\text{el}}} \left(\hat{J}_j(1) - \hat{K}_j(1) \right) \quad (2.20)$$

and describes the interaction of an electron with the *averaged* electrostatic field of all others.

In equation (2.16), the Hartree-Fock equations do not yet represent an eigenvalue problem. The matrix ϵ of Lagrange multipliers, however, being Hermitean, can be brought to diagonal form by unitary transformation (rotation) of the orbitals, whereas the Fock operator is invariant to such a transformation. This leads to the canonical Hartree-Fock equations

$$\hat{F}\varphi'_i = \epsilon'_i \varphi'_i, \quad (2.21)$$

but the primes shall be omitted in the following. Due to the dependence of the Fock operator on the orbitals, they lack the linearity of true eigenvalue equations, but may be viewed as *pseudo*-eigenvalue equations.

The simplest possible application of the HF method is to consider only closed-shell systems by restricting the total spin to zero, and is hence termed *restricted* Hartree-Fock (RHF). The next step that is required in order to facilitate handling of the above equations then consists of performing the integration over the spin variable ω (see equation (2.11)), hence obtaining a formulation in terms of spatial orbitals ϕ_i only. The closed-shell Fock operator is thus obtained as

$$\hat{F}(1) = \hat{H}(1) + \sum_j^{N_{\text{occ}}} \left(2\hat{J}_j(1) - \hat{K}_j(1) \right), \quad (2.22)$$

where $N_{\text{occ}} = N_{\text{el}}/2$, each spatial orbital being doubly occupied, and the expectation value of the energy after spin integration is given by

$$E_{\text{HF}} = 2 \sum_i^{N_{\text{occ}}} (i|\hat{H}|i) + \sum_{i,j}^{N_{\text{occ}}} (2(ii|jj) - (ij|ji)). \quad (2.23)$$

The Mulliken notation $(ii|jj) = \langle ij|ij \rangle$ shall from now on be used in expressions involving spatial orbitals only.

2.3.2 LCAO and the Roothaan-Hall equation

For numerical treatment, the Hartree-Fock integro-differential equations are converted into a generalized matrix eigenvalue problem that is solvable using algebraic techniques. Following the procedure of Roothaan [30], the HF-orbitals are expanded in a basis $\{\chi_\nu(\mathbf{r})\}_{\nu \in N}$ of atom-centered spatial orbital functions (Linear Combination of Atomic Orbitals—LCAO), such that the problem is shifted to the determination of optimal expansions coefficients.

For convenience, the basis functions shall be conceived as being arranged into a row vector

$$|\chi\rangle \equiv \left(|\chi_1\rangle, \dots, |\chi_N\rangle \right) \quad (2.24)$$

and we define $\langle \chi| \equiv |\chi\rangle^\dagger$, but since only real-valued basis functions will be considered, adjunction reduces to transposition. In the context of this vector arrangement, Dirac's notation [31]—though with parentheses to indicate spatial functions—is used to suggest that conjunction of an element $\langle \chi_\mu|$ from $\langle \chi|$ with an element from $|\chi\rangle$ implies the inner product $\langle \chi_\mu|\chi_\nu\rangle$. It follows from this definition that the outer product of $\langle \chi|$ with $|\chi\rangle$ yields the *overlap* matrix $\mathbf{S} = \langle \chi|\chi\rangle$, the $N \times N$ Gramian matrix of the AO-basis, i.e. the positive definite matrix which induces the metric in the spanned space.

Now, defining a corresponding vector $|\phi\rangle$ for the MO and a matrix \mathbf{C} of expansion coefficients, the LCAO reads

$$|\phi\rangle = |\chi\rangle \mathbf{C}, \quad (2.25)$$

and upon inserting the above expansion into the Hartree-Fock equations $\hat{F}|\phi\rangle = |\phi\rangle \epsilon$, one obtains

$$\hat{F}|\chi\rangle \mathbf{C} = |\chi\rangle \mathbf{C} \epsilon, \quad (2.26)$$

where ϵ denotes the diagonal matrix of orbital energies. Taking the outer product with $\langle \chi|$ then leads to

$$\langle \chi|\hat{F}|\chi\rangle \mathbf{C} = \langle \chi|\chi\rangle \mathbf{C} \epsilon, \quad (2.27)$$

such that $\mathbf{F} = \langle \chi|\hat{F}|\chi\rangle$ can be identified as the matrix representation of the Fock operator in the AO-basis. The Roothaan-Hall equation [30]

$$\mathbf{F}\mathbf{C} = \mathbf{S}\mathbf{C}\epsilon \quad (2.28)$$

thus finally results. Since the solution of this equation yields a $N \times N$ coefficient matrix, it follows that in addition to the N_{occ} occupied MO $|\phi_o\rangle$, this generates a set of $N_{\text{virt}} = N - N_{\text{occ}}$ *virtual*, i.e. unoccupied, orbitals $|\phi_v\rangle$. Due to the mutual orthogonality of the MO, the total space \mathcal{F} spanned by

$$\begin{aligned} |\phi\rangle &= |\phi_o\rangle \oplus |\phi_v\rangle \\ &\equiv \left(|\phi_1\rangle, \dots, |\phi_{N_{\text{occ}}}\rangle, |\phi_{N_{\text{occ}}+1}\rangle, \dots, |\phi_{N_{\text{occ}}+N_{\text{virt}}}\rangle \right) \end{aligned} \quad (2.29)$$

thus consists of the two complementary occupied and virtual subspaces

$$\mathcal{F} = \mathcal{O} \oplus \mathcal{V}. \quad (2.30)$$

Typically, the LCAO basis sets consist of contracted Gaussian functions, because they allow for the most efficient evaluation of the occurring integrals. Their radial part has the general form

$$\chi_\mu(\mathbf{r}) \equiv G(\mathbf{r}; \mathbf{R}_\mu, \mathbf{l}) = \sum_k^C K_{\mu k} N_k (\mathbf{r} - \mathbf{R}_\mu)_x^{l_x} (\mathbf{r} - \mathbf{R}_\mu)_y^{l_y} (\mathbf{r} - \mathbf{R}_\mu)_z^{l_z} e^{-\zeta_k(\mathbf{r} - \mathbf{R}_\mu) \cdot (\mathbf{r} - \mathbf{R}_\mu)}, \quad (2.31)$$

where $K_{\mu k}$, N_k , \mathbf{R}_μ , \mathbf{l} and ζ_k respectively denote the contraction coefficients, normalization factors, atomic center, vector of angular momentum components, and exponent of the primitive Gaussians. However, plane-wave, wavelet and other bases also find application, the former especially for calculations of periodic systems.

2.3.3 Symmetric orthonormalization

The basis functions commonly used in quantum chemistry are normalized, but not orthogonal, such that $\mathbf{S} \neq \mathbf{1}$ and the Roothaan-Hall equation in its above form is not in the standard form of an eigenvalue problem. Still, for any basis set it is always possible to find a transformation matrix \mathbf{X} which orthonormalizes the basis, i.e. having the property

$$\mathbf{X}^\dagger \mathbf{S} \mathbf{X} = \mathbf{1}. \quad (2.32)$$

The simplest choice for this transformation matrix, where $\mathbf{X} = \mathbf{S}^{-\frac{1}{2}}$, is Löwdin's *symmetric* orthonormalization procedure [32]. Defining new matrices $\mathbf{C}' = \mathbf{X}^{-1} \mathbf{C} = \mathbf{S}^{\frac{1}{2}} \mathbf{C}$ and correspondingly $\mathbf{F}' = \mathbf{X}^\dagger \mathbf{F} \mathbf{X} = \mathbf{S}^{-\frac{1}{2}} \mathbf{F} \mathbf{S}^{-\frac{1}{2}}$, the Roothaan-Hall equation is converted to a standard matrix eigenvalue equation

$$\mathbf{F}' \mathbf{C}' = \mathbf{C}' \epsilon, \quad (2.33)$$

which can be solved by diagonalization of \mathbf{F}' , followed by reverse transformation to the original basis according to $\mathbf{C} = \mathbf{S}^{-\frac{1}{2}} \mathbf{C}'$.

This procedure, however, may become numerically unstable when the basis functions approach linear dependence with increasing basis size, or especially when diffuse basis functions are included. For then some eigenvalues of \mathbf{S} tend to zero, such that the computation of the complete

inverse square root becomes impossible. In that case, eigenvalues falling below a certain threshold are neglected, and the inverse square root is computed from the diagonal matrix $\mathbf{s}'^{-\frac{1}{2}}$ of reduced rank. Further developments in this field use Newton-Schultz iterations to compute the inverse overlap square root [33].

2.3.4 The Fock matrix and the Hartree-Fock energy

From the definition (2.22) for the closed-shell Fock operator, the Fock matrix is constituted of

$$\mathbf{F} = (\boldsymbol{\chi}|\hat{F}|\boldsymbol{\chi}) = (\boldsymbol{\chi}|\hat{H}|\boldsymbol{\chi}) + \sum_i^{N_{\text{occ}}} \left(2(\boldsymbol{\chi}|\hat{J}_i|\boldsymbol{\chi}) - (\boldsymbol{\chi}|\hat{K}_i|\boldsymbol{\chi}) \right) \quad (2.34)$$

$$\equiv \mathbf{H} + \mathbf{J} - \mathbf{K}, \quad (2.35)$$

with the core Hamiltonian matrix \mathbf{H} , the Coulomb matrix \mathbf{J} and the exchange matrix \mathbf{K} . Introducing the expansion (2.25) into equations (2.18) and (2.19) for the Coulomb and exchange operators, the Fock matrix elements are then given explicitly by

$$F_{\mu\nu} = H_{\mu\nu} + \sum_i^{N_{\text{occ}}} \sum_{\lambda,\sigma}^N (\mu\nu|(2 - \mathcal{P}_{(\nu\sigma)})\hat{C}|\lambda\sigma)C_{\lambda i}C_{\sigma i}. \quad (2.36)$$

For convenience, the ERI occurring in this expression are grouped into a matrix $\boldsymbol{\Gamma}$, the matrix representation of the antisymmetrized Coulombic repulsion operator in the AO-basis

$$\boldsymbol{\Gamma} = (\boldsymbol{\chi} \otimes \boldsymbol{\chi} | (2 - \mathcal{P})\hat{C} | \boldsymbol{\chi} \otimes \boldsymbol{\chi}) = (\boldsymbol{\Omega} | (2 - \mathcal{P})\hat{C} | \boldsymbol{\Omega}). \quad (2.37)$$

A *density* matrix whose properties are discussed in more detail in 2.4 is further defined according to

$$\mathbf{P}_o = \mathbf{C}_o \mathbf{C}_o^T, \quad (2.38)$$

where the matrix \mathbf{C}_o is the $N \times N_{\text{occ}}$ matrix of LCAO coefficients of the occupied MO. The Fock matrix can thus be expressed in terms of \mathbf{P}_o and AO-entities as

$$\begin{aligned} \text{vec}(\mathbf{F}) &= \text{vec}(\mathbf{H}) + \boldsymbol{\Gamma} \text{vec}(\mathbf{P}_o) \\ \mathbf{f} &= \mathbf{h} + \boldsymbol{\Gamma} \mathbf{p}, \end{aligned} \quad (2.39)$$

where the vectorization $\text{vec}(\mathbf{X}) = \mathbf{x}$ denotes the stacking of all the columns of the $N \times N$ argument matrix into a single N^2 -vector. Finally, upon comparison with the expectation value (2.13), the Hartree-Fock energy in terms of the Fock and density matrices is obtained as

$$\begin{aligned} E_{\text{HF}} &= \mathbf{p}^T \mathbf{h} + \mathbf{p}^T \mathbf{f} \\ &= 2\mathbf{p}^T \mathbf{h} + \mathbf{p}^T \boldsymbol{\Gamma} \mathbf{p} \end{aligned} \quad (2.40)$$

which, through reversion of the vectorization, is equivalent to

$$\begin{aligned} E_{\text{HF}} &= \text{Tr}(\mathbf{P}_o \mathbf{H} + \mathbf{P}_o \mathbf{F}) \\ &= \text{Tr}(2\mathbf{P}_o \mathbf{H} + (\mathbf{P}_o \otimes \mathbf{P}_o) \boldsymbol{\Gamma}). \end{aligned} \quad (2.41)$$

2.3.5 The Self-Consistent Field method

Due to the dependence of the Fock operator on its own eigenfunctions, or in the Roothaan-Hall formalism $\mathbf{F} = \mathbf{F}(\mathbf{C}_o) \equiv \mathbf{F}(\mathbf{P}_o)$, direct solution of the HF equations is precluded. Instead, an iterative procedure is required, where starting from an input density matrix the solution is gradually refined until convergence—*self-consistency* of the HF-potential—has been reached, whence the denomination as the Self-Consistent Field method. The SCF-algorithm is thus formulated recursively as:

- Build Fock matrix $\mathbf{F}^{(i+1)}(\mathbf{P}_o^{(i)})$.
- If converged: $E_{\text{HF}} = \text{Tr}(\mathbf{P}_o \mathbf{h} + \mathbf{P}_o \mathbf{F})$. Stop.
- Solve $\mathbf{F}^{(i+1)} \mathbf{C}^{(i+1)} = \mathbf{S} \mathbf{C}^{(i+1)} \boldsymbol{\epsilon}^{(i+1)}$
- Build $\mathbf{P}_o^{(i+1)} = \mathbf{C}_o^{(i+1)} \mathbf{C}_o^{(i+1)T}$.

The criterion of convergence might simply be chosen as the difference between two consecutive energies to be smaller than a predefined threshold value, i.e. $E_{\text{HF}} - E'_{\text{HF}} < \theta$, but from a theoretical point of view, it is preferable to require fulfilment of the commutation relation

$$\mathbf{F} \mathbf{P}_o \mathbf{S} - \mathbf{S} \mathbf{P}_o \mathbf{F} = \mathbf{0}, \quad (2.42)$$

which, resulting from the derivative of E_{HF} with respect to the density matrix [34], indicates that a stationary state has been reached. As to the initial value of the density matrix, different possibilities are given. One widely applied is the superposition of atomic densities (SAD), where the initial density is formed as the *direct sum*

$$\mathbf{P}_o = \bigoplus_A^{N_{\text{Atom}}} \mathbf{P}_A \quad (2.43)$$

of the single-atom density matrices, i.e. the block-diagonal matrix $\text{diag}(\mathbf{P}_1, \dots, \mathbf{P}_{N_{\text{Atom}}})$. In the above naive form, however, the SCF-scheme may converge only slowly or even diverge, such that methods have been developed to improve this. One of the most popular approaches is the *direct inversion in the iterative subspace* (DIIS) procedure of Pulay [35, 36], where the Fock matrices of a few consecutive iterations are combined according to

$$\bar{\mathbf{F}} = \sum_n w_n \mathbf{F}_n, \quad (2.44)$$

the weights being determined so as to minimize some particular error vector, and $\bar{\mathbf{F}}$ is then used as input for the next iteration.

2.4 The density operator and matrix

The description of a quantum mechanical system in terms of a wavefunction (state vector) $\psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$ is possible only for so-called *pure* states. The formalism of density operators and matrices, originating from statistical mechanics, as first introduced by von Neumann [37] in the context of quantum mechanics and further formalized by McWeeny [38], represents a generalization of the wavefunction picture to *mixed* states, i.e. not specified by a particular state vector. The presentation will here be limited to pure electronic states though.

Following Born's statistical interpretation, the probability for an arbitrary electron to be in a configuration space volume element at \mathbf{x}_1 —which, recalling section 2.3, denotes the combined spatial and spin coordinates of electron 1—given by

$$N_{\text{el}} \int \psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) \psi^*(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) d\mathbf{x}_2 \dots d\mathbf{x}_n = \rho'(\mathbf{x}_1) d\mathbf{x}_1, \quad (2.45)$$

defines the electron *spin-density function* $\rho'(\mathbf{x}_1)$. The factor N_{el} arising from the indistinguishability of electrons shall henceforth be omitted in order for the density function to be normalized, i.e. $\rho(\mathbf{x}) = \rho'(\mathbf{x})/N_{\text{el}}$. The coordinate \mathbf{x}_1 does not pertain to any particular electron, and the subscript in \mathbf{x}_1 is dropped.

Being an observable, the electron density may be associated with a *density operator* that is symbolically represented as $\hat{\rho} = |\psi\rangle\langle\psi|$, and can be interpreted as the projection operator onto the one-electron space spanned by $\psi(\mathbf{x})$. Namely, the application of $\hat{\rho}$ on any state $\xi(\mathbf{x})$ yields a multiple of $\psi(\mathbf{x})$, according to

$$\hat{\rho}|\xi\rangle = |\psi\rangle\langle\psi|\xi\rangle \equiv \psi(\mathbf{x}) \int \psi^*(\mathbf{x}')\xi(\mathbf{x}')d\mathbf{x}' = C\psi(\mathbf{x}). \quad (2.46)$$

In accordance with the characteristic idempotency of projection operators, we have

$$\hat{\rho}^2 = |\psi\rangle\langle\psi|\psi\rangle\langle\psi| = \hat{\rho}. \quad (2.47)$$

With this definition of $\hat{\rho}$, the expectation value of any one-electron operator $\hat{\mathcal{A}}$ can be calculated as

$$\langle\mathcal{A}\rangle = \text{Tr}(\hat{\mathcal{A}}\hat{\rho}) \equiv \int \left[\hat{\mathcal{A}}\psi(\mathbf{x}') \right]_{\mathbf{x}' \rightarrow \mathbf{x}} \psi^*(\mathbf{x}) d\mathbf{x}, \quad (2.48)$$

which is to be understood in the sense that the integration is carried out after the action of $\hat{\mathcal{A}}$ on the primed variable.

Now, in the context of HF theory, where the wavefunction is expanded in a basis of one-electron functions, the electron density is obviously given by the sum of squared moduli of the occupied spin-orbitals

$$\rho(\mathbf{x}) = \sum_i^{N_{\text{el}}} |\varphi_i(\mathbf{x})|^2. \quad (2.49)$$

Its corresponding density operator is thus given by the projection operator onto the subspace \mathcal{O}

$$\hat{\rho} = \sum_i^{N_{\text{el}}} |\varphi_i\rangle\langle\varphi_i| = |\varphi_o\rangle\langle\varphi_o|, \quad (2.50)$$

where in the last step the previously established notation of arranging the $|\varphi_p\rangle_{p \in N}$ into a vector has been used, such that $|\varphi_o\rangle$ denotes the vector containing the occupied spin-orbitals.

As for any operator, its representation with respect to a particular basis gives rise to a corresponding matrix, such that the *density matrix* in the MO basis is obtained as

$$\rho_{\text{MO}} = \langle\varphi|\hat{\rho}|\varphi\rangle = \langle\varphi|\varphi_o\rangle\langle\varphi_o|\varphi\rangle = \begin{pmatrix} \mathbf{1}_{N_{\text{el}}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad (2.51)$$

which is obviously Hermitean, idempotent, and has the property that $\text{Tr}\rho_{\text{MO}} = N_{\text{el}}$. For closed-shell systems, spin-integration leads to

$$\rho_{\text{MO}} = 2(\phi|\phi_o)(\phi_o|\phi) = \begin{pmatrix} \mathbf{2}_{N_{\text{occ}}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad (2.52)$$

recalling that $N_{\text{occ}} = N_{\text{el}}/2$. Upon comparison with (2.48), it follows that

$$\text{Tr}(\hat{\mathcal{A}}\hat{\rho}) = \text{Tr}(\hat{\mathcal{A}}|\phi_o)\mathbf{2}_{N_{\text{occ}}}(\phi_o|) = \text{Tr}((\phi_o|\hat{\mathcal{A}}|\phi_o)\mathbf{2}_{N_{\text{occ}}}) = \text{Tr}(\mathcal{A}_{\text{MO}}\rho_{\text{MO}}). \quad (2.53)$$

When the MOs are expanded according to the LCAO approach, insertion of the relation (2.25) in the definition (2.50) of the density operator leads to

$$\hat{\rho} = 2|\chi\rangle\mathbf{C}_o\mathbf{C}_o^T\langle\chi| = 2|\chi\rangle\mathbf{P}_o\langle\chi|, \quad (2.54)$$

such that the previous definition (2.38) of the matrix \mathbf{P}_o is recovered. The factor 2, however, is usually omitted by setting $\hat{\rho} \equiv \frac{1}{2}\hat{\rho}$, in order for the idempotency relation

$$\hat{\rho}^2 = |\chi\rangle\mathbf{P}_o\langle\chi|\chi\rangle\mathbf{P}_o\langle\chi| = |\chi\rangle\mathbf{C}_o\mathbf{C}_o^T\mathbf{S}\mathbf{C}_o\mathbf{C}_o^T\langle\chi| = \hat{\rho}, \quad (2.55)$$

to hold, from which it also follows that when projecting, the metric needs to be taken into account due to the non-orthogonality of AO basis functions. The matrix representation of $\hat{\rho}$ in the AO basis is further obtained as

$$\rho_{\text{AO}} = \langle\chi|\hat{\rho}|\chi\rangle = \langle\chi|\chi\rangle\mathbf{P}_o\langle\chi|\chi\rangle = \mathbf{S}\mathbf{P}_o\mathbf{S}. \quad (2.56)$$

However, \mathbf{P}_o uniquely defines the projector, and considering relation (2.53) for the expectation value of operators, it results that

$$2\text{Tr}(\hat{\mathcal{A}}\hat{\rho}) = 2\text{Tr}(\hat{\mathcal{A}}|\chi\rangle\mathbf{P}_o\langle\chi|) = 2\text{Tr}((\chi|\hat{\mathcal{A}}|\chi)\mathbf{P}_o) = 2\text{Tr}(\mathcal{A}_{\text{AO}}\mathbf{P}_o). \quad (2.57)$$

such that \mathbf{P}_o is defined as the AO occupied one-electron density matrix. The trace property and idempotency constraint of the AO-density matrix are thus expressed as

$$\text{Tr}(|\chi\rangle\mathbf{P}_o\langle\chi|) = \text{Tr}(\mathbf{P}_o\mathbf{S}) = N_{\text{el}}/2 \quad (2.58)$$

$$\mathbf{P}_o\mathbf{S}\mathbf{P}_o = \mathbf{P}_o. \quad (2.59)$$

It is further useful to define a *virtual* density matrix associated with the orthogonal complement $\hat{\rho}_\perp = \hat{1} - \hat{\rho}$ of the projector onto the occupied subspace, according to $\mathbf{P}_v = \mathbf{C}_v\mathbf{C}_v^T$, such that

$$|\chi\rangle\mathbf{P}_v\langle\chi| = 1 - |\chi\rangle\mathbf{P}_o\langle\chi|. \quad (2.60)$$

2.5 Linear scaling Hartree-Fock

The complexity, or *scaling*, of an algorithm is the computational effort, in terms of resources such as running time or storage, that is associated with its execution. It is usually measured by expressing the effort's growth rate as a function $f(n)$ of the input size n and in this context, "big O notation" $f(n) = \mathcal{O}(g(n))$ is used to denote that asymptotically ($n \rightarrow \infty$), the effort $f(n)$ grows at most as $g(n)$, i.e. there exists a certain input size n_0 and a constant K [39], for which

$$f(n) < K \cdot g(n), \quad \forall n > n_0. \quad (2.61)$$

Since quantum chemical calculations always involve a finite basis set whose size is directly proportional to the actual number of particles in the system under consideration, the scaling is commonly expressed with respect to the basis size N . Then, a polynomial scaling $\mathcal{O}(N^\alpha)$ means theoretically that once a certain N_0 has been exceeded, an increase of the system size by a factor k , i.e. $N' = kN_0$, entails an increase of computational effort by at most k^α . Ideally, the computational effort should asymptotically grow linearly with the system size, i.e. be $\mathcal{O}(N^1)$, such that doubling the system size only doubles the computation effort.

The actual scaling exponent α of a given resource R with respect to the basis size for two systems indicated by a, b is calculated as

$$\alpha = \frac{\ln(R_b/R_a)}{\ln(N_b/N_a)}, \quad (2.62)$$

In the following sections, the scaling of the HF-method is analyzed alongside a brief presentation of well-established schemes that have been devised to achieve linear scaling, hereby laying the foundation for the latter discussion of scaling reduction for correlation methods in chapters 3 and 5. The presentation begins with numerical screening methods indispensable to calculate only non-negligible integrals, rather than the entire ERI matrix. Then follow density matrix based HF-methods, where the costly diagonalization of the Fock matrix is avoided, and the chapter concludes with a brief discussion of how the Fock matrix is built with linearly scaling effort.

2.5.1 Integral screening and direct SCF

The HF method in its most straightforward *indirect* version formally scales $\mathcal{O}(N^4)$, due to the required computation and storage of the $N^2 \times N^2$ matrix of two-electron integrals. However, the products $\chi_\mu(\mathbf{r})\chi_\nu(\mathbf{r}) = \Omega_{\mu\nu}(\mathbf{r})$ of basis functions defined in equation (2.37), which may be viewed as charge distributions, decay exponentially with the separation $|\mathbf{R}_\mu - \mathbf{R}_\nu|$ between the centers of the individual Gaussians. Thus, for a given χ_μ , only a relatively small number of χ_ν in close proximity forms significant distributions, such that the number of non-negligible ERI is asymptotically quadratical. Now, due to the fact that the elements

$$(\mu\nu|\lambda\sigma) = \int \int \Omega_{\mu\nu}(\mathbf{r}_1) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \Omega_{\lambda\sigma}(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2 \quad (2.63)$$

of the matrix representation of the Coulombic repulsion operator fulfil all requirements of an inner product, the Cauchy-Schwarz-Bunyakovsky inequality [40, 41]

$$|(\mu\nu|\lambda\sigma)| \leq (\mu\nu|\mu\nu)^{\frac{1}{2}} (\lambda\sigma|\lambda\sigma)^{\frac{1}{2}} = Q_{\mu\nu} Q_{\lambda\sigma}, \quad (2.64)$$

provides an upper bound to their absolute values [15, 42, 43]. This cheap and efficient tool is nowadays universally used to preselect—this procedure is usually denoted as *Schwarz* screening, from the aforementioned inequality—significant ERI in integral-*direct* SCF implementations and has further been introduced into correlation methods, as is discussed in section 5.1.3.

2.5.2 Density-matrix based Hartree-Fock

Since the scaling of the ERI evaluation is easily made quadratic, the actually rate-determining step in the solution of the Roothaan-Hall equation is the diagonalization of the Fock matrix, which scales as $\mathcal{O}(N^3)$. Yet, equation (2.41) establishes the HF-energy as being uniquely specified by the one-electron density matrix \mathbf{P}_o , rather than the MO. This is quite logical, given that the orbitals have no interpretation in themselves, but only the electronic density that they represent, or more generally, that expectation values are quadratic in the wavefunction and thus invariant to unitary transformations thereof. In order to calculate the HF energy, it is thus not inevitable to solve the Roothaan-Hall equation (2.28) for the eigenvalues and eigenvectors (canonical MO) of the Fock matrix, but also possible to determine the density matrix that minimizes the energy expectation value. This is further motivated by the fact that the density matrix decays exponentially with inter-electronic distance for non-metallic, i.e. electronically localized systems [44–46]. The number of significant density matrix elements therefore is asymptotically linear, such that in contrast to the cubic matrix diagonalization, the solution for the density matrix may be obtained with linearly scaling effort.

The problem at hand is thus to minimize the energy functional $E_{HF}[\mathbf{P}_o]$ with respect to the density matrix under the constraint that the density matrix fulfil the idempotency and trace

conditions given in equations (2.58) and (2.59). This leads to the Lagrangian formulation

$$\frac{d\mathcal{L}[\mathbf{P}_o]}{d\mathbf{P}_o} = 0. \quad (2.65)$$

This matter shall not be further elaborated, but it should be mentioned that the imposition of the above conditions on the density matrix can be achieved through polynomial matrix functions, such as the purification transformation of McWeeny [38] or the trace resetting scheme of Niklasson and coworkers [47]. Other approaches are based on exponential parametrizations of the density matrix, such as the curvy steps method [48, 49].

2.5.3 Linear scaling Fock matrix construction

For overall linear scaling HF calculations, aside from circumventing the diagonalization, it is imperative to build the Fock matrix in an efficient way, in particular the Coulomb and exchange matrices \mathbf{J} and \mathbf{K} defined in section 2.3.4. It is discussed in section 2.5.1 how the number of ERI to be evaluated can be reduced by preselection of significant elements according to the Schwarz upper bound. This, however, provides quadratic scaling at best, and further modifications are required for linear scaling.

Linear scaling exchange

The linear scaling formation of the exchange matrix can, for instance, be achieved with the Link method [50], which combines Schwarz screening with the natural sparsity of the density matrix. Recalling equation (2.36), the exchange matrix is given by

$$K_{\mu\nu} = \sum_{\lambda\sigma} P_{\lambda\sigma}(\mu\sigma|\lambda\nu), \quad (2.66)$$

and it follows from equation (2.64) that the value of the individual summands is bound by

$$|P_{\lambda\sigma}(\mu\sigma|\lambda\nu)| \leq |P_{\lambda\sigma}|Q_{\mu\sigma}Q_{\lambda\nu}. \quad (2.67)$$

Thus, in addition to the exponential decay between basis functions, accounted for by the Schwarz estimates, the sparsity of the density matrix can be exploited to couple the indices between both charge distributions $\Omega_{\mu\sigma}$ and $\Omega_{\lambda\nu}$. Namely, for a given λ only a constant number of σ has non-vanishing density matrix elements, such that the number of significant summands asymptotically scales linearly. In association with a pre-ordering of the estimates with decreasing magnitude, such that loops can be terminated as soon as an element falls below a given threshold, this scheme allows for efficient construction of the exchange matrix.

Multipole Methods

In contrast to the exchange matrix formation, the use of the density matrix to couple between both charge distributions is not possible for the formation of the Coulomb matrix. Instead, one has to resort to multipole methods, whose strength lies in the fact that the interaction between well-separated groups of charges can be simplified to the interaction of the overall potentials of the charge distributions.

The *a priori* computational scaling of such methods is still quadratic, since for N charges there are $N(N+1)/2$ pairwise interactions, and even the grouping of individual charges into distributions merely reduces the prefactor. In order to reduce the scaling exponent, so-called fast multipole methods (FMM) [51–54] introduce a subdivision of the considered system into a hierarchy of boxes of different size—*coarseness*—containing varying numbers of particles. For each box, a separation criterion is then used to split the remaining boxes into the so-called near-field (NF) and far-field (FF), such that at each level of the boxing hierarchy, the NF size of any given box is constant. The NF interactions are then calculated explicitly, with constant effort, and the FF by multipole expansion. The advantage of this procedure is that with increasing distance, the interactions with FF boxes can be treated at an ever higher boxing level (coarser grain). Ultimately, the number of FF interactions for each box becomes constant, resulting in overall linear scaling of the total computational effort.

When applying these techniques to the Coulomb matrix construction, however, it has to be taken into account that the convergence of multipole series is guaranteed only for non-overlapping charge distributions, whereas the $\Omega_{\mu\nu}$ formed as products of Gaussians strictly extend over all space. The adaptation of FMM methods to continuous distributions thus requires the definition of extent criteria which can be used to ensure that the multipole approximation is applied only to well-separated $\Omega_{\mu\nu}$ and $\Omega_{\lambda\sigma}$ distributions, i.e. having numerically insignificant overlap. The error of this approximation can be shown to be of the order of numerical accuracy and thus poses no problem in practice. Various multipole-based hierarchical methods have been proposed to reach linear scaling of the Coulomb matrix evaluation, of which only the continuous fast multipole method (CFMM) of White *et al.* [55] shall be mentioned. Further information may be found in the review given in reference [56].

Chapter 3

Electron correlation

The Hartree-Fock method has long been established as an indispensable tool of quantum chemistry. However, even in the complete basis set limit, its very formulation prevents it from being able to recover the exact electronic energy of a many-electron system. This difference between the exact (non-relativistic) energy E_{exact} and the HF-limit energy has been defined by Löwdin as the *correlation* energy [57]

$$E_{\text{corr}} = E_{\text{exact}} - E_{\text{HF}}, \quad (3.1)$$

which is always negative, since the HF method is variational. The correlation energy is typically divided into two contributions of distinct origins. The so-called *dynamical* correlation arises from the mean-field approximation underlying HF-theory, due to its inability to accurately account for the mutual avoidance of electrons. Though the correlation between electrons is partially described for same spin electrons, because HF obeys the Pauli exclusion principle, the electronic interaction energy is typically overestimated due to the insufficient description of the Coulombic repulsion for electrons of opposite spin. On the other hand, *static* correlation effects become important in situations involving near-degenerate states, such as metals or bond-breaking configurations, which cannot be correctly described by a single ground state determinant but require multi-reference methods. Here again, methods such as CASSCF allow partial recovery of the static correlation contribution, but these are not in our focus.

Except for systems dominated by static correlation effects, HF typically recovers about 98-99% of the total electronic energy. Nevertheless, for an accurate analysis of chemical effects, it is crucial to account for the electron correlation. The hierarchy of methods developed to that end allows for arbitrary precision in the solution of the electronic Schrödinger equation, though of course at the cost of ever higher computational effort.

This chapter gives a brief presentation of some standard correlation methods along with a discussion of their respective scalings. The exposition begins with Configuration Interaction (CI) which, aside from its conceptual simplicity, in principle allows to recover the exact non-relativistic

electronic energy. The closely related coupled-cluster method (CC) is then described.

Then follows a concise presentation of the Random Phase Approximation (RPA). Though formally a post-Kohn-Sham method, since usually KS-orbitals are employed, RPA has received much attention in the context of *ab initio* quantum chemistry lately, due to a variety of attractive properties, and it might well achieve wider application in molecular quantum chemistry when further developments have reduced its computational cost.

We then turn our attention to Møller-Plesset perturbation theory (MP n), as it forms the basis of the present work. The essential steps of the derivation are recalled, after which the parametrization schemes that have been proposed to improve upon plain MP2 are presented. The presentation continues with the local correlation approach of Pulay and Saebø, where the computational cost is reduced by exploiting the intrinsically short-ranged character of correlation effects. In that context, an overview of established as well as recently proposed orbital localization procedures is given.

Finally, the *explicitly correlated* approaches pioneered by Hylleraas, which aim at overcoming the slow basis set convergence by explicitly including the inter-electronic separation into the correlation treatment, are briefly described.

3.1 Configuration Interaction and Coupled Cluster

In principle, though not in practice, Configuration Interaction (CI) is the simplest method to improve upon Hartree-Fock theory, it being the logical continuation of the one-determinantal approach. Namely, given a complete set $\{\varphi_p(\mathbf{x})\}$ of spin-orbitals, any electronic wavefunction can be expanded exactly in the basis of all possible Slater-determinants ("configurations") formed from this orbital set [27, 58]. Taking the—finite, and thus approximate—set $\{\varphi_p(\mathbf{x})\}_{p \in 2N}$ of HF spin-orbitals resulting from an SCF-calculation, the *full-CI* wavefunction (FCI) is then defined as

$$\Psi_{\text{FCI}} = C_0 \Phi_{\text{HF}} + \sum_i \sum_a^{N_{\text{el}} \ 2N - N_{\text{el}}} C_i^a \Phi_i^a + \sum_{i,j < i} \sum_{a,b < a}^{N_{\text{el}} \ 2N - N_{\text{el}}} C_{ij}^{ab} \Phi_{ij}^{ab} + \dots, \quad (3.2)$$

i.e. the linear-combination of the HF ground-state determinant and all n -tuply "excited" determinants $\Phi\{n\}$ formed by substituting n occupied orbitals with virtual ones. Applying the variational principle [28] to this wavefunction to determine optimal coefficients leads to the matrix eigenvalue problem

$$\mathbf{HC} = \mathbf{CE}, \quad (3.3)$$

where the elements $H_{kl}\{n, m\} = \langle \Phi_k\{n\} | \hat{\mathcal{H}} | \Phi_l\{m\} \rangle$ of the Hamiltonian matrix are given by the Slater-Condon rules. The expansion coefficients (amplitudes) \mathbf{C} and energy eigenvalues \mathbf{E} can then be determined through diagonalization of \mathbf{H} . As in general one might only be interested in specific energies, say for the ground state, there also exist special algorithms designed for the

extraction of selected eigenvalues, such as the Jacobi-Davidson method [59,60].

Yet, the practical applicability of full CI is severely limited by its enormous (combinatorial) computational cost. Namely, for a system of N_{el} electrons and $2N$ spin-orbitals, there are $\binom{N_{\text{el}}}{n} \binom{2N-N_{\text{el}}}{n}$ possible n -tuply excited configuration state functions [27], such that their number rapidly grows beyond all bounds and full CI calculations of molecules with more than but a few electrons become impossible. For reasonably sized molecules and basis sets, the excitation space must necessarily be restricted, for instance to at most double excitations $\Phi\{2\}$ (CISD), which constitute the most important contribution to the correlation energy, as it may be shown that [27]

$$E_{\text{corr}} = \sum_{i,j}^{N_{\text{el}}} \sum_{a,b}^{2N-N_{\text{el}}} C_{ij}^{ab} \langle \Phi_{\text{HF}} | \hat{\mathcal{H}} | \Phi_{ij}^{ab} \rangle. \quad (3.4)$$

although the doubles-coefficients are still coupled with higher excitations. This truncation, in turn, leads to the loss of *size consistency*, which is defined as the requirement that the energy of a system consisting of two non-interacting parts should be equal to the sum of their individual energies. This property is of particular importance for the computation of reaction energies or the correct description of dissociation curves.

This severe flaw of truncated CI is overcome in the conceptually related Coupled-Cluster (CC) method [61], where instead of the additive incorporation of excitations, an exponential representation of the wavefunction is chosen according to

$$\Psi_{\text{CC}} = e^{\hat{T}} \Psi_{\text{HF}}, \quad (3.5)$$

where $\hat{T} = \hat{T}_1 + \hat{T}_2 + \hat{T}_3 + \dots$ is the sum of single, double, n -tuple excitation operators. The advantage herein lies in the fact that even though the number of included excitations needs to be truncated, just as in CI, the exponential form of the operator ensures the retention of size consistency. In spite of being limited to relatively small systems due to its high computational scaling, coupled cluster—particularly as CCSD(T)—is now well-established as the "gold standard of quantum chemistry" for the computation of high-accuracy reference data.

