Internally Contracted Multireference Coupled-Cluster Theories With Automated Code Generation

Dissertation

zur

Erlangung des Doktorgrades (Dr. rer. nat.)

der

Mathematisch-Naturwissenschaftlichen Fakultät

der

Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Marvin Harald Lechner

aus

München

Bonn, Oktober 2022

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät
der Rheinischen Friedrich-Wilhelms-Universität Bonn

1. Gutachter: Prof. Dr. Frank Neese

2. Gutachter: Prof. Dr. Stefan Grimme

Tag der Promotion: 28. April 2023

Erscheinungsjahr: 2023

Abstract

Developing a coupled-cluster theory based on a multiconfigurational reference wave function still is one of the most challenging problems in quantum chemistry, both from a theoretical and implementational perspective. Hence, no clear scientific consensus has been reached yet on the aspects of such theories. The main reason for this is that many different parameterizations are possible based on several theoretical choices that can be made, e.g., whether to use a contracted or uncontracted *ansatz*, which residual conditions to employ, how to treat the available excitation classes, whether to use a single or sequential similarity transformations of the Hamiltonian, and more. In this thesis, we further elucidate some aspects of this broad topic, to pave a path towards a theoretically rigorous, generally accepted multireference coupled-cluster method. To this end, we focus especially on an efficient implementation, the residual conditions, perturbative approximations, and a way to reduce the dimensionality of the involved tensors, i.e., foremost, density matrices.

To implement such theories, especially the internally contracted approaches, automated tools are required since the theories contain upwards of hundreds of thousands of terms, posing a formidable challenge. Consequently, we wrote a highly performant toolchain, ORCA-AGE II, which can derive and implement even the most complicated variants of multireference coupled-cluster theory. The toolchain consists of an optimized code generation part that keeps the code generation time as short as possible, as well as sophisticated algorithms to find optimal transformations of the tensor contractions so that they can be evaluated close to peak CPU efficiency.

In the multireference theories, more specifically on the topic of residual conditions, we propose a cumulant-based expansion that connects the many-body to the projective residual conditions and clearly demonstrates the more complicated nature of the projective variant. This expansion justifies the truncations present in the many-body expansions used in multireference equation-of-motion theories. These findings are then used in multireference equation-of-motion perturbation theory, which we developed as a perturbative transform-then-diagonalize method. From benchmarking the novel method on various organic and inorganic systems, we find it has an accuracy comparable to that of NEVPT2 theory, while being significantly cheaper and more stable than its parent method, multireference equation-of-motion coupled-cluster theory. On a different aspect of internally contracted methods with projective residual conditions, we developed an automated reduction scheme for high-order density matrices that can be applied to any method. The scheme works exceptionally well on multireference coupled-cluster theory, being always faster than the unreduced implementation through a combination of asymptotic

CPU cost reductions and more efficient usage of the CPU caches through lower-dimensional tensors. These savings are realized even though the maximum order of the density matrices can be reduced by at most one since the structure of the equations does not allow for higher reductions. Finally, we also report highly accurate transition energies computed at the single-reference level benchmarked on indigo dyes, showcasing the applicability yet to be reached with the more complicated multireference approaches.

List of Publications

Published Parts of This Thesis

Parts of this thesis have appeared in the following publications:

- **M. H. Lechner**, F. Neese, and R. Izsák, *An excited state coupled-cluster study on indigo dyes*, Mol. Phys. **119**, e1965235 (2021). doi: 10.1080/00268976.2021.1965235.

  Following a project idea by R. Izsák, I performed the calculations and analyzed all computational data while being advised by my coauthors.

- **M. H. Lechner**, R. Izsák, M. Nooijen, and F. Neese, *A perturbative approach to multireference equation-of-motion coupled cluster*, Mol. Phys. **119**, e1939185 (2021). doi: 10.1080/00268976.2021.1939185.

  Following an idea by M. Nooijen, I worked with my coauthors to derive the variants of MR-EOMPT theory, implemented the variants, performed all calculations and analyzed all computational results while being advised by my coauthors.

Manuscripts in Preparation

Parts of this thesis will be published in the near future:

- **M. H. Lechner**, A. Papadopoulos, K. Sivalingam, A. A. Auer, A. Koslowski, U. Becker, F. Wennmohs and F. Neese, *Generated Code in ORCA: Progress, Efficiency and Tight Integration*.

  The software, ORCA-AGE II, being described in this article is loosely based on prior work by Martin Krupička *et al.* [1]. I designed and implemented the core library as well as the majority of equation processing and code generation functionality. My coathors provided advice and/or implemented parts of the generation toolchain under my guidance

(`reorder`, `sumint`, parts of `memory_opt`). I further wrote the article except for the sections on the gradient and CCSDT calculations.

The content is mainly presented in Ch. 4.

- **M. H. Lechner**, K. Sivalingam, F. Neese, *Automated Scaling Reduction for Internally Contracted Theories*.

  I worked out the automated scheme for scaling reduction, implemented the reduced-scaling scheme in ORCA-AGE, as well as designed, conducted, and evaluated the study while being advised by my coauthors.

  The content is mainly presented in Ch. 7.

Further Publications

- Daniel G. A. Smith, Lori A. Burns, Dominic A. Sirianni, Daniel R. Nascimento, Ashutosh Kumar, Andrew M. James, Jeffrey B. Schriber, Tianyuan Zhang, Boyi Zhang, Adam S. Abbott, Eric J. Berquist, **Marvin H. Lechner**, Leonardo A. Cunha, Alexander G. Heide, Jonathan M. Waldrop, Tyler Y. Takeshita, Asem Alenaizan, Daniel Neuhauser, Rollin A. King, Andrew C. Simmonett, Justin M. Turney, Henry F. Schaefer, Francesco A. Evangelista, A. Eugene DePrince III, T. Daniel Crawford, Konrad Patkowski, and C. David Sherrill, *Psi4NumPy: An Interactive Quantum Chemistry Programming Environment for Reference Implementations and Rapid Development*, J. Chem. Theory Comput. **14**, 3504–3511 (2018). doi: 10.1021/acs.jctc.8b00286.

Conference Talks Related to This Work

- Marvin H. Lechner, "Perturbative Approximations in MR-EOMCC Theory," at *New Developments in Coupled-Cluster Theory*, Telluride Science Research Center, July 23, 2021.

  In this talk, I focused on presenting the results from a previous publication [2]. Additionally, I highlighted the importance of code generation in advanced correlated methods for multireference theories such as MR-EOMPT, MR-EOMCC, and fic-MRCC. To this end, I briefly outlined the structure of my implementation of the ORCA-AGE toolchain (second, rewritten version).

Acknowledgments

First, I want to express my gratitude to Prof. Dr. Frank Neese for proposing the challenging, but fun project that eventually led to this thesis. I enjoyed the freedom I had in exploring viable solutions to the problems in multireference coupled-cluster theory, and always knew who I could turn to with even the most difficult of questions.

The project also would not have been possible without Dr. Róbert Izsák, who was a reliable source of motivation, inspiration, and support during the whole time, as well as a good friend, and more than once helped me clarify my thoughts so that the project could progress further.

Working with the ORCA team has also been highly rewarding. Dr. Frank Wennmohs has been an invaluable resource when it came to programming-related questions, and I particularly enjoyed the philosophical programming discussions with him. I further thank Kantharuban Sivalingam for his invaluable help on the old code generator, explaining the "ORCA way" to me, being an exemplary collaborator on implementations, and for reminding me to keep things as simple as possible. Collaborating with Ute Becker has also been rewarding as she has always been quick to grasp what I coded so that a parallel implementation of the methods shown in this thesis is also on the horizon.

Special thanks go to Prof. Dr. Marcel Nooijen, who patiently answered my questions on MR-EOMCC theory and many-body residual conditions, and to Romain Berraud-Pache for acquiring the experimental spectra in Ch. 5. I further appreciate clarifying comments by Alexander A. Auer, Ondřej Demel, Mihály Kállay, Andreas Köhn, Eric Neuscamman, Masaaki Saitow, Toru Shiozaki, and Alexander Y. Sokolov on the features of code generators in Tab. 4.1.

Two foundations financially supported this project. First, the *Stiftung Stipendien-Fonds des Verbandes der chemischen Industrie e. V.*, which graciously awarded me with a Kekulé scholarship. Second, I was happy to receive a PhD scholarship by the *Studienstiftung des deutschen Volkes e. V.* Apart from the financial support, the highlights were two academies I was able to participate in, one on gravitational waves and one on the topic of blockchains. I am also indebted to the *Max Planck Society* and the *Max Planck Institut für Kohlenforschung*, which provided financial support and a world-class working environment (especially the extremely stable computing infrastructure—thank you Horst Lenk!), which allowed me to entirely focus on my research.

These acknowledgments would not be complete without mentioning my friends and colleagues in Mülheim, who made my time much more enjoyable: Romain Berraud-Pache, Michael Buchsteiner, Pauline Colinet, Nicolas Foglia, Benjamin Helmich-Paris, Małgorzata Krasowska, Lucas Lang, Spencer Léger, Lorenz Löffler, Anastasios Papadopoulos, Eduardo Schiavo, Georgi and Steffy Stoychev, Tim Tetenberg, Van Anh Tran, and Hang Xu. Six of these also served as

proofreaders for this thesis, more precisely, Róbert, Spencer, Anh, Kantharuban, Anastasios, and Pauline.

Special thanks to the 1. BV Mülheim, where I very much enjoyed playing badminton for four years.

Finally, I thank my parents for their loving support during all my time in Mülheim and before. Without them, life and the project would have been infinitely harder.

Short Contents

Contents

List of Figures

List of Tables

1 Introduction

In this thesis, we mainly deal with multireference (MR) coupled-cluster (CC) theories, including different formulations and solution schemes, as well as their implementation through automated tools. We start, however, with a single-reference equation-of-motion (EOM-)CC study of excitation energies on indigo dyes to demonstrate the excellent results obtainable with single-reference (SR) methods. Furthermore, some EOM-CC methods border the region of single- and multireference theories, making the multireference regime accessible even from a single-reference wave function. Compared to the single-reference CC methods, multireference CC theories allow a genuinely multiconfigurational wave function as their zeroth-order reference, which allows them to target the most problematic electronic structures. This flexibility, unfortunately, comes at the expense of highly complicated working equations, which often necessitates automated code generation to implement the theory. Hence, we developed the ORCA-AGE II code generation toolchain, which can generate C++ source code for even the most complicated multireference theories. In this thesis, though, we focus on internally contracted (ic) *ansatzes*, MR-EOMCC and its perturbative variant, MR-EOMPT theory, as well as fully internally contracted (fic-)MRCC theory. Although both MR-EOMCC theory and fic-MRCC theory share an internally contracted *ansatz*, they differ in their solution criteria, which are either *many-body* or *projective* residuals and are theoretically compared in this thesis. Finally, we propose a scheme to *automatically* reduce the scaling of the internally contracted theories with the size of the active space, which reduces the order of the maximum density by one, thereby eliminating the most expensive step for large active spaces, i.e., the construction of the highest-order density.

In the following, we introduce the most important results for the aforementioned research areas and then go on to lay out the theoretical foundations for all theories presented in this thesis. After that, we report our findings for the similarity-transformed equation-of-motion (STEOM) calculations on indigo dyes, the theoretical comparison between many-body and projective residual conditions in MRCC theories, and some details of the ORCA-AGE II toolchain, before discussing the MR-EOMPT method and the automated scheme to reduce the scaling of internally contracted multireference methods. Lastly, we conclude this thesis and give an outlook for the future of automated code generation in quantum chemistry.

1.1 Single-Reference Coupled-Cluster Methods

Coupled-cluster theory is one of the most popular wave function methods, with several hundred studies pertaining to it published every year [3]. In short, this is due to the excellent accuracy

of the results, wide availability of state-of-the-art implementations, and desirable theoretical properties, such as size-consistency [4–6] and -extensivity [7–9]. Consequently, in this thesis, we will use an excited-state variant—STEOM-CC theory [10–12]—to benchmark a family of substituted indigo dyes (Ch. 5). Although CC theory has been reviewed and has appeared in textbooks many times [13–17], we still give a brief overview of the historical and more recent developments and refer to the references for further information.

The CC *ansatz* originates from nuclear physics, where it was conceived by Coester and Kümmel [18, 19]. A few years later, it was introduced to quantum chemistry by the early works of Čížek [20], who was joined by Paldus and Shavitt soon thereafter [21]. A historical perspective on these initial stages was given by both Čížek [22] and Kümmel [23]. The origins of CC theory can further be traced back to many electron theory (MET) by Sinanoğlu [24–26]. Extensive research ensued, which has been reviewed by Bartlett and Musiał [16]. To highlight a few publications, we mention the works of Stanton, Gauss, Bartlett *et al.* [27–30], who were instrumental in popularizing CC theory, in addition to studies on the search for an optimal factorization of the working equations [31, 32] as well as several works focused on an optimal computer implementation [33–36].

In most use cases, CC theory is limited to single and double excitations (SD) as a good compromise between accuracy and the already quite unfavorable scaling of $\mathcal{O}(N^6)$ with system size N. However, to achieve chemical accuracy, triples effects must be included in the treatment. In this regard, the non-iterative perturbative triples-corrected CCSD(T) [37–39] method performs so well that it is widely recognized as the "gold standard" in quantum chemistry, achieving chemical accuracy for bond distances, angles, atomization energies and reaction enthalpies [14, 40, 41]. Even so, the accuracy comes at the price of a higher scaling, $\mathcal{O}(N^7)$, which is further exacerbated by the fact that large basis sets of *at least* triple-ζ quality are required to attain chemical accuracy. The reason for this lies in the poor sampling of the electron cusps in the wave function with Gaussian basis sets, with the theoretical limit given by and named after Kato [42]. For configuration interaction (CI) theory, at least, this simple explanation has been questioned [43]. Two non-exclusive solutions were explored to solve this problem, i.e., local correlation methods—of which we focus on the popular pair natural orbital (PNO)-based ones below—and explicitly correlated approaches.

PNOs as a concept were first formulated by Kutzelnigg [44] as well as Edmiston and Krauss [45, 46]. Further developments involving localization approaches allowed the correlation energy, which can be written as a sum over electron pair contributions, to be divided into contributions of *strong pairs*, which are treated exactly, in addition to *medium* and *weak pairs*, for which either approximate treatments are used, or which are totally neglected. This leads to a compression of information with potentially linear scaling. For a complete discussion, we refer to a review by Werner *et al.* [47]. Specifically for CC theory, early approaches have emerged [48–51], although the PNO-based methods only caught on later, catalyzed through the works of Neese and coworkers [36, 52], who also used resolution of the identity (RI) for the integral transformations

and stored the PNO integrals on disk. These developments culminated in the domain-based, local pair natural orbital (DLPNO) family of methods [53–56].

In explicitly correlated methods, we try to improve the convergence of the method with basis set size by explicit introduction of an inter-electronic r_{12} dependence, first developed by Kutzelnigg, Noga, and Klopper for CC methods [57–59]. Additional work using automated techniques led to the implementation of higher-order methods, for example CCSDTQ-R12 theory by Shiozaki *et al.* [60]. Automated techniques were also used in the development of the alternative F12 family of methods [61–65]. However, as can be seen from the use of automated techniques, these methods can be complicated to implement, and need to be carefully calibrated for cost-effectiveness [66]. The excellent results, though, justify their existence. A fuller account of this research area is given by Valeev and co-workers [67]. Recently, methods that combine both local and explicit correlation have emerged, with the efforts in this direction summarized by Ma and Werner [68].

Let us now turn our attention from ground-state approaches to methods for excited states, which are of particular importance in chemistry, e.g., for spectroscopic properties. To this end, two schools of thought have emerged, equation-of-motion (EOM) CC theory and (linear) response methods, which, at least for excitation energies, give the same results and have recently been reviewed [69–71]. We begin with the EOMCC *ansatz* [72, 73], which was initially described by Emrich [74], although only the standard implementation of Stanton and Bartlett [75] made it popular. Other developments in this area are spin-flip (SF) [76], ionization potential (IP) [77, 78], double ionization potential (DIP) [79], and electron attachment (EA) [80, 81] EOMCC, which can even be used to describe states that otherwise would require a multireference method [82]. Unfortunately, EOMCC still scales as ground-state CC theory, $\mathcal{O}(N^6)$, which precludes application to many systems of interest. Consequently, lower-scaling perturbative methods were formulated [83–87].

For this reason, Nooijen and co-workers devised STEOM-CCSD theory [10–12], a fully non-perturbative method that even includes a higher-order triples term not present in EOM-CCSD [12, 88]. The defining idea lies in the decoupling of the doubles block in the transformed Hamiltonian, which allows for a much smaller diagonalization space as the final step. Thus, the last step effectively becomes a CIS problem, same as in hybrid TD-DFT theory, with a correspondingly lower scaling of $\mathcal{O}(N^4)$. Regardless, some development effort was still required since the mandatory ground-state calculation still scales as $\mathcal{O}(N^6)$, and the ionization potential (IP) [77, 78] and electron attachment (EA) [80] steps both scale as $\mathcal{O}(N^5)$. The scaling of the ground-state calculation was easily solved by using a linear-scaling ground-state method, e.g., the locally correlated DLPNO-CCSD method [54]. This left the IP and EA problems to also be reformulated for the use with PNOs. Several choices of the PNO expansion are possible for the excited states in the IP and EA problems, namely state-specific PNOs [89, 90], state-averaged PNOs [91], or, as used in DLPNO-STEOM, simply the ground-state PNOs [92]. Thus, the transformation steps were recently cast in the locally correlated form for IP [93, 94] and EA [95,

96]. To summarize, the ground-state calculation as well as the IP and EA transformation steps are performed in the DLPNO basis, whereas the cheap final diagonalization remains in the canonical basis. Consequently, DLPNO-STEOM theory has the same formal scaling as TD-DFT, albeit with a larger prefactor, and an accordingly broad applicability. Naturally, the DLPNO-STEOM method was carefully compared to the canonical implementation, with no noteworthy loss in accuracy [92, 93, 95, 97] and has been used to remarkable success for several studies [98–105].

The other approach for excited states, based on response theory, was first presented by Monkhorst [106] in 1977, and later worked out in more detail by Dalgaard and Monkhorst [107]. The first implementation, with an improved theoretical formulation, is due to Sekino and Bartlett [108]. Subsequent developments by Koch *et al.* [109, 110], for example, eventually led to the well-known CC3 model [111, 112], which includes iterative triples and is known to be within the chemical accuracy of full CI calculations [111, 113–116]. We note that response theory is generally not the method of choice for the computation of excitation energies since it comes at an extra cost compared to EOMCC theory, but works better for calculations of other properties [117–119]. Lastly, Sneskov and Christiansen [69] give a fuller picture, and further elucidate the interconnections of EOMCC and linear response theory.

1.2 Introduction to Multireference Coupled-Cluster Theory

After discussing the single-reference CC methods in the previous section, we now turn to the multireference equivalents. Tackling the complexity inherent in multireference correlation approaches is a daunting challenge, which has so far precluded consensus on what features a robust, fast and widely applicable multireference correlation approach should have. This is in stark contrast to the "gold standard," CCSD(T), in single-reference theory. In this thesis, we hope to shed some light on one aspect of the complex multireference problem: the different residual conditions that can be employed in internally contracted multireference coupled-cluster approaches.

Before that, let us turn to the different possible parameterizations. Internal contraction in multireference methods was first thoroughly discussed by Meyer [120] and later by Siegbahn [121] within the multireference configuration interaction (MRCI) framework, although an earlier work by Mukherjee *et al.* already features the essentials in the context of a multireference coupled-cluster (MRCC) method [122]. Since then, it has provided a complementary alternative to the uncontracted parameterizations that have been developed even earlier [123, 124]. In short, the uncontracted approaches have specific excitation operators for each determinant in the multiconfigurational expansion, whereas internal contraction only uses a single operator that affects the multiconfigurational expansion as a whole, acting in the same way on each determinant.

Let us briefly summarize the milestones in the uncontracted approaches. First, the uncontracted MRCI, with several early implementations reported [121, 123–126], is straightforward in structure and has been used to compute highly accurate benchmark data [127, 128]. Apart from

being a variational method, MRCI shares the same theoretical drawbacks as single-reference CI, such as not being size-consistent and -extensive [129, 130], which however may be partially corrected for by approximately size-extensive *pair theories* such as MR-ACPF [131, 132], MR-AQCC [133–136], as well as several uncontracted MR-CEPA approaches [131, 137, 138], among others [135, 139] and Davidson-type corrections [140]. Despite their results being an improvement over MRCI theory [141, 142], these methods did not find widespread adoption. It is worth mentioning that partly contracted formulations also exist that have found wide adoption, such as the formulation of Werner and Knowles [143], which keeps the (semi-)internal classes uncontracted and contracts all other excitation classes.

There has also been an active research community working on uncontracted MRCC methods since its inception by Jeziorski and Monkhorst (JM) [144]. Numerous variants of the JM-type MRCC approach have been published and investigated (we refer to reviews [139, 145, 146] for details), although Mukherjee's state-specific variant (Mk-MRCC) [147–149] has received by far the most attention, both in terms of extensions, such as analytic gradients [150, 151], linear response [152, 153] or local correlation approaches [154, 155] and applications. Several other variants only differ in the *sufficiency conditions* [156], e.g. Brillouin-Wigner BW-MRCC [157–159] and single-root sr-MRCC [160]. Moreover, Malrieu and co-workers [161] describe perturbative theories based on this *ansatz*. A somewhat different approach that cures some theoretical deficiencies shared by the aforementioned methods is Hanrath's MRexpT theory [162, 163]. Nevertheless, these approaches have severe theoretical deficiencies in terms of lacking orbital invariance [164, 165] and proper residual conditions [156, 162], making them unattractive in our opinion. On the one hand, uncontracted approaches have (comparatively) simple working equations, which do not differ too greatly from their single-reference counterparts. On the other hand, a downside shared by all uncontracted theories is that the number of parameters grows proportionally to the number of determinants in the active space, i.e., exponentially with the number of active orbitals.

Internally contracted methods, in contrast to uncontracted approaches, can be rigorously orbital invariant and also have a much-reduced parameter set that does not depend on the number of determinants in the multiconfigurational expansion. Their main drawback does not lie so much in theoretical limitations, but rather in the formidable complexity of the working equations. Only the advent of code generators in quantum chemistry [166] made it feasible to study these equations thoroughly without approximations or allowing only carefully controlled ones.

Turning to the different correlation methods developed with internal contraction, the predominant approaches employ perturbation theory in an internally contracted fashion relying on a complete active space self-consistent field (CASSCF) multiconfigurational wave function. These include complete active space perturbation theory to second order (CASPT2) [167, 168] and the internally (also called partially) contracted variant of n-electron valence perturbation theory to second order (fic-NEVPT2) [169–171]. Since these methods are relatively inexpensive, they have

been applied in a wide range of studies and were also extended to cover explicit correlation [172] and analytic gradients [173]. Moreover, local approaches such as DLPNO-NEVPT2 [55] further extend their scope.

Apart from the early internally contracted MRCI (ic-MRCI) approaches [120, 121, 174], newer developments [175, 176] have focused on using density matrix renormalization group (DMRG) reference wave functions with cumulant reconstruction [177, 178], explicit F12 correlation [179], spin-orbit matrix elements [180], and more efficient implementations [181]. A recent review by Szalay *et al.* discusses these methods at length [182]. Internally contracted MRCI is still not size-consistent and -extensive, though, so Fulde and Stoll [183], Fink and Staemmler [184], and Black and Köhn [185] established several formulations of internally contracted MR-CEPA theory.

Several internally contracted MRCC approaches also exist. One of the earliest methods came about in the form of valence-universal Fock-space (FS-)CC theory [186–189], which is often classified as a separate field of research [139], yet still features an internally contracted cluster operator. ic-MRCC theory was conceived at around the same time [190–193], although further development stalled until much later [194, 195] since deriving and implementing the working equations requires automated symbolic algebra and computational tools due to its enormous complexity. The early approaches contained severe approximations, through truncating the BCH expansion and limiting the excitation classes. Linear dependencies in the projection manifold were also omitted from the discussion entirely [196]. A further implementation was reported by Bartlett and co-workers [197] including only up to linear terms. The *ansatz* was also studied theoretically on several occasions [148, 198, 199].

Despite these early efforts, the first rigorous implementations were reported by Evangelista and Gauss [194], as well as Hanauer and Köhn [195]. Evangelista and Gauss showed for the first time the effects of truncating the BCH expansion and concluded that keeping terms up to quadratic ones offers the best cost-benefit ratio; a conclusion that we also draw upon in our implementation. Additionally, they rigorously proved that the method is invariant with respect to orbital rotations in their respective subspaces. Hanauer and Köhn first succeeded in implementing a polynomial-time version of ic-MRCC and investigated several cluster operators and schemes for discarding linearly dependent amplitudes from the equations. Köhn and co-workers went on to include perturbative triples (T) in the fashion well-known from single-reference CC theory [200] and add F12 explicit correlation [201], response theory and properties [202, 203], among other developments [204–207].

Internal contraction is also at the core of several different theories that may differ in their intended goals from the approaches introduced above. To begin with, Nooijen and co-workers adopted an internally contracted *ansatz* in the state-specific (SS-)EOMCC method [208, 209], partially internally contracted (*p*IC)-MRCC [210] theory, and eventually in the MR-EOMCC [211–213] method. Also, a sequence of similarity transforms with different cluster operators is employed instead of a single one, which has also been investigated in the ic-MRCC

method in the groups of Gauss and Köhn [214]. Nevertheless, MR-EOMCC theory is not entirely contracted, for it ends with an uncontracted MRCI-type diagonalization of the similarity-transformed Hamiltonian over a small excitation manifold to efficiently obtain many excited states; thus being a *transform-then-diagonalize* approach pioneered by Nooijen and Bartlett in the context of the STEOM method [10, 12]. The excitation energies obtained from the MREOM-CC method have been found to provide excellent accuracy [212, 215–217], even for core-valence excitations [218].

The MREOM method also extensively makes use of a general normal ordering with respect to arbitrary reference wave functions, as conceived by Mukherjee [198] and fully developed by Kutzelnigg and Mukherjee (KM) [219, 220], where cumulants arise from $2n$-legged contractions among the operators. This also allowed Datta and Nooijen to employ a set of *many-body residuals* [210], initially proposed by Nooijen and Bartlett [221], which are much simpler in structure than the usual projective conditions used, e.g., in the works of Gauss or Köhn and co-workers [194, 195]. In addition to the reduced cost, the many-body residuals further do not require focus on a particular target state [211]. Spin summation rules have also been reported [222, 223].

Unitary variants of ic-MRCC have also been proposed. Hoffmann and Simons investigated these methods early on [224, 225]. Currently, Yanai and Chan's canonical transformation (CT) theory [226–230] is presumably the most widely known. Since the BCH expansion does not truncate in a unitary formulation, the authors proposed a recursive approximation to the full expansion [226], that later was re-formulated in terms of KM normal order [227]. To further handle the complexity involved in this scheme, the authors also relied on automated tools during the implementation process. The commutator approximation was also investigated by Evangelista and Gauss [231] in the single-reference regime, with proposed corrections for higher accuracy.

The driven similarity-renormalization group (DSRG) approach of Evangelista and co-workers also falls into the internally contracted category [232, 233]. It is related to CT theory inasmuch as it sets the off-diagonal elements of the transformed Hamiltonian to zero, although it is an integration-based approach. Interestingly, the authors decided to use KM normal order and many-body residual conditions [232], much in the same way as in MREOM theory.

The block-correlated coupled-cluster approach by Li and co-workers [234–237] also has an internally contracted cluster operator, although here it is defined with respect to orbital blocks and the active block is described using a Fock space formulation. The authors also relied on a software tool to derive the working equations.

In this thesis, we elucidate the connection between different residual conditions in Ch. 3. Furthermore, the necessity of automated tools when working with MRCC theories is described along with an exposition of our code generation toolchain, ORCA-AGE II, in Ch. 4. We then introduce a perturbative approximation to MR-EOMCC theory—MR-EOMPT theory—in Ch. 6 and work extensively with fic-MRCC theory in Ch. 7.

1.3 Code Generation in Quantum Chemistry

As alluded to above, some developments in the multireference field were only made possible through automated tools. Their advent was in turn aided by the advancements in modern microprocessor technology. Over the last decades, feasible ab initio wave function-based calculations have come from simple, self-consistent field (SCF)-level calculations in small basis sets [238, 239] to regular application of advanced correlated methods, especially the coupled-cluster model, on medium to large molecules in triple-ζ basis sets [240], or even very large molecules when local approximation techniques are employed [54, 241]. This development has been fostered in great part by the advances in CPU processing power, commonly described by Moore's law [242]. It is our opinion that these advances bring more and more computationally demanding theories into the realm of "routinely feasible" computations. These theories include internally contracted [120, 121] multireference theories and gradients of higher-order CC models, the implementation of which is in fact beyond human capacity. Traditionally, a quantum chemical theory would be reformulated for and implemented in computer code entirely manually, which is an onerous and error-prone [243] approach. In light of these challenges, tools have been developed that either simplify or completely automate the implementation process, which we will refer to as automatic code generation. These tools automate at least one of the following general steps needed to go from theory on paper to computer code: (i) derivation of the working equations from an ansatz (equation generation), (ii) manipulation thereof to reduce the computational cost (factorization), and (iii) the actual code generation. We will adopt this distinction throughout this thesis and expand on each of the steps below. For a more detailed introduction to the field of automatic code generation, we also recommend the excellent review of Hirata [166].

Automatic derivation of the working equations from a theoretical ansatz is the most straightforward part, since it mostly relies on a fixed set of rules that can be deterministically applied. Most existing toolchains [60, 64, 177, 178, 181, 199, 209, 228, 229, 244–249] rely on Wick's theorem [17, 250] to obtain the contractions, which was already used very early on in the development of equation generation [251]. Despite this prevalence of Wick's theorem, the first automated equation generators used diagrammatic approaches, mainly to avoid the tedious derivation as well as human errors [252–254]. These early developments, however, lack further equation processing and code generation. A major benefit of diagrammatic generators is that only topologically different contractions are generated, i.e., less work needs to be done in finding equivalent terms in the factorization step [60]. More recent examples of such equation generators are Smith [60] and Kállay and Surján's arbitrary-order CC program [255]. Instead of diagrams or Wick's theorem, the plain (anti-)commutation rules between second-quantized operators can also be straightforwardly applied [14]. This strategy was adopted in the first version of the ORCA-AGE toolchain [1].

The factorization step is arguably the most crucial in the toolchains, since it ensures the proper, minimal computational scaling with system size and significantly reduces the computational cost of the generated code. Unfortunately, finding the global minimum in terms of

computational cost constitutes an NP-hard problem [256]. Hence, virtually all toolchains rely on heuristics to reduce the complexity of this problem. Core concepts were developed early on, e.g., by Janssen and Schaefer [251] or Kállay and Surján [255]. However, the perhaps most complete exposition is given by the tensor contraction engine (TCE) [244, 245], which sequentially uses detection of duplicates, strength reduction, factorization (i.e., application of the distributive law), and common subexpression elimination to arrive at the final working equations before generating code [244]. To aid in detecting duplicate terms and common subexpressions, especially when taking tensor symmetry into account, the tensor contractions are often canonicalized by relabeling and permuting the indices and tensors [244, 251, 256, 257]. A detailed analysis of common subexpression elimination was published by Hartono *et al.* [258]. Overall, these heuristics perform quite well [259], although they will generally not reach the same efficiency as hand-optimized [32] code. More advanced schemes have been discussed [256], but, to the best of our knowledge, only a single tool that uses a genetic algorithm to sample the complete factorization space has been presented to date [259].

Finally, the equations that have been derived and factorized in the previous steps must be evaluated (in the correct order) to arrive at the desired result, which may, for example, be an energy or a residual. To this end, we can either generate code (generally for a compiled programming language) or use an interpreter to evaluate the tensor contractions. Generated code frequently relies on further libraries, most often on the basic linear algebra subroutines (BLAS) [260] to speed up the evaluation of the tensor contractions. BLAS can be extended to arbitrary (binary) tensor contractions [261] and even faster algorithms have been developed for the same sake [262]. As an intermediate between low-level generated code and interpreters, specialized tensor contraction libraries have emerged that more or less completely take care of the computational kernel such that code generation can be greatly simplified. Examples of such libraries include the CTF [263, 264], libtensor [265], LITF [266], TCL [262, 267], and TiledArray [268, 269]. Interpreters even further remove the connection of the contractions to the (compiled) code or hardware by fully abstracting away the latter two, requiring just the contractions and the input quantities. This concept is perhaps best illustrated with the super instruction assembly language (SIAL) of ACES III, which is a fully-fledged virtual machine and parallel programming model [270] used to evaluate generated and handwritten contractions [271, 272]. An integrated tensor framework (ITF) has been reported by Werner and co-workers for the implementation of an internally contracted multireference configuration interaction (ic-MRCI) method [181, 273]. Other toolchains with interpreters using string-based methods [254, 274] include the general contraction engine (GeCCo) by Köhn and co-workers, first to appear in the context of explicitly correlated coupled-cluster methods [64, 249], and Kállay and Surján's arbitrary-order CC program [255, 275].

To conclude this section, we now briefly discuss applications of existing code generators. One of the most complete implementations remains the TCE [244, 245], which encompasses all the steps outlined above. Generated code exists in NWChem [276, 277], and examples

include higher-order CC methods (up to EOM-CCSDTQ) [278] or combined CC theory and MBPT [279]. The Smith generator by Shiozaki and co-workers [60, 246], which can be viewed as the successor to the TCE [280], is another feature-complete tool that was initially used to implement CCSD-R12 [280], and later extended for state-specific [246] and multistate [281] CASPT2 gradients. Around the same time, the SQA, used for the automatic implementation of CT theory [228, 229] and later for the MR-ADC(2) method [282, 283], and FEMTO codes [177, 178, 247] were introduced. FEMTO has specialized features to work with cumulant reconstructions that appear in internally contracted DMRG-MRCI theory, and has been extended for PNOs [247]. Also noteworthy are the GeCCo [64, 249] and APG [199, 209, 248] codes, which have been used to implement the highly complicated ic-MR methods of ic-MRCC [195] theory and MR-EOMCC [210–213] theory, respectively. For a long time, the APG and SQA were the only tools that supported Kutzelnigg-Mukherjee normal order [219], with generalized normal order (GNO) being taken up more recently by the Wick&d program of Evangelista [284]. Last but not least, the first version of the ORCA-AGE toolchain was introduced in 2017 [1]. Different multireference contraction schemes were compared with its support [176], and the STEOM method based on an unrestricted Hartree-Fock (UHF) reference has been implemented [285] with the aid of the APG into the ORCA [286–289] program package.

In this thesis, we describe the completely rewritten ORCA-AGE toolchain (Ch. 4), henceforth referred to as ORCA-AGE II, which has been used extensively for the implementation of fic-MRCC theory as also reported herein (Chs. 3, 7). We must emphasize again that essentially none of the work presented in this thesis would have been possible without this automated toolchain. Furthermore, we will compare its features to other existing solutions and describe a few internal algorithms in App. A.

1.4 Multireference Equation-of-Motion Perturbation Theory

Electron correlation effects play a vital role in calculating molecular properties and taking them into account accurately is often inevitable even to obtain a merely qualitative description of a system. Apart from dynamical correlation effects that arise from the inadequate treatment of short-range electron repulsion at the single-reference level, static correlation effects must also be considered in many important chemical systems. Examples include long-range dissociation problems and situations with near electronic degeneracy [290]. The standard treatment in these cases accounts for the static effects by a relatively small multiconfigurational expansion, which also serves as a starting point for quantitatively accurate multireference methods [291] that may rely on perturbation, configuration interaction or coupled-cluster theory. In contrast to current density functional theory approaches, such treatment can be systematically improved upon, possibly until the point of reaching chemical accuracy [292].

Recently, the MR-EOMCC method [211–213] was proposed for the accurate treatment of strongly correlated electronic states. While this method holds the promise of efficiently cal-

culating multiple states at the same time, some transformation steps involved therein require a significant amount of computational effort to carry through. To address this problem, we propose a multireference perturbation theory (MRPT) approximation to MR-EOMCC, to be referred to as MR-EOMPT. At their core, both MR-EOMCC and its MRPT variant are *transform-then-diagonalize* approaches, where a transformed Hamiltonian is diagonalized over a compact manifold to obtain the states of interest. However, where MR-EOMCC uses iteratively determined amplitudes, the MR-EOMPT method foregoes the iterations and uses estimates from MRPT to yield comparable results at a lower computational cost.

An early approach that shares a close connection [293] to the EOMCC formalism, namely the valence-universal [193, 294, 295] Fock-space coupled-cluster (FSCC) method [186–188], allows the computation of excited as well as ionized or electron-attached states from a single reference state. At its core, FSCC relies on the theory of effective Hamiltonians, where a similarity-transformed Hamiltonian is set up in a reduced model space to reproduce a few eigenvalues of the exact Hamiltonian once the former has been diagonalized. The model space must be judiciously chosen to avoid intruder states, which generally cause serious convergence issues in FSCC calculations [187], although developments such as variants based on the use of intermediate Hamiltonians [296–300] or Mukherjee and co-workers' eigenvalue-independent partitioning technique [301–303] greatly alleviate these issues. This theory further requires a reference for the cluster operators, which is a fully occupied core determinant [187]. The concepts [189, 304–306] introduced with FSCC theory are integral to all *transform-then-diagonalize* methods, of which the single-reference STEOM method has already been mentioned and of which the MR-EOMCC method can be regarded as a multireference extension.

In contrast to the single-reference approaches, genuine multireference methods require a reference that is not just a single determinant or configuration state function (CSF). For example, bond dissociation in N_2 with its triple bond is notoriously difficult to describe with single-reference methods, unless higher-order cluster operators including full triples and quadruples are also considered. For an overview of the competing developments, we refer the reader to a review by Lyakh *et al.* [139] as well as to Sec. 1.2. Here, we are interested in methods geared at computing excited states, which are often multistate approaches. This means that several states are obtained in one go, as opposed to state-specific approaches, in which each state requires a separate calculation. Many of these excited state methods are obtained by extending single-reference techniques to the genuine multireference regime.

As mentioned above, the MR-EOMCC method can be viewed as a multireference extension of the FSCC approach in that it uses a state-averaged (SA) CASSCF wave function as its reference state. Using such a multideterminantal reference became possible after the introduction of generalized normal order, which we reviewed in Sec. 1.2. This many-body aspect of MR-EOMCC (as well as FSCC and STEOM) further removes the need for higher-rank density matrices [210]. Additionally, the many-body residuals can be viewed as the natural extension of FSCC to multireference reference states [221].

The final ingredient to the MR-EOMPT method proposed in this study is multireference perturbation theory. Arguments from MRPT have already been used in an orbital selection scheme for the MR-EOMCC method to reduce computational cost, which is nonetheless different from the novel MR-EOMPT method since the selection scheme retained its iterative nature [217]. In general, perturbation theory is a well-tested approach to improving results compared to SCF calculations, especially when the SCF solution is already close to the exact wave function of the desired state. This assumption can be expected to hold for many CASSCF-based reference wave functions, since CASSCF itself already recovers a large part of static and, to a lesser extent, dynamic electron correlation. To discuss the relation of MR-EOMPT to the other MRPT approaches, we will classify them according to the order of diagonalization or perturbation steps. In this scheme, the most popular methods, state-specific CASPT2 [167, 168] and NEVPT2 [169–171], can be called *diagonalize-then-perturb* methods and generally yield good accuracy. Furthermore, they have been extensively applied to compute excitation energies [307–310], and also properties such as g-tensors of transition metal complexes [311, 312]. CASPT2 was also studied in the extensive benchmark by Thiel and co-workers [313, 314], with results from NEVPT2 reported later on [308]. A major drawback is that CASPT2 and NEVPT2, by construction, rely on the four-body density matrix, which requires significant computational power to construct. To overcome this, both methods have been combined with approximations to enable larger active spaces, such as DMRG [315–317] or cumulant approximations [318–320], a full discussion of which is given below (Sec. 1.5). Local correlation approaches to reduce their overall scaling with system size have also been successfully implemented [55, 321, 322].

From the viewpoint of MR-EOMPT, a more closely related approach is the *diagonalize-then-perturb-then-diagonalize* scheme, which generally describes quasidegenerate (QD) perturbation theory. First, the degenerate subspace of the Hamiltonian is diagonalized to avoid diverging perturbation expressions and the resulting eigenfunctions are then used in the subsequent perturbative treatment and final diagonalization to obtain the states of interest. Examples of this school are multiconfigurational MC-QDPT [323–325], multistate (MS-)CASPT2 [326] and QD-NEVPT2 [327], which have been extensively applied in the investigation of conical intersections [328, 329], among others. In contrast to these methods, however, MR-EOMPT only uses a single set of amplitudes for all the states obtained in the final diagonalization.

The conceptually closest approach to the *transform-then-diagonalize* methods discussed above are the *perturb-then-diagonalize* formulations consisting of a single perturbative step and a subsequent diagonalization step. More precisely, after the effective Hamiltonian is constructed using perturbation theory methods, the final diagonalization mixes the perturbed model space functions into the final states. This approach has benefits over *diagonalize-then-perturb* methods, in which dynamic correlation would mix the CASCI functions significantly, for example in ionic-covalent curve crossings [330, 331] (see also Sec. 6.5). Compared to the more popular methods discussed earlier, relatively few approaches have been developed, including NOCI-MP2 [332, 333] and DCD-CAS [334].

Ultimately, however, the MR-EOMPT method should be considered as a multistate *transform-then-diagonalize* approach. The non-iterative amplitudes are estimated such that the transformed Hamiltonian only couples to few determinants outside of the CAS space, which is then accounted for by a diagonalization over the CAS and remaining first-order interacting space. As will be demonstrated in the following, the results are still satisfactory in terms of accuracy despite these approximations, while being cheaper to compute than fully iterative MR-EOMCC results.

In this thesis, we delineate the MR-EOMPT method in Chapter 6. In that chapter, we first recapitulate the basic ideas behind MR-EOMCC theory, with particular focus on the sequential similarity transformations. Then, we go on to explain how perturbatively estimated amplitudes may be used instead of the iterated equations to obtain the final, transformed Hamiltonian that can be diagonalized over the manifold of holes and particles to obtain the states of interest. We then present benchmarks and test systems from several problem types to validate the performance of MR-EOMPT against accurate reference values, its parent method, MR-EOMCC, and fic-NEVPT2 [171] theory.

1.5 Reduction of Density Matrices in Internally Contracted Theories

As discussed above, multireference self-consistent field (SCF) approaches such as CASSCF [335–338], RASSCF [339, 340], and DMRG [341–347] theory, and even perturbative approaches such as the widely popular CASPT2 [167, 168, 326] and NEVPT2 [169–171, 327, 348] methods, do not always give the desired accuracy, e.g., when computing transition energies [308, 313]. Hence, highly correlated methods have been developed early on, or even prior to the perturbative approaches. Unfortunately, in the multireference field, they all suffer from either theoretical or practical shortcomings that have so far precluded their widespread use. In internally contracted theories specifically, high-order density matrices pose a significant computational challenge, which we will address through a fully automated scheme capable of reducing their order by at least one.

Before going on to discuss reduced-scaling schemes for internally contracted theories, we first reiterate the advantages and disadvantages of internally contracted approaches, compared against uncontracted approaches. The specific developements in those fields have already been discussed at length above (Secs. 1.2, 1.4). First, uncontracted MRCC methods are much simpler to derive and implement than the contracted approaches, their working equations essentially being coupled single-reference problems. However, they have severe theoretical deficiencies in terms of lacking orbital invariance [164, 165] and proper residual conditions [156, 162]. In contrast, internally contracted MRCC theory trivially fulfills these conditions [194]. They also have the major benefit that the number of parameters is largely independent of the active space, or determinants in the reference wave function [143]. However, as alluded to above, a major issue with internally contracted formulations are high-order density matrices, which

scale exponentially [349] in computational time and polynomially in size with growing active spaces [143, 181].

Several techniques that aim at alleviating this computational bottleneck have already been proposed. Broadly speaking, we can either approximate the high-order densities or attempt to obviate their computation by novel methodologies. The most successful approximate scheme is undoubtedly based on cumulant reconstruction techniques, in which higher-order densities are expanded through density cumulants while neglecting the highest-order terms [350–357], which found early adoption in contracted Schrödinger equation approaches to avoid explicit knowledge of the 4-RDM [351, 352, 358, 359]. Cumulant approximations were also successfully used in conjunction with DMRG theory, e.g., in several DMRG-CASPT2 implementations [318, 319, 360] or in DMRG-MRCI theory by Saitow *et al.* [177] to allow for large active spaces. Without the use of DMRG, Zgid *et al.* [320] investigated a cumulant-approximated sc-NEVPT2 implementation and a full comparison of the cumulant approach to prescreening approximations was reported by Guo *et al.* [361], which led to the development of the full-rank (FR-)NEVPT2 method [362]. Also, iterative methods such as CT [226–229] and MR-EOMCC theory [210, 211, 213] use this technique to limit the order of the densities or cumulants to at most two, neglecting all higher orders. A similar scheme is also used for the driven similarity renormalization group (DSRG) method [232, 363–365].

We can also use different approaches to avoid the computation of high-order densities in the first place. The simplest such scheme is to decontract the offending excitation classes leading to the high-order densities, resulting in a mixed contraction scheme. A prominent example is Werner and Knowles's MRCI [143], where only the excitation classes with excitations into two virtual orbitals are contracted, and the rest are left uncontracted. Additionally, Celani and Werner used such an approach for their MRPT2 theory [366]. Further perturbative methods falling into this category are t-NEVPT2 theory [367, 368] and matrix product state perturbation theory (MPS-PT) [369], based on matrix product state linearized coupled-cluster (MPS-LCC) theory [370], which uses conventional internal contraction for all but two excitation classes, which are treated with MPSs.

The MPS representation in DMRG also allows approaches that avoid the explicit computation of the four-body density by implicitly contracting the density with other tensors [315, 316, 370–372]. A projection scheme developed by Roemelt *et al.* for DMRG references also circumvents the construction of the four-body density, but introduces a projection error and slower convergence rate in exchange [373]. The computation can also be avoided by using stochastic approaches, such as Sharma and co-workers' full CI quantum Monte Carlo (FCI-QMC)-based MRCI and -PT [374] and an FCI-QMC-based sc-NEVPT2 implementation by Booth and co-workers [375]. Perturbation theory based on the DSRG method also features a reduction in the order of the density matrices by one order compared to conventional internally contracted perturbation theories [233, 376–379]. We further note that multireference adiabatic connection (AC) can be used towards the same end [349, 380].

Our approach is based on the work by Werner and Knowles [34, 143], who contracted appropriately factorized coupling coefficients with molecular integrals, which avoids the explicit computation of high-order coupling coefficients, while also not introducing any approximations. Chatterjee and Sokolov [381] use a similar trick for NEVPT2 theory, developed in the context of MR-ADC(2) theory. They shift the excitation operators in the densities such that lower-order "intermediate states" can be formed with molecular integrals, with an accordingly reduced scaling. Tangentially, they also use automatic equation derivation through SQA [228, 229] for some aspects of MR-ADC(2), albeit not for automatically reducing the scaling. These ideas were recently picked up in our group for NEVPT2 and CASPT2 theory [382], where the shifting technique of the excitation operators was used to reduce the scaling of the terms containing the highest-order densities. This was, however, achieved manually.