3.2 The Random Phase Approximation

The Random Phase Approximation (RPA), originally introduced by Pines and Bohm in the context of plasma theory for the electron gas [62–64], has a long history of application within density functional theory. In recent years, however, it has also received much attention as a correlation method, due to its appealing properties of being applicable to small or zero gap systems, and accounting for long-range dispersion. Though Kohn-Sham-orbitals are typically used, the RPA correlation energy can nevertheless also be used to improve upon HF-calculations

according to

$$E = E_{\text{HF}} + E_{\text{corr}}^{\text{RPA}}. \quad (3.6)$$

The computation of RPA correlation energies formally requires the solution of the non-Hermitian eigenvalue problem [65]

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ -\mathbf{B} & -\mathbf{A} \end{pmatrix} \begin{pmatrix} \mathbf{X} & \mathbf{Y} \\ \mathbf{Y} & -\mathbf{X} \end{pmatrix} = \begin{pmatrix} \mathbf{X} & \mathbf{Y} \\ \mathbf{Y} & -\mathbf{X} \end{pmatrix} \begin{pmatrix} \boldsymbol{\Omega} & \mathbf{0} \\ \mathbf{0} & -\boldsymbol{\Omega} \end{pmatrix}. \quad (3.7)$$

Within direct RPA (dRPA), the most-widely applied variant of RPA, exchange contributions are neglected and the above matrices are given by

$$A_{ia,jb} = (\epsilon_a - \epsilon_i)\delta_{ij}\delta_{ab} + (ia|jb) \quad (3.8)$$

$$B_{ia,jb} = (ia|jb). \quad (3.9)$$

The symplectic eigenvalue problem (3.7) can be converted to the Hermitian form

$$\mathbf{M}\mathbf{Z} = \mathbf{Z}\boldsymbol{\Omega}^2, \quad (3.10)$$

with $\mathbf{M} = (\mathbf{A} - \mathbf{B})^{\frac{1}{2}}(\mathbf{A} + \mathbf{B})(\mathbf{A} - \mathbf{B})^{\frac{1}{2}}$. As shown by Furche [66], the dRPA correlation energy can then be succinctly expressed by the *plasmon* equation

$$E_{\text{corr}}^{\text{dRPA}} = \frac{1}{2}\text{Tr}(\boldsymbol{\Omega} - \mathbf{A}) = \frac{1}{2}\text{Tr}(\mathbf{M}^{\frac{1}{2}} - \mathbf{A}). \quad (3.11)$$

Through the use of the RI approximation for integral evaluation and frequency integration to avoid the explicit diagonalization of \mathbf{M} , the dRPA correlation energy can be calculated with $\mathcal{O}(N^4 \log N)$ effort, as shown by Eshuis *et al.* in [67], where an interesting relation to the SOS-MP2 scheme of Jung *et al.* [8] (see section 3.3.3) is also pointed out.

As a noteworthy aside, it has been proven analytically [68, 69] that the dRPA correlation energy (3.11) is equivalent to the direct ring coupled-cluster doubles (drCCD) correlation energy

$$E_{\text{corr}}^{\text{drCCD}} = \frac{1}{2}\text{Tr}\mathbf{B}\mathbf{T} = E_{\text{corr}}^{\text{dRPA}}, \quad (3.12)$$

where the amplitudes \mathbf{T} are determined from a CCD amplitudes equation

$$\mathbf{B} + \mathbf{A}\mathbf{T} + \mathbf{T}\mathbf{A} + \mathbf{T}\mathbf{B}\mathbf{T} = \mathbf{0}. \quad (3.13)$$

3.3 Many-Body Perturbation Theory

It has already been mentioned that usually the correlation energy is relatively small in comparison to the total energy, such that it can be described as a perturbation in the spirit of Rayleigh-Schrödinger perturbation theory (RSPT) [70]. The application of RSPT to the treatment of electron correlation in many-body systems is called Møller-Plesset perturbation theory

(MP n) [1] and is nowadays routinely used in quantum chemistry. Apart from being straightforward to implement on top of conventional HF code, MP n does further have the attractive traits of being size consistent, in contrast to truncated CI methods, and invariant to unitary orbital transformations (see also sections 3.4 and 4.1.1). It is, however, not variational and some investigations have shed considerable doubt on its convergence behaviour upon inclusion of higher-order terms [71, 72], such that in practice, it is usually restricted to low—typically second—order (MP2). Still, and in spite of difficulties regarding the correct description of strongly correlated systems such as diradicals or compounds containing metals, its good balance of computational cost and quality of results has made MP2 one of the universally applied correlation methods, and to this day, it is subject to much research effort, such as in the present work, that aims at reducing its computational cost in view of application to large molecules and basis sets.

3.3.1 Møller-Plesset Perturbation Theory

The starting point of RSPT is the partitioning

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_0 + \hat{\mathcal{H}}' \quad (3.14)$$

of the Hamiltonian into an unperturbed part $\hat{\mathcal{H}}_0$ and a perturbation $\hat{\mathcal{H}}'$. In the context of MP n theory [1], the unperturbed part is defined as the sum

$$\hat{H}_0 = \sum_i^{N_{\text{el}}} \hat{F}(i) \quad (3.15)$$

over one-electron Fock-operators, such that this naturally leads to the perturbation being given by the difference

$$\hat{\mathcal{H}}' = \sum_i^{N_{\text{el}}} \sum_{j < i} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_i^{N_{\text{el}}} \hat{V}^{\text{HF}}(i) \quad (3.16)$$

of the exact electronic interaction and the HF-potential, as mentioned in the definition of the correlation energy. With these definitions, it is readily established that the sum of zeroth and first order perturbed energies is equal to the HF-energy

$$E_0^{(0)} + E_0^{(1)} = \sum_i^{N_{\text{el}}} \epsilon_i - \frac{1}{2} \sum_{i,j}^{N_{\text{el}}} \langle ij || ij \rangle = E_{\text{HF}}, \quad (3.17)$$

such that the first correction to HF occurs in second order. Here and in the following, the order of perturbation shall be denoted by the superscripts in parentheses, whereas the 0 subscript indicates entities pertaining to the ground-state.

3.3.2 Second order Møller-Plesset theory (MP2)

From RSPT, the second order perturbed energy is given in general form by [27]

$$E_0^{(2)} = \sum_{n \neq 0} \frac{|\langle \Psi_0^{(0)} | \hat{\mathcal{H}}' | \Psi_n^{(0)} \rangle|^2}{E_0^{(0)} - E_n^{(0)}}, \quad (3.18)$$

where the sum is to be taken over all states except the ground state. Again, as in section 3.1, the wavefunctions $\Psi_0^{(n)}$ are given by the n -tuply "excited" determinants $\Phi\{n\}$ formed through substitution of n occupied orbitals by virtuals, whereas the zeroth-order wavefunction is the HF-determinant Φ_{HF} . Recalling that $\hat{\mathcal{H}}' = \hat{\mathcal{H}} - \hat{\mathcal{H}}_0$, it then follows that all terms involving single excitations vanish, since

$$\langle \Phi_{\text{HF}} | \hat{H} | \Phi_i^a \rangle = 0 \quad (3.19)$$

by Brillouin's theorem [73], and also

$$\langle \Phi_{\text{HF}} | \hat{H}_0 | \Phi_i^a \rangle = \langle \Phi_{\text{HF}} | \sum_k^{N_{\text{el}}} \hat{F}(k) | \Phi_i^a \rangle = F_{ia} = 0, \quad (3.20)$$

the Fock matrix being diagonal in the basis of canonical MO. Further, due to the two-electron nature of the Hamiltonian, all triples and higher substitutions vanish as well, such that the second order energy involves only double substitutions according to

$$\begin{aligned} E_0^{(2)} &= \frac{1}{4} \sum_{i,j}^{N_{\text{el}}} \sum_{a,b}^{2N-N_{\text{el}}} \frac{|\langle \Phi_{\text{HF}} | \sum_k^{N_{\text{el}}} \sum_{l < k} \frac{1}{|\mathbf{r}_k - \mathbf{r}_l|} | \Phi_{ij}^{ab} \rangle|^2}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \\ &= \frac{1}{4} \sum_{i,j}^{N_{\text{el}}} \sum_{a,b}^{2N-N_{\text{el}}} \frac{|\langle ij || ab \rangle|^2}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b}, \end{aligned}$$

where the factor 1/4 accounts for the symmetry of the expression with respect to i, j and a, b . After carrying out the spin integration, the second order energy for a closed-shell system is finally obtained as

$$E_0^{(2)} = \sum_{i,j}^{N_{\text{occ}}} \sum_{a,b}^{N_{\text{virt}}} \frac{(ia|jb) [2(ia|jb) - (ib|ja)]}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b}, \quad (3.21)$$

where indices a, b now denote the $N_{\text{virt}} = N - N_{\text{occ}}$ unoccupied spatial orbitals. The terms

$$E_{ij}^{(0)} = \sum_{a,b}^{N_{\text{virt}}} \frac{(ia|jb) [2(ia|jb) - (ib|ja)]}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (3.22)$$

represent the electron-pair contributions to the second-order correlation energy that also occur in coupled-pair theories [27]. In analogy to the Coulomb and exchange terms occurring in HF-theory, it is convenient to define a Coulomb-type part

$$E_{\text{J}}^{(2)} = 2 \sum_{i,j}^{N_{\text{occ}}} \sum_{a,b}^{N_{\text{virt}}} \frac{(ia|jb)^2}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (3.23)$$

and an exchange-type part,

$$E_{\text{K}}^{(2)} = \sum_{i,j}^{N_{\text{occ}}} \sum_{a,b}^{N_{\text{virt}}} \frac{(ia|jb)(ib|ja)}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b}. \quad (3.24)$$

such that

$$E_0^{(2)} = E_{\text{J}}^{(2)} - E_{\text{K}}^{(2)}. \quad (3.25)$$

3.3.3 Parametrized MP2

In a pragmatic approach to improve upon plain MP2, Grimme introduced the spin-component scaled MP2 (SCS-MP2) method [74], based on a separate scaling of the same spin and opposite spin components of the MP2 energy, according to

$$E_{\text{SCS}}^{(2)} = C_{\text{OS}} E_{\text{OS}}^{(2)} + C_{\text{SS}} E_{\text{SS}}^{(2)}. \quad (3.26)$$

The coefficients are then calibrated in order to accurately reproduce energies for a set of high-end benchmark calculations. With the definitions (3.23) and (3.24) of Coulomb and exchange parts of MP2, the opposite spin and same spin components are respectively given as

$$E_{\text{OS}}^{(2)} = \frac{1}{2} E_{\text{J}}^{(2)} \quad E_{\text{SS}}^{(2)} = \frac{1}{2} E_{\text{J}}^{(2)} - E_{\text{K}}^{(2)}. \quad (3.27)$$

Following Grimme’s pioneering work, a series of similar schemes has been proposed, such as the scaled opposite spin (SOS) MP2 of Jung *et al.* [8], where only the opposite spin component is taken into account according to

$$E_{\text{SOS}}^{(2)} = C_{\text{SOS}} E_{\text{OS}}^{(2)}, \quad (3.28)$$

or the S2-PT approach of Fink [75], which represents a generalization of SCS-MP2.

However, though on average these approaches represent a certain improvement, they forsake the *ab initio* character of MP2 theory and some of these schemes have the severe flaw of not recovering various effects to the same extent as unparametrized MP2, e.g. long-range dispersion interactions. In effect, for reasonable reliability, this means that the parameters either have to be specifically tailored for the case of interest, or that yet further parameters need to be introduced, for instance in a modification of the SOS-MP2 approach that employs an attenuated Coulomb metric [76, 77], and where the attenuation parameter needs to be optimized.

3.4 Local correlation methods

The computational cost of MP2 calculations is determined by the quadruple transformation

$$(ia|jb) = \sum_{\mu}^N C_{\mu i} \left(\sum_{\nu}^N C_{\nu a} \left(\sum_{\lambda}^N C_{\lambda j} \left(\sum_{\sigma}^N C_{\sigma b}(\mu\nu|\lambda\sigma) \right) \right) \right) \quad (3.29)$$

of ERI over atomic orbitals to the MO-basis. Due to the fact that the canonical MO generally spread over the whole system, i.e. the coefficient matrix \mathbf{C} is dense, this results in $\mathcal{O}(N^5)$ scaling. Though this transformation can be simplified and the overall prefactor be decreased, for instance by applying 'resolution of the identity' (RI) techniques (see chapter 6), conventional MP2 calculations rapidly become prohibitively expensive.

However, electronic correlation effects are rather short-ranged, displaying the same long-distance decay of r_{ij}^{-6} as dispersion effects, such that in non-delocalized systems, for each electron only a relatively small (asymptotically constant) number of spatially close electrons lead to significant pair interaction energies (3.22). It is thus intuitive that the correlation energy could be calculated in a linear-scaling manner by exploiting this behaviour through the use of a basis of localized MO (LMO). This section gives a short overview of orbital localization methods, after which the essential traits of Pulay's local MP2 formulation [78, 79] are outlined, it being the counterpart to the Laplace-transform approach described in the next chapter.

3.4.1 Orbital localization methods

The disadvantages incurred from the use of the delocalized canonical MO for the treatment of electron correlation have been recognized early on and some localization schemes were proposed relatively long ago [80, 81]. Yet, this topic still constitutes an active field of research, such that today a broad array of different methods is available.

Technically, MO locality is expressed by a diagonal-dominant structure of the coefficient matrix, which means that for each orbital, only a small subset of spatially close basis functions has non-negligible expansion coefficients (cf. equation (2.25)). The procedure for localizing the orbitals, however, is not uniquely defined, but is usually formulated as the minimization of some particular functional $\zeta(i)$ of the orbitals, and requires an iterative solution. The common objective of all localization procedures is thus the computation of a matrix \mathbf{L}_o of localized coefficients, connected to the canonical MO by a unitary transformation $\mathbf{L}_o = \mathbf{C}_o \mathbf{U}_o$, in order to preserve the orthonormality. In this section, both established and recently proposed schemes are briefly described, following in part the review given by Pipek and Mezey [82]. Another short recapitulation of localization methods is also found in [83].

The method of Boys [80, 84] aims at obtaining LMO of minimal spatial extent by maximizing the distance of orbital centroids. This can be shown to be equivalent to maximization of the functional

$$\zeta_{\text{Boys}}(i) = \sum_i^{N_{\text{occ}}} (i|\mathbf{r}|i)^2, \quad (3.30)$$

where \mathbf{r} is the position operator. Though this approach bears the slight disadvantage of producing artificial τ -orbitals as linear combinations of σ - and π -orbitals, which however is of no concern for the purpose of methodical scaling reduction, its relatively low $\mathcal{O}(N^3)$ scaling is an

attractive feature. Through the use of pre-localized orbitals, as discussed below, the scaling of this procedure can be further reduced to linear.

In the approach chosen by Edminston and Ruedenberg (ER) [81, 85], the orbital extents are minimized by maximizing the Coulombic self-interaction energy

$$\zeta_{\text{ER}}(i) = \sum_i^{N_{\text{occ}}} (ii|ii), \quad (3.31)$$

leading to minimal exchange energy between them. Though this method preserves the σ - and π -separation, in contrast to the Boys method, the required integral transformation causes the computational cost to scale as $\mathcal{O}(N^5)$, which is unfortunate since the primary incentive for using LMO is to speed up correlation methods. Another scheme proposed by von Niessen [86] is conceptually related to the Edminston-Ruedenberg method.

The method of Pipek and Mezey [82, 87] advocates the minimization of the sum of quadratic gross atomic Mulliken populations to minimize the number of atoms over which the MO spread, i.e. the functional

$$\zeta_{\text{PM}}(i) = \sum_i^{N_{\text{occ}}} \sum_A^{N_{\text{atom}}} \left| (i|\hat{P}_A|i) \right|^2, \quad (3.32)$$

where $\hat{P}_A = \sum_{\mu \in A} |\mu\rangle\langle\mu|$ is the projection operator onto the space spanned by the AOs centered on atom A . Because it combines preservation of the σ - π -separation with a favourable $\mathcal{O}(N^3)$ scaling, the PM method is nowadays one of the most widely applied localization methods.

On a somewhat different footing, the scheme of Aquilante *et al.* [88] exploits the natural sparsity of the one-electron density matrix \mathbf{P}_o to generate LMO by performing a Cholesky decomposition [9] (see also section 5.1.1)

$$\mathbf{P}_o = \mathbf{L}_o \mathbf{L}_o^T. \quad (3.33)$$

In contrast to the previously presented methods of Boys, ER and PM, this method has the significant advantage of being non-iterative. Further, it does not require an initial MO set, but merely a density matrix, such that it is applicable in conjunction with density-based HF algorithms. Though it formally scales as $\mathcal{O}(N^3)$, a linear-scaling implementation of the Cholesky decomposition for positive semi-definite matrices has recently been proposed in conjunction with further use of the obtained Cholesky MO as input for the Boys procedure [89]. It is also straightforwardly applied to obtain local virtual MO by decomposing the virtual density matrix

$$\mathbf{P}_v = (\mathbf{S}^{-1} - \mathbf{P}_o) = \mathbf{L}_v \mathbf{L}_v^T \quad (3.34)$$

defined in equation (2.60), whereas standard localization functions, though in principle applicable to the virtual MO as well, have been discussed to perform rather poorly for that purpose [90]. Being sparse and normally positive definite, the overlap matrix can be inverted in a linear-scaling fashion by inverse Cholesky decomposition [91].

Another localization procedure recently proposed by Jansík *et al.* [83] is based on minimization of powers of the orbital variance Ω_p

$$\zeta_n(p) = \sum_p^N \Omega_p^n = \sum_p^N ((p|\mathbf{r}^2|p) - (p|\mathbf{r}|p)^2)^n, \quad (3.35)$$

where the general MO index p is deliberately used to indicate that this method has been shown to perform well both for occupied and virtual orbitals. For $n = 1$, this approach is equivalent to the Boys method, though a different optimization algorithm is employed, and locality is improved upon increasing n .

Instead of using virtual LMO, yet another way to span the virtual space \mathcal{V} is to use a set of projected atomic orbitals (PAO) obtained by projecting out \mathcal{O} , in order to maintain the orthogonality of both subspaces. Recalling section 2.4, the orthogonal complement of the density operator is given by

$$\hat{\rho}_\perp = 1 - |\chi\rangle\mathbf{P}_o\langle\chi|, \quad (3.36)$$

so that the PAO, which shall be denoted with tildes, are obtained as

$$|\tilde{\chi}\rangle = \hat{\rho}_\perp|\chi\rangle = (1 - |\chi\rangle\mathbf{P}_o\langle\chi|)|\chi\rangle = |\chi\rangle - |\chi\rangle\mathbf{P}_o\mathbf{S} = |\chi\rangle(1 - \mathbf{P}_o\mathbf{S}). \quad (3.37)$$

Though orthogonal to the occupied MO, the PAO are not orthonormal among themselves.

3.4.2 Local MP2

In section 3.3.2, the traditional derivation of MP2 was followed. Through minimization of the Hylleraas functional, it can also be obtained in the more general form [92]

$$E_0^{(2)} = \sum_{k,l \in \mathcal{O}} \sum_{r,s \in \mathcal{V}} \mathcal{T}_{kl}^{rs}(kr|ls) = \sum_{k,l \in \mathcal{O}} \sum_{r,s \in \mathcal{V}} \mathcal{T}_{kl}^{rs} M_{kl}^{rs} = \text{Tr} \mathcal{T} \mathbf{M}, \quad (3.38)$$

where k, l and r, s indicate arbitrary basis functions for the occupied and virtual subspaces, respectively, and the antisymmetrized amplitudes \mathcal{T} are given by $\mathcal{T}_{kl}^{rs} = 2T_{kl}^{rs} - T_{kl}^{sr}$. For the case of the canonical MO basis, a comparison with equation (3.21) directly yields that

$$\mathcal{T}_{ij}^{ab} = \frac{2(ia|jb) - (ib|ja)}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} = \frac{(ia||jb)}{f_{ii} + f_{jj} - f_{aa} - f_{bb}}. \quad (3.39)$$

Their calculation, though computationally expensive, is thus a straightforward matter, as the MO Fock matrix \mathbf{f} is diagonal. In a basis of LMO and PAO, however, as used in the local MP2 (LMP2) method proposed by Pulay and Saebø [78, 79], the LMO and PAO Fock matrices \mathbf{f}' and $\tilde{\mathbf{F}}$ are non-diagonal and the amplitudes need to be determined iteratively for each electron pair from the set of equations [92–94]

$$\mathbf{M}_{ij} + \tilde{\mathbf{F}}\mathbf{T}_{ij}\tilde{\mathbf{S}} + \tilde{\mathbf{S}}\mathbf{T}_{ij}\tilde{\mathbf{F}} - \sum_k^{N_{\text{occ}}} \tilde{\mathbf{S}} [f'_{ik}\mathbf{T}_{kj} + \mathbf{T}_{ik}f'_{kj}] \tilde{\mathbf{S}} = \mathbf{0}, \quad (3.40)$$

where $\tilde{\mathbf{S}}$ denotes the PAO overlap matrix and \mathbf{T}_{ij} is the amplitude sub-matrix of the ij -th electron pair. The advantage herein is that on the basis of the locality of the LMO and PAO, approximations can now be introduced to exploit the short-range character of correlation in order to reduce the computational scaling.

First, the space into which electrons are substituted (excited) is truncated by assigning to each LMO $|i\rangle$ an excitation *domain*, i.e. a subset $[i]$ of spatially close PAO to which substitution is restricted. These domains are usually constructed according to the method of Boughton and Pulay [95], where for each LMO, 3-4 atoms are determined by a least-squares residual minimization such that their combined PAO span the considered LMO to a defined extent, typically 98%. For an electron pair described by $|i\rangle$ and $|j\rangle$ the pair excitation domain $[ij]$ is then given by the union $[i] \cup [j]$ of individual orbital domains.

Further, a hierarchy of (ij) pairs is defined according to the minimum distance R in Bohr (a.u.) between any atoms in their respective domains. The contraction to the MP2 energy thus becomes

$$E_0^{(2)} = \sum_{i,j \in (ij)} \sum_{\tilde{\mu}, \tilde{\nu} \in [ij]} T_{ij}^{\tilde{\mu}\tilde{\nu}}(i\tilde{\mu}|j\tilde{\nu}). \quad (3.41)$$

In the LMP2 implementations of Werner *et al.* [93, 96], the so-called *strong* ($R \in (0; 1]$) and *weak* ($R \in (1; 8]$) pairs are treated explicitly. Strong pairs may even be treated at a higher theoretical level, for instance CCSD, such that a larger portion of the correlation energy is recovered. *Distant* pairs ($R \in (8; 15]$) may either be treated in the same manner as weak pairs, but approximating the ERI by multipole methods as mentioned in section 2.5.3, or neglected, whereas *very distant* pairs ($R \in (15; \infty)$) are ignored altogether.

This method has been successfully applied in a variety of linear scaling LMP2 implementations [93, 94, 96], and has further found application in coupled-cluster theory. However, due to the *a priori* neglect of electronic interactions, these local approximations have been shown to produce discontinuous potential curves [97]. This is problematic in view of molecular geometry optimizations and calculation of other properties, but progress on the solution of this issue has been reported [98].

3.5 Explicitly correlated methods

One of the major problems of MP2, and other correlation methods for that matter, is the slow convergence of the correlation energy with basis set size. For instance, within Dunning’s series of correlation consistent basis sets cc-pVnZ ($n = 2, 3, \dots$) [99–102], the deviation from the basis set limit correlation energies decreases only as $\mathcal{O}(n^{-3})$, since the basis set size N_n grows as $\mathcal{O}(n^3)$. However, the corresponding rise of $\mathcal{O}(n^{12})$ in the number of integrals and thus computational cost [103] is in no relation to the obtained improvement. This behaviour is due to the

inability of orbital product expansions of the wavefunction (Slater determinants), moreover in the typically used Gaussian basis sets, to accurately model the cusp of the wavefunction [104]. This (finite) discontinuity in the wavefunction arises from the fact that the Coulomb interaction operator $1/r_{12}$ diverges in the limit $r_{12} \rightarrow 0$, such that the kinetic energy must compensate for this divergence in order to maintain the local energy constant.

In order to overcome this disadvantageous behaviour, a variety of methods, going back to Hylleraas' pioneering work on the helium atom [105, 106], have been developed which explicitly introduce the interelectronic distance into the correlation treatment by including a *correlation factor* $F_{12}(r_{12})$ of the electronic separation. In the so-called *R12* methods, this is simply $F_{12} = r_{12}$, but it has been shown that Slater-type factors $F_{12} = e^{-\gamma r_{12}}$ used in *F12* approaches perform better [107, 108]. Other schemes have also been proposed, but we refer to [103, 109] for further detail.

One disadvantage of these explicitly correlated methods is the occurrence of three- and four-electron integrals whose evaluation is computationally very demanding, and significant progress in view of applying explicitly correlated methods to systems of moderate size has only been made upon Kutzelnigg and Klopper's introduction of special auxiliary bases to approximate these integrals using RI techniques [110, 111].

Chapter 4

Laplace-Transform Møller-Plesset Perturbation Theory

We now turn to the description of Almlöf and Häser’s Laplace-transform formalism to remove energy denominators [2, 3], upon which the work presented in the next chapter is based. This chapter is dedicated to a recapitulation of interesting theoretical implications that the Laplace-transform formulation has for Møller-Plesset theory. The intuitive fact that the MP n energy should be independent from the choice of orbital basis is proven, from which it logically follows that Laplace-transform MP corresponds to a (non-iterative) version of local MP schemes introduced by Pulay. As an aside, an interesting relation between Laplace-transform MP2 and Lyapunov equations is briefly pointed out.

We then turn to atomic-orbital MP theory (AO-MP), as first introduced by Häser for MP2 [4], which allows to express the MP energy as a functional of the one-electron density matrix only, thus providing the natural extension of density-based HF-schemes to correlation methods. Due to the inherent locality of the AO-basis, Häser’s formulation has laid the foundation for linear-scaling implementations based on the preselection of integrals according to schemes previously described in section 2.5.1, in order to extend perturbative correlation treatment to large molecules intractable with the canonical formulation of MP theory.

4.1 Laplace-Transform MP2

We start by recalling equation 3.21 for the MP2 energy of a closed-shell system:

$$E_0^{(2)} = - \sum_{i,j}^{N_{\text{occ}}} \sum_{a,b}^{N_{\text{virt}}} \frac{(ia|jb) [2(ia|jb) - (ib|ja)]}{\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j}, \quad (4.1)$$

where indices i, j and a, b denote occupied and virtual spatial molecular orbitals (MO), respectively. Following the notation established in the first chapter, all entities pertaining to the

AO-basis are indicated with greek letters.

As discussed in section 3.4, the fact that the canonical MOs are delocalized (i.e. the coefficient matrix \mathbf{C} is dense), leads to the transformation

$$(ia|jb) = \sum_{\mu\nu\lambda\sigma}^N C_{\mu i} C_{\nu a} (\mu\nu|\lambda\sigma) C_{\lambda j} C_{\sigma b}, \quad (4.2)$$

from the AO to the MO basis to scale as $\mathcal{O}(N^5)$. Since the MP2 energy is independent of the particular choice of the MO basis (see section 4.1.1), it is desirable to express it in a localized basis in order to exploit the short-range character of electron correlation effects. The non-separability of the canonical orbital energies in the denominator of the MP2 expression, however, precludes a straightforward change of basis, such that the amplitudes of a local MP2 calculation need to be determined iteratively, as discussed in section 3.4.2.

In another approach to overcome this problem, Almlöf and Häser proposed to remove the denominator $\Delta_{iajb} = \epsilon_a - \epsilon_i + \epsilon_b - \epsilon_j$ by expressing it through the integral [2, 3]

$$\frac{1}{\Delta_{iajb}} = \int_0^\infty e^{-\Delta_{iajb}t} dt, \quad (4.3)$$

which, according to the definition of the Laplace transform [112],

$$F(s) = \mathcal{L}(f(t)) \equiv \int_0^\infty f(t) e^{-st} dt, \quad \text{Re}(s) > 0, \quad (4.4)$$

corresponds to a transform of the constant function $f(t) = 1$. The integral (4.3) can be approximated by a finite sum of exponentials

$$\int_0^\infty e^{-\Delta_{iajb}t} \approx \sum_{\alpha}^{\tau} w_{\alpha} e^{-\Delta_{iajb}t_{\alpha}}, \quad (4.5)$$

where $5 \leq \tau \leq 8$ quadrature points have been shown to be sufficient for μ Hartree accuracy [3, 6]. The orbital energies in the exponent thus become separable, giving

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{i,j}^{N_{\text{occ}}} \sum_{a,b}^{N_{\text{virt}}} e^{-(\epsilon_a - \epsilon_i)t_{\alpha}} (ia|jb) [2(ia|jb) - (ib|ja)] e^{-(\epsilon_b - \epsilon_j)t_{\alpha}}. \quad (4.6)$$

For notational simplicity, the ERI and orbital energy differences are arranged into matrices

$$\begin{aligned} \mathbf{A} &\equiv A_{\mu\nu,\lambda\sigma} &= (\mu\nu|\lambda\sigma) \\ \mathbf{\Gamma} &\equiv \Gamma_{\mu\nu,\lambda\sigma} &= 2A_{\mu\nu,\lambda\sigma} - A_{\mu\sigma,\lambda\nu} \\ \mathbf{M} &\equiv M_{ia,jb} &= (ia|jb) \\ \mathbf{G} &\equiv G_{ia,jb} &= 2M_{ia,jb} - M_{ib,ja} \\ \mathbf{D} &\equiv D_{ia,kc} &= (\epsilon_a - \epsilon_i)\delta_{ik}\delta_{ac} \end{aligned} \quad (4.7)$$

with dimension $N^2 \times N^2$ or $N_{\text{occ}}N_{\text{virt}} \times N_{\text{occ}}N_{\text{virt}}$, such that equation (4.6) reads in matrix form

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} (e^{-\mathbf{D}t_{\alpha}} \mathbf{M} e^{-\mathbf{D}t_{\alpha}} \mathbf{G}). \quad (4.8)$$

The matrix exponential is defined by the extension of the corresponding Taylor series from scalars to matrices, i.e.

$$e^{\mathbf{X}} = \sum_{n=0}^{\infty} \frac{\mathbf{X}^n}{n!} = \mathbf{1} + \mathbf{X} + \frac{1}{2}\mathbf{X}^2 + \frac{1}{6}\mathbf{X}^3 + \dots \quad (4.9)$$

and always converges. Note that under similarity transformations

$$\mathbf{T}^{-1} e^{\mathbf{X}} \mathbf{T} = e^{\mathbf{T}^{-1} \mathbf{X} \mathbf{T}} \quad (4.10)$$

holds, i.e. the similarity-transformed matrix exponential is equal to the exponential of the similarity-transformed exponent.

Inserting the LCAO expansion (2.25) into the definition (2.37) of the ERI matrix, the transformation (4.2) reads

$$\mathbf{M} = (\mathbf{C}_o \otimes \mathbf{C}_v)^T \mathbf{A} (\mathbf{C}_o \otimes \mathbf{C}_v), \quad (4.11)$$

recalling the elementary properties of the Kronecker product

$$(\mathbf{X} \otimes \mathbf{Y})^T = (\mathbf{X}^T \otimes \mathbf{Y}^T) \quad (4.12)$$

$$(\mathbf{X} \otimes \mathbf{Y})(\mathbf{V} \otimes \mathbf{W}) = (\mathbf{XV} \otimes \mathbf{YW}). \quad (4.13)$$

The subscripts o and v denote matrices belonging to the occupied and virtual subspaces, respectively. We proceed by pointing out that the above expression for \mathbf{D} corresponds to the definition of the Kronecker sum (not to be confused with the *direct* sum)

$$\mathbf{D} = -\epsilon_o \oplus \epsilon_v = -\epsilon_o \otimes \mathbf{1}_v + \mathbf{1}_o \otimes \epsilon_v, \quad (4.14)$$

for which the following relation holds with respect to the matrix exponential

$$e^{\mathbf{X} \oplus \mathbf{Y}} = e^{\mathbf{X}} \otimes e^{\mathbf{Y}}. \quad (4.15)$$

Finally, by recalling that the orbital energies are recovered by transforming the Fock matrix to the basis of canonical MOs, we have

$$\epsilon_o = \mathbf{C}_o^T \mathbf{F} \mathbf{C}_o \quad \epsilon_v = \mathbf{C}_v^T \mathbf{F} \mathbf{C}_v. \quad (4.16)$$

4.1.1 Orbital invariance and relation to local MP2

The invariance of the MP2 energy under orbital rotations has already been pointed out by Almlöf upon his introduction of the Laplace transform [2]. The proof is now given in the most general way. Inserting relation (4.15) into the energy (4.8), noting that the exponential matrices of orbital energies are symmetric, leads to

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((e^{t_{\alpha} \epsilon_{\text{o}}} \otimes e^{-t_{\alpha} \epsilon_{\text{v}}}) \mathbf{M} (e^{t_{\alpha} \epsilon_{\text{o}}} \otimes e^{-t_{\alpha} \epsilon_{\text{v}}}) \mathbf{G} \right). \quad (4.17)$$

By taking the square root of the matrix exponentials according to

$$(e^{t_{\alpha} \epsilon_{\text{o}}} \otimes e^{-t_{\alpha} \epsilon_{\text{v}}}) = (e^{\frac{t_{\alpha}}{2} \epsilon_{\text{o}}} \otimes e^{-\frac{t_{\alpha}}{2} \epsilon_{\text{v}}}) (e^{\frac{t_{\alpha}}{2} \epsilon_{\text{o}}} \otimes e^{-\frac{t_{\alpha}}{2} \epsilon_{\text{v}}}) = (\mathbf{e}_{\text{o}}^{(\alpha)} \otimes \mathbf{e}_{\text{v}}^{(\alpha)}) (\mathbf{e}_{\text{o}}^{(\alpha)} \otimes \mathbf{e}_{\text{v}}^{(\alpha)}), \quad (4.18)$$

the cyclic invariance of the trace can be used to rearrange the previous expression to the symmetric form

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\mathbf{e}_{\text{o}}^{(\alpha)} \otimes \mathbf{e}_{\text{v}}^{(\alpha)}) \mathbf{M} (\mathbf{e}_{\text{o}}^{(\alpha)} \otimes \mathbf{e}_{\text{v}}^{(\alpha)}) (\mathbf{e}_{\text{o}}^{(\alpha)} \otimes \mathbf{e}_{\text{v}}^{(\alpha)}) \mathbf{G} (\mathbf{e}_{\text{o}}^{(\alpha)} \otimes \mathbf{e}_{\text{v}}^{(\alpha)}) \right). \quad (4.19)$$

Now, considering unitary rotations $\mathbf{L}_{\text{o}} = \mathbf{C}_{\text{o}} \mathbf{U}_{\text{o}}$ and $\mathbf{L}_{\text{v}} = \mathbf{C}_{\text{v}} \mathbf{U}_{\text{v}}$ of the occupied and virtual orbitals, and defining a matrix $\mathbf{U} = (\mathbf{U}_{\text{o}} \otimes \mathbf{U}_{\text{v}})$, the identity relation $\mathbf{U} \mathbf{U}^T = \mathbf{1}$ may simply be inserted at the appropriate places, such that

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\mathbf{e}_{\text{o}}^{(\alpha)'} \otimes \mathbf{e}_{\text{v}}^{(\alpha)'}) \mathbf{M}' (\mathbf{e}_{\text{o}}^{(\alpha)'} \otimes \mathbf{e}_{\text{v}}^{(\alpha)'}) (\mathbf{e}_{\text{o}}^{(\alpha)'} \otimes \mathbf{e}_{\text{v}}^{(\alpha)'}) \mathbf{G}' (\mathbf{e}_{\text{o}}^{(\alpha)'} \otimes \mathbf{e}_{\text{v}}^{(\alpha)'}) \right). \quad (4.20)$$

The primes here denote matrices in a general basis of non-canonical, thus potentially local orbitals, i.e. from equation (4.11)

$$\mathbf{M}' = (\mathbf{U}_{\text{o}}^T \otimes \mathbf{U}_{\text{v}}^T) \mathbf{M} (\mathbf{U}_{\text{o}} \otimes \mathbf{U}_{\text{v}}) = (\mathbf{L}_{\text{o}}^T \otimes \mathbf{L}_{\text{v}}^T) \mathbf{A} (\mathbf{L}_{\text{o}} \otimes \mathbf{L}_{\text{v}}) \quad (4.21)$$

and for the exponential energy matrices, invoking relations (4.10) and (4.16)

$$\mathbf{e}_{\text{o}}^{(\alpha)'} = \mathbf{U}_{\text{o}}^T \mathbf{e}_{\text{o}}^{(\alpha)} \mathbf{U}_{\text{o}} = \mathbf{U}_{\text{o}}^T e^{\frac{t_{\alpha}}{2} \epsilon_{\text{o}}} \mathbf{U}_{\text{o}} = e^{\frac{t_{\alpha}}{2} \mathbf{L}_{\text{o}}^T \mathbf{F} \mathbf{L}_{\text{o}}}, \quad (4.22)$$

with corresponding relations for the matrix \mathbf{G} of antisymmetrized ERI and the virtual exponential energy matrix $\mathbf{e}_{\text{v}}^{(\alpha)}$. The equality of the previous energy expression to the canonical MP2 energy (4.19) is obvious.

The matrix \mathbf{G}' in (4.20) can further be directly related to the amplitude matrix in equation (3.39). For localized MO, the Laplace transform approach can thus be interpreted as a variant of the local MP2 formulation of Pulay, where instead of the iterative solution for the amplitudes, the ERI matrices have to be transformed with the exponential weighting matrices $\mathbf{e}_{\text{o}}^{(\alpha)}$ and $\mathbf{e}_{\text{v}}^{(\alpha)}$ for each quadrature point. This relationship between local and Laplace transform MP2 has first been pointed by Schütz [113], though for the particular case of spanning the virtual subspace with PAO.