We present our scheme in full detail in Ch. 7 by first recapitulating the origin of the high-order density matrices that appear in internally contracted methods and summarize the foundations of ic-MRCI and ic-MRCC theories. Afterwards, we introduce our fully automated scheme to reduce the order of the density matrices by at least one and describe its implementation in the ORCA-AGE II code generation package. Both fic-MRCI and fic-MRCC theories are then benchmarked on unsaturated alkene chains in terms of both scaling with system and active space sizes.

2 Theoretical Background

2.1 Conventions

In this section, we introduce conventions on how common quantum chemical quantities will be denoted throughout the entire thesis. Starting with molecular orbitals, which are denoted by ϕ or φ, we use the indices i, j, k, l to indicate doubly occupied, inactive orbitals; t, u, v, w to represent active orbitals of varying occupations; a, b, c, d for unoccupied, virtual orbitals; and p, q, r, s as general labels. Notwithstanding, we will sometimes also use the general labels p, q, r, s to refer to *active* orbitals to avoid tedious double-primed labels, e.g., t'', appearing in contractions with high-order densities. Atomic orbitals are referred to by χ. Given no further information, we always assume spatial orbitals. Spin orbitals may also be denoted by appending either α, β, or generic spin, σ, τ. For the spin variable, we generally use ω.

Many-electron quantities, such as wave functions, generally use capitalized letters. From molecular orbitals, we can construct many-electron determinants or configuration state functions (CSFs), which are indicated by Φ, possibly with a capitalized index I, J if referring to a set of them, as in Φ_I. Self-consistent field solutions, which are normally taken to be zeroth-order wave functions, are denoted by Ψ_0, or just 0 for short. Correlated wave functions are referred to as Ψ.

Operators of any sort are indicated by a hat, for example, $\hat{T}$. Tensorial quantities, including vectors and matrices, are denoted by a boldface symbol when referring to it as a whole, as in $\boldsymbol{A}$. However, we will more commonly use indexed notation instead, for example, A_{ij}. We further use atomic units throughout the entire thesis unless stated otherwise.

We also introduce a simplified notation for general n-body quantities X, which may be general tensors or operators with $2n$ indices. If we do not wish to fully specify indices on X, especially when we want to focus on which orders of a quantity can appear in a term, we can write

$$\underbrace{X^{pq\cdots}_{rs\cdots}}_{2n \text{ indices}} \equiv X_n. \tag{2.1}$$

For example, $\hat{E}^p_q \equiv \hat{E}_1$ and $\hat{E}^{pq}_{rs} \equiv \hat{E}_2$. When we use this notation in equations or contractions, we assume that index permutations are implicitly included for simplicity.

2.2 Second Quantization

While quantum mechanics is generally taught in first-quantized form, meaning an explicit representation including position and momentum operators for the involved particles, the language of second quantization if far better suited to formulating advanced correlated theories since it provides a notation free from the number of particles involved. Hence, the entire theory in this thesis is laid out in second quantization.

We therefore introduce the basics of second quantization in this section. After defining the basic operators and the molecular Hamiltonian, we discuss the rules that are used to manipulate the expressions, namely the basic commutation rules as well as normal order in conjunction with Wick's theorem, both in its "simple" and generalized variant.

2.2.1 Spin-Orbital and Spin-Free Operators

The language of second quantization operates in Fock space, which is in turn spanned by the molecular orbitals as the basis functions [14]. A vector in Fock space, called the occupation number vector (ONV), represents our wave function, or Slater determinant,

$$|\Phi\rangle = |x_1, x_2, \ldots, x_n\rangle\,, \tag{2.2}$$

where x_i is the occupation number of the ith *spin* orbital, with a total of n orbitals. For example, the ground state of the H_2 molecule in a minimal basis may be represented as

$$|\Phi_{H_2}\rangle = |1, 1, 0, 0\rangle\,, \tag{2.3}$$

with the molecular orbitals defined as

$$\phi_1(\boldsymbol{r}, \omega) = \varphi_1(\boldsymbol{r})\alpha(\omega) = N_+^{-1/2}[\chi_{1s,A}(\boldsymbol{r}) + \chi_{1s,B}(\boldsymbol{r})]\alpha(\omega), \tag{2.4}$$

$$\phi_2(\boldsymbol{r}, \omega) = \varphi_1(\boldsymbol{r})\beta(\omega) = N_+^{-1/2}[\chi_{1s,A}(\boldsymbol{r}) + \chi_{1s,B}(\boldsymbol{r})]\beta(\omega), \tag{2.5}$$

$$\phi_3(\boldsymbol{r}, \omega) = \varphi_2(\boldsymbol{r})\alpha(\omega) = N_-^{-1/2}[\chi_{1s,A}(\boldsymbol{r}) - \chi_{1s,B}(\boldsymbol{r})]\alpha(\omega), \tag{2.6}$$

$$\phi_4(\boldsymbol{r}, \omega) = \varphi_2(\boldsymbol{r})\beta(\omega) = N_-^{-1/2}[\chi_{1s,A}(\boldsymbol{r}) - \chi_{1s,B}(\boldsymbol{r})]\beta(\omega). \tag{2.7}$$

Here, $N_+^{-1/2}$ and $N_-^{-1/2}$ are the appropriate normalization factors, $\chi_{1s,A}$ and $\chi_{1s,B}$ are the minimal basis atomic orbitals on atoms A and B, respectively, and $\alpha(\omega)$ as well as $\beta(\omega)$ are the spin functions. We further used φ to denote the spatial part of the spin orbitals ϕ. In first quantization, the above ground state of the H_2 molecule may be written in form of an equivalent Slater determinant,

$$|\Phi_{H_2}\rangle = |\phi_1\phi_2\rangle = 2^{-1/2}\begin{vmatrix} \phi_1(\boldsymbol{r}_1, \omega_1) & \phi_2(\boldsymbol{r}_1, \omega_1) \\ \phi_1(\boldsymbol{r}_2, \omega_2) & \phi_2(\boldsymbol{r}_2, \omega_2) \end{vmatrix}. \tag{2.8}$$

As with Slater determinants, the antisymmetry of the fermionic electrons is encoded in the ONV through a phase factor upon swapping the order of two orbitals,

$$|\Phi_{H_2}\rangle = |\phi_1\phi_2\rangle = -|\phi_2\phi_1\rangle\,. \tag{2.9}$$

This phase factor will be used to define the anticommutation rules in Sec. 2.2.4.

Before that, let us define the spin-orbital creation and annihilation operators, $\hat{a}^p \equiv \hat{a}_p^\dagger$ and $\hat{a}_p$, respectively, through their actions on an ONV,

$$\hat{a}^p\,|x_1,\ldots,0_p,\ldots\rangle = X_p\,|x_1,\ldots,1_p,\ldots\rangle\,, \tag{2.10}$$
$$\hat{a}^p\,|x_1,\ldots,1_p,\ldots\rangle = 0, \tag{2.11}$$
$$\hat{a}_p\,|x_1,\ldots,0_p,\ldots\rangle = 0, \tag{2.12}$$
$$\hat{a}_p\,|x_1,\ldots,1_p,\ldots\rangle = X_p\,|x_1,\ldots,0_p,\ldots\rangle\,. \tag{2.13}$$

Here, the indices $p, q, \ldots$ refer to spin orbitals, but we dropped the explicit spin label for readability, i.e., $p \equiv p\sigma$. A creation operator acting on an already occupied orbital as well as an annihilation operator acting on an empty orbital annihilate the entire wave function. In contrast, a creation operator acting on an empty orbital will create it, after which it is occupied, and an annihilation operator acting on an occupied orbital will annihilate it, leaving it unoccupied. In both cases, the complex conjugate transpose versions of Eqs. (2.10)–(2.13) also hold. Furthermore, a phase factor X_p needs to be taken into account for the antisymmetry of the electrons,

$$X_p = \prod_{i=1}^{p-1} (-1)^{x_p}. \tag{2.14}$$

Staying with the minimal basis H_2 system from above, we have, e.g.,

$$\hat{a}^3\,|\Phi_{H_2}\rangle = \hat{a}^3\,|1,1,0,0\rangle = |1,1,1,0\rangle\,, \tag{2.15}$$
$$\hat{a}_2\,|\Phi_{H_2}\rangle = \hat{a}_2\,|1,1,0,0\rangle = -|1,0,0,0\rangle\,. \tag{2.16}$$

Alternatively, this may also be more compactly represented as

$$\hat{a}^3\,|\Phi_{H_2}\rangle = \hat{a}^3\,|\phi_1\phi_2\rangle = |\phi_1\phi_2\phi_3\rangle\,, \tag{2.17}$$
$$\hat{a}_2\,|\Phi_{H_2}\rangle = \hat{a}_2\,|\phi_1\phi_2\rangle = -|\phi_1\rangle\,. \tag{2.18}$$

While the spin-orbital operators form the basis of second quantization, they do not conserve the number of electrons. For example, in Eq. (2.17), a state describing H_2 is turned into one describing H_2^-. Unless our goal is to describe ionization, number-conserving operators are preferred, i.e., operators that have an equal number of creators and annihilators. The prototypical case is the excitation operator $\hat{a}^p\hat{a}_q$, which takes a single electron from spin-orbital q and places it on p. Unfortunately, in the spin-orbital basis, these operators are not in general spin adapted,

i.e., they do not produce a total spin eigenstate when they act on an eigenstate. For example, the operator $\hat{a}^3\hat{a}_1$ acting on $|\Phi_{H_2}\rangle$ produces an impure spin state. This feature is undesirable for spin-adapted theories. Of special interest in such theories are operators that preserve the total spin. This can be achieved by working with spin-free (or spin-traced) excitation operators,

$$\begin{aligned}\hat{E}_q^p &= \sum_\sigma \hat{a}^{p\sigma}\hat{a}_{q\sigma} \\ &= \hat{a}^{p\alpha}\hat{a}_{q\alpha} + \hat{a}^{p\beta}\hat{a}_{q\beta}.\end{aligned} \tag{2.19}$$

By summing over both α and β spin, these operators are rigorously spin-adapted, and hence will not change the spin state. Using these operators, we can excite electrons from the $S = 0$ ground state to the antibonding φ_2 molecular orbital,

$$\begin{aligned}\hat{E}_1^2 |\Phi_{H_2}\rangle &= \hat{E}_1^2 |\varphi_1\alpha, \varphi_1\beta\rangle \\ &= |\varphi_2\alpha, \varphi_1\beta\rangle + |\varphi_1\alpha, \varphi_2\beta\rangle .\end{aligned} \tag{2.20}$$

If we now use CSFs, which are essentially linear combinations of Slater determinants such that a given spin state is achieved [14, 383], along with a theory completely written in terms of the spin-free excitation operators, we have robustly spin-adapted the entire theory. Moreover, for a given configuration, i.e., spatial orbital ONV, the number of CSFs with a given total spin is always less than the number of determinants. This is the main approach followed in the ORCA software as well as the ORCA-AGE code generator. Hence, in the following body of this thesis, we will always assume quantities to be spin-adapted.

2.2.2 The Molecular Hamiltonian

In this section, we will make the connection between first and second quantization for molecular operators, especially for the Hamiltonian. Starting from the full molecular Hamiltonian in first quantization, we first introduce the Born-Oppenheimer approximation, and then partition it into its one-electron (core Hamiltonian) and two-electron parts. These parts can then be represented in the language of second quantization, and the connection to Slater's rules can be made.

The molecular, non-relativistic Hamiltonian for a system of interacting electrons i, j and nuclei A, B reads

$$\hat{H} = -\frac{1}{2}\sum_i \nabla_i^2 - \frac{1}{2}\sum_A \frac{1}{M_A}\nabla_A^2 - \sum_{iA}\frac{Z_A}{r_{ia}} + \sum_{A>B}\frac{Z_A Z_B}{r_{AB}} + \sum_{i>j}\frac{1}{r_{ij}}, \tag{2.21}$$

with (absolute) particle distances indicated by r, and nuclear mass and charge denoted by M and Z, respectively. However, in quantum chemistry, we almost exclusively use the molecular Hamiltonian in the Born-Oppenheimer (clamped nuclei) approximation, which allows us to separate out nuclear motion into a parametric dependence, and thus also gives rise to the

concept of a potential energy surface (PES). The electronic, or Born-Oppenheimer Hamiltonian, simply no longer contains the kinetic energy term of the nuclei,

$$\hat{H}_{\mathrm{BO}} = -\frac{1}{2}\sum_i \nabla_i^2 - \sum_{iA}\frac{Z_A}{r_{ia}} + \sum_{A>B}\frac{Z_A Z_B}{r_{AB}} + \sum_{i>j}\frac{1}{r_{ij}}, \tag{2.22}$$

and can further be separated into its one-electron,

$$\hat{h} = -\frac{1}{2}\sum_i \nabla_i^2 - \sum_{iA}\frac{Z_A}{r_{ia}}, \tag{2.23}$$

and two-electron parts,

$$\hat{V} = \sum_{i>j}\frac{1}{r_{ij}}. \tag{2.24}$$

Now, we can formulate the corresponding one- and two-electron operators in their second-quantized form [14]. First, we give the expressions for the spin-orbital operators (where we tacitly drop the spin label, $p \equiv p\sigma$), beginning with the core Hamiltonian,

$$\hat{h} = \sum_{pq} h_{pq}\hat{a}^p\hat{a}_q, \tag{2.25}$$

$$h_{pq} = \langle p|\hat{h}|q\rangle = \int \mathrm{d}\boldsymbol{r}\, \phi_p^*(\boldsymbol{r})\hat{h}(\boldsymbol{r})\phi_q(\boldsymbol{r}). \tag{2.26}$$

The two-electron operators can be written as

$$\hat{V} = \frac{1}{2}\sum_{pqrs} v_{pqrs}\hat{a}^p\hat{a}^q\hat{a}_s\hat{a}_r, \tag{2.27}$$

$$v_{pqrs} = \langle pq|\hat{V}|rs\rangle = \iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2\, \phi_p^*(\boldsymbol{r}_1)\phi_q^*(\boldsymbol{r}_2)\hat{V}\phi_r(\boldsymbol{r}_1)\phi_s(\boldsymbol{r}_2). \tag{2.28}$$

Note that, even though we have chosen to use the one- and two-electron parts of the Hamiltonian in the definition of the second-quantized operators, the above translation from first to second quantization holds for any one- and two-electron operator, respectively.

The one- and two-electron operators in Eqs. (2.25) and (2.27) reproduce Slater's rules [14, 384, 385]. For example, evaluating the one-electron operator with determinants that differ in up to two orbital replacements, indicated by $\hat{a}^a\hat{a}_i\,|0\rangle = |_i^a\rangle$ and $\hat{a}^a\hat{a}^b\hat{a}_j\hat{a}_i\,|0\rangle = |_{ij}^{ab}\rangle$, we obtain

$$\langle 0|\hat{h}|0\rangle = \sum_i h_{ii}, \tag{2.29}$$

$$\langle 0|\hat{h}|_i^a\rangle = h_{ia}, \tag{2.30}$$

$$\langle 0|\hat{h}|_{ij}^{ab}\rangle = 0. \tag{2.31}$$

For two-electron operators, we obtain

$$\langle 0|\hat{V}|0\rangle = \frac{1}{2}\sum_{ij} v_{ijij} - v_{ijji}, \tag{2.32}$$

$$\langle 0|\hat{V}|_i^a\rangle = \sum_j v_{ijaj} - v_{ijja}, \tag{2.33}$$

$$\langle 0|\hat{V}|_{ij}^{ab}\rangle = v_{ijab} - v_{ijba}. \tag{2.34}$$

To summarize, the molecular, electronic Hamiltonian in second quantization is given by

$$\begin{aligned}\hat{H} &= \sum_{pq} h_{pq}\hat{a}^p\hat{a}_q + \frac{1}{2}\sum_{pqrs} v_{pqrs}\hat{a}^p\hat{a}^q\hat{a}_s\hat{a}_r \\ &= \sum_{pq} h_{pq}\hat{a}^p\hat{a}_q + \frac{1}{2}\sum_{pqrs} \langle pq|rs\rangle\,\hat{a}^p\hat{a}^q\hat{a}_s\hat{a}_r,\end{aligned} \tag{2.35}$$

where we introduced the conventional representation of the two-electron repulsion integrals in physicist's notation, $\langle pq|rs\rangle$. The other popular, chemist's notation, $(pq|rs) = \langle pr|qs\rangle$, is organized by electron number instead of complex conjugation, $(11|22) = \langle 12|12\rangle$.

Since this thesis is concerned with spin-adapted theories, it is beneficial to eliminate spin from the molecular Hamiltonian as well. This can easily be achieved by spin summation, which allows us to write the Hamiltonian in terms of the spin-free excitation operators introduced in Sec. 2.2.1,

$$\hat{H} = \sum_{pq} h_{pq}\hat{E}_q^p + \frac{1}{2}\sum_{pqrs} \langle pq|rs\rangle\,(\hat{E}_r^p\hat{E}_s^q - \delta_r^q\hat{E}_s^p), \tag{2.36}$$

where the indices $p, q, \ldots$ again denote purely the spatial orbitals instead of spin orbitals.

2.2.3 Quantum Chemistry and Density Matrices

Quantum chemistry as a field builds on the axioms of quantum physics, which are[1]

1. An observable is represented by self-adjoint operators, indicated by a hat, e.g., $\hat{O}$, in this thesis.

2. A state of a system is represented by a wave function Ψ, i.e., a unit vector in Hilbert space.

3. The expectation value of $\hat{O}$, if the system is in state Ψ, is the inner product of Ψ and $\hat{O}\Psi$, i.e., $\langle\Psi|\hat{O}|\Psi\rangle$.

[1] This list could be amended by the time-dependent Schrödinger equation as the quantum equation of motion, from which the eigenproblem of the time-independent form derives, and the Pauli (or (anti-)symmetrization) principle.

Each of these axioms can be connected to what we discuss in this section. The operators for observables are constructed from tensors, e.g., molecular integrals, and excitation operators (Secs. 2.2.1 and 2.2.2). States are usually given in terms of an *ansatz* in quantum chemistry, i.e.,

$$\Psi = \hat{\Omega} \left|0\right\rangle , \tag{2.37}$$

where $|0\rangle$ is the reference state and $\hat{\Omega}$ is the ansatz or wave operator. The main criterion for the reference state is that it should have a good overlap with the exact solution, and we will discuss how to find good reference states in Sec. 2.3. The *ansatz* operator generates a better approximation to the true wave function and is specific to a given theory. Based on the reference they act on, they can be single reference approaches discussed in Sec. 2.4 and multireference approaches introduced in Sec. 2.6. In second quantization, $\hat{\Omega}$ is again written in terms of excitation operators acting on these reference wave functions. Consequently, expectation values of excitation operator strings play an important role in quantum chemistry, since their knowledge allows the computation of any observable, e.g., energies, through contractions of the tensorial quantities with expectation values, $\langle 0| \cdots |0\rangle$, of operator strings. These central quantities in quantum chemistry hence have their own names, called *density matrices*.

We define these density matrices, or densities for short, as

$$\gamma_n \equiv \gamma^{tv\cdots}_{uw\cdots} = \langle 0| \hat{E}^t_u \hat{E}^v_w \cdots |0\rangle , \tag{2.38}$$

with n denoting the number of index *pairs*, e.g., $\gamma^{tv}_{uw} \equiv \gamma_2$. These density matrices are *reduced*, since they only refer to n-body quantities. However, reduced density matrices (RDMs) are conventionally defined slightly differently [386–388],

$$\Gamma^{tv\cdots}_{uw\cdots} = \sum_{\sigma\tau\cdots} \langle 0| \hat{a}^{t\sigma} \hat{a}^{v\tau} \cdots \hat{a}_{w\tau} \hat{a}_{u\sigma} |0\rangle . \tag{2.39}$$

The definitions in Eqs. (2.38) and (2.39) are closely related, since either can be written in terms of the other and lower-order density matrices (see Eq. (2.77) for an example). Our choice in this thesis is a purely practical one, motivated by our code generation tool, which generates densities according to Eq. (2.38), not by any further theoretical justification.

In the following, we will give more details on how to evaluate and manipulate operator strings. In later subsections, we also cover the commonly used reference wave functions, before eventually discussing correlated single- and multireference theories.

2.2.4 Commutation Rules

With the definition of the operators of second quantization in Secs. 2.2.1 and 2.2.2, we can now turn to how they can be used to evaluate quantum chemical expressions. To this end, we can employ (anti-)commutation rules among the operators. When repeated in a strategic fashion, these rules can be used to remove all operators from the expressions to be left with only tensorial quantities, such as molecular integrals, densities, and Kronecker deltas.

For the spin-orbital creation and annihilation operators, we have a set of anti-commutation rules,

$$[\hat{a}^{p\sigma}, \hat{a}_{q\tau}]_+ = \delta^p_q \delta^\sigma_\tau, \tag{2.40}$$
$$[\hat{a}^{p\sigma}, \hat{a}^{q\tau}]_+ = 0, \tag{2.41}$$
$$[\hat{a}_{p\sigma}, \hat{a}_{q\tau}]_+ = 0, \tag{2.42}$$

which encode the antisymmetry of the wave function upon interchange of the orbitals [14]. The commutation rule for the spin-free operators may be derived from Eqs. (2.19) and (2.40),

$$[\hat{E}^p_q, \hat{E}^r_s] = \hat{E}^p_s \delta^r_q - \hat{E}^r_q \delta^p_s. \tag{2.43}$$

The commutation rules, by themselves, are of limited usefulness. However, they can be used in conjunction with termination conditions, which remove an operator from the expression. For the spin-orbital operators, these conditions have already been set out in Eqs. (2.11) and (2.12). Note that Eqs. (2.10) and (2.13) do not constitute termination conditions in our sense as they also change the wave function, which is generally undesired when formulating the working equations of a theory. Therefore, to remove spin-orbital operators, we

- move all annihilation operators of unoccupied and all creation operators of occupied orbitals towards the ket, and
- move all annihilation operators of occupied and all creation operators of unoccupied orbitals towards the bra.

More importantly, these rules also apply to the spin-free excitation operators, in a slightly modified form,

$$\hat{E}^i_p \left|0\right\rangle = 2\delta^i_p \left|0\right\rangle, \qquad \left\langle 0\right| \hat{E}^p_i = 2\delta^p_i \left\langle 0\right|, \tag{2.44}$$
$$\hat{E}^p_a \left|0\right\rangle = 0, \qquad \left\langle 0\right| \hat{E}^a_p = 0. \tag{2.45}$$

Eqs. (2.44) and (2.45) form the basis of the equation generator [1] used for all generated theories in this thesis, except for the multireference equation-of-motion (MR-EOM) theories, which use the generalized Wick's theorem instead (Sec. 2.2.6).

The aforementioned rules can be easily demonstrated for both the spin-orbital operators,

$$\langle 0|\hat{a}^{i\sigma}\hat{a}_{j\tau}|0\rangle = -\langle 0|\hat{a}_{j\tau}\hat{a}^{i\sigma}|0\rangle + \langle 0|[\hat{a}^{i\sigma}, \hat{a}_{j\tau}]_+|0\rangle = \delta^i_j \delta^\sigma_\tau, \tag{2.46}$$

and the spin-free excitation operators,

$$\begin{aligned}\langle 0|\hat{E}^j_b \hat{E}^a_i|0\rangle &= \langle 0|\hat{E}^a_i \hat{E}^j_b|0\rangle + \langle 0|[\hat{E}^j_b, \hat{E}^a_i]|0\rangle \\ &= \langle 0|\hat{E}^j_i \delta^a_b|0\rangle - \langle 0|\hat{E}^a_b \delta^j_i|0\rangle = 2\delta^j_i \delta^a_b.\end{aligned} \tag{2.47}$$

The main benefit of using the commutation rules is that they are highly general. Any ansatz that can be formulated in second quantization can be turned into working equations following the approach outlined above. Furthermore, implementing the commutation rules in a software is straightforward, and the computer can easily keep track of the indices and swapped operators. However, when hand-deriving equations, the large number of (intermediate) terms that may appear makes this strategy rather unwieldy, and hence more powerful approaches like Wick's theorem, to be discussed in the following sections, and diagrammatic techniques have been developed.

2.2.5 Normal Order and Wick's Theorem

To simplify the evaluation of strings of excitation operators, we can introduce the concept of *normal order* [15, 17], which, in conjunction with Wick's theorem [250], allows us to directly evaluate the result without having to commute operators multiple times. In this section, we first introduce normal order for strings of excitation operators, both for a true vacuum reference and a closed-shell Slater determinant (particle-hole normal order). After defining the contractions for both formalisms, we outline Wick's theorem and how it can be applied to the evaluation of operator strings. We will also address how normal order and Wick's theorem can be used with spin-free operators.

For a true vacuum reference state in which no orbitals are occupied, $|\rangle$, a string of creation or annihilation operators is said to be in normal order when all creation operators are to the left of all annihilation operators. This order, which is indicated by curly braces around the operator string, guarantees that the expectation value of such a string is zero through annihilation of the reference wave function,

$$\hat{a}^p \hat{a}^q \cdots \hat{a}_r \hat{a}_s \cdots |\rangle = \{\hat{a}^p \hat{a}^q \cdots \hat{a}_r \hat{a}_s \cdots\}_{\text{vac}} |\rangle = 0. \tag{2.48}$$

Here, we again dropped the spin function for readability, $p \equiv p\sigma$.

While the above already gives a working definition of normal order, it is generally more useful in quantum chemistry to choose a (closed-shell) determinant as the reference wave function $|0\rangle$ instead of true vacuum $|\rangle$, as this keeps the operator strings short. This is called particle-hole normal order, and is generally referred to as *normal order.* Here, hole and particle creation operators, $\hat{a}_i$ and $\hat{a}^a$, respectively, are moved to the left of the string, $\langle 0| \hat{a}_i = \langle 0| \hat{a}^a = 0$, whereas hole and particle annihilation operators, $\hat{a}^i$ and $\hat{a}_a$, are moved to the right, $\hat{a}^i |0\rangle = \hat{a}_a |0\rangle = 0$,

$$\hat{a}_i \hat{a}^a \cdots \hat{a}^j \hat{a}_b \cdots |0\rangle = \{\hat{a}_i \hat{a}^a \cdots \hat{a}^j \hat{a}_b \cdots\} |0\rangle = 0. \tag{2.49}$$

Note that, in normal-ordered strings, the usual commutation rules (Sec. 2.2.4) do not apply. Instead, operators can always be swapped with a phase factor of -1 for fermions,

$$\hat{X} = \{\hat{a}^i \hat{a}_a\} = -\{\hat{a}_a \hat{a}^i\}, \tag{2.50}$$

$$\langle 0|\hat{X}|0\rangle = 0. \tag{2.51}$$

This property simplifies reordering of the second-quantized operators, and is used to determine the final phase in contracted operator strings, shown below.

Before discussing Wick's theorem, we first introduce the concept of contractions. A contraction between two operators $\hat{X}$ and $\hat{Y}$ is the difference between the operator string and its normal-ordered form,

$$\hat{X}\hat{Y} = \hat{X}\hat{Y} - \{\hat{X}\hat{Y}\}, \tag{2.52}$$

and is conventionally indicated by a connecting bracket over the contracted operators. In a sense, a contraction can thus be viewed as the result that would be obtained by evaluating the action of the uncontracted operators on the reference wave function. For vacuum normal order, there is only one non-zero contraction,

$$\hat{a}_p\hat{a}^q = \hat{a}_p\hat{a}^q - \{\hat{a}_p\hat{a}^q\}_{\text{vac}} = \hat{a}_p\hat{a}^q + \hat{a}^q\hat{a}_p = \delta_p^q. \tag{2.53}$$

For particle-hole normal order, two non-zero contractions exist,

$$\hat{a}^i\hat{a}_j = \hat{a}^i\hat{a}_j - \{\hat{a}^i\hat{a}_j\} = \hat{a}^i\hat{a}_j + \hat{a}_j\hat{a}^i = \delta_j^i, \tag{2.54}$$

$$\hat{a}_a\hat{a}^b = \hat{a}_a\hat{a}^b - \{\hat{a}_a\hat{a}^b\} = \hat{a}_a\hat{a}^b + \hat{a}^b\hat{a}_a = \delta_a^b, \tag{2.55}$$

which can be generalized to

$$\hat{a}^p\hat{a}_q = x_p\delta_q^p, \tag{2.56}$$

$$\hat{a}_p\hat{a}^q = (1 - x_p)\delta_p^q. \tag{2.57}$$

where x_p denotes the occupation of orbital p (see also Sec. 2.2.1).

Now, we can define Wick's theorem, which states that [17, 250]

> *an operator string can be written as its normal-ordered string, plus the sum of all possible contractions.*

In mathematical notation, this can be represented as

$$\begin{aligned}\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots = &\{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\} \\ &+ \sum_{\text{single}}\{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\} \\ &+ \sum_{\text{double}}\{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\} \\ &+ \text{triple, quadruple, … contractions.}\end{aligned} \tag{2.58}$$

In a more general form for a product of two normal ordered operator strings, the contractions can only run between operators from different normal-ordered strings [15, 389],

$$\begin{aligned}\{\hat{A}\hat{B}\hat{C}\cdots\}\{\hat{M}\hat{N}\cdots\} = \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\} \\ + \sum_{\text{single}} \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\} \\ + \sum_{\text{double}} \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\} \\ + \text{triple, quadruple, … contractions.}\end{aligned} \tag{2.59}$$

Note that Eq. (2.59) is often referred to as a *generalized* form of Wick's theorem in the literature [15]. In this thesis, however, we make no distinction in nomenclature between Eqs. (2.58) and (2.59). Instead, we use the terms *generalized* normal order and *generalized* Wick's theorem to refer to the multireference generalizations put forward by Mukherjee and Kutzelnigg [219], to be detailed in Sec. 2.2.6.

Wick's theorem is a useful tool to quickly evaluate the expectation value of (products of) operator strings. This is achieved by the fact that only the fully contracted terms can contribute to the expectation value, as all other remaining normal-ordered strings annihilate the reference wave function. For example, the expectation value of a one-electron operator in particle-hole normal order is

$$\begin{aligned}\sum_{pq} h_{pq} \langle 0|\hat{a}^p \hat{a}_q|0\rangle &= \sum_{pq} h_{pq} \langle 0|\{\hat{a}^p \hat{a}_q\}|0\rangle \\ &= \sum_{pq} h_{pq} x_p \delta^p_q = \sum_i h_{ii}.\end{aligned} \tag{2.60}$$

Although Wick's theorem is primarily used with spin-orbital creation and annihilation operators, equivalent formulations for spin-free excitation operators have been derived [390, 391]. Kutzelnigg also worked out the respective equations in particle-hole normal order [392, 393]. However, a major drawback of these formalisms is that they are more complicated to use than the simpler spin-orbital formalism, and thus have not gained widespread traction. An alternative path, which first converts all operators to the spin-orbital operators through Eq. (2.19), then contracts the operators according to the rules in Eqs. (2.53)–(2.57), and finally performs the spin summation will give the same results and may even be easier to execute.

In passing, we remark that even more powerful diagrammatic techniques [394–396] have been developed based on Wick's theorem and have found success especially in the development of coupled-cluster [17, 20, 389, 397, 398] and many-body perturbation theory [17, 398, 399]. Automated approaches based on diagrammatic techniques have also been reported [252–255, 400]. In our opinion, however, it is more feasible to rely on the commutation rules (and Wick's theorem) for a generally applicable equation generator, especially considering the multireference regime. Consequently, we will not further discuss diagrammatic methods.

2.2.6 Generalized Normal Order and Generalized Wick's Theorem

While Wick's theorem is highly useful in evaluating operator strings, it is unfortunately limited to a single reference determinant. This shortcoming was addressed by Mukherjee and Kutzelnigg, who developed a *generalized* normal order and the corresponding contractions and *generalized* Wick's theorem [219, 220, 222, 401], which allows arbitrary linear combinations as reference wave functions, for example, those appearing in a CASSCF wave function. In this section, we will outline the foundations of generalized normal order, and refer the reader to the original works for the derivations.

While the theory has been worked out in Refs. [219, 220, 222, 401], the contraction rules as we present them have been worked out in Ref. [402] and summarized in Ref. [233]. Most importantly, we have single contractions, which now result in a density γ_q^p or hole density η_q^p,

$$\hat{a}^p \hat{a}_q = \gamma_q^p, \tag{2.61}$$

$$\hat{a}_p \hat{a}^q = \delta_p^q - \gamma_p^q = \eta_p^q, \tag{2.62}$$

and additionally $2k, k = 2, 3, \ldots$-legged contractions, which are associated with cumulants,

$$\hat{a}^p \hat{a}^q \cdots \hat{a}_s \hat{a}_r = \lambda_{rs\cdots}^{pq\cdots}. \tag{2.63}$$

Although strictly speaking cumulants are defined for $2 \leq k$, we will also use the convention $\lambda_q^p = \gamma_q^p = \delta_q^p - \eta_q^p$. Compared to vacuum or particle-hole normal order, the quantities resulting from the contractions are more complicated objects, since they may either be matrices with off-diagonal entries, or, in the case of cumulants, higher-dimensional tensors.

The cumulants arise as a consequence of the multideterminantal expansion of the wave function [219]. Initially, they were thought of as the "irreducible" parts of the higher-order densities, i.e., those parts which cannot be constructed as combinations of lower-order density matrices [354]. This view may have been influenced by the structure of the defining equations, e.g., for the two-body cumulant [219],

$$\lambda_{qs}^{pr} = \gamma_{qs}^{pr} - \gamma_q^p \gamma_s^r + \gamma_s^p \gamma_q^r. \tag{2.64}$$

This view, however, has been revised more recently. Instead, the cumulants are now interpreted as measures of correlation between orbital occupation numbers [403–405]; a view that, unfortunately, precludes a simple pictorial interpretation. Moreover, this revised interpretation is more in line with the observation that higher-order cumulants need not necessarily have small norms in genuine multireference situations [405].

The computation of higher-order cumulants is essentially similar to that for the two-body cumulant (Eq. 2.64). It has been formalized by using Grassmann algebra, which gives access to cumulants of arbitrary order [358, 405–408].

Despite the new contraction rules for $2k$-legged contractions, Wick's theorem essentially remains unchanged from its presentation in Sec. 2.2.5. All that is required is to include the new contraction types into its formulation,

$$\begin{aligned}\{\hat{A}\hat{B}\hat{C}\cdots\}_{\text{GNO}}\{\hat{M}\hat{N}\cdots\}_{\text{GNO}} = \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\}_{\text{GNO}} \\ + \sum_{\text{single}} \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\}_{\text{GNO}} \\ + \sum_{\text{double}} \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\}_{\text{GNO}} \\ + \sum_{\text{4-legged}} \{\hat{A}\hat{B}\hat{C}\cdots\hat{M}\hat{N}\cdots\}_{\text{GNO}} \\ + \text{triple, 6-legged, quadruple, … contractions.} \end{aligned} \tag{2.65}$$

Although the additions compared to Eq. (2.59) are minor, we find it useful to refer to this form as *generalized Wick's theorem* to emphasize when we refer to multireference theories later on in this thesis.

Generalized normal order has also been worked out for the spin-free cases [219], but is once again more complicated to use than in the spin-orbital formalism. In addition to more complicated contraction rules, spin-adaptation of the cumulants is still an open problem for cumulants of order > 4 [222, 223, 354, 409, 410]. A major reason for the difficulties is that cumulants, in contrast to reduced density matrices, are not expectation values of molecular operators [223].

The relative novelty of generalized normal order, along with its unclear spin adaptation for cumulants, has so far led to but a few applications. Most notable are multireference equation-of-motion coupled-cluster (MR-EOMCC) theory [211–213] and the driven similarity renormalization group (DSRG) method [232, 233, 363, 365]. When it comes to automated code generation, only few codes can generate equations based on generalized normal order, e.g., the automated program generator (APG) [199, 248, 411–413] and Wick&d [284]. However, due to the more complicated spin-free formalism, the APG internally converts all expressions to the spin-orbital formalism, and recovers the spin-free working equations only after a final spin summation step [210]. Furthermore, the APG limits the order of the appearing cumulants to second order at most [209, 211]. Higher-order density matrices are not entirely neglected, though, but rather represented through a cumulant expansion [353–355, 414, 415] truncated at second order [209].

2.3 Self-Consistent Field Theory

All highly accurate, correlated methods are based on zeroth-order wave functions, which are computed with a self-consistent field method. In this section, we reiterate the basics of (restricted

closed-shell) Hartree-Fock (RHF) theory as the basis for all single-reference calculations presented in this thesis, as well as the complete active space self-consistent field (CASSCF) method, which is the most popular multiconfigurational approach and exclusively used in this thesis.

2.3.1 Hartree-Fock Theory

In 1928, Hartree developed the concept of a *self-consistent field,* an iterative procedure that can be used to approximately solve the Schrödinger equation for many-electron systems [416, 417]. However, the initial theory was lacking the exchange contributions that arise from the physical indistinguishability of electrons, a fact that was pointed out by both Fock [418, 419] and Slater [420]. These developments were subsequently consolidated into Hartree-Fock theory [421], which still is one of the most regularly used methods in quantum chemistry, be it as a zeroth-order guess for correlated methods or, in modified form, in density functional theory through the Kohn-Sham equations [422]. In this section, we give a brief introduction, mainly since we use restricted Hartree-Fock as the starting point for our single-reference coupled-cluster calculations. Further details can be found in the excellent textbook by Szabo and Ostlund [384].

Hartree-Fock theory uses a single Slater determinant (Eq. 2.8) to represent a wave function, which automatically incorporates electron exchange into the equations. The energy of such a determinant (in spin-orbital basis) can be derived using the Slater-Condon rules [423, 424],

$$E_{\mathrm{HF}} = \sum_i h_{ii} + \frac{1}{2} \sum_{ij} \langle ij \,||\, ij \rangle \,. \tag{2.66}$$

The main goal is then to variationally minimize the energy expression, Eq. (2.66), with respect to the molecular orbitals ϕ_i. The molecular orbitals, in turn, are given as a linear combination of atomic orbitals, which was a critical development by Roothan [425] that made computer implementation feasible,

$$\phi_i = \sum_\mu c_{\mu i} \chi_\mu . \tag{2.67}$$

The condition that will be required of a variational minimum is that the Hartree-Fock energy is stationary with respect to the orbital coefficients,

$$\frac{\partial E_{\mathrm{HF}}}{\partial c_{\mu i}} = 0 \quad \forall c_{\mu i} . \tag{2.68}$$

For details of the derivation, we refer the reader to Refs. [14, 384]. The central result of the derivation is an eigenvalue equation involving the Fock matrix $\boldsymbol{F}$,

$$\boldsymbol{FC} = \boldsymbol{ESC}, \tag{2.69}$$

with $\boldsymbol{S}$ being the orbital overlap matrix and the Fock matrix depending on the orbital coefficients $\boldsymbol{C}$ itself [384],

$$F_{\mu\nu} = h_{\mu\nu} + \sum_{i} \sum_{\lambda\sigma} c_{\lambda i} c^{*}_{\sigma i} \left[2(\mu\nu|\sigma\lambda) - (\mu\lambda|\sigma\nu) \right]. \tag{2.70}$$

This dependence, $\boldsymbol{F} = \boldsymbol{F}(\boldsymbol{C})$, of the Fock matrix on the solutions of Eq. (2.69), is at the heart of the self-consistent field procedure. All calculations begin with a guess at the MO coefficients $\boldsymbol{C}$, or rather the density matrix [384]

$$\boldsymbol{P} = 2\boldsymbol{C}\boldsymbol{C}^{\dagger}. \tag{2.71}$$

Afterwards, Eq. (2.69) is solved for a new set of orbitals $\boldsymbol{C}'$, then these new MO coefficients are used to build a new density and Fock matrix, at which point the procedure starts again with solving Eq. (2.69). Normally, but this is hard to guarantee, this iterative procedure is carried out until the calculation converges to a *self-consistent* solution, i.e., the MO coefficients obtained by solving Eq. (2.69) are the same that entered the construction of the Fock matrix $\boldsymbol{F}$ in that iteration. Eventually, the energy and other quantities of interest are computed from the final, self-consistent set of orbitals.

2.3.2 Complete Active Space Self-Consistent Field Theory

Single-reference approaches like Hartree-Fock theory, or even correlated methods such as (single-reference) CI or CC methods, have well-documented shortcomings, e.g., when dissociating bonds or in (low-spin) transition metal complexes. These situations, in which several configurations contribute with large weights to the wave function, are usually referred to as problems with large *static* or *nondynamic* correlation. Quite commonly, the terms static and nondynamic are used interchangeably, with some authors further defining that static correlation refers to the fixed linear combination of determinants to obtain proper spin states, and nondynamic correlation to the remaining contributions [13]. In this thesis, we will generally talk about nondynamic correlation as we consistently use CSFs throughout, which already address the issue of static correlation.

In essence, these situations appear straightforward to deal with, since the deficiencies of single-reference approaches can easily be cured by simply constructing a wave function as a linear combination of the important configurations,

$$|0\rangle = \sum_{I} c_I \, |\Phi_I\rangle \,, \tag{2.72}$$

where c_I are the expansion coefficients. However, the important configurations, or CSFs $|\Phi_I\rangle$, are generally not known *a priori*. Choosing the important configurations, especially for expensive (MR)CI calculations in the early days of quantum chemistry [426, 427], used to be more of an art than science.

Therefore, researchers do not pick configurations by hand nowadays, but rather use multi-configurational self-consistent field (MCSCF) methods that will variationally optimize both the orbitals as well as the configurations or CSFs in the expansion from Eq. (2.72). Perhaps the most popular method is the complete active space self-consistent field (CASSCF) method, developed by Roos, Taylor, and Siegbahn [335], who, among others, also worked on the convergence schemes for this method [336–338]. Its popularity is also due in part to the CASSCF method serving as the basis for correlated methods such as CASPT2 [167, 168, 326] and NEVPT2 [169–171, 327, 348]. The historical developments leading up to this theory have been summarized by Roos [428].

In CASSCF calculations, we do not directly choose the important configurations, as discussed above. Instead, an *active space* is chosen, which encompasses a certain number of active electrons, n_{el}, and active orbitals, n_{act}, denoted as CAS($n_{\mathrm{el}}, n_{\mathrm{act}}$). This spans an expansion as in Eq. (2.72), which includes all configurations or CSFs that satisfy any of the electronic arrangements of the specified multiplicities possible in the active space. Their respective weights are then simultaneously optimized with the orbitals, resulting in a variational energy that is stationary with respect to both optimization parameters,

$$E_{\mathrm{CASSCF}} = \sum_{pq} \Gamma_q^p h_{pq} + \sum_{pqrs} \Gamma_{qs}^{pr} (pq|rs), \tag{2.73}$$

$$\frac{\partial E_{\mathrm{CASSCF}}}{\partial c_I} = 0 \quad \forall c_I, \tag{2.74}$$

$$\frac{\partial E_{\mathrm{CASSCF}}}{\partial c_{\mu p}} = 0 \quad \forall c_{\mu p}. \tag{2.75}$$

In the above equations, we used the one- and two-body reduced density matrices Γ_q^p and Γ_{qs}^{pr}, respectively,

$$\Gamma_q^p = \langle 0 | \hat{E}_q^p | 0 \rangle , \tag{2.76}$$

$$\Gamma_{qs}^{pr} = \frac{1}{2} \langle 0 | \hat{E}_q^p \hat{E}_s^r - \hat{E}_s^p \delta_q^r | 0 \rangle . \tag{2.77}$$

The density matrices are important quantities in essentially any quantum chemical calculation. In CASSCF, perhaps even more so since they do not trivially reduce to Kronecker deltas in the active space, but instead compactly represent the information from the multiconfigurational wave function. Their importance will become more obvious in multireference correlated methods, where even higher-order density matrices appear.

Unfortunately, converging CASSCF calculations is more complicated than Hartree-Fock calculations for they need to minimize the energy with respect to two coupled sets of parameters, the expansion coefficients and the orbitals, which leads to many (spurious) minima. In short, we are not guaranteed to obtain the global minimum, and are required to carefully assess each CASSCF wave function with respect to the final active space. Hence, CASSCF is not considered to be a black-box method, requiring substantial input from the researcher.

A major drawback of the CASSCF method is inherent to the simplicity of the complete active space, i.e., a factorial scaling of CSFs in the expansion (2.72), which limits the size of the active space to generally below 16 active orbitals in routine calculations. For example, a singlet CAS(16,16) calculation involves millions of configurations and tens of millions of CSFs.

This matter motivated the search for other multiconfigurational methods that do not show this scaling with the active space. These are the density matrix renormalization group (DMRG) [341–347] and full CI quantum Monte Carlo (FCIQMC) [429] methods, as well as several iterative CI approaches [129, 430–439], which were pioneered by the "configuration interaction by perturbation with multiconfigurational zeroth-order wave function selected by iterative process" (CIPSI) [430] method. These methods no longer consider the full CASSCF expansion, but rather attempt to approximate it to varying degrees.

Consequently, these methods only produce approximated densities, which can lead to numerical issues in subsequent calculations. In both CASPT2 and NEVPT2 theory, false intruder states have been documented [318, 320, 361, 440]. A recent investigation of this issue for the NEVPT2 method further narrowed down the reason for the false intruders to the eigenvalue equation of the zeroth-order Hamiltonian, which no longer explicitly holds with approximate methods, and proposed the full-rank (FR-)NEVPT2 method as a solution [362].

Lastly, we make a distinction between state-specific and state-averaged CASSCF calculations. A state-specific calculation only targets a single state as in Eq. (2.72), which is normally confined through a given multiplicity or symmetry. On the other hand, the wave function of a state-averaged calculation is optimized such that all included states k are best described by a single set of orbitals and the same model space, but different CI expansion coefficients c_I^k,

$$|k\rangle = \sum_I c_I^k \, |\Phi_I\rangle \, . \tag{2.78}$$

This, in turn, leads to reduced density matrices for each of the states k,

$$\Gamma_q^{p\,(k)} = \langle k | \hat{E}_q^p | k \rangle \, , \tag{2.79}$$

$$\Gamma_{qs}^{pr\,(k)} = \frac{1}{2} \langle k | \hat{E}_q^p \, \hat{E}_s^r - \hat{E}_s^p \delta_q^r | k \rangle \, , \tag{2.80}$$

and weighted, averaged density matrices

$$\overline{\Gamma}_q^p = \sum_k w_k \Gamma_q^{p\,(k)}, \quad \overline{\Gamma}_{qs}^{pr} = \sum_k w_k \Gamma_{qs}^{pr\,(k)}. \tag{2.81}$$

that can then be used in the evaluation of the energy expression from Eq. (2.73), as well as for state-specific, correlated methods such as fic-NEVPT2, fic-MRCI, and fic-MRCC. Naturally, the weights of the individual states w_k must sum to unity, $\sum_k w_k = 1$. We note that the state-averaged CASSCF calculation is variational with respect to the averaged state, but not with respect to the individual states k.

2.4 Single-Reference Coupled-Cluster Theory

With the groundwork on self-consistent field methods now laid out, we go on to discuss advanced, correlated methods based on restricted Hartree-Fock reference wave functions, more specifically, single-reference CC theory.