4.1.2 Relation to Lyapunov equations

In this section, we wish to briefly point out an elegant relation between the Laplace MP2 energy and Lyapunov equations, which to the best of our knowledge has not been mentioned so far. For given matrices \mathbf{A} and Hermitean \mathbf{C} , the equation

$$\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{A}^T = \mathbf{C}, \quad (4.23)$$

which occurs in certain areas of control theory, is called the continuous time Lyapunov equation. The equation has a unique Hermitean solution \mathbf{X} if and only if \mathbf{A} and $-\mathbf{A}^T$ have no eigenvalues in common, and if \mathbf{C} is positive (semi)definite the solution is also positive (semi)definite [114–116]. For asymptotically stable \mathbf{A} , it can further be shown that the solution to the above Lyapunov equation may be written as

$$\mathbf{X} = - \int_0^\infty e^{\mathbf{A}t} \mathbf{C} e^{\mathbf{A}^T t} dt. \quad (4.24)$$

The condition for asymptotic stability of a matrix is that all its eigenvalues have real parts in the open left half plane of \mathbb{C} , i.e. that they be negative in the case of real eigenvalues [114].

Recalling eq. (4.6) for the Laplace transform MP2 energy, we have

$$E_0^{(2)} = - \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} \int_0^\infty e^{-(\epsilon_a - \epsilon_i)t} (ia|jb) [2(ia|jb) - (ib|ja)] e^{-(\epsilon_b - \epsilon_j)t} dt, \quad (4.25)$$

which, using the definitions from equation (4.7), can also be written as

$$E_0^{(2)} = - \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} \left(\int_0^\infty e^{-\mathbf{D}t} (\mathbf{M} \circ \mathbf{G}) e^{-\mathbf{D}t} dt \right)_{ia,jb}, \quad (4.26)$$

where \circ denotes the Hadamard product (i.e. entry-wise) of \mathbf{M} and \mathbf{G} .

By setting $\mathbf{C} = \mathbf{M} \circ \mathbf{G}$ and $\mathbf{A} = -\mathbf{D}$ in equation (4.24), and taking into account that $-\mathbf{D}$ is diagonal, hence symmetric, and by definition negative definite, such that it fulfils the above condition of being asymptotically stable, it follows that the MP2 energy is obtained as the sum

$$E_0^{(2)} = \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} X_{ia,jb}. \quad (4.27)$$

over the elements of the matrix \mathbf{X} that solves the corresponding Lyapunov equation

$$-\mathbf{D}\mathbf{X} - \mathbf{X}\mathbf{D} = \mathbf{M} \circ \mathbf{G}. \quad (4.28)$$

Due to \mathbf{D} being diagonal, the above equation directly yields

$$X_{ia,jb} = \frac{(ia|jb) [2(ia|jb) - (ib|ja)]}{\epsilon_i - \epsilon_a + \epsilon_j - \epsilon_b}, \quad (4.29)$$

such that aside from the purely mathematical elegance of this relation which closes the circle from the Laplace transform approach to conventional MP2 by way of the theory of Lyapunov

equations, this does not seem to indicate any straightforward possibility of improvements over conventional MP2 implementations. However, for cases where the solution to the Lyapunov equation is known to be non-negative definite, there exist algorithms [115–117] to compute a low rank decomposition $\mathbf{X} = \mathbf{B}\mathbf{B}^T$ without first explicitly evaluating \mathbf{X} . Therefore, due to $\mathbf{M} \circ \mathbf{G}$ and hence \mathbf{X} being positive semidefinite in the present context, this relation hints at the possibility to obtain a low rank decomposition of the amplitude matrix from the previous equation using methods developed for the solution of large Lyapunov equations. This would represent an alternative to approaches based on the decomposition of merely the ERI matrix (see also chapter 7), and could also possibly be useful in cases involving very similar expressions, maybe some coupled cluster approaches [118]. Conversely, the methods developed in the context of Laplace-transform MP-theory could prove useful for the solution of Lyapunov equations.

4.2 Atomic-orbital MP theory (AO-MP)

In his seminal work on the application of Møller-Plesset theory to large molecules, Häser used the Laplace transform formulation to express the MP2 energy purely in terms of AO-entities [4] and has thus paved the way for linear-scaling implementations based on exploiting the inherent locality of the AO-basis [5, 6, 119]. Namely, inserting relation (4.11) for the ERI transformation into equation (4.17), leads to

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((e^{t_{\alpha}\epsilon_o} \otimes e^{-t_{\alpha}\epsilon_v}) (\mathbf{C}_o^T \otimes \mathbf{C}_v^T) \mathbf{A} (\mathbf{C}_o \otimes \mathbf{C}_v) (e^{t_{\alpha}\epsilon_o} \otimes e^{-t_{\alpha}\epsilon_v}) (\mathbf{C}_o^T \otimes \mathbf{C}_v^T) \mathbf{\Gamma} (\mathbf{C}_o \otimes \mathbf{C}_v) \right) \quad (4.30)$$

which, again invoking the invariance of the trace under cyclic permutations and using property (4.13), can straightforwardly be rearranged to

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\mathbf{C}_o e^{t_{\alpha}\epsilon_o} \mathbf{C}_o^T \otimes \mathbf{C}_v e^{-t_{\alpha}\epsilon_v} \mathbf{C}_v^T) \mathbf{A} (\mathbf{C}_o e^{t_{\alpha}\epsilon_o} \mathbf{C}_o^T \otimes \mathbf{C}_v e^{-t_{\alpha}\epsilon_v} \mathbf{C}_v^T) \mathbf{\Gamma} \right) \quad (4.31)$$

$$= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\underline{\mathbf{P}}^{(\alpha)} \otimes \overline{\mathbf{P}}^{(\alpha)}) \mathbf{A} (\underline{\mathbf{P}}^{(\alpha)} \otimes \overline{\mathbf{P}}^{(\alpha)}) \mathbf{\Gamma} \right) \quad (4.32)$$

$$= - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{\mu \cdots \sigma'} \underline{P}_{\mu\mu'}^{(\alpha)} \overline{P}_{\nu\nu'}^{(\alpha)} A_{\mu\nu\lambda\sigma} \underline{P}_{\lambda\lambda'}^{(\alpha)} \overline{P}_{\sigma\sigma'}^{(\alpha)} \Gamma_{\mu\nu\lambda\sigma} \quad (4.33)$$

This AO-based formulation can be extended to arbitrary order n of MP theory through the use of individual integral representations for each orbital energy term $\Delta_{kr,ls}$ occurring in the canonical MP n expression, as has been shown explicitly for MP3 in [7, 23]. Due to the high algorithmic complexity that this entails, however, and given that the cost to quality ratio of higher order MP n terms progressively deteriorates, the implementation of AO-MP3 and higher orders has so far not received much attention in our group, beyond the preliminary investigations mentioned above. The discussion will therefore be limited to MP2 in the following.

4.2.1 Pseudo-density matrices and operators

The matrices introduced in equation (4.32),

$$\underline{\mathbf{P}}^{(\alpha)} = \mathbf{C}_o e^{t\alpha\epsilon_o} \mathbf{C}_o^T \quad \overline{\mathbf{P}}^{(\alpha)} = \mathbf{C}_v e^{-t\alpha\epsilon_v} \mathbf{C}_v^T, \quad (4.34)$$

where the underline and overline serve to denote the occupied and virtual matrix, bear a certain resemblance to the energy weighted density matrices that occur in the context of HF energy gradients, albeit with an exponential weighting of the orbital energies. For numerical reasons, it is convenient to introduce a scaling parameter ϵ_F according to

$$\underline{\mathbf{P}}^{(\alpha)} = \mathbf{C}_o e^{t\alpha(\epsilon_o - \epsilon_F \mathbf{1}_o)} \mathbf{C}_o^T \quad \overline{\mathbf{P}}^{(\alpha)} = \mathbf{C}_v e^{-t\alpha(\epsilon_v - \epsilon_F \mathbf{1}_v)} \mathbf{C}_v^T \quad (4.35)$$

which is generally chosen as the Fermi-level $\epsilon_F = (\epsilon_{\text{LUMO}} + \epsilon_{\text{HOMO}})/2$ [4]. This parameter, however, has no effect on the following theoretical considerations and will thus be omitted, along with the explicit dependence on the quadrature point α , as it is implied by the underlines and overlines.

As discussed in section 2.4 for the occupied and virtual one-particle density matrices, the above matrices may be associated with corresponding operators

$$\hat{\rho}(t) = |\phi_o\rangle e^{\epsilon_o t} \langle \phi_o| = |\chi\rangle \underline{\mathbf{P}}(\chi) \quad (4.36)$$

$$\hat{\bar{\rho}}(t) = |\phi_v\rangle e^{-\epsilon_v t} \langle \phi_v| = |\chi\rangle \overline{\mathbf{P}}(\chi), \quad (4.37)$$

which, in contrast to their analogues, lack the idempotency of projection operators, for one has

$$\hat{\rho}^2(t) = |\phi_o\rangle e^{\epsilon_o t} \langle \phi_o | \phi_o \rangle e^{\epsilon_o t} \langle \phi_o| = |\phi_o\rangle e^{2\epsilon_o t} \langle \phi_o| \quad (4.38)$$

$$\hat{\bar{\rho}}^2(t) = |\phi_v\rangle e^{-\epsilon_v t} \langle \phi_v | \phi_v \rangle e^{-\epsilon_v t} \langle \phi_v| = |\phi_v\rangle e^{-2\epsilon_v t} \langle \phi_v|. \quad (4.39)$$

They are therefore called *pseudo-density operators* or *matrices*, respectively. We deduce from the previous relation that upon repeated application

$$\lim_{n \rightarrow \infty} (\hat{\rho})^n(t) = \lim_{n \rightarrow \infty} |\phi_o\rangle e^{n\epsilon_o t} \langle \phi_o| = 0 \quad (4.40)$$

$$\lim_{n \rightarrow \infty} (\hat{\bar{\rho}})^n(t) = \lim_{n \rightarrow \infty} |\phi_v\rangle e^{-n\epsilon_v t} \langle \phi_v| = 0, \quad (4.41)$$

since ϵ_o and $-\epsilon_v$ are negative definite. They do, however, have the property of being mutually orthogonal,

$$\hat{\rho}(t) \hat{\bar{\rho}}(t) = |\phi_o\rangle e^{\epsilon_o t} \langle \phi_o | \phi_v \rangle e^{-\epsilon_v t} \langle \phi_v| = 0, \quad (4.42)$$

or for their corresponding matrices $\underline{\mathbf{P}}\overline{\mathbf{P}} = \mathbf{0}$, which is of particular importance in the context of integral screening, as is discussed in section 5.1.3. Upon comparison with equation (4.32), it further follows that by defining an operator according to

$$\hat{\rho}(t) = \hat{\rho}(t) \otimes \hat{\bar{\rho}}(t) = |\chi \otimes \chi\rangle (\underline{\mathbf{P}} \otimes \overline{\mathbf{P}}) (\chi \otimes \chi), \quad (4.43)$$

the MP2 energy may be expressed in rough analogy to an expectation value (cf. section 2.4) as

$$E_0^{(2)} = \int_0^\infty \text{Tr} \left(\hat{\rho}(t) \hat{C} \hat{\rho}(t) (2 - \mathcal{P}) \hat{C} \right) dt, \quad (4.44)$$

though of course t is not a configuration space variable. As in section 2.3, \mathcal{P} denotes the permutation accounting for the antisymmetry and \hat{C} is the Coulombic repulsion operator.

4.2.2 The MP n energy as a functional of the density matrix

In equation (2.38), it is shown that the Hartree-Fock energy is completely specified by the density matrix \mathbf{P}_o and can thus be calculated without diagonalization of the Fock matrix. In view of the combination of AO-MP with such implementations, the question naturally arises whether the canonical orbital energies occurring in the definition of the pseudo-density matrices or operators can be avoided. As has first been pointed out by Surján [120] in the present context, this is indeed the case, as can be readily established by inserting relation (4.16) into the definition (4.34) for the pseudo-densities (ignoring the scaling with the Fermi-energy which is of no concern here)

$$\underline{\mathbf{P}} = \mathbf{C}_o e^{t_\alpha \mathbf{C}_o^T \mathbf{F} \mathbf{C}_o} \mathbf{C}_o^T. \quad (4.45)$$

Explicitly considering the series defining the matrix exponential

$$\underline{\mathbf{P}} = \mathbf{C}_o \left(\sum_{n=1}^{\infty} \frac{(t_\alpha \mathbf{C}_o^T \mathbf{F} \mathbf{C}_o)^n}{n!} \right) \mathbf{C}_o^T, \quad (4.46)$$

the matrix \mathbf{C}_o can be shifted back and forth, such that it follows from the definition $\mathbf{P}_o = \mathbf{C}_o \mathbf{C}_o^T$ for the occupied one-electron density matrix that

$$\underline{\mathbf{P}} = \mathbf{C}_o \mathbf{C}_o^T \left(\sum_{n=1}^{\infty} \frac{(t_\alpha \mathbf{F} \mathbf{C}_o \mathbf{C}_o^T)^n}{n!} \right) = \mathbf{P}_o e^{t_\alpha \mathbf{F} \mathbf{P}_o}. \quad (4.47)$$

Given that $\underline{\mathbf{P}}$, as well as \mathbf{P}_o and \mathbf{F} , are Hermitean, i.e. real-symmetric in the cases that are of interest for us, it is also obvious that

$$\underline{\mathbf{P}} = \underline{\mathbf{P}}^T = (\mathbf{P}_o e^{t_\alpha \mathbf{F} \mathbf{P}_o})^T = e^{t_\alpha \mathbf{P}_o \mathbf{F} \mathbf{P}_o}. \quad (4.48)$$

For the virtual pseudo-density matrix, the corresponding relation is given by

$$\bar{\mathbf{P}} = \mathbf{P}_v e^{-t_\alpha \mathbf{F} \mathbf{P}_v} = e^{-t_\alpha \mathbf{P}_v \mathbf{F} \mathbf{P}_v} \quad (4.49)$$

Recalling the complementarity relation $\mathbf{P}_v \mathbf{S} = \mathbf{1} - \mathbf{P}_o \mathbf{S}$, this evidently establishes the AO-MP n energy as a functional $E_0^{(n)}[\mathbf{P}_o]$ of the occupied one-electron density matrix.

Chapter 5

Cholesky Decomposed Densities in MP2

Following the preliminaries established in the previous chapter, we now arrive at the main subject of this dissertation, the Cholesky decomposed pseudo-densities in MP2 method (CDD-MP2). We begin with a discussion of the limitations and formal disadvantages of atomic-orbital based MP2 as previously implemented by Doser *et al.* [5–7], which motivated the investigation of symmetrically decomposed pseudo-density matrices occurring in AO-MP2 theory. Following a brief analysis of the relation between CDD-MP2 and local MP theory, the pre-selection scheme that is essential to determine significant contributions to the energy and perform the corresponding transformation of the ERI matrix in an efficient way is presented. After a detailed description of the implementation of the CDD-MP2 algorithm, results are presented, beginning with the locality of the Cholesky orbitals, efficiency of the screening procedure to reduce integral numbers and the numerical accuracy afforded hereby, followed by a discussion of the method’s performance and the encountered difficulties.

5.1 Cholesky-decomposed pseudo-density matrices (CDD-MP2)

In Häser’s initial AO-MP2 formulation [4] and Ayala’s subsequent implementation [119], the transformation with the pseudo-densities in equation (4.32) was performed according to

$$\mathbf{A}_{\text{FT}} = (\underline{\mathbf{P}} \otimes \overline{\mathbf{P}}) \mathbf{A} (\underline{\mathbf{P}} \otimes \overline{\mathbf{P}}) \quad (5.1)$$

from which the MP2 energy is then calculated as

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} (\mathbf{A}_{\text{FT}} \mathbf{\Gamma}) = - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{\mu\nu\lambda\sigma}^N (\underline{\mu\nu} | \underline{\lambda\sigma}) [2(\mu\nu | \lambda\sigma) - (\mu\sigma | \lambda\nu)], \quad (5.2)$$

where the underlined and overlined indices denote transformation with the pseudo-densities. In contrast to this, the implementation of Doser [7] is based on a half-transformation of ERI matrices according to

$$\mathbf{A}_{\text{HT}} = (\mathbf{P} \otimes \overline{\mathbf{P}})\mathbf{A}. \quad (5.3)$$

The immediate advantage herein lies in the greatly simplified integral transformation and contraction, albeit only for the case of the opposite spin part

$$E_{\text{OS}}^{(2)} = \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr}(\mathbf{A}_{\text{HT}}\mathbf{A}_{\text{HT}}) = \sum_{\alpha}^{\tau} w_{\alpha} \sum_{\mu\nu\lambda\sigma}^N (\underline{\mu}\overline{\nu}|\lambda\sigma)(\underline{\lambda}\overline{\sigma}|\mu\nu), \quad (5.4)$$

Namely, in terms of half-transformed ERI the exchange term required for the same-spin component (see equation 3.24) is explicitly given by

$$E_{\text{K}}^{(2)} = \sum_{\alpha}^{\tau} w_{\alpha} \sum_{\mu\nu\lambda\sigma}^N (\underline{\mu}\nu|\underline{\lambda}\sigma)(\underline{\mu}\overline{\sigma}|\lambda\overline{\nu}) \quad (5.5)$$

which again leads to great difficulties for an efficient implementation, as one has to perform two different types of transformations and handle two ERI-arrays. This implementation is therefore only practicable in conjunction with the scaled-opposite spin parametrization scheme (SOS-MP2) [8]

$$E_{\text{SOS}}^{(2)} = C_{\text{SOS}}E_{\text{OS}}^{(2)}, \quad (5.6)$$

and is thus subject to its shortcomings, as discussed previously in section 3.3.3. It was further noticed that summands of different sign in the contraction (5.4) cancel each other to a large extent, i.e.

$$\sum_{\alpha}^{\tau} w_{\alpha} \sum_{\mu\nu\lambda\sigma}^N |(\underline{\mu}\overline{\nu}|\lambda\sigma)(\underline{\lambda}\overline{\sigma}|\mu\nu)| \gg \left| \sum_{\alpha}^{\tau} w_{\alpha} \sum_{\mu\nu\lambda\sigma}^N (\underline{\mu}\overline{\nu}|\lambda\sigma)(\underline{\lambda}\overline{\sigma}|\mu\nu) \right|. \quad (5.7)$$

from which it was anticipated that the number of required integrals, hence the prefactor, might be reduced if this cancellation was avoided by summing up products of equally transformed integrals, which would all be positive.

These two facts motivated the exploration of symmetrically decomposed pseudo-density matrices, which could then be used as transformation matrices for the integrals, thus in principle solving both problems just mentioned. For that end, both the Cholesky decomposition and matrix square root may be envisaged. However, initial investigations showed that the latter presents no particular improvement [23], such that subsequent work was based solely on the Cholesky decomposition.

5.1.1 Cholesky decomposition

The Cholesky decomposition [9] of a positive definite $N \times N$ Hermitean matrix \mathbf{V} yields a lower-triangular matrix \mathbf{L} with the property

$$\mathbf{V} = \mathbf{L}\mathbf{L}^T. \quad (5.8)$$

The pseudo-density matrices, however, are positive *semi*-definite by definition, being formed as the outer product of only N_{occ} and N_{virt} linearly independent vectors from the MO-coefficient matrix \mathbf{C} (equation (4.34)). A pivot algorithm therefore has to be employed, as described in section 5.2.1, yielding the decomposition:

$$\underline{\mathbf{P}} = \underline{\mathbf{L}}\underline{\mathbf{L}}^T \quad \overline{\mathbf{P}} = \overline{\mathbf{L}}\overline{\mathbf{L}}^T, \quad (5.9)$$

where $\underline{\mathbf{L}}$ and $\overline{\mathbf{L}}$ have dimension $N \times N_{\text{occ}}$ and $N \times N_{\text{virt}}$, respectively. Upon insertion into equation 4.32, we obtain [10, 11]

$$\begin{aligned} E_0^{(2)} &= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\underline{\mathbf{L}}\underline{\mathbf{L}}^T \otimes \overline{\mathbf{L}}\overline{\mathbf{L}}^T) \mathbf{A} (\underline{\mathbf{L}}\underline{\mathbf{L}}^T \otimes \overline{\mathbf{L}}\overline{\mathbf{L}}^T) \mathbf{\Gamma} \right) \\ &= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\underline{\mathbf{L}}^T \otimes \overline{\mathbf{L}}^T) \mathbf{A} (\underline{\mathbf{L}} \otimes \overline{\mathbf{L}}) (\underline{\mathbf{L}}^T \otimes \overline{\mathbf{L}}^T) \mathbf{\Gamma} (\underline{\mathbf{L}} \otimes \overline{\mathbf{L}}) \right) \end{aligned} \quad (5.10)$$

$$= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} (\mathbf{M}\mathbf{G}), \quad (5.11)$$

where the rearrangement in the second equation is based on property (4.13) and invariance of the trace under cyclic permutations. The last term is explicitly written as

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{i,j}^{N_{\text{occ}}} \sum_{a,b}^{N_{\text{virt}}} (i\bar{a}|j\bar{b}) [2(i\bar{a}|j\bar{b}) - (i\bar{b}|j\bar{a})], \quad (5.12)$$

from which we notice that apart from fulfilling the initial purpose, this approach has the formal advantage of reintroducing the restriction of indices to the occupied and virtual subspaces. Due to N_{occ} being constant for a given molecule, this approach is expected to display a smaller increase of the prefactor with basis set size than the AO-based expression where indices run over the entire basis set.

5.1.2 CDD as localized pseudo-MO

It is shown in section 5.3.1 that the Cholesky factors inherit the sparsity of the pseudo-densities, and as such they may be interpreted as exponentially energy-weighted localized MO-coefficient matrices. This also becomes evident by a closer analysis of their structure. In equation (4.47), we can take the square root of the exponential matrix, and with the shifting property used in (4.46), we can write

$$\underline{\mathbf{P}} = e^{\frac{t_{\alpha}}{2} \mathbf{P}_o \mathbf{F}} \mathbf{P}_o e^{\frac{t_{\alpha}}{2} \mathbf{F} \mathbf{P}_o}. \quad (5.13)$$

Inserting the Cholesky decomposition of \mathbf{P}_o (see also section 3.4.1)—or any localized coefficient matrix \mathbf{L}_o , for that matter—we obtain

$$\underline{\mathbf{P}} = e^{\frac{t_\alpha}{2} \mathbf{L}_o \mathbf{L}_o^T \mathbf{F}} \mathbf{L}_o \mathbf{L}_o^T e^{\frac{t_\alpha}{2} \mathbf{F} \mathbf{L}_o \mathbf{L}_o^T}. \quad (5.14)$$

The corresponding relation holds for the virtual pseudo-density as well, and we may thus identify the pseudo-density Cholesky factors as [10]

$$\underline{\mathbf{L}}' = e^{\frac{t_\alpha}{2} \mathbf{L}_o \mathbf{L}_o^T \mathbf{F}} \mathbf{L}_o = \mathbf{L}_o e^{\frac{t_\alpha}{2} \mathbf{L}_o^T \mathbf{F} \mathbf{L}_o} \quad (5.15)$$

$$\overline{\mathbf{L}}' = e^{-\frac{t_\alpha}{2} \mathbf{L}_v \mathbf{L}_v^T \mathbf{F}} \mathbf{L}_v = \mathbf{L}_v e^{-\frac{t_\alpha}{2} \mathbf{L}_v^T \mathbf{F} \mathbf{L}_v}, \quad (5.16)$$

at least up to a "phase factor", i.e. unitary rotation $\underline{\mathbf{L}} = \underline{\mathbf{L}}' \mathbf{U}'_o$ and $\overline{\mathbf{L}} = \overline{\mathbf{L}}' \mathbf{U}'_v$. This phase factor is of no concern in the present context, however, and if required, it could be recovered from

$$e^{-\frac{t_\alpha}{2} \mathbf{L}_o^T \mathbf{F} \mathbf{L}_o} \mathbf{L}_o^T \mathbf{S} \underline{\mathbf{L}} = \mathbf{U}'_o \quad (5.17)$$

$$e^{\frac{t_\alpha}{2} \mathbf{L}_v^T \mathbf{F} \mathbf{L}_v} \mathbf{L}_v^T \mathbf{S} \overline{\mathbf{L}} = \mathbf{U}'_v, \quad (5.18)$$

by virtue of the orthonormality of the local MOs and the fact that the inverse of an exponential matrix always exists. These considerations establish CDD-MP2 as an instance of local Laplace MP2 as discussed in general terms in section 4.1.1. It follows from the previous equations that these Cholesky MO are not orthonormal, since we have

$$\underline{\mathbf{L}}^T \mathbf{S} \underline{\mathbf{L}} = e^{\frac{t_\alpha}{2} \mathbf{L}_o^T \mathbf{F} \mathbf{L}_o} \mathbf{L}_o^T \mathbf{S} \mathbf{L}_o e^{\frac{t_\alpha}{2} \mathbf{L}_o^T \mathbf{F} \mathbf{L}_o} = e^{t_\alpha \mathbf{L}_o^T \mathbf{F} \mathbf{L}_o} \neq \mathbf{1}_o \quad (5.19)$$

$$\overline{\mathbf{L}}^T \mathbf{S} \overline{\mathbf{L}} = e^{-\frac{t_\alpha}{2} \mathbf{L}_v^T \mathbf{F} \mathbf{L}_v} \mathbf{L}_v^T \mathbf{S} \mathbf{L}_v e^{-\frac{t_\alpha}{2} \mathbf{L}_v^T \mathbf{F} \mathbf{L}_v} = e^{-t_\alpha \mathbf{L}_v^T \mathbf{F} \mathbf{L}_v} \neq \mathbf{1}_v, \quad (5.20)$$

but that the occupied and virtual subspaces are still orthogonal,

$$\underline{\mathbf{L}}^T \mathbf{S} \overline{\mathbf{L}} = \mathbf{0}, \quad (5.21)$$

which is the more important property in view of integral screening.

5.1.3 Integral screening

It has been discussed in the previous section that the Cholesky factors of the pseudo-densities correspond to exponentially weighted local MO coefficient matrices. Accordingly, the number of significant integral products in the contraction to the CDD-MP2 energy should asymptotically be linear. Now, it is described in section 2.5.1 how the Cauchy-Schwarz-Bunyakovsky inequality (2.64) can be used to preselect the elements $(\mu\nu|\lambda\sigma)$ of the ERI matrix according to their absolute value. It has further been shown by Häser [4] that this can be extended to the transformed ERI as well, such that in our case the upper bounds

$$\begin{aligned} |(\underline{i}\nu|\lambda\sigma)| &\leq (\underline{i}\nu|\underline{i}\nu)^{\frac{1}{2}} (\lambda\sigma|\lambda\sigma)^{\frac{1}{2}} = X_{i\nu} Q_{\lambda\sigma} \\ |(\mu\overline{a}|\lambda\sigma)| &\leq (\mu\overline{a}|\mu\overline{a})^{\frac{1}{2}} (\lambda\sigma|\lambda\sigma)^{\frac{1}{2}} = Y_{a\mu} Q_{\lambda\sigma} \\ |(\underline{i}\overline{a}|\lambda\sigma)| &\leq (\underline{i}\overline{a}|\underline{i}\overline{a})^{\frac{1}{2}} (\lambda\sigma|\lambda\sigma)^{\frac{1}{2}} = Z'_{ia} Q_{\lambda\sigma} \\ |(\underline{i}\overline{a}|\underline{j}\overline{b})| &\leq (\underline{i}\overline{a}|\underline{i}\overline{a})^{\frac{1}{2}} (\underline{j}\overline{b}|\underline{j}\overline{b})^{\frac{1}{2}} = Z'_{ia} Z'_{jb}, \end{aligned} \quad (5.22)$$

provide an estimate for any type of (partially) transformed integral. The matrices \mathbf{Q} , \mathbf{X} and \mathbf{Y} are rather straightforward to compute, with $\mathcal{O}(N^3)$ effort, but the situation is more complicated for \mathbf{Z}' . Namely, the explicit evaluation of the $(\bar{i}a|\bar{i}a)$ would scale as $\mathcal{O}(N^5)$, which is unacceptable for efficiency reasons. On the other hand, in order to be reliable as screening matrices, their evaluation, as well as that of the previous screening matrices, cannot be based on the exploitation of the sparsity of the LMO-coefficient matrices to reduce transformation time, since this would induce a numerical error that may be of the order of magnitude of the integrals that are to be preselected. In order to keep the computational cost of this step within reasonable bounds, a matrix \mathbf{Z} is thus built as [4]

$$Z_{ia} = \min \left(\sum_{\lambda}^N X_{i\lambda} |\bar{L}_{\lambda a}|, \sum_{\lambda}^N Y_{a\lambda} |\underline{L}_{\lambda i}| \right) \geq Z'_{ia}, \quad (5.23)$$

which is an upper bound to \mathbf{Z}' by virtue of the Cauchy inequality.

This screening procedure, however, only exploits the decay of charge distributions with the distance between the centers of their constituting orbitals, and the number of preselected integrals can thus at best be reduced to quadratic. In order to achieve linear scaling, one needs to account for the fact that ERI depend on the separation R between the respective charge distributions. For that end, the Schwarz upper bounds are modified by a reduced distance parameter R' , as recently proposed by Maurer *et al.* [12]. Although these modified Schwarz estimates do not represent strict upper bounds to the ERI any more, it has been shown on a wide selection of molecular systems that this procedure causes no significant loss of reliability. The rationale behind this approach is based on the transformation behaviour of the multipole expansion of the ERI, which, as has already been mentioned in section 2.5.3, is valid for well-separated charge distributions. In order to separate the variables acted upon by the Coulombic repulsion operator, it can be expanded as a series (cf. ref. [56] and references given therein for more details on this subject)

$$\hat{C} = \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} = \frac{1}{|\mathbf{r}' - \mathbf{r}'' + \mathbf{R}|} \quad (5.24)$$

$$= \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{k=0}^{\infty} \sum_{l=-k}^k (-1)^k \mathcal{R}_{nm}(\mathbf{r}') \mathcal{I}_{n+k,m+l}^*(\mathbf{R}) \mathcal{R}_{kl}(\mathbf{r}''), \quad (5.25)$$

where $\mathbf{r}' = \mathbf{r}_1 - \mathbf{R}'$ and $\mathbf{r}'' = \mathbf{r}_2 - \mathbf{R}''$ denote the coordinates of electrons 1 and 2 with respect to the centers \mathbf{R}' and \mathbf{R}'' of some charge distributions $\Omega(\mathbf{r}_1)$ and $\Omega(\mathbf{r}_2)$, and $\mathbf{R} = \mathbf{R}' - \mathbf{R}''$. Further, \mathcal{R} and \mathcal{I} denote the scaled regular and irregular solid harmonics, whose explicit form is of secondary importance here (see reference given above), but behave as

$$\mathcal{R}_{nm}(\mathbf{r}') \propto |\mathbf{r}'|^n = r'^n \quad (5.26)$$

$$\mathcal{I}_{n+k,m+l}^*(\mathbf{R}) \propto |\mathbf{R}|^{-(n+k+1)} = R^{-(n+k+1)}. \quad (5.27)$$

In the following, we shall also not be concerned with particular aspects of the transformation from primitive to contracted Gaussians, as this doesn't affect the point to be made here. Arranging the harmonics into a vector and a matrix—both infinite-dimensional—according to

$$|\mathcal{R}\rangle \equiv |\mathcal{R}_{nm}(\mathbf{r}')\rangle \quad (5.28)$$

$$\mathcal{T} \equiv \mathcal{T}_{nm,kl}(\mathbf{R}) = (-1)^k \mathcal{T}_{n+k,m+l}^*, \quad (5.29)$$

where as in previous cases, the notation in terms of kets serves to denote the operator character of the expression, the above expansion may be suggestively summarized as

$$\hat{C} = |\mathcal{R}\rangle \mathcal{T} \langle \mathcal{R}|. \quad (5.30)$$

From this, recalling equations (4.7) and (2.37), the ERI matrix is obtained as

$$\mathbf{A} = (\Omega | \hat{C} | \Omega) = (\Omega | \mathcal{R}) \mathcal{T} (\mathcal{R} | \Omega) = \mathcal{Q}^T \mathcal{T} \mathcal{Q}, \quad (5.31)$$

the elements of the $\infty \times N^2$ matrix $\mathcal{Q} \equiv \mathcal{Q}(\mathbf{R}')$ of spherical multipoles being given by

$$\mathcal{Q}_{kl,\lambda\sigma}(\mathbf{R}') = \int \mathcal{R}_{kl}(\mathbf{r}') \Omega_{\lambda\sigma}(\mathbf{r}) d\mathbf{r}. \quad (5.32)$$

Considering the monopoles, i.e. the first row of \mathcal{Q} ($k = 0$ and thus $l = 0$), we have

$$\mathcal{Q}_{00,\lambda\sigma}(\mathbf{R}') = \int \mathcal{R}_{00}(\mathbf{r}') \Omega_{\lambda\sigma}(\mathbf{r}) d\mathbf{r}, \quad (5.33)$$

such that from the proportionality relation (5.26), it follows that apart from a scaling factor and angular part, this row is just equal to the vectorized overlap matrix \mathbf{S} , since

$$\mathcal{Q}_{00,\lambda\sigma} \propto \int |\mathbf{r}'|^0 \Omega_{\lambda\sigma}(\mathbf{r}) d\mathbf{r} = \int \chi_\lambda(\mathbf{r}) \chi_\sigma(\mathbf{r}) d\mathbf{r} = S_{\lambda\sigma}. \quad (5.34)$$

Thus, upon performing the first half-transformation to the MO-basis according to

$$\mathbf{A}_{\text{HT}} = \mathbf{A}(\underline{\mathbf{L}} \otimes \bar{\underline{\mathbf{L}}}) = \mathcal{Q}^T \mathcal{T} \mathcal{Q}(\underline{\mathbf{L}} \otimes \bar{\underline{\mathbf{L}}}) = \mathcal{Q}'^T \mathcal{T} \mathcal{Q}' \quad (5.35)$$

and using that $\text{vec}^T(\mathbf{X})(\mathbf{V} \otimes \mathbf{W}^T) = \text{vec}^T(\mathbf{W}\mathbf{X}\mathbf{V})$, it follows by virtue of the orthogonality relation (5.21) that

$$\text{vec}^T(\mathbf{S})(\underline{\mathbf{L}} \otimes \bar{\underline{\mathbf{L}}}) = \text{vec}^T(\bar{\underline{\mathbf{L}}}^T \mathbf{S} \underline{\mathbf{L}}) = \text{vec}^T(\mathbf{0}), \quad (5.36)$$

such that the first row of $\mathcal{Q}' = \mathcal{Q}(\underline{\mathbf{L}} \otimes \bar{\underline{\mathbf{L}}})$, which represents the transformed monopoles, vanishes. Accordingly, since an element of \mathbf{A}_{HT} is obtained from

$$(\mu\nu | \underline{j}\bar{b}) = \sum_{n=0}^{\infty} \sum_{m=-n}^n \sum_{k=0}^{\infty} \sum_{l=-k}^k \mathcal{Q}_{\mu\nu,nm} \mathcal{T}_{nm,kl} \mathcal{Q}'_{kl,j\bar{b}}, \quad (5.37)$$

and given that the terms for which $k = l = 0$ drop out and only $k \geq 1$ remains, it follows from the proportionality relation (5.27) that the leading order of R in the half-transformed integrals is R^{-2} . After the second half-transformation

$$\mathbf{A}_{\text{FT}} \equiv \mathbf{M} = (\underline{\mathbf{L}}^T \otimes \bar{\mathbf{L}}^T) \mathbf{Q}^T \mathcal{T} \mathbf{Q}' = \mathbf{Q}'^T \mathcal{T} \mathbf{Q}, \quad (5.38)$$

the same argument leads to the realization that only terms of $\mathcal{O}(R^{-3})$ remain in the multipole expansion. This behaviour is invoked to justify the modification of the Schwarz estimates of half-transformed and fully transformed ERI according to

$$|(\mu\nu|\underline{j}\bar{b})| \lesssim \frac{Q_{\mu\nu} Z_{jb}}{R'^2} \quad (5.39)$$

$$|(i\bar{a}|\underline{j}\bar{b})| \lesssim \frac{Z_{ia} Z_{jb}}{R'^3}, \quad (5.40)$$

for the case of well-separated distributions. The criterion of well-separatedness is defined as $R' = R_{ia,jb} - (\text{Ext}_{ia} + \text{Ext}_{jb}) > 1$, where $R_{ia,jb}$ stands for the distance between the centers of the charge distributions and Ext denotes their respective extents. These centers are determined as the weighted means of the centers of untransformed distributions

$$\text{Cen}_{ia}^{(q)} = \frac{\sum_{\mu\nu} |\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}| \text{Cen}_{\mu\nu}^{(q)}}{\sum_{\mu\nu} |\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|}, \quad (5.41)$$

where $q = x, y, z$. The idea behind including the overlap elements into the weighting factors $|\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|$ in addition to the transformation matrix elements is that they provide a measure of a distribution's significance, since mutually distant Gaussians will have vanishing overlap.

The definition of the transformed extents, however, leaves more possibilities. Initially, the transformed extents were defined in the same way as the centers, i.e. as the weighted means of the extents of untransformed distributions plus their distance $R_{ia,\mu\nu}$ from the respective distribution's center

$$\text{Ext}_{ia} = \frac{\sum_{\mu\nu} |\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}| (R_{ia,\mu\nu} + \text{Ext}_{\mu\nu})}{\sum_{\mu\nu} |\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|}. \quad (5.42)$$

However, it was realized that this definition might not be able to correctly account for orbital "bulges" distant from the centers. It was thus modified so as to contain the farthest untransformed distribution whose weight factor exceeds a predefined threshold value:

$$\text{Ext}_{ia} = \max_{\mu,\nu} \left\{ \frac{|\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|}{\sum_{\mu\nu} |\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|} \text{Ext}_{\mu\nu} + R_{ia,\mu\nu} \mid \frac{|\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|}{\sum_{\mu\nu} |\underline{L}_{\mu i} S_{\mu\nu} \bar{L}_{\nu a}|} > \theta \right\}. \quad (5.43)$$

5.2 Implementation of CDD-MP2

This section is dedicated to the presentation of algorithmic details of the CDD-MP2 method as it has been implemented in a development version of the quantum chemical programme package

Q-CHEM [121]. In the following, the different types of (partially) transformed integrals will be abbreviated as

UTI	(untransformed integrals):	$(\mu\nu \lambda\sigma)$
STI	(single transformed integrals):	$(i\nu \lambda\sigma)$
HTI	(half transformed integrals):	$(i\bar{a} \lambda\sigma)$
TTI	(triple transformed integrals):	$(i\bar{a} \underline{j}\sigma)$
FTI	(fully transformed integrals):	$(i\bar{a} \underline{j}\bar{b})$

The algorithm consists of the following steps (for each quadrature point):

- Cholesky decomposition of pseudo-density matrices
- Formation of screening prerequisites (screening matrices, centers and extents of transformed charge distributions)
- Screening for significant FTI products $(i\bar{a}|\underline{j}\bar{b})(i\bar{a}|\underline{j}\bar{b})$ and/or $(i\bar{a}|\underline{j}\bar{b})(i\bar{b}|\underline{j}\bar{a})$
- Computation of pre-selected FTI:
 - Screening for required HTI $(i\bar{a}|\lambda\sigma)$
 - Computation of HTI according to screening record
 - ▷ Screening for required STI $(i\nu|\lambda\sigma)$, then for UTI $(\mu\nu|\lambda\sigma)$
 - ▷ Transformations to HTI
 - Transformations to FTI
- Contraction of FTI to MP2 energy

The chosen approach is purely integral-direct, i.e. no integrals or other major arrays are written to disk. Accordingly, due to the rapid increase of integral numbers with system size, a batching scheme has to be employed. It should also be noted that even though AO-indices are retained in the following for the sake of notational clarity, all AO-based quantities are in practice grouped into shells of basis functions for computational efficiency, as described in 5.2.2.

5.2.1 The Cholesky decomposition algorithm

The pseudo-densities used as input for the Cholesky decomposition are obtained from Doser's SOS-AO-MP2 implementation [7], where they are computed according to their definition given in 4.2.1, but with direct inclusion of the numerical quadrature weighting factors w_α , i.e.

$$\underline{\mathbf{P}} = |w_\alpha|^{\frac{1}{4}} \mathbf{C}_o e^{t_\alpha(\epsilon_o - \epsilon_F \mathbf{1}_o)} \mathbf{C}_o^T \quad \overline{\mathbf{P}}^{(\alpha)} = |w_\alpha|^{\frac{1}{4}} \mathbf{C}_v e^{-t_\alpha(\epsilon_v - \epsilon_F \mathbf{1}_v)} \mathbf{C}_v^T. \quad (5.44)$$

Even though they could in principle be obtained without the canonical orbital energies, as shown in section 4.2.2, this has so far not been subject to optimization endeavours.

As mentioned in section 5.1.1, the pseudo-density matrices are positive semi-definite by definition, i.e. their rank—the number of non-zero eigenvalues—is smaller than their dimension N (except for exceptional cases involving minimal bases), and their Cholesky decomposition

$$\underline{\mathbf{P}} = \underline{\mathbf{L}}\underline{\mathbf{L}}^T \quad \overline{\mathbf{P}} = \overline{\mathbf{L}}\overline{\mathbf{L}}^T \quad (5.45)$$

must therefore be performed with a pivoting technique. For this, a slightly modified version of an $\mathbf{L}'\mathbf{D}\mathbf{L}'^T$ -decomposition algorithm described in [122] is used, which has been adapted for direct computation of the Cholesky decomposition according to $\mathbf{L}\mathbf{L}^T$, with $\mathbf{L} = \mathbf{L}'\mathbf{D}^{\frac{1}{2}}$. In essence, this should correspond to the algorithm proposed by Beebe and Linderberg in the context of the decomposition of the ERI matrix [123].

Formulated recursively, the Cholesky decomposition of an input matrix \mathbf{V} is computed as [122]:

- If $\dim(\mathbf{V}) = 1$, define $\mathbf{L} = \mathbf{V}^{\frac{1}{2}}$
- Determine $\hat{\mathbf{P}}$ as the permutation of 1 and j such that $V_{jj} = \max_{i \in N} \{V_{ii}\}$. Define then the $N \times N$ matrix \mathbf{U} as

$$\hat{\mathbf{P}}\mathbf{V}\hat{\mathbf{P}}^T =: \mathbf{U} = \begin{pmatrix} U_{11} & \mathbf{u}^T \\ \mathbf{u} & \tilde{\mathbf{U}} \end{pmatrix} \quad (5.46)$$

- a) If $U_{11} > \theta_{CD}$: compute the Cholesky decomposition of the $(N-1) \times (N-1)$ matrix $\tilde{\mathbf{U}} - \frac{1}{U_{11}}\mathbf{u}\mathbf{u}^T$:

$$\tilde{\mathbf{P}}(\tilde{\mathbf{U}} - \frac{1}{U_{11}}\mathbf{u}\mathbf{u}^T)\tilde{\mathbf{P}}^T = \tilde{\mathbf{L}}'\tilde{\mathbf{L}}'^T \quad (5.47)$$

Then define

$$\mathbf{L}' := \begin{pmatrix} \sqrt{U_{11}} & 0 \\ \frac{1}{\sqrt{U_{11}}}\tilde{\mathbf{P}}\mathbf{u} & \tilde{\mathbf{L}}' \end{pmatrix}, \quad \mathbf{P} := \begin{pmatrix} 1 & 0 \\ 0 & \tilde{\mathbf{P}} \end{pmatrix} \hat{\mathbf{P}} \quad (5.48)$$

- b) If $U_{11} \leq \theta_{CD}$, define $\mathbf{P} = \hat{\mathbf{P}}$ and $\mathbf{L}' = \mathbf{0}$. Stop.

The matrix \mathbf{L}' thus obtained obeys the relation $\mathbf{P}\mathbf{V}\mathbf{P}^T = \mathbf{L}'\mathbf{L}'^T$, and it finally has to be permuted back according to $\mathbf{L} = \mathbf{P}^T\mathbf{L}'$, such that $\mathbf{L}\mathbf{L}^T = \mathbf{V}$ holds.

More explicitly, our implementation of this algorithm had the structure:

- Create record array P of original order of \mathbf{V} 's diagonal elements, i.e. $P[i] = i$, for $i \in N$.
- Loop $i \in (N-1)$
 - Determine $V_{jj} = \max_{k>i} \{V_{kk}\}$.
 - If $V_{jj} \leq \theta_{CD}$, set $M = \text{rank}(\mathbf{V}) = i$ and exit loop.

- Else, perform permutation $\hat{\mathbf{P}}$ of i -th and j -th row and column, keeping record of permutation by swapping $P[i]$ and $P[j]$.
- Get \mathbf{v} as $\{V_{ki} \mid k > i\}$, i.e. the sub-diagonal part of the i -th column of \mathbf{V} .
- Form submatrix $\tilde{\mathbf{V}} = \mathbf{V} - \frac{1}{V_{ii}}\mathbf{v}\mathbf{v}^T$.
- Determine $\tilde{V}_{jj} = \max_{j>i}\{\tilde{V}_{ii}\}$ to obtain the permutation $\tilde{\mathbf{P}}$.
- Form i -th column of preliminary, i.e. unpermuted \mathbf{L}'' according to $L''_{ii} = \sqrt{V_{ii}}$ and $L''_{ki} = \frac{v_k}{\sqrt{V_{ii}}}$, $k > i$.
- Perform permutation $\tilde{\mathbf{P}}$ of $(i+1)$ -th and j -th row of \mathbf{L}'' to obtain \mathbf{L}' .

End i loop

- If algorithm has not stopped for $i \in (N-1)$, which implies that $\text{rk}(\mathbf{V}) \geq N-1$:
 - If $V_{NN} > \theta_{CD}$, set $L'_{NN} = \sqrt{V_{NN}}$.

Finally, perform permutation $\mathbf{L} = \mathbf{P}^T \mathbf{L}'$ by permuting the rows back to the original order according to the record P , i.e. switching the i -th and $P[i]$ -th row for $i \in N$. The decomposition threshold is taken as

$$\theta_{CD} = \Delta_{\text{rel}} \cdot N \cdot \max_{i \in N} \{V_{ii}\}, \quad (5.49)$$

as proposed in [122], where $\Delta_{\text{rel}} = 10^{-15}$ is the relative precision of a variable of `double` type. Since it has been observed that in general most elements of $\bar{\mathbf{L}}$ are larger than those in $\underline{\mathbf{L}}$, the Cholesky factors are then scaled with the maximum element of the virtual Cholesky factor according to

$$\underline{\mathbf{L}} = \mathbf{L} \cdot \bar{L}_{\text{max}}, \quad \bar{\mathbf{L}} = \bar{\mathbf{L}} / \bar{L}_{\text{max}}, \quad (5.50)$$

in order to equilibrate the distribution of magnitudes in both matrices, in view of their use in the integral screening procedure.

Although the current implementation scales as $\mathcal{O}(N^3)$, decomposition times have so far been negligible for all investigated system, so that no optimization effort has been made. However, since the pseudo-densities are sparse—for non-metallic systems—a linear scaling implementation is feasible, as has been shown recently for the practically identical case of the one-particle density matrix \mathbf{P}_o [89].

5.2.2 Screening prerequisites

The Schwarz matrices are computed in accordance with their definitions given in 5.1.3, but AO indices are then converted to shell indices (cf. introduction in section 5.2), for instance

```

Loop I ∈ NShl
  Loop J ∈ NShl
    QIJ = maxμ,ν∈I,J{|Qμν|}
    XIj = maxμ∈I{|Xμj|}
  End J
End I
    
```

In the same way, the centers and extents of untransformed charge distributions occurring in the definitions (5.41) and (5.43) are given for shells IJ instead of AOs $\mu\nu$.

5.2.3 External screening for FTI

The pre-selection of significant FTI products according to the procedure described in section (5.1.3), the so-called *external* screening, is performed rather straightforwardly:

```

Loop ia ∈ {ia | Zia ≥ θsparse}
  Loop jb ∈ {jb | Zjb ≥ θsparse}, j ≤ i for bra-ket-symmetry

    R' = Ria,jb - (Extia + Extjb)
    R'' = Rib,ja - (Extib + Extja)

    Near-field (R' ≤ 1 and/or R'' ≤ 1)
    if (ZiaZjbZiaZjb > ½θext or ZiaZjbZibZja > θext)

      Far-field (R' > 1 and/or R'' > 1)
      if (ZiaZjbZiaZjb / R'6 > ½θext or ZiaZjbZibZja / R'3R''3 > θext)

        → record FTidx of i, a, j, b indices
        if reached FTBatchSize → leave screening
      End jb
    End ia
  End ia
End ia
    
```

Note that for Coulomb-type products half the threshold θ_{ext} is used, since $E_J^{(2)}$ is weighted by a factor of 2 in the MP2 energy. The *FTidx* is built in such a way as to distinguish indices for which the exchange-type estimate was significant from those where only the Coulomb estimate was significant. Integrals are batched over the index i , such that for a given i all bras $(i\bar{a}|$ and their associated kets $|j\bar{b})$ are recorded. The maximum allowed batch size is set to half the total RAM, in order to allow sufficient memory for treatment of HTI and associated arrays. In view of the screening of HTI, an array of occurring (ij) -pairs is extracted from the *FTidx*.

Subsequent to this screening, an auxiliary array $Kidx$ is recorded and sorted so as to allow for direct assignment between integrals $(\underline{i}\bar{b}|\underline{j}\bar{a})$ and $(\underline{i}\bar{a}|\underline{j}\bar{b})$, in order to perform the contraction of $(\underline{i}\bar{a}|\underline{j}\bar{b})(\underline{i}\bar{b}|\underline{j}\bar{a})$ products in an efficient manner. This is done according to:

```

Loop  $iajb \in FTidx_K$ 
  record  $Tmp$  array of indices coded as  $iN^3 + bN^2 + jN + a$ ,
  along with  $Kidx$  array of running number  $k$ 
    
```

End

Perform QuickSort of Tmp to increasing order
with parallel reordering of $Kidx$

$FTidx_K$ here denotes the subset of indices from $FTidx$ for which the exchange estimate was significant. Due to the reordering of $Kidx$ in accord with the sorting of Tmp back to increasing $iN^3 + aN^2 + jN + b$ (i.e. the order in which the indices $iajb$ occur in $FTidx$), for the k -th FTI, the corresponding exchange integral $(\underline{i}\bar{b}|\underline{j}\bar{a})$ is simply that at position $Kidx[k]$.

5.2.4 Internal screening for HTI

Having obtained the list $FTidx$ of significant FTI, screening has to be performed for the HTI required in their formation. In contrast to the *external* screening for the fully transformed integrals, i.e. the final quantities that are required for the energy calculation, this screening deals with the intermediary quantities required for the FTI, and is thus termed *internal* screening. Here too, a batching scheme is necessary, and the maximum batch size is determined as 90% of RAM remaining after allocation of all arrays from the external screening, in order to allow for allocation of other minor arrays.

Ideally, the selection of required HTI would be performed by directly screening *down* from FTI to HTI according to

```

Loop  $ia \in FTidx$ 
  Loop  $jb \in FTidx_{ia}$ 
    Loop  $\sigma \in \{\sigma \mid \bar{L}_{\sigma b} \geq \theta_{sparse}\}$ 
      if  $Z_{ia}X_{j\sigma}\bar{L}_{\sigma b} \geq \theta_{int}$ : record  $j, \sigma$ 
    End  $\sigma$  and  $jb$ 

  Loop  $j\sigma \in \{j\sigma \mid \text{recorded}\}$ 
    Loop  $\lambda \in \{\lambda \mid \underline{L}_{\lambda j} \geq \theta_{sparse}\}$ 
      if  $Z_{ia}Q_{\lambda\sigma}\underline{L}_{\lambda j} \geq \theta_{int}$ : record  $i, a, \lambda, \sigma$ 
    End  $\lambda$ 
  End  $j\sigma$ 
End  $ia$ 
    
```

However, though this would in principle allow to preselect the minimum number of HTI for the particular choice of internal screening threshold, this approach leads to significant technical difficulties. Namely, given the similarity of the $(\underline{i\bar{a}}|\lambda\sigma)$ integrals for CDD-MP2 and the $(\underline{\mu\bar{\nu}}|\lambda\sigma)$ from the AO-MP2 implementation of Doser [6, 7], it is desirable to use the highly efficient HTI generation code previously developed in that context, after adaptation for our particular purpose, as discussed in the next section. This requires that the $HTidx$ be written as $(\lambda\sigma|\underline{i\bar{a}})$ and batched over bras, i.e. that all $|\underline{i\bar{a}}\rangle$ ket indices be recorded for a subset of the $(\lambda\sigma|$. This clearly cannot be obtained—at any rate not in an efficient, direct and dynamic way—using the above loop structure, where the untransformed indices form the inner loop.

An alternative approach therefore had to be chosen, were screening is performed *up* from all potentially significant HTI. In order to keep the number of integrals to a minimum, the transformation matrix \underline{L} is used to preselect the HTI not merely by their estimated value, but according to their weight in the transformation to TTI. Additionally to the Schwarz/R estimates described in section 5.1.3, the (ij) -pairs extracted from the $FTidx$ are used for coupling between left and right charge distributions. Screening is then performed with the following loop structure:

```

Loop  $\lambda\sigma \in ShellPairs, \quad \sigma \leq \lambda$ 
  Loop  $ia \in \{ia \mid Z_{ia} \geq \theta_{sparse}\}, \forall i \in FTidx$ 
    if  $Z_{ia}Q_{\lambda\sigma} \geq \theta_{int}$ :  $R'_{\lambda\sigma,ia} = R_{\lambda\sigma,ia} - (Ext_{\lambda\sigma} + Ext_{ia})$ 
  End  $ia$ 

  Loop  $j \in \{j \mid \underline{L}_{\lambda j} \geq \theta_{sparse}\}$ 
    Loop  $i \in (ij) - pairs$ 
      Loop  $a \in \{a \mid R'_{\lambda\sigma,ia} \text{ recorded}\}$ 

        Check for well-separatedness:  $R' > 1$ 

        if  $\frac{Z_{ia}Q_{\lambda\sigma}}{R'^2} |\underline{L}_{\lambda j} \bar{L}_{\sigma,max}| \geq \theta_{int}$ : record  $HTidx$  of  $\lambda, \sigma, i, a$ 

        if reached HTBatchSize:  $\rightarrow$  leave screening
      End  $a, i, j$ 
    End  $\lambda\sigma$ 

```

In spite of being a quadratic step, the determination of the reduced distances $R'_{\lambda\sigma,ia}$ before entering the actual screening loop has such a small prefactor that it is by far more efficient than repeated computation within the second loop. It should be obvious that R' is used in the estimate of a given HTI only in the case of well separatedness of the corresponding charge distributions.

5.2.5 Formation of the HTI

As mentioned in the previous section, the similarity of HTI in the CDD-MP2 and AO-MP2 methods called for the use of Doser’s very efficient implementation [6] of the first two transformations, after adaptations for our purpose. For technical reasons, in Doser’s implementation the transformations from UTI to HTI are performed on the kets, i.e. from $(\mu\nu|\lambda\sigma)$ to $(\mu\nu|\underline{\lambda}\bar{\sigma})$, since this allows for maximum efficiency. In the case of AO-MP2, the integrals thus formed can then be directly contracted to the MP2 energy according to equation (5.4), such that no further problems arise. For CDD-MP2, however, in view of the further two transformations that are required from HTI to FTI and which, for the reason just mentioned, are best performed on the kets, the HTI need to be stored as $(\underline{i}\bar{a}|\lambda\sigma)$. Yet, as described in the previous section, the *HTidx* has to be initially built for $(\lambda\sigma|\underline{i}\bar{a})$, in order to be compatible with the aforementioned ket-transformation scheme from UTI to HTI. In order to avoid an explicit transposition step from $(\lambda\sigma|\underline{i}\bar{a})$ to $(\underline{i}\bar{a}|\lambda\sigma)$, the code from Doser’s implementation had to be adapted for direct recording of the HTI in the bra-transformed order upon transformation. A corresponding indexing array *HTidxT* has thus to be written simultaneously to the recording. Further modifications of the code were of a rather marginal character and mostly consisted of adapting the array handling in the screening and transformation routines to our purpose. We thus give an overview of the algorithm’s structure, but refer to [7] for further details

Loop $\lambda\sigma \in HTidx$

Screening for required STI and UTI,
with formation of auxiliary arrays for transformation

1st transformation to $(\lambda\sigma|\underline{i}\nu)$

2nd transformation to $(\lambda\sigma|\underline{i}\bar{a})$, but direct recording as $(\underline{i}\bar{a}|\lambda\sigma)$
with simultaneous *HTidxT* formation

End $\lambda\sigma$

5.2.6 Last transformations to FTI

The last transformations of the HTI to the FTI are in principle similar to the first and second, although, in contrast to the screening *down* from HTI to STI and then UTI, the desired FTI and present HTI are known, such that different possibilities open up for the screening scheme of TTI intermediaries.

The ideal procedure, here again, would be the first one described above in section 5.2.4, with a corresponding recording of auxiliary arrays of indices required in the transformation, i.e. an

array $T3aux$ of j to which a given $|\lambda\sigma\rangle$ is transformed, according to

$$(\underline{i\bar{a}}|\underline{j\sigma}) = \sum_{\lambda \in HTidxT_{ia}} (\underline{i\bar{a}}|\lambda\sigma)\underline{L}_{\lambda j}, \quad \forall \sigma \in HTidxT_{ia}, \quad \forall j \in T3aux_{\lambda\sigma} \quad (5.51)$$

and an array $T4aux$ of σ occurring in the last transformation

$$(\underline{i\bar{a}}|\underline{j\bar{b}}) = \sum_{\sigma \in T4aux_{jb}} (\underline{i\bar{a}}|\underline{j\sigma})\bar{L}_{\sigma b}, \quad \forall jb \in FTidx_{ia}. \quad (5.52)$$

Initial tests, however, revealed that this causes the screening step to be too expensive due to the numerous allocations of individual auxiliary arrays, such that it was abandoned.

Given that the indices of HTI and FTI are known, an alternative to such a screening procedure would be simply to get the lists $J = \{j \mid j \in FTidx_{ia}\}$ and $\Sigma = \{\sigma \mid \sigma \in HTidxT_{ia}\}$ and to form the list of TTI as $\{j\sigma\} = J \otimes \Sigma$, i.e. all possible $j\sigma$ combinations. However, though this saves significant screening time, this approach artificially increases the numbers of TTI to such an extent that transformation time becomes dominant.

The procedure finally chosen is a compromise of both these approaches, where for the last transformation an array $T4aux$ is formed as defined previously, whereas the list of $\{j \mid j \in FTidx_{ia}\}$ is recorded in view of the third transformation. The overall structure of the algorithm for the transformation from HTI to FTI is then

Loop $ia \in FTidx$

Screening:

Loop $jb \in FTidx_{ia}$

Loop $\sigma \in \{\sigma \mid \bar{L}_{\sigma b} \geq \theta_{sparse}\}$

if $Z_{ia}X_{j\sigma}\bar{L}_{\sigma b} > \theta_{int}$: record σ to $T4aux_{jb}$

End σ, jb

Get list $J = \{j \mid j \in FTidx_{ia}\}$

3rd trafo:

Loop $\lambda\sigma \in HTidxT_{ia}$

$(\underline{i\bar{a}}|\underline{j\sigma}) = \sum_{j \in J} (\underline{i\bar{a}}|\lambda\sigma)\underline{L}_{\lambda j}$

End $\lambda\sigma$

4th trafo:

Loop $jb \in FTidx_{ia}$

$(\underline{i\bar{a}}|\underline{j\bar{b}}) = \sum_{\sigma \in T4aux_{jb}} (\underline{i\bar{a}}|\underline{j\sigma})\bar{L}_{\sigma b}$

End jb

End ia

5.2.7 Contraction to MP2 energy

With the assignment array $Kidx$ as defined in section 5.2.3, the contraction to the MP2 energy reduces to a straightforward matter:

```

Loop  $iajb \in FTidx$ 
  running number  $k$ 
   $(i\bar{a}|j\bar{b}) = FTidx[k]$ 
   $(i\bar{b}|j\bar{a}) = FTidx[Kidx[k]]$ 

   $E_J^{(2)} + = (i\bar{a}|j\bar{b})(i\bar{a}|j\bar{b}), \times 2$  if  $i \neq j$ 
   $E_K^{(2)} + = (i\bar{a}|j\bar{b})(i\bar{b}|j\bar{a}), \times 2$  if  $i \neq j$ 
End

```

$$E_0^{(2)} = 2E_J^{(2)} - E_K^{(2)}$$

5.3 Results and discussion

The following sections give an account of the performance of the CDD-MP2 algorithm, as tested on linear alkanes, chains of α -D-glucose (amylose), chains of DNA base pairs (adenine-thymine), as well as the S22 test set of Hobza [124]. Although it is originally provided with benchmark results to evaluate the accuracy of new methods devised for the calculation of interaction energies, the S22 set has here merely been used for its eclectic mix of systems in order to compare the accuracy of CDD-MP2 interaction energies with respect to canonical MO-MP2.

Calculations have been performed on a computer cluster equipped with Intel Xeon E5620 (2.40 GHz) processors. The code was compiled using the C++/Fortran compilers provided by the Intel Composer XE 2011.2.137 suite and its associated MKL library, the optimization level for compilation was set to O3. All calculations were further based on the frozen core approximation, a threshold of 10^{-10} for the determination of significant shell pairs, and the convergence criterion of all underlying HF calculations was chosen as 10^{-7} . The sparsity threshold below which matrix elements were discarded was set to $\theta_{sp} = 10^{-9}$, unless indicated otherwise. For the evaluation of transformed extents required in the Schwarz/R screening procedure, the untransformed extents occurring in definition (5.43) were determined using a truncation criterion of 10^{-1} [12], and the neglect threshold θ was chosen as 10^{-3} .

The accuracy provided by the number of quadrature points in the numerical evaluation of the Laplace transform integral has already been amply documented (see for instance [6, 7, 13]), such that this matter has not been systematically investigated in the present work, and unless noted otherwise, a fixed number of 5 quadrature points has been used as default setting in all cal-

culations. Further, all data depending on the quadrature points, such as matrix sparsities and integral numbers, has been taken from the first quadrature point, since it is usually the "worst case".

The presentation begins with the sparsities of the Cholesky-factors, screening matrices and pseudo-densities and then proceeds with integral numbers, accuracies, along with timings, after which the gained insights are discussed and conclusions are drawn.

5.3.1 Matrix sparsities

The first issue that needs to be addressed for our present objective of reducing the scaling exponent of MP2 is whether the (exponentially weighted) orbitals obtained through Cholesky decomposition of the pseudo-densities matrices are actually local and thus have the potential for application of integral screening. In order to demonstrate that this is so, table 5.1 gives the numbers of significant elements of the scaled (see section 5.2.1) occupied and virtual Cholesky factors $\underline{\mathbf{L}}$ and $\overline{\mathbf{L}}$, as well as the screening matrix \mathbf{Z} defined in section 5.1.3, for a variety of test systems. We see that for linear alkanes or amylose chains with the basis sets 6-31G* or SVP, the numbers of significant elements in the Cholesky factor matrices rapidly grow linearly with the system size, which confirms that, as discussed in section 5.1.2, they represent energy-weighted localized molecular orbitals. Accordingly, the screening matrix \mathbf{Z} displays the same behaviour and it can thus be anticipated that in conjunction with the distance-dependent estimates introduced in section 5.1.3, the number of FTI products that will be pre-selected for computation of the MP2 energy should rapidly display linear scaling with increasing system size. However, we also notice that for the larger cc-pVTZ basis, as well as for the more bulky DNA systems even with the smaller basis sets, the occupied Cholesky orbitals do not quite display the same locality as the virtuals, and linear scaling of the significant matrix elements is not achieved in these cases.

For comparison, table 5.2 shows the corresponding numbers of significant elements in the pseudo-density matrices that served as input for the Cholesky decomposition. In essence, the exact same trends are observed as for the sparsity of the Cholesky factors. The main difference merely lies in the fact that in the problematic cases discussed above, it is the virtual rather than the occupied pseudo-density that does not scale linearly. This is to be attributed to the aforementioned "normalization" of the virtual Cholesky factor $\overline{\mathbf{L}}$ by the maximum element \overline{L}_{\max} (section 5.2.1), such that the magnitudes of elements in the scaled $\overline{L}_{\max}\overline{\mathbf{L}}$ are shifted up. Aside from this, however, we notice that the Cholesky decomposition indeed preserves the sparsity of the input matrix, and in cases where the pseudo-density is sufficiently local, so are the obtained Cholesky orbitals.

5.3.2 Integral numbers and accuracy

Following the validation of principle of the Cholesky decomposition approach that was given in the previous section, it has to be established whether the integral screening procedures described in sections 5.2.3 and 5.2.4 allow to reduce the numbers of FTI as well as integrals required in the transformation steps to linear, while retaining numerical accuracy.

To begin with, tables 5.3 and 5.4 display the numbers of preselected FTI products and their scaling exponents with respect to the previous molecule, along with the numbers of (ij) -pairs which are used as described in section 5.2.4. The tables further display the numbers of UTI, STI and HTI that occur in the consecutive transformations to the FTI. The numbers of TTI are not given, because due to the HTI being formed subset-wise, the third transformation cannot be fully carried out to obtain explicitly formed $(\underline{i}\bar{a}|\underline{j}\bar{b})$ integrals, but can only be performed partially for each HTI batch. The FTI are thus formed successively according to

$$(\underline{i}\bar{a}|\underline{j}\bar{b}) = \sum_B^{N_{Batch}} \sum_{\lambda, \sigma \in B} (\underline{i}\bar{a}|\lambda\sigma) \underline{L}_{\lambda j} \bar{L}_{\sigma b} = \sum_B^{N_{Batch}} \sum_{\sigma \in B} (\underline{i}\bar{a}|\underline{j}\sigma)' \bar{L}_{\sigma b}, \quad (5.53)$$

where the prime on $(\underline{i}\bar{a}|\underline{j}\sigma)'$ serves to indicate that the TTI are actually only intermediary summation entities.

Now, as can be seen from table 5.3, within the series of linear alkanes the number of preselected FTI products and their associated numbers of (ij) -pairs rapidly achieve linear scaling. The corresponding partially transformed integrals of the transformation steps from UTI to FTI ultimately also scale linearly with system size, though significantly later than the FTI products. Furthermore, as the integral transformation degree decreases, i.e. more indices run over the AO basis rather than the Cholesky orbital basis, the integral numbers increase tremendously. For instance, the numbers of UTI are two or three orders of magnitude larger than the FTI numbers. Considering the less idealized test series consisting of amylose chains, whose integral data is displayed in table 5.4, the observed trends are similar for the partially transformed integrals, but differ substantially for the UTI. Though the fully transformed integral products as well as the formed HTI and STI asymptotically scale linearly for the small 6-31G* and SVP basis sets, the untransformed integrals do not achieve better than quadratic scaling, and their scaling exponent eventually increases again. This seemingly anomalous behaviour of the UTI numbers is to be ascribed to the batching scheme over FTI, which is described in section 5.2.3. Namely, according to the transformation

$$(\underline{i}\bar{a}|\underline{j}\bar{b}) = \sum_{\mu\nu\lambda\sigma} \underline{L}_{\mu i} \bar{L}_{\nu a} (\mu\nu|\lambda\sigma) \underline{L}_{\lambda j} \bar{L}_{\sigma b} \quad (5.54)$$

a given integral $(\mu\nu|\lambda\sigma)$ enters into all $(\underline{i}\bar{a}|\underline{j}\bar{b})$, or within a screening scheme such as in the present case, at least into a subset of the FTI. Therefore, as soon as the batching of FTI sets

in, repeated computation of UTI that are required in the transformation to FTI belonging to different batches becomes inevitable within our integral direct approach. Though the numbers of UTI formed for a given FTI batch should eventually attain linear scaling in the asymptotic regime, the prefactor of their total number is not the same between a system for which all FTI can all be treated as a whole and a system requiring subdivision into two or more FTI batches. Evaluation of the scaling exponent of the UTI number according to equation (2.62), which presupposes a *constant* prefactor, thus leads to the observed inconsistency. In such a case involving non-constant prefactors, equation (2.62) should be replaced by

$$\alpha = \frac{\ln(T_b/T_a) - \ln(K_b/K_a)}{\ln(N_b/N_a)}, \quad (5.55)$$

such that the evaluation of the true scaling exponent α is impossible as it requires knowledge of the unknown prefactors K_a and K_b . Thus, for lack of any better data, the evaluation of all scaling exponents here and in following cases shall yet be based on equation (2.62), bearing in mind the limited meaningfulness of the data thus obtained.

All in all, we are forced to the conclusion that, though linear scaling of the number of UTI is in principle possible, as could be shown for the case of linear alkanes, it is hardly achievable for realistic systems, where multiple FTI batches are required. For large basis sets, the situation is further complicated by the fact that in these cases, as discussed in the previous section, the pseudo-densities, and as a consequence the Cholesky orbitals, are only poorly localized. The numerous successive screening steps that are necessary in the transformations to the FTI thus lose efficiency in reducing the numbers of transformation intermediaries, and as can be seen from the integral data for Amylose with cc-pVTZ, all integrals scale very unfavourably for the systems that could be investigated.

As concerns the accuracy of the CDD-MP2 approach, tables 5.5, 5.6 and 5.7 display the errors in MP2 total energies, as well as resolved into the Coulomb and exchange components $E_J^{(2)}$ and $E_K^{(2)}$, for different numerical settings and basis sets. For the 6-31G* basis, the results given in table 5.5 show that errors are already within very reasonable bounds for the loose thresholds. Merely the large amylose systems and DNA pairs call for the tighter thresholds to obtain sub- m Hartree accuracy. Also, because the errors in the Coulomb-type and exchange-type component have the same sign in all these cases, they partially compensate each other upon combination of both terms to total energies, due to the exchange part entering with a different sign than its Coulomb counterpart (cf. equation (3.25)). We further notice that the error in $E_J^{(2)}$ is larger than the corresponding error in $E_K^{(2)}$, such that the total error is dominated by the former. Considering the results for the larger basis sets SVP and cc-pVTZ, given in tables 5.6 and 5.7, respectively, different trends become apparent for the errors in the Coulomb-like and exchange-like term. For the latter, not only are the deviations reasonable for both threshold settings, but they are very well-behaved in the sense that they increase almost linearly with system size, and tightening of

the screening thresholds by an order of magnitude also reduces the error by roughly an order of magnitude. For the Coulomb-type component, however, the errors are not quite so well-behaved, and we observe a rather unsystematic behaviour, which may be explained as follows. On the one hand, as is to be seen in equation (5.12), the CDD-MP2 Coulomb type energy should theoretically represent a lower bound to the exact $E_J^{(2)}$. Namely, since only positive summands of the form $(\underline{i\bar{a}}|\underline{j\bar{b}})^2$ are added up, neglect of small elements in the summation causes

$$E_{J,\text{CDD}}^{(2)} < E_{J,\text{exact}}^{(2)}, \quad (5.56)$$

such that one would expect the error in the Coulomb part to be always negative. On the other hand, however, the FTI used in the contraction to the energy are not exact, being formed themselves through an internal screening procedure that discards small contributions in the consecutive transformations from UTI to FTI. They are thus affected by a certain numerical error and we may write the FTI matrix (see equation (5.11)) as

$$\mathbf{M} = \mathbf{M}_{\text{exact}} + \mathbf{\Delta}, \quad (5.57)$$

where both the FTI matrix and the matrix $\mathbf{\Delta}$ of errors in the FTI are symmetric. Now, upon calculation of the Coulomb-type energy according to

$$E_J^{(2)} = 2\text{Tr}(\mathbf{M}^2) = 2\text{Tr}(\mathbf{M}_{\text{exact}}^2) + 4\text{Tr}(\mathbf{M}_{\text{exact}}\mathbf{\Delta}) + 2\text{Tr}(\mathbf{\Delta}^2) \quad (5.58)$$

a strictly positive error $\text{Tr}(\mathbf{\Delta}^2) = \|\mathbf{\Delta}\|_F^2 \geq 0$ ensues, the squared Frobenius norm of the error matrix. The overall result of these two antagonistic effects is a high sensitivity of the Coulomb-like energy to the combination of screening thresholds. The number of FTI discarded by the external screening and the precision with which the preselected FTI are formed appear to be difficult to be balanced against each other, such that the results are numerically not very stable. The first signs of this may be seen in table 5.6 for the case of amylose systems with the SVP basis set. Going over from Amylose₂ to Amylose₄, for the looser threshold combination (left column), we notice that the error in the Coulomb energy becomes smaller, which hints at the fact that the additive error $\|\mathbf{\Delta}\|_F^2$ starts to gain weight in compensating for the negative error arising in $\text{Tr}(\mathbf{M}^2)$ due to FTI neglect. For the very same systems with lower screening thresholds (right column), the error in $E_J^{(2)}$ actually becomes positive, indicating that in spite of a tighter internal screening, the additive error now even prevails over the negative error. In this particular case, the total accuracy provided by lower thresholds is actually poorer. This overcompensation of the negative by the additive error component is further observed in table 5.7 for all systems with the cc-pVTZ basis. In contrast to the well-behaved cases with the 6-31G* basis, where the errors of $E_J^{(2)}$ and $E_K^{(2)}$ have same signs and thus partially cancel each other upon formation of total energies, the Coulomb and exchange errors of different signs here mutually add up, resulting in relatively poor overall accuracy of CDD-MP2 total absolute energies.

Yet, the situation is improved for relative energies, which are more relevant for chemistry anyway, as may be seen in table 5.8, where the MO-MP2 interaction energies of complexes from the S22 test set [124] are listed together with the errors of the CDD-MP2 method with respect to these reference values, along with the mean and standard deviation of the errors. For the data set pertaining to the SVP basis, we notice that the loosest threshold combination already provides chemical accuracy. Even for the largest complexes of the test set, such as 5, 6, 7 or 15, the deviation from the MO-MP2 reference interaction energy is well below 1 kcal/mol. Going over to the next tighter combination, the errors become practically negligible, especially when considering that they are smaller than the inherent error of MP2 as compared to high-end methods like coupled cluster. For the case of the cc-pVTZ basis, the loose thresholds are evidently inappropriate as is indicated by the large average error and standard deviation. We notice in particular that the errors in the results are unacceptable for systems 5, 6, 7, 20 and especially 21. Tightening of the thresholds by two orders of magnitude, however, substantially improves the results and chemical accuracy is achieved for all systems except 5 and 7. For both basis sets, we further notice from the last columns of both data sets, respectively, that upon increasing the number of quadrature points in the Laplace integral evaluation, and imposing very tight screening conditions, the numerical error can for all practical purposes be considered as null.

We may thus draw the conclusion that the large errors affecting the absolute CDD-MP2 energies due to the numerical imbalance of the Coulomb-like energy component largely cancel each other out in the formation of energy differences. This is actually not so surprising, since it is to be expected that the numerical error due to the screening scheme in two systems of similar dimensions should be of similar magnitude and particularly of same sign, given that the additive error in $E_J^{(2)}$ showed the tendency to dominate over its negative counterpart. Upon formation of energy differences, it thus seems logical that the overall errors should mutually cancel each other to a large extent. In effect, this behaviour is analogous, though not as systematic, to situations where RI-techniques are applied. Even though the deviations of RI-approximated absolute energies are quite substantial, due to the incompleteness of the auxiliary basis (see also section 6.4.2), experience has revealed that they compensate each other in relative energies, such that the application of RI is yet meaningful.

5.3.3 Timings

Finally, the central issue of calculation times shall be addressed. The discussion begins with a comparison of total wall times for CDD-MP2 and conventional MO-MP2, the data of which is given in table 5.9, in order to outrightly point out the problems that arise. This is followed by a closer analysis of timings of the individual steps of the CDD-MP2 algorithm.

Now, as may be seen in table 5.9, for the homologous series of linear alkanes near-linear scaling

is achieved with both the 6-31G* and SVP basis sets, though admittedly at a relatively late point, considering that alkanes are idealized one-dimensional systems. Accordingly, the cross-overs with MO-MP2 are situated at relatively large molecular sizes, namely between $C_{40}H_{82}$ and $C_{80}H_{162}$ for 6-31G*, whereas it is barely below $C_{80}H_{162}$ for the SVP basis.

For alkanes with the larger cc-pVTZ basis and the somewhat bulkier amylose chains with all three bases, however, the situation is far less favourable. In both cases the computational scaling of CDD-MP2 does not get better than quadratic for the system sizes that could be investigated. In the former case of alkanes with cc-pVTZ, the cross-over point again lies between $C_{40}H_{82}$ and $C_{80}H_{162}$, but due to the still high scaling exponent and prefactor the gains are rather meager, as the overall calculation time for $C_{80}H_{162}$ is merely halved as compared to MO-MP2. This relatively small gain is further mitigated by the fact that because of memory problems due to the tremendously increasing numbers of integrals with tight numerical thresholds, these calculations had to be performed using thresholds (7, 7, 9) (see definition in table 5.9) which, though totally acceptable for 6-31G* and SVP, have been discussed in section 5.3.2 as producing quite large errors in the absolute energies with cc-pVTZ.