2.4.1 Ground-State Theory

To begin with, let us first briefly introduce a linear parameterization, CI theory, first. We define the linear excitation operator containing single and double excitations in the usual way,

$$\hat{T} = \sum_{ia} t_a^i \hat{E}_i^a + \frac{1}{2} \sum_{ijab} t_{ab}^{ij} \hat{E}_i^a \hat{E}_j^b. \tag{2.82}$$

One of the most straightforward ways of dealing with the electron correlation makes use of this truncated excitation operator to define the configuration-interaction with singles and doubles, CISD, expansion of the wavefunction as

$$|\Psi_{\mathrm{CI}}\rangle = (1 + \hat{T})|0\rangle. \tag{2.83}$$

The parameters in $\hat{T}$ are then determined variationally by making the expectation value of the energy stationary under the constraint that the wavefunction is (intermediate) normalized. This results in the following residual conditions

$$\langle 0|(\hat{H} - E)(1 + \hat{T})|0\rangle \stackrel{!}{=} 0, \tag{2.84}$$

$$r_i^a = \langle 0|\tilde{E}_a^i(\hat{H} - E)(1 + \hat{T})|0\rangle \stackrel{!}{=} 0, \tag{2.85}$$

$$r_{ij}^{ab} = \langle 0|\tilde{E}_{ab}^{ij}(\hat{H} - E)(1 + \hat{T})|0\rangle \stackrel{!}{=} 0. \tag{2.86}$$

These equations are projections of the Schrödinger equation for Ψ_{CI} to the excitation spaces involved in the wave function. The operators $\hat{E}$ define the covariant basis for the ket-functions in our formulation. Since this basis is not orthonormal, it is often necessary to define a contravariant basis for bra functions such that

$$\langle 0|\tilde{E}_{cd\ldots}^{kl\ldots}\hat{E}_{ij\ldots}^{ab\ldots}|0\rangle = \langle\tilde{\Phi}_{kl\ldots}^{cd\ldots}|\hat{E}_{ij\ldots}^{ab\ldots}|0\rangle = \delta_{kl\ldots}^{ij\ldots}\delta_{cd\ldots}^{ab\ldots}, \tag{2.87}$$

where $\hat{E}_{ij\ldots}^{ab\ldots}$ is a product of several $\hat{E}$ operators defined in Eq. (2.19), $\tilde{E}_{ij\ldots}^{ab\ldots}$ denotes the generator of the corresponding contravariant functions $\tilde{\Phi}_{kl\ldots}^{cd\ldots}$, and $\delta_{kl\ldots}^{ij\ldots}, \delta_{cd\ldots}^{ab\ldots}$ are products of Kronecker deltas.

This method has fallen into disfavor due to its lack of size-consistency and size-extensivity, despite having the advantage of being a variational method. To amend these shortcomings, the coupled-cluster method was introduced, relying on an exponential parametrization,

$$|\Psi_{\mathrm{CC}}\rangle = \mathrm{e}^{\hat{T}}|0\rangle. \tag{2.88}$$

In this case, a variational solution is impractical since the expansion of the exponential leads to an infinite series. It is thus usual to define the similarity-transformed Hamiltonian as

$$\bar{H} = \mathrm{e}^{-\hat{T}} \hat{H} \mathrm{e}^{\hat{T}}, \tag{2.89}$$

which has a commutator expansion terminating at fourth order, $N = 4$, in the closed shell (more generally, single-reference spin-unrestricted) case,

$$\begin{aligned} \bar{H} &= \sum_{k=0}^{N} \frac{1}{k!} [\hat{H}, \hat{T}]_k = \hat{H} + [\hat{H}, \hat{T}] + \frac{1}{2}[[\hat{H}, \hat{T}], \hat{T}] + \cdots \\ &= \sum_{k=0}^{N} \left(\hat{H} \hat{T}^k \right)_C . \end{aligned} \tag{2.90}$$

Eq. (2.90) is well known as the Baker-Campbell-Hausdorff (BCH) expansion, from which it is obvious that only the connected parts survive, indicated by the subscript C. This expansion also appears in the multireference case (Sec. 2.6.3), although the expansion only terminates at higher orders N.

The Schrödinger equation of the transformed Hamiltonian can still be solved using the projection technique,

$$\langle 0 | (\bar{H} - E) | 0 \rangle \stackrel{!}{=} 0, \tag{2.91}$$

$$r_i^a = \langle 0 | \tilde{E}_a^i (\bar{H} - E) | 0 \rangle \stackrel{!}{=} 0, \tag{2.92}$$

$$r_{ij}^{ab} = \langle 0 | \tilde{E}_{ab}^{ij} (\bar{H} - E) | 0 \rangle \stackrel{!}{=} 0, \tag{2.93}$$

However, in this case, the projection solution is not obtained from the variation of the expectation value of the Hamiltonian with Ψ_{CC}. Although it is possible to relate these projection conditions to the variational CC ones and it is also possible to formulate a variational problem from which these projection equations follow [441], they are not unique in the sense the projection solutions are unique for CI. Other methods of solution have also been investigated [442], although the traditional projection approach remains the method of choice.

Normal ordering does not change this picture much in the single-reference case, although it certainly makes the evaluation of the second quantized strings easier. The creators and annihilators in the excitation operators of the closed shell case in Eq. (2.82) commute, and thus they are essentially normal ordered. This is no longer the case with open shell excitation operators, which is why they are more difficult to deal with. The Hamiltonian can of course be written in a normal ordered form,

$$\hat{H} = h_0 + \hat{H}_N, \tag{2.94}$$

$$h_0 = \langle 0 | \hat{H} | 0 \rangle = E_0, \tag{2.95}$$

$$\hat{H}_N = h_p^q \{ \hat{E}_q^p \} + \frac{1}{2} h_{pr}^{qs} \{ \hat{E}_{qs}^{pr} \}. \tag{2.96}$$

Switching to $\hat{H}_N$ in the above equation thus amounts to a shift with the (HF) reference energy E_0, and then the rules of evaluating normal-ordered products may be applied in evaluating matrix elements. In practice, this evaluation is usually performed in the simpler spin-orbital form.

For further details, we refer to a detailed introduction by Crawford and Schaefer [15], as well as to textbooks [14, 17].

2.4.2 Equation-of-Motion Coupled-Cluster Theory

Coupled-cluster theory, as outlined in Sec. 2.4.1, is a ground-state method. To access excited states, the EOMCC method has been developed, which uses a linear parameterization to this end. Because of the projection conditions, the singles and doubles manifolds are decoupled from the ground state and the excited states can be obtained by diagonalizing the Hamiltonian in the space of single and double excitations. This would yield an eigenproblem for the energy of the κth excited state, E_κ and the linear excitation operator $\hat{R}_\kappa$ that produces the κth excited state from the CC ground state. Instead, the EOM equations are formulated in a way to yield the κth excitation energy $\omega_\kappa = E_\kappa - E_0$ directly,

$$\langle 0|\tilde{L}_\lambda^\dagger[\bar{H}, \hat{R}_\kappa]|0\rangle = \omega_\kappa \delta_{\lambda\kappa}, \quad \langle 0|\tilde{L}_\lambda^\dagger \hat{R}_\kappa|0\rangle = \delta_{\lambda\kappa}. \tag{2.97}$$

Since this is a non-symmetric eigenproblem, the left-hand solutions $\hat{L}_\lambda$ are not the same as the right-hand ones $\hat{R}_\kappa$, but it is not necessary to calculate both if only the excitation energy is needed. Depending on the precise definition of $\hat{R}_\kappa$, the above eigenvalue problem can be solved for excitation energies (EE), ionization potentials (IP) and electron affinities (EA),

$$\hat{R}_\kappa^{\mathrm{EE}} = \sum_{ia} r(\kappa)_a^i \hat{E}_i^a + \frac{1}{2}\sum_{ijab} r(\kappa)_{ab}^{ij} \hat{E}_i^a \hat{E}_j^b, \tag{2.98}$$

$$\hat{R}_\kappa^{\mathrm{IP}} = \sum_{i} r(\kappa)^i \hat{a}_i + \sum_{ijb} r(\kappa)_b^{ij} \hat{E}_j^b \hat{a}_i, \tag{2.99}$$

$$\hat{R}_\kappa^{\mathrm{EA}} = \sum_{a} r(\kappa)_a \hat{a}^a + \sum_{jab} r(\kappa)_{ab}^{j} \hat{E}_j^b \hat{a}^a, \tag{2.100}$$

where the various $r(\kappa)_{a\ldots}^{i\ldots}$ are the EOM excitation/ionization amplitudes. The $\tilde{L}_\lambda$ operators have a similar structure except they are formulated in terms of contravariants. In the EE case, the computational cost scales as the sixth power, in the IP/EA cases as the fifth power of the system size.

2.4.3 Similarity-Transformed Equation-of-Motion Coupled-Cluster Theory

The STEOM method utilizes a second similarity transformation, compared to EOM-CC theory, to decouple the singles manifold from the doubles and higher-order excitations,

$$\hat{G} = \{\mathrm{e}^{\hat{S}}\}^{-1} \bar{H} \{\mathrm{e}^{\hat{S}}\}, \tag{2.101}$$

where the operator $\hat{S}$ is obtained from the solutions of the IP/EA equations

$$\hat{S} = \hat{S}^{\mathrm{IP}} + \hat{S}^{\mathrm{EA}}, \tag{2.102}$$

$$\hat{S}^{\mathrm{IP}} = \sum_{i'm} s_m^{i'} \{\hat{E}_{i'}^m\} + \sum_{ijmb} s_{mb}^{ij} \{\hat{E}_j^b \hat{E}_i^m\}, \tag{2.103}$$

$$\hat{S}^{\mathrm{EA}} = \sum_{a'e} s_{a'}^e \{\hat{E}_e^{a'}\} + \sum_{jabe} s_{ab}^{ej} \{\hat{E}_j^b \hat{E}_e^a\}. \tag{2.104}$$

Here, the indices m and e correspond to active orbitals in the occupied and virtual spaces, respectively, while the corresponding inactive orbitals are denoted as i' and a'. The selection of the active orbitals within the STEOM-CCSD method is discussed elsewhere, and it can be made entirely automatic [97, 443]. Finally, the connection between the STEOM amplitudes $s_{\cdots}^{\cdots}$ and the IP/EA amplitudes is given by

$$s_m^{i'} = -\sum_\kappa r^{-1}(\kappa)_m \, r(\kappa)^{i'}, \quad s_{mb}^{ij} = -\sum_\kappa r^{-1}(\kappa)_m \, r(\kappa)_b^{ij}, \tag{2.105}$$

$$s_{a'}^e = \sum_\kappa r^{-1}(\kappa)^e \, r(\kappa)_{a'}, \quad s_{ab}^{ej} = \sum_\kappa r^{-1}(\kappa)^e \, r(\kappa)_{ab}^{\ j}, \tag{2.106}$$

where the summation includes a handful of states κ in the active space. Once the transformation is carried out, $\hat{G}$ can be diagonalized in the subspace of single excitations. On this basis, the STEOM working equations can be derived and implemented along the lines first indicated by Nooijen and Bartlett [12].

2.5 Domain-Based Pair Natural Orbitals

The DLPNO scheme has been applied with success to obtain ground state CC energies for molecules containing hundreds of atoms [54]. To this end, it is necessary to reduce the size of the amplitudes and integrals involved in the CC equations; a process [53, 54] that will be broadly outlined here. At the basis of PNO schemes is the realization that the CC correlation energy E_{cor} can be written as a sum of pair energies ε_{ij} for each distinct occupied orbital pair ij,

$$E_{\mathrm{cor}} = \sum_{i \leq j} \varepsilon_{ij}. \tag{2.107}$$

If the occupied orbitals are localized using an appropriate localization scheme [444, 445], then the number of pair energies above a certain threshold scales linearly with system size. These pairs can be treated at the CC level, while the smaller contributions can be accounted for at a lower level of theory [54].

The next step is to choose a compact representation for the virtual space. Since atomic orbitals (AO) are local by construction, a suitable basis for the virtual space may be found by a

projection that removes the occupied components from the AO basis. The resulting projected atomic orbitals (PAO) [446] are given by

$$|\tilde{\mu}\rangle = \sum_{\mu} \tilde{P}_{\mu\tilde{\mu}} |\mu\rangle \,, \tag{2.108}$$

where μ is an AO and $\tilde{\mu}$ is a (redundant) PAO. $\tilde{P}_{\mu\tilde{\mu}}$ is an element of the PAO coefficient matrix defined as

$$\tilde{\boldsymbol{P}} = \mathbf{1} - \boldsymbol{C}_O \boldsymbol{C}_O^\dagger \boldsymbol{S} \,, \tag{2.109}$$

with $\mathbf{1}$ being the unit matrix, $\boldsymbol{S}$ the AO overlap matrix and $\boldsymbol{C}_O$ the occupied block of the molecular orbital (MO) coefficient matrix. A set of linearly independent or nonredundant PAOs can be conveniently constructed from the functions $\tilde{\mu}$ by standard orthonormalization/canonicalization techniques. Next, a list of PAOs can be generated for each occupied orbital i, called the domain of i. This will include all members of those shells of PAOs for which at least one PAO has a differential overlap integral (DOI) with i above a certain threshold,

$$\mathrm{DOI}_{i\tilde{\mu}} = \sqrt{\int \mathrm{d}\boldsymbol{r} \, |i(\boldsymbol{r})|^2 |\tilde{\mu}(\boldsymbol{r})|^2} . \tag{2.110}$$

The domain for an occupied pair ij is then simply the union of the domains of i and j.

The pair domains obtained in the previous step can be used to span the virtual subspace that interacts significantly with a given pair. To achieve this, a perturbative (or other approximate) pair density $\boldsymbol{D}^{ij}$ is needed,

$$\boldsymbol{D}^{ij} = \tilde{\boldsymbol{T}}^{ij} \boldsymbol{T}^{ij\dagger} + \tilde{\boldsymbol{T}}^{ij\dagger} \boldsymbol{T}^{ij}, \quad \tilde{\boldsymbol{T}}^{ij} = \frac{1}{1+\delta_{ij}} (4\boldsymbol{T}^{ij} - 2\boldsymbol{T}^{ij\dagger}), \tag{2.111}$$

where we use the first-order Møller-Plesset amplitudes for $\boldsymbol{T}^{ij}$, which is the simplest choice. The PNOs can be obtained by diagonalizing the pair density spanned in the basis of PAOs,

$$\tilde{D}^{ij}_{\tilde{\mu}_{ij}\tilde{\nu}_{ij}} = \sum_{\tilde{a}_{ij}} d^{ij}_{\tilde{a}_{ij}\tilde{\mu}_{ij}} n^{ij}_{\tilde{a}_{ij}} d^{ij}_{\tilde{a}_{ij}\tilde{\nu}_{ij}} , \tag{2.112}$$

where $\tilde{\mu}_{ij}$ is a nonredundant PAO belonging to the pair domain of ij, and $d^{ij}_{\tilde{a}_{ij}\tilde{\mu}_{ij}}$ is the matrix that transforms the PAOs into the PNOs $\tilde{a}_{ij}$ of a given pair ij. The fact that for any given occupied pair there are only a few PNOs that are larger than a threshold is the basis on which the CC equations can be made linear scaling. The steps outlined above provide a series of transformations that transform the canonical equations into the DLPNO basis. Used together with the resolution of identity integral decomposition scheme [447, 448], all quantities can be reduced to a size as compact as the physics encoded in the pair density matrix allows them to be. The various thresholds necessary to define these transformations have been studied elsewhere and predefined defaults are available [449]. The settings we will use in this study have also been described in an earlier work [103].

The procedure described so far makes the ground state linear scaling, but that still leaves the IP and EA steps that scale with the fifth power of the system size. It is relatively straightforward to implement the IP equations in the DLPNO basis, and the resulting code will be nearly linear scaling since all intermediates can be expressed using ground state DLPNO quantities. The detailed description of this implementation can be found elsewhere [93]. In contrast, the EA equations are more problematic since the EA singles are difficult to assign to any occupied pair or orbital. Thus, these equations are partly kept in the canonical basis for a better control of accuracy. As a result, the algorithm is a quadratic scaling one. Some terms in the EA equations also need special treatment, and they are expanded in the (larger) diagonal PNO basis. The accuracy and efficiency of this algorithm has also been investigated in the literature and it does remove the bottleneck from the DLPNO-STEOM process [95]. Since the final step of DLPNO-STEOM is essentially a configuration interaction singles (CIS) calculation, it is also very easy to obtain approximate properties and the required one-body densities for DLPNO-STEOM [103, 450]. The CIS approximation used with success in the past [103] will also be used in Ch. 5 to account for solvent effects within the conductor-like polarizable continuum model (CPCM) [451]. A somewhat more involved formalism [450] can be used to calculate transition dipoles that are needed for spectral intensities.

2.6 Multireference Methods

Now, we can move on to methods based on multiconfigurational reference wave functions. To this end, we first recapitulate the concept of internal contraction and discuss the fic-MRCI and fic-MRCC theories, before stating the salient features of MR-EOMCC theory. We finish this section by briefly describing how high-order density matrices appear in internally contracted approaches.

2.6.1 Internal Contraction

Internal contraction was first developed for MRCI theory by Meyer and Siegbahn [120, 121], and its general form can be written as

$$|\Psi\rangle = \hat{\Omega}\,|0\rangle = \hat{\Omega}\sum_I c_I\,|\Phi_I\rangle\,, \tag{2.113}$$

where $|0\rangle$ is a multiconfigurational zeroth-order wave function, most likely obtained from a CASSCF calculation, c_I are the expansion coefficients, and $|\Phi_I\rangle$ are individual configuration state functions (CSFs), or alternatively determinants. $\hat{\Omega}$ is a general wave operator depending on the exact theory, e.g., $\hat{\Omega} = \exp(\hat{T})$ for coupled-cluster theory. In summary, the defining feature of internal contraction is that a single wave operator acts *as a whole* on the zeroth-order

wave function, rather than having one wave operator per CSF as in uncontracted approaches,

$$|\Psi\rangle = \sum_I \hat{\Omega}^{(I)} c_I \, |\Phi_I\rangle \, . \tag{2.114}$$

Eqs. (2.113)–(2.114) further illustrate that internally contracted theories have a much smaller parameter set for long CI expansions, i.e., calculations with large active spaces. Hence, given that we can find a way to eliminate high-order densities, internally contracted methods appear to be the natural choice for systems with large active spaces.

2.6.2 Fully Internally Contracted Multireference Configuration Interaction Theory

As stated above, the first reports of internally contracted MRCI theory are due to Meyer and Siegbahn [120, 121]. However, since this theory shares the lack of size-consistency and -extensivity with its single-reference counterparts, it has never become widely used within the computational chemistry community. The most popular approach is arguably the *partially* contracted MRCI implementation by Werner and Knowles [143], which leaves all but the doubly external classes uncontracted.

In our implementation, we use the linear parameterization with $\hat{\Omega} = 1 + \hat{T}$, and

$$\begin{aligned}
\hat{T} = &\frac{1}{2} \sum_{ijab} t_{ij}^{ab} \hat{E}_i^a \hat{E}_j^b + \sum_{ijta} t_{ij}^{at} \hat{E}_i^a \hat{E}_j^t + \frac{1}{2} \sum_{ijtu} t_{ij}^{tu} \hat{E}_i^t \hat{E}_j^u \\
&+ \sum_{itab} t_{it}^{ab} \hat{E}_i^a \hat{E}_t^b + \sum_{itau} t_{it}^{au} \hat{E}_i^a \hat{E}_t^u + \sum_{itua} t_{it}^{ua} \hat{E}_i^u \hat{E}_t^a \\
&+ \sum_{ituv} t_{it}^{uv} \hat{E}_i^u \hat{E}_t^v + \frac{1}{2} \sum_{tuab} t_{tu}^{ab} \hat{E}_t^a \hat{E}_u^b + \sum_{tuav} t_{tu}^{av} \hat{E}_t^a \hat{E}_u^v,
\end{aligned} \tag{2.115}$$

including all eight excitation classes from CASPT2 and NEVPT2 theory [167, 171]. Eq. (2.115) can also be graphically depicted to more clearly show the different orbital ranges involved in the excitation classes (Fig. 2.1).

The energy and residual equations are virtually identical to the single-reference case, save for the fact that the reference wave function $|0\rangle$ is a CASSCF wave function,

$$E = \langle 0|\hat{H}|\Psi\rangle \, , \tag{2.116}$$

$$r_P = \langle \tilde{\Phi}_P|\hat{H} - E|\Psi\rangle \stackrel{!}{=} 0, \tag{2.117}$$

with a tilde indicating the (possible) use of contravariant projection functions [33, 452]. In addition to contravariant projection functions being used, the projection functions are not linearly independent in the multireference case. The scheme used by us to remove redundant excitations was published by Sivalingam *et al.* [176].

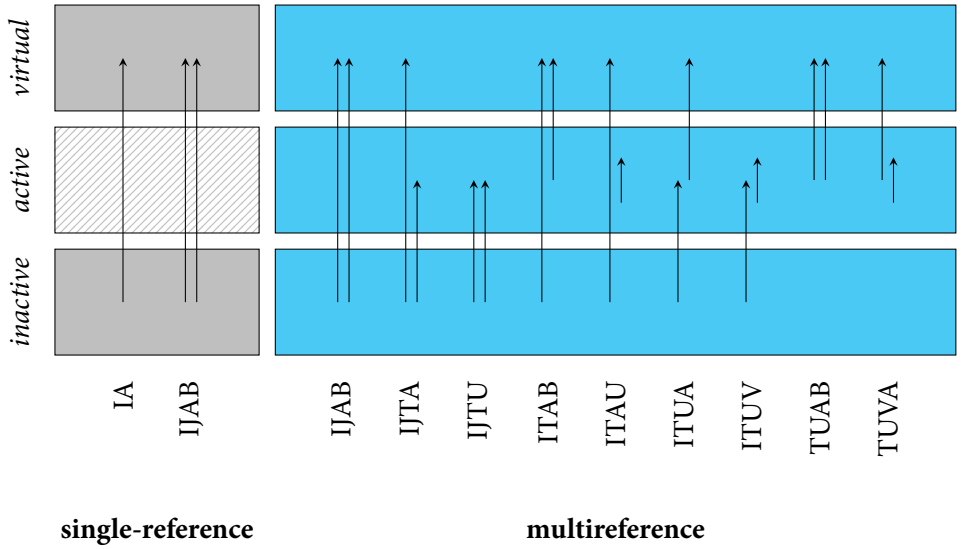

Figure 2.1 Graphical representation of the excitation operators from Eq. (2.115).

2.6.3 Fully Internally Contracted Multireference Coupled-Cluster Theory

Throughout quantum chemistry, we always strive for faster and more accurate theories to aid in our understanding of chemical systems. As mentioned before, CCSD(T) has been quite universally established as the "gold standard" in the single reference domain. In the multireference regime, however, no clear-cut "best" approach exists yet. Various approaches to transfer the coupled-cluster *ansatz* to multireference wave functions have been introduced [139], some dating back to the late 1970s and early 80s [144, 188, 190, 453]. In our group, we focus on the internally contracted variant since it has several desirable properties, such as orbital invariance [194], a limited parameter space [195], and size-extensivity [454]. Furthermore, ic-MRCC theory is known to be significantly more accurate than ic-MRCI theory, even when the BCH expansion of the similarity-transformed Hamiltonian is truncated at quadratic terms [185].

Given a zeroth-order wave function $|0\rangle$ obtained from a run of a complete active space self-consistent field (CASSCF) calculation, the *ansatz* for the coupled-cluster wave function can be written as

$$|\Psi\rangle = \mathrm{e}^{\hat{T}}\,|0\rangle\,, \tag{2.118}$$

with the cluster operator $\hat{T}$ identical to the one we used for fic-MRCI theory, Eq. (2.115). The energy and residuals are computed as in single-reference CC theory as the expectation value of the similarity-transformed Hamiltonian or by projecting the transformed Hamiltonian with

functions from all excitation classes P,

$$E = \langle 0|\bar{H}|0\rangle\,, \tag{2.119}$$

$$r_P = \langle \tilde{\Phi}_P|\bar{H}|0\rangle \stackrel{!}{=} 0. \tag{2.120}$$

In Eqs. (2.119)–(2.120), we truncate the BCH expansion of the similarity-transformed Hamiltonian after quadratic terms,

$$\bar{H} \approx \hat{H} + [\hat{H}, \hat{T}] + \frac{1}{2}[[\hat{H}, \hat{T}], \hat{T}], \tag{2.121}$$

although the expansion would only naturally terminate after octuple powers, $\hat{T}^8$ [195]. The reason for this long expansion lies in the active indices, which can easily be contracted in chains among themselves, delaying the termination of the BCH expansion when compared to closed shell single-reference CC theory (where it terminates after $\hat{T}^4$, see Sec. 2.4.1). Computing the full expansion, however, would be prohibitively expensive, so Evangelista and Gauss [194] investigated the issue and reported that truncating the BCH after quadratic terms leads to minimal losses in accuracy.

MRCC theory further requires a scheme to deal with linear dependencies among the projection functions, which is introduced by the overlaps of functions with active orbital indices. The problem is much more complicated than in SR-CC theory, where linear dependencies can trivially be removed by restricting the summation indices in the cluster operator, $i < j$ (Eq. 2.82), and with contravariant projection functions. In MRCC theory, on the other hand, a new, orthonormal basis must be found for the excitation operators, which is conventionally done through Löwdin's canonic orthonormalization [385]. The scheme that we use throughout this thesis was presented by Sivalingam *et al.* [176], and a fuller discussion of the possible alternative schemes is given by Köhn and co-workers [195, 455].

Moreover, implementing fic-MRCC theory is much more involved than SR-CC theory, requiring automated tools for the sheer number of terms that arise from Eqs. (2.119)–(2.120). A full discussion of our solution is given in Sec. 4.3. We also note the somewhat different behavior of normal order in MR theories compared to the SR case discussed in Sec. 2.4.1. The Hamiltonian can essentially be written in the form of Eqs. (2.94)–(2.96), except now using GNO. This implies that the matrix elements h will contain up to two-body cumulants. In the MR case, the excitation operator in Eq. (2.115) is not in normal order either. For a generic cluster operator of excitation level N, the normal order can be given formally as

$$\hat{T} = \sum_{p\ldots q\ldots}^{N} t^{q\cdots}_{p\ldots} \hat{E}^{p\cdots}_{q\ldots} = \sum_{p\ldots q\ldots}^{N} \overline{t}^{q\cdots}_{p\ldots} \{\hat{E}^{p\cdots}_{q\ldots}\}_{GNO}, \tag{2.122}$$

where the cluster amplitudes $\overline{t}$ in normal order consist of products of the original amplitudes t and at most N-body cumulants. While we do not make use of normal order in our implementation, these properties will be useful for the discussion in Ch. 3.

2.6.4 Multireference Equation-of-Motion Coupled-Cluster Theory

The MR-EOMCC approach is an excellent tool to target multiple excited states simultaneously, with highly accurate results [210–213, 215, 217]. Over time, several variations on the main concept have been reported, but an extensive benchmark by Huntington *et al.* [215] points to the MR-EOM-T|SXD|U procedure as the overall best in accuracy. Since this variant was also chosen as the basis for the MR-EOMPT method reported in Ch. 6, we reiterate its salient features here. Furthermore, in the following we will use the term "MR-EOMCC" to refer to this specific variant, if not stated otherwise.

The method relies on several, sequential similarity transformations with different cluster operators $\hat{O}$,

$$\bar{H}' = \{\mathrm{e}^{\hat{O}}\}^{-1} \bar{H} \{\mathrm{e}^{\hat{O}}\}, \tag{2.123}$$

where $\bar{H}$ is a transformed Hamiltonian from a previous step, or just the Born-Oppenheimer Hamiltonian in the first step. Eq. (2.123) also implies an internally contracted approach, for there is only one set of amplitudes for the state-averaged CASSCF reference wave function (Eq. 2.78). These transformation steps are the hallmark of *transform-then-diagonalize* methods, to which both MR-EOMCC and STEOM (Sec. 2.4.3) belong. Their main aim is to decouple blocks of the Hamiltonian, which includes all double excitations in case of MR-EOMCC theory (see also Fig. 2.1, multireference classes). Then, the final states are obtained by a diagonalization of the transformed Hamiltonian in a small space of only singles excitations, making MR-EOMCC a multistate method. Further details of the equations are discussed in Sec. 6.1.1.

Compared to fic-MRCC theory (Sec. 2.6.3), MR-EOMCC theory uses *many-body* instead of projective residual conditions (Eq. 3.18). This ties in with MR-EOMCC theory being a multistate method and further reduces the maximum order of density matrices that appear in the working equations (see Sec. 2.6.5). The many-body residuals are extensively discussed and compared to the projective residuals, for which we refer to Ch. 3 for full details.

2.6.5 Density Matrices and Their Appearance in Internally Contracted Methods

In internally contracted theories, the highest-order densities arise from matrix elements forming the σ-vector or residuals. These matrix elements can generally be written as

$$r_P = \langle 0 | (\hat{\tau}'_P)^\dagger \hat{H} \hat{\tau}^n | 0 \rangle \,, \tag{2.124}$$

where $\hat{\tau}, \hat{\tau}'_P$ are (potentially contravariant [33, 452]) excitation operators, e.g., $\hat{\tau}^{ab}_{ij} = \hat{E}^a_i \hat{E}^b_j$, and n indicates the order of the expansion ($n = 1$ for CI theory, $n > 1$ for CC theory [195]). We further use the spin-free molecular Hamiltonian in the Born-Oppenheimer (clamped nuclei) approximation, Eq. (2.36). In the matrix elements from Eq. (2.124), the index ranges to which the indices on the excitation operators $\hat{\tau}$ belong depend on the excitation class, with more active indices directly translating to higher-order densities, for inactive and virtual labels only give

Kronecker deltas upon contraction. Additional active labels can arise from the Hamiltonian unless they are forced to be inactive or virtual such as to obtain non-zero matrix elements. All these active indices then give rise to high-order densities, provided that the number of upper and lower indices match. For example, terms involving the four-body density γ_4 are contained ($\leftarrow$) in the linear matrix element

$$\begin{aligned} r_{tu}^{ab} &= \langle 0|(\hat{E}_t^a \hat{E}_u^b)^\dagger \hat{H} \hat{E}_{t'}^{a'} \hat{E}_{u'}^{v'}|0\rangle \\ &\leftarrow \sum_{t''u''v''} \delta_b^{a'} (at''|u''v'')\gamma_{t't''v''u'}^{ut''u''v'}, \end{aligned} \tag{2.125}$$

with a' contracted with b and a contracted with an index from the Hamiltonian, thus appearing in the two-electron integral. The remaining active indices, including the double-primed indices from the Hamiltonian, then form (at most) the four-body density. Nevertheless, not all indices from the Hamiltonian are required to be active, thus densities from γ_2 to γ_4 appear in the full set of equations.

As a final point, high-order densities can be avoided to a certain degree in internally contracted theories by using a many-body approach [188, 221], examples of which can be found in MR-EOMCC theory [210] and the DSRG method [232, 233]. The key step is to avoid projection as in Eqs. (2.124) and (2.125), which adds further excitation operators, $\hat{\tau}'_P$, to the residuals, by expanding the transformed Hamiltonian in many-body form, Eq. (3.5), and setting the many-body elements to zero, $\overline{h}_{qs}^{pr} = 0$. For similarly constructed *ansatzes* with up to quadratic amplitudes $\hat{T}^2$, this leads to only the four-body density being required in many-body approaches [209, 210], compared to the five-body density in projective schemes. [195] However, many-body theories cannot be considered state-specific any longer, and intruder states arise more easily [211, 232].

2.7 Vibronic Effects

So far, we have exclusively discussed methods for determining the electronic part Ψ of the total wave function $\overline{\Psi}$ of a system, which also includes a nuclear contribution. In other words, Ψ is an (approximate) eigenfunction of the Born-Oppenheimer Hamiltonian, Eq. (2.22), whereas the total wave function $\overline{\Psi}$ is an (approximate) eigenfunction of the full molecular Hamiltonian, Eq. (2.21). Since chemical behaviour depends primarily on electronic structure, much of quantum chemistry focuses on finding solutions of the electronic part. However, when comparing to experimental reference data, e.g., thermochemistry, kinetics or in the form of spectra, vibrational and rotational effects must be included at least via the simplest rigid-rotor-harmonic-oscillator model [456]. When calculating electronic spectra in particular, vibronic, i.e., vibro-electronic, coupling effects must often be accounted for. In the following, we will outline a simple scheme to include the most important vibronic effects in spectroscopic calculations and thus provide a link between some of the involved electronic structure theory discussed in this thesis and measurable data.

The starting point for our treatment [457–459] of vibronic effects is the Fermi-Golden rule for periodic interactions which establishes that the rate of transition between an initial state $|\bar{\Psi}_i\rangle$ and a final state $|\bar{\Psi}_f\rangle$ coupled by an operator $\hat{W}$ is

$$k(\omega)_{if} = \frac{2\pi}{\hbar}|\langle\bar{\Psi}_i|\hat{W}|\bar{\Psi}_f\rangle|^2\delta(E_i - E_f \pm \hbar\omega), \tag{2.126}$$

where the plus sign corresponds to absorption leading to a final state with energy $E_f = E_i + \hbar\omega$ and the minus sign to emission with $E_f = E_i - \hbar\omega$ as a result of interacting with a photon of energy $\hbar\omega$. Within the dipole approximation, the coupling operator becomes a dipole operator, and the entire spectrum $\sigma_\pm$ is given by summing over thermally weighted initial and all final states,

$$\sigma_\pm(\omega) = \alpha_\pm(\omega)\sum_{if} P_i(T)|\langle\bar{\Psi}_i|\hat{\mu}|\bar{\Psi}_f\rangle|^2\delta(E_i - E_f \pm \hbar\omega), \tag{2.127}$$

where $\alpha_\pm$ is a frequency dependent prefactor [459] that differs for absorption and emission, and $P_i(T)$ is the Boltzmann population of the initial states. The Dirac delta function, δ, may be replaced by various line shape functions to account for line broadening. Integrating this expression with respect to ω yields the rate of the radiative processes,

$$k_\pm = \int \mathrm{d}\omega\, \sigma_\pm(\omega). \tag{2.128}$$

To evaluate this formula, the inverse Fourier transform of the Dirac delta is substituted into the rate expression and the wavefunction is separated into an electronic (Ψ) and nuclear (vibrational) part (Θ). Denoting final states by a prime, and assuming a single initial and a single final electronic state, the total wavefunctions can be written as $|\bar{\Psi}_i\rangle = |\Psi\Theta_i\rangle$ and $|\bar{\Psi}_f\rangle = |\Psi'\Theta'_f\rangle$, which leads to the following expression,

$$\sigma_\pm(\omega) = \alpha'_\pm(\omega)\int \mathrm{d}t\ \chi_\pm(t)\mathrm{e}^{\pm i\omega t}, \tag{2.129}$$

where all constants are contained in $\alpha'_\pm$ [459]. This expression can be discretized and evaluated using the discrete Fourier transformation technique. The correlation function $\chi_\pm(t)$ is given by

$$\chi_\pm(t) = \sum_i \langle\Theta_i|\ |\mu_e|^2\mathrm{e}^{-i\hat{\mathscr{H}}\tau}\mathrm{e}^{-i\hat{\mathscr{H}}'\tau'}|\Theta_i\rangle\mathrm{e}^{\mp i\Delta E t}, \tag{2.130}$$

where ΔE is the electronic energy difference between initial and final states, τ and τ' are appropriately shifted time-like variables [459], and $\hat{\mathscr{H}}$ and $\hat{\mathscr{H}}'$ are the vibrational Hamiltonians of the initial and final states. The integral $\mu_e = \langle\Psi|\hat{\mu}|\Psi'\rangle$ is the electric transition dipole moment, which also depends on nuclear positions. The normal coordinates $\boldsymbol{Q}$ of the initial and $\boldsymbol{Q}'$ of the final states are related by a linear transformation $\boldsymbol{Q} = \boldsymbol{J}\boldsymbol{Q}' + \boldsymbol{K}$, where $\boldsymbol{J}$ is the Duschinsky rotation matrix and $\boldsymbol{K}$ is the displacement vector. Thus, the integral μ_e may be regarded as a function of either of these coordinate sets and it can be expanded as a power

series. Truncating this series at zeroth order yields the Franck-Condon approximation, while keeping the first-order terms as well yields the (first-order) Herzberg-Teller approximation. To specify $\boldsymbol{J}$ and $\boldsymbol{K}$, it is necessary to choose a model for the initial and final state potential energy surfaces (PES) [460]. Within the harmonic approximation, the initial state PES V can be expanded around its equilibrium geometry using the diagonal matrix of frequencies $\boldsymbol{\Omega}$ calculated at that geometry, $V(\boldsymbol{Q}) = \frac{1}{2}\boldsymbol{Q}^{\dagger}\boldsymbol{\Omega}^{2}\boldsymbol{Q}$. This still leaves us with a choice for the final state PES V'. Adiabatic models expand $\overline{V}$ around the final state equilibrium geometry, while in vertical models the expansion is performed around the initial state geometry. In one of the simplest models, the so-called vertical gradient (VG) approximation, it is further assumed that the final state Hessian is identical to the initial state one. This yields $\boldsymbol{J} = \mathbf{1}$ and $\boldsymbol{K} = -\boldsymbol{\Omega}^{-2}\boldsymbol{g}'$, where $\boldsymbol{g}'$ is the gradient of the excited state surface at the initial state geometry. This means that $\boldsymbol{Q}$ and $\boldsymbol{Q}'$ are simply displaced, and that quadratic differences between V and V' are neglected.

For the indigo dyes considered in Ch. 5, Barone and co-workers [461] have compared the VG model to the adiabatic Hessian (AH) approach, in which the final state Hessian is also calculated at the final state equilibrium geometry, and found that the VG model results agree well with AH ones for this family of dyes. Since the AH model is considerably more expensive and often less stable than VG, we will adopt the VG model for our purposes. The necessary ground state geometries and frequencies can be obtained at the density functional theory (DFT) level, while the vertical excitation energy and the dipole integrals can be calculated at the DLPNO-STEOM level. The 0–0 transition energy can then be simply obtained by correcting the DLPNO-STEOM energy with the VG energy shift obtained from DFT frequencies and displacements, and the entire spectrum can also be calculated and compared to its experimental counterpart.

3 Comparison of Projection and Many-Body Solution Criteria

In the following, we will compare projection and many-body techniques that yield solutions for various correlated wavefunction *ansatzes*. We will also discuss the relevance of normal ordering in projection techniques, while normal ordering is assumed throughout for the many-body approach. We begin our discussion with the simple case of single-reference methods before moving on to the more relevant multireference approaches.

3.1 Single-Reference Methods

The basic theory which is necessary to discuss the projective and many-body residual conditions has been recapitulated in Sec. 2.4.1 (single-reference CC theory), Sec. 2.2.5 (normal order and Wick's theorem) and Sec. 2.2.6 (generalized normal order and Wick's theorem). Together, these concepts can be used to expand the similarity-transformed Hamiltonian $\bar{H}$ from CC theory in a many-body form, given below. The discussion here will be given at a general level, but in this section we will focus on the consequences for SR methods. However, before doing so, we first need to slightly recast Wick's theorem to make the connection of the many-body form to normal order more obvious.

Starting from Eq. (2.58) in the SR case, or more generally from Eq. (2.65), we consider now a string $\hat{C}^{prt\ldots}_{qsu\ldots}$ which is an arbitrary product of creators $\hat{a}^p, \hat{a}^r, \hat{a}^t, \ldots$ and annihilators $\hat{a}_q, \hat{a}_s, \hat{a}_u, \ldots$ The generalized Wick's theorem (we will drop the GNO index of curly braces from now on) makes it possible to rewrite such a string in terms of normal ordered strings and constants

$$\begin{aligned}\hat{C}^{prt\ldots}_{qsu\ldots} = \{\hat{a}^p\hat{a}_q\hat{a}^r\hat{a}_s\hat{a}^t\hat{a}_u\ldots\} + P(^p_q|^{rt}_{su}\ldots)C^p_q\{\hat{a}^r\hat{a}_s\hat{a}^t\hat{a}_u\ldots\} \\ + P(^{pr}_{qs}|^t_u\ldots)C^{pr}_{qs}\{\hat{a}^t\hat{a}_u\ldots\} + \ldots + C^{prt\ldots}_{qsu\ldots}. \quad (3.1)\end{aligned}$$

Here, the n-body coefficients $C^{p\ldots}_{q\ldots}$ contain all possible products of cumulants with a particle rank of at most n such that the total particle rank of the product is n,

$$C^{p\ldots r\ldots t\ldots}_{q\ldots s\ldots u\ldots} = \lambda^{p\ldots}_{q\ldots}\cdots\lambda^{r\ldots}_{s\ldots}\cdots\lambda^{t\ldots}_{u\ldots}\cdots \quad (3.2)$$

This further simplifies in the single-reference case as then only the one-particle cumulants ($\lambda^p_q = \gamma^p_q = \delta^p_q - \eta^p_q$) survive [219], and those can in fact be reduced to Kronecker deltas

(Sec. 2.2.5). The fully contracted term $C^{prt\cdots}_{qsu\cdots} = \langle 0|\hat{C}^{prt\cdots}_{qsu\cdots}|0\rangle$ is often denoted simply as C_0. The permutation operators $P(\ldots|\ldots)$ create all possible swaps between the index groups separated by | with the appropriate sign, essentially accounting for all possible contraction patterns. This formulation can be extended to the generators of the unitary group to enable a spin-free formulation, while products of normal ordered operators, as in in Eq. (2.59), are also conveniently evaluated, since only contractions connecting different terms need to be considered.

With Eq. (3.1), we can now define a normal-ordered (many-body expanded) form of the similarity-transformed Hamiltonian. From Eq. (2.90),

$$\bar{H} = \sum_{k=0}^{N} \left(\hat{H}\hat{T}^k\right)_C = \sum_{k=0}^{N} \sum_{\substack{pqr\\stu}} h^{q\cdots}_{p\cdots} \underbrace{\overline{t}^{s\cdots}_{r\cdots} \cdots \overline{t}^{u\cdots}_{t\cdots}}_{k \text{ times}} \{\hat{a}^p \hat{a}_q \ldots\}\{\hat{a}^r \hat{a}_s \ldots\}\cdots\{\hat{a}^t \hat{a}_u \ldots\}, \tag{3.3}$$

where we reverted to the spin-orbital form for simplicity and used general labels for the amplitudes $\overline{t}$ as in Eq. (2.122). In the SR case, the special form in Eq. (2.82) applies, i.e., the amplitudes do not contain cumulants ($\overline{t} = t$). Thus, for the purposes of this chapter, we may assume that the cluster operator is normal-ordered and set $\overline{t} = t$ from now on. The normal-ordered Hamiltonian contains no cumulants as a result of normal ordering in the SR case, see Eq. (2.96), while in the MR case we consider these as part of h_0, h_1 and h_2. With these conventions, using Eq. (2.59), substituting Eq. (3.1) and rearranging the terms leads to the following structure,

$$\bar{H} = \left(\sum_{\cdots} h^{\cdots}_{\cdots} t^{\cdots}_{\cdots} \cdots t^{\cdots}_{\cdots} P(\cdots|\cdots) C^{\cdots}_{\cdots}\right)_0 + \left(\sum_{\cdots} h^{\cdots}_{\cdots} t^{\cdots}_{\cdots} \cdots t^{\cdots}_{\cdots} P(\cdots|\cdots) C^{\cdots}_{\cdots}\right)^q_p \{\hat{a}^p \hat{a}_q\} + \left(\sum_{\cdots} h^{\cdots}_{\cdots} t^{\cdots}_{\cdots} \cdots t^{\cdots}_{\cdots} P(\cdots|\cdots) C^{\cdots}_{\cdots}\right)^{qs}_{pr} \{\hat{a}^p \hat{a}_q \hat{a}^r \hat{a}_s\} + \ldots. \tag{3.4}$$

Readopting the spin-free notation, the transformed matrix elements $\overline{h}$ can now be easily identified as a product of untransformed integrals h, amplitudes t and cumulants λ (collected in C). Thus, we may write

$$\bar{H} = \overline{h}_0 + \overline{h}^q_p \{\hat{E}^p_q\} + \overline{h}^{qs}_{pr} \{\hat{E}^{pr}_{qs}\} + \ldots, \tag{3.5}$$

$$\overline{h}_0 = \langle 0|\bar{H}|0\rangle = E. \tag{3.6}$$

Note that this expression does not terminate at two-body terms. Using the product rule in Eq. (2.59), the projection conditions in Eq. (2.92) and (2.93) take the following simple form in the SR case,

$$r^a_i \equiv \overline{h}^i_a = 0, \quad r^{ab}_{ij} \equiv \overline{h}^{ij}_{ab} = 0. \tag{3.7}$$

Since we now set conditions on the n-body matrix elements of $\bar{H}$, $\overline{h}_n$, directly, we will refer to this approach as the many-body solution technique, which was first discussed by Nooijen and Bartlett [221]. In the SR case, this does not lead to simplifications, in fact, we have established that in the single determinant case, the projection technique and the many-body approach lead to the same results [210].

3.2 Multireference Methods

In the multireference regime, the well-known *ansatzes* from the single-reference limit can also be applied, albeit at higher complexity since the working equations now also need to encode the multideterminantal nature of the wave function. We will not discuss multireference CI approaches here, for which the residual conditions are formally very similar to the single reference case and for which an excellent review has been published [182]. Instead, we will solely focus on the internally contracted CC approach, and the flexibility in choosing the residual conditions for this method.

3.2.1 Similarity-Transformed Hamiltonian Expansions

The similarity-transformed Hamiltonian $\bar{H}$ in coupled-cluster theory is given in Eq. (2.89). For the multireference case presented here, we do not make any assumptions about the cluster operator $\hat{T}$ contained therein. In particular, the cluster operator may contain all-active excitations, which are the only ones that may lead to the maximum expansion length in the cumulant expansion below, since cumulants with virtual or inactive labels vanish [219].

As any operator, the similarity-transformed Hamiltonian $\bar{H}$ can be expanded into a sum of normal-ordered operators of increasing ranks [221],

$$\bar{H} = \sum_{j=0}^{N(M-1)+2} \bar{h}_j \{\hat{E}_j\}, \tag{3.8}$$

where N is the order of the commutator expansion (Eq. 2.90), M is the maximum excitation level of the coupled-cluster theory, e.g., $M = 2$ for CCD or CCSD, and 1 is subtracted to account only for the connected part of the Hamiltonian. This *many-body expansion* is formally identical to the that in the single-reference case, Eq. (3.5), except for the fact that here we use normal-order with respect to a multideterminantal reference using the formalism of Kutzelnigg and Mukherjee [219].

The expansion elements $\bar{h}_j$ in Eq. (3.8) deserve some attention. The similarity-transformed Hamiltonian $\bar{H}$ is constructed from the Fock matrix elements and two-electron integrals from the Born-Oppenheimer Hamiltonian (denoted as h_1 and h_2 in the following, respectively) as well as the n-tuple amplitudes t_n, along the lines indicated in Eq. (3.4). Using these quantities, the *cumulant expansion of the $\bar{H}$ matrix elements* can be expressed in a more compact form

$$\bar{h}_j = \sum_{k=0,1,2} \sum_{l=N_0}^{N} (h_k \prod_{i=0}^{l} t_{n_i} \Lambda_{k+L-j})_C . \tag{3.9}$$

Here, $t_{n_0} = 1$, $L = \sum n_i$, $1 \leq n_i \leq M$ and N_0 is the minimal order of the nested commutator contributing to h_j obtained as $N_0 = \max(0, \lceil \frac{j-2}{M} \rceil)$. For a given order of commutator expansion l, there are many permutations of operator strings with same total order L. For simplicity,

summation over these is implicitly assumed but not explicitly denoted. The quantity Λ connects the integrals and amplitudes in all possible label permutations, except for the disconnected terms. Λ_m itself consists of the sum of all allowed products of cumulants of the form Eq. (3.2), the collective rank of which is m, and follows directly from the application of Wick's theorem (Eq. 3.1) and the KM normal order contraction rules [219] on the operators in the BCH expansion (Eq. 2.90).