For amylose chains, as already mentioned, only quadratic scaling is achieved with the 6-31G* and SVP basis sets and in both cases the cross-over to conventional MO-MP2 is located between Amylose₈ and Amylose₁₆. Although it has to be taken into consideration that the scaling that was determined for the MO-MP2 references is smaller than the expected fifth order, such that the timings of the large amylose chains that could not be computed explicitly are extrapolated rather generously, the advantage of CDD-MP2 is again relatively small. With 6-31G*, the CDD-MP2 calculation time of Amylose₁₆ is roughly half as large as for the conventional method, and only slightly improved with the SVP basis. For cc-pVTZ, the high scaling exponent of 3.6 results in an explosion of computation time and the cross-over to MO-MP2 could not be reached.

At this point, computation times of DNA chains as representatives of model systems which might be encountered in realistic situations would be of further interest. However, as is already hinted at by the poor locality of the Cholesky orbitals for DNA systems (see section 5.3.1), the problems that arise in the aforementioned test systems of alkanes and amylose chains are here even worse. Though calculations on DNA₁ and DNA₂ were performed in order to determine the method's numerical accuracy, the rather delocalized nature of the Cholesky orbitals and small extent of these systems logically do not allow for the distance based screening procedure to discard enough integrals so as to reduce the computational scaling. Due to the ensuing explosion of memory requirements and computation times, DNA₄, DNA₈ and higher elements of the sequence of DNA base pairs, where the effects of integral screening on computational scaling could be expected to reveal themselves, become utterly intractable even with the small 6-31G* and SVP bases. It is therefore not possible to discuss the trends that would arise within this homologous series, such that the listing of calculation times of merely the first two elements

is not deemed to be of particular use either. Suffice it to say that in these two cases for which calculations could be performed, CDD-MP2 performs significantly slower than canonical MP2, and that the cross-over point cannot be reached with the present implementation.

For a thorougher analysis of the situation, tables 5.10 to 5.15 display timings of the individual steps of the CDD-MP2 algorithm for alkanes and amylose chains with the different basis sets. At the outset, it has to be remarked that due to the roundoff of wall times to full seconds, very fast steps such as the Cholesky decomposition or contraction to the MP2 energy display irregular fluctuations that only level out when this roundoff becomes negligible in calculations on larger systems. Furthermore, though in the current implementation the Cholesky decomposition and external screening have complexities of $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$, respectively, their contributions to the total calculation times are negligible in all systems that could be investigated, such that no optimization of these steps was deemed necessary so far.

For linear alkanes, we notice that all steps except the aforementioned decomposition and external screening asymptotically achieve near-linear scaling. In the case of 6-31G* (table 5.10), the internal screening for UTI and their evaluation clearly dominate the calculation time. These are followed by the first two transformations, whereas the portion of the remaining steps in the calculation time is relatively small. As is to be expected from the previously discussed behaviour of integral numbers within the sequence of transformations from UTI to FTI, transformation times decrease gradually from the first to the last transformation. Yet, we notice that in contrast to the first three transformations, the last one does not achieve fully linear scaling. Going over to the slightly larger SVP basis (table 5.11), the trends in the individual scaling exponents remain largely the same, though it is the integral screening for UTI that here becomes the most time-demanding step due to its not fully linear scaling. Here again, all transformations nearly achieve linear scaling, except for the fourth transformation (see explanation below) which, however, still represents less than 2% of the total time.

The corresponding data for the series of amylose chains with the 6-31G* and SVP bases is given in tables 5.13 and 5.14, respectively. For both cases, in accordance with the previously discussed overall scalings given in table 5.9, all steps representing significant portions of the total time achieve only quadratic asymptotic scaling. With the 6-31G* basis, merely the first and third transformation attain distinctly sub-quadratic scaling. Moreover, in contrast to the previous trend of decreasing transformation time, the last transformation here experiences a dramatic surge of execution time and exhibits a nearly cubic scaling exponent. Together with the second internal screening step for TTI intermediaries, these steps alone make up for nearly 30% of the total time. Similarly to the case of multiple UTI formations discussed previously, this behaviour is caused by the batching scheme over HTI. Namely, the screening procedure described in section 5.2.6, which is required to perform the last transformations as efficiently as possible, has to be gone through for every batch. Therefore, although this screening is implemented in such a

way that it should scale linearly for sparse enough screening matrices, the prefactor of its total execution time is not constant but depends on the number of HTI batches that are required. Accordingly, as has already been discussed in section 5.3.2 in the context of integral numbers, the true scaling exponents cannot be determined, and the observed inconsistencies are due to the use of equation 2.62 for their evaluation.

The fact that HTI are handled one subset at a time has further repercussions on the last transformation. Namely, for a given bra $(\underline{i}\bar{a}|$ with its associated set of $(\underline{i}\bar{a}|\lambda\sigma)$ integrals in a given HTI batch B , the set $J = \{j \mid j \in FTidx_{ia}\}$ for which transformations

$$(\underline{i}\bar{a}|\underline{j}\sigma)' = \sum_{\lambda \in B} (\underline{i}\bar{a}|\lambda\sigma) \underline{L}_{\lambda j}$$

are performed is unique. Recalling equation (5.53), $(\underline{i}\bar{a}|\underline{j}\sigma)'$ represents an intermediary summation entity rather than a fully formed TTI. Now, due to J being unique, for a linearly scaling total number of HTI and sparse enough occupied Cholesky factor \underline{L} , the third transformation, even performed successively, scales linearly, as is observed. However, some of the σ indices do occur repeatedly in different HTI batches, such that when more than one HTI batch is needed, some combinations of $(\underline{i}\bar{a}|\underline{j}\sigma)'$ -indices arise multiple times, and thus need to be addressed repeatedly in the final transformation

$$(\underline{i}\bar{a}|\underline{j}\bar{b}) = \sum_{\sigma \in T4aux_{jb}} (\underline{i}\bar{a}|\underline{j}\sigma)' \bar{L}_{\sigma b}$$

which is performed successively for each HTI batch. In effect, the prefactor of this transformation step also varies with the number of HTI batches, such that equation (2.62) is here again only of limited validity, and we accordingly observe an unsystematic, increasing behaviour of the scaling exponents. Going over to the SVP basis (table 5.14), the same shift of the computational time of individual steps is observed as in the case of alkanes, and the screening for UTI assumes a dominant role in the total computation time, followed by the UTI evaluation. Similarly to the 6-31G* basis, the first and third transformations here again scale more efficiently than the second and fourth, although the differences are less pronounced, especially for the two last transformations. In particular, we note that the fourth transformation here takes considerably less time than with 6-31G*.

For the cc-pVTZ basis, timings of individual steps for linear alkanes as well as amylose chains (tables 5.12 and 5.15) exhibit the same tendencies that have just been discussed. In both cases, the high scaling exponents of the UTI screening, UTI evaluation and the first two transformations, which by themselves constitute up to 75% of the total computational time, account very well for the overall scaling exponent. Also, the second internal screening for TTI intermediaries and last transformation again display abnormally high scaling exponents which, as discussed in the previous paragraph, are due to the dependence of these steps' prefactors on the numbers of HTI batches.

5.3.4 Conclusion

The overall picture that emerges from the very mitigated results presented in the foregoing sections is that, in spite of promising features, the CDD-MP2 method does not live up to expectations and—at least in its current implementation—is clearly not fit for MP2 calculations of large molecules. Though it could be shown for the case of linear alkanes as idealized one-dimensional systems and with small basis sets that the approach is in principle valid to achieve asymptotic linear scaling of the computational effort, the applicability to systems of realistic size with large bases is severely hampered by the technical problems that arise. In spite of the fact that the numbers of preselected FTI in general rapidly scale efficiently, such that the time required for contraction to the energy in itself becomes negligible, the successive integral preselection steps required for the transformation intermediaries from UTI to FTI lead to long "screening tails" and to a late onset of overall linear scaling even for the one-dimensional case. For compact systems such as amylose chains or DNA base pairs, and particularly in the case of large basis sets, the poorer locality of the pseudo-density matrices, and accordingly of the Cholesky orbitals, further adds to the difficulties and in these cases the overwhelming computational effort associated with the internal screenings and the quadruple integral transformation makes it impossible to reach the linear scaling regime in practice. Due to the associated large computational prefactor for large systems and basis sets, the performance gains as compared to conventional MO-MP2 are relatively small, and in some cases cross-over points could not even be reached.

Apart from these performance problems, the approach was discovered to have the additional flaw of being numerically unstable for the formation of the Coulomb-type component $E_J^{(2)}$ of the MP2 energy. Namely, whereas the negative error component of $E_J^{(2)}$ due to the discarding of negligible FTI products seems to correlate well with the choice of external screening threshold, the error induced by the internal screening in the finally formed FTI propagates tremendously over the four consecutive transformations, resulting in a large overall deviation in $E_J^{(2)}$ and thus insufficient accuracy of total absolute energies with practicable screening thresholds. Even though the overall accuracy is in principle fully tunable through the numerical settings, the minimization of this additive error component would call for internal screening thresholds and enlargements of the distribution extents that utterly annihilate any prospects of reducing the computational effort and obtaining performance improvements over conventional implementations. Despite the fact that this error was shown to partially cancel upon forming energy differences, such that relative energies on the average display very reasonable accuracy even with relatively loose screening thresholds, this trait remains problematic, be it only because it severely limits the reliability of CDD-MP2 absolute energies for use in basis set limit extrapolation schemes.

As to possible approaches to tackle the encountered problems, one might be inclined to point

out that according to section 5.1.2, the transformation matrices used within CDD-MP2 need not necessarily be obtained through Cholesky decomposition of the pseudo-density matrices, but can also be formed from any other set of localized molecular orbitals. This might result in sparser transformation matrices that would allow for more efficient screening. Additional improvements of the performance can also be obtained from a parallel implementation, which has already been undertaken but will only be concluded upon combination with a parallelized version of the UTI evaluation algorithm, which will be the object of future work. Furthermore, some of the arising difficulties are directly related to the integral batching that is inevitable within the chosen integral direct approach, such that an integral indirect implementation would circumvent some of these problems.

However, in light of the numerical problems inherent to the approach, the potential improvements brought by alternative transformation matrices obtained from other local molecular orbitals are questionable, since a reduction of integral numbers through the use of sparser screening matrices can only come at the cost of further losses in numerical accuracy. Also, though the parallelized CDD-MP2 algorithm can be anticipated to be more efficient, it would merely reduce the overall prefactor and hardly touch at the deeper rooted problems due to the observed high scaling exponents with large systems. Finally, due to the large numbers of UTI and tremendous effort involved in the quadruple integral transformation, an integral indirect algorithm only makes sense in conjunction with a decomposition of the integral matrices, such that both issues of storage requirements and transformation cost are solved. The logical consequence of this is the RI-CDD-MP2 method that is discussed in chapter 6. Also, a promising alternative to the RI integral decomposition approach is the method specific Cholesky decomposition investigated in chapter 7.

Molecule	6-31G*						SVP						cc-pVTZ					
	$\underline{\mathbf{L}}$	Sc.	$\overline{\mathbf{L}}$	Sc.	\mathbf{Z}	Sc.	$\underline{\mathbf{L}}$	Sc.	$\overline{\mathbf{L}}$	Sc.	\mathbf{Z}	Sc.	$\underline{\mathbf{L}}$	Sc.	$\overline{\mathbf{L}}$	Sc.	\mathbf{Z}	Sc.
C ₁₀ H ₂₂	67	-	197	-	47	-	76	-	249	-	65	-	183	-	1741	-	176	-
C ₂₀ H ₄₂	242	1.8	717	2.0	185	2.0	271	2.0	893	2.0	249	2.0	701	1.9	6389	1.9	675	2.0
C ₄₀ H ₈₂	915	1.9	2650	2.0	655	1.8	1022	2.0	3230	1.9	934	1.9	2701	2.0	18808	1.6	2645	2.0
C ₈₀ H ₁₆₂	3171	1.8	7855	1.6	1678	1.4	3908	2.0	9262	1.6	2959	1.7	10076	1.9	47672	1.4	10177	2.0
C ₁₆₀ H ₃₂₂	7903	1.3	18511	1.2	3724	1.2	13513	1.8	21795	1.2	7150	1.3	30415	1.6	111118	1.2	30927	1.6
Amylose ₁	67	-	197	-	56	-	76	-	249	-	65	-	184	-	1385	-	173	-
Amylose ₂	242	1.8	717	2.0	203	2.0	271	2.0	893	2.0	232	2.0	656	2.0	4938	2.0	617	2.0
Amylose ₄	915	1.9	2650	2.0	768	2.0	1022	2.0	3230	1.9	875	2.0	2472	2.0	17828	1.9	2325	2.0
Amylose ₈	3171	1.8	7855	1.6	2751	1.9	3908	2.0	9262	1.6	3377	2.0	9564	2.0	53258	1.6	9018	2.0
Amylose ₁₆	7903	1.3	18511	1.2	7025	1.4	13513	1.8	21795	1.2	11453	1.8	37137	2.0	130154	1.3	35214	2.0
Amylose ₃₂	17375	1.1	39828	1.1	15570	1.2	34625	1.4	46788	1.1	28824	1.3	-	-	-	-	-	-
DNA ₁	505	-	1561	-	427	-	549	-	1791	-	470	-	1312	-	9777	-	1233	-
DNA ₂	2368	2.0	7118	2.0	1982	2.0	2532	2.0	7845	2.0	2147	2.0	6002	2.0	42145	1.9	5616	2.0
DNA ₄	10201	1.9	29123	1.9	8509	2.0	10834	2.0	31039	1.9	9142	2.0	25587	2.0	168312	1.9	23891	2.0
DNA ₈	39566	1.9	92606	1.6	33840	1.9	44085	2.0	96923	1.6	37355	2.0	-	-	-	-	-	-

Table 5.1: Numbers [in 10^2] of elements of magnitude $\geq 10^{-9}$ in the scaled (cf. section 5.2.1) Cholesky factors $\underline{\mathbf{L}}$ and $\overline{\mathbf{L}}$ and screening matrix \mathbf{Z} , along with scaling exponents with respect to previous molecule, for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the indicated basis sets; First quadrature point.

Molecule	6-31G*				SVP				cc-pVTZ			
	$\underline{\mathbf{P}}$	Sc.	$\overline{\mathbf{P}}$	Sc.	$\underline{\mathbf{P}}$	Sc.	$\overline{\mathbf{P}}$	Sc.	$\underline{\mathbf{P}}$	Sc.	$\overline{\mathbf{P}}$	Sc.
C ₁₀ H ₂₂	285	-	285	-	535	-	547	-	3318	-	3642	-
C ₂₀ H ₄₂	1099	2.0	1116	2.0	2034	2.0	2105	2.0	11772	1.9	13822	2.0
C ₄₀ H ₈₂	3403	1.6	4133	1.9	6510	1.7	8014	2.0	33846	1.6	53035	2.0
C ₈₀ H ₁₆₂	8070	1.3	10983	1.4	16285	1.3	24228	1.6	84609	1.3	184130	1.8
C ₁₆₀ H ₃₂₂	17405	1.1	24687	1.2	35679	1.1	57551	1.3	195605	1.2	501876	1.5
Amylose ₁	416	-	416	-	520	-	520	-	2787	-	2788	-
Amylose ₂	1513	2.0	1513	2.0	1866	2.0	1866	2.0	9942	2.0	9960	2.0
Amylose ₄	5717	2.0	5761	2.0	6958	2.0	7056	2.0	36624	2.0	37558	2.0
Amylose ₈	17051	1.6	21147	1.9	21934	1.7	26440	1.9	116683	1.7	145262	2.0
Amylose ₁₆	39808	1.2	55413	1.4	57023	1.4	79588	1.6	312060	1.4	518443	1.9
Amylose ₃₂	85341	1.1	123975	1.2	128743	1.2	192148	1.3	-	-	-	-
DNA ₁	3328	-	3352	-	3830	-	3906	-	19839	-	20392	-
DNA ₂	15460	2.0	15674	2.0	17260	2.0	17739	2.0	87546	2.0	90962	2.0
DNA ₄	65138	2.0	67480	2.0	71614	2.0	75369	2.0	359407	2.0	383398	2.0
DNA ₈	208597	1.6	270171	2.0	242471	1.7	303476	2.0	-	-	-	-

Table 5.2: Numbers [in 10^2] of elements of magnitude $\geq 10^{-9}$ in the pseudo-density matrices $\underline{\mathbf{P}}$ and $\overline{\mathbf{P}}$, along with scaling exponents with respect to previous molecule, for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the indicated basis sets; First quadrature point.

Basis	Molecule	N	UTI	Sc.	STI	Sc.	HTI	Sc.	FTI	Sc.	(ij)	Sc.
6-31G*	C ₁₀ H ₂₂	194	235	-	61	-	47	-	6	-	496	-
	C ₂₀ H ₄₂	384	1095	2.3	401	2.8	298	2.7	20	1.8	1779	1.9
	C ₄₀ H ₈₂	764	4174	1.9	1552	2.0	1025	1.8	47	1.3	4558	1.4
	C ₈₀ H ₁₆₂	1524	11120	1.4	3980	1.4	2507	1.3	103	1.1	10091	1.2
	C ₁₆₀ H ₃₂₂	3044	24928	1.2	8829	1.2	5467	1.1	215	1.1	21167	1.1
SVP	C ₁₀ H ₂₂	250	549	-	120	-	96	-	13	-	496	-
	C ₂₀ H ₄₂	490	2505	2.3	783	2.8	559	2.6	49	1.9	1852	2.0
	C ₄₀ H ₈₂	970	9528	2.0	3180	2.1	1932	1.8	122	1.3	5414	1.6
	C ₈₀ H ₁₆₂	1930	26212	1.5	8341	1.4	4759	1.3	269	1.1	12467	1.2
	C ₁₆₀ H ₃₂₂	3850	59702	1.2	18688	1.2	10396	1.1	561	1.1	26563	1.1
cc-pVTZ	C ₁₀ H ₂₂	608	14483	-	1447	-	1229	-	72	-	496	-
	C ₂₀ H ₄₂	1188	63258	2.2	9297	2.8	6613	2.5	226	1.7	1847	2.0
	C ₄₀ H ₈₂	2348	230774	1.9	35949	2.0	21122	1.7	530	1.2	5119	1.5
	C ₈₀ H ₁₆₂	4668	614045	1.4	92185	1.4	50940	1.3	1133	1.1	11624	1.2

Table 5.3: Numbers of integrals [in 10^6] for all partial transformations and numbers of (ij) -pairs, along with scaling exponents with respect to previous molecule, for linear alkanes with the indicated basis sets; First quadrature point; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{sp} = 10^{-9}$.

Basis	Molecule	N	UTI	Sc.	STI	Sc.	HTI	Sc.	FTI	Sc.	(ij)	Sc.
6-31G*	Amylose ₁	204	498	-	115	-	87	-	15	-	666	-
	Amylose ₂	389	2771	2.7	914	3.2	693	3.2	110	3.1	2346	2.0
	Amylose ₄	759	12278	2.2	5842	2.8	4271	2.7	374	1.8	8758	2.0
	Amylose ₈	1499	47627	2.0	24524	2.1	16041	1.9	868	1.2	25539	1.6
	Amylose ₁₆ ^a	2979	250208	2.4	65294	1.4	40434	1.3	1875	1.1	59188	1.2
SVP	Amylose ₁	228	640	-	142	-	110	-	19	-	666	-
	Amylose ₂	432	3468	2.6	1059	3.1	822	3.1	131	3.0	2346	2.0
	Amylose ₄	840	15205	2.2	6049	2.6	4511	2.6	454	1.9	8771	2.0
	Amylose ₈	1656	56974	1.9	23169	2.0	15658	1.8	1068	1.3	25974	1.6
	Amylose ₁₆ ^a	3288	280618	2.3	59713	1.4	38822	1.3	2322	1.1	60571	1.2
cc-pVTZ	Amylose ₁	528	13781	-	1501	-	1292	-	123	-	666	-
	Amylose ₂	998	70349	2.6	11050	3.1	9102	3.1	850	3.0	2346	2.0
	Amylose ₄ ^a	1938	818125	3.7	67373	2.7	49820	2.6	2948	1.9	8778	2.0

^a Onset of FTI batching

Table 5.4: Numbers of integrals [in 10^6] for all partial transformations and numbers of (ij) -pairs, along with scaling exponents with respect to previous molecule, for α -D-glucose chains (amylose) with the indicated basis sets; First quadrature point; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{sp} = 10^{-9}$.

5.3. RESULTS AND DISCUSSION

Molecule	N	ΔE (7, 7, 9)			ΔE (8, 8, 9)		
		TOTAL	Coulomb	Exchange	TOTAL	Coulomb	Exchange
C ₁₀ H ₂₂	194	-45.91	-80.92	-35.01	-13.06	-16.18	-3.12
C ₂₀ H ₄₂	384	-171.19	-270.77	-99.57	-62.05	-72.66	-10.61
C ₄₀ H ₈₂	764	-427.78	-657.54	-229.76	-163.52	-189.66	-26.15
C ₈₀ H ₁₆₂	1524	-934.66	-1423.46	-488.79	-364.75	-421.58	-56.82
Amylose ₁	204	-5.77	-22.47	-16.70	-0.33	-1.49	-1.16
Amylose ₂	389	-78.64	-148.84	-70.21	-9.52	-15.47	-5.95
Amylose ₄	759	-446.49	-636.33	-189.83	-150.37	-169.29	-18.92
Amylose ₈	1499	-1305.11	-1753.60	-448.48	-530.38	-577.58	-47.19
DNA ₁	579	-259.68	-396.78	-137.11	-102.97	-128.32	-25.35
DNA ₂	1252	-1335.67	-1919.00	-583.33	-471.24	-555.70	-84.46

Table 5.5: Errors [μ Hartree] of CDD-MP2 total energy as well as Coulomb and exchange parts, with respect to canonical MO-MP2 reference ($\Delta E_X = E_{X,Ref}^{(2)} - E_{X,CDD}^{(2)}$) for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the 6-31G* basis set; 5 quadrature points; numbers (k, l, m) indicate numerical thresholds $\theta_{ext} = 10^{-k}$, $\theta_{int} = 10^{-l}$, $\theta_{sp} = 10^{-m}$.

Molecule	N	ΔE (7, 7, 9)			ΔE (8, 8, 9)		
		TOTAL	Coulomb	Exchange	TOTAL	Coulomb	Exchange
C ₁₀ H ₂₂	250	-34.92	-85.79	-50.87	-9.18	-13.50	-4.32
C ₂₀ H ₄₂	490	-117.18	-260.29	-143.11	-31.66	-46.97	-15.31
C ₄₀ H ₈₂	970	-308.18	-634.72	-326.54	-105.87	-143.57	-37.70
Amylose ₁	228	-13.34	-43.94	-30.61	-1.13	-3.89	-2.76
Amylose ₂	432	-44.30	-136.21	-91.91	-6.12	-15.53	-9.41
Amylose ₄	840	179.39	-55.45	-234.84	194.49	168.63	-25.86
DNA ₁	625	-6.68	-169.25	-162.57	-14.09	-37.97	-23.88
DNA ₂	1332	337.60	-361.06	-698.66	371.59	275.97	-95.61

Table 5.6: Errors [μ Hartree] of CDD-MP2 total energy as well as Coulomb and exchange parts, with respect to canonical MO-MP2 reference ($\Delta E_X = E_{X,Ref}^{(2)} - E_{X,CDD}^{(2)}$) for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the SVP (VDZ**) basis set; 5 quadrature points; numbers (k, l, m) indicate numerical thresholds $\theta_{ext} = 10^{-k}$, $\theta_{int} = 10^{-l}$, $\theta_{sp} = 10^{-m}$.

Molecule	N	ΔE (7, 7, 9)			ΔE (8, 8, 9)		
		TOTAL	Coulomb	Exchange	TOTAL	Coulomb	Exchange
$C_{10}H_{22}$	608	88.86	-139.17	-228.03	-16.32	-37.69	-21.37
$C_{20}H_{42}$	1188	837.54	221.05	-616.50	513.01	445.70	-67.30
Amylose ₁	528	-14.95	-183.38	-168.43	-8.60	-26.33	-17.73
Amylose ₂	998	1197.14	718.69	-478.44	168.02	115.01	-53.01
Amylose ₄	1938	5836.13	4714.62	-1121.51	-	-	-
DNA ₁	1428	5190.63	4411.45	-779.18	2250.67	2102.36	-148.31

Table 5.7: Errors [μ Hartree] of CDD-MP2 total energy as well as Coulomb and exchange parts, with respect to canonical MO-MP2 reference ($\Delta E_X = E_{X,Ref}^{(2)} - E_{X,CDD}^{(2)}$) for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the cc-pVTZ basis set; 5 quadrature points; numbers (k, l, m) indicate numerical thresholds $\theta_{ext} = 10^{-k}$, $\theta_{int} = 10^{-l}$, $\theta_{sp} = 10^{-m}$.

Complex	E_{int}	ΔE_{int} (SVP)			E_{int}	ΔE_{int} (cc-pVTZ)		
		5; 6, 6, 9	5; 8, 8, 9	6;10,10,12		5; 6, 6, 9	5; 8, 8, 9	6;10,10,12
01. Ammonia dimer	-2.011	0.011	0.001	0.000	-2.774	-0.130	-0.002	0.000
02. Water dimer	-4.263	0.004	0.001	0.000	-4.434	0.001	-0.003	0.000
03. Formic acid dimer	-14.723	0.013	-0.005	-0.001	-16.891	0.065	-0.013	-0.003
04. Formamide dimer	-12.042	0.016	0.000	0.000	-14.247	0.073	-0.010	0.000
05. Uracil dimer (H-bonded)	-16.571	-0.173	-0.068	-0.016	-18.741	-4.706	-1.778	-0.032
06. 2-Pyridoxine 2-aminopyridine	-14.220	-0.258	-0.069	-0.011	-15.907	-4.382	-0.335	-0.021
07. Adenine-thymine (Watson-Crick)	-13.033	-0.302	-0.093	-0.012	-14.919	-4.644	-1.668	-0.038
08. Methane dimer	-0.033	0.024	0.001	0.000	-0.317	-0.128	-0.001	0.000
09. Ethene dimer	-0.243	-0.008	-0.006	-0.001	-1.150	-0.326	-0.018	-0.002
10. Benzene-methane	-0.500	0.023	-0.016	-0.003	-1.410	-0.466	-0.013	-0.007
11. Benzene dimer (parallelly displaced)	-1.534	-0.067	-0.056	-0.007	-3.760	-1.666	-0.095	-0.024
12. Pyrazine dimer	-3.049	-0.075	-0.092	-0.014	-5.435	-0.243	-0.135	-0.043
13. Uracil dimer (stack)	-5.693	-0.051	-0.052	-0.015	-9.089	-0.925	-0.058	-0.020
14. Indole-benzene (stack)	-3.292	-0.095	-0.062	-0.008	-6.428	-3.039	-0.087	-0.018
15. Adenine-thymine (stack)	-7.736	-0.100	-0.063	-0.013	-12.287	-2.409	-0.073	-0.021
16. Ethene-ethyne	-0.917	0.030	-0.002	0.000	-1.426	-0.149	-0.009	-0.001
17. Benzene-water	-1.964	-0.028	-0.038	-0.011	-2.927	-0.129	-0.033	-0.016
18. Benzene-ammonia	-1.188	-0.017	-0.027	-0.007	-2.127	-0.236	-0.020	-0.009
19. Benzene-HCN	-3.021	0.013	-0.040	-0.007	-4.600	-0.524	-0.056	-0.009
20. Benzene dimer (T-shape)	-1.633	0.007	-0.053	-0.006	-3.004	-1.864	-0.095	-0.019
21. Indole-benzene (T-shape)	-4.024	-0.015	-0.053	-0.005	-6.155	-6.076	-0.109	-0.021
22. Phenol dimer	-5.640	-0.111	-0.079	-0.014	-6.707	-1.978	-0.132	-0.026
Mean error		-0.053	-0.040	-0.007		-1.540	-0.216	-0.015
Sample standard deviation		0.091	0.032	0.006		1.884	0.494	0.013

Table 5.8: MO-MP2 interaction energies [kcal/mol] (counterpoise corrected) and errors of CDD-MP2 interaction energies with respect to MO-MP2 reference ($\Delta E_{\text{int}} = E_{\text{int,Ref}}^{(2)} - E_{\text{int,CDD}}^{(2)}$), along with mean errors and standard deviations, for the S22 test set with the SVP and cc-pVTZ basis sets and the indicated numerical settings. The numbers $m; n, o, p$ stand for m quadrature points and thresholds $\theta_{\text{ext}} = 10^{-n}$, $\theta_{\text{int}} = 10^{-o}$, $\theta_{\text{sp}} = 10^{-p}$, respectively.

Basis	Molecule	N	t_{CDD}	Sc.	t_{MO}	Sc.
6-31G*	C ₁₀ H ₂₂	194	470	-	42	-
	C ₂₀ H ₄₂	384	2483	2.4	290	2.8
	C ₄₀ H ₈₂	764	9900	2.0	5473	4.3
	C ₈₀ H ₁₆₂	1524	26477	1.4	104442 ^a	4.4
	C ₁₆₀ H ₃₂₂	3044	68748	1.4	2192294 ^a	4.4
	Amylose ₁	204	745	-	84	-
	Amylose ₂	389	5714	3.2	849	3.6
	Amylose ₄	759	35259	2.7	10243	3.7
	Amylose ₈	1499	173700	2.3	118606	3.6
	Amylose ₁₆	2979	778829	2.2	1404545 ^a	3.6
SVP	C ₁₀ H ₂₂	250	1620	-	48	-
	C ₂₀ H ₄₂	490	9161	2.6	575	3.7
	C ₄₀ H ₈₂	970	38122	2.1	8693	4.0
	C ₈₀ H ₁₆₂	1930	120290	1.7	134110 ^a	4.0
	C ₁₆₀ H ₃₂₂	3850	333675	1.5	2090245 ^a	4.0
	Amylose ₁	228	1798	-	84	-
	Amylose ₂	432	13659	3.2	1036	3.9
	Amylose ₄	840	84796	2.7	16154	4.1
	Amylose ₈	1656	450376	2.5	251884 ^a	4.1
	Amylose ₁₆	3288	1658424	1.9	4157621 ^a	4.1
cc-pVTZ	C ₁₀ H ₂₂	608	27427	-	1639	-
	C ₂₀ H ₄₂	1188	154391	2.6	14174	3.2
	C ₄₀ H ₈₂	2348	715438	2.3	284035 ^a	4.4
	C ₈₀ H ₁₆₂	4668	2886367 ^b	2.0	5840847 ^a	4.4
	Amylose ₁	528	27152	-	1515	-
	Amylose ₂	998	211859	3.2	25065	4.4
	Amylose ₄	1938	2327338	3.6	467098 ^a	4.4

^a Extrapolated from previous molecule with indicated scaling exponent

^b Extrapolated from previous molecule using indicated scaling exponent from first quadrature point

Table 5.9: Total Wall times [s] of CDD-MP2 and conventional MO-MP2 along with scaling exponents with respect to previous molecule for calculations on linear alkanes and α -D-glucose chains (amylose) with the indicated basis sets; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{sp} = 10^{-9}$.

5.3. RESULTS AND DISCUSSION

	C ₁₀ H ₂₂	C ₂₀ H ₄₂	C ₄₀ H ₈₂	C ₈₀ H ₁₆₂	C ₁₆₀ H ₃₂₂ ^a	Scaling			
	(194)	(384)	(764)	(1524)	(3044)				
TOTAL	470	2483	9900	26477	68748	2.4	2.0	1.4	1.4
Cholesky dec.	0	1	7	45	347	-	2.8	2.7	3.0
External scr. FTI	4	18	75	288	1508	2.2	2.1	1.9	2.4
Internal scr. HTI	5	87	579	1742	4948	4.2	2.8	1.6	1.5
Internal scr. UTI	56	329	1588	4460	12591	2.6	2.3	1.5	1.5
Make UTI	198	807	3203	7495	17064	2.1	2.0	1.2	1.2
1st transfo.	48	398	1479	3849	8898	3.1	1.9	1.4	1.2
2nd transfo.	54	386	1466	3922	8649	2.9	1.9	1.4	1.1
Screening TTI	9	55	166	709	1712	2.7	1.6	2.1	1.3
3rd transfo.	22	113	368	799	1837	2.4	1.7	1.1	1.2
4th transfo.	8	41	119	325	946	2.4	1.5	1.5	1.5
Contraction	1	0	2	4	9	-	-	1.0	1.2

^a Onset of HTI batching

Table 5.10: Wall times [s] of total calculation and essential steps of the CDD-MP2 algorithm for linear alkanes, along with scaling exponents between consecutive molecules, for the 6-31G* basis set. The numbers in parentheses denote numbers N of basis functions; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{int} = 10^{-9}$.

	C ₁₀ H ₂₂ (250)	C ₂₀ H ₄₂ (490)	C ₄₀ H ₈₂ (970)	C ₈₀ H ₁₆₂ ^a (1930)	C ₁₆₀ H ₃₂₂ (3850)	Scaling			
TOTAL	1620	9161	38122	120290	333675	2.6	2.1	1.7	1.5
Cholesky dec.	0	0	10	86	712	-	-	3.1	3.1
External scr. FTI	7	32	141	601	3760	2.3	2.2	2.1	2.7
Internal scr. HTI	18	309	2566	9560	28948	4.2	3.1	1.9	1.6
Internal scr. UTI	224	1652	8321	31350	93520	3.0	2.4	1.9	1.6
Make UTI	668	3057	11091	29255	69662	2.3	1.9	1.4	1.3
1st transfo.	262	1667	7075	19563	47224	2.7	2.1	1.5	1.3
2nd transfo.	154	1085	4067	12636	31848	2.9	1.9	1.6	1.3
Screening TTI	27	126	471	2435	9160	2.3	1.9	2.4	1.9
3rd transfo.	42	301	1100	2784	6722	2.9	1.9	1.3	1.3
4th transfo.	22	120	299	1536	4987	2.5	1.3	2.4	1.7
Contraction	0	2	5	10	19	-	1.3	1.0	0.9

^a Onset of HTI batching

Table 5.11: Wall times [s] of total calculation and essential steps of the CDD-MP2 algorithm for linear alkanes, along with scaling exponents between consecutive molecules, for the SVP (VDZ**) basis set; The numbers in parentheses denote numbers N of basis functions; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{int} = 10^{-9}$.

	C ₁₀ H ₂₂ (608)	C ₂₀ H ₄₂ ^a (1188)	C ₄₀ H ₈₂ (2348)	Scaling	
TOTAL	26670	153471	829885	2.6	2.5
Cholesky dec.	2	17	145	3.2	3.1
External scr. FTI	95	418	1765	2.2	2.1
Internal scr. HTI	157	2996	31514	4.4	3.5
Internal scr. UTI	2435	19581	136609	3.1	2.9
Make UTI	10381	45892	172215	2.2	1.9
1st transfo.	7102	39037	164153	2.5	2.1
2nd transfo.	3217	22783	122302	2.9	2.5
Screening TTI	345	4386	75190	3.8	4.2
3rd transfo.	471	3630	12736	3.0	1.8
4th transfo.	480	5758	74097	3.7	3.7
Contraction	4	22	53	2.5	1.3

^a Onset of HTI batching

Table 5.12: Wall times [s] of total calculation and essential steps of the CDD-MP2 algorithm for linear alkanes, along with scaling exponents between consecutive molecules, for the cc-pVTZ basis set; The numbers in parentheses denote numbers N of basis functions; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{int} = 10^{-9}$.

	Amylose ₁	Amylose ₂	Amylose ₄ ^a	Amylose ₈	Amylose ₁₆	Scaling			
	(204)	(389)	(759)	(1499)	(2979)				
TOTAL	745	5714	35259	173700	778829	3.2	2.7	2.3	2.2
Cholesky dec.	0	1	6	41	352	-	2.7	2.8	3.1
External scr. FTI	11	81	305	1161	5461	3.1	2.0	2.0	2.3
Internal scr. HTI	5	101	1857	13574	47362	4.7	4.4	2.9	1.8
Internal scr. UTI	70	601	4385	20574	68552	3.3	3.0	2.3	1.8
Make UTI	287	1517	6279	23441	96434	2.6	2.1	1.9	2.1
1st transfo.	102	925	6035	27943	80333	3.4	2.8	2.3	1.5
2nd transfo.	81	937	6194	25312	99258	3.8	2.8	2.1	2.0
Screening TTI	22	353	2709	24978	188104	4.3	3.0	3.3	2.9
3rd transfo.	34	424	3267	9704	24979	3.9	3.1	1.6	1.4
4th transfo.	38	370	2299	18736	131560	3.5	2.7	3.1	2.8
Contraction	1	4	13	30	66	2.1	1.8	1.2	1.1

^a Onset of HTI batching

Table 5.13: Wall times [s] of total calculation and essential steps of the CDD-MP2 algorithm for α -D-glucose chains (amylose), along with scaling exponents between consecutive molecules, for the 6-31G* basis set; The numbers in parentheses denote numbers N of basis functions; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{int} = 10^{-9}$.

	Amylose ₁	Amylose ₂	Amylose ₄ ^a	Amylose ₈	Amylose ₁₆	Scaling			
	(228)	(432)	(840)	(1656)	(3288)				
TOTAL	1798	13659	84796	450376	1658424	3.2	2.7	2.5	1.9
Cholesky dec.	0	0	5	59	462	-	-	3.6	3.0
External scr. FTI	12	84	344	1416	6366	3.0	2.1	2.1	2.2
Internal scr. HTI	15	288	5352	45813	179207	4.6	4.4	3.2	2.0
Internal scr. UTI	238	2447	17690	100997	404995	3.6	3.0	2.6	2.0
Make UTI	681	3730	16122	63771	253001	2.7	2.2	2.0	2.0
1st transfo.	345	2898	16860	73160	220750	3.3	2.6	2.2	1.6
2nd transfo.	178	1567	9625	38166	125248	3.4	2.7	2.0	1.7
Screening TTI	26	434	5049	50821	217547	4.4	3.7	3.4	2.1
3rd transfo.	64	788	5617	18205	45520	3.9	3.0	1.7	1.3
4th transfo.	45	400	3364	26298	95416	3.4	3.2	3.0	1.9
Contraction	2	6	13	35	76	1.7	1.2	1.5	1.1

^a Onset of HTI batching

Table 5.14: Wall times [s] of total calculation and essential steps of the CDD-MP2 algorithm for α -D-glucose chains (amylose), along with scaling exponents between consecutive molecules, for the SVP (VDZ**) basis set; The numbers in parentheses denote numbers N of basis functions; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{int} = 10^{-9}$.

	Amylose ₁ (228)	Amylose ₂ ^a (432)	Amylose ₄ (840)	Scaling	
TOTAL	27152	211859	2327338	3.2	3.6
Cholesky dec.	1	11	85	3.8	3.1
External scr. FTI	113	727	2846	2.9	2.1
Internal scr. HTI	115	2205	109733	4.6	5.9
Internal scr. UTI	2295	24222	240578	3.7	3.5
Make UTI	10006	50268	410247	2.5	3.2
1st transfo.	7707	51799	384470	3.0	3.0
2nd transfo.	3357	29093	235341	3.4	3.2
Screening TTI	375	14231	396295	5.7	5.0
3rd transfo.	607	7499	55508	3.9	3.0
4th transfo.	611	21442	414029	5.6	4.5
Contraction	9	41	129	2.4	1.7

^a Onset of HTI batching

Table 5.15: Wall times [s] of total calculation and essential steps of the CDD-MP2 algorithm for α -D-glucose chains (amylose), along with scaling exponents between consecutive molecules, for the cc-pVTZ basis set; The numbers in parentheses denote numbers N of basis functions; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-7}$, $\theta_{int} = 10^{-7}$, $\theta_{int} = 10^{-9}$.

Chapter 6

Resolution of the Identity CDD-MP2

In many situations, the evaluation of the $N^2 \times N^2$ matrix of four-center two-electron integrals is one of the central computational obstacles. One approach to alleviate this problem is to reduce it to the computation of three- and two-center integrals through the application of resolution of the identity (RI) techniques [14–18] for approximation of the ERI. This procedure has been applied to the CDD method by Zienau *et al.* in a resolution of the identity CDD-MP2 (RI-CDD-MP2) implementation [10, 11, 13]. Being originally limited to the calculation of the Coulomb-type component $E_J^{(2)}$ within the SOS-MP2 scheme [8] (cf. section 3.3.3), this algorithm has in the present work been completed for calculation of the total MP2 energy $E_0^{(2)}$. In addition, certain steps of the algorithm that weren't already inherently parallel through their use of MKL library routines have been parallelized with the OpenMP (OMP) procedure [125].

This chapter begins with a review of RI-theory and how it is applied within RI-CDD-MP2, followed by a brief presentation of the concepts of parallel computing. The RI-CDD-MP2 algorithm is then recalled in order to discuss the contributions that have been made to it, after which performance and accuracy of the extended method are discussed, as well as the efficiency of the OMP parallelization.

6.1 Resolution of the Identity and Density Fitting

For any given complete orthonormal set of basis functions $\{\xi(\mathbf{r})\}$ spanning a space \mathcal{X} , which, using the previously established notation, shall be arranged into the vector $|\boldsymbol{\xi}\rangle$, the projection operator

$$\hat{\rho}_\xi = |\boldsymbol{\xi}\rangle\langle\boldsymbol{\xi}| = 1 \tag{6.1}$$

corresponds to the identity operator, since the action of $\hat{\rho}_\xi$ on any function projects it onto its (exact) representation in the basis $\{\xi(\mathbf{r})\}$. The application of this closure relation is thus termed "resolution of the identity" (RI). When using the RI in practice, the introduced auxiliary basis $\{\xi_q(\mathbf{r})\}_{q \in N_{\text{aux}}}$ is necessarily finite (typically $3N \lesssim N_{\text{aux}} \lesssim 7N$, where N denotes the size

of the AO-basis), such that the closure relation cannot be strictly fulfilled. In general, the auxiliary basis is also non-orthogonal, but it may still be orthonormalized, for instance through symmetric orthonormalization (see section 2.3.3) $|\xi'\rangle = |\xi\rangle\Xi^{-\frac{1}{2}}$, such that the corresponding projector reads

$$\hat{\rho}_{\xi'} = |\xi'\rangle\langle\xi'| = |\xi\rangle\Xi^{-1}\langle\xi| \approx 1, \quad (6.2)$$

where $\Xi = (\Xi_{pq})_{p,q \in N_{\text{aux}}}$ denotes the Gramian matrix, i.e. metric, of the auxiliary basis functions. Now, recalling equation (5.31), \mathbf{A} is the matrix representation of the Coulombic repulsion operator $\hat{C} = 1/|\mathbf{r}_1 - \mathbf{r}_2|$ in the basis of the two-center one-electron distributions $\Omega(\mathbf{r}) = \chi(\mathbf{r}) \otimes \chi(\mathbf{r})$, i.e.

$$\mathbf{A} = (\Omega|\hat{C}|\Omega) \equiv (\Omega|\Omega)_C, \quad (6.3)$$

where in the last step we have established the notation that shall be used throughout this section to distinguish integrals involving the operator \hat{C} from normal inner products. Insertion of the projector (6.2) onto the space \mathcal{X} into the expression for matrix \mathbf{A} according to

$$\mathbf{A} \approx (\Omega|\hat{\rho}_{\xi'}|\Omega)_C = (\Omega|\xi)\Xi^{-1}(\xi|\Omega)_C \quad (6.4)$$

then leads to the aforementioned three- and two-center integrals. Upon second insertion of $\hat{\rho}_{\xi'}$, the symmetric relation

$$\mathbf{A} \approx (\Omega|\xi)\Xi^{-1}(\xi|\xi)_C\Xi^{-1}(\xi|\Omega) \quad (6.5)$$

results. The use of RI in this particular form, however, has been shown to be too inaccurate in practice, so that another definition is required. Namely, an alternative approach consists of expanding the two-center charge densities in the auxiliary basis, i.e. to determine a coefficient matrix \mathbf{D} such that

$$\Delta = |\Omega\rangle - |\xi\rangle\mathbf{D}^T = \mathbf{0}, \quad (6.6)$$

whence the denomination of *density fitting* (DF) for this procedure [126]. It may be shown that optimal coefficients are obtained through minimization of the least-squares error of the electric field [16, 127], i.e. minimization of

$$(\Delta|\Delta)_C = (\Omega|\Omega)_C - (\Omega|\xi)_C\mathbf{D}^T - \mathbf{D}(\xi|\Omega)_C + \mathbf{D}(\xi|\xi)_C\mathbf{D}^T, \quad (6.7)$$

such that upon derivation with respect to \mathbf{D} the optimal coefficients are determined as

$$\mathbf{D} = (\Omega|\xi)_C(\xi|\xi)_C^{-1} = (\Omega|\xi)_C\Xi_C^{-1}. \quad (6.8)$$

The significant difference of this approach is that the resulting metric is defined by the overlap $\Xi_C = (\xi|\hat{C}|\xi)$ of the auxiliary basis functions with respect to the Coulombic repulsion operator. Inserting the density fitting relation (6.6) into expression (6.3) for the ERI matrix, and using

the expansion coefficients of the previous equation, the optimal RI-approximation of the ERI matrix is obtained as

$$\mathbf{A} \approx (\boldsymbol{\Omega}|\boldsymbol{\xi})_C \boldsymbol{\Xi}_C^{-1} (\boldsymbol{\xi}|\boldsymbol{\Omega})_C. \quad (6.9)$$

Performing a symmetric factorization of $\boldsymbol{\Xi}_C^{-1}$, typically by forming the matrix square root (the Cholesky decomposition is applicable as well), the RI-decomposition of the ERI matrix may be brought to the final form

$$\mathbf{A} \approx \left((\boldsymbol{\Omega}|\boldsymbol{\xi})_C \boldsymbol{\Xi}_C^{-\frac{1}{2}} \right) \left(\boldsymbol{\Xi}_C^{-\frac{1}{2}} (\boldsymbol{\xi}|\boldsymbol{\Omega})_C \right) = \mathbf{B}' \mathbf{B}'^T, \quad (6.10)$$

where the orthonormalized matrix \mathbf{B}' of three-center integrals is formally of dimension $N^2 \times N_{\text{aux}}$. Within the RI-CDD-MP2 approach, the matrix \mathbf{M} of pseudo-MO integrals in the energy expression (5.11) is thus constructed according to

$$\mathbf{M} \approx (\underline{\mathbf{L}}^T \otimes \bar{\mathbf{L}}^T) \mathbf{B}' \mathbf{B}'^T (\underline{\mathbf{L}} \otimes \bar{\mathbf{L}}) = \mathbf{B} \mathbf{B}^T. \quad (6.11)$$

6.2 The RI-CDD-MP2 algorithm

The RI-CDD-MP2 algorithm developed and implemented by Zienau *et al.* is extensively described in [13], and only a brief overview of its structure shall be given here in order to discuss the modifications added to it in the context of the present work. The preliminary steps, i.e. Cholesky decomposition of the pseudo-density matrices and formation of screening prerequisites are identical to the CDD-MP2 case. Also, as the MBIE integral screening protocol [5, 128] which was initially employed in RI-CDD-MP2 has been revealed as erroneously implemented and too conservative [129, 130], it has here been replaced by the Schwarz/R procedure [12] that is used within CDD-MP2 (see section 5.1.3). Due to the use of the RI approximation, the main difference between RI-CDD-MP2 and the CDD-MP2 implementation described in the previous chapter lies in the integral formation procedure, and the overall structure of the algorithm is [13]:

- Computation of matrix $(\boldsymbol{\Omega}|\boldsymbol{\xi})_C$ of untransformed three-center integrals (Disk)
- Computation of inverse auxiliary overlap matrix square root $\boldsymbol{\Xi}_C^{-\frac{1}{2}}$ (Disk)
- Construction of orthonormalized matrix $\mathbf{B}' = (\boldsymbol{\Omega}|\boldsymbol{\xi})_C \boldsymbol{\Xi}_C^{-\frac{1}{2}}$ (Disk)
- Loop over quadrature points:
 - Cholesky decomposition of pseudo-densities and formation of screening prerequisites (RAM)
 - External screening for *FTidx* of required FTI (Disk)
 - Internal screening for intermediaries in transformations from \mathbf{B}' to \mathbf{B} (RAM)

- Perform consecutive transformations according to $\mathbf{B} = (\mathbf{1} \otimes \bar{\mathbf{L}}^T)(\underline{\mathbf{L}}^T \otimes \mathbf{1})\mathbf{B}'$ (Disk)
- Construct FTI matrix $\mathbf{M} = \mathbf{B}\mathbf{B}^T$ (RAM) and contract to energy

End

The algorithm is integral-indirect, such that apart from those steps annotated with (RAM), all arrays are written to disk and read in as required. As mentioned previously, Zienau's implementation was limited to the computation of the Coulomb part $E_J^{(2)}$ of the energy, and has here been complemented for calculation of the exchange part in order to allow full MP2 calculations. The central steps of generating the matrix \mathbf{B} of transformed three-center integrals and its assembly to the FTI matrix are totally unaffected by this extension, since both Coulomb and exchange type parts use the same integrals. The complementation rather consisted of adapting peripheral steps such as the screening protocol, for the reasons mentioned above, as well as the contraction to the MP2 energy.

6.3 OMP Parallelization

The objective of parallel programming is to exploit modern multi-core processor architectures to achieve reductions of computational time by going over from sequential to parallel execution, i.e. simultaneous execution of programme parts by different cores. For instance, in order to parallelize a standard control structure such as a loop, the set over which the loop variable runs is simply divided into subsets that are treated individually by a *thread*. The OpenMP (OMP) application programming interface (API) [125] is an easily applicable tool for this purpose, and in general the parallelization of present code merely requires the adaptation of the loops in question, with corresponding redeclarations of variables that are to be private for each particular thread or shared by all threads.

However, not all programmes may be straightforwardly parallelized, and some not at all. Namely, in order to avoid conflicts between different threads and ensure *thread safety*, steps involving communication among threads, for instance during the modification of shared variables by separate threads, may require to be placed within so-called *critical* environments. Critical code can only be executed sequentially, such that unpredictable programme behaviour due to simultaneous modification of a shared variable by different threads is precluded. When critical steps occur in the innermost control structure, parallelization loses all potential for performance improvement and would actually impede the flow of execution.

Considering the sequential and parallel programme execution times t_1 and t_p , performance gains may be assessed by evaluating the parallelization *efficiency* η , which is defined as the ratio

$$\eta = \frac{S}{N_{\text{thr}}} = \left(\frac{t_1}{t_p} \right) \cdot \frac{1}{N_{\text{thr}}} \quad (6.12)$$

of speed-up $S = t_1/t_p$ to the number N_{thr} of threads. Optimal performance is thus indicated by an efficiency value of 1, whereas the worst case of obtaining no speed-up is expressed by an efficiency of $1/N_{\text{thr}}$.

6.4 Results and discussion

The value of the completed RI-CDD-MP2 algorithm lies in the ability to calculate full MP2 energies, instead of merely the Coulomb-like component $E_J^{(2)}$ to which it was previously restricted. As far as computational scaling and performance are concerned, the modifications made for inclusion of the exchange-type energy $E_K^{(2)}$ lead to no fundamental differences as compared to the original sequential RI-CDD-MP2 algorithm, since only peripheral steps needed to be adapted. Therefore, due to the fact that this aspect has already been thoroughly discussed by Zienau *et al.* [11,13], the subject of timings is broached only briefly, and stronger emphasis is placed on the numerical accuracy of RI-CDD-MP2 energies with the corrected external screening procedure and inclusion of the exchange-type term. The discussion then proceeds with an analysis of the efficiency of the parallelized RI-CDD-MP2 algorithm.

All calculations were performed on the same computer cluster as the results presented for CDD-MP2, and for comparability, are all based on the same standard settings for underlying HF calculations and other numerical settings that are given in section 5.3. Except for the investigations on parallelization efficiency, all calculations were performed purely sequentially, i.e. with one core.

6.4.1 Timings

All relevant data pertaining to the computational performance of the completed sequential RI-CDD-MP2 algorithm has been summarized in table 6.1, which contains numbers of significant FTI preselected in the external screening and calculation times both for RI-CDD-MP2 and conventional RI-MO-MP2, for comparison, along with scaling exponents of the individual quantities.

Unsurprisingly, the behaviour of the numbers of FTI that are preselected in the external screening procedure is the same as that observed in section 5.3.2 for the case of CDD-MP2, since this step is performed identically. For all but the bulky DNA systems, the FTI numbers rapidly achieve linear scaling with system size. As to the corresponding calculation times, it has to be kept in mind that even for a linear scaling number of FTI, the use of the RI approximation at best allows for quadratically scaling effort of the integral assembly step $\mathbf{M} = \mathbf{B}\mathbf{B}^T$. Being typically the most time-consuming, this step therefore determines the overall scaling of RI-CDD-MP2 calculations. On the other hand, however, the matrix inversion required to obtain $\Xi_C^{-\frac{1}{2}}$

scales cubically. This is also the case for the matrix product $\mathbf{B}' = (\mathbf{\Omega}|\boldsymbol{\xi})_C \boldsymbol{\Xi}_C^{-\frac{1}{2}}$, which, though formally of quartic scaling, is asymptotically cubic, since the exponential coupling of Gaussian basis functions constituting the $\Omega_{\mu\nu}$ distributions leads to the number of significant $(\Omega_{\mu\nu}|\xi_p)$ three-center integrals to scale quadratically. Accordingly, in spite of relatively small prefactors, these steps will asymptotically dominate the calculation time. Inspection of the timings given in table 6.1 reveals precisely these tendencies. For linear alkanes and amylose chains with both basis sets, although the computation time tends towards quadratic, true quadratic scaling is not achieved in spite of near-linear scaling of the FTI number. On the one hand, this indicates that the truly quadratic scaling regime of the FTI assembly step could not be reached with the investigated systems, but also shows that with increasing molecular size, the aforementioned cubically scaling steps have an impact on the total scaling exponent. For the compact DNA systems, the observed high scaling exponents of the computational time are in good agreement with the large numbers of preselected FTI.

In terms of performance as compared to $\mathcal{O}(N^5)$ canonical RI-MO-MP2, the lower scaling of the RI-CDD-MP2 approach leads to the cross-over points to be located at moderate system sizes. For alkanes, both with SVP and cc-pVTZ, the cross-over point is situated between $\text{C}_{20}\text{H}_{42}$ and $\text{C}_{40}\text{H}_{82}$, and for amylose chains with both bases, it is located between Amylose_4 and Amylose_8 . As to the DNA systems that were investigated, the cross-over point of RI-CDD-MP2 is between DNA_2 and DNA_4 , but it could not be reached with the cc-pVTZ basis.

6.4.2 Accuracy

The issue of more interest, namely the accuracy provided for MP2 energies by the completed RI-CDD-MP2 algorithm with the Schwarz/R integral preselection procedure, shall now be addressed. Errors of RI-CDD-MP2 absolute energies with respect to canonical RI-MO-MP2 are given in tables 6.2 and 6.3 for the SVP and cc-pVTZ basis set, respectively. In both cases, the loosest numerical thresholds lead to arguably too large errors, particularly for the DNA systems. With the tighter threshold combination (8, 8, 9) (see definition in table), the situation is improved and the obtained accuracies are reasonable, whereas (9, 9, 10) provides very accurate results. Considering the individual accuracies of the Coulomb-type and exchange-type component $E_J^{(2)}$ and $E_K^{(2)}$ of the MP2 energy, in part similar trends are observed as for the CDD-MP2 method discussed in section 5.3.2. We notice in particular that, just as in the CDD-MP2 case, the error of the exchange component behaves very systematically, i.e. that it diminishes by roughly an order of magnitude upon lowering the screening threshold by an order of magnitude, and further that it grows almost linearly with system size. The Coulomb-type component, on the other hand, is not so well-behaved because the error grows stronger than linearly with the system size, in particular for the DNA systems. Also, upon lowering the external screening threshold, the

improvements in accuracy are less pronounced than for its exchange counterpart.

However, in stark contrast to the numerical problems that arise with CDD-MP2 (see discussion in section 5.3.2), the errors of $E_J^{(2)}$ are here consistently negative, such that the RI-CDD-MP2 Coulomb-type component fulfils the theoretical condition of representing a lower bound to the exact $E_J^{(2)}$. Although the additive error contribution in the Coulomb component, arising from the inaccuracy induced in the FTI by the internal screening, is here also in principle present, this fundamental difference of the observed accuracies is due to the different ways in which the integral transformation is performed. Namely, in the case of CDD-MP2, the numerical error induced by the neglect of insignificant transformation intermediaries can strongly propagate over the course of the four consecutive transformations from UTI to final FTI. In the RI-CDD-MP2 approach, on the other hand, where the FTI matrix is assembled according to $\mathbf{M} = \mathbf{B}\mathbf{B}^T$, the matrix \mathbf{B} of transformed three-center integrals is obtained from the untransformed \mathbf{B}' matrix after only two transformations, such that there is far less room for error propagation as for CDD-MP2. Furthermore, since in the RI-approach the computational effort associated with these transformation steps is significantly smaller than for CDD-MP2, where the minimization of the transformation effort demanded a highly selective internal screening procedure, the internal screening used in Zienau’s RI-CDD-MP2 implementation is far less selective. This further reduces the numerical error in the transformed integrals, such that for RI-CDD-MP2 the negative deviation of $E_J^{(2)}$, arising from the neglect of insignificant FTI products, clearly dominates over its additive counterpart. Therefore, due to the numerical error in the exchange component being also always negative in all the considered cases and the fact that $E_K^{(2)}$ enters into the MP2 energy with a different sign than $E_J^{(2)}$, the errors in both terms partially cancel in the final energy, and since the error in $E_K^{(2)}$ is in general about an order of magnitude smaller, the overall error is mostly determined by the Coulomb error.

In order to compare the error due to the integral screening in RI-CDD-MP2 with the systematic error induced by the RI-approximation, tables 6.4 and 6.5 display deviations of conventional RI-MO-MP2 and RI-CDD-MP2 energies from their respective non-RI canonical MO-MP2 references, for SVP and cc-pVTZ. For both basis sets, the RI error (left part) is far more significant than the screening errors discussed previously, such that the overall errors of RI-CDD-MP2 with respect to MO-MP2 (right part) are only marginally different from the deviations of RI-MO-MP2. Furthermore, since the screening errors were observed to behave systematically, i.e. be always of the same sign, there is no reason why the RI-CDD-MP2 overall errors should not compensate each other to a similar extent as the RI error upon formation of energy differences for relative energies. To demonstrate this, accuracies of RI-CDD-MP2 interaction energies determined for the S22 test set [124] are given in table 6.6. For both basis sets, the very loose threshold combination (6, 6, 9) already provides chemical accuracy for all complexes, even for systems 5, 6, 7, 11, 14, 15, 20, 21 or 22, where CDD-MP2 produced unacceptably high errors

with cc-pVTZ (see table 5.8). Further, upon tightening of the numerical thresholds, the improvement of the accuracy is in general relatively small and the error even fluctuates in some cases, for instance 11, 12 or 22 with cc-pVTZ. This indicates that the results with loose thresholds may already be considered as converged, since although the absolute errors are larger with loose thresholds, upon formation of relative energies they cancel each other to practically the same extent as the errors of calculations using tighter screening conditions. Thus, no particular improvement is obtained through the imposition of stricter screening, which merely leaves room for slight variations in the individual absolute errors to increase or decrease the relative errors in a fortuitous rather than systematic way.

6.4.3 Parallel RI-CDD-MP2

As has been mentioned in the introduction, the RI-CDD-MP2 method has been parallelized using the OMP API for various steps that weren't outright parallel through their use of library routines, such as the inversion of the auxiliary overlap matrix Ξ_C or the FTI assembly step $\mathbf{M} = \mathbf{B}\mathbf{B}^T$, which invoke Intel MKL library routines [131]. For reasons of thread safety (see section 6.3), however, some sequential steps could not be parallelized, for instance the external screening where the progressive build-up of the *FTidx* array precludes parallelization of the occurring loops. In addition, the parallelization of steps involving the generation of particular integrals, such as the formation of the metric Ξ_C , the matrix $(\Omega|\xi)_C$ of three-center integrals, and the evaluation of the screening matrices from equation (5.22), is part of the larger undertaking of parallelizing the entire integral evaluation code and was outside the scope of the present work. In the preliminary RI-CDD-MP2 algorithm presented here, parallelization has been applied to the calculation of the centers and extents of transformed charge distributions defined in equations (5.41) and (5.43), which are required for the screening procedure described in section 5.1.3, and to the transformation $\mathbf{B} = (\mathbf{1} \otimes \bar{\mathbf{L}}^T)(\mathbf{L}^T \otimes \mathbf{1})\mathbf{B}'$ of the orthonormalized matrix \mathbf{B}' of three-center integral matrix. Though the former step of evaluating extents is typically not very computationally demanding, it scales cubically and the prefactor reduction accomplished by the parallelization will eventually be of importance for larger systems where this step constitutes a larger portion of the total computation time.

For an analysis and discussion of the parallelized RI-CDD-MP2 algorithm, tables 6.7 to 6.12 contain total calculation timings as well as for individual steps with different numbers of threads, along with the corresponding efficiencies, for the same model systems and basis sets that have been used in previous contexts. Due to the algorithm being only partially parallelized, the evaluation of efficiencies for the total calculation times may not be meaningful in a strict sense. Nevertheless, they shall serve as an indication of the influence that the speed-ups of individual parallel steps have on the total time. The rationale for including efficiencies for the FTI assembly matrix product $\mathbf{M} = \mathbf{B}\mathbf{B}^T$ is that, being based on an Intel MKL matrix multiplication routine,

they serve as reference data on parallelization efficiency.

The tendencies that reveal themselves for the various computational steps are very similar for all the different model systems with both bases. Considering the calculation of centers and extents of the transformed charge distributions, the efficiencies observed with two threads are near-ideal, usually well above 90%. There are a few exceptional cases, for instance $C_{40}H_{82}$ with cc-pVTZ (table 6.8) or Amylose₂ with SVP (table 6.9) where $\eta > 1$ seems to indicate super-linear speed-up. These might result from cache effects, or also the fact that timings are yet subject to a certain variability between different calculations on identical systems, especially for relatively short steps. At any rate, these should be considered as outliers, and they do not particularly alter the picture that emerges from the data. Increasing the thread numbers to four and eight, however, the near-ideal behaviour is lost and we observe a gradual decrease of the efficiency. Furthermore, in most cases this loss of efficiency also occurs with increase of the system size, the effect of which is most pronounced for calculations with eight threads, as may be seen for instance for linear alkanes with SVP (table 6.7). Thus, in the worst cases, this step of transforming centers and extents achieves hardly 50% efficiency for large systems.

For the transformation of the three-center integral matrix \mathbf{B}' , the determined efficiencies behave very similarly to the aforementioned situation. Even though with two threads near-ideal speed-ups are obtained, upon increasing the numbers of threads to eight the efficiency here again diminishes drastically, and in the worst cases almost drops to a mere 30%. Although, of course, the increase of the number of threads still results in a certain speed-up of execution times, except for few exceptional cases, the benefits of parallelization are not fully reaped with large threads numbers.

Turning our attention to the reference efficiencies afforded by the Intel MKL routines for the FTI assembly step $\mathbf{M} = \mathbf{B}\mathbf{B}^T$, we notice that the situation is analogous as for the previous steps. For calculations with two and four threads, the observed speed-ups are practically ideal and the corresponding efficiencies come very close to 100%. Increasing the number of threads to eight, however, efficiencies suddenly drop and in most cases vary between only 40% and 50%. These poor efficiencies registered on calculation involving eight threads are to be ascribed to the employed computers' double-quadcore architecture, i.e. two processors of four cores each. Namely, for calculations with more than four threads the communication required among both processors inevitably slows down the programme's execution, which prevents ideal efficiencies to be achieved with mere OMP parallelization. In this case, a deeper-level MPI parallelization might be required to fully exploit higher core numbers.

In light of this last argument, it can be concluded that up to the maximum meaningful number of four threads, for the computers on which calculations were performed, the parallelized RI-CDD-MP2 performs very satisfactorily. For two threads, the observed efficiencies of all parallel steps practically achieve optimal values, and calculations with four threads also display very

reasonable efficiencies of at least 80% for the systems large enough to allow for the benefits of parallelization to reveal themselves.

6.4.4 Conclusion

The overall conclusion that may be drawn from the results presented here is that the RI-CDD-MP2 algorithm of Zienau performs very well and, in its extended form for evaluation of the exchange-type component of the energy, now allows for efficient calculations of MP2 energies on systems of moderate to large size with large basis sets. Due to its quadratically scaling computational effort—for system sizes where the asymptotically cubic steps do not yet dominate—RI-CDD-MP2 could be shown to be competitive over canonical RI-MO-MP2 for a variety of model systems with basis sets up to cc-pVTZ. The errors induced in the absolute MP2 energies by the integral preselection procedure are within reasonable bounds with numerical screening thresholds that still lead to performance gains. In addition, RI-CDD-MP2 absolute energies display the attractive property of being upper bounds to the exact RI-MO-MP2 energies—at least for all systems that were investigated. Regarding relative energies, it was demonstrated for the S22 test set that even rather loose threshold combinations afford chemical accuracy, due to the mutual cancellation of errors in energy differences. In contrast to the non-RI based CDD-MP2 approach presented in the previous chapter, which was revealed to suffer from the overwhelming effort of performing the quadruple integral transformation and further displayed a pronounced numerical instability (see section 5.3.2), the overall prefactor reduction and simplification of the integral transformation brought about by the RI-approach makes RI-CDD-MP2 far more efficient and numerically stable.

In addition to being extended to total MP2 energies, a parallelized version of the RI-CDD-MP2 algorithm was presented, which was demonstrated to achieve high efficiencies, at least with thread numbers for which the particular architecture of the used computers does not create further bottlenecks (see section 6.4.3). Although the parallelization yet has to be completed for the integral evaluation steps, the speed-ups brought about by the preliminary parallelization already have a significant impact on total computation times. For four threads, total computation times are already reduced by a factor of about two, on the average. As the steps that do not allow for parallelization for reasons of thread safety have rather marginal influence on the total time, it is to be anticipated that upon finalization of the parallelization, total speed-ups will be very satisfactory. This will further contribute to making RI-CDD-MP2 applicable to large systems arising in realistic situations of interest.

6.4. RESULTS AND DISCUSSION

Basis	Molecule	N	N_{FTI}	Sc.	$t_{\text{RI-CDD}}$	Sc.	$t_{\text{RI-MO}}$	Sc.
SVP	$\text{C}_{10}\text{H}_{22}$	250	16	-	0.0	-	0.0	-
	$\text{C}_{20}\text{H}_{42}$	490	74	2.2	0.1	2.9	0.0	4.5
	$\text{C}_{40}\text{H}_{82}$	970	196	1.4	0.9	2.6	1.1	4.6
	$\text{C}_{80}\text{H}_{162}$	1930	441	1.2	4.6	2.4	31.1	4.9
	$\text{C}_{160}\text{H}_{322}$	3850	929	1.1	23.6	2.4	933.1 ^a	4.9
	Amylose ₁	228	21	-	0.0	-	0.0	-
	Amylose ₂	432	179	3.4	0.2	3.6	0.0	4.3
	Amylose ₄	840	737	2.1	2.0	3.3	0.8	4.6
	Amylose ₈	1656	1855	1.4	11.5	2.6	20.4	4.9
	Amylose ₁₆	3288	4127	1.2	59.4	2.4	566.7 ^a	4.9
	Amylose ₃₂	6552	8667	1.1	280.2	2.3	16058.7 ^a	4.9
	DNA ₁	625	218	-	0.5	-	0.2	-
	DNA ₂	1332	2184	3.0	10.5	4.0	7.1	4.8
	DNA ₄	2746	10374	2.2	143.6	3.6	232.5	4.8
	DNA ₈	5574	27950	1.4	1003.8 ^b	2.7	7085.9 ^a	4.8
	cc-pVTZ	$\text{C}_{10}\text{H}_{22}$	608	102	-	0.3	-	0.0
$\text{C}_{20}\text{H}_{42}$		1188	390	2.0	2.2	2.8	0.5	4.5
$\text{C}_{40}\text{H}_{82}$		2348	965	1.3	12.7	2.5	13.1	4.8
$\text{C}_{80}\text{H}_{162}$		4668	2106	1.1	68.0	2.4	361.0 ^a	4.8
$\text{C}_{160}\text{H}_{322}$		9308	4383	1.1	328.4	2.3	10090.8 ^a	4.8
Amylose ₁		528	142	-	0.3	-	0.0	-
Amylose ₂		998	1230	3.4	3.1	3.6	0.4	4.5
Amylose ₄		1938	5258	2.2	28.0	3.3	8.7	4.7
Amylose ₈		3818	13679	1.4	180.0	2.7	216.6 ^a	4.7
Amylose ₁₆		7578	30596	1.2	922.7 ^b	2.4	5589.1 ^a	4.7
DNA ₁	1428	1723	-	7.3	-	1.9	-	
DNA ₂	3016	16284	3.0	165.9	4.2	18.9	3.1	

^a Extrapolated from previous molecule with indicated scaling exponent

^b Extrapolated from previous molecule using scaling exponent from first quadrature point

Table 6.1: Numbers N_{FTI} of significant FTI products [in 10^6] (from first quadrature point), total Wall times [h] of RI-CDD-MP2 and conventional RI-MO-MP2 along with scaling exponents with respect to previous molecule for calculations on linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the indicated basis sets; 5 quadrature points; numerical thresholds: $\theta_{\text{ext}} = 10^{-8}$, $\theta_{\text{int}} = 10^{-8}$, $\theta_{\text{sp}} = 10^{-9}$.

Molecule	ΔE (7, 7, 9)			ΔE (8, 8, 9)			ΔE (9, 9, 10)		
	Total	Coulomb	Exchange	Total	Coulomb	Exchange	Total	Coulomb	Exchange
C ₁₀ H ₂₂	-58.97	-89.80	-30.83	-11.32	-13.95	-2.63	-2.16	-2.38	-0.22
C ₂₀ H ₄₂	-211.32	-302.89	-91.58	-58.32	-68.17	-9.85	-19.79	-20.85	-1.05
C ₄₀ H ₈₂	-520.11	-731.94	-211.83	-155.15	-179.70	-24.54	-57.36	-60.21	-2.85
C ₈₀ H ₁₆₂	-1131.71	-1583.43	-451.72	-347.84	-401.58	-53.74	-132.43	-138.85	-6.42
Amylose ₁	-27.85	-46.39	-18.54	-2.29	-3.98	-1.69	-0.44	-0.81	-0.37
Amylose ₂	-110.90	-169.14	-58.24	-11.48	-17.51	-6.02	-1.47	-2.46	-0.99
Amylose ₄	-460.87	-612.30	-151.42	-119.23	-136.08	-16.85	-34.12	-36.75	-2.63
Amylose ₈	-1326.47	-1680.65	-354.18	-452.27	-493.38	-41.11	-182.27	-188.51	-6.24
Amylose ₁₆	-3030.29	-3767.05	-736.75	-1132.37	-1219.32	-86.95	-502.62	-515.76	-13.14
DNA ₁	-294.50	-403.53	-109.03	-87.91	-106.24	-18.33	-34.70	-43.77	-9.07
DNA ₂	-1428.13	-1903.76	-475.63	-406.43	-474.57	-68.15	-138.44	-158.37	-19.93
DNA ₄	-5651.50	-6919.02	-1267.52	-2128.91	-2316.35	-187.43	-784.24	-830.49	-46.25

Table 6.2: Errors [μ Hartree] of RI-CDD-MP2 total energy as well as Coulomb and exchange parts, with respect to canonical RI-MO-MP2 reference ($\Delta E_X = E_{X,Ref}^{(2)} - E_{X,RI-CDD}^{(2)}$) for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the SVP (VDZ**) basis set; 5 quadrature points; numbers (k, l, m) indicate numerical thresholds $\theta_{ext} = 10^{-k}$, $\theta_{int} = 10^{-l}$, $\theta_{sp} = 10^{-m}$.

Molecule	ΔE (7, 7, 9)			ΔE (8, 8, 9)			ΔE (9, 9, 10)		
	Total	Coulomb	Exchange	Total	Coulomb	Exchange	Total	Coulomb	Exchange
C ₁₀ H ₂₂	-226.46	-367.44	-140.98	-41.90	-55.45	-13.55	-11.49	-14.64	-3.15
C ₂₀ H ₄₂	-705.17	-1100.75	-395.58	-163.97	-208.24	-44.27	-55.56	-64.01	-8.45
C ₄₀ H ₈₂	-1671.65	-2572.07	-900.42	-415.12	-521.11	-105.99	-147.99	-167.33	-19.34
C ₈₀ H ₁₆₂	-3578.64	-5482.38	-1903.74	-909.76	-1138.00	-228.24	-330.74	-371.85	-41.11
Amylose ₁	-125.16	-226.89	-101.73	-15.53	-27.63	-12.11	-7.06	-12.36	-5.31
Amylose ₂	-445.37	-745.14	-299.77	-57.65	-93.43	-35.78	-16.69	-28.23	-11.55
Amylose ₄	-1382.91	-2105.54	-722.64	-308.39	-399.05	-90.66	-103.09	-128.56	-25.47
Amylose ₈	-3424.43	-5024.65	-1600.23	-951.30	-1156.44	-205.15	-387.76	-442.21	-54.45
DNA ₁	-903.62	-1427.49	-523.87	-277.85	-400.50	-122.65	-164.45	-248.40	-83.95
DNA ₂	-3770.48	-5639.82	-1869.34	-1053.39	-1412.09	-358.70	-482.17	-665.19	-183.02

Table 6.3: Errors [μ Hartree] of RI-CDD-MP2 total energy as well as Coulomb and exchange parts, with respect to canonical RI-MO-MP2 reference ($\Delta E_X = E_{X,\text{Ref}}^{(2)} - E_{X,\text{RI-CDD}}^{(2)}$) for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the cc-pVTZ basis set; 5 quadrature points; numbers (k, l, m) indicate numerical thresholds $\theta_{ext} = 10^{-k}$, $\theta_{int} = 10^{-l}$, $\theta_{sp} = 10^{-m}$.