Although the idea behind the cumulant expansion is the expectation that higher rank cumulants become smaller, this assumption does not hold true in all cases [405]. However, it is possible to argue independently, based on Eq. (3.9), that for a given M, the many-body terms $\overline{h}_j$ become progressively smaller for higher orders j [213]. This is evident, since the minimal contraction length N_0 increases with j. For $M = 2$, N_0 is 1 for $j = 3$ and it is at least 2 for larger j. Thus, there is at least one t amplitude in every contribution to $\overline{h}_j$, $j > 2$. If the magnitude of the amplitudes can be assumed to be $||t|| \ll 1$, then that the magnitude of $\overline{h}_j$ will decrease with increasing j. It should be noted that the *commutator expansion* in the BCH formula Eq. 2.90 is a third way of controlling the complexity of the Hamiltonian, since it affects the matrix elements $\overline{h}_j$. All these, the many-body expansion of the Hamiltonian, the cumulant expansion of matrix elements, and the cumulant expansion of the Hamiltonian can serve as a basis for truncation schemes, and all of which are associated with the transformed Hamiltonian. A fourth possibility related to the residual conditions is discussed in the next section.

3.2.2 Projective and Many-Body Residual Conditions

The usual residual conditions employed in CC theory are the projective residuals,

$$r_n = \langle 0|\{\hat{E}_n\}^\dagger \bar{H}|0\rangle, \tag{3.10}$$

which result in $||n||$ non-linear equations for $||n||$ cluster parameters. In contrast to CI theory, this does not constitute a variational condition, although one can view this as the zeroth-order term of such a condition [441]. Note that residual Eqs. (3.10) are generally linearly dependent in the multireference regime. Consequently, a scheme to remove the dependencies is needed, which may be achieved through defining the metric

$$S_{ij} = \langle 0|\hat{E}_i^\dagger \hat{E}_j|0\rangle \tag{3.11}$$

for $i, j \leq n$ and using Löwdin's canonic orthogonalization procedure to obtain an orthonormal basis in which the CC ansatz can be written. Since here we discuss concepts, we will not consider this issue any further, and refer to Refs. [176, 195, 385] for details.

The projective residual equations can be cast using the many-body expansion $\bar{H}$ (Eq. 3.8) as

$$r_n = \sum_{j=1}^{nN+2} \langle 0|\{\hat{E}_n\}^\dagger \overline{h}_j \{\hat{E}_j\}|0\rangle = \sum_{j=1}^{nN+2} \overline{h}_j \Lambda_{n+j}, \tag{3.12}$$

where Λ_{n+j} again contains products of different orders of cumulants (*vide supra*) that lead to the total rank $n + j$. We may refer to this expression as the *cumulant expansion of the projection conditions*.

Yet, in contrast to Λ_m in the previous section, we can here impose further restrictions on the possible products of cumulants in Λ_m since we only have two excitation operators and only obtain the fully contracted part in Eq. (3.12). In particular, we remark that a term containing only powers of the one-body cumulant, $(\lambda_1)^m$, can only appear when the operators in Eq. (3.12) are of the same order. For example,

$$\langle 0|\{\hat{E}_2\}\{\hat{E}_2\}|0\rangle = \lambda_4 + \lambda_3\lambda_1 + \lambda_2(\lambda_2 + \lambda_1^2) + \lambda_1^4, \tag{3.13}$$

whereas for

$$\langle 0|\{\hat{E}_2\}\{\hat{E}_1\}|0\rangle = \lambda_3 + \lambda_2\lambda_1. \tag{3.14}$$

It follows that only a single term in Eq. (3.12) can have products of purely one-body cumulants, namely when $n = j$ we get $(\lambda_1)^n$.

As $\lambda_1 = \gamma_1$ and η_1 are the only survivors in the single determinant limit, they are also expected to be among the largest quantities in the multireference case as well. Moreover, higher-order terms $\bar{h}_{j>2}$ in the many-body expansion are also small since they contain at least one amplitude t. Thus, separating the leading order term in Λ_{n+j} in Eq. (3.12), we find

$$r_n = \bar{h}_n(\lambda_1)^n + R_n, \tag{3.15}$$

with the higher order terms collected in R_n.

Eq. (3.15) can be simplified by using the fact that a CASSCF wave function is invariant with respect to arbitrary rotations within the inactive, active, and virtual orbital blocks. We can use this flexibility to assume active natural orbitals, i.e., diagonal active blocks of the densities, $\lambda_q^p = n_i \delta_q^p \delta_p^i$ and $\eta_q^p = n_{a'} \delta_q^p \delta_p^{a'}$. In this case, Eq. (3.15) reduces to

$$r_n = \prod_{i=1}^{n} n_i \bar{h}_n + R_n \tag{3.16}$$

$$= n\bar{h}_n + R_n. \tag{3.17}$$

The orbital occupations $0 \leq n_i \leq 2$ are defined such that for virtual orbitals we have $n_a = 2$, a result of virtual labels being associated with a hole density η (see, e.g., Eq. 3.28). Effectively, the term $(\lambda_1)^n$ reduces to a simple prefactor of the n-body elements $\bar{h}_n$. Neglecting all higher-order terms R_n, we find that solving for the residual equations r_n is equivalent to setting the many-body elements $\bar{h}_n$ themselves directly to zero,

$$r_n = \bar{h}_n \stackrel{!}{=} 0, \tag{3.18}$$

for non-vanishing occupation numbers (see discussion below). The residual condition in Eq. (3.18) has been termed *many-body residuals* by Nooijen and co-workers [210]. In their approach, instead of taking this as the leading-order term of the projected residual conditions, they have arrived at these conditions through similar conditions as used in Fock-space CC theory (FSCC) [210, 221, 462].

As discussed elsewhere, these residual conditions are formally free from linear dependencies that plague the projective ic-MRCC residual conditions [210]. It should be remarked that these conditions must not be used for the singles residuals,

$$r_i^a = \overline{h}_a^i = 0, \tag{3.19}$$

$$r_t^a = \overline{h}_a^t = 0, \tag{3.20}$$

$$r_i^t = \overline{h}_t^i = 0, \tag{3.21}$$

since these conditions would violate the generalized Brillouin theorem [210]. The higher-order residuals do not have this restriction. Therefore, the two-body elements $\overline{h}_2$ are well suited to replace the projective doubles residual conditions from Eq. (3.10). These many-body conditions are much cheaper than the projective residual equations, which is evident from Eq. (3.15), and an analysis published by Datta *et al.* [210]. To summarize, we have established that satisfying the many-body conditions in the multireference case correspond to putting the leading terms in the projection conditions to zero and leaving only residual terms behind, which may be expected to be small. We will investigate the validity of this expectation in Sec. 3.3.

3.2.3 A Simplified Ansatz for Comparing Many-Body and Projection Conditions

In this section, we construct two simplified internally contracted MRCC theories that only differ in the residual condition; one uses the conventional projective residuals, henceforth denoted as *pr*-MRCC, and the other uses the many-body residuals, called *mb*-MRCC in what follows. To the best of our knowledge, these two conditions have only been compared in the work of Datta, Kong, and Nooijen [210] as *mixed* residuals and were always followed by an uncontracted MRCI-type diagonalization.

In multireference coupled-cluster theories, it is customary to include all but the active-active excitations into the cluster operator $\hat{T}$ [194, 195], simply for retaining intermediate normalization and avoiding even higher-order densities,

$$\langle 0|\Psi\rangle = 1. \tag{3.22}$$

If necessary, the effect of active-active excitations can be included through relaxing the reference CI expansion coefficients in the presence of $\hat{T}$ by solving an effective eigenvalue problem [194, 195]. Different schemes have already been discussed in the literature [195].

In this chapter, we restrict the cluster operator to doubles excitations into the virtual space

(i.e., two-particle excitations),

$$\hat{T} = \frac{1}{2}\sum_{ijab} t_{ab}^{ij}\hat{E}_i^a\hat{E}_j^b + \sum_{itab} t_{ab}^{it}\hat{E}_i^a\hat{E}_t^b + \frac{1}{2}\sum_{tuab} t_{ab}^{tu}\hat{E}_t^a\hat{E}_u^b. \tag{3.23}$$

This implies that all excitation operators commute among themselves, in other words, we have a normal-ordered $\hat{T}$. We remark that this cluster operator is the same as used in MR-EOMCC theory [211], except for the singles and the fact that we use state-specific instead of state-averaged densities throughout.

The ramifications of such a restricted cluster operator are as follows. First, due to having only commuting excitation operators in $\hat{T}$, the BCH expansion terminates exactly after fourfold nested commutators. However, in the residual and energy expressions of both the projective and the many-body variant of fic-MRCC, the truncation already occurs after the twofold commutators,

$$\bar{H} = \hat{H} + [\hat{H},\hat{T}] + \frac{1}{2}[[\hat{H},\hat{T}],\hat{T}]. \tag{3.24}$$

This follows from the excitation ranks of the Hamiltonian, the projection and the cluster operators in the case of the projective variant. For the many-body variant, due to the constrained $\hat{T}$, terms with at least $\hat{H}\hat{T}^3$ can only appear in many-body elements of rank four $\bar{h}_4$ or higher. These terms neither appear in the energy expression ($= \bar{h}_0$) nor in the residuals ($= \bar{h}_2$). Secondly, the rank of the cumulants or densities is also limited. For the projective equations, at most terms of four-body character λ_4 can appear. In the many-body case, all cumulants are strictly limited to two-body terms λ_2 since we only have two-particle excitation classes in Eq. (3.23).

For the chosen cluster operator (Eq. 3.23), the many-body expanded Hamiltonian in Eq. (3.8) takes the form

$$\begin{aligned}\bar{H} = \bar{h}_0 &+ \bar{h}_p^q\{\hat{E}_q^p\} + \bar{h}_{pq}^{rs}\{\hat{E}_{rs}^{pq}\} + \bar{h}_{pqa}^{rsi}\{\hat{E}_{rsi}^{pqa}\} \\ &+ \bar{h}_{pqab}^{rsij}\{\hat{E}_{rsij}^{pqab}\} + \bar{h}_{pabcd}^{qijkl}\{\hat{E}_{qijkl}^{pabcd}\} + \bar{h}_{abcdef}^{ijklmn}\{\hat{E}_{ijklmn}^{abcdef}\},\end{aligned} \tag{3.25}$$

where i, j, k, l, m, n are taken to be either inactive or active indices for the sake of brevity. Inactive/Active and virtual indices appear in Eq. (3.25) instead of all-general indices for the simple reason that the cluster operator (Eq. 3.23) only contains two-particle excitation classes. The restricted $\hat{T}$ also limits the total order of cumulants in the cumulant expansion of the $\bar{H}$ matrix elements (Eq. 3.9) to Λ_2 since cumulants with non-active indices must vanish [219].

For some indices i', j', a', b' appearing in the excitation classes in Eq. (3.23), and applying a contravariant projection to eliminate permutations, the relation between projection and the

many-body conditions thus becomes

$$\langle 0|(\bar{H} - E)|0\rangle = 0, \tag{3.26}$$

$$\begin{aligned} r^{a'b'}_{i'j'} &= \langle 0|\{\tilde{E}^{a'b'}_{i'j'}\}^\dagger \bar{H}|0\rangle \\ &= \left(\gamma^{i'}_r \gamma^{j'}_s \eta^p_{a'} \eta^q_{b'} + \lambda^{i'j'}_{rs} \eta^p_{a'} \eta^q_{b'}\right) \bar{h}^{rs}_{pq} \\ &\quad + \left(\eta^a_{b'} \eta^q_{a'} \gamma^{i'}_r \lambda^{j'p}_{si} + \eta^a_{b'} \eta^q_{a'} \lambda^{i'j'p}_{rsi}\right) \bar{h}^{rsi}_{pqa} \\ &\quad + \left(\eta^a_{a'} \eta^b_{b'} \gamma^{i'}_r \lambda^{j'pq}_{sij} + \eta^a_{a'} \eta^b_{b'} \lambda^{i'q}_{rs} \lambda^{j'p}_{ij} + \eta^a_{a'} \eta^b_{b'} \lambda^{i'j'pq}_{rsij}\right) \bar{h}^{rsij}_{pqab} \stackrel{!}{=} 0, \end{aligned} \tag{3.27}$$

which can be written in a more compact fashion (cf. Eq. 3.15) as

$$r^{a'b'}_{i'j'} = \gamma^{i'}_r \gamma^{j'}_s \eta^p_{a'} \eta^q_{b'} \bar{h}^{rs}_{pq} + R^{a'b'}_{i'j'} = 0, \tag{3.28}$$

in which $R^{a'b'}_{i'j'}$ denotes the higher order terms. As we noted earlier, the projection and many-body conditions are equivalent in the single determinant case (all higher-order cumulants vanish and the one-body particle and hole densities reduce to Kronecker deltas), whereas they are obviously different in the current context. By following the analysis from Sec. 3.2.2 in using natural orbitals, the projection conditions under the present assumptions become

$$r^{a'b'}_{i'j'} = n_{i'} n_{j'} n_{a'} n_{b'} \bar{h}^{i'j'}_{a'b'} + R^{a'b'}_{i'j'} = 0, \tag{3.29}$$

which lends itself naturally to the approximation

$$r^{a'b'}_{i'j'} \approx n_{i'} n_{j'} n_{a'} n_{b'} \bar{h}^{i'j'}_{a'b'} = 0, \iff \bar{h}^{i'j'}_{a'b'} = 0, \quad \text{if } n_{i'}, n_{j'} \neq 0. \tag{3.30}$$

Thus, if higher order terms can be neglected, and the occupation numbers are non-zero, the projection and many-body conditions are equivalent under the present assumptions. The active orbital occupation numbers may become zero as the system approaches the single reference limit (e.g., for antibonding orbitals included in the active space). In this case, the two conditions are not equivalent, and it is precisely the case where *mb*-MRCC has convergence problems. In *pr*-MRCC, $h^{i'j'}_{a'b'}$ is multiplied by the occupation number (the density in general), which becomes singular and is thus removed as a linear dependency. In the more general case, Eq. (3.28) still applies, and the fulfillment of the many-body conditions also implies that some of the largest components of the first term in this equation vanish, while in the terms that do survive the matrix elements are multiplied with off-diagonal elements of the one body densities that vanish in the single determinant limit.

3.3 Comparing Projection and Many-Body Conditions

To assess how close the projective and the many-body solutions are numerically, we evaluated the residual conditions of each method using the converged amplitudes of the other, according to Eqs. (3.10) and (3.18), respectively.

We first turn to the single reference limit (Tab. 3.1), i.e., we choose a fully occupied active space, CAS($2n, n$). In this case, as detailed above, both residual conditions reduce to the same equations since the many-body and projective residuals are *identical* in the single-reference case, and consequently result in the same set of amplitudes for both *pr*-MRCC and *mb*-MRCC methods. Accordingly, the computed correlation energy from all methods and the cross-substituted variants agrees to within machine precision, which numerically demonstrates the equivalence of *pr*-MRCC and *mb*-MRCC theory in the single-reference case. Moreover, all residuals lie below the chosen convergence threshold of $\max(||\boldsymbol{r}||) < 10^{-6}$, indicating that also the cross-substituted variants are converged.

Table 3.1 Single reference limit of projective and many-body methods. *pr*-MRCC(T_{mb}): projective code evaluated with many-body amplitudes, *mb*-MRCC(T_{pr}): many-body code evaluated with projective amplitudes. The N_2 molecule had an inter-atomic distance of $r_{\mathrm{NN}} = 1.0$ Å.

		***pr*-MRCC**	***mb*-MRCC**	***pr*-MRCC($\boldsymbol{T}_{\mathbf{mb}}$)**	***mb*-MRCC($\boldsymbol{T}_{\mathbf{pr}}$)**
N_2 CAS(2,1)	$\max(\|\|\boldsymbol{r}\|\|)$	2.10E-07	9.30E-08	1.60E-07	9.00E-08
	E_{cor}	−0.308358	−0.308358	−0.308358	−0.308358
N_2 CAS(4,2)	$\max(\|\|\boldsymbol{r}\|\|)$	9.80E-07	9.20E-08	1.80E-07	7.40E-07
	E_{cor}	−0.308358	−0.308358	−0.308358	−0.308358

In the multireference limit (Tab. 3.2), the situation is more complicated as the residual equations are no longer the same (Eq. 3.27). We base the comparison on very tightly converged MRCC calculations ($\max(||\boldsymbol{r}||) < 10^{-10}$) and then cross-substitute the amplitudes.

Table 3.2 Multireference limit of projective and many-body methods. *pr*-MRCC(T_{mb}): projective code evaluated with many-body amplitudes, *mb*-MRCC(T_{pr}): many-body code evaluated with projective amplitudes. Also scaled with orbital occupation numbers to assess magnitude of higher order terms, *mb*-MRCC(T_{pr})* (see text). Results were computed at fragment or atomic separations of 3–5 Å.

		***pr*-MRCC**	***mb*-MRCC**	***pr*-MRCC($\boldsymbol{T}_{\mathbf{mb}}$)**	***mb*-MRCC($\boldsymbol{T}_{\mathbf{pr}}$)**	***mb*-MRCC($\boldsymbol{T}_{\mathbf{pr}}$)***
Li_2 CAS(2,2)	$\max(\|\|\boldsymbol{r}\|\|)$	5.14E-11	9.21E-11	2.03E-03	2.99E-02	1.57E-01
	E_{cor}	−0.071298	−0.070872	−0.070872	−0.071298	−0.071298
C_2H_2 CAS(6,6)	$\max(\|\|\boldsymbol{r}\|\|)$	4.44E-11	4.17E-11	5.11E-03	2.77E-02	1.05E-01
	E_{cor}	−0.141370	−0.137662	−0.137662	−0.141370	−0.141370
C_2H_4 CAS(4,4)	$\max(\|\|\boldsymbol{r}\|\|)$	7.12E-11	7.96E-11	4.33E-03	2.52E-02	1.08E-01
	E_{cor}	−0.230057	−0.227539	−0.227539	−0.230057	−0.230057
C_2H_6 CAS(2,2)	$\max(\|\|\boldsymbol{r}\|\|)$	7.14E-11	9.83E-11	2.60E-03	5.27E-02	2.59E-01
	E_{cor}	−0.323238	−0.322022	−0.322022	−0.323238	−0.323238

It is obvious that the correlation energy is the same if the same set of amplitudes is used, as $\bar{H}$ is the same in the case of state-specific densities used here, i.e., *pr*-MRCC and the many-body code with *pr*-MRCC amplitudes (*mb*-MRCC(T_{pr})), and *mb*-MRCC and *pr*-MRCC with many-body amplitudes (*pr*-MRCC(T_{mb})), both give the same value of the correlation energy E_{cor}. It is also worth noting that the correlation energy E_{cor} is consistently lower in the *pr*-MRCC method than in the *mb*-MRCC method.

Let us turn to the magnitude of the residuals. The projective and many-body calculations that preceded the cross-substitution were converged to almost machine precision so that the magnitude of the cross-substituted residuals can be judged correctly. The *pr*-MRCC(T_{mb}) are on the order of $2\text{–}5 \times 10^{-3}$. This signals that the cross-substituted *pr*-MRCC is not converged, however, the T_{mb} solution is already better than the perturbative fic-NEVPT2 guess for the amplitudes used in the *pr*-MRCC calculations.

The many-body residuals, when evaluated with the *pr*-MRCC amplitudes (*mb*-MRCC(T_{pr})), are further away from zero than the other way around. In this case, the residuals are on the order of $2\text{–}5 \times 10^{-2}$, roughly one magnitude more than for *pr*-MRCC(T_{mb}). This seems to indicate that *the converged many-body amplitudes are a better guess of the solution of the projection method than the other way around.* This can be understood by considering that *pr*-MRCC never zeroes out elements of the transformed Hamiltonian $\bar{H}$ itself, but rather only zeroes the sum from Eq. (3.27). Seeing as there are no constraints on the individual terms, the many-body elements may in principle be arbitrarily large. In contrast, the *pr*-MRCC(T_{mb}) residuals are smaller for the fact that the many-body condition explicitly zeroes out the largest elements of the many-body expanded Hamiltonian, and thus also makes the leading-order terms in the projection condition from Eq. (3.27) vanish. Nonetheless, higher-order elements still contribute to the projective residual conditions.

Let us put these findings into the context of higher-order terms according to Eq. (3.27). Substituting the many-body amplitudes into the projective (*pr*-MRCC(T_{mb})) conditions in Eq. (3.29) yields the higher order terms

$$r_{i'j'}^{a'b'}(T_{\mathrm{mb}}) = R_{i'j'}^{a'b'}(T_{\mathrm{mb}}), \tag{3.31}$$

since $\overline{h}_{a'b'}^{i'j'}(T_{\mathrm{mb}}) = 0$. On the other hand, the *pr*-MRCC amplitudes satisfy Eq. (3.29), i.e.,

$$-n_{i'}n_{j'}n_{a'}n_{b'}\overline{h}_{a'b'}^{i'j'}(T_{\mathrm{pr}}) = R_{i'j'}^{a'b'}(T_{\mathrm{pr}}). \tag{3.32}$$

The left hand side contains the many-body conditions evaluated at the *pr*-MRCC solution ($\overline{h}_{a'b'}^{i'j'}(T_{\mathrm{pr}})$, *mb*-MRCC($T_{\mathrm{pr}}$)) multiplied by the orbital occupation numbers. Thus, the cross substituted residual conditions contain information about the higher order terms $R_{i'j'}^{a'b'}$. To assess their relative magnitude for both methods, the *mb*-MRCC(T_{pr}) residuals need to be multiplied by the orbital occupation numbers as shown in Eq. (3.32). These scaled residuals are denoted as *mb*-MRCC(T_{pr})* in Tab. 3.2 and correspond exactly to the higher-order terms.

We find that the scaled *mb*-MRCC(T_{pr})* residuals are on the order of 0.1, which is larger by two magnitudes than the higher order terms evaluated with many-body amplitudes (*pr*-

MRCC(T_{mb})). This observation underpins our earlier finding that the many-body solution is a better guess of the projection solution than the other way around. In fact, the many-body solution not only removes the largest contributions from the projection condition but also makes the remaining terms relatively small. Since the cumulants in these higher order terms are not in general small, these must be small because the higher order many-body elements that contain at least one t amplitude are small. Similar conclusions cannot be reached about the projection solution when used as a guess in the many-body approach. In the final analysis, the many-body conditions are not only cheaper than the projection solution, but they also provide a good approximation to it.

This behavior of the many-body solution, in turn, perfectly justifies the truncated sequential transformation procedure in MR-EOMCC theory. The goal of MR-EOMCC, see Sec. 2.6.4, is to decouple blocks of the Hamiltonian, which is equivalent to satisfying projection conditions but MREOM uses many-body conditions instead. Once the many-body solutions for a given excitation class are obtained, this solution is then used to similarity transform the Hamiltonian. This procedure is repeated sequentially for several excitation classes and in each step the previous similarity-transformed Hamiltonian is truncated after the two-body terms since the higher order terms obtained from the many-body solutions are small. For more details, see Sec. 6.1.1.

4 Implementation with ORCA-AGE

In this chapter, we describe the updated code generation toolchain, ORCA-AGE II, of which I wrote the vast majority of the code. We begin with an overview of the toolchain, which emphasizes the main improvements over the previous toolchain, ORCA-AGE [1], before comparing it to other code generators. Then, we show how we implemented the fic-MRCC method using ORCA-AGE II, and demonstrate how even complicated contractions can be efficiently implemented.

4.1 Software Description

The ORCA-AGE II toolchain has been written from scratch, a decision that has been motivated by the experience gained while working with its first version [1]. In this section, we first recapitulate the layout of ORCA-AGE II, which has mainly undergone streamlining, before we discuss the new internal algorithms, equation and code generators. Further information on the code generation process can be found in the preceding publication [1].

From an architectural standpoint, ORCA-AGE II still uses the modular structure of its predecessor, organized into three groups: equation generation, equation processing (factorization), and code generation. Each of these three steps is further split into smaller executables that only perform a very specific subtask, e.g., merging several contractions into one by exploiting tensor symmetry. This layout allows highly flexible workflows, where each step can be turned on or off depending on the user's preference, although all steps are enabled by default to achieve the best possible performance of the generated code. In turn, each of these steps may have optional features like debug output or different contraction engines, which are controlled through easy-to-understand command line options as well. Furthermore, the user can easily verify or modify all intermediate stages, which are simply text files containing equations divided into a single contraction per line, as necessary. An example of such an equation file is given in Listing 4.1. No further configuration files are required to avoid additional dependencies on additional libraries or codes, in contrast to the first version of the toolchain.

While the general structure of the toolchain has remained largely unchanged, the code itself is an entirely new development. The main reason for redeveloping the software lies in the long code generation times of the previous version, which required about one month to generate highly complicated theories such as fic-MRCC, leading to a slow overall development and debugging process. This was addressed by improving the internal algorithms throughout all factorization

Listing 4.1 Example for an equation file used in the ORCA-AGE II toolchain. This particular equation file contains an excerpt of the full fic-MRCC equations.

```
#S pairIJAB 2 0 pairIJ[AUTOCI_IJAB].GetRows() i0=pairIJ[AUTOCI_IJAB]({0},0) i1=pairIJ[AUTOCI_IJAB]({0},1) pairIJ[AUTOCI_IJAB] 400
#S pairIJTU 2 0 pairIJ[AUTOCI_IJTU].GetRows() i0=pairIJ[AUTOCI_IJTU]({0},0) j0=pairIJ[AUTOCI_IJTU]({0},1) pairIJ[AUTOCI_IJTU] 200
#T TMatrixContainer Ctuva 4 Ctuva[{1},{3}]({0}-FirstVirtual,{2}-FirstActive)
#T TMatrixContainer Cituv 4 Cituv[{1},{3}]({0}-FirstActive,{2}-FirstActive)
#T TMatrixContainer Citau 4 Citau[{1},{3}]({0}-FirstVirtual,{2}-FirstActive)
#T [ ... ]
Sijab(a0,i0,b0,j0) += 1.0 I(b0,j0,a0,i0)
Sijab(a0,i0,b0,j0) += 0.5 DC1(P1,P2) Cijtu(P0,j0,P1,I1) I(I1,P2,I2,P0) Cijab(b0,I2,a0,i0)
Sijta(t0,i0,a0,j0) += -1/3 Citab(A1,i0,A2,P0) I(I1,A1,a0,A2) DC2(P1,P0,P2,t0) Cijtu(P1,j0,P2,I1)
Sijtu(t0,i0,u0,j0) += -0.5 Citua(P0,i0,A0,P1) I(I1,A0,P3,P0) DC3(P2,u0,P3,t0,P4,P1) Cijtu(P2,j0,P4,I1)
Situv(u0,i0,t0,v0) += 0.5 I(I1,P0,P4,P5) Cituv(P3,I1,P6,P7) DC5(v0,t0,P3,u0,P4,P5,P6,P7,P1,P2) Cituv(P0,i0,P1,P2)
[ ... ]
```

and generation steps. For example, we introduced a hash-based data compaction method for eliminating duplicate intermediates generated by the initial, term-wise factorization step. This reduces the scaling from a quadratic to an expected linear-time algorithm, and consequently to a speedup of about two orders of magnitude for the fic-MRCC method in this step alone. While this example shows the biggest improvement over the old toolchain, other parts, such as the detection of tensor symmetry, have received improvements of up to one order of magnitude as well. To further ensure good use of multicore architectures, we use OpenMP to allow for thread-based parallelism in the toolchain. Lastly, the new toolchain has been written in the C++ programming language, compared to the previous implementation in the Python programming language. Thus, we further obtain small performance improvements, with the major benefit that ORCA-AGE II fits better into the ORCA ecosystem through a more homogeneous and cleaner software structure.

Internally, the executables rely on a core library that provides the required basic functionality to read/write, represent, and manipulate tensor contractions as a set of high-level classes and functions. This library, and by extension the entire toolchain, has no limitations in terms of number of indices per tensor, allowed tensor symmetries, or number of source tensors per contraction, to name a few features. Consequently, any theory that can be written as a set of tensor contractions can be implemented with ORCA-AGE II. Also, despite this generality, even costly operations such as determining the tensor symmetry of an intermediate based on its source tensors are so fast that they can be routinely used for hundreds of tensors with, for example, ten indices, as is the case in fic-MRCC theory (see Sec. 4.3).

In addition to the improvements to the underlying infrastructure, the functionality (especially of the code generator) has been extended as well. In terms of performance improvements, the recent work falls into the two categories of compute and I/O optimizations. We first discuss the compute optimizations, which are mainly geared towards making BLAS calls more pervasive in the generated code. To this end, we use the fact that essentially every binary tensor contraction

can be written as a matrix-matrix multiplication,

$$C_{ij} = \sum_k A_{ik} B_{kj}, \tag{4.1}$$

in which the *compound* indices i, j, k may refer to none, one, or multiple actual indices. As a result, outer products and matrix-vector and vector-vector multiplications can be viewed as special cases thereof. Such a scheme has already been discussed in the literature, and is generally referred to as *Transpose-Transpose-DGEMM-Transpose* (TTGT, see App. A.3 for details), since (in the worst case) we might need to reorder (transpose) the axes (indices) on both source tensors and the target tensor [261, 262]. To the best of our knowledge, none of the previous implementations are as general as the current implementation in ORCA-AGE II, as our implementation also allows for further edge cases (trace operations, repeated indices, …) as well as an I/O-aware strategy for indices associated with a disk read/write penalty. Moreover, indices can be individually pulled out from the tensor contraction to be processed in outer for-loops, which is highly useful for integrating the TTGT engine with other specialized functions as well as enabling easy parallelization over these outer loops. Effectively, such situations are handled automatically by reordering and *tiling* [244] the tensors such that the chunks can be treated most efficiently with large DGEMM calls and are trivial to fetch from disk or memory. This scheme is especially useful for compute-bound scenarios [262], e.g., contractions involving large and high-dimensional objects such as the four- and five-body densities, γ_4 and γ_5, respectively. Naïve loop code will almost certainly lead to many cache misses for such objects and can thus perform well over a factor of 100 slower for multi-MB or -GB tensors.

Also in the category of compute optimizations are further hand-coded routines. Additional contraction patterns have been added to the `ContractionEngine` functionality, which works with 4-index quantities stored in matrix containers discussed below, to perform tasks of the type

$$C_{rs}^{pq} = \sum_{p'q'} A_{pq}^{p'q'} B_{rs}^{p'q'}, \tag{4.2}$$

which complement the features already available since the first version of the code generator [1].

On the note of I/O improvements, we should first delineate how ORCA-AGE II deals with a mixed strategy of keeping some tensors on disk and some entirely in memory. Ultimately, the decision is up to the end user, but by default 4-index quantities, called *matrix containers*, are stored on disk with two indices encoding disk access, for which a matrix can be retrieved. For example, given X_{pq}^{rs} and using the lower indices for access, we can load matrices of size $r \times s$ for a given index pair p, q. These quantities are stored on disk to integrate nicely with the rest of the ORCA infrastructure as well as to avoid memory bottlenecks (e.g., the 4-external integrals $(ab|cd)$ already exceed 10 GB for $n_{\text{virt}} > 188$ virtual orbitals in double precision).

We optimized the I/O performance by both hand-coded contractions and a new on-the-fly resorting scheme. The hand-coded functions are mostly geared towards copy/add operations,

$$C_{ij\cdots} = A_{\pi(ij\cdots)}, \tag{4.3}$$

which repeatedly occur through the application of the distributive law. Here, $\pi(ij\cdots)$ denotes an arbitrary permutation of indices on the source tensor. These functions are tailored to the 4-index tensors that are stored on disk. Throughout the generated code, we try to minimize contractions that have many different indices associated with I/O operations in them by determining the best index order on the intermediates, but I/O-intensive operations such as

$$X_{pq}^{rs} = Y_{rs}^{pq}, \tag{4.4}$$

where none of the I/O-associated (lower) indices match, cannot always be avoided. In a naïve implementation, this will lead to I/O operations being done in inner loops, with the associated increase in runtime due to disk latency and unnecessary repetitions of I/O operations. To alleviate the I/O bottleneck of these additions, special functions were coded that read batches of the tensors up to the allowed memory limit, do the addition in memory, and then efficiently write the results back to disk.

A more general scheme was introduced by resorting the on-disk quantities on the fly. The underlying inefficiency is the same as discussed above, namely, that non-matching indices are associated with I/O operations. In the general case, however, we do not opt for a batching scheme, but rather exploit the fact that reordering such a 4-index tensor only scales as $\mathcal{O}(N^4)$ and can be done itself with efficient handwritten functions. Virtually all contractions that use such stored quantities scale higher than that, and hence the additional reordering step is negligible in cost. Furthermore, these resorted matrix containers can be kept until the end of the computation to be reused in multiple contractions, at the expense of additional disk space being used.

Compared to the first ORCA-AGE toolchain, the updated version has also been integrated much more tightly with the main ORCA source code. Developers can now automatically trigger a complete (re-)generation of specific, or all, generated code modules from within the main build environment for the following reasons: First, this ensures that the autogenerated modules are consistent and reliable, since the entire procedure required to produce the module has been programmed. This greatly simplifies maintenance as well as any breaking changes can be readily compared to a previous state that is known to be good in the version control system. Second, it guarantees that the modules always feature the latest additions or improvements to the generated code, e.g., additional hand-coded contraction functions, if a complete rebuild with the latest version of ORCA-AGE II is triggered before each release build. Effectively, this automatically brings any new (performance) improvements to the older compute modules, and keeps the code used in different modules consistent with each other.

4.2 Comparison to Other Toolchains

To summarize the features of ORCA-AGE II and show how it relates to other code generators used in quantum chemistry, we summarized the main code generation features in Tab. 4.1. To keep the size of the table manageable, we only included other codes that are sufficiently complete or have been more widely used.

Table 4.1 Overview and comparison of several code generators used in quantum chemistry with ORCA-AGE II. An "x" indicates that the feature of that column is supported by the software, a "—" signifies the absence, and "n/a" stands for "not applicable." For the equation generation columns, we indicate whether commutator algebra, Wick's theorem with normal order, or diagrams are used to derive the working equations. In the factorization step, the codes may support detection of duplicates, strength reduction for tensor contractions, using the distributive law (sometimes called "factorization") to reduce the prefactor, and (sub)expression elimination. For code generation, we indicate whether the generator directly produces "plain" code, i.e., code not relying on external libraries, or whether it uses external libraries or interprets the equations directly. We further indicate if a string-based approach is used, and whether the BLAS library is employed. Lastly, we specify whether parallelized code can be produced. The main references for the programs are: FEMTO [177], SQA [229], Smith3 [246, 281], Smith [60, 280], TCE [244], GeCCo [64, 195, 249], APG [199, 248], ORCA-AGE [1], ORCA-AGE II (this work), MRCC [255].

	Equation Generation			**Factorization**				**Code Generation**						
Name	*Comm.*	*Wick*	*Diagrams*	*Duplicates*	*Strength reduction*	*Distrib. law*	*Subexp. elim.*	*Code*	*Tensor libraries*	*Interpreter*	*String*	*BLAS*	*Parallel*	**Year**
FEMTO	—	x	—	x	x	partial	partial	x	—	—	—	x	x	2013
SQA	—	x	—	x	x	limited	—	x	—	—	—	x	—	2009
Smith3	—	x	internal[†]	n/a	x	x	x	x	—	—	—	x	x	2015
Smith	—	—	x	n/a	x	x	x	x	—	—	—	x	x	2008
TCE	—	x	—	x	x	x	x	x	—	—	—	x	x	2003
GeCCo	—	x	internal[†]	x	x	x	x	—	—	x	x	x	—	2008
APG	—	x	—	x	x	x	x	x	—	—	—	x	—	2001
ORCA-AGE (II)	x	x[‡]	—	x	x	x	x	x	—	—	—	x	x	2017, 2022
MRCC	—	—	x	n/a	x	x	x	—	—	x	x	x	x	2001

[†] While generating terms with Wick's theorem, topologically equivalent terms (diagrams) can be recognized and combined. Thus, ideas from diagrammatic approaches are effectively used internally.

[‡] Development in progress.

We now highlight a few salient features of the tools from Tab. 4.1. First, only ORCA-AGE relies on a commutator-based engine to derive the working equations for simplicity and generality. However, using a Wick's theorem or diagram-based engine can be significantly faster and reduce the time required to remove equivalent contractions. Hence, efforts to develop such an engine is ongoing for ORCA-AGE II. Second, great care is placed by all tools on a complete factorization toolchain, which is a testament to the importance of this step in obtaining a performant implementation. Third, the actual evaluation part either relies on generation of actual code or simply interprets the factorized equations akin to a specialized virtual machine. To this end, all tools rely on BLAS to maximize computational efficiency. Parallelization support, however, is not implemented for all toolchains since doing so in a completely automated fashion is highly non-trivial.

4.3 Internally Contracted Multireference Coupled-Cluster Theory

fic-MRCC theory (see Sec. 2.6.3) is difficult to implement due to its large number of terms [195]; so many terms, in fact, that an implementation by hand is beyond human capacity. These terms arise from the non-commuting operators in the excitation classes, Eq. (2.115), and the commutation rule for spin-free excitation operators, Eq. (2.43). In our formulation, fic-MRCC theory has 159,173 terms before removing contractions redundant by symmetry and 49,186 terms after removal thereof. After factorization, the resulting 65,848 tensor contractions are translated to 3,765,810 lines of code.

A theory having so many terms places a high computational workload on the code generator, which was the primary reason for redeveloping the ORCA-AGE infrastructure. In particular, removing duplicate intermediates in the equation factorization step proved to be highly demanding, as that step creates 93,803 intermediates in the case of fic-MRCC theory, which are eventually pruned to 15,987 intermediates. In the previous version of ORCA-AGE, a simple algorithm for detecting duplicates was used that scaled as $\mathcal{O}(N^2)$, thus leading to impractically long generation times of about one month. In the updated version, this algorithm uses a hash table-based data compression technique with expected linear time instead, reducing the time to about 4 hours for the entire toolchain. In both cases, the toolchain, which has time-limiting steps parallelized with OpenMP threads, was run on 16 cores with 32 hyperthreads.

If we now look at the generator from an end user's perspective, generating the code is also much simpler. All that is required is to define the excitation classes from Eq. (2.115) in a user-friendly input file, an example of which is given in Listing 4.2. Each line corresponds to an excitation class with an identifier, the amplitudes, and the excitation operators for that class. Optionally, contravariant projection operators [33] can be defined such that the projection functions are biorthonormal with functions in the excitation space. The generator itself features code that can create all necessary terms from the definition of the excitation classes, be it the similarity-transformed Hamiltonian through application of the Baker-Campbell-Hausdorff

(BCH) expansion, residuals through projection or the energy equations. Hence, everything from changing the cluster operator to the number of terms retained in the BCH expansion is simple to program and quick to extend to other methods as well.

Listing 4.2 Definition of the excitation classes for fic-MRCC theory (see also Eq. 2.115). Each line provides an identifier, an amplitude, the excitation operators contained in the cluster operator $\hat{T}$, and (optionally) contravariant projection functions [33] for the residuals.

```
$class IJAB Cijab(ab,ij) 0.5 E(a,i)E(b,j); 1/6 (2 E(a,i)E(b,j) + E(b,i)E(a,j))
$class IJTA Cijta(ta,ij)     E(t,i)E(a,j); 1/3 (2 E(t,i)E(a,j) + E(a,i)E(t,j))
$class IJTU Cijtu(tu,ij) 0.5 E(t,i)E(u,j); E(t,i)E(u,j)
$class ITAB Citab(ab,it)     E(a,i)E(b,t); 1/3 (2 E(a,i)E(b,t) + E(b,i)E(a,t))
$class ITAU Citau(at,iu)     E(a,i)E(t,u)
$class ITUA Citua(ta,iu)     E(t,i)E(a,u)
$class ITUV Cituv(ut,iv)     E(u,i)E(t,v)
$class TUAB Ctuab(ab,tu) 0.5 E(a,t)E(b,u); E(a,t)E(b,u)
$class TUVA Ctuva(at,uv)     E(a,u)E(t,v)
```

To exemplify the efficiency of the generated code, we consider two limiting cases in fic-MRCC theory: first, large systems with small active spaces, and second, systems with large active spaces. We further note that the contractions given below are asymptotically limiting, i.e., they need not be the most expensive contractions for small systems. Large systems with small active spaces, e.g., using triple-ζ basis sets and a CAS(2,2), are dominated by the 4-external term, which also appears in single-reference CC theory,

$$r_{ij}^{ab} = \sum_{cd} (ac|bd) t_{ij}^{cd} . \tag{4.5}$$

This term scales as $\mathcal{O}(n_{\text{inact}}^2 n_{\text{virt}}^4)$ and thus dominates the computational time for $n_{\text{virt}} \gg n_{\text{inact}}, n_{\text{act}}$. Therefore, it is crucial to fully optimize this term, which is complicated by the fact that the 4-external integrals must be stored on disk for their large size. This is where the hand-coded `ContractionEngine` functionality is employed: it minimizes I/O through reading large batches of the amplitudes and integrals (up to a given memory limit) and then applies a large-scale DGEMM operation to have maximum computational efficiency [1]. This strategy is also followed in the entirely handwritten CC code in ORCA.

The other extreme is given by systems with large active spaces. In these cases, the runtime is dominated by contractions such as

$$Y_{qsa}^{prt} = \sum_{p'q'uvw} X_{vp'a}^{uwq'} \gamma_{rpuqq'}^{tswvp'} , \tag{4.6}$$

which scales as $\mathcal{O}(n_{\text{act}}^{10} n_{\text{virt}})$ and consequently dominates the other contractions for $n_{\text{virt}} > n_{\text{inact}}$, $n_{\text{act}} \gtrsim 6$. The tensor γ_5 is difficult for the CPU to manage because of its complicated structure

Listing 4.3 Generated code for the time-limiting step of fic-MRCC theory involving a five-body density γ_5. This is an example of a BLAS call generated by the TTGT engine (Sec. A.3).

```
for (int p3 = 0; p3 < NActive; ++p3)
 for (int p6 = 0; p6 < NActive; ++p6)
  for (int t0 = 0; t0 < NActive; ++t0)
   for (int u0 = 0; u0 < NActive; ++u0)
    for (int v0 = 0; v0 < NActive; ++v0)
     for (int P0 = 0; P0 < NActive; ++P0)
      for (int P5 = 0; P5 < NActive; ++P5)
       for (int P4 = 0; P4 < NActive; ++P4)
        for (int P1 = 0; P1 < NActive; ++P1)
         for (int P2 = 0; P2 < NActive; ++P2)
          DC5_T({p3,p6,t0,u0,v0,P0,P5,P4,P1,P2}) = DC5({v0,t0,u0,p3,P4,P0,P5,p6,P1,P2});

BLAS_Add_Mat_x_Mat(false, false, NActive * NActive * NActive * NActive * NActive, NVirtual,
↪  NActive * NActive * NActive * NActive * NActive, 1.0, &DC5_T({0,0,0,0,0,0,0,0,0,0}),
↪  NActive * NActive * NActive * NActive * NActive, X({0,0,0,0,0,0}), NVirtual, 1.0,
↪  Y({0,0,0,0,0,0}), NVirtual);
```

with ten dimensions as well as its size (e.g., 80 GB for $n_{\text{act}} = 10$), thus leading to frequent cache misses if implemented naively with eleven nested loops. Such contractions are ideal candidates for the TTGT engine discussed above, as BLAS libraries carefully optimize both algorithms and cache usage. In this case, the generated code is reproduced in Listing 4.3.

The TTGT engine first reorders the five-body density `DC5` to optimally align and group the indices together, thus enabling a single large-scale DGEMM instruction to digest the contraction. As a result, explicit loops are only used to reorder the tensors, while the actual contraction is done with implicit loops in the DGEMM matrix-matrix multiplication. Compared to the same contraction implemented with plain `for`-loops, the TTGT scheme easily leads to a speedup by a factor > 100. All tensors are kept in memory since the space requirements are not as high as for the 4-external term and the number of active indices is limited. The TTGT engine is also capable of accounting for indices associated with I/O and will dynamically adapt to load (sub-)tensors from disk in an optimal fashion and perform a DGEMM over the remaining non-I/O indices, if required.

The improvements made to the code generator in ORCA-AGE II now allow quite large systems to be computed with the fic-MRCC implementation. We will now present two calculations of tetradecene, $C_{14}H_{28}$, and naphthalene, $C_{10}H_8$, for which the requirements are representative of the computational resources for large systems with small active spaces and systems with large active spaces, respectively. Both systems were run in serial on an AMD EPYC™ 75F3 processor and use the def2-SVP basis set [463]. The tetradecene system uses a small CAS(2,2) active space. Thus, its 112 electrons are filled into 14 frozen core, 41 inactive, 2 active, and 279 virtual

orbitals (336 in total). On average, a single iteration takes 6.58 hours. The calculation converges smoothly in 12 iterations for a total runtime of 3.3 days, a timeframe that is manageable for routine calculations. The naphthalene system, on the other hand, uses a CAS(10,10), and distributes 68 electrons among 10 frozen core, 19 inactive, 10 active, and 141 virtual orbitals (180 in total). Each iteration requires 11.35 hours to complete, and an additional 17.46 hours are spent computing the five-body density, which takes up 80.00 GB of memory. The total runtime for all 16 iterations is 8.3 days. After full parallelization support, we expect systems of 600–700 basis functions to be practical with fic-MRCC theory (cf. Ch. 7 and Tab. B.6).

From these two examples, we see that increasing the size of the active space severely limits the system size that remains computationally feasible. This is both due to the asymptotic scaling, $\mathcal{O}(n_{\mathrm{act}}^{10} n_{\mathrm{virt}})$, as well as the difficulty of handling a large, high-dimensional tensor like γ_5 for the CPU, as discussed above. Calculations of systems with 12 active orbitals are unfeasible for fic-MRCC theory as computing the five-body density in that case takes 28.4 days, while requiring 495.34 GB of memory. To meet this challenge, we developed a version of fic-MRCC theory that only requires the four-body density. Hence, the new implementation allows even larger active spaces, e.g., CAS(12,12) on biphenyl to be computed in 5.4 days per iteration. We present the full reduced-scaling scheme in Ch. 7.

In summary, ORCA-AGE II has been proven to be a performant toolchain that generates near-optimal code for the time-limiting steps of fic-MRCC theory. The code optimization strategies illustrated here are, of course, also applied to the non-limiting contractions in the code. However, despite the performant implementation, fic-MRCC theory remains a very time-intensive method due to the large number of contractions it contains.

5 Study on the Excited States of Indigo Dyes

5.1 Indigo Dyes and Prior Work

In this chapter, we present a benchmark study of indigo dyes where we predict the optical transition energy including vibronic effects and compare the results against experimental reference values. To begin with, indigo dyes form an entire class of compounds with different substitution patterns, the most well-known representatives of which are indigo and Tyrian purple. Indigo is one of the oldest known organic dyes, its use dating back to the 18th dynasty in ancient Egypt [464]. It was previously produced from several species of plants, which contain the precursors indican along with isatan B and C, which yield indigo upon hydrolysis and oxidation [465]. The first total synthesis was reported by von Baeyer in 1870 [466], well before he also reported the structure of indigo in 1883 [467]. Nowadays, indigo is almost exclusively synthesized industrially [468] based on a process originally due to Heumann [469, 470]. Industrial production reached 50,000 metric tons in 2011, most of which is being used for dyeing garments [471]. Tyrian purple, or 6,6'-dibromoindigo, has no commercial importance contemporarily, but is still infamous for its high price when produced the traditional way from the secretions of a few species of shellfish [472]. Its synthesis was first reported in 1903 by Sachs and Kempf [473], again prior to structural elucidation by Friedländer in 1909 [474].