Molecule	ΔE (RI-MO-MP2)			ΔE (RI-CDD-MP2 8, 8, 9)		
	Total	Coulomb	Exchange	Total	Coulomb	Exchange
C ₁₀ H ₂₂	-843.00	-2301.14	-1458.14	-854.33	-2315.09	-1460.76
C ₂₀ H ₄₂	-1662.83	-4602.95	-2940.12	-1721.15	-4671.11	-2949.97
C ₄₀ H ₈₂	-3302.50	-9206.60	-5904.10	-3457.65	-9386.30	-5928.64
Amylose ₁	-619.15	-2752.36	-2133.21	-621.44	-2756.34	-2134.90
Amylose ₂	-1148.91	-5173.06	-4024.15	-1160.39	-5190.57	-4030.17
Amylose ₄	-2209.01	-10015.71	-7806.70	-2328.24	-10151.78	-7823.55
DNA ₁	-1556.07	-7879.29	-6323.22	-1643.98	-7985.53	-6341.55
DNA ₂	-2489.26	-16859.93	-14370.67	-2895.69	-17334.51	-14438.81

Table 6.4: Deviation [μ Hartree] of RI-MO-MP2 and RI-CDD-MP2 total energies, as well as individual components, with respect to conventional MO-MP2 reference for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the SVP (VDZ**) basis; 5 quadrature points; numerical thresholds for RI-CDD-MP2: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Molecule	ΔE (RI-MO-MP2)			ΔE (RI-CDD-MP2 8, 8, 9)		
	Total	Coulomb	Exchange	Total	Coulomb	Exchange
C ₁₀ H ₂₂	-301.63	-869.99	-568.36	-343.53	-925.44	-581.91
C ₂₀ H ₄₂	-599.67	-1740.43	-1140.76	-763.65	-1948.67	-1185.02
Amylose ₁	-227.28	-1120.02	-892.74	-242.81	-1147.65	-904.85
Amylose ₂	-421.83	-2134.36	-1712.52	-479.48	-2227.79	-1748.30
Amylose ₄	-811.00	-4163.72	-3352.73	-1119.39	-4562.78	-3443.39
DNA ₁	-646.66	-3274.24	-2627.58	-924.50	-3674.74	-2750.24

Table 6.5: Deviation [μ Hartree] of RI-MO-MP2 and RI-CDD-MP2 total energy, as well as individual components, with respect to conventional MO-MP2 reference for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the cc-pVTZ basis; 5 quadrature points; numerical thresholds for RI-CDD-MP2: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Complex	E_{int}	ΔE_{int} (SVP)			E_{int}	ΔE_{int} (cc-pVTZ)		
		5; 6, 6, 9	5; 8, 8, 9	6;10,10,12		5; 6, 6, 9	5; 8, 8, 9	6;10,10,12
01. Ammonia dimer	-2.011	0.009	0.001	0.000	-2.775	0.026	-0.002	0.000
02. Water dimer	-4.262	0.004	0.001	0.000	-4.434	0.019	-0.003	0.000
03. Formic acid dimer	-14.721	0.012	-0.005	-0.001	-16.891	0.095	-0.013	-0.003
04. Formamide dimer	-12.042	0.020	0.000	0.000	-14.248	0.096	-0.010	0.000
05. Uracil dimer (H-bonded)	-16.571	0.022	-0.050	-0.016	-18.742	0.288	-0.057	-0.021
06. 2-Pyridoxine 2-aminopyridine	-14.220	-0.021	-0.068	-0.011	-15.907	0.175	-0.101	-0.021
07. Adenine-thymine (Watson-Crick)	-13.032	0.040	-0.055	-0.012	-14.919	0.262	-0.075	-0.022
08. Methane dimer	-0.033	0.044	0.001	0.000	-0.317	0.165	-0.002	0.000
09. Ethene dimer	-0.244	0.001	-0.006	-0.001	-1.152	0.011	-0.019	-0.002
10. Benzene-methane	-0.501	0.023	-0.016	-0.003	-1.412	0.097	-0.014	-0.007
11. Benzene dimer (parallelly displaced)	-1.540	-0.031	-0.057	-0.007	-3.765	0.037	-0.112	-0.024
12. Pyrazine dimer	-3.055	-0.075	-0.093	-0.014	-5.440	-0.038	-0.144	-0.043
13. Uracil dimer (stack)	-5.701	-0.033	-0.053	-0.015	-9.095	0.129	-0.070	-0.020
14. Indole-benzene (stack)	-3.302	-0.044	-0.063	-0.008	-6.435	0.049	-0.109	-0.018
15. Adenine-thymine (stack)	-7.747	-0.015	-0.064	-0.013	-12.297	0.181	-0.091	-0.021
16. Ethene-ethyne	-0.917	0.050	-0.002	0.000	-1.427	0.066	-0.010	-0.001
17. Benzene-water	-1.964	-0.021	-0.038	-0.011	-2.927	0.023	-0.034	-0.016
18. Benzene-ammonia	-1.189	-0.002	-0.027	-0.007	-2.129	0.040	-0.021	-0.009
19. Benzene-HCN	-3.022	0.013	-0.041	-0.007	-4.601	0.054	-0.056	-0.009
20. Benzene dimer (T-shape)	-1.635	0.071	-0.054	-0.006	-3.007	0.195	-0.097	-0.019
21. Indole-benzene (T-shape)	-4.026	0.019	-0.054	-0.005	-6.159	0.135	-0.108	-0.021
22. Phenol dimer	-5.641	-0.031	-0.080	-0.014	-6.710	0.044	-0.133	-0.026
Mean error		0.003	-0.038	-0.007		0.098	-0.058	-0.014
Sample standard deviation		0.034	0.030	0.006		0.085	0.047	0.012

Table 6.6: RI-MO-MP2 interaction energies [kcal/mol] (counterpoise corrected) and errors of RI-CDD-MP2 with respect to RI-MO-MP2 reference ($\Delta E_{\text{int}} = E_{\text{int,Ref}}^{(2)} - E_{\text{int,RI-CDD}}^{(2)}$), along with mean errors and standard deviations, for the S22 test set with the SVP and cc-pVTZ basis sets and the indicated numerical settings. The numbers $m; n, o, p$ stand for m quadrature points and thresholds $\theta_{\text{ext}} = 10^{-n}$, $\theta_{\text{int}} = 10^{-o}$, $\theta_{\text{sp}} = 10^{-p}$, respectively.

Molecule	TOTAL							Centers+extents						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
C ₂₀ H ₄₂	536	401	332	302	0.67	0.40	0.22	28	15	9	5	0.93	0.78	0.70
C ₄₀ H ₈₂	3160	2248	1921	1911	0.70	0.41	0.21	264	135	79	51	0.98	0.84	0.65
C ₈₀ H ₁₆₂	16479	11022	9118	8574	0.75	0.45	0.24	1394	730	454	252	0.95	0.77	0.69
C ₁₆₀ H ₃₂₂	85041	61919	47148	45423	0.69	0.45	0.23	4610	2376	1549	1056	0.97	0.74	0.55
	Transf. $\mathbf{B}' \rightarrow \mathbf{B}$							$\mathbf{M} = \mathbf{B}\mathbf{B}^T$						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
C ₂₀ H ₄₂	56	28	15	17	1.00	0.93	0.41	186	107	56	37	0.87	0.83	0.63
C ₄₀ H ₈₂	334	153	95	82	1.09	0.88	0.51	1197	650	328	368	0.92	0.91	0.41
C ₈₀ H ₁₆₂	1624	775	478	377	1.05	0.85	0.54	5736	2969	1494	1662	0.97	0.96	0.43
C ₁₆₀ H ₃₂₂	8944	6132	3307	2852	0.73	0.68	0.39	24388	12860	6581	7012	0.95	0.93	0.43

Table 6.7: Wall times [s] of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm with 1, 2, 4 and 8 threads, respectively, along with corresponding efficiencies η (see text), for calculations on linear alkanes with the SVP (VDZ**) basis set; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Molecule	TOTAL							Centers+extents						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
C ₁₀ H ₂₂	1219	1087	947	943	0.56	0.32	0.16	21	12	5	4	0.88	1.05	0.66
C ₂₀ H ₄₂	8062	6193	5257	5270	0.65	0.38	0.19	240	124	66	43	0.97	0.91	0.70
C ₄₀ H ₈₂	45741	34740	27143	26398	0.66	0.42	0.22	2938	1305	687	492	1.13	1.07	0.75
C ₈₀ H ₁₆₂	244908	173958	140916	137708	0.70	0.43	0.22	22492	11114	6120	4533	1.01	0.92	0.62
C ₁₆₀ H ₃₂₂	1182154	875283	696396	663788	0.68	0.42	0.22	89802	50884	30641	25001	0.88	0.73	0.45
	Transf. $\mathbf{B}' \rightarrow \mathbf{B}$							$\mathbf{M} = \mathbf{B}\mathbf{B}^T$						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
C ₁₀ H ₂₂	80	47	39	20	0.85	0.51	0.50	200	94	45	59	1.06	1.11	0.42
C ₂₀ H ₄₂	666	293	179	126	1.14	0.93	0.66	2495	1301	664	591	0.96	0.94	0.53
C ₄₀ H ₈₂	4179	3554	1413	896	0.59	0.74	0.58	14438	7401	3813	4134	0.98	0.95	0.44
C ₈₀ H ₁₆₂	24373	13597	8613	7228	0.90	0.71	0.42	67024	34212	17881	18576	0.98	0.94	0.45
C ₁₆₀ H ₃₂₂	86451	45514	31139	31663	0.95	0.69	0.34	289434	150489	78920	80457	0.96	0.92	0.45

Table 6.8: Wall times [s] of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm with 1, 2, 4 and 8 threads, respectively, along with corresponding efficiencies η (see text), for calculations on linear alkanes with the cc-pVTZ basis set; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Molecule	TOTAL							Centers+extents						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
Amylose ₂	828	625	543	462	0.66	0.38	0.22	25	13	6	4	0.96	1.04	0.78
Amylose ₄	7191	4776	3704	3679	0.75	0.49	0.24	292	122	74	47	1.20	0.99	0.78
Amylose ₈	41489	26642	20132	19535	0.78	0.52	0.27	2437	1278	656	461	0.95	0.93	0.66
Amylose ₁₆	213809	144425	104625	103994	0.74	0.51	0.26	14382	7686	4160	2962	0.94	0.86	0.61
Amylose ₃₂ ^a	205936	130070	93004	105166	0.79	0.55	0.24	9517	5146	2981	2126	0.92	0.80	0.56
	Transf. $\mathbf{B}' \rightarrow \mathbf{B}$							$\mathbf{M} = \mathbf{B}\mathbf{B}^T$						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
Amylose ₂	78	39	24	19	1.00	0.81	0.51	334	171	95	64	0.98	0.88	0.65
Amylose ₄	599	258	165	143	1.16	0.91	0.52	3872	2006	995	1182	0.97	0.97	0.41
Amylose ₈	3250	1549	1020	760	1.05	0.80	0.53	22886	11688	6025	6038	0.98	0.95	0.47
Amylose ₁₆	17926	9471	5487	4031	0.95	0.82	0.56	106385	55199	28285	30969	0.96	0.94	0.43
Amylose ₃₂ ^a	16842	9213	4683	5460	0.91	0.90	0.39	140956	72092	36875	39427	0.98	0.96	0.45

^a Calculation times for 1st quadrature point only

Table 6.9: Wall times [s] of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm with 1, 2, 4 and 8 threads, respectively, along with corresponding efficiencies η (see text), for calculations on α -D-glucose chains (amylose) with the SVP (VDZ**) basis set; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Molecule	TOTAL							Centers+extents						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
Amylose ₂	11033	8576	7039	6740	0.64	0.39	0.20	180	93	47	36	0.97	0.96	0.63
Amylose ₄	100661	67986	50668	48098	0.74	0.50	0.26	2063	1073	541	399	0.96	0.95	0.65
Amylose ₈	648056	429442	322118	323384	0.75	0.50	0.25	27299	12884	6978	4769	1.06	0.98	0.72
Amylose ₁₆ ^a	684616	417941	279264	268386	0.82	0.61	0.32	32798	17821	9845	7702	0.92	0.83	0.53
	Transf. $\mathbf{B}' \rightarrow \mathbf{B}$							$\mathbf{M} = \mathbf{B}\mathbf{B}^T$						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
Amylose ₂	840	728	324	179	0.58	0.65	0.59	4056	2103	1080	1158	0.96	0.94	0.44
Amylose ₄	6931	5581	2520	1739	0.62	0.69	0.50	57636	29328	15111	14018	0.98	0.95	0.51
Amylose ₈	37328	23626	12831	9984	0.79	0.73	0.47	378803	193985	100075	104124	0.98	0.95	0.45
Amylose ₁₆ ^a	41768	19831	11991	7338	1.05	0.87	0.71	523166	266165	138194	144562	0.98	0.95	0.45

^a Calculation times for 1st quadrature point only

Table 6.10: Wall times [s] of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm with 1, 2, 4 and 8 threads, respectively, along with corresponding efficiencies η (see text), for calculations on α -D-glucose chains (amylose) with the cc-pVTZ basis set; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Molecule	TOTAL							Centers+extents						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
DNA ₁	1837	1316	1089	919	0.70	0.42	0.25	80	43	22	12	0.93	0.91	0.83
DNA ₂	37901	24182	17775	16904	0.78	0.53	0.28	1451	769	393	243	0.94	0.92	0.75
DNA ₄	516862	318530	204776	206441	0.81	0.63	0.31	26618	13867	7092	4921	0.96	0.94	0.68
DNA ₈ ^a	1060398	641176	426311	438397	0.83	0.62	0.30	61715	33524	18034	13511	0.92	0.86	0.57
	Transf. $\mathbf{B}' \rightarrow \mathbf{B}$							$\mathbf{M} = \mathbf{B}\mathbf{B}^T$						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
DNA ₁	174	74	79	47	1.18	0.55	0.46	778	417	212	151	0.93	0.92	0.64
DNA ₂	2622	1237	824	626	1.06	0.80	0.52	23480	12065	6208	5962	0.97	0.95	0.49
DNA ₄	25272	13682	8016	5422	0.92	0.79	0.58	356296	183412	93369	100283	0.97	0.95	0.44
DNA ₈ ^a	37307	17244	10798	7282	1.08	0.86	0.64	849069	435220	225165	244676	0.98	0.94	0.43

^a Calculation times for 1st quadrature point only

Table 6.11: Wall times [s] of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm with 1, 2, 4 and 8 threads, respectively, along with corresponding efficiencies η (see text), for calculations on DNA bases pairs (adenine-thymine) with the SVP (VDZ**) basis set; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Molecule	TOTAL							Centers+extents						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
DNA ₁	26286	18676	14567	14355	0.70	0.45	0.23	599	312	158	116	0.96	0.95	0.65
DNA ₂	597397	375576	256807	265743	0.80	0.58	0.28	11441	5926	3050	2154	0.97	0.94	0.66
	Transf. $\mathbf{B}' \rightarrow \mathbf{B}$							$\mathbf{M} = \mathbf{B}\mathbf{B}^T$						
	t_1	t_2	t_4	t_8	η_2	η_4	η_8	t_1	t_2	t_4	t_8	η_2	η_4	η_8
DNA ₁	2184	1770	899	486	0.62	0.61	0.56	12193	6192	3182	3538	0.98	0.96	0.43
DNA ₂	33114	18968	10161	7332	0.87	0.81	0.56	395751	201233	102947	111523	0.98	0.96	0.44

Table 6.12: Wall times [s] of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm with 1, 2, 4 and 8 threads, respectively, along with corresponding efficiencies η (see text), for calculations on DNA bases pairs (adenine-thymine) with the cc-pVTZ basis set; 5 quadrature points; numerical thresholds: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$.

Chapter 7

Method Specific Cholesky Decomposition

Apart from RI-approximations, which have already been discussed in the previous chapter, there has been an increasing focus on the Cholesky decomposition (CD) to obtain low-rank decompositions of the ERI matrix in recent years. In a secondary project, it has been investigated to what extent the method specific Cholesky decomposition (MSCD) formulated by Boman *et al.* [19] in general terms but investigated only in the context of HF-theory, might be beneficially applied to correlation methods, and specifically within the CDD-MP2 approach. The conceptual basis of MSCD is the input of information about the final form the ERI matrix is to be used in, for instance transformed to a basis of localized molecular orbitals (LMO), rather than in its untransformed form, in order to obtain a Cholesky decomposition with minimum rank.

In order to set this work into the appropriate context, this chapter begins with a brief survey of various ERI decomposition procedures that have been proposed, though with particular emphasis on the relation between RI and CD. We then proceed with a presentation of the MSCD and how it can be applied in the specific context of CDD-MP2, after which the obtained results are discussed.

7.1 Decompositions of the ERI matrix

Due to the fact that the evaluation or handling of the matrix of ERI or wavefunction amplitudes constitutes one of the bottlenecks of many quantum chemical methods, this aspect has been subject to intensive research aiming at reducing its computational cost by applying matrix/tensor decomposition procedures. However, although they are related by their common objective, the methods that have been devised to that effect differ by their approach and scope of applicability. The procedure of resolution of the identity (RI) and density fitting (DF) has already been described in section 6.1. Further approaches advocate the use of a Cholesky decomposition of the

ERI matrix, or propose higher order tensor decompositions for further reductions of complexity and storage requirements.

7.1.1 Resolution of the Identity

It has been discussed in section 6.1 that within the so-called resolution of the identity (RI) [14–18] a lower-rank decomposition to the ERI matrix \mathbf{A} can be straightforwardly obtained through the introduction of an auxiliary basis $\{\xi_q(\mathbf{r})\}_{q \in N_{\text{aux}}}$ and application of the closure relation

$$\hat{\rho}_\xi = |\xi\rangle\langle\xi| = 1, \quad (7.1)$$

leading to a factorization of the form

$$\mathbf{A} = \left((\boldsymbol{\Omega}|\xi)_C \boldsymbol{\Xi}_C^{-\frac{1}{2}} \right) \left(\boldsymbol{\Xi}_C^{-\frac{1}{2}} (\xi|\boldsymbol{\Omega})_C \right) = \mathbf{B}'\mathbf{B}'^T. \quad (7.2)$$

Recall that the $N_{\text{aux}} \times N_{\text{aux}}$ matrix $\boldsymbol{\Xi}_C^{-\frac{1}{2}}$ denotes the inverse square root of the overlap matrix $\boldsymbol{\Xi}_C$ of the auxiliary basis functions with respect to the Coulombic repulsion operator \hat{C} . The $N^2 \times N_{\text{aux}}$ matrix $(\boldsymbol{\Omega}|\xi)_C$ (ignoring integral symmetry) is composed of the three-center two-electron repulsion integrals, i.e. has elements

$$\left((\boldsymbol{\Omega}|\xi)_C \right)_{\mu\nu,p} = (\Omega_{\mu\nu}|\xi_p) = \int \int \Omega_{\mu\nu}(\mathbf{r}_1) \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \xi_p(\mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2. \quad (7.3)$$

RI-techniques are nowadays routinely employed due to the high performance improvements that they provide and their ease of implementation. Despite this wide success, however, they are not entirely flawless. On the one hand, the above closure relation cannot be strictly fulfilled due to the incompleteness of the auxiliary basis set. The error introduced by this approximation is in practice reasonable, and in general cancels to a large extent upon formation of relative energies. Nevertheless, this impairs the reliability of RI-calculated properties for use in basis set limit extrapolation procedures, as the RI-error may be larger than the extrapolation accuracy [19]. Furthermore, the auxiliary basis sets, which typically consist of atom-centered Gaussian functions, have to be tailored *a priori* for the particular AO basis set as well as method that they are to be used with. They are thus biased by construction, similarly to conventional LCAO basis sets, and of limited applicability outside the original scope of their preoptimization.

7.1.2 Cholesky decomposition

Due to linear near-dependence among the $\Omega_{\mu\nu}$ charge distributions, the ERI matrix \mathbf{A} is only positive semidefinite. In a seminal paper from 1977, Beebe and Linderberg [123] proposed to exploit this property to obtain a low rank factorization from a Cholesky decomposition (CD) [9], according to

$$\mathbf{A} = \mathbf{B}'\mathbf{B}'^T. \quad (7.4)$$

As has already been discussed in section 5.1.1 for the pseudo-density matrices, the CD of a semidefinite matrix requires a pivoting technique (see section 5.2.1). The obtained Cholesky factors \mathbf{B}' are therefore $N^2 \times R$ matrices (ignoring integral symmetry), where R is the (numerical) rank of the ERI matrix, i.e. the number of non-vanishing eigenvalues. In contrast to other factorization techniques, the Cholesky factor has the significant advantage of being built column by column, and thus does not require the entire input matrix to be evaluated beforehand.

The reason for using the same symbol \mathbf{B}' as for the RI-approximation is that, as has been shown by Beebe and Linderberg in their original paper [123], the Cholesky decomposition and RI are in principle related, because they both represent *inner projections* in the sense of Löwdin [132]. The fundamental difference, however, is that in contrast to RI, where an external, preoptimized auxiliary basis is used for the projection, the Cholesky decomposition produces an auxiliary basis set that is intrinsic to the considered space, and therefore unbiased. Furthermore, whereas in the RI-case a non-tunable systematic error arises due to the incompleteness of the auxiliary basis, the pivot-CD can be made arbitrarily precise by adjusting the decomposition threshold. Due to these attractive features, the Cholesky decomposition of the ERI matrix has been amply investigated [133–139] for application within Hartree-Fock schemes [140] or correlation methods [141], and further for the analytical evaluation of energy gradients [134, 142]. For a comparison between RI and CD methods, see also reference [143].

One arguable drawback of CD, however, is that with relatively tight decomposition thresholds, the Cholesky auxiliary bases are normally larger than comparable RI basis sets. Therefore, in order to obtain a particularly small Cholesky basis, some of the cited schemes advocate the incorporation of only atom-centered distributions $\Omega_{\mu\nu}$ in a first decomposition step, after which the obtained auxiliary basis may be further refined through inclusion of further distributions. This procedure reestablishes a certain similarity to RI, where the employed functions are also typically atom-centered. Yet, it has to be remarked that although the errors induced by this approach were shown to be of reasonable magnitude, and can in principle be made arbitrarily small by gradually reintroducing all distributions, such a truncation of the Cholesky basis forsakes numerical rigour in the sense that the errors are then not bound by the decomposition threshold any more.

7.1.3 Tensor approximation

Alternatively to the aforementioned procedures, which are most conveniently viewed as matrix decompositions, there also exist approaches that aim at higher order decomposition of the ERI tensor or other wavefunction parameters. For instance, it has recently been proposed by Benedikt *et al.* [24] that an efficient reduction of storage and transformation effort involved with the handling of ERI can be achieved through an expansion of \mathbf{A} using representing vectors in the

canonical product (CP) format

$$\mathbf{A} = \sum_{r=1}^R \mathbf{a}_r^{(\mu)} \otimes \mathbf{a}_r^{(\nu)} \otimes \mathbf{a}_r^{(\lambda)} \otimes \mathbf{a}_r^{(\sigma)}, \quad (7.5)$$

such that the computational cost is reduced from $\mathcal{O}(N^4)$ to $\mathcal{O}(d \cdot R \cdot N)$, where d stands for the dimension of the tensor, i.e. 4 in the case of ERI, and R is the decomposition rank. The superscripts on the $\mathbf{a}_r^{(\mu)}$ are not indices, but serve to denote the dimension to which the different representing vectors belong. The primary advantage of this decomposition scheme is its generality, which makes it extendable to any other tensor, such as amplitudes, that may occur in quantum chemical situations. For illustration, within the Laplace transform approach to MP2 (see chapter 4), which allows to factorize the denominator of orbital energies, this format in principle enables an efficient reformulation of the integral transformation and contraction [24]. However, this procedure at present requires the computation of the entire input tensor \mathbf{A} , followed by an iterative reduction of the initial rank of a trivial CP decomposition in order to approximate it within a predetermined accuracy. Although significant rank reductions which still afford reasonable accuracy have been reported, this method may so far only be beneficially applied within higher correlation methods such as coupled cluster, whose complexity is not determined by the $\mathcal{O}(N^4)$ evaluation and storage of the ERI.

On a very similar footing, as an aside, a proposition is put forth in appendix B how a matrix Kronecker product approximation (KPA) of the ERI matrix, rather than a CP approximation, might also prove favourable for the calculation of the MP2 energy.

7.2 Method specific Cholesky decomposition (MSCD)

In most applications, the ERI matrix \mathbf{A} is not required in raw, untransformed form, but it rather occurs either as an intermediary in the formation of the Coulomb and exchange matrices \mathbf{J} and \mathbf{K} constituting the Fock matrix of HF-theory, or transformed to some particular basis for correlation methods. With this in mind, Boman *et al.* [19] proposed to base the Cholesky decomposition of \mathbf{A} not on its diagonal itself as pivot elements, but on some specific *characteristic matrix* $\mathbf{\Omega}$, i.e. functional of \mathbf{A} , that describes the particular situation one is confronted with. The rationale behind this approach, termed method specific Cholesky decomposition (MSCD), is to obtain a Cholesky decomposition of minimum rank, through the screening induced by the transformation procedure on \mathbf{A} , while retaining arbitrary accuracy of the final expression that is to be evaluated through the sole parameter of the decomposition threshold.

Now, within the CDD-MP2 approach described in chapter 5, and in chapter 6 for its RI-based variant, the ERI matrix needs to be transformed to the basis of Cholesky orbitals according to

$$\mathbf{M} = (\underline{\mathbf{L}}^T \otimes \bar{\underline{\mathbf{L}}}^T) \mathbf{A} (\underline{\mathbf{L}} \otimes \bar{\underline{\mathbf{L}}}), \quad (7.6)$$

where, recalling the definitions given in section 4.1, the matrix \mathbf{M} has elements $M_{ia,jb} = (\underline{i}\bar{a}|\underline{j}\bar{b})$. However, the matrix that has been used as basis for the application of MSCD in the present context is the half-transformed ERI matrix

$$\mathbf{A}_{\text{HT}} = (\underline{\mathbf{L}}^T \otimes \mathbf{1})\mathbf{A}(\underline{\mathbf{L}} \otimes \mathbf{1}), \quad (7.7)$$

i.e. where both occupied indices have been transformed according to $(\underline{i}\nu|\underline{j}\sigma)$. The characteristic matrix whose diagonal elements are used for pivoting in the Cholesky decomposition is in this case given by

$$\mathbf{\Omega} \equiv \left(\underline{L}_{\mu i}(\mu\nu|\lambda\sigma)\underline{L}_{\lambda j} \right)_{\substack{i,j \in N_{\text{occ}} \\ \mu,\nu,\lambda,\sigma \in N}}. \quad (7.8)$$

This procedure thus produces a $N^2 \times R_{\text{MSCD}}$ Cholesky decomposition

$$\mathbf{A} = \mathbf{B}'\mathbf{B}'^T \quad (7.9)$$

for the ERI matrix \mathbf{A} , where the dimension R_{MSCD} of the Cholesky factor, however, is the numerical rank of the characteristic matrix $\mathbf{\Omega}$, which is itself determined by the decomposition threshold θ_{MSCD} . Since for systems where the electronic density is sufficiently localized $\underline{\mathbf{L}}$ represents LMO, as has been discussed for CDD-MP2 in sections 5.1.2 and 5.3.1, the screening based on the above characteristic matrix may be anticipated to lead to a significant rank reduction. After transformation of the reduced rank Cholesky factor \mathbf{B}' according to

$$\tilde{\mathbf{B}} = (\underline{\mathbf{L}}^T \otimes \mathbf{1})\mathbf{B}', \quad (7.10)$$

the half-transformed ERI matrix can then be constructed with minimal effort as

$$\mathbf{A}_{\text{HT}} = \tilde{\mathbf{B}}\tilde{\mathbf{B}}^T \quad (7.11)$$

and with a numerical accuracy determined by θ_{MSCD} . The reason for having based the MSCD on the half-transformed ERI matrix \mathbf{A}_{HT} instead of the fully transformed one \mathbf{M} is that, as had already been pointed out by Boman *et al.* [19], and as was confirmed by our initial investigations in this context, the structure of the characteristic matrix

$$\mathbf{\Omega} \equiv \left(\underline{L}_{\mu i}\bar{L}_{\nu a}(\mu\nu|\lambda\sigma)\underline{L}_{\lambda j}\bar{L}_{\sigma b} \right)_{\substack{i,j \in N_{\text{occ}} \\ a,b \in N_{\text{virt}} \\ \mu,\nu,\lambda,\sigma \in N}}. \quad (7.12)$$

for the MSCD of \mathbf{M} does not lead to a favourable rank reduction. Still, the half-transformed integral matrix \mathbf{A}_{HT} thus obtained can then be straightforwardly used for further transformations. Alternatively, a second transformation can be directly carried out as $\mathbf{B} = (\mathbf{1} \otimes \bar{\mathbf{L}}^T)\tilde{\mathbf{B}}$, after which the FTI matrix can be assembled identically to RI-CDD-MP2 as $\mathbf{M} = \mathbf{B}\mathbf{B}^T$.

7.2.1 Technical details

In spite of the aforementioned property that the Cholesky decomposition does not necessarily require the evaluation of the entire matrix that is to be decomposed, an implementation of this procedure for selective, column-wise formation of the ERI matrix with full exploitation of integral symmetry for successive build-up of the Cholesky factor was beyond the scope of this project. For the preliminary investigations that were performed on this subject, the MSCD has been implemented within a modified version of the CD algorithm that is used for decomposition of the pseudo-density matrices (see sections 5.1.1 and 5.2.1), after adaptation to the pivoting induced by the characteristic matrix given in equation (7.8). Since this non-optimized algorithm requires the input matrix, i.e. the $N^2 \times N^2$ ERI matrix \mathbf{A} , to be formed in its entirety, investigations of this method were possible only on very small systems, due to the ensuing tremendous increase of memory requirements, especially with the large basis set cc-pVTZ.

In order to analyze the potential of the MSCD as applied to the present situation, the numerical ranks obtained from a MSCD were compared to those of a standard CD of the ERI matrix for the series of linear alkanes with a variety of basis sets. The particular decomposition thresholds for both decompositions have been defined as follows. For the standard CD, in analogy to the threshold used in section 5.2.1, θ_{CD} has been chosen as

$$\theta_{\text{CD}} = \mathbf{A}_{\text{max}} \cdot 10^{-9}, \quad (7.13)$$

where \mathbf{A}_{max} denotes the maximum (diagonal) element of \mathbf{A} . This yields a precision of 10^{-9} Hartree of the individual elements of \mathbf{A} as reconstructed from the Cholesky factor hereby obtained. The threshold θ_{MSCD} applied to the diagonal elements of $\mathbf{\Omega}$, on the other hand, has been defined as

$$\theta_{\text{MSCD}} = \mathbf{A}_{\text{max}} \cdot 10^{-8} / \underline{\mathbf{L}}_{\text{max}}^2, \quad (7.14)$$

so that, in view of the transformation (7.10), the individual elements of \mathbf{A}_{HT} may be recomposed with an accuracy of 10^{-8} Hartree. As in section 5.2.1, $\underline{\mathbf{L}}_{\text{max}}$ here denotes the maximum element of the occupied Cholesky factor matrix $\underline{\mathbf{L}}$, but in contrast to (RI)-CDD-MP2, $\underline{\mathbf{L}}$ is here not scaled with the maximum element of the virtual Cholesky factor.

The reason for imposing different target precisions is that, for further use in some transformation procedure, the raw matrix \mathbf{A} ought to be constructed at least with the aforementioned accuracy, whereas from the experience gathered in the context of RI-CDD-MP2, a precision of 10^{-8} should be sufficient for the elements of the half-transformed matrix \mathbf{A}_{HT} .

7.2.2 Results and discussion

The numerical ranks R_{CD} and R_{MSCD} determined as described in the previous section are listed in table 7.1, along with the maximum possible rank $R_{\text{max}} = N(N + 1)/2$ (in case of

linear independence, when integral symmetry is fully exploited), for illustration of the advantage brought by both standard CD and MSCD. Namely, whereas the formal maximum rank grows quadratically, it has already been amply documented (see references given in section 7.1.2) that the linear dependence in the ERI matrix causes the numerical rank to scale linearly. In addition to the numerical rank data, the Frobenius norm $\|\Delta\|_F$ of the difference matrix

$$\Delta = \mathbf{A}_{\text{HT}} - \tilde{\mathbf{B}}\tilde{\mathbf{B}}^T \quad (7.15)$$

between the exact HTI matrix and as recomposed from the transformed MSCD factor matrix is given in table 7.1, as a measure of the overall accuracy afforded by the low-rank approximation. Further, for comparison with the RI-approximation, the size N_{aux} of the corresponding auxiliary basis has been indicated for the basis sets SVP and cc-pVTZ (STO-3G and 6-31G* do not possess an RI auxiliary basis for MP2).

From the behaviour of the numerical ranks with different basis sets, it becomes apparent that the benefits of MSCD over standard CD are most significant for large bases. For the STO-3G minimal basis, the difference between R_{CD} and R_{MSCD} is rather small for the initial molecular sizes, and even vanishes for the smallest molecule CH_4 . Although this difference becomes more pronounced towards larger alkanes, MSCD here doesn't achieve a reduction of R_{CD} beyond a factor of about 1.6 for the largest system $\text{C}_{30}\text{H}_{62}$. With the next larger 6-31G* basis, the situation is already noticeably different and for $\text{C}_{10}\text{H}_{22}$, a rank reduction by a factor 2.4 is obtained. Increasing the basis set size, the rank reduction brought about by MSCD progressively gains further importance. With SVP and cc-pVTZ, for the largest molecule that could be investigated in each case, the rank is reduced by a factor 2.5 and even 3.7, respectively. Considering that in these last cases, the calculations were restricted to very small molecules, it may be surmised that for larger molecular sizes the effects of the MO locality will further enhance the observed rank reduction. Turning our attention to the error measure provided by the Frobenius norm $\|\Delta\|_F$ of the error matrix, we notice that in all cases the HTI matrix is recomposed with a very high and well-controlled overall accuracy.

Finally, a comparison of R_{MSCD} with N_{aux} yields that the RI-approximation yet provides a lower-rank decomposition of \mathbf{A} in all cases both for SVP and cc-pVTZ. However, as has already been alluded to in section 7.1.2, this result has to be weighed against the fact that the incompleteness of the RI auxiliary basis induces a fixed, non-tunable error in the approximated entity. In contrast to this, the precision provided by MSCD can be controlled through the decomposition threshold. It may therefore be argued that, though MSCD would come at higher computational cost, it offers the possibility to compute precise integrals, and thus absolute energies. On the other hand, if one were satisfied with smaller accuracy, the decomposition threshold θ_{MSCD} may also be chosen so as to make R_{MSCD} comparable to N_{aux} , such that MSCD can still be competitive with RI-approaches. One could also consider the application of truncation schemes of the

Cholesky auxiliary basis, which have been mentioned previously. At any rate, the rank reduction of MSCD as compared to standard CD is rather encouraging, and it might be interesting to pursue developments in this direction.

Molecule	STO-3G				SVP				
	R_{\max}	R_{CD}	R_{MSCD}	$\ \Delta\ _F$	R_{\max}	R_{CD}	R_{MSCD}	$\ \Delta\ _F$	N_{aux}
CH ₄	45	45	45	0.0	595	454	207	6.3	104
C ₂ H ₆	136	123	109	0.6	1711	836	347	14.8	180
C ₃ H ₈	276	218	167	1.5	3403	1206	489	23.8	256
C ₄ H ₁₀	465	311	224	1.9	5671	1563	627	31.0	332
C ₅ H ₁₂	703	404	278	3.0	8515	1931	763	36.5	408
C ₁₀ H ₂₂	2628	869	559	4.7	-	-	-	-	-
C ₁₅ H ₃₂	5778	1334	836	6.8	-	-	-	-	-
C ₂₀ H ₄₂	10153	1799	1112	8.1	-	-	-	-	-
C ₃₀ H ₆₂	22578	2729	1663	11.0	-	-	-	-	-
	6-31G*				cc-pVTZ				
	R_{\max}	R_{CD}	R_{MSCD}	$\ \Delta\ _F$	R_{\max}	R_{CD}	R_{MSCD}	$\ \Delta\ _F$	N_{aux}
CH ₄	276	235	131	3.1	3741	1095	302	22.2	201
C ₂ H ₆	903	524	242	14.3	10440	1864	506	52.7	342
C ₃ H ₈	1891	789	347	23.5	20503	2617	703	80.9	483
C ₄ H ₁₀	3240	1056	453	26.6	-	-	-	-	-
C ₅ H ₁₂	4950	1317	553	31.8	-	-	-	-	-
C ₁₀ H ₂₂	18915	2616	1077	50.9	-	-	-	-	-

Table 7.1: Numerical ranks of ERI matrix \mathbf{A} (theoretical maximum R_{\max} , standard CD R_{CD} and method specific CD R_{MSCD}), along with Frobenius norm $\|\Delta\|_F$ [$\mu\text{Hartree}$] of error matrix Δ between exact HTI matrix \mathbf{A}_{HT} and recomposed from MSCD as $\tilde{\mathbf{B}}\tilde{\mathbf{B}}^T$, and size N_{aux} of auxiliary RI-basis for SVP and cc-pVTZ. See discussion in section 7.2.1 for settings of the individual decompositions.

Chapter 8

Conclusion and outlook

The main topic of this dissertation has been the implementation of the CDD-MP2 method, the investigation of which was motivated by a series of promising theoretical features that endow it with the potential to scale linearly, and thus be useful for the evaluation of MP2 energies of large molecules. In particular, the CDD approach has the advantage of yielding full MP2 energies, rather than merely the opposite spin component of it, to which the AO-MP2 implementation that forms its basis is currently restricted for technical reasons. However, although it could be shown on ideal model systems that the CDD-MP2 approach does in principle scale linearly, its practical applicability to molecular systems of sensible size with large basis sets was severely limited and it could not be established as being competitive—both in terms of computational performance and reasonable accuracy—to canonical MP2 for fairly large systems. Namely, although the employed integral preselection scheme does achieve considerable reductions of the numbers of significant contributions to the MP2 energy in themselves, the formation of these contributions through four consecutive transformations is associated with an overwhelming computational effort that could not be efficiently reduced, in spite of an internal screening procedure that minimizes the evaluation cost of the preselected terms by discarding insignificant intermediaries in the formation. In addition to this, the very formulation of the Coulomb-type MP2 energy component in terms of purely positive summands within this CDD approach, in conjunction with the fact that the internal screening induces a certain error in the evaluated integrals, leads to a high numerical imbalance with practicable numerical settings and large basis sets. Accordingly, overall accuracies are rather poor even on relatively small molecules, such that it doesn't appear sensible to continue research on this specific integral-direct CDD-MP2 approach, since improvements merely of the computational performance will not help in coming to terms with the numerical problems.

On the other hand, the RI-CDD-MP2 implementation of Zienau *et al.*, a modification of CDD-MP2 which is based on RI-approximation of the required integrals, has been extended to include the exchange-type term and is now also capable of producing total MP2 energies. Due to its

near-quadratic scaling, this method was revealed as advantageous over fifth-order scaling canonical RI-MP2 for a variety of large test systems with basis sets up to cc-pVTZ. The entirely different integral transformation procedure also causes this approach not to be affected by the numerical problems that plague CDD-MP2, and it was shown to provide very satisfying accuracy with practicable screening thresholds. Apart from the complementation to full MP2, a preliminary OMP parallelized version of this algorithm has been presented, which attains relatively high efficiencies. Although this parallelization yet has to be completed for all steps involving integral generation, the speed-ups brought about by the partially parallel algorithm were shown to already have a substantial impact on the total calculation times, as these were on the average halved in calculations with four threads. The RI-CDD-MP2 method thus represents a serious alternative to canonical RI-MP2 for the evaluation of MP2 energies of moderate to large systems with large bases, and is anticipated to further gain efficiency upon completion of the parallelization.

As an alternative approach to RI-approximated integrals, it has also been explored whether a method specific Cholesky decomposition of the ERI matrix can be fruitfully applied in the CDD context. For implementational reasons, the preliminary investigations on this subject had to be restricted to very small molecules, but significant rank reductions with respect to standard Cholesky decomposition were already observed in these cases. In particular, these reductions were most pronounced with large basis sets. Even though MSCD ranks were still larger than corresponding RI auxiliary bases, this approach presents the advantage of reproducing the approximated entity with arbitrary precision, in contrast to the basis set incompleteness error that affects RI methods. Due to these encouraging results, it would seem worthwhile to pursue efforts along this line and work out an optimized MSCD implementation for selective column formation and full use of integral symmetry in order to study the method's potential also on large systems.

As a concluding retrospective and outlook, it may be said that over the last decades, remarkable progress has been made in the treatment of electron correlation, and a wide array of wavefunction based methods, ranging from coupled cluster approaches to low-scaling perturbational schemes is nowadays at our disposal to provide quantum theoretical treatment of many of the questions arising in contemporary research applications. Yet, the difficulties with which quantum chemistry is faced in that respect are still many and the evaluation of electron correlation energies with higher precision or for very large systems remains as challenging as ever. In the high-accuracy domain, the currently established methods such as coupled cluster are still limited to rather small molecules, due to their high computational cost. On the other hand, even lower-level methods will have to be made more efficient than they are at present in order to be routinely applicable for the investigation of systems that occur in biochemistry and similar situations. Certainly, computational possibilities will continue to benefit from the ongoing tremendous im-

provements of hardware, so that larger systems will progressively become tractable with current methods and algorithms. As concerns these developments, however, the growth rate of processor frequency has somewhat decreased in recent times in favour of a stronger emphasis on parallel computing, such that it will be important to follow this trend and adapt quantum chemical software for that purpose. Nevertheless, despite the undoubtable impact of hardware improvements, the predominant driving force of quantum chemical progress has always been and will remain methodical innovation and the development of new, more efficient algorithms. It is therefore indispensable that purely technical advances go hand in hand with method-theoretical research. In that respect, however, it is hard to predict how the presently applied techniques will evolve and from which direction significant advances will be stimulated. One may even be tempted to the somewhat controversial admission that however valuable the currently established algebraic procedures have been thus far, they also seem to lead into further and further approximations and increasingly tricky algorithms. One could therefore speculate that a fundamentally different mathematical approach will eventually have to emerge in order to provide new possibilities for the formidable endeavour of solving Schrödinger's equation.

List of Tables

5.1	Numbers of significant elements in the scaled Cholesky factors $\underline{\mathbf{L}}$ and $\bar{\mathbf{L}}$ and screening matrix \mathbf{Z}	67
5.2	Numbers of significant elements in the pseudo-density matrices $\underline{\mathbf{P}}$ and $\bar{\mathbf{P}}$	68
5.3	Numbers of integrals and numbers of (ij) -pairs for linear alkanes, along with scaling exponents with respect to previous molecule.	69
5.4	Numbers of integrals and numbers of (ij) -pairs for α -D-glucose chains (amylose), along with scaling exponents with respect to previous molecule.	70
5.5	Errors of CDD-MP2 total energy as well as Coulomb and exchange parts with respect to canonical MO-MP2 reference, with the 6-31G* basis.	71
5.6	Errors of CDD-MP2 total energy as well as Coulomb and exchange parts with respect to canonical MO-MP2 reference, with the SVP basis.	71
5.7	Errors of CDD-MP2 total energy as well as Coulomb and exchange parts with respect to canonical MO-MP2 reference, with the cc-pVTZ basis.	72
5.8	Errors of CDD-MP2 interaction energies with respect to MO-MP2 reference for the S22 test set with the SVP and cc-pVTZ basis sets.	73
5.9	Comparison of total Wall times of CDD-MP2 and conventional MO-MP2.	74
5.10	Wall times of total calculation and essential steps of the CDD-MP2 algorithm for linear alkanes with the 6-31G* basis set.	75
5.11	Wall times of total calculation and essential steps of the CDD-MP2 algorithm for linear alkanes with the SVP basis set.	76
5.12	Wall times of total calculation and essential steps of the CDD-MP2 algorithm for linear alkanes with the cc-pVTZ basis set.	77
5.13	Wall times of total calculation and essential steps of the CDD-MP2 algorithm for α -D-glucose chains (amylose) with the 6-31G* basis set.	78
5.14	Wall times of total calculation and essential steps of the CDD-MP2 algorithm for α -D-glucose chains (amylose) with the SVP basis set.	79
5.15	Wall times of total calculation and essential steps of the CDD-MP2 algorithm for α -D-glucose chains (amylose) with the cc-pVTZ basis set.	80

LIST OF TABLES

6.1	Comparison of total Wall times of RI-CDD-MP2 and conventional RI-MO-MP2 along with numbers of preselected FTI products.	91
6.2	Errors of RI-CDD-MP2 total energy as well as Coulomb and exchange parts with respect to canonical RI-MO-MP2 reference, with the SVP basis set.	92
6.3	Errors of RI-CDD-MP2 total energy as well as Coulomb and exchange parts with respect to canonical RI-MO-MP2 reference, with the cc-pVTZ basis set.	93
6.4	Deviation [μ Hartree] of RI-MO-MP2 and RI-CDD-MP2 total energies, as well as individual components, with respect to conventional MO-MP2 reference for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the SVP (VDZ**) basis; 5 quadrature points; numerical thresholds for RI-CDD-MP2: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$	94
6.5	Deviation [μ Hartree] of RI-MO-MP2 and RI-CDD-MP2 total energy, as well as individual components, with respect to conventional MO-MP2 reference for linear alkanes, α -D-glucose chains (amylose) and DNA base pairs (adenine-thymine) with the cc-pVTZ basis; 5 quadrature points; numerical thresholds for RI-CDD-MP2: $\theta_{ext} = 10^{-8}$, $\theta_{int} = 10^{-8}$, $\theta_{sp} = 10^{-9}$	94
6.6	Errors of RI-CDD-MP2 interaction energies with respect to RI-MO-MP2 reference for the S22 test set with the SVP and cc-pVTZ basis sets.	95
6.7	Wall times of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm for alkanes with the SVP basis.	96
6.8	Wall times of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm for alkanes with the cc-pVTZ basis.	97
6.9	Wall times of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm for α -D-glucose chains (amylose) with the SVP basis.	98
6.10	Wall times of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm for α -D-glucose chains (amylose) with the cc-pVTZ basis.	99
6.11	Wall times of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm for DNA bases pairs (adenine-thymine) with the SVP basis.	100
6.12	Wall times of total calculation and individual parallel steps of the OMP parallelized RI-CDD-MP2 algorithm for DNA bases pairs (adenine-thymine) with the cc-pVTZ basis.	101
7.1	Comparison of numerical ranks from MSCD and standard CD, along with accuracy of MSCD recomposed ERI matrix.	110

Appendix A

Exponentially weighted energy differences in MP2

The primary incentive to express the MP2 energy—and other electron correlation methods, for that matter—in a local basis is the intrinsically delocalized nature of canonical HF molecular orbitals (MO). This precludes the application of integral screening schemes based on spatial separation of charge distributions to exploit the short-range nature of electronic correlation effects, which in principle would lead to a $1/R^3$ decay of electron repulsion integrals (ERI) $(ia|jb)$.

Considering the expression for the canonical MP2 energy

$$E_0^{(2)} = \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} \frac{(ia|jb) [2(ia|jb) - (ib|ja)]}{\epsilon_i - \epsilon_a + \epsilon_j - \epsilon_b} = \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} \frac{M_{ia,jb} G_{ia,jb}}{D_{ia} + D_{jb}}, \quad (\text{A.1})$$

however, a reduction of the computational effort required to evaluate $E_0^{(2)}$ might also be based on the truncation of the substitution space of a given occupied orbital ϕ_i , i.e. the set $\{\phi_a\}_{a \in N_{virt}}$ of virtual MO into which electron substitution $i \rightarrow a$ is permitted, by taking into account the difference $D_{ia} = \epsilon_i - \epsilon_a$ of orbital energies. Namely, with increasing D_{ia} and D_{jb} , the value of an associated amplitude $M_{ia,jb} G_{ia,jb}$ should dramatically decrease. This approach is actually already widely applied in form of the frozen core and frozen virtual approximations, where coreorbitals and high lying virtuals are excluded from the active and virtual spaces, since the occurrence of D_{ia} in the denominator already induces a $1/D_{ia}$ decay of the summands. This damping form, however, should not allow for efficient reduction of the number of significant contributions beyond the frozen core and frozen virtual approximation.

In the expression of the Laplace-transform MP2 energy

$$E_0^{(2)} = - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} e^{-D_{ia}t\alpha} M_{ia,jb} G_{ia,jb} e^{-D_{jb}t\alpha}. \quad (\text{A.2})$$

as discussed in section 4.1, the individual summands decay *exponentially* with the orbital energy difference, which already would lead to a more rapid decrease of significant contributions. However, in this context, this expression bears the disadvantage that no coupling between the substitution energies D_{ia} and D_{jb} is induced.

Another possibility of expressing the energy denominator in (A.1) is to use an integral transform of the type

$$\frac{1}{x} = \frac{2}{\sqrt{\pi}} \int_0^\infty e^{-x^2 t^2} dt \approx \frac{2}{\sqrt{\pi}} \sum_\alpha^\tau w_\alpha e^{-x^2 t_\alpha^2}, \quad (\text{A.3})$$

which, in contrast to the Laplace transform approach, does not decouple the orbital energies in view of a change of basis, but more suits the present purpose of exploiting the damping induced by the energy differences. With this expression, the MP2 energy takes the form

$$E_0^{(2)} = - \sum_\alpha^\tau w'_\alpha \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} M_{ia,jb} G_{ia,jb} e^{-(D_{ia}+D_{jb})^2 t_\alpha^2}, \quad (\text{A.4})$$

where $w'_\alpha = 2w_\alpha/\sqrt{\pi}$, such that with $(D_{ia} + D_{jb})^2 = D_{ia}^2 + 2D_{ia}D_{jb} + D_{jb}^2$ one finally obtains

$$E_0^{(2)} = - \sum_\alpha^\tau w'_\alpha \sum_{i,j}^{N_{occ}} \sum_{a,b}^{N_{virt}} M_{ia,jb} G_{ia,jb} e^{-(D_{ia}^2 + 2D_{ia}D_{jb} + D_{jb}^2) t_\alpha^2}. \quad (\text{A.5})$$

With this functional form, the substitution amplitude should strongly decrease with the differences of orbital energies. For illustration, assuming that the maximum magnitude of ERI is $\mathcal{O}(1)$ a.u. and noting that $e^{-D_{min}^2 t_\alpha^2}$ is always smaller than unity, in order for the term $e^{-D_{ia}^2 t_\alpha^2}$ and hence the whole summand to be smaller in magnitude than 10^{-6} , one obtains the condition

$$D_{ia} > \frac{\sqrt{6 \ln 10}}{t_\alpha} \approx 3.7, \quad (\text{A.6})$$

assuming $t_\alpha \approx 1$ in the last step. From this consideration, the aforementioned frozen core and frozen virtual approximations are certainly directly recovered, since the energy differences between highest lying core orbital and LUMO, as well as between the HOMO and highest lying virtuals should always be large enough for the substitution to be strongly damped by the exponential terms. The above expression bears the further advantage than the term $e^{-2D_{ia}D_{jb}t_\alpha^2}$ couples both charge distributions in the ERI, so that for a given set of D_{ia} only a reduced number of D_{jb} should lead to numerically significant terms.

On the basis of expression (A.5), one may thus envisage a reduction of the substitution space based on successive screening of the excitations that have significant contributions according to the exponential damping. Although for full numerical control the integrals should be known in advance, in order to apply the exponential term to the actual amplitudes rather than relying on the rough estimate used above for illustrative purposes, this procedure could for instance be used in conjunction with resolution of the identity (RI). Since with RI approximated integrals the

computational effort is shifted away from the AO-MO-transformation towards the contraction to the energy, the proposed procedure may help reduce the complexity of this step within canonical RI-MP2 algorithms.

Appendix B

Kronecker Product Approximation of ERI

Along the lines of the canonical product tensor approximation of ERI proposed by Benedikt *et al.* [24] (see section 7.1.3), we wish to propose an alternative, though related approach that may prove beneficial for the calculation of the MP2 energy within the Laplace transform formalism of Almlöf and Häser and with Cholesky decomposed pseudo-density matrices (see sections 4.1 and 5.1). Namely, if instead of a CP factorization a low rank matrix Kronecker product approximation (KPA) of the ERI matrix were obtained as

$$\mathbf{A} \approx \sum_r^R (\mathbf{X}_r \otimes \mathbf{Y}_r), \quad (\text{B.1})$$

where \mathbf{X}_r and \mathbf{Y}_r are of dimension $N \times N$, the transformation of the integrals in equation (4.20) could then be efficiently performed according to

$$\mathbf{M}' = \sum_r^R (\underline{\mathbf{L}}^T \otimes \bar{\mathbf{L}}^T) (\mathbf{X}_r \otimes \mathbf{Y}_r) (\underline{\mathbf{L}} \otimes \bar{\mathbf{L}}) = \sum_r^R (\underline{\mathbf{L}}^T \mathbf{X}_r \underline{\mathbf{L}} \otimes \bar{\mathbf{L}}^T \mathbf{Y}_r \bar{\mathbf{L}}) = \sum_r^R (\mathbf{X}'_r \otimes \mathbf{Y}'_r). \quad (\text{B.2})$$

Using the useful property of the Kronecker product

$$\text{Tr}(\mathbf{A} \otimes \mathbf{B}) = \text{Tr}\mathbf{A} \cdot \text{Tr}\mathbf{B}, \quad (\text{B.3})$$

the contraction to the energy (here only the Coulomb type energy, but these considerations apply to the exchange part as well, though for permuted indices) then becomes

$$\begin{aligned}
 E_J^{(2)} &= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left((\underline{\mathbf{e}}' \otimes \bar{\mathbf{e}}') \mathbf{M}' (\underline{\mathbf{e}}' \otimes \bar{\mathbf{e}}') \right)^2 \\
 &= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left(\sum_r^R (\underline{\mathbf{e}}' \otimes \bar{\mathbf{e}}') (\mathbf{X}'_r \otimes \mathbf{Y}'_r) (\underline{\mathbf{e}}' \otimes \bar{\mathbf{e}}') \right)^2 \\
 &= - \sum_{\alpha}^{\tau} w_{\alpha} \text{Tr} \left(\sum_r^R (\underline{\mathbf{e}}' \mathbf{X}'_r \bar{\mathbf{e}}' \otimes \underline{\mathbf{e}}' \mathbf{Y}'_r \bar{\mathbf{e}}') \right)^2 \\
 &= - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{r,s}^R \text{Tr} (\mathbf{X}''_r \mathbf{X}''_s \otimes \mathbf{Y}''_r \mathbf{Y}''_s) \\
 &= - \sum_{\alpha}^{\tau} w_{\alpha} \sum_{r,s}^R \text{Tr} (\mathbf{X}''_r \mathbf{X}''_s) \cdot \text{Tr} (\mathbf{Y}''_r \mathbf{Y}''_s),
 \end{aligned} \tag{B.4}$$

where the double prime indicates contraction with the $\underline{\mathbf{e}}'$ and $\bar{\mathbf{e}}'$, i.e. $\mathbf{X}''_r = \underline{\mathbf{e}}' \mathbf{X}'_r \bar{\mathbf{e}}'$. The occupied and virtual transformation are hereby completely decoupled, as is seen in equation (B.2), and the contraction to the energy is reduced to forming the traces of $2R^2$ products of—in practice surely sparse— $N_{occ} \times N_{occ}$ and $N_{virt} \times N_{virt}$ matrices, hence allowing for the use of optimized linear algebra library routines as well as parallel implementation.

The KPA of a matrix \mathbf{A} , i.e. the search of matrices \mathbf{X}_r and \mathbf{Y}_r for which the norm

$$\|\mathbf{A} - \tilde{\mathbf{A}}\|_F = \|\mathbf{A} - \sum_r^R (\mathbf{X}_r \otimes \mathbf{Y}_r)\|_F \tag{B.5}$$

is minimized, is closely related to the singular value decomposition (SVD), at least for the case $R = 1$. Namely, it can be shown that for $\tilde{\mathbf{A}}$ having the SVD $\tilde{\mathbf{A}} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$, the minimizers of

$$\|\mathbf{A} - \mathbf{X} \otimes \mathbf{Y}\|_F \tag{B.6}$$

are the matrices \mathbf{X}^* and \mathbf{Y}^* with

$$\text{vec}(\mathbf{X}^*) = \sigma_1 \mathbf{u}_1, \quad \text{vec}(\mathbf{Y}^*) = \mathbf{v}_1, \tag{B.7}$$

where \mathbf{u}_1 and \mathbf{v}_1 are the first columns of \mathbf{U} and \mathbf{V} and σ_1 is the largest element of $\mathbf{\Sigma}$, i.e. the largest singular value of $\tilde{\mathbf{A}}$ [144].

Further, as the ERI matrix that is to be decomposed is highly symmetric and sparse, especially when applying a screening scheme to determine and compute only those elements that are required in the transformations to the MO ERIs, the matrices \mathbf{X}_r and \mathbf{Y}_r would themselves most probably be symmetric and relatively sparse, so that one might possibly apply a constrained KPA algorithm.

Appendix C

List of abbreviations

AO	-	Atomic Orbitals
API	-	Application Programming Interface
CC	-	Coupled Cluster
CD	-	Cholesky Decomposition
CDD	-	Cholesky Decomposed Densities
CFMM	-	Continuous Fast Multipole Method
CI	-	Configuration Interaction
CP	-	Canonical Product Format
DF	-	Density Fitting
DFT	-	Density Functional Theory
DIIS	-	Direct Inversion in the Iterative Subspace
DNA	-	Deoxyribonucleic Acid
ER	-	Edminston-Ruedenberg Localization
ERI	-	Electron Repulsion Integrals
FCI	-	Full Configuration Interaction
FMM	-	Fast Multipole Method
FTI	-	Fully Transformed Integrals ($\underline{i\bar{a}} \underline{j\bar{b}}$)
<i>FTidx</i>	-	Indexing array of FTI
HF	-	Hartree-Fock
HTI	-	Half Transformed Integrals ($\underline{i\bar{a}} \lambda\sigma$)
<i>HTidx</i>	-	Indexing array of HTI
<i>Kidx</i>	-	Assignment array of exchange integrals
KPA	-	Kronecker Product Approximation
LCAO	-	Linear Combination of Atomic Orbitals
LMO	-	Localized Molecular Orbitals

APPENDIX C. LIST OF ABBREVIATIONS

MBPT	-	Many Body Perturbation Theory
MKL	-	Math Kernel Library
MO	-	Molecular Orbitals
MP2	-	Second Order Møller-Plesset Perturbation Theory
MP n	-	n -th Order Møller-Plesset Perturbation Theory
MSCD	-	Method Specific Cholesky Decomposition
OMP	-	Open Multi-Processing
PAO	-	Projected Atomic Orbitals
PM	-	Pipek-Mezey Localization
RAM	-	Random Access Memory
RHF	-	Restricted Hartree-Fock
RI	-	Resolution of the Identity
RPA	-	Random Phase Approximation
RSPT	-	Rayleigh-Schrödinger Perturbation Theory
SAD	-	Superposition of Atomic Densities
SCS	-	Spin Component Scaling
SOS	-	Scaled Opposite Spin
STI	-	Single Transformed Integrals ($i\nu \lambda\sigma$)
TTI	-	Triple Transformed Integrals ($i\bar{a} j\sigma$)
UTI	-	Untransformed Integrals ($\mu\nu \lambda\sigma$)

Bibliography

- [1] C. Møller and M. S. Plesset, *Phys. Rev*, **46**, 618 (1934).
- [2] J. Almlöf, *Chem. Phys. Lett.*, **181**, 319 (1991).
- [3] M. Häser and J. Almlöf, *J. Chem. Phys.*, **96**, 489 (1991).
- [4] M. Häser, *Theor. Chem. Acc.*, **87**, 147 (1993).
- [5] B. Doser, D. S. Lambrecht and C. Ochsenfeld, *J. Chem. Phys.*, **123**, 184102 (2005).
- [6] B. Doser, D. S. Lambrecht, J. Kussmann and C. Ochsenfeld, *J. Chem. Phys.*, **130**, 064107 (2009).
- [7] B. Doser, *Linear-Scaling Møller-Plesset Perturbation Theory for the Calculation of Electron Correlation in Large Molecules*, Dissertation, Eberhard Karls Universität Tübingen, C. Ochsenfeld Group, (2009).
- [8] Y. Jung, R. C. Lochan, A. D. Dutoi and M. Head-Gordon, *J. Chem. Phys.*, **121**, 9793 (2004).
- [9] C. E. Benoît, *Bulletin Géodesique*, **2**, 67 (1924).
- [10] J. Zienau, L. Clin, B. Doser and C. Ochsenfeld, *J. Chem. Phys.*, **130**, 204112 (2009).
- [11] B. Doser, J. Zienau, L. Clin and C. Ochsenfeld, *Z. Phys. Chem.*, **224**, 397 (2010).
- [12] S. Maurer, D. S. Lambrecht, D. Flaig and C. Ochsenfeld, *J. Chem. Phys.*, **136**, 144107 (2012).
- [13] J. Zienau, *Low-Order Scaling Second-Order Møller-Plesset Perturbation Theory Using Auxiliary Basis Sets and Quantum-Chemical Calculations of NMR Chemical Shifts for Supramolecular Systems*, Dissertation, Eberhard Karls Universität Tübingen, C. Ochsenfeld Group, (2009).
- [14] F. E. Harris and R. Rein, *Theor. Chem. Acc.*, **6**, 73 (1966).

BIBLIOGRAPHY

- [15] J. L. Whitten, *J. Chem. Phys.*, **58**, 4496 (1973).
- [16] B. I. Dunlap, J. W. D. Connolly and J. R. Sabin, *J. Chem. Phys.*, **71**, 4993 (1979).
- [17] R. A. Kendall, H. A. Früchtl, *Theor. Chem. Acc.*, **97**, 158 (1979).
- [18] M. Feyereisen, G. Fitzgerald and A. Komornicki, *Chem. Phys. Lett.*, **208**, 359 (1993).
- [19] L. Boman, H. Koch and A. S. de Merás, *J. Chem. Phys.*, **129**, 134107 (2008).
- [20] W. Heisenberg, *Physik und Philosophie*, 6. Aufl., Hirzel, Stuttgart, 2000.
- [21] D. R. Hartree, *Proc. Cambridge Philos. Soc.*, **24**, 89 (1928).
- [22] V. Fock, *Z. Phys.*, **61**, 126 (1930).
- [23] L. Clin, *A Study of Integral Transformation in Atomic Orbital Møller-Plesset Theory*, Diplomarbeit, Eberhard Karls Universität Tübingen, C. Ochsenfeld Group, (2007).
- [24] U. Benedikt, A. A. Auer, M. Espig and W. Hackbusch, *J. Chem. Phys.*, **134**, 054118 (2011).
- [25] E. Schrödinger, *Ann. Phys.*, **384**, 361 (1926).
- [26] L. D. Landau and E. M. Lifschitz, *Lehrbuch der Theoretischen Physik, Band III - Quantenmechanik*, Akademie Verlag, Berlin, 1986.
- [27] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry*, Dover Publications, Mineola, New York, 1996.
- [28] W. Ritz, *Reine. Angew. Math.*, **135**, 1 (1909).
- [29] A. Messiah, *Quantum mechanics*, Dover Publications, Mineola, New York, 1999.
- [30] C. C. J. Roothaan, *Rev. Mod. Phys.*, **23**, 69 (1951).
- [31] P. A. M. Dirac, *The Principles of Quantum Mechanics*, Oxford University Press, Oxford, 1958.
- [32] P. O. Löwdin, *J. Chem. Phys.*, **18**, 365 (1950).
- [33] B. Jansík, S. Høst, P. Jørgensen and J. Olsen, *J. Chem. Phys.*, **126**, 124104 (2007).
- [34] T. Helgaker, P. Jørgensen and J. Olsen, *Molecular Electronic Structure Theory*, Wiley, Chichester, 2000.
- [35] P. Pulay, *Chem. Phys. Lett.*, **73**, 393 (1980).

BIBLIOGRAPHY

- [36] P. Pulay, *J. Comput. Chem.*, **3**, 556 (1982).
- [37] J. v. Neumann, *Göttinger Nachrichten*, page 245 (1927).
- [38] R. McWeeny, *Rev. Mod. Phys.*, **32**, 335 (1960).
- [39] H. S. Wilf, *Algorithms and complexity, Internet edition*, University of Pennsylvania, Philadelphia, 1994.
- [40] V. Y. Bunyakovsky, *Mem. Acad. Sci. St. Petersburg*, **1 (9)**, 1 (1859).
- [41] N. I. Akhiezer and I. M. Glazman, *Theory of Linear Operators in Hilbert Space*, Dover Publications, Mineola, New York, 1993.
- [42] M. Häser and R. Ahlrichs, *J. Comput. Chem.*, **10**, 104 (1989).
- [43] P. M. W. Gill, B. G. Johnson and J. A. Pople, *Chem. Phys. Lett.*, **217**, 65 (1994).
- [44] S. Goedecker, *Rev. Mod. Phys.*, **71**, 1085 (1999).
- [45] J. des Cloizeaux, *Phys. Rev.*, **135**, A685 (1964).
- [46] W. Kohn, *Phys. Rev.*, **115**, 809 (1959).
- [47] A. M. N. Niklasson, C. J. Tymczak and M. Challacombe, *J. Chem. Phys.*, **118**, 8611 (2003).
- [48] T. Helgaker, H. Larsen, J. Olsen and P. Jørgensen, *Chem. Phys. Lett.*, **327**, 397 (2000).
- [49] M. Head-Gordon, Y. Shao, C. Saravanan and C. A. White, *Mol. Phys.*, **101**, 37 (2003).
- [50] C. Ochsenfeld, C. A. White and M. Head-Gordon, *J. Chem. Phys.*, **109**, 1663 (1998).
- [51] L. Greengard and V. Rokhlin, *J. Comput. Phys.*, **60**, 187 (1990).
- [52] L. Greengard, *Science*, **265**, 909 (1994).
- [53] C. A. White and M. Head-Gordon, *J. Chem. Phys.*, **101**, 6593 (1994).
- [54] C. A. White and M. Head-Gordon, *J. Chem. Phys.*, **105**, 5061 (1996).
- [55] C. A. White, B. G. Johnson, P. M. W. Gill and M. Head-Gordon, *Chem. Phys. Lett.*, **230**, 8 (1998).
- [56] C. Ochsenfeld, J. Kussmann and D. S. Lambrecht, *Linear-Scaling Methods in Quantum Chemistry*, Wiley-VCH, New York, Weinheim, 2007.
- [57] P. O. Löwdin, *Phys. Rev.*, **97**, 1509 (1955).

BIBLIOGRAPHY

- [58] C. D. Sherrill, *An Introduction to Configuration Interaction Theory*, 1995.
- [59] G. L. G. Sleijpen and H. A. A. van der Vorst, *SIAM J. Matrix. Anal. Appl.*, **17**, 401 (1996).
- [60] G. L. G. Sleijpen and H. A. A. van der Vorst, *SIAM Review*, **42**, 267 (2000).
- [61] R. J. Bartlett, D. R. Yarkony (ed), *Modern electronic Structure Theory*, World Scientific, 1995.
- [62] D. Bohm and D. Pines, *Phys. Rev.*, **82**, 625 (1951).
- [63] D. Pines and D. Bohm, *Phys. Rev.*, **85**, 338 (1952).
- [64] D. Bohm and D. Pines, *Phys. Rev.*, **92**, 609 (1953).
- [65] A. D. McLachlan and M. A. Ball, *Rev. Mod. Phys.*, **36**, 844 (1964).
- [66] F. Furche, *J. Chem. Phys.*, **129**, 114105 (2008).
- [67] H. Eshuis, J. Yarkony and F. Furche, *J. Chem. Phys.*, **132**, 234114 (2010).
- [68] G. Scuseria, T. M. Henderson and D. C. Sorensen, *J. Chem. Phys.*, **129**, 231101 (2008).
- [69] G. Jansen, R-F. Liu and J. G. Ángyán, *J. Chem. Phys.*, **133**, 154106 (2010).
- [70] E. Schrödinger, *Ann. Phys*, **385**, 437 (1926).
- [71] J. Olsen, O. Christiansen, H. Koch and P. Jørgensen, *J. Chem. Phys.*, **105**, 5082 (1996).
- [72] M. L. Leininger, W. D. Allen, H. F. Schaefer III and C. D. Sherrill, *J. Chem. Phys.*, **112**, 9213 (2000).
- [73] L. Brillouin, *Actual. scient. ind.*, **71**, sections 159, 160 (1933).
- [74] S. Grimme, *J. Chem. Phys.*, **118**, 9095 (2003).
- [75] R. Fink, *J. Chem. Phys.*, **133**, 174113 (2010).
- [76] R. C. Lochan, Y. Jung and M. Head-Gordon, *J. Phys. Chem. A*, **109**, 7598 (2005).
- [77] Y. Jung, Y. Shao and M. Head-Gordon, *J. Comput. Chem.*, **28**, 1953 (2007).
- [78] P. Pulay, *Chem. Phys. Lett.*, **100**, 151 (1983).
- [79] S. Saebø and P. Pulay, *Annu. Rev. Phys. Chem.*, **44**, 213 (1993).
- [80] S. F. Boys, *Rev. Mod. Phys.*, **32**, 296 (1960).

BIBLIOGRAPHY

- [81] C. Edmiston and K. Ruedenberg, *Rev. Mod. Phys.*, **35**, 457 (1963).
- [82] J. Pipek and P. G. Mezey, *J. Chem. Phys.*, **90**, 4916 (1989).
- [83] B. Jansík, S. Høst, K. Kristensen and P. Jørgensen, *J. Chem. Phys.*, **134**, 194104 (2011).
- [84] J. M. Foster and S. F. Boys, *Rev. Mod. Phys.*, **32**, 300 (1960).
- [85] C. Edmiston and K. Ruedenberg, *J. Chem. Phys.*, **43**, S97 (1965).
- [86] W. von Niessen, *J. Chem. Phys.*, **56**, 4290 (1972).
- [87] J. Pipek, *Int. J. Quantum Chem.*, **36**, 487 (1989).
- [88] F. Aquilante, T. B. Pedersen, A. S. de Merás and H. Koch, *J. Chem. Phys.*, **125**, 174101 (2006).
- [89] Y. Guo, W. Li and S. Li, *J. Chem. Phys.*, **135**, 134107 (2011).
- [90] J. E. Subotnik, A. D. Dutoi and M. Head-Gordon, *J. Chem. Phys.*, **123**, 114108 (2005).
- [91] S. Schweizer, J. Kussmann, B. Doser and C. Ochsenfeld, *J. Comput. Chem.*, **29**, 1004 (2008).
- [92] P. Pulay and S. Saebø, *Theor. Chem. Acc.*, **69**, 357 (1986).
- [93] M. Schütz, G. Hetzer and H.-J. Werner, *J. Chem. Phys.*, **111**, 5691 (1999).
- [94] G. Rauhut, P. Pulay and H.J. Werner, *J. Comput. Chem.*, **19**, 1241 (1998).
- [95] J. W. Boughton and P. Pulay, *J. Comput. Chem.*, **14**, 736 (1994).
- [96] H.-J. Werner, F. R. Manby and P. R. Knowles, *J. Chem. Phys.*, **118**, 8149 (2003).
- [97] N. Russ and T. D. Crawford, *J. Chem. Phys.*, **121**, 691 (2004).
- [98] R. A. Mata and H.-J. Werner, *J. Chem. Phys.*, **125**, 184110 (2006).
- [99] T. H. Dunning Jr., *J. Chem. Phys.*, **90**, 1007 (1989).
- [100] D. E. Woon and T. H. Dunning Jr., *J. Chem. Phys.*, **98**, 1358 (1993).
- [101] D. E. Woon and T. H. Dunning Jr., *J. Chem. Phys.*, **100**, 2975 (1994).
- [102] K. A. Peterson, D. E. Woon and T. H. Dunning Jr., *J. Chem. Phys.*, **100**, 7410 (1994).
- [103] W. Klopper, F. R. Manby, S. Ten-no and E. F. Valeev, *Int. Rev. Phys.Chem.*, **25**, 427 (2006).

BIBLIOGRAPHY

- [104] T. Kato, *Commun. Pure Appl. Math.*, **10**, 151 (1957).
- [105] E. A. Hylleraas, *Z. Phys.*, **54**, 347 (1929).
- [106] T. Helgaker and W. Klopper, *Theor. Chem. Acc.*, **103**, 180 (2000).
- [107] S. Ten-no, *Chem. Phys. Lett.*, **398**, 56 (2004).
- [108] D. P. Tew and W. Klopper, *J. Chem. Phys.*, **123**, 074101 (2005).
- [109] T. B. Adler, H.-J. Werner and F. R. Manby, *J. Chem. Phys.*, **126**, 164102 (2007).
- [110] W. Kutzelnigg and W. Klopper, *J. Chem. Phys.*, **94**, 1985 (1991).
- [111] W. Kutzelnigg and W. Klopper, *J. Chem. Phys.*, **94**, 2020 (1991).
- [112] C. B. Lang and N. Pucker, *Mathematische Methoden in der Physik, 2. Aufl*, Spektrum Akademischer Verlag, Heidelberg, 2005.
- [113] D. Kats, D. Usvyat and M. Schütz, *Phys. Chem. Chem. Phys.*, **10**, 3430 (2008).
- [114] A. J. Laub, *Matrix Analysis for Scientists and Engineers*, Soc. Ind. Appl. Math. (SIAM), 2004.
- [115] S. J. Hammarling, *IMA J. Num. Anal.*, **2**, 303 (1982).
- [116] C. C. K. Mikkelsen, *Numerical Methods for Large Scale Lyapunov Equations*, PhD thesis, Purdue University, West Lafayette, Indiana, (2009).
- [117] J.-R. Li and J. White, *SIAM J. Mat. Anal. Appl.*, **24**, 260 (2004).
- [118] G. E. Scuseria and P. Ayala, *J. Chem. Phys.*, **111**, 8330 (1999).
- [119] P. Y. Ayala and G. E. Scuseria, *J. Chem. Phys.*, **110**, 3660 (1999).
- [120] P. R. Surján, *Chem. Phys. Lett.*, **406**, 318 (2005).
- [121] Development version of the programme package Q-CHEM, (www.q-chem.com).
- [122] O. Reiß, *Mathematical Methods for the efficient Assessment of Market and Credit Risk*, Dissertation, Universität Kaiserslautern, (2003).
- [123] N. F. H. Beebe and J. Linderberg, *Int. J. Quantum Chem.*, **XII**, 683 (1977).
- [124] P. Jurečka, J. Šponer, J. Černý, and P. Hobza, *Phys. Chem. Chem. Phys.*, **8**, 1985 (2006).
- [125] OpenMP website: <http://openmp.org/wp/>.

BIBLIOGRAPHY

- [126] O. Vahtras, J. Almlöf and M. W. Feyereisen, *Chem. Phys. Lett.*, **213**, 514 (1993).
- [127] B. I. Dunlap, J. W. D. Connolly and J. R. Sabin, *Int. J. Quantum Chem.*, **Symp. 11**, 81 (1977).
- [128] D. S. Lambrecht and C. Ochsenfeld, *J. Chem. Phys.*, **123**, 184101 (2005).
- [129] D. S. Lambrecht and C. Ochsenfeld, *J. Chem. Phys.*, **136**, 149901 (2012).
- [130] D. S. Lambrecht, B. Doser and C. Ochsenfeld, *J. Chem. Phys.*, **136**, 149902 (2012).
- [131] J. J. Dongarra, J. Du Croz, S. Hammarling and I. S. Duff, *ACM Transactions on Mathematical Software*, **16**, 1 (1990).
- [132] P. O. Löwdin, *Int. J. Quantum Chem.*, **S4**, 231 (1971).
- [133] I. Røeggen and E. Wisløff-Nilssen, *Chem. Phys. Lett.*, **132**, 154 (1986).
- [134] D. W. O'Neal and J. Simons, *Int. J. Quantum chem.*, **XXXVI**, 673 (1989).
- [135] S. Wilson, *Comp. Phys. Comm.*, **58**, 71 (1990).
- [136] H. Koch, A. S. de Merás and T. B. Pedersen, *J. Chem. Phys.*, **118**, 9481 (2003).
- [137] F. Aquilante, R. Lindh and T. B. Pedersen, *J. Chem. Phys.*, **127**, 114107 (2007).
- [138] T. B. Pedersen, F. Aquilante and R. Lindh, *Theor. Chem. Acc.*, **124**, 1 (2009).
- [139] F. Aquilante, L. Gagliardi, T. B. Pedersen and R. Lindh, *J. Chem. Phys.*, **130**, 154107 (2009).
- [140] F. Aquilante, T. B. Pedersen and R. Lindh, *J. Chem. Phys.*, **126**, 194106 (2007).
- [141] T. S. Chwee and E. A. Carter, *J. Chem. Phys.*, **132**, 074104 (2010).
- [142] F. Aquilante, R. Lindh and T. B. Pedersen, *J. Chem. Phys.*, **129**, 034106 (2008).
- [143] F. Weigend, M. Kattannek and R. Ahlrichs, *J. Chem. Phys.*, **130**, 164106 (2009).
- [144] N. P. Pitsianis, *The Kronecker Product in Approximation and Fast Transform Generation*, PhD thesis, Cornell University, (1997).

BIBLIOGRAPHY

Appendix D

List of academic teachers

Prof. Dr. K. Albert	Prof. Dr. H. Mayer
Prof. Dr. T. Chassé	Prof. Dr. A. Meixner
Prof. Dr. D. Christen	Prof. Dr. H.-J. Meyer
Prof. Dr. H. Eckstein	Prof. Dr. U. Nagel
PD Dr. H. Egelhaaf	Prof. Dr. H. Oberhammer
Dr. A. V. Failla	Prof. Dr. C. Ochsenfeld
Prof. Dr. G. Gauglitz	Prof. Dr. H. Pommer
Prof. Dr. Hartschuh	Prof. Dr. V. Schurig
Prof. Dr. W. Jäger	Prof. Dr. E. Schweda
PD Dr. A. Klein	Prof. Dr. J. Strähle
Prof. Dr. W. Knapp	PD Dr. U. Weimar
Prof. Dr. N. Kuhn	Prof. Dr. L. Wesemann
Dr. C. Maichle-Mössmer	Prof. Dr. K.-P. Zeller
Prof. Dr. M. Maier	Prof. Dr. T. Ziegler