The first theoretical studies, as opposed to the synthetic work, were performed by Lüttke and co-workers [475–477], who used the semiempirical Hückel [478–481] and Pariser-Parr-Pople (PPP) [482–484] methods to elucidate spectroscopic properties. Later, using a more rigorous approach with time-dependent density functional theory (TD-DFT), Jacquemin *et al.* [485] published a broadly scoped study, including 86 solvated indigoid dyes, for each of which the main optical transition was computed and compared to a database of experimental values. In its wake, further TD-DFT studies were published, each focusing on a different aspect of the dyes [486–490]. In another TD-DFT study of indigo and its smaller model systems by Barone *et al.* [461], vibronic effects were thoroughly investigated.

Fewer studies using *ab initio* wave function theory have been published to date. The first such investigation by Serrano-Andrés and Roos [491] employed CASSCF and CASPT2 theory to compute electronic transition energies and confirmed that the transition in the visible spectrum is due to the central chromophore of indigo. Domcke and co-workers [492] additionally used second-order approximate coupled-cluster (CC2) theory in their study of the photostability of indigo. To the best of our knowledge, no comprehensive benchmark has been published to date using wave function theory. Therefore, we present the transition energies computed with

the STEOM method [10–12] on the test set originally proposed by Jacquemin *et al.* [485] in the following. As it was shown, for example, by Stanton that both the electronic and nuclear contributions are important to reproduce absorption spectra quantitatively [493], we include vibronic contributions as in Barone *et al.* [461] through time-dependent perturbation theory [459], which has been found to yield highly accurate absorption and emission spectra for BODIPY dyes [103] and chlorophyll *a* [101]. Also, solvation effects are treated implicitly through a conductor-like polarizable continuum solvation model [451].

5.2 Computational Details and Experimental Methods

All calculations in this chapter were carried out with a development version of the ORCA program package [287].

Geometry optimizations were performed at the DFT level using the B3LYP functional [494] with the def2-TZVP basis set [463] and the matching auxiliary basis set [495]. The D3 dispersion correction as parametrized by Grimme and co-workers was also applied [496]. The B3LYP functional has already been used with success to obtain geometries in our earlier studies on other families of dyes [103]. The geometries obtained in this study can be obtained as described in App. B.1.6. The RIJCOSX [497, 498] approach applying the resolution of identity (RI) approximation to the Coulomb part and the chain of spheres (COS) seminumerical integration algorithm to the exchange term was also used to accelerate the optimization process. Harmonic vibrational frequencies were computed at the same level of theory. The `TightSCF` keyword was used to set convergence criteria for energy calculations (thresholds in atomic units: 10^{-8} for energy and 10^{-5} for the orbital gradient) and the `TightOpt` keyword for geometry optimizations (thresholds in atomic units: 10^{-6} for energy and 10^{-4} for the maximum component of the gradient). The DFT grid was set to `DEFGRID2` (Lebedev angular grid with 302 points and a Gauss-Chebyshev radial grid with a radial integration accuracy of 4.67) and all the other parameters were chosen as default.

Vertical excitation energies and transition dipoles were computed with DLPNO-STEOM-CCSD on the DFT geometries. The basis set dependence of DLPNO-STEOM-CCSD calculations was studied using the def2-SVP, def2-TZVP and def2-QZVPP basis sets [463, 495]. The RIJCOSX [497, 498] approximation was used in the STEOM integral dressing step. Five roots were requested using ORCA's `TightPNO` settings [449], i.e., setting the following thresholds: $T_{\texttt{CutPairs}} = 10^{-5}$, $T_{\texttt{CutDO}} = 5 \times 10^{-3}$, $T_{\texttt{CutPNO}} = 10^{-7}$, $T_{\texttt{CutMKN}} = 10^{-3}$. Since the quality of the singles (diagonal) PNOs is especially important for the EA problem [95], these are generated as a separate set using a tighter than usual threshold of $T_{\texttt{CutPNOSingles}} = 10^{-11}$. Furthermore, to have a sufficiently large active character of the computed states (%active), the cutoff values for the natural orbital occupation numbers in the active space selection procedure for the occupied and virtual orbitals (`Othresh` and `Vthresh`, respectively) were set to 5.0×10^{-3} [443]. The CPCM model [451] was used throughout to account for implicit solvent effects. The list of the

parameters used for the solvents can be found in the Appendix (Tab. B.1).

For all dyes studies here, the computed energies are associated with the HOMO → LUMO transition (weight above 80%). The lowest-lying electronic state is well separated from the other ones, thus it is enough to compare this with experiment. Based on Fig. B.1, we found that the (first-order) Herzberg-Teller approximation does not significantly improve the normalized absorption spectra of indigoid dyes, while it is significantly more expensive than a Franck-Condon calculation. For this reason, vibronic effects were computed at the Franck-Condon level using the `ESD` module of ORCA [459] from the STEOM excitation energy and transition dipole moment and DFT frequencies. A Gaussian line shape is used to model the absorption spectra together with a linewidth of 700 cm^{-1}. The theoretical and experimental spectra were renormalized. The VG model was used without recalculating the STEOM energy at the displaced geometry, and the correction for the 0–0 transition energy was obtained from the excited state TDDFT gradient and ground state DFT frequencies (`VGFC` keyword).

In the following discussion, the statistical parameters "mean absolute error" (MAE), "mean error" (ME), "maximal positive and negative deviations" (Max(+) and Max(−)), and "standard deviation" (SD) will be used in the statistical characterization of our results.

For a quantitative prediction of absorption spectra, it is desirable that the computational method be able to reproduce the positions of band maxima within about 0.1 eV. Our earlier experience shows that the DLPNO-STEOM method can deliver such accuracy for various dyes [100, 103]. In the present study, the settings have also been chosen in such a way that the expected accuracy should be below 0.1 eV and this threshold value will also be used in judging whether deviations are significant or not.

The UV-VIS experiments have been carried out using an Ocean Optics Deuterium-Tungsten Halogen DH-2000-BAL light source and an Ocean Optics Flame detector (FLMS00699) with Ocean Optics QP400-1-UV-VIS glass fibers for acquisition. The data were collected with the OceanView 2.0.7 software in absorption measurement mode, and an average of 100 scans of 3 ms integration time was performed for each measurement. The numerical data for the plotted spectra can be obtained as stated in App. B.1.7.

For the spectra, commercially available dyes were used: indigo (95%) from Sigma-Aldrich, France and genuine Tyrian purple from Kremer Pigmente, Germany. These materials were used without further purification.

5.3 Multireference Effects

Prior to computing the STEOM transition energies on the indigoids, we assessed whether these systems exhibit significant multireference character. Earlier studies using multiconfigurational wave functions include those carried out by Serrano-Andrés and Roos [491] as well as Domcke and co-workers [492], including perturbative corrections to second order.

For our investigation, we limited ourselves to the so-called "H chromophore" (Scheme 5.1),

which has already been used in semi-empirical studies by Lüttke *et al.* [477] and Dähne and Leupold [499], and was also considered by Serrano-Andrés and Roos [491]. Given a small basis set, this compound can be treated with highly correlated multireference methods using a reasonably large active space of six electrons in six active orbitals (CAS(6,6), including the full central π-system). We did not consider including the n-electrons in the active space since these do not contribute significantly to the $S_0 \rightarrow S_1$ transition [491].

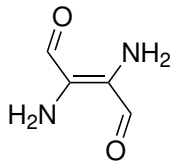

Scheme 5.1 Simplified central "H chromophore" of indigo.

The computed energies for the S_0 and S_1 states can be found in Tab. 5.1. We should note that all multireference correlated methods, fic-NEVPT2 [169, 171], uncontracted (uc-)MRCI, and fic-MRCC (see Sec. 2.6.3) use the same SA-CASSCF wave function averaged over the S_0 and S_1 states. Furthermore, all methods presented in Tab. 5.1 are truncated after doubles excitations. The fic-NEVPT2 results are comparable to previous reports [491] and match the results from a DLPNO-STEOM calculation very well. A higher energy difference is reported by the fic-MRCC and the single-reference EOMCC methods at about 3.5 eV and 3.4 eV, respectively, with the results from the uc-MRCI falling roughly between these and the STEOM results.

Table 5.1 Total energies for the S_0 and S_1 states of the H chromophore (E_h). The excitation energy to the S_1 state, ω, is computed in eV. Here, we understand "STEOM" to mean "DLPNO-STEOM" with identical settings as in the remaining calculations in this chapter. All other methods are canonical.

	CASSCF	**fic-NEVPT2**	**uc-MRCI**	**fic-MRCC**	**EOMCC**	**STEOM**
S_0	−413.304138	−414.457377	−414.457254	−414.507567	−414.492813	−414.492887
S_1	−413.129797	−414.346606	−414.341664	−414.379725	−414.368914	−414.381096
ω	4.744	3.014	3.145	3.479	3.371	3.042

Thus, we conclude that a genuine multireference treatment is not required. This is mainly evidenced by the fact that, in the state-averaged multiconfigurational wave function, the S_0 and S_1 states are dominated ($\approx$ 90%) by a single configuration, and the transition can be characterized as singles-dominated $\pi \rightarrow \pi^*$. This could also explain why the fic-MRCC and EOMCC results are rather close and the DLPNO-STEOM results are lower by $\sim$ 0.4 eV: STEOM-CCSD contains an implicit triples term absent in both EOM-CCSD [12, 88] and in fic-MRCCSD. Since the wave function is dominated by two configurations, implicit higher-order excitations do not play a significant role in the fic-MRCC results.

Furthermore, EOMCC methods can be used to treat multiconfigurational situations given that

certain parameters are fulfilled. For EOM-CCSD, this usually means that the singles character is required to be greater than $\%T_1 > 90\%$, and in STEOM calculations the active space character should be $> 95\%$ [97, 443]. Both criteria were fulfilled in our calculations.

5.4 Basis Set Effects

In accordance with our earlier experience [103], triple-ζ basis sets seem to be a good compromise between accuracy and efficiency when calculating the optical spectra of dyes. The Max(+) values in Tab. 5.2 are calculated using the def2-QZVPP values as reference and they indicate that the def2-SVP basis set does not produce converged results, with its Max(+) being 0.19 eV. While these calculations are fast (about 20 min on 8 cores for Indigo) and yield a tolerable accuracy for some purposes, it is still above our target threshold of 0.1 eV. The def2-TZVP values, on the other hand, deviates only up to Max(+) = 0.03 eV and 0.01 eV on average from the def2-QZVPP results, which is not only well below the desired threshold but also still cheap. A def2-TZVP calculation on Indigo takes about 2 hours on 8 cores, while obtaining the slightly better def2-QZVPP values takes about 7.5 hours on 16 cores. Thus, the def2-TZVP basis set will be used in the remainder of this chapter for all geometry optimization, STEOM and ESD calculations.

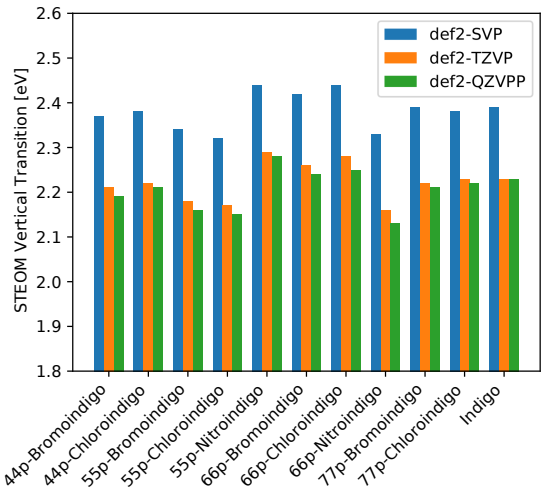

Figure 5.1 Vertical transition energies computed by the DLPNO-STEOM method for a set of symmetrically substituted indigo dyes. Data plotted from Tab. 5.2.

Table 5.2 Vertical transition energies computed by the DLPNO-STEOM method for a set of symmetrically substituted indigo dyes. All values are given in eV.

Compound	**def2-SVP**	**def2-TZVP**	**def2-QZVPP**
Indigo	2.39	2.23	2.23
4,4'-Bromoindigo	2.37	2.21	2.19
4,4'-Chloroindigo	2.38	2.22	2.21
5,5'-Bromoindigo	2.34	2.18	2.16
5,5'-Chloroindigo	2.32	2.17	2.15
5,5'-Nitroindigo	2.44	2.29	2.28
6,6'-Bromoindigo	2.42	2.26	2.24
6,6'-Chloroindigo	2.44	2.28	2.25
6,6'-Nitroindigo	2.33	2.16	2.13
7,7'-Bromoindigo	2.39	2.22	2.21
7,7'-Chloroindigo	2.38	2.23	2.22
Deviation from def2-QZVPP			
Max(−)	0.16	0.00	
ME	0.17	0.01	
Max(+)	0.19	0.03	
SD	0.01	0.01	
MAE	0.17	0.01	

5.5 Solvent Effects

As a last step before calculating the full benchmark set from Jacquemin *et al.* [485], the magnitude of the CPCM implicit solvation effects on the STEOM calculations will also be assessed. To this end, we selected a few symmetrically substituted indigo dyes from the full set, including Tyrian purple (6,6'-dibromoindigo), and optimized the geometries at the B3LYP-D3(BJ)-CPCM/def2-TZVP level of theory. Subsequently, we computed the DLPNO-STEOM vertical transition energies for each of the optimized geometries, again including solvation effects through CPCM. The results are presented in Tab. 5.3. As expected, the average effect grows as the relative permittivity increases and ranges between 0.09 eV and 0.16 eV. For the least polar solvents, CCl_4, Xylene and Benzene, the effect is about 0.10 eV, for the solvents of medium polarization, $CHCl_3$ and TCE, it is about 0.135 eV, and for the most polar ones, ethanol and DMSO, the effect is about 0.155 eV. Tab. B.2 of the Appendix also indicates the magnitude of the indirect solvent effects that are due to solvating the molecular orbitals and the ground state CC amplitudes. It turns out that on average 75% of the solvent shift is due to indirect effects and this ratio slightly increases with polarity. Nevertheless, the direct contributions from the STEOM step are not negligible. As

also discussed by Jacquemin *et al.* [485], the effects could be larger for polar solvents than the CPCM numbers indicate. This is especially true for ethanol, where the formation of hydrogen bonds introduces specific interactions that are not included in the CPCM model. Although we do not plan to account for such specific interactions in this study, the inclusion of solvent effects using CPCM is certainly necessary. All subsequent calculations contain the CPCM correction for the various solvents considered.

Table 5.3 Solvation effects compared to gas-phase STEOM calculations. Solvents sorted with increasing dielectric constant ε from left to right. All values are given in eV.

	CCl_4	**Xylene**	**Benzene**	**$CHCl_3$**	**TCE**	**EtOH**	**DMSO**
ε	2.24	2.27	2.28	4.9	8.2	24.3	47.2
Indigo	−0.099	−0.101	−0.101	−0.137	−0.153	−0.155	−0.164
4,4'-Bromoindigo	−0.100	−0.102	−0.102	−0.137	−0.156	−0.165	−0.173
4,4'-Chloroindigo	−0.102	−0.104	−0.104	−0.140	−0.155	−0.166	−0.175
5,5'-Bromoindigo	−0.093	−0.095	−0.095	−0.124	−0.138	−0.145	−0.153
5,5'-Chloroindigo	−0.096	−0.097	−0.098	−0.121	−0.135	−0.142	−0.150
5,5'-Nitroindigo	−0.082	−0.083	−0.075	−0.109	−0.123	−0.128	−0.138
6,6'-Bromoindigo	−0.100	−0.102	−0.103	−0.135	−0.151	−0.156	−0.168
6,6'-Chloroindigo	−0.102	−0.104	−0.104	−0.138	−0.153	−0.162	−0.170
6,6'-Nitroindigo	−0.106	−0.127	−0.127	−0.157	−0.175	−0.185	−0.195
7,7'-Bromoindigo	−0.078	−0.080	−0.081	−0.108	−0.120	−0.126	−0.134
7,7'-Chloroindigo	−0.087	−0.089	−0.089	−0.116	−0.132	−0.134	−0.142
Average	−0.095	−0.099	−0.098	−0.129	−0.145	−0.151	−0.160

5.6 Benchmark Results

The main results of our study are presented in Tab. 5.4. To facilitate comparison with experimental absorption band maxima, the CPCM-corrected DLPNO-STEOM values still need to be corrected for vibronic effects to arrive at a computational estimate of the 0–0 transition energy. At the VG level, this correction can be obtained from DFT frequencies and displacements. This is a relatively small correction: while the STEOM energies are all close to 2.0 eV within a few tenths of eVs, the VG shift, ω_{VG}, is typically −0.05 eV, see Tab. 5.5. Summing up these values yields the DLPNO-STEOM 0–0 energies, $\omega_{0\text{–}0}$, which can be compared to experiment.

The deviations between predicted and experimental values ($\Delta\omega$; averaged experimental values computed by Jacquemin *et al.* [485]; for individual measurements see references therein) are presented in Tab. 5.4 as well as in Fig. 5.2, while the statistics can be found in Tab. 5.5. On average, the STEOM error is −0.02 eV for this family of dyes, the maximum deviation being −0.11 eV. This indicates that in general, the computed values are lower than the measured ones. The deviation from experiment tends to be larger for polar solvents, as expected. Nevertheless,

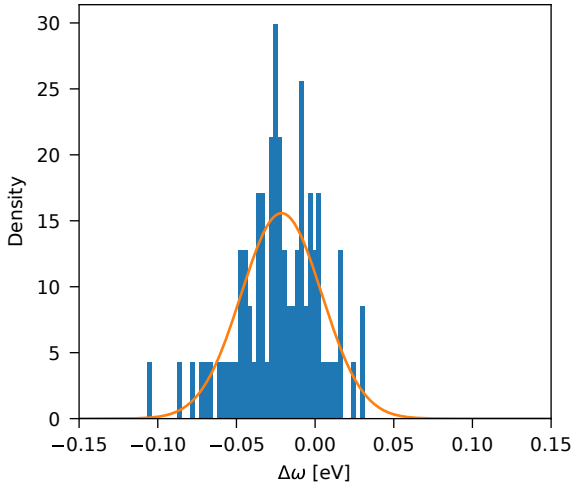

Figure 5.2 Normalized error density distribution with respect to experiment, $\Delta\omega$, from Tab. 5.4, with a Gaussian probability density function overlaid.

even for the most polar EtOH and DMSO solvents, the average deviation from experiment is only −0.04 eV. The protocol we present here consists of many components including the DLPNO-STEOM vertical excitation energies, the CPCM correction calculated from STEOM amplitudes using the CIS formula and the VG shift obtained at the DFT level. The performance of these have been evaluated independently, to the extent this is possible. Thus, the eventual excellent agreement between the predicted and measured 0–0 transitions indicates the reliability of the proposed protocol and it certainly agrees with our initial goal of providing an affordable computational method that delivers excitation energies within 0.1 eV for a large number of chemically relevant molecules.

Table 5.4 Benchmark results: the CPCM-corrected DLPNO-STEOM excitation energy (ω), the VG correction ω_{VG}, the 0–0 transition energy ($\omega_{0\text{–}0}$), the experimental absorption band maxima (ω_{exp}) [485] and the error of the total STEOM estimate with respect to experiment ($\Delta\omega$). All values reported in eV.

Name	Solvent	ω	ω_{VG}	$\omega_{0\text{–}0}$	ω_{exp}	$\Delta\omega$
Indigo	Benzene	2.13	−0.05	2.09	2.08	0.00
Indigo	CCl_4	2.13	−0.05	2.09	2.09	0.00
Indigo	$CHCl_3$	2.10	−0.05	2.05	2.05	0.00
Indigo	DMSO	2.07	−0.05	2.02	2.00	0.02
Indigo	EtOH	2.08	−0.05	2.03	2.04	−0.01
Indigo	TCE	2.08	−0.05	2.03	2.02	0.01

continued

Table 5.4 Benchmark results (*continued*).

Name	Solvent	ω	ω_{VG}	ω_{0-0}	ω_{exp}	$\Delta\omega$
4-Bromoindigo	TCE	2.08	−0.05	2.03	2.03	−0.01
5-Bromoindigo	Benzene	2.12	−0.05	2.07	2.05	0.02
5-Bromoindigo	TCE	2.07	−0.05	2.02	2.02	0.00
6-Bromoindigo	TCE	2.08	−0.05	2.03	2.06	−0.03
7-Bromoindigo	TCE	2.08	−0.05	2.03	2.05	−0.02
4,4'-Bromoindigo	TCE	2.06	−0.05	2.00	2.03	−0.02
4,4'-Chloroindigo	$CHCl_3$	2.08	−0.05	2.03	2.05	−0.02
4,4'-Chloroindigo	EtOH	2.06	−0.05	2.00	2.03	−0.02
4,4'-Chloroindigo	TCE	2.07	−0.05	2.01	2.02	−0.01
4,4'-Chloroindigo	Xylene	2.12	−0.05	2.07	2.07	0.00
4,4'-Methoxycarbonylindigo	EtOH	2.03	−0.05	1.98	2.03	−0.04
4,4'-Methoxyindigo	TCE	2.09	−0.06	2.03	2.06	−0.02
5,5'-Bromoindigo	$CHCl_3$	2.06	−0.05	2.00	2.03	−0.03
5,5'-Bromoindigo	EtOH	2.03	−0.05	1.98	2.01	−0.03
5,5'-Bromoindigo	TCE	2.04	−0.05	1.99	2.00	−0.01
5,5'-Bromoindigo	Xylene	2.08	−0.05	2.03	2.05	−0.02
5,5'-Chloroindigo	EtOH	2.02	−0.05	1.97	2.02	−0.04
5,5'-Chloroindigo	TCE	2.03	−0.05	1.98	2.00	−0.02
5,5'-Chloroindigo	Xylene	2.07	−0.05	2.02	2.05	−0.03
5,5'-Fluoroindigo	$CHCl_3$	2.02	−0.05	1.97	2.01	−0.04
5,5'-Fluoroindigo	TCE	2.01	−0.05	1.96	2.02	−0.06
5,5'-Methoxyindigo	TCE	1.90	−0.07	1.83	1.91	−0.08
5,5'-Methylindigo	$CHCl_3$	2.04	−0.05	1.99	2.01	−0.02
5,5'-Methylindigo	TCE	2.03	−0.05	1.97	2.00	−0.03
5,5'-Nitroindigo	TCE	2.16	−0.05	2.11	2.14	−0.03
6,6'-Bromoindigo	TCE	2.11	−0.05	2.06	2.11	−0.05
6,6'-Bromoindigo	Xylene	2.16	−0.04	2.11	2.10	0.01
6,6'-Chloroindigo	EtOH	2.12	−0.05	2.07	2.18	−0.11
6,6'-Chloroindigo	TCE	2.13	−0.05	2.08	2.10	−0.02
6,6'-Chloroindigo	Xylene	2.18	−0.04	2.13	2.22	−0.08
6,6'-Fluoroindigo	TCE	2.20	−0.05	2.14	2.18	−0.03
6,6'-Methoxyindigo	TCE	2.24	−0.06	2.18	2.16	0.02
6,6'-Methylindigo	TCE	2.11	−0.05	2.06	2.08	−0.02
6,6'-Methylindigo	Xylene	2.15	−0.05	2.10	2.11	−0.01
6,6'-Nitroindigo	TCE	1.98	−0.05	1.93	1.95	−0.02

continued

Table 5.4 Benchmark results (*continued*).

Name	Solvent	ω	ω_{VG}	ω_{0-0}	ω_{exp}	$\Delta\omega$
6,6'-Nitroindigo	Xylene	2.03	−0.04	1.99	1.96	0.03
7,7'-Bromoindigo	TCE	2.10	−0.06	2.05	2.05	0.00
7,7'-Chloroindigo	EtOH	2.09	−0.06	2.04	2.10	−0.07
7,7'-Chloroindigo	TCE	2.09	−0.06	2.04	2.06	−0.02
7,7'-Fluoroindigo	$CHCl_3$	2.09	−0.05	2.04	2.00	0.03
7,7'-Fluoroindigo	TCE	2.08	−0.06	2.02	2.05	−0.03
7,7'-Methoxyindigo	TCE	1.97	−0.08	1.89	1.93	−0.03
7,7'-Methylindigo	$CHCl_3$	2.08	−0.06	2.02	2.03	−0.01
7,7'-Methylindigo	Xylene	2.11	−0.06	2.06	2.05	0.00
5,5',7-Bromoindigo	Xylene	2.09	−0.05	2.04	2.03	0.01
4,4',5,5'-Bromoindigo	Benzene	2.07	−0.05	2.02	2.02	0.00
4,4',5,5'-Chloroindigo	TCE	2.03	−0.05	1.98	1.99	−0.01
4,4',6,6'-Chloroindigo	TCE	2.09	−0.05	2.04	2.08	−0.04
4,4',6,6'-Methylindigo	Xylene	2.13	−0.05	2.08	2.09	0.00
4,4',7,7'-Chloroindigo	TCE	2.06	−0.06	2.01	2.03	−0.03
4,4',7,7'-Chloroindigo	Xylene	2.11	−0.05	2.05	2.08	−0.02
4,4',7,7'-Methylindigo	Xylene	2.07	−0.06	2.01	2.06	−0.05
5,5',6,6'-Chloroindigo	TCE	2.05	−0.05	2.00	2.05	−0.05
5,5',6,6'-Methylindigo	Xylene	2.10	−0.05	2.04	2.05	−0.01
5,5',7,7'-Bromoindigo	Benzene	2.08	−0.06	2.03	2.02	0.01
5,5',7,7'-Bromoindigo	$CHCl_3$	2.06	−0.06	2.00	2.02	−0.01
5,5',7,7'-Bromoindigo	EtOH	2.04	−0.06	1.98	2.00	−0.02
5,5',7,7'-Bromoindigo	TCE	2.05	−0.06	1.99	2.00	0.00
5,5',7,7'-Bromoindigo	Xylene	2.08	−0.06	2.03	2.02	0.00
5,5',7,7'-Chloroindigo	EtOH	2.03	−0.06	1.97	2.02	−0.04
5,5',7,7'-Chloroindigo	TCE	2.04	−0.06	1.98	2.00	−0.02
5,5',7,7'-Chloroindigo	Xylene	2.07	−0.06	2.01	2.03	−0.02
6,6',7,7'-Chloroindigo	TCE	2.12	−0.05	2.07	2.10	−0.03
6,6',7,7'-Methylindigo	Xylene	2.08	−0.05	2.03	2.08	−0.06
4,4'-Chloro-5,5'-bromoindigo	$CHCl_3$	2.04	−0.05	1.99	2.03	−0.04
4,4'-Chloro-5,5'-bromoindigo	TCE	2.02	−0.05	1.97	2.00	−0.03
4,4'-Chloro-5,5'-bromoindigo	Xylene	2.07	−0.05	2.02	2.02	0.00
5,5'-Chloro-7,7'-bromoindigo	Xylene	2.07	−0.06	2.02	2.03	−0.01
4,4',5,5',6,6'-Chloroindigo	TCE	2.03	−0.05	1.98	2.03	−0.05
4,4',5,5',7,7'-Bromoindigo	$CHCl_3$	2.03	−0.06	1.98	1.99	−0.01

continued

Table 5.4 Benchmark results (*continued*).

Name	Solvent	ω	ω_{VG}	ω_{0-0}	ω_{exp}	$\Delta\omega$
4,4',5,5',7,7'-Bromoindigo	Xylene	2.05	−0.05	1.99	2.01	−0.02
4,4',5,5',7,7'-Chloroindigo	TCE	2.03	−0.06	1.98	2.02	−0.04
4,4',5,5',7,7'-Chloroindigo	Xylene	2.07	−0.05	2.02	2.02	−0.01
4,4',6,6',7,7'-Chloroindigo	TCE	2.09	−0.05	2.03	2.08	−0.05
5,5',6,6',7,7'-Chloroindigo	TCE	2.05	−0.05	2.00	2.07	−0.07
5,5',7,7'-Bromo-6,6'-aminoindigo	Xylene	2.31	−0.07	2.25	2.22	0.03
4,4',5,5',6,6',7,7'-Chloroindigo	EtOH	2.03	−0.05	1.98	2.05	−0.07
4,4',5,5',6,6',7,7'-Chloroindigo	TCE	2.04	−0.05	1.98	2.02	−0.04
4,4',5,5',6,6',7,7'-Chloroindigo	Xylene	2.06	−0.05	2.01	2.04	−0.03

Table 5.5 Statistical evaluation of the VG corrections (ω_{VG}) and the errors in DLPNO-STEOM 0–0 transition energies (ω_{0-0}) from Tab. 5.4 with respect to the experimental values (eV).

	ω_{VG} [†]	ω_{0-0}
Max(−)	−0.08	−0.11
ME	−0.05	−0.02
Max(+)	−0.04	0.03
SD	0.01	0.03
MAE	0.05	0.03

[†] For ω_{VG}, the smallest, average, maximum, … *shifts* are reported instead of errors.

5.7 Chemical Substitution

As a further illustration, we now turn to investigating the effect of chemical substitutions on the CPCM-corrected DLPNO-STEOM 0–0 transition energy. Fig. 5.3 shows how the excitation energy of indigo changes as various substituents are added to the core structure. Indigo in TCE is one of the few dyes for which our protocol predicts a slightly too high value. Most derivatives lie in the band stretching between perfect agreement and −0.1 eV deviation from experiment. Among these is 6,6'-dibromoindigo in TCE. It is interesting to note that the shift caused by the Br substituents is relatively small compared especially to the methoxy and Cl substituents at the same position. In fact, 6,6'-dichloroindigo is the only dye in Fig. 5.3 that is dissolved in Xylene and is shown mainly because it features the largest measured positive shift compared to indigo. The STEOM error for 6,6'-dichloroindigo is also quite large in Xylene (−0.08 eV), and it is the largest in the test set when EtOH is used as the solvent (−0.11 eV). The methoxy groups are

notable because they are responsible both for some of the largest positive and negative shifts. In the 6,6' positions, these substituents cause a measured positive shift of about 0.14 eV with respect to the experimental absorption maximum of indigo at 2.02 eV, while in the 5,5' positions they are responsible for a −0.11 eV shift. By contrast, two nitro groups in the 6,6' positions cause only a −0.07 eV shift. Nevertheless, for all these different substituents, our protocol remains within the 0.1 eV error bar and as a consequence, it is perfectly suitable for predicting the effects of chemical changes.

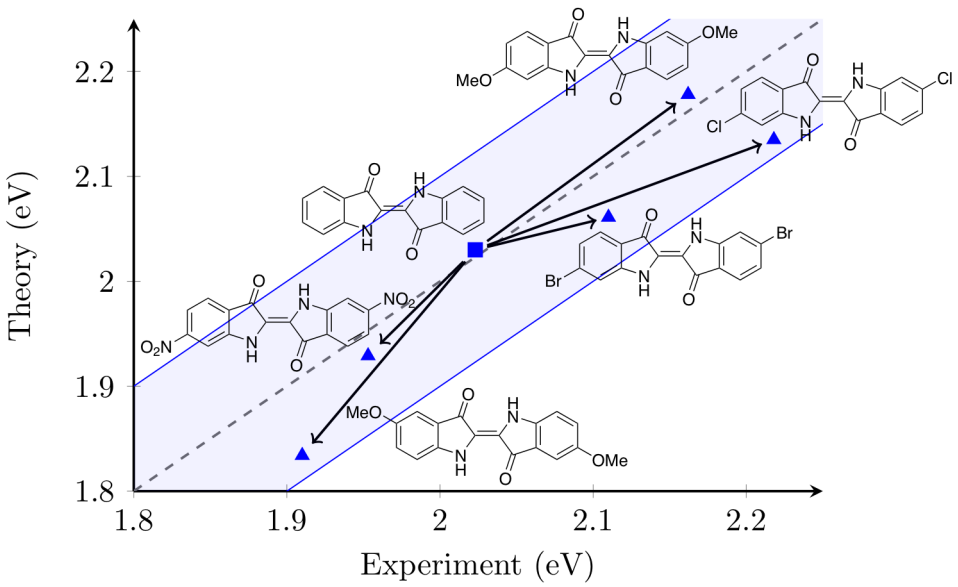

Figure 5.3 The effect of chemical substitutions on excitation energies. The dashed gray line indicates perfect agreement between theory and experiment, the blue zone signifies the region within which the agreement is within ±0.1 eV.

5.8 Absorption Spectra

Finally, the protocol under discussion allows for the calculation of entire absorption spectra. As an illustration, Fig. 5.4 shows the computed and measured spectra of indigo, 6-bromoindigo and 6,6'-dibromoindigo in DMSO. The spectra predicted by our protocol agrees remarkably well with the measured ones *without* any empirical shift. The only empirical parameter in the three computed spectra is the linewidth, which was determined as a common parameter for all three of them from the half width at half maximum of the low-energy side of the experimental curves. This was done to avoid the broadening of the spectra due to a shoulder on the high-energy side. The stick plots below the computed curves indicate the contributions of vibrational modes. The highest peak at $\sim$ 16000 cm^{-1} corresponds to the 0–0 transition, since at 0 K this is the lowest

energy transition. As the spectra in Fig. 5.4 are computed at room temperature, transitions from a vibrationally excited ground state (hot bands) can also just be observed at energies lower than the 0–0 transition. The shoulder towards higher energies ($\sim$ 18000 cm^{-1}) is mainly due to two similar fundamental modes that feature a strong symmetric stretching motion of the CC and CO bonds in the central chromophore, while the remaining system also participates in a symmetric stretching/rocking motion (the `xyz` trajectories can be obtained as described in App. B.1.6). Vibrational overtones are present in the computed spectra at even higher energies, although their intensity is extremely low. For a more detailed discussion of the vibrational progression and temperature effects, we refer to Sec. 2.7 and to App. B.1.4.

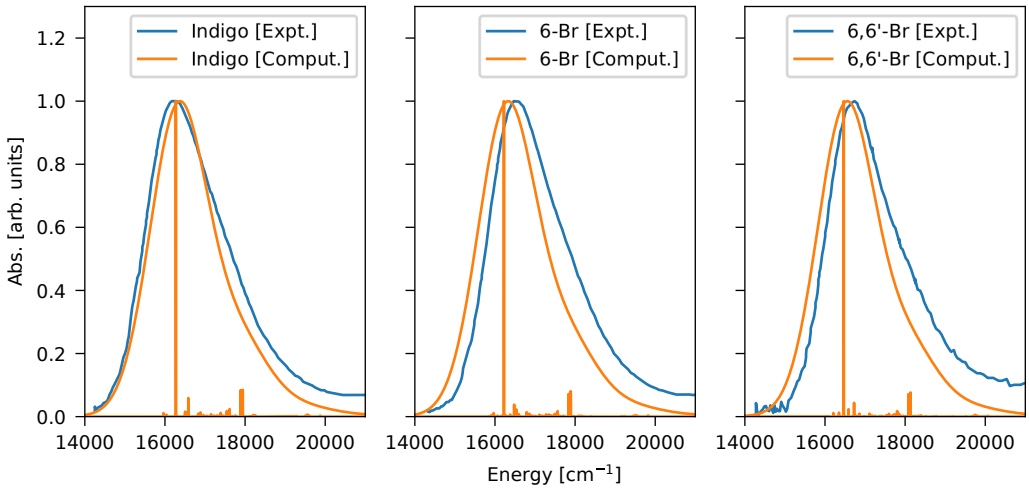

Figure 5.4 Normalized VIS absorption spectra for indigo, 6-bromoindigo, and 6,6'-dibromoindigo in DMSO. The unshifted computed spectra are presented with an inhomogeneous line width of 700 cm^{-1} as well as sticks for each vibrational mode. The experimental spectra were obtained from Ref. [500].

As a final touch, we also measured the absorption spectrum of a commercially available Tyrian purple dye in chloroform. The natural dye is a mixture of different isatinoids, indigoids, and indirubinoids [500], whose composition may naturally vary between different batches and types of shellfish used in the production process. Based on our results, the shoulder in Fig. 5.5 at about $\sim$ 16500 cm^{-1} can be assigned to the indigoids, for which we indicated the 0–0 transitions (see Fig. B.3 for measured and computed spectra of indigo in chloroform). Based on the spectral data from Ref. [500], the highest peak in Fig. 5.5 corresponds to the region where the indirubinoids absorb visible light. Note that the isatinoids do not absorb in the energy range shown in Fig. 5.5. Thus, to reproduce the full spectrum of this particular mixture, the spectra of the indirubinoids would also need to be calculated and the concentrations of the components would also need to be known. While this is out of the scope of this thesis, we feel that the accuracy of our method has been demonstrated and its practical utility is well illustrated by our contribution to the study

of indigoids in general and Tyrian purple in particular.

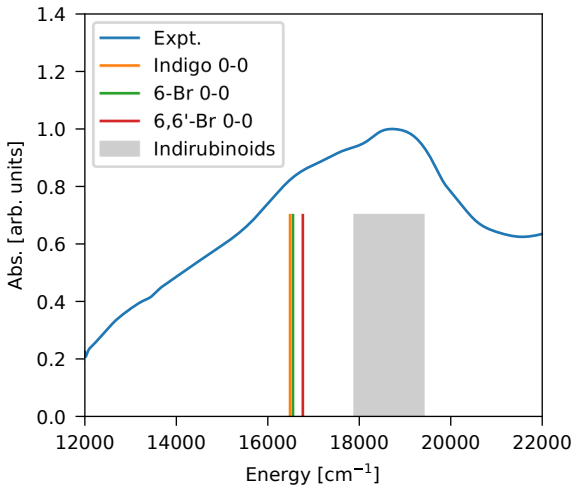

Figure 5.5 The spectrum of Tyrian purple and the 0–0 transition energies of three of its components studied here. The shaded area corresponds to the region where indirubin dyes absorb (its width chosen from the spread between the band maxima of indigorubin dyes in Ref. [500] with an additional 500 cm^{-1} in both directions).

6 Multireference Equation-of-Motion Perturbation Theory

MR-EOMPT theory, a perturbative variant of the iterative MR-EOMCC method, is presented in this chapter. To this end, we first delineate the working equations and relate the new method to other perturbative approaches. Subsequently, six benchmark sets are used to determine the accuracy and efficacy of the perturbative approximations, in the context of both the parent method, MR-EOMCC theory, and another popular perturbative approach, fic-NEVPT2 theory.

6.1 Theory

6.1.1 MR-EOMCC and its Perturbative Variant, MR-EOMPT

The electronic structure method described in this chapter is closely related to the MR-EOMCC approach in that it retains the sequential similarity transformations as well as the final, uncontracted multireference configuration interaction including singles (MRCIS) diagonalization over a compact manifold. We only recapitulate the relevant features in this section and refer the reader to the published articles [211–213, 215] for further details, especially regarding technicalities of the normal-ordered expansions. Multiple variants of the similarity transformations and the final diagonalization step have been described, but in the following, "MR-EOMCC" should be taken to mean the MR-EOM-T|T†|SXD|U-h-v method using the notation of Ref. [215]. In brief, the vertical lines indicate the sequential transformation steps, and the suffix "h-v" denotes that the final Hamiltonian is symmetrized (*vide infra*). We have already discussed the basic ideas behind MR-EOMCC in Sec. 2.6.4, and in Sec. 3 we have also discussed the connection between the many-body solution criterion that MR-EOMCC uses to the projection conditions more commonly used in MRCC methods.

The similarity transformation, as shown in Eq. (2.123), plays a vital role in MR-EOMCC theory. In fact, it uses four sequential similarity transformations, which has also been investigated in the context of internally contracted (ic-)MRCC [214]. Of course, since MR-EOMCC is a genuine MRCC method, the reference state is no longer closed-shell, but rather a (SA-)CASSCF wave function. Furthermore, the similarity-transformed Hamiltonian is expanded using Kutzelnigg and Mukherjee's generalized normal order (GNO) [219, 220, 401] in a many-body form, see

Ch. 3,

$$\begin{aligned}\bar{H} &= \mathrm{e}^{-\hat{T}}\,\hat{H}\,\mathrm{e}^{\hat{T}} \\ &= \bar{h}_0 + \sum_{pq} \bar{h}^p_q \{\hat{E}^q_p\} + \sum_{pqrs} \bar{h}^{pq}_{rs} \{\hat{E}^{rs}_{pq}\} + \ldots,\end{aligned} \tag{6.1}$$

where the cluster operator $\hat{T}$ is defined as

$$\hat{T} = \sum_{ia} t^i_a \hat{E}^a_i + \sum_{ta} t^t_a \hat{E}^a_t + \frac{1}{2}\sum_{ijab} t^{ij}_{ab} \hat{E}^{ab}_{ij} + \frac{1}{2}\sum_{itab} t^{it}_{ab} \hat{E}^{ab}_{it} + \frac{1}{2}\sum_{tuab} t^{tu}_{ab} \hat{E}^{ab}_{tu}. \tag{6.2}$$

We note that the cluster operator $\hat{T}$ is normal-ordered since the (de-)excitation orbital spaces are disjoint. Consequently, we did not use curly braces $\{\hat{E}^{\cdots}_{\cdots}\}$ around the spin-free excitation operators in Eq. (6.2) to indicate normal-ordered operators. In line with the arguments in Ch. 3, the many-body expansion in Eq. (6.1) is also truncated after the two-body term for all subsequently applied excitation classes discussed in this chapter.

In the fully iterative version, MR-EOMCC, the amplitudes for the cluster operator $\hat{T}$ would be obtained by requiring projective residual conditions for the singles and many-body residuals [210, 221] for the doubles to be fulfilled (see Ch. 3). In the MR-EOMPT method described in this chapter, most of the amplitudes are computed from many-body perturbation theory to second order (MBPT2). In the following, we will list the perturbative formulae in the relevant cases, while dedicating Sec. 6.1.2 to their theoretical justification. To begin with, the virtual-inactive block of $\hat{T}$ is obtained as

$$t^{ij}_{ab} = \frac{(ia|jb)}{\epsilon_i + \epsilon_j - (\epsilon_a + \epsilon_b)}, \tag{6.3}$$

where ϵ_p denotes the orbital energies of the inactive and virtual orbitals of the CASSCF reference.

Amplitudes involving active orbitals are computed using extended Koopmans' theorem (EKT) [501–503],

$$t^{tu}_{ab} = \sum_{\tilde{v}\tilde{w}} c^{\mathrm{IP}}_{t\tilde{v}} c^{\mathrm{IP}}_{u\tilde{w}} \frac{(\tilde{v}a|\tilde{w}b)}{\epsilon^{\mathrm{IP}}_{\tilde{v}} + \epsilon^{\mathrm{IP}}_{\tilde{w}} - (\epsilon_a + \epsilon_b)}, \tag{6.4}$$

where $c^{\mathrm{IP}}_{t\tilde{v}}$ are the EKT orbital coefficients and are the EKT orbital energies for ionization potentials (IP). We also defined the transformed integrals,

$$(\tilde{v}a|\tilde{w}b) = \sum_{tu} c^{\mathrm{IP}}_{t\tilde{v}} c^{\mathrm{IP}}_{u\tilde{w}} (ta|ub), \tag{6.5}$$

where a tilde indicates an orbital in the EKT-IP basis.

The t^{it}_{ab} amplitudes are computed analogously to the t^{tu}_{ab} amplitudes. Explicit equations can be found in App. B.2.1.

The singles amplitudes t_a^i, t_a^t are computed to second instead of first order because the Brillouin condition requires that the first order amplitudes $t_a^{i\,(1)} = 0$ [217]. Hence, the singles amplitudes are initially set to zero and the doubles estimate is performed according to Eq. (6.3) and its analogs for the t_{ab}^{it} and t_{ab}^{tu} amplitudes. Subsequently, the doubles-only state is used to compute the singles residuals,

$$r_a^i = \langle \Phi_0 | \hat{E}_a^i \bar{H} | \Phi_0 \rangle \,, \tag{6.6}$$

which are then employed to estimate the second order singles amplitudes,

$$t_a^{i\,(2)} = \frac{r_a^i}{\epsilon_i - \epsilon_a}, \quad t_a^{t\,(2)} = \sum_{\tilde{v}} c_{t\tilde{v}}^{\mathrm{IP}} \frac{r_a^{\tilde{v}}}{\epsilon_{\tilde{v}}^{\mathrm{IP}} - \epsilon_a}. \tag{6.7}$$

In Eqs. (6.7), the transformed residual matrix elements $r_a^{\tilde{v}}$ are defined analogously to the transformed integrals from Eq. (6.5), again in the EKT-IP basis.

At this point, the Hamiltonian $\bar{H}$ is transformed with the de-excitation operator $\hat{T}^\dagger$ in complete analogy to Eq. (6.1) to give the Hamiltonian $\tilde{H}$. This non-iterative step approximates the de-excitation amplitudes by setting

$$t_i^a = t_a^i, \quad t_t^a = t_a^t, \tag{6.8}$$

$$t_{ij}^{ab} = t_{ab}^{ij}, \quad t_{it}^{ab} = t_{ab}^{it}, \quad t_{tu}^{ab} = t_{ab}^{tu}. \tag{6.9}$$

The $\hat{T}^\dagger$ transformation is included to make the transformed Hamiltonian more Hermitian. For more details, we refer to Ref. [215].

We then proceed with a third similarity transformation on top of the previous transformed Hamiltonian,

$$\begin{aligned} \bar{F} &= \{\mathrm{e}^{\hat{S}+\hat{X}+\hat{D}}\}^{-1} \tilde{H}_2 \{\mathrm{e}^{\hat{S}+\hat{X}+\hat{D}}\} \\ &= \bar{f}_0 + \bar{f}_q^p \{\hat{E}_p^q\} + \bar{f}_{rs}^{pq} \{\hat{E}_{pq}^{rs}\} + \ldots, \end{aligned} \tag{6.10}$$

where $\tilde{H}_2$ is the Hamiltonian from after the $\hat{T}$, $\hat{T}^\dagger$ transformations truncated at two-body terms. For details on how the inverse of the normal-ordered exponential $\{\mathrm{e}^{\hat{S}+\hat{X}+\hat{D}}\}^{-1}$ is computed we refer to Ref. [462]. The individual cluster operators used in this transformation are defined as

$$\begin{aligned} \hat{S} &= s_{at}^{ij} \{\hat{E}_{ij}^{at}\}, \\ \hat{X} &= x_{au}^{ti} \{\hat{E}_{ti}^{au}\}, \\ \hat{D} &= d_{au}^{it} \{\hat{E}_{it}^{au}\}. \end{aligned} \tag{6.11}$$

The amplitudes of excitations into active orbitals are again obtained perturbatively in MR-EOMPT using the electron attachment (EA) results from EKT for the orbital energies, e.g.,

$$s_{at}^{ij} = \sum_{\tilde{v}} c_{t\tilde{v}}^{\mathrm{EA}} \frac{(ia|j\tilde{v})}{\epsilon_i + \epsilon_j - (\epsilon_a + \epsilon_{\tilde{v}}^{\mathrm{EA}})} \tag{6.12}$$

with the corresponding integral being defined here in the EKT-EA basis,

$$(ia|j\tilde{v}) = \sum_t c^{\mathrm{EA}}_{t\tilde{v}} (ia|jt). \tag{6.13}$$

The guess amplitudes for $\hat{X}$ and $\hat{D}$ are computed equivalently and may be found in App. B.2.1.

The remaining amplitudes $\hat{U}$ of the last similarity transformation are computed exactly as in MR-EOMCC theory. By replacing the first three similarity transformations of the $\hat{T}$, $\hat{T}^\dagger$ and $\hat{S} + \hat{X} + \hat{D}$ blocks with perturbatively estimated amplitudes, the MR-EOMPT method will, naturally, be faster than its parent method, MR-EOMCC, although at the cost of accuracy. This tradeoff has already been partially attempted with an iterative orbital selection scheme, which we mention here since it used the same definition for the non-iterative amplitude estimates therein [217]. One of our main foci here will be to investigate the time-accuracy tradeoff. Furthermore, it will be easier to perturbatively obtain the amplitudes for the aforementioned similarity transformations since the iterative MR-EOMCC method may face convergence difficulties in the case of nearly singular amplitudes [209, 210].

The last similarity transformation is described by

$$\begin{aligned} \bar{G} &= \mathrm{e}^{-\hat{U}} \bar{F}_2 \mathrm{e}^{\hat{U}} \\ &= \bar{g}_0 + \bar{g}^p_q \{\hat{E}^q_p\} + \bar{g}^{pq}_{rs} \{\hat{E}^{rs}_{pq}\} + \ldots, \end{aligned} \tag{6.14}$$

with the two-body operator

$$\hat{U} = \frac{1}{2} u^{ij}_{tu} \{\hat{E}^{tu}_{ij}\}. \tag{6.15}$$

Amplitudes for this transformation are not computed perturbatively, but are solved for iteratively by requiring that the corresponding elements of the many-body expanded Hamiltonian $\bar{G}$, $\bar{g}^{ij}_{tu}$, must be zero as in MR-EOMCC theory,

$$r^{ij}_{tu} = \bar{g}^{ij}_{tu} \stackrel{!}{=} 0. \tag{6.16}$$

Although these amplitudes could also be estimated from MBPT2, they are iterated because they are never the rate-limiting step of the entire calculation. Furthermore, we normally observe fast convergence within up to five iterations.

After symmetrizing the similarity-transformed Hamiltonian $\bar{G}$ according to Ref. [215], the only remaining step is the final diagonalization of $\bar{G}$ over a suitable manifold. This step will then give access to the total energies of the desired states as well as the transition energies between them. The optimal setup of this final diagonalization step has been investigated several times by Nooijen and co-workers [212, 213, 215, 216]. In this study, we have chosen to diagonalize the transformed Hamiltonian $\bar{G}$ over the reference CAS space including single hole and particle excitations with respect to the reference space. In the case of MR-EOMPT theory, there is a strong argument to use the smallest reasonable diagonalization space to minimize the computational costs. Extending the final diagonalization manifold is usually very costly with limited benefits for accuracy [215].

6.1.2 Discussion of the Perturbative Approximations

Since we will relate MR-EOMPT theory to other perturbative approaches in the next section, we will briefly comment on the perturbative concept of MR-EOMPT theory in this section. MR-EOMPT theory is the lowest-order, iterative approximation to MR-EOMCC theory, as further discussed below. As such, we do not start from the partitioning of a Hamiltonian into a zeroth-order and perturbative part, i.e., $\hat{H} = \hat{H}_0 + \hat{V}$, but rather consider MR-EOMPT through an iterative expansion.

We can define such an iterative expansion that applies to essentially all iterative methods since the update of the amplitudes, collectively denoted as $\boldsymbol{T}$, is generally based on the previous amplitudes along with an incremental update,

$$\boldsymbol{T}^{(n+1)} = \boldsymbol{T}^{(n)} + \boldsymbol{f}(\boldsymbol{r}(\boldsymbol{T}^{(n)})), \tag{6.17}$$

for the nth order of iteration. The residuals $\boldsymbol{r}$ are a function of the current set of amplitudes, and the additive update step is another function $\boldsymbol{f}$ of them, in turn. The update function most often is a denominator, $\boldsymbol{f}(\boldsymbol{x}) = \Delta^{-1}\boldsymbol{x}$, which is chosen such that the iterative sequence converges quickly, i.e., the residuals tend to zero, $\boldsymbol{r}(\boldsymbol{T}^{(n\to\infty)}) = \boldsymbol{0}$. In MR-EOMPT theory, we use the iterative sequence to define orders of perturbation, aborting after a single step for the doubles (first order) and two steps (second order) for the singles amplitudes (see Sec. 6.1.1). To ensure accurate results, we need to choose the update function, denominators in our case, appropriately.

The scheme for choosing the denominators has already delineated by Nooijen and co-workers [209]. Accordingly, we will only outline the underlying principle here, and refer to the publication for further details. The denominators for both electron-attached and ionized states, $\hat{p}^\dagger |0\rangle$ and $\hat{q} |0\rangle$, respectively, are obtained from extended Koopmans states [501–505] chosen such that the coupling between active and virtual or inactive orbitals, respectively, is removed. This choice leads to large denominators, as the difference in orbital energies between the electron-attached and ionized states *for the same active label* is on the order of multiple eV. However, this necessitates that we transform the residuals to the new one-particle basis defined by the extended Koopmans states to compute the denominator update, see, e.g., Eq. (6.7).

We conclude by noting that deriving a perturbation $\hat{V}$ is possible, albeit complicated. If we were to derive such a perturbation $\hat{V}$, we would start from the initial guess, $\boldsymbol{T}^{(0)} = \boldsymbol{0}$, and use the fact that the full Hamiltonian enters MR-EOMPT theory through the MR-EOMCC working equations. This happens at first order since we abort after just one step. The MR-EOMCC working equations effectively couple active space determinants to excitations out of the active space to achieve a decoupling (see Secs. 2.6.4 and 6.1.3). A perturbation operator that has the same effect can naturally be found, but is associated with significant theoretical challenges, making the iterative view our preferred choice for presenting MR-EOMPT theory.

6.1.3 Relations to Other MRPT approaches

It seems appropriate to elaborate on the differences between the fic-NEVPT2 approach as an example of a state-specific, *diagonalize-then-perturb* MRPT and the MR-EOMPT method, which is a state-universal, *transform-then-diagonalize* approach. Both NEVPT2 and MR-EOMPT rely on the solution of the CASSCF problem (cf. Sec. 2.3.2),

$$\hat{H}_{\mathrm{CAS}} = \hat{P}_{\mathrm{CAS}} \hat{H} \hat{P}_{\mathrm{CAS}}, \quad \hat{P}_{\mathrm{CAS}} = \sum_{I} |\Phi_I\rangle \langle \Phi_I| , \tag{6.18}$$

$$\hat{H}_{\mathrm{CAS}} |\Psi_k\rangle = E_k |\Psi_k\rangle = E_k \sum_{I} c_{I,k} |\Phi_I\rangle , \tag{6.19}$$

where $|\Phi_I\rangle$ are individual CSFs from the CASCI space combined with coefficients $c_{I,k}$ into the solution for the kth state. For fic-NEVPT2, these solutions are then taken to define the internally contracted perturber functions for MRPT by applying the excitation operators from the perturbation operator $\hat{V}$ to the CASSCF solution. For example, the perturber function space for a two-hole, one-particle situation is given by [169, 171]

$$\overline{S}^{+1}_{ija} = \hat{E}^a_i \hat{E}^t_j |\Psi_m\rangle \ , \quad \dim(\overline{S}^{+1}_{ija}) = n_a. \tag{6.20}$$

Here, n_a denotes the number of active orbitals. In a similar fashion for all the remaining seven excitation classes, these perturber functions are state-specific in that they already explicitly contain the state of interest. The perturber functions then enter the expressions for the energy correction to second order [171]. Another consequence of the state-specific, internally contracted perturbers is that they ensue the presence of four-body (active) densities in the working equations for the $\overline{S}^{(-1)}_i$ and $\overline{S}^{(-1)}_a$ spaces [171].

For MR-EOMPT, an equation like (6.20) is conspicuously missing. Instead, a similarity-transformed Hamiltonian is built in the CASCI model space using perturbatively estimated amplitudes (*vide supra*). The working equations depend on the reference state in two ways: First, the one- and two-body reduced density matrices (RDMs) enter through the EKT orbital energies, e.g., Eq. (6.4). Second, the one-body RDM is also required to construct the similarity-transformed Hamiltonians $\bar{H}$, $\bar{F}$, and $\bar{G}$. Hence, this method is state-universal in character for all states in the CASCI space, as opposed to the state-specific fic-NEVPT2 method. The purpose of the similarity transformations is to decouple the model space from the outer space, as thoroughly discussed in both STEOM [10–12] and MR-EOMCC [211–213] theory. Eventually, the MRCIS diagonalization gives access to the final states of the transformed Hamiltonian. We note that the diagonalization is not only over the model CASCI space, but rather over the CASCI space and all single-hole and single-particle excitations relative to it, which is the first-order interacting subspace based on the (large elements of the) transformed Hamiltonian. The rationale behind this choice is that any residual coupling to the outer space should be captured in such a fashion. It is, however, not intended to significantly change the character of the model space states, which is why MR-EOMCC calculations generally require that the weight of the CSFs from the CASCI space $|\Phi_I\rangle$ exceed 90% in the final states [217].

6.2 Computational Details

All calculations presented in this chapter use a (SA-)CASSCF reference wave function and were computed using a development version of the ORCA [286–288] program package. The detailed settings regarding the number of active electrons, number of active orbitals, or which roots were included in the reference state can be found in Tab. B.4 of the Appendix. We normally used a "standard choice" of the active space, meaning that we included all (anti-)bonding partners, complete π-systems, and other orbitals involved in excited state transitions, e.g., the n-orbitals.

When comparing the results of the MR-EOMCC and MR-EOMPT methods to pure CASSCF or fic-NEVPT2 calculations (referred to in the following as CASSCF and NEVPT2, respectively), we note that the MR-EOM methods do generally *not* use the same reference wave function as the CASSCF/NEVPT2 calculations. For the CASSCF and NEVPT2 results, we often simply included all roots of interest with equal weights in the CASSCF reference, with possible splitting into separate blocks of different multiplicity and irreducible representation (irrep).

The setup of the MR-EOMCC and MR-EOMPT calculations may be somewhat counterintuitive in that the aim of their SA-CASSCF reference is to decouple the active space by virtue of the similarity transformations; and then obtain the states of interest from the final MRCIS step, while also taking any residual coupling into account. In more practical terms, our general procedure was as follows. First, an initial CASCI or CASSCF calculation over a number ($\lesssim 5$) of low-lying singlet and triplet roots each is done to obtain a "spectrum" of states from which the MR-EOM reference will be selected. In the second step, we choose all roots beneath a threshold of a few eV ($\lesssim 3$ eV), while taking care that the selection does not break wave function symmetry. This will be the reference state used in the MR-EOMCC and -PT calculations. In case of convergence difficulties in the MR-EOMCC method or final roots with a too low contribution of the CASCI space (*vide supra*), the reference state may be modified to achieve convergence. While a detailed study regarding root selection is out of scope of this study, the MR-EOMCC method (as well as the MR-EOMPT variant in our testing) has been shown to be relatively invariant with respect to the choice of the reference state [215]. The exact setup is detailed in Tab. B.4.

When it comes to comparing the MR-EOMCC and MR-EOMPT results, we emphasize that the reference as well as the calculation settings were strictly the same in both cases, the only difference being that perturbative amplitudes were used for the MR-EOMPT calculations. The convergence threshold for the iterative amplitude equations was always set to a maximum absolute residual of 10^{-6}.

6.3 Diatomic Systems

The first test set that we use to gauge the accuracy of the MR-EOMPT method is a set of small, diatomic molecules first published as a means to assess the contraction error of a fic-MRCI implementation [176]. These diatomic molecules are entirely from the second period and are

computed in the def2-SVP [463] basis set. Hence, we are able to obtain near-full CI (FCI) results from the iterative configuration expansion (ICE-)CI method [438, 439], a selected CI approach in the vein of CIPSI (configuration interaction by perturbation with multiconfigurational zeroth-order wavefunction selected by iterative process) [430], by including all electrons and orbitals in the active space (ICE-FCI). The iterative cutoffs that were used in the ICE-FCI reference calculations ($T_{\mathrm{gen}} = 10^{-5}$, $T_{\mathrm{var}} = 10^{-12}$) are expected to yield average deviations of less than 0.1 mE_{h} when compared to the total energy of a corresponding FCI calculation [439]. Consequently, the aim for this test set is to demonstrate how closely the MR-EOMPT method can reproduce nearly exact, FCI-quality transition energies.

In all calculations, a full valence CASSCF reference was chosen. The primary reason for such a comparatively large active space is that some transitions, for example the $X\,^2\Sigma^+ \rightarrow 2\,^2\Sigma^+$ transition on both CN and CO^+, describe a $2\sigma^* \rightarrow 3\sigma$ excitation out of the antibonding orbital formed from the atomic 2*s* orbitals, and neither of these methods can describe excitations not contained in the active space. For CASSCF, it is immediately clear that states outside of the active space cannot be described, and therefore also state-specific fic-NEVPT2 cannot perturbatively improve upon such state. In the case of MR-EOMCC and MR-EOMPT, the similarity transformations also only account for the states contained in the active space, and the final MRCIS diagonalization evaluates the residual coupling to the outside space (see Sec. 6.1.3). The detailed setup is summarized in Tab. B.4.

The results of the CASSCF, NEVPT2, MR-EOMPT, and MR-EOMCC calculations are presented in Table 6.1. Due to the large, full-valence active space, all methods report rather accurate results, with the largest deviation in the vertical transition energy (VTE) being 0.48 eV by the CASSCF method. As expected, the CASSCF method is outperformed by the remaining methods since it does not account for most of dynamic electron correlation, even despite the large active space. We wish to emphasize that these results may be of limited applicability to larger calculations, since the full valence active space comprises an exceptionally large fraction of the total number of orbitals in the systems, which therefore captures a larger than usual amount of dynamic correlation in the reference.

The statistical analysis of the errors reported in Table 6.1 shows that the correlated methods NEVPT2, MR-EOMPT, and MR-EOMCC perform similarly across this test set, judging by the maximum absolute error (MAE) or the root-mean-square deviation (RMSD). Out of all methods, fic-NEVPT2 reports the lowest standard deviation. In combination with an average deviation of 0.05 eV, we can conclude that it yields the most consistent results, despite the MR-EOMCC method having an average deviation of only 0.01 eV from the ICE-FCI reference. Even so, the RMSD for all methods beyond CASSCF is below 0.1 eV, which is generally considered to be in excellent agreement.

We further report the total energies for all the states used in this section in Tab. B.3, as well as a statistical analysis of the deviations compared to the ICE-FCI reference data in Table 6.2. While the MR-EOM methods are primarily targeted at computing transition energies on a single

Table 6.1 Errors in vertical transition energies (VTEs) with respect to ICE-FCI results (reference values) on diatomic systems from the second period. The VTEs are given relative to the ground state (GS) in parentheses. All values given in eV.

			ΔVTE			
		VTE			*MR-EOM*	
System (GS)	**Exc. State**	*ICE-FCI*	*CASSCF*	*NEVPT2*	*-PT*	*-CC*
CH ($^2\Pi$)	$^2\Delta$	3.04	0.07	−0.01	−0.02	0.00
	$^2\Sigma^-$	3.40	−0.10	−0.02	−0.02	−0.04
	$^2\Sigma^+$	4.15	−0.04	0.04	0.10	0.14
CN ($^2\Sigma^+$)	$^2\Pi$	1.46	0.09	0.02	0.06	0.01
	$2\,^2\Sigma^+$	3.03	0.07	0.04	0.05	0.02
CO ($^1\Sigma^+$)	1Π	8.78	0.47	0.08	0.01	−0.12
	$^1\Sigma^-$	10.3	0.15	0.07	0.06	0.00
	$^1\Delta$	10.42	0.09	0.07	0.05	−0.02
	$^3\Pi$	6.33	0.20	0.03	−0.01	−0.08
	$^3\Sigma^+$	8.62	0.12	0.09	0.15	0.12
	$^3\Delta$	9.55	0.14	0.08	0.11	0.06
	$^3\Sigma^-$	10.04	0.04	0.08	0.10	0.04
CO^+ ($^2\Sigma^+$)	$^2\Pi$	3.59	0.24	0.05	0.01	0.01
	$2\,^2\Sigma^+$	5.83	0.31	0.08	0.02	0.02
N_2 ($^1\Sigma^+$)	$^1\Pi$	9.65	0.44	0.07	0.06	0.06
	$^1\Sigma^-$	10.41	0.47	0.08	0.01	0.00
	$^1\Delta$	10.8	0.39	0.09	0.06	0.06
	$^3\Sigma^+$	7.94	0.16	0.06	0.09	0.06
	$^3\Pi$	8.26	0.48	0.04	0.08	0.10
	$^3\Delta$	9.26	0.32	0.07	0.05	0.04
	$^3\Sigma^-$	10.06	0.21	0.09	0.11	0.11
O_2 ($^3\Sigma^+$)	$^1\Delta$	1.07	−0.12	−0.04	−0.04	−0.06
	$^1\Sigma^+$	1.73	−0.26	−0.05	−0.03	−0.06
	$^3\Delta$	6.34	−0.19	0.00	−0.01	−0.06
OH ($^2\Pi$)	$^2\Sigma^+$	4.29	0.34	0.02	−0.03	−0.09
		Statistical Evaluation of Errors				
	Min.		−0.26	−0.05	−0.04	−0.12
	Avg.		0.16	0.05	0.04	0.01
	Max.		0.48	0.09	0.15	0.14
	Std.		0.20	0.04	0.05	0.07
	MAE		0.22	0.05	0.05	0.05
	RMSD		0.26	0.06	0.07	0.07

system, the total energies may, e.g., also be used to compute energy differences between different geometries (see also Sec. 6.6). Unsurprisingly, the iterative MR-EOMCC method is the closest to the ICE-FCI reference, with an RMSD of 0.15 eV. Next are the perturbative methods MR-EOMPT and NEVPT2 with RMSDs of 0.46 eV and 0.77 eV, respectively. Thus, MR-EOMPT outperforms NEVPT2 in this regard, despite the RMSD across the transition energies from Table 6.1 being slightly worse than that of NEVPT2. As usual, the CASSCF results are far off since even the large full-valence active spaces used in this section are not enough to recover a substantial amount of dynamic electron correlation.

Table 6.2 Statistical evaluation of the differences of the total energies compared to the ICE-FCI reference values (given in eV). The total energies of the 32 states can be found in Tab. B.3.

			MR-EOM	
	CASSCF	**NEVPT2**	**-PT**	**-CC**
Min.	2.03	0.43	0.09	−0.03
Avg.	5.04	0.76	0.43	0.13
Max.	7.56	1.05	0.64	0.32
Std.	1.44	0.16	0.16	0.07
MAE	5.04	0.76	0.43	0.13
RMSD	5.24	0.77	0.46	0.15

6.4 Thiel Test Set

Continuing to focus on excitation energies, Thiel's extensive test set [313, 314, 506] lends itself to further benchmarking. This test set contains highly accurate reference data in the form of CC3 results with a high singles-excitation dominated character ($\%T_1 > 90\%$) [313]. In the case of the Thiel benchmark set, multireference methods are not *per se* required since the included molecules are of closed-shell type, at least in their respective ground states. Nevertheless, the reported CC3 values may be used as reference values against which the MR-EOMPT method can be benchmarked, provided that the same calculation settings and basis set are used. Note that CC3 is known to have MAEs of well below 0.1 eV for singles-dominated transitions [313].

The reason why this test set is suitable for benchmarking the MR-EOMPT method is that it contains reference data for accurate excitation energies, which is also what its parent method, MR-EOMCC, was designed to calculate. Thiel's test set, or a subset thereof, has already been used to benchmark other variants of the MR-EOMCC approach [212, 215]. We note that our choice of the zeroth-order CASSCF reference for the MR-EOMCC and MR-EOMPT calculations here is not necessarily identical to the references in the published literature. The choice of the reference was found to have only a minor impact on the calculations [215], which we were able to

confirm in our testing with the MR-EOMPT and MR-EOMCC methods. The reference state(s) used were obtained through the procedure outlined in Sec. 6.2 and can be found in Tab. B.4.

Our results are summarized in Tab. 6.3, where we report the reference values (CC3) from Ref. [313] along with the deviation of the CASSCF, fic-NEVPT2, MR-EOMPT, and MR-EOMCC methods from our calculations. In line with expectations, the CASSCF results again fail to reproduce the reference values accurately and consistently. If we turn our attention to the perturbational approaches, we see that the description is much improved over the CASSCF base results. Furthermore, extreme outliers are rare, although absolute deviations of up to 0.68 eV can be found for the MR-EOMPT method. The fully iterative results from MR-EOMCC theory are closer to the reference values than either perturbative method, as expected from the results reported in Ref. [215].

The statistical summary of the results on the Thiel benchmark set is presented in Tab. 6.4. As in the previous publications [215, 313], we compute the statistics over singlets and triplets separately to see any systematic error in either of the subsets.

The statistical analysis shows that the MR-EOMPT method clearly performs better on the subset of transitions to triplet excited states, compared to singlet excited states. In the case of a triplet final state, the RMSD across all transitions is merely 0.11 eV and can be considered in excellent agreement with the reference values, whereas the RMSD for the singlet transitions is about four times as large at 0.43 eV. This contrasts with the results from the NEVPT2 method, which reports almost equal RMSDs for the singlet and triplet transitions. The same effect is also present in the MR-EOMCC results, albeit much less pronounced at RMSDs of 0.18 eV and 0.07 eV for the singlet and triplet transitions, respectively.

These results may seem surprising initially when they are not carefully compared to the correct variant of the MR-EOMCC approach in the published literature. The closest variant from Ref. [215] is called "MR-EOM-T|T†|SXD-h-v," which is implemented in the ORCA program package [287] along with the variant used here, which could be analogously called "MR-EOM-T|T†|SXD|U-h-v" (see Sec. 6.1.1). The published MR-EOM-T|T†|SXD-h-v results lack the final $\hat{U}$ transformation and therefore use a larger MRCI diagonalization space that also includes 2h excitations from the CASCI space. Nevertheless, the MR-EOM-T|T†|SXD-h-v results are comparable to our own MR-EOMCC results. Furthermore, Ref. [215] shows that especially the $\hat{T}^\dagger$ transformation step deteriorates the quality of the singlet transitions, while it has negligible impact on the triplet transitions. For the MR-EOMPT method, this effect is amplified, making these results less robust compared to the iterative MR-EOMCC method. In the future, other variants of the MR-EOMPT approach, e.g., without the $\hat{T}^\dagger$ transformation, can be explored. We further remark that a previous study [212] also found that the MR-EOMCC results tend to perform favorably for the triplet over singlet transitions, at least when compared to CC3 reference values. Even with this caveat, the MR-EOMPT method compares well with the reference data for both singlet and triplet states.

Table 6.3 Errors in transition energies (eV) with respect to CC3 results [313] on molecules from the Thiel test set, TZVP [507] basis set.

			ΔVTE			
		VTE			*MR-EOM*	
Molecule	**Transition**	*CC3* (Ref.)	*CASSCF*	*NEVPT2*	*-PT*	*-CC*
Ethene	$1\,^1A_g \rightarrow 1\,^1B_{1u}$	8.37	1.07	0.27	0.41	0.09
	$1\,^1A_g \rightarrow 1\,^3B_{1u}$	4.48	−0.20	0.13	−0.05	−0.01
Butadiene	$1\,^1A_g \rightarrow 1\,^1B_u$	6.58	1.52	−0.31	0.61	0.24
	$1\,^1A_g \rightarrow 2\,^1A_g$	6.77	−0.18	0.05	−0.10	−0.17
	$1\,^1A_g \rightarrow 1\,^3B_u$	3.32	−0.14	0.12	0.02	0.01
	$1\,^1A_g \rightarrow 1\,^3A_g$	5.17	−0.24	0.16	−0.02	−0.02
Cyclopropene	$1\,^1A_1 \rightarrow 1\,^1B_1$	6.90	0.39	−0.06	0.01	−0.09
	$1\,^1A_1 \rightarrow 1\,^1B_2$	7.10	1.52	−0.03	0.58	0.25
	$1\,^1A_1 \rightarrow 1\,^3B_2$	4.34	−0.14	0.21	−0.03	−0.05
	$1\,^1A_1 \rightarrow 1\,^3B_1$	6.62	0.22	−0.05	0.00	−0.07
Benzene	$1\,^1A_{1g} \rightarrow 1\,1B_{2u}$	5.07	−0.11	0.15	−0.02	−0.03
	$1\,^1A_{1g} \rightarrow 1\,^1B_{1u}$	6.68	1.28	−0.26	0.40	0.13
	$1\,^1A_{1g} \rightarrow 1\,^1E_{1u}$	7.45	1.92	−0.40	0.68	0.23
	$1\,^1A_{1g} \rightarrow 2\,^1E_{2g}$	8.43	−0.28	−0.03	−0.13	−0.27
	$1\,^1A_{1g} \rightarrow 1\,^3B_{1u}$	4.12	−0.27	0.20	−0.03	−0.02
	$1\,^1A_{1g} \rightarrow 1\,^3E_{1u}$	4.90	0.11	0.06	0.05	0.00
	$1\,^1A_{1g} \rightarrow 1\,^3B_{2u}$	6.04	1.19	−0.37	0.36	0.09
	$1\,^1A_{1g} \rightarrow 1\,^3E_{2g}$	7.49	−0.32	0.11	−0.04	−0.17
Pyrrole	$1\,^1A_1 \rightarrow 2\,^1A_1$	6.40	0.05	0.16	0.10	−0.01
	$1\,^1A_1 \rightarrow 1\,^1B_2$	6.71	0.91	−0.11	0.46	0.16
	$1\,^1A_1 \rightarrow 3\,^1A_1$	8.17	1.35	0.02	0.67	0.19
	$1\,^1A_1 \rightarrow 1\,^3B_2$	4.48	−0.44	0.27	−0.02	−0.03
	$1\,^1A_1 \rightarrow 1\,^3A_1$	5.51	−0.15	0.20	0.06	−0.03

Table 6.4 Statistical evaluation of the errors on the subset of the Thiel benchmark reported in Tab. 6.3. All values are given in eV.

			MR-EOM	
	CASSCF	**fic-NEVPT2**	**-PT**	**-CC**
		ΔVTE, Singlets		
Min.	−0.28	−0.40	−0.13	−0.27
Avg.	0.79	−0.05	0.31	0.06
Max.	1.92	0.27	0.68	0.25
Std.	0.74	0.19	0.30	0.17
MAE	0.88	0.16	0.35	0.16
RMSD	1.08	0.20	0.43	0.18
		ΔVTE, Triplets		
Min.	−0.44	−0.37	−0.05	−0.17
Avg.	−0.03	0.09	0.03	−0.03
Max.	1.19	0.27	0.36	0.09
Std.	0.42	0.17	0.11	0.06
MAE	0.31	0.17	0.06	0.05
RMSD	0.43	0.19	0.11	0.07

6.5 LiF Avoided Crossing

In this section, we present how MR-EOMPT fares against the other multireference methods in situations with weakly avoided crossings, such as in the dissociation of the LiF system. In this system, two potential energy surfaces cross at about $d_{\mathrm{LiF}} \approx 13\ a_0$. Close to its equilibrium geometry, the system can be best described as $\mathrm{Li}^+\mathrm{F}^-$, i.e., it is an ionic state. As the system is pulled apart, the electrostatic penalty of charge separation continues to grow, thus leading to a lower energy state of neutral Li and F atoms.

Weakly avoided crossings can be difficult to describe. For example, it is well known [330] that state-specific NEVPT2 theories lead to unphysical double crossings for the LiF system when employing a minimal CAS(2,2) reference wave function. In this case, the underlying reference (which is diagonalized before perturbational corrections are considered) is too far off the correct description of the electronic states such that the perturbational description cannot recover the correct behavior. The situation changes when we switch to multistate methods since the final diagonalization, which results in the final states of the system, already includes the perturbational description. In contrast to the previous section, where we compared the MR-EOMPT results to state-specific fic-NEVPT2, we here compare against (strongly contracted) QD-NEVPT2 [327], partly for the known failure of state-specific fic-NEVPT2 to describe the LiF dissociation and further because this system was also extensively discussed in Ref. [327].

The potential energy surfaces for MR-EOMCC, MR-EOMPT, QD-NEVPT2 (in the strongly contracted variant) and the reference, ICE-FCI, are shown in Fig. 6.1. The underlying CAS(2,2) zeroth-order wave function, which included the lithium $2s$ and fluorine $2p_z$ orbitals, was shared among all correlation methods such that any difference must come from the correlation treatment. Furthermore, the same basis set was employed throughout, with cc-pVDZ on the lithium atom and aug-cc-pVDZ [508] on the fluorine atom to account for its anionic nature in one of the potential energy surfaces.

All presented methods accurately predict a single, weakly avoided crossing at about $d_{\mathrm{LiF}} \approx 13\ a_0$ and correctly predict the shape of the PES, compared to the ICE-FCI reference. From Fig. 6.1, it is obvious that the QD-NEVPT2 results are closest to the ICE-FCI reference in terms of total energy, apart from an artifact at about $d_{\mathrm{LiF}} \approx 8\ a_0$. This artifact is located where state-specific fic-NEVPT2 would exhibit its first, spurious crossing, thus hinting that the QD-NEVPT2 approach also has difficulties in correcting the deficiencies in the underlying zeroth-order reference wave function. In contrast, MR-EOMCC and MR-EOMPT do not have this issue and are both highly parallel to the reference (Tab. 6.5). In fact, out of all methods, MR-EOMCC is most parallel to the ICE-FCI reference, with the non-parallelity error (NPE) [509] being 1.289 mE_h on the neutral state and even better on the ionic state.

A detailed look at the avoided crossing (inset of Fig. 6.1) shows that both MR-EOMCC and MR-EOMPT accurately predict the location of the crossing, whereas QD-NEVPT2 slightly underestimates its location at $d_{\mathrm{LiF}} \approx 12.5\ a_0$. We further realize that the energy difference between both ${}^1\Sigma^+$ states at any given separation d_{LiF} is overestimated for separations larger than the

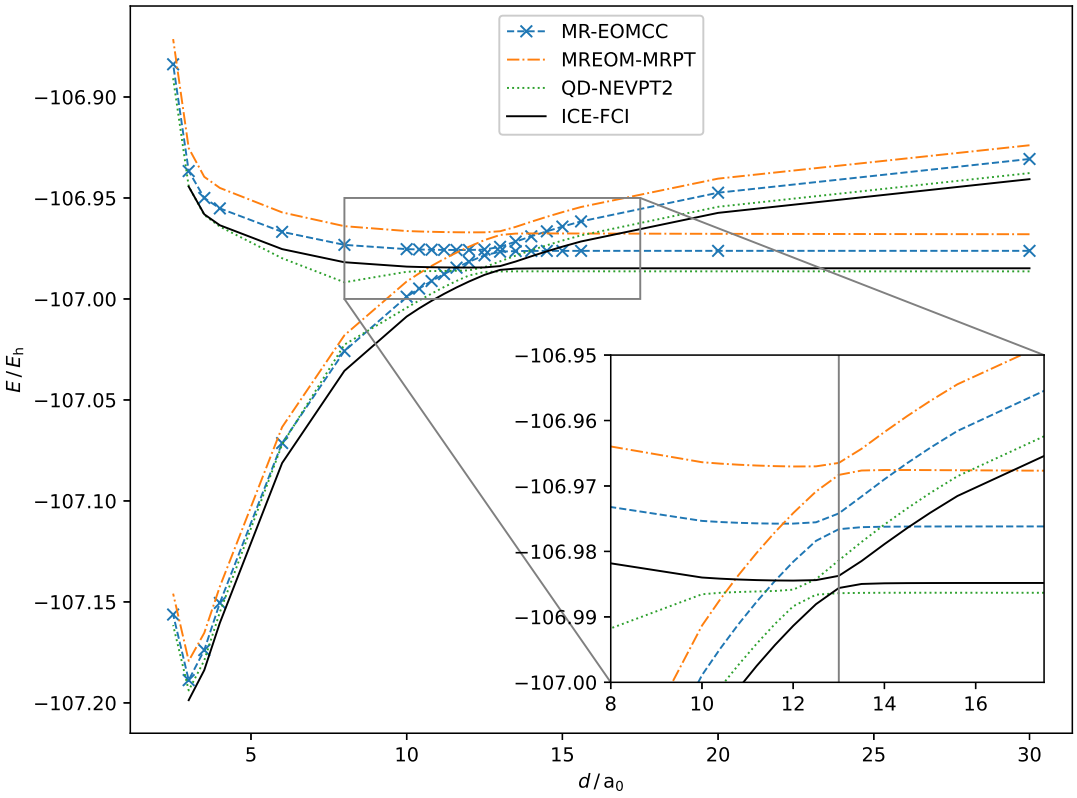

Figure 6.1 Neutral and ionic potential energy surfaces of the two lowest $^1\Sigma^+$ states of LiF for several computational methods. The ICE-FCI solution can be considered nearly exact, and thus the reference, within the given basis set and settings. The inset is a magnification around the region of the avoided crossing, with a vertical gray line indicating its position for reference.

Table 6.5 Non-parallelity errors across the PESs for the neutral and ionic states of $^1\Sigma^+$ LiF. All values are reported in mE_h.

State	QD-NEVPT2	MR-EOMPT	MR-EOMCC
Ionic	10.618	2.647	0.463
Neutral	10.051	2.213	1.289

avoided crossing in the QD-NEVPT2 results and underestimated for smaller separations. A statistical evaluation of the transition energies from the neutral to the ionic state, computed at every point of the PES scan, is shown in Tab. 6.6. It is noteworthy that due to fortuitous error cancellation, MR-EOMPT performs even better than MR-EOMCC, which slightly underestimates the transitions. QD-NEVPT2, also partly due to the artifact at $d_{\mathrm{LiF}} \approx 8\ a_0$, underestimates the transitions and has the largest deviation.

Table 6.6 Statistical evaluation of the errors in the transition energies from the neutral to the ionic state on the LiF PES. All values are reported in mE_{h}.

	QD-NEVPT2	**MR-EOMPT**	**MR-EOMCC**
Min.	−22.916	−0.364	−2.217
Avg.	−6.285	0.178	−1.264
Max.	−3.152	0.501	−0.548
Std.	4.343	0.160	0.329
RMSD	7.640	0.239	1.306

6.6 CH_2 and SiH_2 Singlet-Triplet Splitting

Singlet-triplet splittings in organic diradicals (carbenes) have been historically challenging to compute [510, 511]. However, with the advent of highly correlated methods using large basis sets, the situation has greatly improved. Nowadays, these systems are also used to assess electronic structure methods, for example in the development of the NEVPT2 method [512].

In the context of benchmarking the MR-EOMPT method, the singlet-triplet splittings reported in this section differ from the results reported previously in this chapter in that they are not transitions between different electronic states of a frozen geometry, but rather require computations on two distinct geometries since the ground-state geometries of singlet and triplet carbenes are not identical, especially when it comes to the bond angle on the atom carrying the two unpaired electrons. Generally speaking, singlet carbenes feature significantly smaller bond angles than do triplet carbenes [510]. Hence, the singlet-triplet splittings can be regarded as a test of the quality of the ground-state energy reported by various computational methods.

This test has already been used in Ref. [512] for the NEVPT2 method. In contrast to that publication, we use a larger basis set, def2-SVP, for our computations and use a FCI or ICE-FCI reference for methylene and silylene, respectively. The results are presented in Tab. 6.7.

We can see from Tab. 6.7 that the iterative MR-EOMCC outperforms all other methods, being off by < 0.1 kcal mol^{-1} for both systems. MR-EOMPT betters fic-NEVPT2 on these systems, with errors of about 0.7 kcal mol^{-1} and 1.3 kcal mol^{-1}, respectively. Nonetheless, it must be emphasized that all these errors are small, being well below 0.05 eV in all cases.

Table 6.7 Singlet-triplet ($^1A_1$–$^3B_1$) separation (kcal mol^{-1}) for methylene and silylene, computed at the equilibrium geometries from Ref. [512].

Method	CH_2	SiH_2
FCI†	14.46	−18.18
CASSCF	14.61	−16.51
fic-NEVPT2	15.75	−17.00
MR-EOMPT	15.10	−17.48
MR-EOMCC	14.54	−18.24

†Full CI for methylene, ICE-FCI ($T_{\text{gen}} = 10^{-5}$) for silylene.

6.7 Excited States of the Co and Cr Atoms

In the development of the MR-EOMCC method, a benchmark paper on electronic transitions on bare metal atoms and ions has been published to demonstrate the accuracy of several variants of the MR-EOMCC approach [216]. These systems are well suited to illustrate the main benefit of transform-then-diagonalize methods such as the MR-EOM variants, since a small active space followed by a correlation treatment can give access to basically all states with d-d transitions. For the purposes of this publication, we will pick two atoms, cobalt and chromium, and compare the performance of the perturbative variant, MR-EOMPT, and the other methods to the experimental reference values from the NIST atomic spectra database (NIST ASD) [513]. It must, however, be emphasized that different electronic structure programs (ORCA [286, 287] vs. ACES II [514]) were used to compute the MR-EOMCC values. Furthermore, we note that what we denote as "MR-EOMCC" most closely resembles "MREOM-T|SXD|U-min" from Ref. [216], with an additional $\hat{T}^\dagger$ transformation and the transformed Hamiltonian being symmetrized as discussed in Sec. 6.1.1.

All calculations in this section were done with Douglas-Kroll-Hess (DKH) scalar relativistic treatment [515, 516] and the def2-TZVPP basis set [463] with the exponents recontracted for the DKH procedure. Furthermore, the active space consisted of the 4s and 3d orbitals, with reference states being chosen such that the average occupation of the 4s orbital is 1.5, as described in Ref. [216]. In contrast to Ref. [216], we do not align the average energies of the calculated and experimental states prior to comparing the transition energies, but rather simply zero-align the lowest computed or experimental state and then compute the transition energies with respect to that state. We do, nonetheless, use the same J-averaging procedure to obtain a single energy for the non-degenerate sublevels of the experimental L-S multiplets [216],

$$E_{LS} = \frac{\sum_J (2J+1) E_J}{\sum_J (2J+1)}. \tag{6.21}$$

For both systems in this section, the setup of the CASSCF and NEVPT2 calculations differs from the general procedure described in Sec. 6.2. The other CASSCF and NEVPT2 calculations in

this chapter generally include all desired roots in the SA-CASSCF calculation (possibly split into different multiplicity and irrep blocks), which would lead to detrimental results as we would average over tens of states at the same time. Hence, we effectively report CASCI results, and NEVPT2 values computed on top of that, with the input orbitals from the MR-EOMPT and -CC reference states described below.

The transition energies for a neutral cobalt atom with a CASSCF reference of equal weights of the two lowest-lying ^{4}F ($4s^1 3d^8$) and ^{4}F ($4s^2 3d^7$) states are summarized in Tab. 6.8. In the Co atom results, the ^{2}D ($4s^2$) state was possibly not well described by the either the MR-EOMCC or the MR-EOMPT method, as indicated by reference weights of 89% and 82%, respectively [215]. Consequently, this state is removed from the final statistical evaluation of the errors in the transition energies with respect to the experimental reference values. For consistency, it is also not included in the evaluation of the CASCI and NEVPT2 methods. However, we note that including the state in the statistical evaluation would have improved the accuracy of the methods.

On the Co atom and with this particular CASSCF reference, fic-NEVPT2 and MR-EOMCC perform the best, with errors always below 0.17 eV for all transitions. The performance of the MR-EOMPT method is reasonable and certainly much better than the CASCI results, but nonetheless a bit disappointing with errors as large as 0.59 eV on the ^{4}F ($4s^2$) $\rightarrow$ ^{2}P ($4s^2$) transition. We suspect that the reason for this behavior lies in the fact that the CASSCF reference mixes the $4s$ and $3d$ orbitals among themselves in a detrimental way that cannot be corrected for as in the iterative MR-EOMCC method.

On a different note, we found that the current implementation of MR-EOMPT results in a slight degeneracy breaking of the L-S states of up to 3 meV, despite the SA-CASSCF reference being averaged over complete multiplets. However, due to the *s*-*d* mixing described above, the active orbitals differ in their energies depending on how much *s*-character they possess. Such a degeneracy breaking is not observed for the MR-EOMCC results.

Our second system is a neutral chromium atom with a CASSCF reference including equal root weights for the ^{7}S ($4s^1 3d^5$), ^{5}S ($4s^1 3d^5$) and ^{5}D ($4s^2 3d^4$) states [216]. The individual transition energies are reported in Tab. 6.9, including the statistical evaluation of the errors with respect to the experimental reference values.

From the reported errors, MR-EOMCC outperforms all other methods on this system with an RMSD of 0.07 eV over 18 states (118 roots), as expected from Ref. [216]. In contrast to the Co atom system from above, MR-EOMPT does not perform much worse at an RMSD of 0.12 eV with respect to the NIST values. In fact, the roles are reversed with fic-NEVPT2, which has deviations of up to 0.47 eV on the X ^{7}S ($4s^1$) $\rightarrow$ ^{3}D ($4s^1$) transition. In combination with Ref. [216], we may again conclude that MR-EOMCC performs accurately on these systems, and that the MR-EOMPT method is less robust in terms of consistency, but still gives the expected accuracy of well-established perturbational approaches on these systems.

Table 6.8 Transition energies (eV) between different states on neutral cobalt. Red values indicate unreliable results by a too low reference contribution. The errors were computed with respect to the NIST reference data [513], save for the unreliable ^{4}F ($4s^2$) $\rightarrow$ ^{2}P ($4s^2$) transition.

					MR-EOM	
Config.	**Term**	**NIST**	**CASSCF**	**NEVPT2**	**-PT**	**-CC**
$4s^2$	^{4}F	0.00	0.00	0.00	0.00	0.00
$4s^1$	^{4}F	0.42	1.78	0.38	0.47	0.28
$4s^1$	^{2}F	0.88	2.43	0.88	0.46	0.71
$4s^2$	^{4}P	1.63	2.16	1.76	1.52	1.73
$4s^1$	^{4}P	1.83	3.94	1.86	2.22	1.77
$4s^2$	^{2}G	1.99	2.16	2.08	2.40	2.14
$4s^1$	^{2}D	1.97	3.11	1.94	1.60	1.88
$4s^1$	^{2}P	2.20	2.86	2.28	1.94	2.13
$4s^2$	^{2}P	2.47	4.62	2.48	3.06	2.62
$4s^2$	^{2}H	2.64	2.88	2.74	3.21	2.76
$4s^2$	^{2}D	2.68	3.79	2.84	2.55	2.83
$4s^1$	^{2}G	2.78	4.78	2.63	2.63	2.70
		Statistical Evaluation of Errors				
	Min.		0.17	−0.15	−0.42	−0.17
	Avg.		1.19	0.02	0.07	−0.01
	Max.		2.14	0.13	0.59	0.15
	Std.		0.72	0.08	0.37	0.12
	MAE		1.19	0.07	0.33	0.11
	RMSD		1.39	0.08	0.38	0.12

Table 6.9 Transition energies (eV) between different states on neutral chromium. The errors were computed with respect to the NIST reference data [513].

					MR-EOM	
Config.	**Term**	**NIST**	**CASSCF**	**NEVPT2**	**-PT**	**-CC**
$4s^1$	^{7}S	0.00	0.00	0.00	0.00	0.00
$4s^1$	^{5}S	0.94	1.28	1.07	1.06	0.92
$4s^2$	^{5}D	1.00	1.10	1.15	1.20	1.05
$4s^1$	^{5}G	2.54	3.57	2.63	2.52	2.50
$4s^1$	^{5}P	2.71	4.09	2.94	2.81	2.81
$4s^2$	^{3}P	2.95	3.54	3.40	3.15	3.03
$4s^2$	^{3}H	2.99	3.31	3.27	3.09	2.95
$4s^1$	^{5}D	3.01	4.41	3.11	3.05	3.01
$4s^1$	^{3}G	3.09	4.47	3.42	3.13	3.03
$4s^2$	^{3}F	3.11	3.63	3.49	3.29	3.15
$4s^1$	^{3}P	3.37	4.95	3.66	3.56	3.47
$4s^2$	^{3}G	3.44	3.86	3.60	3.59	3.42
$4s^1$	^{3}D	3.55	5.30	3.91	3.63	3.52
$4s^1$	^{3}D	3.85	5.90	4.32	3.93	3.86
$4s^1$	^{3}I	3.85	5.21	3.94	3.75	3.73
$4s^1$	^{5}F	3.89	5.87	4.07	4.06	4.02
$4s^2$	^{1}G	3.97	4.57	4.31	3.99	3.87
$4s^2$	^{1}I	3.98	4.42	4.30	3.99	3.87
	Statistical Evaluation of Errors					
	Min.		0.09	0.09	−0.10	−0.12
	Avg.		1.01	0.25	0.09	0.00
	Max.		2.06	0.47	0.20	0.13
	Std.		0.61	0.12	0.08	0.07
	MAE		1.01	0.25	0.11	0.06
	RMSD		1.19	0.28	0.12	0.07

6.8 Comparison to LR-ic-MRCC

To show the differences between the approaches discussed in this chapter and linear response methods, we now compare some results against the linear response internally contracted multireference coupled-cluster (LR-ic-MRCC) method developed in the group of Köhn and coworkers [202]. However, we must emphasize that the LR-ic-MRCC method is expected to yield results that are of superior accuracy to the perturbative MR-EOMPT method. A more fitting comparison would be against the MR-EOMCC results, which are also iterative in nature. Still, it is illuminating to understand in which situations either of the methods may be more applicable.

Response theory describes how a wave function behaves under an external, periodic influence, e.g., electromagnetic radiation using a series expansion in terms of response functions. From the response functions, the energies of the excited states can be determined as they correspond to poles of the response function. As such, the excited states and transition energies are not limited to lie within an active space, but rather require a satisfactory description of the ground-state wave function. It is also sufficient to consider only the linear response function if only excitation energies and transition moments between ground and excited states are desired. A more detailed exposition can be found in Refs. [69, 517], of which Ref. [69] also details the relation to EOM methods in the single-reference case.

Response theory is therefore in stark contrast to the other methods discussed in this chapter, including the MR-EOMPT method. As discussed in Sec. 6.1.3, these methods require that all states of interest lie within the CASCI space spanned by the active electrons and orbitals. For example, given a full-valence active space, it is not possible to describe a valence-to-Rydberg transition, as the leading configurations of the final state would certainly lie outside the CASCI space. Hence, to compare to the LR-ic-MRCC results, we generally need to choose larger active spaces than the minimal spaces in the published reference [202].

We first present the singlet states of the methylene system. A full description of this test system has already been given in Refs. [113, 202]. Naturally, we use the same geometry and basis set in our calculations (see Sec. B.2.4). We must, however, remark that the reported [113] FCI transition energies $1\ ^1\mathrm{A}_1 \rightarrow 2\ ^1\mathrm{A}_2$ and $1\ ^1\mathrm{A}_1 \rightarrow 3\ ^1\mathrm{B}_1$ appear to be in error. While our own FCI calculation perfectly matched the other reported transition energies, these two transitions are in error of -0.627 eV and -0.104 eV, respectively. Using our corrected values also improves the statistics [202] of the LR-ic-MRCC method for this test system (Tab. 6.10).

For our calculations (Tab. 6.10), we need to choose a larger active space than the minimal CAS(2,2) that may be used with the LR-ic-MRCC method (*vide supra*). For such a small system, a full valence active space is in order, which gives six active electrons in six active orbitals: all four carbon atomic orbitals from the second shell as well as two hydrogen 1*s* orbitals. The previous study [202] also reports calculations with a CAS(6,6) active space, including states outside of the CASCI space. We were forced to omit the $1\ ^1\mathrm{A}_1 \rightarrow 3\ ^1\mathrm{B}_1$ and $1\ ^1\mathrm{A}_1 \rightarrow 4\ ^1\mathrm{A}_1$ transitions in Tab. 6.10 since the leading configurations of the excited states lie outside of the CASCI space, and hence cannot be treated by our present methods. Otherwise, for the CASSCF

and fic-NEVPT2 results, each of the four irreps of the C_{2v} point group were computed in a separate calculation, whereas in the MR-EOMCC and MR-EOMPT calculations, all singlet roots were obtained in a single run. The SA-CASSCF reference for the latter two methods consisted of the lowest seven singlet roots with equal weights, regardless of the irrep. This reference was necessary since otherwise the model space was not decoupled sufficiently from the outside space and nearly singular $\hat{T}$ amplitudes led to convergence difficulties.

Table 6.10 Transition energies (eV) between different states on singlet methylene, all with respect to the ground state $1\ ^1A_1$. Transitions may be either singles- or doubles-dominated, denoted s and d, respectively. We provide two reference values: the FCI results from Ref. [113], and our own FCI results. The LR-ic-MRCCSD values were sourced from Ref. [202] and use $N_{\text{com}} = 2$, $\eta = 10^{-5}$ and a CAS(6,6). For the statistical evaluation of the errors, our own FCI results were chosen as the reference.

						MR-EOM		
Exc. State	**Type**	**FCI** (Pub.)	**FCI**	**CASSCF**	**NEVPT2**	**-PT**	**-CC**	**LR-ic-MRCCSD**
$2\ ^1A_1$	d	4.656	4.655	4.614	4.693	4.629	4.679	4.721
$3\ ^1A_1$	s	6.514	6.514	6.398	6.500	6.498	6.645	6.449
$1\ ^1B_1$	s	1.793	1.794	2.500	1.939	1.729	1.778	1.803
$2\ ^1B_1$	d	8.906	8.907	8.861	8.995	8.924	9.108	9.332
$1\ ^1B_2$	s	7.704	7.703	7.232	7.681	7.685	7.790	7.647
$2\ ^1B_2$	d	8.016	8.016	8.259	8.086	7.972	8.052	8.136
$1\ ^1A_2$	s	5.853	5.853	6.032	5.980	5.821	5.885	5.870
$2\ ^1A_2$	d	9.410[†]	10.036	9.714	10.009	10.073	10.235	10.063
Statistical Evaluation of Errors								
	Min.			−0.471	−0.027	−0.064	−0.015	−0.065
	Avg.			0.016	0.050	−0.019	0.087	0.068
	Max.			0.707	0.146	0.037	0.201	0.425
	Std.			0.341	0.064	0.030	0.077	0.146
	MAE			0.266	0.067	0.032	0.090	0.098
	RMSD			0.342	0.081	0.035	0.116	0.161

[†] This appears to be an error in the original reference [113].

All methods save for the CASSCF calculations show good accuracy and, somewhat surprisingly, the perturbative approaches perform best in this benchmark. In the case of the fic-NEVPT2 results, this has been aided by the fact that each irrep block was treated separately, as opposed to the MR-EOMCC and MR-EOMPT calculations, where all roots are obtained in a one calculation from a single reference. Otherwise, we assume that fortuitous error cancellation at work, especially since in other benchmark sets discussed in this chapter, non-perturbative methods generally show better performance.

All methods tend to overestimate transition energies, except MR-EOMPT, which underestimates them by −0.019 eV, on average. When comparing the iterative coupled-cluster methods, MR-EOMCC and LR-ic-MRCC, we find that the linear response theory is on average closer to the FCI results, but the larger standard deviation also leads to a higher overall RMSD. This is largely due to the $1\ ^1\mathrm{A}_1 \rightarrow 2\ ^1\mathrm{B}_1$ transition, a high-lying doubles-dominated transition, which is in error of 0.423 eV compared to the FCI reference. Had this transition been left out of the evaluation, then the MR-EOMCC and LR-ic-MRCC results would have been of comparable accuracy. We further emphasize that, in contrast to single-reference CC theory [69], there is no simple relation between the multireference EOM-CC and LR-CC theories.

We now present some additional calculations for the all-(*E*)-hexatriene molecule. The active space was chosen to be a CAS(6,6) of the π-system orbitals (three each of a_u and b_g symmetry). We focus on only two singlet transitions, which have already been discussed in Thiel's benchmark set [313] and in the context of the LR-ic-MRCC method [202]. The transitions are the $1\ ^1\mathrm{A}_g \rightarrow 2\ ^1\mathrm{A}_g$ transition to the optically "dark" state and the $1\ ^1\mathrm{A}_g \rightarrow 1\ ^1\mathrm{B}_u$ transition to the "bright" state.

Although MR-EOMCC calculations have already been published for hexatriene, compared to the results reported here, the calculation from Ref. [212] uses a different MR-EOMCC implementation that not only differs in the definitions of the cluster operators, but also uses a larger diagonalization space. Huntington *et al.* [215] have reported values for these transitions with a more similar implementation (*vide supra*), but again chose a different CASSCF reference with only five active orbitals and reference states of $^3\mathrm{A}_g$ and $^3\mathrm{B}_u$ symmetry. This choice was motivated by the fact that convergence with the MR-EOMCC method was problematic due to nearly singular amplitudes when exciting out of the 3-b_g orbital. Nonetheless, projecting out the offending amplitudes and choosing a reference of a single $^3\mathrm{B}_u$ state resulted in successful convergence of the many-body residual conditions. The same reference state was also chosen for the MR-EOMPT calculation.

Our results and the cited reference values are summarized in Tab. 6.11. Unfortunately, choosing the CC3 results [313] to be our reference values is not so simple since the $1\ ^1\mathrm{A}_g \rightarrow 2\ ^1\mathrm{A}_g$ is to a great extent a double excitation, $1a_u^2 1b_g^2 2a_u^2 \rightarrow 1a_u^2 1b_g^2 2b_g^2$, which is also evidenced by the rather low $\%T_1$ of 65.8%. When comparing the MR-EOMPT results to the MR-EOMCC and LR-ic-MRCC results, we find that the order of the excited states is the same, with the $2\ ^1\mathrm{A}_g$ state having lower total energy than the $1\ ^1\mathrm{B}_u$ state. This contrasts with CC3 results,

in which the ordering of the states is the reverse. Furthermore, the MR-EOMPT method yields higher transition energies than both MR-EOMCC and LR-ic-MRCC theories. If we compare the transition energies of the singles-dominated $1\ ^1\mathrm{A}_g \rightarrow 1\ ^1\mathrm{B}_u$ transition to the CC3 values, we find that the MR-EOMPT method is in error of 0.55 eV, compared to only 0.22 eV and 0.08 eV for the MR-EOMCC and LR-ic-MRCC results, respectively. It does, however, outperform the fic-NEVPT2 result on this transition, which has an error of 1.19 eV.

Table 6.11 Bright and dark states of all-(*E*)-hexatriene (eV). The LR-ic-MRCC values were obtained from Ref. [202], and the CC3 values from Ref. [313].

Transition	**CAS**	**NEVPT2**	**MR-EOMPT**	**MR-EOMCC**	**LR-ic-MRCC**	**CC3**
$1\ ^1\mathrm{A}_g \rightarrow 2\ ^1\mathrm{A}_g$	5.45	5.59	5.52	5.43	5.44	5.72
$1\ ^1\mathrm{A}_g \rightarrow 1\ ^1\mathrm{B}_u$	6.59	6.77	6.13	5.80	5.66	5.58

7 Reduced Scaling for Internally Contracted Multireference Theories

7.1 Theory

To convey our reduced-scaling scheme, we first recapitulate salient features of density matrices, before giving a rigorous description of the scheme to reduce the scaling of terms involving densities, including a fully worked example. We then conclude by introducing both our variants of fic-MRCI and fic-MRCC theory, while also describing how the reduced-scaling scheme was applied to them. Finally, we benchmark the reduced-scaling implementations against the standard, unreduced codes on a set of linear alkenes and polyenes.

7.1.1 Properties of Density Matrices

In this section, we discuss the two most important theoretical foundations of the reduced-scaling scheme presented in this chapter, i.e., the ability to arbitrarily shift the operators, or index pairs, in densities, as well as the resolution of the identity (RI), which is applied in the construction of the densities but can also be used to break the densities into smaller parts. The reduced-scaling strategy is then set forth in Sec. 7.1.2.

Since densities are just expectation values of operator strings, Eq. (2.38), the operators can be swapped according to the known commutation rule for spin-free orbital replacement operators, Eq. (2.43). In fact, through repeated application of Eq. (2.43), they can be permuted in a totally arbitrary fashion, given that additional *commutator terms* are added to the final expression,

$$\langle 0|\hat{E}_1 \cdots \hat{E}_N|0\rangle = \langle 0|\pi(\hat{E}_1 \cdots \hat{E}_N)|0\rangle + \sum_{i=1}^{M} \langle 0|\hat{p}_i[\hat{E}_n, \hat{E}_{n+1}]\hat{s}_i|0\rangle \,, \tag{7.1}$$

leading to densities of an order reduced by one for the commutator terms,

$$\gamma_n = \gamma_n^{(\pi)} + \sum_{i=1}^{M} \gamma_{n-1}^{(i)}, \tag{7.2}$$

which can be seen directly from Eqs. (2.43) and (7.1). Above, π is a permutation of the operators $\hat{E}_n$, with $n = 1, \ldots, N$ denoting the position of the operator in the string. To determine the

commutator terms, we use the fact that every permutation can be written as a product of transpositions under the condition that only adjacent operators may be permuted,

$$\pi = \prod_{i=1}^{M} (n_i, n_i + 1), \tag{7.3}$$

using the cycle notation for the M transpositions of adjacent operators required to form the permutation π. In each of the transposition steps, we have a prefix $\hat{p}_i = \hat{E}_1 \cdots \hat{E}_{n_i-1}$ and a suffix $\hat{s}_i = \hat{E}_{n_i+2} \cdots \hat{E}_N$, which are the operator strings that are not permuted by the respective transposition step, and may be empty strings. Both $\hat{p}_i$ and $\hat{s}_i$ depend on the order of the transpositions and are thus not uniquely defined. The result, however, will be the same if the decomposition of the permutation π into adjacent transpositions is correct. A full example is presented below in Eqs. (7.5)–(7.10).

As discussed by Knowles and Werner [518], with foundational work on coupling coefficients by Siegbahn [519], among others [520–524], the densities can be split *exactly* at an arbitrary position through a resolution of the identity,

$$\begin{aligned} \gamma^{t\cdots v\ v'\cdots t'}_{u\cdots w\ w'\cdots u'} &= \langle 0|\hat{E}^t_u \cdots \hat{E}^v_w \hat{E}^{v'}_{w'} \cdots \hat{E}^{t'}_{u'}|0\rangle \\ &= \sum_I \langle 0|\hat{E}^t_u \cdots \hat{E}^v_w|\Phi_I\rangle \langle \Phi_I|\hat{E}^{v'}_{w'} \cdots \hat{E}^{t'}_{u'}|0\rangle \,. \end{aligned} \tag{7.4}$$

The $|\Phi_I\rangle$ are the CSFs in the RI space, which is equivalent to the CAS space here. The RI is also used to construct the high-order densities from the coupling coefficients $\langle \Phi_I|\hat{E}^t_u|\Phi_J\rangle$ through repeated application, which can be made highly efficient through a matrix-based form [518]. However, constructing the high-order densities can become limiting despite efficient BLAS operations because of the steep scaling of $\mathcal{O}(n_{\mathrm{CSF}} n_{\mathrm{act}}^{2m})$, with m denoting the order of the density γ_m. Here, it is generally tacitly ignored that computing the densities scales exponentially with the size of the active space, since $n_{\mathrm{CSF}} \sim \exp(n_{\mathrm{act}})$ for CASSCF reference wave functions.

7.1.2 A Scheme for Reduced Scaling

With the necessary groundwork laid out in the previous section, we now present the full scheme that we implemented in ORCA-AGE II such that we can automatically reduce the scaling of internally contracted *ansatzes*. We discuss how the excitation operators are shifted, and how the subsequent introduction of the RI together with a different contraction pattern leads to the reduction in scaling. Finally, a fully worked example is given.

The goal can be summarized in that we aim to contract away indices as soon as possible, as that directly leads to lower-order intermediates and a lower scaling. To this end, the first step is to shift the operators in the densities according to Sec. 7.1.1 such that summed indices that are shared with another tensor are grouped together and moved to one side of the density. Note that, to be useful, the summed indices *must* be on one excitation operator, e.g., $\hat{E}^t_u$, as opposed

to multiple operators, $\hat{E}^t_* \hat{E}^u_*$. Sharing the same operator with non-summed indices leads to no savings as the non-summed indices will be introduced to the intermediate instead. In summary, the number of the operators involved must be smaller than the number of indices, $n_{\text{op}} < n_{\text{idx}}$. Then, the scaling reduction is from $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^{2m})$ to $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^{2m-2o})$, where o is the number of operators with both indices shared that can be removed by contraction.

Next, the RI is introduced between the group of summed indices and the remaining operators in the density. Importantly, we only introduce the RI once into a density, since multiple splits would lead to transition densities that scale at least as $\mathcal{O}(n_{\text{CSF}}^2)$. For large active spaces, would quickly become unmanageable. Further, our automated scheme so far allows insertion of the RI after one, two, and three $\hat{E}$ operators, dependent on the contraction pattern, which we found sufficient for all use cases so far. More variability may be required if densities higher than the five-body γ_5 appears in the working equations.

Last, the partial densities must be contracted with another tensor that ideally shared many summed indices with the grouped excitation operators, as discussed above. This can be regarded as a "compression" of information into lower-order intermediates, thus realizing the scaling reduction.

The foregoing is perhaps best illustrated with an example. Some concrete examples have already been presented by Chatterjee and Sokolov [381] and Kollmar *et al.* [382], although with a focus on a specific theory. In the following, we aim to demonstrate the concepts laid out above, and present real-world examples from our automated software below. For this purpose, the example is simplified, with only two tensors in the contraction, as opposed to working equations from fic-MRCI and fic-MRCC theory, which contain more tensors.

The binary contraction for which we present the reduced-scaling scheme is

$$Y^{tva}_{uwt'} = \sum_{u'v'w'} t^{u'a}_{v'w'} \gamma^{tv\ v'\ t'}_{uu'w'w}, \tag{7.5}$$

which scales as $\mathcal{O}(n_{\text{act}}^8 n_{\text{virt}})$. Here, we highlighted the summed indices in blue to indicate the position of the $\hat{E}$ operators in the density. Again, the scaling of this contraction appears to be better than expected because it is generally ignored that it requires the computation of γ_4, which scales as $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^8)$. We begin by shifting the index pairs with the summed indices arbitrarily to the right-hand side of the density, which is described by the permutation $(2,3,4) = (2,3) \cdot (3,4)$ in cycle notation. Since two transpositions are required to form the full permutation and each transposition leads to two commutator terms, we get a total of four commutator terms, Eqs. (7.7)–(7.10), in addition to the contraction with the permuted density,

Eq. (7.6),

$$Y^{tva}_{uwt'} \leftarrow \sum_{u'v'w'} t^{u'a}_{v'w'} \gamma^{tt'\,v\,v'}_{uwu'w'}, \tag{7.6}$$

$$Y^{tva}_{uwt'} \leftarrow \sum_{u'v'} t^{u'a}_{v't'} \gamma^{tvv'}_{uu'w}, \tag{7.7}$$

$$Y^{tva}_{uwt'} \leftarrow -\sum_{u'v'} t^{u'a}_{ww'} \gamma^{tvt'}_{uu'w'}, \tag{7.8}$$

$$Y^{tva}_{uwt'} \leftarrow \sum_{v'w'} t^{t'a}_{v'w'} \gamma^{tvv'}_{uww'}, \tag{7.9}$$

$$Y^{tva}_{uvt'} \leftarrow -\sum_{u'v'w'} t^{u'a}_{v'w'} \gamma^{tt'v'}_{uu'w'}. \tag{7.10}$$

In Eqs. (7.6)–(7.10) and below, the left arrow, $\leftarrow$, means "add to," since the single contraction from Eq. (7.5) is broken up into multiple contributions in the reduced-scaling approach. In this series, we can easily see that all four commutator terms scale as $\mathcal{O}(n_{\text{act}}^7 n_{\text{virt}})$, and more importantly, only contain the three-body density γ_3, which only requires $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^6)$ operations to be computed. The Kronecker deltas in the commutator terms have been removed through summation.

We can then insert the RI,

$$Y^{tva}_{uwt'} = \sum_{u'v'w'} \sum_{I} t^{u'a}_{v'w'} \langle 0|\hat{E}^t_u \hat{E}^{t'}_w|\Phi_I\rangle \langle \Phi_I|\hat{E}^v_{u'} \hat{E}^{v'}_{w'}|0\rangle\,, \tag{7.11}$$

and form the first intermediate with a partial density, henceforth referred to as a "CSF intermediate" for the appearance of the index I enumerating the CSFs in RI space,

$$X^a_{v,I} = \sum_{u'v'w'} t^{u'a}_{v'w'} \langle \Phi_I|\hat{E}^v_{u'} \hat{E}^{v'}_{w'}\rangle\,. \tag{7.12}$$

Computing Eq. (7.12) scales as $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^4 n_{\text{virt}})$, slightly less than forming the resulting tensor,

$$Y^{tva}_{uwt'} \leftarrow \sum_{I} X^a_{v,I} \langle 0|\hat{E}^t_u \hat{E}^{t'}_w|\Phi_I\rangle\,, \tag{7.13}$$

which scales as $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^5 n_{\text{virt}})$.

All the steps involved in both the "standard" implementation, as well as the reduced-scaling formulation, are summarized in Tab. 7.1. For large active spaces, $n_{\text{CSF}} \gg n_{\text{inact}}, n_{\text{act}}, n_{\text{virt}}$, it is obvious that either the computation of γ_4 or Eq. (7.13) become rate-determining. The dependence of Eq. (7.13) on the number of virtual orbitals is potentially problematic, yet it only delays the crossover when the reduced-scaling scheme performs better than the standard implementation. Even for large virtual orbital spaces of $n_{\text{virt}} = 1000$, this is the case for $n_{\text{act}} > 10$. Besides, we note that the intermediate in Eq. (7.12) depends on an amplitude, i.e., it must be recomputed in every iteration.

Table 7.1 Summary of the computational steps required to compute $Y^{tva}_{uwt'}$. The "standard" implementation includes the full computation of γ_4, whereas the reduced-scaling formulation avoids it, at the cost of additional commutator terms and a dependence of intermediates on the number of CSFs.

Equation	Standard	Reduced
γ_3	—	$\mathcal{O}(n_{\mathrm{CSF}} n_{\mathrm{act}}^6)$
γ_4	$\mathcal{O}(n_{\mathrm{CSF}} n_{\mathrm{act}}^8)$	—
(7.5)	$\mathcal{O}(n_{\mathrm{act}}^8 n_{\mathrm{virt}})$	—
(7.7)–(7.10)	—	$\mathcal{O}(n_{\mathrm{active}}^7 n_{\mathrm{virtual}})$
(7.12)	—	$\mathcal{O}(n_{\mathrm{CSF}} n_{\mathrm{act}}^4 n_{\mathrm{virt}})$
(7.13)	—	$\mathcal{O}(n_{\mathrm{CSF}} n_{\mathrm{act}}^5 n_{\mathrm{virt}})$

7.1.3 fic-MRCI Theory

The basics of our formulation of fic-MRCI theory are outlined in Sec. 2.6.2 and in Ref. [176]. For both the standard and reduced-scaling implementations, we use the rank reduction trick by Dyall [525], which is also used in NEVPT2 theory [170, 361, 362], to reduce the order of the densities from γ_5 at face value to γ_4, *before* we further reduce the scaling to at most γ_3. We did not continue to reduce the scaling beyond γ_3, as we do not consider γ_3 to become the rate-limiting contraction, although we might need to revisit this when combining the theory with approximations to CASSCF, e.g., the iterative configuration expansion (ICE) approach available in ORCA [438, 439]. The standard implementation in this thesis is identical in *ansatz* to the one presented earlier [176], although the implementation has been updated with ORCA-AGE II. We further remark that in fic-MRCI theory, we can guarantee that all CSF-intermediates have at most two additional indices, $X^{p}_{q,I}$, which is why the storage requirements for the CSF intermediates is rather low.

7.1.4 fic-MRCC Theory

Our implementation of fic-MRCC theory was introduced in Sec. 2.6.3, where we give the expressions for the energy, the residuals, and discuss details necessary for the solution thereof. In this section, we instead focus specifically on the density matrices that appear in the working equations of the theory.

The highest orders of density matrices appear in the quadratic residual elements,

$$r_P = \langle \tilde{\Phi}_P | [[\hat{H}, \hat{T}], \hat{T}] | 0 \rangle \,, \tag{7.14}$$

with up to the five-body density γ_5 [195]. Higher expansions of the BCH would incur even higher orders of densities, until all general indices in the Hamiltonian are contracted with inactive or virtual indices from the cluster operator. For instance, with terms up to $\hat{T}^3$, we would need the six-body density γ_6.

In the reduced-scaling variant, we only eliminate γ_5 from the working equations, as fic-MRCC theory has *irreducible* contractions with γ_4 that *always* scale worse when the RI trick is applied. For example, in

$$r_{ij}^{pq} \leftarrow \frac{1}{2} \sum_{\substack{rstu \\ vwab}} t_{is}^{ar} (ta|ub) t_{jw}^{bv} \gamma_{qpws}^{utvr} \tag{7.15}$$

the indices are arranged such that each tensor shares two summed indices with the density γ_4, but also has two unshared indices. This would lead to the introduction of additional indices in the CSF intermediates, and accordingly no computational savings. By keeping the terms with γ_4, we also simplify the generation of the equations, since we do not need to make recursive calls, as reducing γ_5 produces commutator terms containing γ_4. In the resulting working equations, the CSF intermediates carry at most four additional indices, $X_{rs,I}^{pq}$, an increase by two from fic-MRCI theory, with correspondingly larger storage requirements. As a final remark, please note that in Eq. (7.15), we used the indices p, q, r, s to refer to active indices to avoid tedious double-primed labels.

7.2 Computational Methods

In this section we explain how the reduced-scaling versions of fic-MRCI and fic-MRCC theory were implemented and which computational settings were used to produce the numerical results. For the implementation, which uses the ORCA-AGE II code generation toolchain, we only describe the updates that were made specifically for the reduced-scaling scheme and refer to Ch. 4 for further information. The resulting code was integrated into the ORCA program package [288, 289], a development version of which was used for all numerical calculations.

7.2.1 Implementation through Automatic Code Generation

The reduced-scaling scheme presented in this thesis has been fully automated, with no user intervention required. (I) It begins with an additional step after the generation of the canonical, "standard" working equations, which applies the reduced-scaling scheme in a completely method-agnostic way to the working equations. To this end, all generated tensor contractions are first checked and filtered for high-order densities. (II) Then, the full scheme is applied independently to each of these contractions by doing an exhaustive, global search of the full space of sensible operator permutations and RI splits, of which each permutation/split is factorized in a termwise optimal fashion (see below). (III) Eventually, we pick the lowest-cost factorization, as estimated through a simple cost model, which currently assumes 10 active orbitals and 20,000 CSFs in RI space. In cases such as Eq. (7.15), where no satisfactory reduced-scaling variant exists, the contraction is left unmodified, and a warning is printed instead. (IV) As the last step, we check for duplicate contractions or intermediates and continue with the normal code generation procedure.

We should emphasize that we found a simple, linear factorization tool to be insufficient, i.e., the factorization tool must guarantee that the best possible pattern is found for each contraction. This is best illustrated with an example from fic-MRCI theory,

$$\sigma_{ij}^{tu} = \sum_{\substack{vwt' \\ u'v'w'}} (vj|wt')\gamma_{ut'tw'}^{vwu'v'} t_{iw'}^{u'v'}. \tag{7.16}$$

For Eq. (7.16), the only sensible reduced-scaling factorization is

$$X_{j,I}^{u} = \sum_{vwt'} (vj|wt') \langle 0|\hat{E}_u^v \hat{E}_{t'}^w|\Phi_I\rangle, \tag{7.17}$$

$$Y_{i,I}^{t} = \sum_{u'v'w'} \langle \Phi_I|\hat{E}_t^{u'} \hat{E}_{w'}^{v'}|0\rangle t_{iw'}^{u'v'}, \tag{7.18}$$

$$\sigma_{ij}^{tu} = \sum_{I} Y_{i,I}^{t} X_{j,I}^{u}, \tag{7.19}$$

which corresponds to a pattern of $(AB)(CD)$, with $A = (vj|wt')$, $D = t_{iw'}^{u'v'}$, and B and C being the partial densities from Eqs. (7.17) and (7.18), respectively. A simple linear factorization tool can only generate patterns such as $((AB)C)D$ and permutations thereof, and consequently would never find the best pattern, leading to unnecessary inefficiencies in the generated code. The updated factorization tool in ORCA-AGE II can find the optimal termwise factorization through a global search of all unique factorization patterns. Although this scales as a double factorial,[1]

$$N_{\text{patterns}} = (2N_{\text{tensors}} - 3)!!, \tag{7.20}$$

the overall cost of this global search is limited as the maximum number of tensors per contraction in fic-MRCI and fic-MRCC theory is three and four, respectively. If a density is split, this increases to four and five, leading to at most 15 and 105 possible factorization patterns per term, which is rapid to evaluate in our toolchain.

Once the working equations have been finalized, the code generation part takes over for the actual implementation into C++ code. Overall, only three minor modifications were required for this part since the ORCA-AGE II toolchain is kept as general as possible. The first modification concerns the contractions of partial densities with other tensors, as, e.g., in Eq. (7.17). These contractions are handled through hand-coded functions in ORCA that contract a one-, two-, or three-body partial density with any number of arbitrary tensors through a flexible interface based on lambda functions [526]. All other contraction types, including CSF intermediates with other CSF intermediates (Eq. 7.19), are handled through the existing functionality of ORCA-AGE II, including automatic application of BLAS.

[1] An analogy to the number of all unique rooted binary trees can be made since every factorization pattern can be graphically represented as such a tree.

The second optimization was required to optimize the performance of the contractions with partial densities, which is achieved by preloading some quantities normally stored on disk (partially) into memory. The reason for this lies in the construction of the partial densities, which always starts at one end of the density and expands them through a RI, as discussed in Sec. 7.1.1. Therefore, the order of traversal of the indices in the hand-coded function described above is fixed. This potentially leads to high IO overhead or cache misses, which is why the tensors are preloaded into memory with the axes aligned for optimal CPU performance.

Third, once the CSF intermediates are computed, they are stored on disk as separate tensor chunks for every electronic configuration, with each chunk only containing the CSFs belonging to that configuration. This accounts for the possibly substantial size of these intermediates for large actives spaces, keeping them out of expensive memory until needed. Since we already accounted for different storage and tensor layouts as well as special contraction functions in ORCA-AGE II, no further changes to deal with them were necessary.

7.2.2 Computational Details

The numerical results reported below were run on unsaturated organic molecules, which were optimized at the B3LYP [494, 527–529], D3BJ [496, 530] / def2-TZVP [463] level of theory, with the def2/J [495] auxiliary basis set used for integral RI. The resulting geometries were checked for the existence of imaginary frequencies.

In all cases, the standard settings of the ORCA program package [288, 289] were used, which include use of the RIJCOSX [497, 498] approximation with the default integration grid, `DEFGRID2` (Lebedev angular grid of 302 points and Gauss-Chebyshev radial grid, integration accuracy of 4.67). Both the SCF procedure and the geometry optimization were performed using tight thresholds, `TightSCF` (10^{-8} E_h for energy and 10^{-5} a.u. for the orbital gradient) and `TightOpt` (10^{-6} E_h for energy and 10^{-4} a.u. for the largest element of the gradient).

The CASSCF calculations were always started from a MP2 natural orbital guess and use a CAS space indicated with the results. For the highly correlated calculations, we used the def2-SVP and def2-TZVP [463] basis sets and always performed a full MO transformation. The iterative methods were considered converged once the largest element of the residual vector falls below an absolute value of 10^{-6}. We already note here that the reduced-scaling implementations give *identical* results to the "standard" implementations and require no further parameters.

The timing benchmarks were all computed on a system containing dual AMD EPYC™ 75F3 processors.

7.3 Results and Discussion

In this section, we present the benchmark results for the standard and reduced-scaling implementations on two types of systems: linear alk-1-enes and polyunsaturated alkenes. The first set of systems is to determine if the scaling with system size in a small CAS(2,2) active space

is still as expected for the reduced-scaling implementation, whereas the second set of polyenes is to demonstrate the scaling with growing, large active spaces. Finally, we compare how the implementations scale with both increasing active spaces and system size.

7.3.1 Linear Alk-1-ene Chains

The main motivation behind investigating the reduced-scaling implementations on linear alkenes is to demonstrate that the scaling with system size for both fic-MRCI and fic-MRCC theory, $\mathcal{O}(N^6)$, is not changed by the reduced-scaling scheme. The benchmark is computed on linear alk-1-enes from ethene ($n_{\text{carbons}} = 2$) to through octadecene ($n_{\text{carbons}} = 18$), a subset of which is presented in Fig. 7.1. All calculations were done in both the def2-SVP and def2-TZVP basis sets [463] and include the (anti-)bonding π orbitals in the CAS(2,2) active space. The full system descriptions, including the number of electrons and the exact sizes of the orbital ranges, can be found in Tab. B.5.

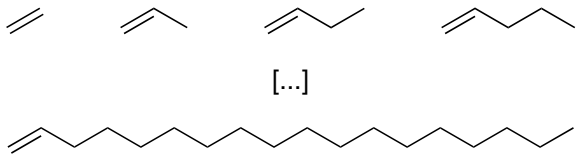

Figure 7.1 Linear alk-1-enes used for benchmarking of the reduced-scaling implementations. In the top row, ethene to pentene are shown, $n_{\text{carbons}} = 2, \ldots, 5$, and octadecene is shown in the bottom row ($n_{\text{carbons}} = 18$).

The timings for the standard and reduced-scaling implementations of fic-MRCI and fic-MRCC theory are shown in Fig. 7.2 for both basis sets (top and bottom row for def2-SVP and def2-TZVP, respectively). Each set of timings further consists of the pre-iterative and iterative timings. More specifically, the pre-iterative timings include the computation of the densities, the fic-NEVPT2 amplitude guess, and the "constant intermediates," i.e., contractions that do not depend on amplitudes and can hence be computed once prior to the iterative part. We note that our implementation of fic-MRCC theory does *not* use constant intermediates as the speedup was negligible, since virtually all factorized equations depend on an amplitude, and consequently cannot be pre-computed. The iterative part includes all other contractions that depend on amplitudes, and thus are recomputed in every iteration. The raw numerical data is reported in Tab. B.6.

As presented in Fig. 7.2, the pre-iterative times are negligible in the calculation since they are much faster than a single iteration of either fic-MRCI and fic-MRCC theory and only need to be computed once, compared to normally 10–15 times for the iterative parts. The only time when they are of some significance is in fic-MRCI theory, where the constant intermediates scale as $\mathcal{O}(N^4)$.[2] However, since the constant intermediates scale less than the iterative part, $\mathcal{O}(N^4)$

[2] The most expensive contractions in the standard fic-MRCI constant intermediates scale as $n_{\text{inact}}^2 n_{\text{virt}}^2$, $n_{\text{inact}} n_{\text{act}}^6 n_{\text{virt}}^2$,

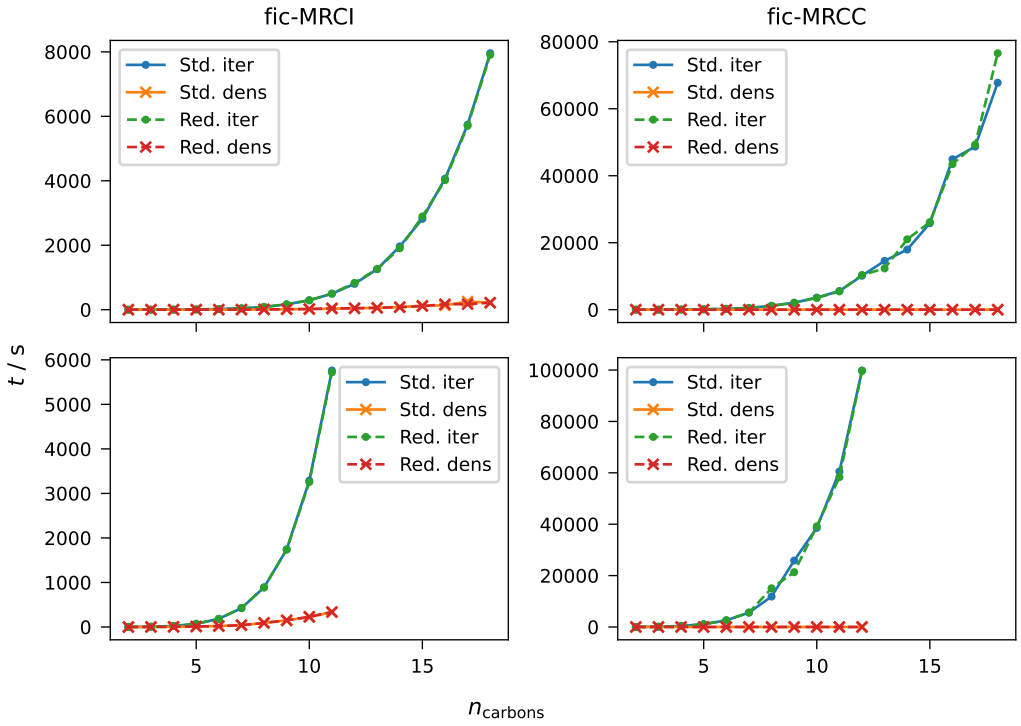

Figure 7.2 Scaling of "standard" and reduced-scaling implementations of fic-MRCI (left column) and fic-MRCC theory (right column) with system size on linear alk-1-enes (e.g., $n_{\text{carbons}} = 5$ for pentene). The top and bottom rows show the results in the def2-SVP and def2-TZVP basis sets [463], respectively. In each panel, the timings of both the average time per iteration as well as the time for the pre-iterative densities for the standard and reduced implementations are shown. Irregularities due to IO congestion issues.

vs. $\mathcal{O}(N^6)$, they will become less and less important in the overall timings the larger the system size is. In case of the fic-MRCC method, the fic-NEVPT2 guess dominates the pre-iterative timings, since the densities, even up to γ_5, all take less than 0.1 s to compute.

Most importantly, the iteration times of the standard and reduced-scaling implementations do not differ for either fic-MRCI or fic-MRCC theory. This confirms that the scaling of the reduced-scaling implementations is the same as the standard implementation for small active spaces, $\mathcal{O}(N^6)$.

In more detail, this scaling can be observed in Fig. 7.2 both with respect to the number of carbon atoms, n_{carbons}, as well as with respect to the increase in basis set size, by comparing the top and bottom rows. The steep scaling in the def2-TZVP basis set limits the achievable system size for the fic-MRCC systems to dodecaene ($n_{\text{carbons}} = 12$), compared to octadecene ($n_{\text{carbons}} = 18$) in the def2-SVP basis set. These calculations further show what is "routinely achievable" with our fic-MRCC implementation, since an average time of 10^5 s $\approx$ 28 h per iteration for these systems corresponds to a total runtime of $\approx$ 14 days for typically twelve iterations. It must, however, be emphasized that all these results are from single-threaded computations, as the parallelization has not yet been completed. Once the generated codes have been parallelized, we expect systems such as pentadecene ($n_{\text{carbons}} = 15$) to be viable targets for our fic-MRCC implementation in a triple-ζ basis set.

If we compare fic-MRCI and fic-MRCC theory (left and right columns of Fig. 7.2, respectively), we can immediately see the much higher prefactor of the fic-MRCC method. In the def2-SVP basis set, the prefactor is about ten times higher, which grows to about 15 for the def2-TZVP basis set. A reason for this lies in the constant intermediates of fic-MRCI theory, which noticeably reduce the iterative time in fic-MRCI theory and are absent in fic-MRCC theory. Moreover, the higher IO and memory pressure that fic-MRCC theory puts on the compute nodes slows down those calculations slightly.

The high IO pressure of fic-MRCC theory is also evidenced by the distinct irregularities in the iterative timings for both the standard and reduced-scaling implementation in Fig. 7.2. This is not due to the fic-MRCC method itself, but rather due to an overloaded shared RAID IO-system for the dual-CPU compute nodes because two calculations were run simultaneously on one of two NUMA-nodes on the same server. This does not lead to CPU or memory resource contention, but the IO system cannot handle heavy IO operations from both processes simultaneously, thus resulting in irregular iteration times. A hardware issue was also responsible for some missing calculations of the fic-MRCC method.

and $n_{\text{act}}^2 n_{\text{virt}}^3$, the first of which will be limiting for $n_{\text{inact}} \lesssim n_{\text{virt}} \gg n_{\text{act}}$, i.e., large systems with small active spaces like the alk-1-enes. None of the 4-external terms with $\mathcal{O}(N^6)$ scaling appear here since they always contain an amplitude, and are thus not constant.

7.3.2 Polyunsaturated Alkenes

We now turn to systems with large active spaces, on which we can also investigate the effect of an increasing active space size with the other orbital ranges kept at a similar size. To this end, we chose all-(*E*)-polyunsaturated alkenes (polyenes) from ethene to tetradecaheptaene in the def2-SVP basis set [463] (Fig. 7.3).

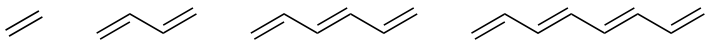

Figure 7.3 Linear, all-(*E*)-polyunsaturated alkenes from ethene to octatetraene. Each polyene can be treated with active spaces ranging from CAS(2,2) to CAS(2*n*,2*n*), where *n* is the number of double bonds.

The polyenes allow us to investigate active spaces ranging from a small CAS(2,2) for a single double bond up to CAS($2n, 2n$) in steps of 2, where n is the number of double bonds. Each active space contains both the π and corresponding π^* bonds. Hence, as the active space size increases, the orbital ranges change according to the following pattern,

$$\begin{aligned} \mathrm{CAS}(n_{\mathrm{act}}, n_{\mathrm{act}}): &\quad n_{\mathrm{inact}}, n_{\mathrm{act}}, n_{\mathrm{virt}}, \\ \mathrm{CAS}(n_{\mathrm{act}}+2, n_{\mathrm{act}}+2): &\quad n_{\mathrm{inact}}-1, n_{\mathrm{act}}+2, n_{\mathrm{virt}}-1, \end{aligned} \tag{7.21}$$

which is almost *ceteris paribus* regarding the number of inactive and virtual orbitals. The largest active space corresponds to the full π-system of the fully unsaturated polyene, and the smaller active spaces to a chemically sound subset thereof. In all cases, we used an MP2 natural orbital guess as the starting orbitals for the subsequent CASSCF calculation. The full system descriptions can be found in Tab. B.7.

Although we computed a full set of polyunsaturated alkenes from ethene to all-(*E*)-dodecahexaene, in the following we will focus on decapentaene and the results from fic-MRCC theory, as that system allows up to ten active orbitals, i.e., a large active space. The remaining calculations with smaller active spaces essentially confirm the findings of Sec. 7.3.1, and the corresponding, full dataset can be found in Tab. B.8.

The timings of the standard and reduced-scaling implementation of fic-MRCC theory on decapentaene are plotted in Fig. 7.4. From left to right, the panels show the timings of the pre-iterative step, the iterations, and the total runtime, each of which we will further discuss in this order below.

The timings of the non-iterative steps contain the computation of the densities, γ_5 and γ_4 for the standard and reduced-scaling implementation, respectively, as well as the fic-NEVPT2 guess. However, the fic-NEVPT2 guess is entirely negligible for the calculations with large active spaces. Conversely, the construction of the densities are negligible for small active spaces. Note that the timings below a few seconds are not entirely reliable, as they were obtained from a single run on a system subject to random fluctuations.

In the semilogarithmic plot of the pre-iterative times (left panel of Fig. 7.5), we can easily see the exponential scaling of both γ_4 and γ_5 with the size of the active space. The line of γ_4 has a

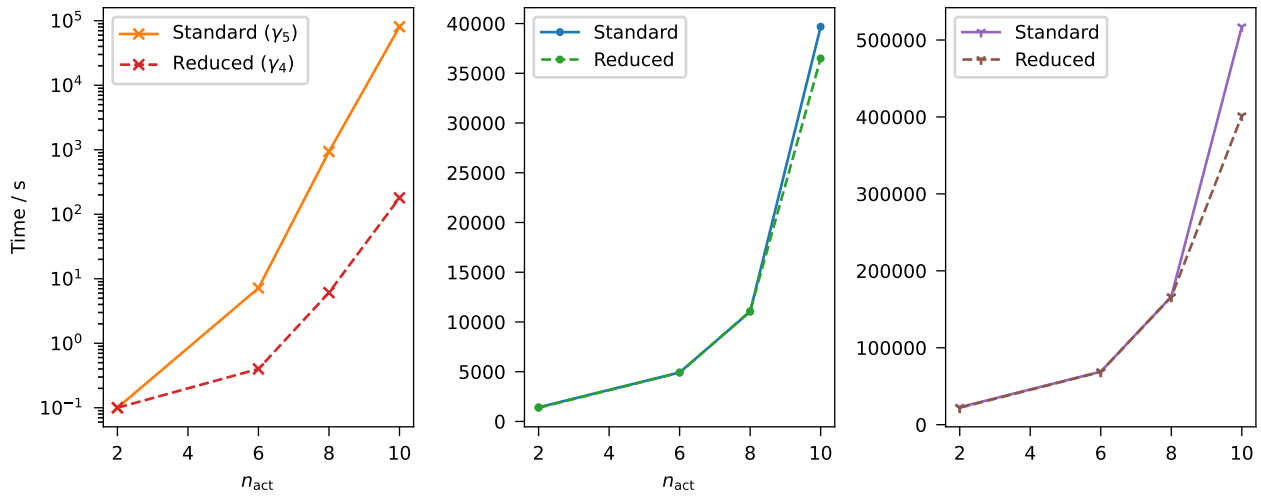

Figure 7.4 Scaling of the density computation and fic-MRCC iterations on decapentaene with symmetric actives spaces, CAS(n_{act},n_{act}). The left panel shows the pre-iterative timings, the center panel the average iteration time, and the right panel the total time.

lower slope than that of γ_5, which is consistent with the theoretical scaling (see also Sec. 7.1.1),

$$n_{CSF} \sim a e^{b n_{act}}, \tag{7.22}$$

$$\log(n_{CSF} n_{act}^{2m}) \sim b n_{act} + 2m \log n_{act} + \log a. \tag{7.23}$$

This finding is corroborated by the plots from Fig. 7.5, which shows the same scaling in the left panel with almost identical results since the computation of the densities is independent of the size of the system.

Continuing with the iterative part, the time per iteration of standard and reduced-scaling fic-MRCC (Fig. 7.3, center panel) are essentially indistinguishable for active spaces up to $n_{act} \leq 8$, although both show a steep scaling with n_{act}. Only for the largest active space, CAS(10,10), does the *reduced-scaling* variant show even shorter iteration times than the standard implementation. This is unexpected since the scaling of the (limiting contractions of the) reduced-scaling variant is $\mathcal{O}(n_{CSF} n_{active}^6 n_{virtual})$ (Eq. 7.25), compared to $\mathcal{O}(n_{inactive} n_{active}^{10})$ (Eq. 7.24) for the standard implementation. Given that $n_{CSF} = 19404 \approx 2n_{act}^4$ for this system, it should be *slightly slower* in this case, and the gains should come entirely from avoiding the computation of γ_5.

The explanation lies in the very slow computation of the terms containing γ_5: Looking at the most expensive terms from that computation, the first 23 of 30 terms contain γ_5. The worst offender is

$$(Y_{92738})^{iqsu}_{prtv} \leftarrow \sum_{p'q'w} t^{wp'}_{iq'} \gamma^{vuwp'q}_{tpsq'r}, \tag{7.24}$$

which not only contains the five-body density γ_5, but also the eight-dimensional target tensor Y_{92738}. To make matters worse, the tensors are also highly disordered, being stored in memory in row-major order as `Y92738(i,p,q,r,s,t,u,v)` and `G5(v,t,u,p,w,s,p',q',q,r)`, making simultaneous traversal impossible. For large, highly dimensional tensors as in this contraction (the size of γ_5 is 80.00 GB, well beyond the size of the CPU caches), this leads to frequent cache misses, which incur high memory latency and thus a big slowdown.[3] The slowdown is not due to a highly inefficient implementation, but rather inherent to the memory architecture of current compute hardware. To alleviate this issue, the high-order densities are aligned prior to computation for optimal usage in high-dimensional BLAS dgemm calls, which operate at peak CPU efficiency, in our code. For details on the TTGT scheme, we refer to App. A.3. A recent publication by Matthews [531] gives a full discussion of the problem, including alternative approaches to its solution.

In comparison, the slowest contraction of the reduced-scaling variant is

$$(X_{305})^{prt}_{qsa} \leftarrow \sum_{u,I} (X_{304})^{u}_{a,I} \, \langle \Phi_I | E^{ust}_{qpr} \rangle, \tag{7.25}$$

which only takes 56% of the computational time of Eq. (7.24). This contraction has at most 6-dimensional tensors, limited to a size of 127.2 MB, which is much smaller than γ_5. The CSF intermediate X_{304} is stored in $n_{\mathrm{CFG}} = 8953$ chunks of tensors of average size of 28 KB (combined size: 246.8 MB) since it is stored per configuration on disk, see also Sec. 7.2.1. The TTGT scheme is also used for these contractions, and they even have the disadvantage that the indices are fixed in the order of `t,r,w,p,u,q` in Eq. (7.25), for the densities must be constructed from the edge of the zeroth-order wave function, thus dictating the order of four `for` loops.[4]

As a final note, the speedup will be greater the fewer configurations there are, since it will make the reduced-scaling variant even cheaper. The same also holds true for the computation of the five-body density γ_5, which then also depends on fewer configurations or CSFs. The iterative part of the standard fic-MRCC implementation, however, will *not* experience any speedup since the size of the density only depends on n_{act}, still leading to large memory requirements and the corresponding inefficient computation.

Looking at the total computation time (Fig. 7.3, right panel), we find that they are essentially identical for $n_{\mathrm{act}} \leq 8$, since the computational cost of the pre-iterative steps is negligible and each iteration is timed about the same. However, for $n_{\mathrm{act}} = 10$ the reduced-scaling variant runs faster by 30%, both due to the reduction in the pre-iterative timings due to a reduction in the density order and the slightly faster iterative timings as well. We should note that we focused on the time per iteration above because the total timings have the drawback that they also depend

[3] The other 6 of 30 most expensive contractions do not contain γ_5, but γ_4 and two other six-dimensional tensors with two virtual labels, scaling as $\mathcal{O}(n^8_{\mathrm{active}} n^2_{\mathrm{virtual}})$.

[4] This contraction requires reordering of tensors as well, but on lower-dimensional, smaller quantities, which is much faster.

on the number of iterations, i.e., introduce another variable not inherently tied to the method, but rather dependent on the system and the starting orbitals.

These results are reinforced by the series of growing polyenes, all shown with their largest active spaces, CAS(2,2) for ethene, up to CAS(12,12) for dodecahexaene (Fig. 7.5). The main difference to the decapentaene plot (Fig. 7.3) is that the scaling of the iterative and total part is even worse, since both active space and system size grow.

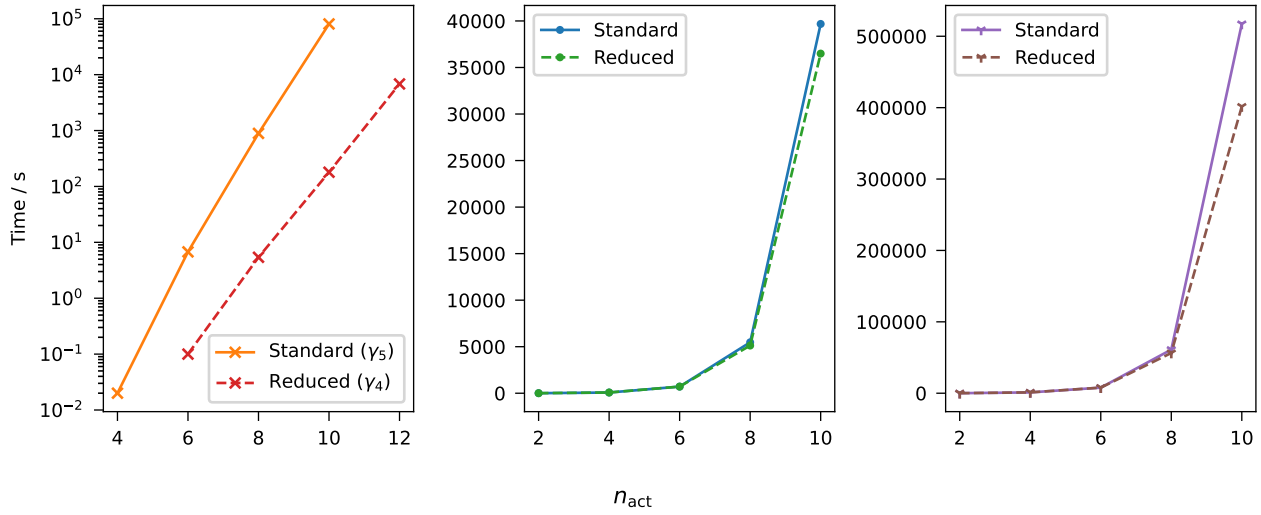

Figure 7.5 Scaling of the density computation (left panel), fic-MRCC iterations (center panel), and total timings (right panel) on growing polyenes in full π-systems, CAS(n_{act}, n_{act}), from ethene ($n_{act} = 2$) to dodecahexaene ($n_{act} = 12$).

The situation is more nuanced with the results from the standard and reduced-scaling implementations of fic-MRCI theory. Since fic-MRCI computations are faster than those of fic-MRCC theory, we discuss a larger system, all-(*E*)-dodecahexaene, instead of decapentaene here, with active spaces up to CAS(12,12). We first look at the timings for the construction of the densities, which appear up to γ_4 and γ_3 for the standard and reduced-scaling implementations, respectively (Fig. 7.6, left panel). As expected, the construction of the four-body density is more expensive than of the three-body density by roughly an order of magnitude, which translates to one-time savings of $\approx$ 1.7 hours for the CAS(12,12) active space on dodecahexaene. It is immediately obvious that the savings due to constructing only γ_3 instead of γ_4 are much less pronounced compared to the saving of avoiding the computation of γ_5 in fic-MRCC theory. Consequently, the reduced-scaling fic-MRCI method must be implemented in a highly efficient manner, with a correspondingly low prefactor, to be faster for medium-sized active spaces, or it should be used with much larger active spaces, e.g., those accessible with approximate methods such as the ICE method [438, 439], to benefit from the scaling reduction.

The computation of the constant intermediates is a significant factor in the fic-MRCI calcu-

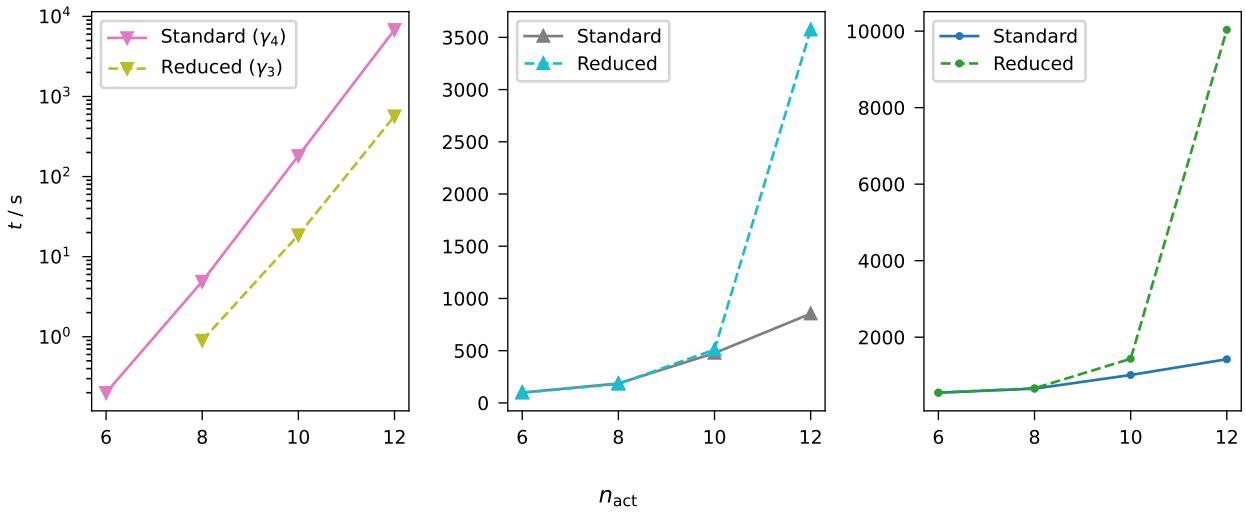

Figure 7.6 Scaling of the density computation and fic-MRCI iterations on dodecahexaene with symmetric actives spaces, CAS(n_{act},n_{act}). The left panel shows the timings of the density construction, the center panel the time for the constant intermediates, and the right panel the average time per iteration.

lations (Fig. 7.6, center panel). For the standard implementation, we only have a polynomial scaling (see also Sec. 7.3.1), leading to moderately longer timings with the size of the active space. For twelve active orbitals, the timing is 853.8 s, well below that of the computation of γ_4, which requires 6699.9 s. The situation is reversed for the reduced-scaling implementation, where the constant intermediates show a steep scaling with n_{act}. This is due to some constant intermediates being CSF intermediates, which introduces a dependence on the number of CSFs, n_{CSF}, in the scaling.[5] Therefore, especially for $n_{act} = 12$, the timings reflect the predicted exponential scaling, taking 3572.4 s to compute the reduced-scaling constant intermediates, or 53% of the time taken to compute γ_4 in the standard implementation.

In some regard, this distinction between the densities and what we call constant intermediates here is somewhat artificial, since one can make the argument that the densities are "constant intermediates" as well. A benefit of the densities is that they can indeed be entirely precomputed, in contrast to the CSF intermediates, which also appear in the iterative part (see below and Sec. 7.1.2).

A very similar picture presents itself for the average time per iteration (Fig. 7.6, right panel). Up to a CAS(8,8), the timings between the standard and reduced-scaling implementation are virtually identical, but the reduced-scaling variant becomes more expensive for CAS(10,10) by a factor of 1.4, and more expensive for a CAS(12,12) by a factor of 7.1, which is also due to the exponential scaling with n_{CSF}. The full timings can be found in Tab. B.8.

[5] The constant intermediates contractions with the worst scaling are $\mathcal{O}(n_{CSF}n_{inact}n_{act}^5)$, $\mathcal{O}(n_{CSF}n_{act}^6)$, and $\mathcal{O}(n_{CSF}n_{act}^4n_{virt})$.

Looking at the total computation times, we find that the reduced-scaling method takes longer to compute, at 37.4 h for 13 iterations, compared to only 7.3 h for the standard implementation. This is a consequence of the three effects discussed above: The reduced-scaling implementation only saves 1.7 h on the construction of the densities, and takes 0.8 h and 2.4 h more time for the computation of the constant intermediates and in every iteration, respectively.

To further elucidate this behavior, we now analyze the most expensive contractions for both the standard and reduced-scaling implementations of fic-MRCI theory, respectively,

$$\sigma_{pq}^{ab} \leftarrow \sum_{cd} (Z_{333})_{cd}^{pq}\, (bd|ac), \tag{7.26}$$

$$(X_{73})_{a,I}^{p} \leftarrow \sum_{qrst} (X_{72})_{qsa}^{prt}\, \langle E_{rt}^{qs}|\Phi_I\rangle, \tag{7.27}$$

which scale as $\mathcal{O}(n_{\text{act}}^2 n_{\text{virt}}^4)$ and $\mathcal{O}(n_{\text{CSF}} n_{\text{act}}^5 n_{\text{virt}})$. Interestingly, Eq. (7.26) *does not contain any density,* which is in stark contrast to the results reported for fic-MRCC theory above, where virtually all of the first 30 contractions depend on high-order densities. This is an indication that the contractions with the densities are (by far) not limiting in the iterative part of fic-MRCI theory, in contrast to fic-MRCC theory.

To better assess the relative performance of the limiting contractions, Eqs. (7.26)–(7.27), as well as the construction of the four-body density, we compare the measured and predicted, relative timings in Tab. 7.2. The timings of Eqs. (7.26)–(7.27) are given per iteration, as γ_4 is also computed only once. However, keep in mind that X_{73} and σ_{pq}^{ab} must be recomputed in every iteration, i.e., they must be computed $n_{\text{iter}} = 13$ times until the fic-MRCI iterations on dodecahexaene have converged. The predicted relative timings were computed from a simple estimation from the asymptotic scaling and the known orbital ranges for dodecahexaene (Tab. B.7), with the underlying assumption that we have already reached the asymptotic scaling regime, irrespective of implementation details and efficiency of the compute hardware. According to the asymptotic estimate, the computation of Eq. (7.27) and γ_4 should be 58 and 530 times more expensive than computation of Eq. (7.26), respectively. In this case, a single iteration of the reduced-scaling implementation on dodecahexaene would be faster than the computation of γ_4 by a factor of ≈ 7, for γ_4 would clearly be the limiting term. Overall, the entire calculation with $n_{\text{iter}} = 13$ would still be slower by 60% than the standard implementation. Regrettably, this is in stark contrast to the measured relative factors of 9.3 and 21.9. Perhaps most surprising is that the computation of γ_4 is only 2.4 times slower than the most time-consuming contraction of the reduced-scaling implementation, Eq. (7.27), in a single iteration, despite being predicted to be about an order of magnitude slower. The main reason for this is that the computation of γ_4 is truncated with `D4TPre 1.0e-14`, which reduces the number of electronic configurations in the CI space by about 50%, and correspondingly the computational effort, by only considering those with CI coefficients larger than `D4TPre`. Such a reduction is not yet implemented for the CSF intermediates. Another reason is that, while high-dimensional tensors do appear in the computation of γ_4, these can always be traversed perfectly in storage order through BLAS operations,

leading to an optimal computer implementation. For the CSF intermediate in Eq. (7.27), we are required to traverse the intermediate X_{72} out-of-order, leading to a more inefficient operation despite use of the BLAS-based TTGT engine (see Sec. A.3). Moreover, the CSF intermediates are stored on disk (Sec. 7.2.1), compared to the four-body density γ_4, which is stored entirely in high-speed memory, slightly skewing the comparison in favor of the standard implementations. For very large active spaces, though, there is no alternative to keeping the CSF intermediates on disk, and also the densities would eventually need to be removed from memory.

Table 7.2 Comparison of theoretical scaling, predicted, and measured, *relative* performance of the limiting contractions in fic-MRCI theory on dodecahexaene in a CAS(12,12). The most expensive contraction from the standard implementation, Eq. (7.26), was normalized to 1.0, and all other timings reported relative to it. The predicted factors were obtained from a simple estimate based on the asymptotic scaling of the terms.

Equation	Scaling	Measured	Predicted
(7.26)	$\mathcal{O}(n^2_{\text{active}} n^4_{\text{virtual}})$	1.0	1.0
(7.27)	$\mathcal{O}(n_{\text{CSF}} n^5_{\text{active}} n_{\text{virtual}})$	9.3	58.0
γ_4	$\mathcal{O}(n_{\text{CSF}} n^8_{\text{act}})$	21.9	530.0

Apart from only presenting the timings for the CAS(12,12) active space, we can also compare the scaling of Eq. (7.27) and the computation of the four-body density to assess when a crossover will occur (Fig. 7.7). For active spaces of up to $n_{\text{act}} \leq 10$, computation of X_{73} appears to scale worse than computation of γ_4. As mentioned above, this will be in part due to the highly efficient implementation of γ_4. For $n_{\text{act}} = 12$, however, the computation of X_{73} is faster that of γ_4. However, since X_{73} must be recomputed in every iteration, more time is spent computing X_{73} during the entire calculation on dodecahexaene than on constructing γ_4. All in all, we conclude that we have not reached the asymptotic scaling limit with this system yet.

To summarize, the reduction in asymptotic scaling does not show for the fic-MRCI method since computing γ_4 is highly efficient and occurs only once, whereas the CSF intermediates are *comparatively* inefficient and must be recomputed in every iteration. Nevertheless, the asymptotic scaling should guarantee that the reduced-scaling fic-MRCI implementation becomes faster for sufficiently large active spaces. Since the fic-MRCI method is significantly faster than the fic-MRCC method and will soon be fully parallelized, we suggest investigating this matter further, especially in conjunction with approximate schemes such as the ICE method.

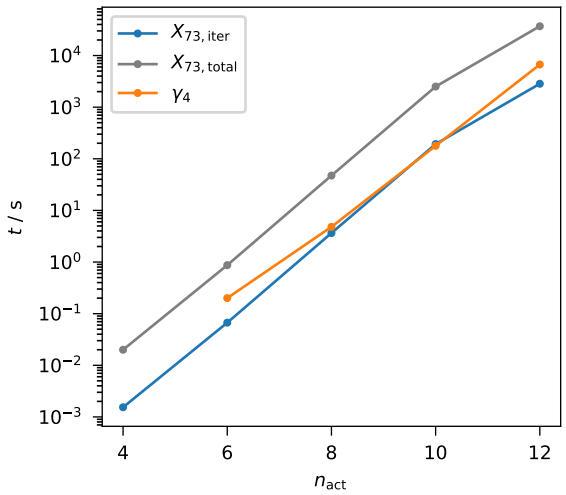

Figure 7.7 Scaling of X_{73} (Eq. 7.27) and γ_4 with the number of active orbitals, n_{act}, on dodecahexaene. All in all, $X_{73,\mathrm{iter}}$ must be recomputed in every of the $n_{\mathrm{iter}} = 13$ iterations ($X_{73,\mathrm{total}}$), making its total computation time during the calculation slower than that of γ_4.

8 Conclusion

In this thesis, we mainly shed light on open issues in multireference coupled-cluster theories in quantum chemistry. We state the main conclusions here, before further details and nuances are given in the following sections.

On the topic of the residual conditions in multireference theories, we developed a new relation between the many-body and projective residuals. The expansion we propose connects these two residual conditions in the form of a cumulant expansion, which in turn allows us to justify the truncations in the many-body expansions apparent in MR-EOMCC theory.

Still on the subject of MR-EOM methods, we also investigated the effects of turning the MR-EOMCC method, a transform-then-diagonalize approach, into a perturbative theory. While this does cause some loss in accuracy compared to the parent theory, the novel method is as accurate as another perturbative approach, i.e., NEVPT2 theory.

We also proposed an automated scheme to reduce the order of the density matrices in internally contracted multireference theories that do use projective residual conditions, e.g., fic-MRCI and fic-MRCC theory. While the scheme works exceptionally well for fic-MRCC theory, the high prefactor only justifies its use with fic-MRCI theory for the largest of active spaces and warrants further research.

Since all the above methods are highly complicated multireference theories, automatic code generation was required to perform the studies reported in this thesis. To this end, we implemented the ORCA-AGE II toolchain from scratch. It is a major improvement over the old version [1] in terms of functionality, extensibility, usability, and performance, requiring under five hours to generate the full version of fic-MRCC theory on a current desktop computer.

Lastly, we report the excellent results of a single-reference STEOM study on excited states of indigo dyes. Despite not being a multireference approach, STEOM shares a close connection with MR-EOMCC theory as a transform-then-diagonalize method. The results can be viewed as a celebration of what can currently be achieved with single-reference CC theories, setting a bar for multireference approaches both in terms of accuracy and usability.

8.1 Many-Body and Projective Residual Conditions

In Ch. 3, we provide insight into how well the many-body conditions perform when compared to the standard projection conditions in solving for a set of internally contracted multireference coupled-cluster amplitudes. To this end, we use two simplified MR-CCD approaches that differ

only in the residual conditions, having either many-body or projective equations. For the many-body equations, we remark that those equations are almost identical to the first $\hat{T}$ transformation of the MREOM method [215]. We provide a strict comparison of both approaches that details underlying assumptions and approximations in the parameterizations.

We explicitly relate the equations that result from many-body conditions to the projective equations by expressing the latter in terms of Hamiltonian elements and the associated prefactors of density matrices and cumulants. These terms can be arranged into a hierarchy of low- and higher-rank contributions in the projective conditions suitable for numerical comparison. The differences in the magnitude of the higher-order terms are analyzed by cross-substituting converged amplitudes from the many-body or projective conditions into the residual equations of the opposite condition, which results in a numerical gauge on the importance of the higher-order terms. The results show that the many-body conditions, which directly zero out elements of the transformed Hamiltonian, perform much better in keeping these higher-order terms small compared to the corresponding projective conditions, which just require a sum of terms to be zero. This, in turn, clearly connects to the idea of the sequential transformations used in MREOM theory, where higher-order Hamiltonian terms are neglected in the subsequent similarity transformations.

Our efforts should be viewed in line with the research by Gauss, Evangelista, Köhn, Nooijen and Chan on how to construct a generally applicable, fast and efficient multireference coupled-cluster method. Several important aspects such as general parameterization, amplitude removal, sequential transformations, transform-then-diagonalize strategies and unitary approaches have been discussed in the literature with their advantages and shortcomings, while the focus of this thesis was on characterizing the differences between many-body and projective conditions using a simplified MRCC approach. We hope that the considerations presented in this thesis will pave the way to understanding what aspects are truly important in designing a "gold-standard" multireference method.

8.2 Automatic Code Generation in Quantum Chemistry

In Ch. 4, we describe a newer version of ORCA-AGE, ORCA-AGE II, which is a complete overhaul of the previously introduced toolchain [1]. Due to algorithmic improvements throughout the code (especially in the time-limiting steps) and a shift in programming language from Python to C++, the toolchain is faster by a factor of $\approx$ 150 as measured while generating fic-MRCC theory. Its highly modular layout has also been refreshed, enabling us to add new code engines such as the TTGT scheme easily. It is also simpler to use than the old toolchain as external dependencies were removed and because of an even tighter integration with the main ORCA source code: all computational modules can be regenerated with a simple command, which automatically brings new features and better code to existing modules as well.

The new toolchain has proven to be useful in generating even the most complicated theories,

as demonstrated for fic-MRCC theory, for which an enormous number of terms was successfully generated through the commutation rule-based equation generator. The generated code is highly performant despite the large number of terms that must be computed.

In the end, developing automated code generators is beneficial for researchers as it allows them to spend their resources on their critical projects while at the same time transparently accounting for an efficient, production-level implementation. We believe that this is the direction in which *ab initio* quantum chemistry is evolving and will hence continue to align our development efforts on ORCA and ORCA-AGE with these trends.

We will now give an outlook of how we think quantum chemical software will develop in the future. While it is true that this is subjective to a certain degree, we believe that the underlying trends in ever more complex theories as well as consistent increases in (parallel) computing power, especially in the field of *ab initio* method development, require better and more powerful tools than just "handwritten code" to achieve an efficient scientific workflow. First, some theories such as fic-MRCC are too complex to be implemented in the conventional way in the first place, and thus obviously require higher-level tools. But even for less demanding theories, an automated toolchain can lead to higher productivity on the researcher's end, as tedious implementation details are transparently automatized, thus freeing time to spend on quantum chemistry instead of programming. Moreover, any performance improvements or new theoretical developments in the generator toolchain also transfer to older method implementations as well, without having to rewrite older codes.

Effectively, we envision a clearer partitioning of quantum chemistry codes in the future into low-level, mid-level, and high-level parts, to be explicated below. While the exact boundaries between the hierarchy of codes will be somewhat fluid, we nonetheless believe that the three main areas identified here successfully capture future directions of code development in quantum chemistry.

Under "low-level," we understand the foundational code of any quantum chemical software package that provides basic quantities such as molecular integrals and other general-purpose libraries used throughout the codes, e.g., a library for arbitrary-dimensional tensors. Integral libraries already exist (for example, Libint [532] and SHARK [288, 289, 533]) and have been extensively tuned for shortest possible runtimes. Interestingly, specialized code generation has also been employed for the integral libraries Libint [532] and SHARK [288, 289], as especially loop unrolling allows the compiler to generate the most efficient code. Other low-level libraries include BLAS, which allows peak efficiency for the ubiquitous matrix multiplications in modern-day quantum chemistry codes. As a side note, general BLAS-like libraries for arbitrary tensors that may become more pervasive in the future would also fall in this category. Noteworthy examples are TBLIS [531] and GETT [262], among others.

Under our classification, the mid-level code would mainly comprise common algorithms that are implemented by hand, such as general solvers and drivers. Given a flexible interface, these common algorithms enable the (iterative) computation of the desired molecular properties

without needing to copy-and-paste code, possibly with minor modifications depending on the actual method. Doing so guarantees that these core algorithms are free from errors since there is only the library that needs to be tested; and not multiple, slightly different implementations. Also, any improvements also apply across all implemented methods consistently.

The highest level is reserved for the actual implementation of quantum chemical theories such as CCSD, CCSDT, fic-MRCC, gradients, properties, … This is ideally enabled through a high-level interface that more closely resembles the mathematical equations or ansatzes rather than computer code. This can be achieved in many possible ways, each with its own benefits and drawbacks. In our group, we decided to focus on code generation since all contractions and tensors (with all their dimensions) are already known at compile time. Thus, there is no need to interpret the data on-the-fly, and the generated code can still be tailored to different architectures and to different framework backends such as the ones mentioned in the introduction (CTF [263, 264], libtensor [265], LITF [266], TCL [262, 267], and TiledArray [268, 269]). Other approaches may include interpreters such as SIAL [271, 272, 534, 535] or direct usage of the aforementioned libraries. Eventually, only the mathematical equations would be stored in the code repository, which makes them easy to understand for new researchers as well as easy to validate and modify.

In ORCA, first steps have been taken in this direction [289, 533]. In Ch. 4, we outlined the latest developments in ORCA-AGE's code generation capabilities, which showcases what we think will be representative of high-level interfaces in quantum chemical codes in the future.

8.3 Similarity Transformed Equation-of-Motion Study of Indigo Dyes

In Ch. 5, we reported an excited state coupled cluster study of indigoid dyes. After discussing the significance of these dyes, we investigated multireference effects and found that a single reference method, such as STEOM, should be sufficient to provide accurate results for this family of dyes. After studying basis set effects, the def2-TZVP basis set was chosen as an excellent compromise between accuracy and efficiency. Implicit solvent effects also had to be taken into account in order to achieve our target accuracy of 0.1 eV. We then presented the DLPNO-STEOM results for 0–0 transition energies and compared them to experimental absorption band maxima. We found that in virtually all cases, our predictions agree with experiment within 0.1 eV and that STEOM also reproduces the effects of chemical substitutions on the excitation energy. Finally, we presented calculated spectra for some of the main components of Tyrian purple and also compared these findings to a measured spectrum of the same dye. It is especially important to emphasize that our protocol was able to reproduce the experimental spectra of individual components *without* any empirical shift. Thus, we are confident that the VG-corrected DLPNO-STEOM method will become a useful tool for computational chemists studying photochemical processes.

8.4 Multireference Equation-of-Motion Perturbation Theory

In Ch. 6, we presented a new multireference perturbation theory, called MR-EOMPT, which is based on the MR-EOMCC method. The approach is best classified as a *transform-then-diagonalize* method in complete analogy to MR-EOMCC since the general theoretical structure is the same, except for the fact that perturbatively estimated amplitudes are used for the sequential similarity transformations. MR-EOMPT also shares the multistate character (which is closely related to their many-body aspect) with its parent method, and only requires a minimal MRCI diagonalization space due to the decoupling of the model CASCI space from the external space through the similarity transformations. Furthermore, MR-EOMPT and -CC only require one- and two-body RDMs of the CAS reference states, in stark contrast to other multireference perturbation theories, which formally require up to four-body RDMs. Our main goal was to show that using perturbative amplitudes in the MR-EOMPT method decreases computational cost while retaining at least the accuracy of other popular perturbational methods. Moreover, since the MR-EOMPT method foregoes the many-body residual equations, there are no convergence difficulties as in the MR-EOMCC method for nearly singular amplitudes.

The benchmarks presented mainly focus on vertical transition energies, as calculating these is the major application of EOM methods in which the excited states can all be obtained from a single state-universal calculation. The first test set of small diatomic molecules in a double-ζ basis set shows that the MR-EOMPT method performs on par with fic-NEVPT2 at an RMSD of 0.07 eV over a set of 25 transition energies. This test set of small molecules with FCI-quality reference data is complemented by a subset from Thiel and co-workers' benchmark [313] with larger molecules in a triple-ζ basis set. Compared to the published [313] CC3 reference values with a singles contribution of $\%T_1 > 90\%$, we found that the triplet transitions were treated more accurately than the transitions to singlet final states. For the triplet transitions, the RMSD across 12 transitions is 0.11 eV, about half of what was found for state-specific fic-NEVPT2 with an RMSD of 0.19 eV. The situation is the opposite for the singlets, where the RMSD across 11 transitions is 0.43 eV, about twice as much as for fic-NEVPT at 0.20 eV. The MR-EOMCC results with the same parent state are much more robust at an RMSD of only 0.18 eV for the singlet transitions, but still higher than for the triplet transitions (0.07 eV). For the MR-EOMCC results, this has already been found to be due to the $\hat{T}^\dagger$ transformation in ORCA [215]. We should note that in the subset of Thiel's benchmark, the fic-NEVPT2 calculations use orbitals optimized for each multiplicity and irrep, whereas the MREOM methods only use a single reference state.

The lithium fluoride PES scan was included to demonstrate the ability of the MR-EOMPT method to accurately reproduce weakly avoided crossings, which is straightforward with transform-then-diagonalize approaches since only the final step gives the states as combinations of the model space functions. When comparing the results to QD-NEVPT2, we found that the PES computed by the MR-EOMPT method is much more parallel to the FCI-quality reference without any artifacts located in the region of the avoided crossing. The MR-EOMPT method also performs better when it comes to the vertical transition energies between the neutral and

ionic states of LiF at each point of the PES, with the results being effectively identical to the ICE-FCI reference data.

To assess the accuracy of ground-state energy calculation, we chose to compute the singlet-triplet gaps on methylene and silylene. Despite not being the focus of EOM-like methods, the ground state calculations on the singlet and triplet geometries demonstrate that the singlet-triplet gap is reproduced accurately to within 0.7 kcal mol^{-1}, an error comparable to that of the fic-NEVPT2 method at 1.3 kcal mol^{-1}. This finding is also supported by the quality of the total energies for the diatomic test systems, where the MR-EOMPT method reproduces the FCI-quality total energies better than fic-NEVPT2.

A paramount example of the MREOM methods are transition energies on transition metals, since most of the *d-d* transitions can be described from a single parent state in a reasonably small CAS(n,5) calculation. We computed a total of 190 state energies on the cobalt and chromium atom combined and found again that the perturbative approach is not as robust as the fully iterative MR-EOMCC with RMDSs of 0.38 eV and 0.12 eV versus 0.12 eV and 0.07 eV, respectively. The performance is, nonetheless, similar to that of fic-NEVPT2 on the transition metals. Also, since we used the same orbitals in these fic-NEVPT2 and MR-EOMPT calculations, the setup of the Co and Cr calculations was more similar than the setup of the Thiel benchmark calculations, where CASSCF and NEVPT2 used optimized orbitals for each multiplicity and irrep block.

Furthermore, we included a comparison against the LR-ic-MRCC method [202]. Linear response methods have the advantage that the states of interest need not lie in the CASCI space, which in turn allows smaller active spaces that only capture the necessary amount of static electron correlation. The MR-EOMPT method, as well as fic-NEVPT2, cannot treat states outside of the CASCI space. Hence, a larger active space is required to compare against the LR-ic-MRCC results on higher-lying singlet roots of methylene and hexatriene [202]. We found that the MR-EOMPT method outperforms even the MR-EOMCC method on the methylene singlet roots, which is attributed to fortuitous error cancellation since normally the MR-EOMCC method is more accurate. Conversely, on hexatriene, the iterative MR-EOMCC results are superior to the perturbative results, with MR-EOMPT and NEVPT2 performing similarly.

Overall, we have demonstrated that the accuracy of the MR-EOMPT approach is comparable to established multireference perturbation theories such as state-specific fic-NEVPT2. When benchmarked against its parent method, MR-EOMCC, we found that, despite sometimes even outperforming it, the MR-EOMPT method is less robust, i.e., less precisely reproduces the reference values. This comes at the benefit of lower computational cost and avoided convergence difficulties. The new method is especially useful in situations where dynamic electron correlation significantly mixes the model space functions among themselves, as showcased in the PES of the LiF system. We therefore recommend this new approach in these situations as well as for systems where the high accuracy of the MR-EOMCC method is not strictly required, but perturbational results of NEVPT2-like quality are sufficient.

8.5 Automatic Scaling Reduction for Internally Contracted Multireference Theories

For the first time, to the best of our knowledge, we presented in Ch. 7 an *automated* reduced-scaling scheme that can be applied to *any* theory with high-order densities, completely irrespective of any specifics. Our strategy relies on the fact that index pairs in density matrices can be permuted in an arbitrary fashion, and the permuted density be split with RI into smaller parts and then contracted with other tensors such that many indices are contracted "away." Thus, we can completely avoid the computation of the highest-order densities, whose computation scales steeply with the size of the active space and the number of CSFs in the CASCI space.

Our automated scheme is straightforward to implement, given an existing code generation toolchain that allows for easy manipulation of the working equations on-the-fly. The global search of index permutations and RI split positions is not computationally intensive since the number of (sensible) possibilities is quite limited, even for contractions containing five-body densities in the case of fic-MRCC theory. A minor caveat while processing the reduced working equations is that a termwise optimal factorizer is required, instead of simpler tools that only do a left-to-right, linear factorization. Sometimes, a linear factorizer cannot find the only sensible contraction, and consequently may even leave the contraction "unreduced." We further conclude that the resulting, large CSF intermediates should be stored on disk, for they contain an index that scales with the number of CSFs, and thus exponentially with the size of the active space.

We assessed the performance of the reduced-scaling implementations of fic-MRCI and fic-MRCC theory on two benchmark sets, which are aimed at demonstrating the scaling with the size of the systems as well as the size of the active space. For the scaling of the standard and reduced-scaling variants with respect to system size in a series of linear alkenes using a small, CAS(2,2) active space, we do not observe any meaningful difference in runtime, neither for the preparation steps, which include the computation of the densities, constant intermediates, and the initial guess, nor for the iterative part of the calculations. This conclusion holds for increasing the system size both in terms of adding more carbon and hydrogen atoms, as well as using larger, triple-ζ basis sets; the main reason simply being that neither the computation of the densities nor the contractions involving the densities are time-critical.

For the benchmarks on polyunsaturated alkenes with various active spaces up to a CAS(12,12), substantial performance differences become obvious. Turning first to fic-MRCC theory, we observed a clear win for the reduced-scaling implementation in all cases, no matter the size of the active space. In fact, the reduced-scaling variant was even *faster* per iteration than the standard implementation. Normally, the iterative part is expected to be *slower* than for the standard implementation, since we introduce a scaling dependence on the number of CSFs there, and the computational gains should come entirely from avoiding the computation of the high-order density matrices. This unexpected, but welcome result can, however, be explained by the fact that the contractions involving a five-body density in the standard implementation are

very inefficient on modern CPU / memory architectures, which do not handle high dimensional, multi-GB tensors well. Of course, avoiding the computation of γ_5 further reduces the runtime of the reduced-scaling fic-MRCC implementation.

The reduced-scaling scheme, unfortunately, does not work as well for fic-MRCI theory. For all our polyunsaturated benchmark systems, the standard implementation is always faster than the reduced-scaling implementation. This is a combination of three main factors: First, the construction of the four-body density γ_4 is still relatively fast compared to the subsequent fic-MRCI calculation, due to a highly efficient implementation. Second, the iterative part of the reduced-scaling fic-MRCI calculations, which, again, *is expected to be slower* than that of the standard implementation, incurs an additional factor of n_{iter} compared to the densities that we avoid (which only need to be computed once). Third, the reduced-scaling implementation also has a higher prefactor than the standard version since there are more contractions to evaluate (see Tab. 7.1). In short, we have not reached the truly asymptotic limit for this method. As a redeeming point we do, however, see from an analysis of the limiting contraction in the reduced-scaling case that a crossover should occur for very large active spaces, when the computation of γ_4 becomes completely unfeasible.

These points lead us to the outlook for our reduced-scaling scheme. Obviously, the reduced-scaling fic-MRCI variant requires some further investigation. To this end, we propose using both the fic-MRCC and fic-MRCI with approximate solvers for the CASCI problem, e.g., the DMRG [341, 343–347], FCIQMC [429], or the ICE-CI [438, 439] methods. In these approximate schemes, fewer CSFs are involved, and thus also the computational effort of the CSF intermediates is reduced, as well as the storage space. Conversely, the size of the density matrices is not reduced by these approximate methods, i.e., the standard implementations do not benefit in the same way as the reduced-scaling variants do. Moreover, by going up to active spaces of, e.g., 30 active orbitals, we should be guaranteed to finally be in the asymptotically scaling limit even for fic-MRCI theory. However, this requires that the method to which the reduced-scaling scheme is applied is relatively insensitive to approximate densities, which is not true, e.g., for fic-NEVPT2 theory [361, 362]. A possible solution to this problem was proposed by Kollmar *et al.* [382]. Changes to the reduced-scaling scheme as presented in this thesis are not required to use approximate CASCI solvers.

A Detailed Description of the ORCA-AGE Toolchain

A.1 Overview of ORCA-AGE

To give a broader overview of the ORCA-AGE II toolchain, we show a high-level overview in Fig. A.1. Words in `typewriter` font are names of individual programs, e.g., `cgen`, or of internal functionality such as the `ContractionEngine`. `cgen` is the default tool to generate the working equations, and the `ContractionEngine` is explained in Sec. 4.1.

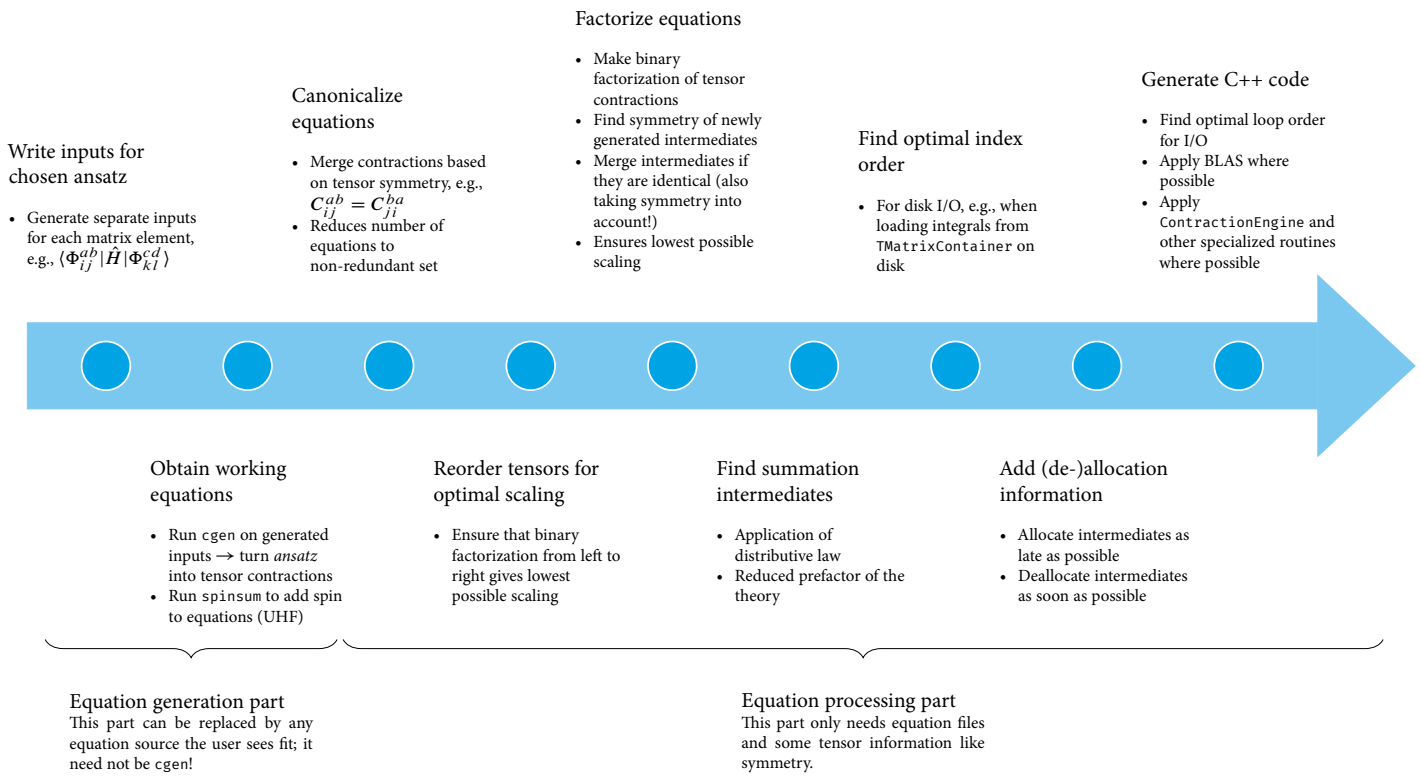

Figure A.1 Overview of all steps involved in generating a functional `orca_autoci` module. The working equations are derived in the first two steps, which are then further processed along the arrow to make the final implementation more performant and computer friendly. The last step is where the final code will be generated.

A.2 Canonicalizer Enhancements

A.2.1 Introduction

The *canonicalizer* is an important part of the ORCA-AGE toolchain, since it removes duplicate contractions by merging the prefactors of identical contractions, e.g.,

$$A \leftarrow \frac{1}{2} BC, \tag{A.1}$$

$$A \leftarrow \quad BC, \tag{A.2}$$

$$\Rightarrow A \leftarrow \frac{3}{2} BC, \tag{A.3}$$

which immediately reduces the prefactor of the overall method being generated.

Given a pairwise comparison function that takes two contractions as arguments, a naive quadratic-time algorithm, $\mathcal{O}(N^2)$, can be implemented that compares each contraction against all others, and merges them if they are found to be identical. Doing so, however, quickly becomes unfeasible for large theories such as fic-MRCC or higher-order CC methods since the algorithm scales quadratically with the number of contractions, and each comparison step is rather expensive since all possible index permutations must be checked.

A.2.2 Current Algorithm

To avoid the quadratically-scaling algorithm outlined in Sec. A.2.1, the canonicalizer first partitions the entire set of contractions into *bins* based on

1. the number of indices on the target tensor,
2. the name of the target tensor,
3. the number of source tensors, and
4. the names of the source tensors.

This process can be viewed as a form of hashing, which leads to a much faster canonicalization since the expensive, explicit pairwise comparison only needs to be done within each of the bins for few (or maybe just a single) contractions, leading to an expected linear time algorithm, $\mathcal{O}(N)$, provided a good hash function with few collisions is used.

Most of the time, the hash function described above performs quite well, leading to speedups exceeding a factor of 200 compared to the naive Python implementation in the first version of ORCA-AGE. Unfortunately, the hash function is imperfect for theories such as CCSDTQ, where many more contractions get hashed to the same bin through the large permutational symmetry of the quadruples amplitudes, effectively making the algorithm run in $\mathcal{O}(N^2)$ complexity again.

A.2.3 Canonicalization in the Literature

Strictly speaking, *canonicalization* does not refer to the process of removing duplicate contractions in the literature,[244, 251] but rather to reducing a given contraction to a certain *canonical form*, i.e., a certain naming and ordering of tensors and indices for a contraction. However, the ultimate goal of transforming a contraction to a canonical form is to ensure that the pairwise comparison of two contractions in canonical form becomes trivial, since (ideally) no index or tensor permutations must be performed to check their equivalence. For ORCA-AGE II, we refrained from implementing such a scheme because it can still lead to ambiguities that need more expensive checks, e.g., when a contraction contains two equivalent $\hat{T}^2$ amplitudes.

It might still be worth thinking about implementing "proper" canonicalization in ORCA-AGE. First, it can still be used to prescreen contractions for equality, and then to use the usual duplicate detection algorithm for in-depth checks. This is expected to be highly useful for higher-order CC methods, since most of the triples, quadruples, … amplitudes will be trivially identical, and thus easy to merge. Second, the contractions should become more human-readable as well. Nevertheless, canonicalization does *not* reduce the quadratic scaling of naive equivalence checks, it just reduces the cost of a single pairwise comparison.

A.2.4 Requirements for Hash Functions

Generally speaking, we are looking for a hash function that maps symmetry-equivalent tensor contractions to a single value, and symmetry-nonequivalent contractions to a different value. In other words, given sets of symmetry-equivalent contractions $X, Y, \ldots$ in the domain of the hash function, which may not be known beforehand,

$$\forall x, x' \in X : x \equiv x', \tag{A.4}$$

$$\forall x \in X, y \in Y : x \not\equiv y, \tag{A.5}$$

we want the hash function $f : X \to A$ to fulfill the following conditions,

$$\forall x, x' \in X : f(x) = f(x'), \tag{A.6}$$

$$\forall x \in X, y \in Y : f(x) \neq f(y). \tag{A.7}$$

Note that the codomain A of the hash function need not be identical to the domain X, the set of all contractions. For example, A might be the set of integer numbers, $A = \mathbb{Z}$, which would be ideal for a computer, or perhaps the set of all strings. Mapping the contractions to a string can simplify computer implementation as well, since hash functions that map strings to integer numbers readily exist in all major programming languages, thus allowing the native hash map implementations to be reused.

If all conditions (A.4)–(A.7) are met, we actually need not perform an expensive pairwise check for contraction equivalence. Instead, we can just merge the prefactors for all contractions from the sets $X, Y, \ldots$ and return the final, single contraction from each set.

If condition (A.6) cannot be met, the resulting theory will have a higher than necessary prefactor overall, since not all equivalent contractions will have been merged, as they ended up in different bins.

In case that Eq. (A.7) is relaxed, each bin will be larger than necessary, and may contain contractions from non-equivalent sets X, Y. Consequently, each contraction that maps to the same bin $f(x) = a \in A$ must be checked against all other contractions in that bin in quadratic time complexity. This is the case with the rather simplistic hash function defined in Sec. A.2.2.

A.2.5 Proposed Enhancements

We suggest enhancing the hash function from Sec. A.2.2 to restore expected linear time for the duplicate removal step. The hash function below effectively uses canonicalization as well to reduce the computational load.

First of all, we need to transform the contractions to a canonical form. This is, in principle, arbitrary. A rather intuitive scheme is presented by Hirata [244]. If this canonicalization fulfills Eqs. (A.6)–(A.7), there is no need to do any further duplicate checks, as explained in Sec. A.2.4. As a possible performance enhancement, especially for many equivalent contractions with quadruples amplitudes, it may be faster to perform the pruning through a hash map. To do so, we could translate the canonicalized contractions to strings, and then using C++'s hash function for strings to check for string equivalence.

However, if Eqs. (A.6)–(A.7) are not fulfilled or, in the case of triple or quadruple excitations, the number of equivalent contractions, or the size of the bins, will become very large. Then, we propose to do a two-staged checking to first remove trivially equivalent contractions, and only then to revert to the expensive pairwise checking.

In the case that no suitable canonicalization can be found for certain equivalent tensors or symmetry-related indices, we suggest mapping them to the same value, e.g.,

$$\sigma_{ij}^{ab} = \sum_{cd} (ac|bd) t_{ij}^{cd} \quad \rightarrow \quad \sigma_{ii}^{aa} = \sum_{cd} (ac|ad) t_{ii}^{cd} . \tag{A.8}$$

While such a canonicalization for sure violates Eq. (A.7), it might be the only solution to do an expensive pairwise checking. Although, as described above, it should already suffice to remove most of the expensive checks through the aforementioned procedure.

Regardless, any hash function that is better than the current state from Sec. A.2.2 will lead to faster times for the canonicalization step, through smaller bins.

A.3 Transpose-Transpose-DGEMM-Transpose Engine

A.3.1 Introduction

In theory, any tensor contraction can be evaluated with BLAS DGEMM calls, given that the tensor axes are properly aligned. Such a scheme, of which TTGT is an example, has been

discussed in several publications [261, 262], including its limitations [262]. Nonetheless, such a scheme will always beat naive loop code, especially when addressing of the tensor axes is highly out-of-order or done through addressing arrays. Consequently, such a general TTGT contraction engine is crucial to ensuring high-performance tensor operations if no other, possibly manually tailored, contraction algorithms exist.

In the following, we will assume that the binary tensor contractions treated by the TTGT algorithm always have the form of

$$Z_{c_1 c_2 \ldots c_m} = \sum_{s_1 \ldots s_k} X_{a_1 \ldots a_n} Y_{b_1 \ldots n_o}. \tag{A.9}$$

In the worst case, no useful statements can be made about the index sets $\{ c_i \}$, $\{ s_i \}$, $\{ a_i \}$, and $\{ b_i \}$. However, this is usually not the case for "well-behaved" quantum-chemical theories. Yet, the algorithm presented in the following section still allows for arbitrary tensor contractions as shown in Eq. (A.9).

A.3.2 Algorithm

We assume that we are given a general tensor contraction as shown in Eq. (A.9) as the input. Our goal is to produce code for this contraction according to the TTGT contraction scheme.

1. Determine the index sets $C = \{ c_i \}$, $S = \{ s_i \}$, $A = \{ a_i \}$, and $B = \{ b_i \}$. From there, determine the *positional* (P), *summed* (S), and *remaining* (R) indices for each source tensor (X and Y, only shown for X):

$$P_A = A \cap C \tag{A.10}$$
$$S_A = A \cap S \tag{A.11}$$
$$R_A = A \setminus (P_A \cup S_A) \tag{A.12}$$

   For the target tensor, we only need to know which indices are found on any of the source tensors (P) and which are not (R).

$$P_C = C \cap (A \cup B) \tag{A.13}$$
$$R_C = C \setminus P_C \tag{A.14}$$

   Summed indices *cannot* appear on the target tensor, as that would be an ill-defined operation.

2. The next step is to take care of *special cases*. The general strategy to tackle these cases is to transpose any of the "misbehaving" tensors such that the "special indices" (defined below) are the slowest-running indices in memory, open `for`-loops over the special indices and address the well-behaved subtensors, and then restart from step 1, while passing on the information about which loops have already been opened.

We note that the actual implementation is not written in a reentrant way, since that could lead to a tensor being reordered twice: once for moving the special indices to the front, and once for aligning the remaining indices. Instead, we simply "delete" the special indices and continue with the procedure given below, which will tell us how to reorder the other indices as well. Then, when generating code, we only need to transpose once to move the special indices to the front, and then add the reordered indices. Finally, in the BLAS call, we only need to set a pointer to the correct slice of the reordered tensor. Nevertheless, we do separate the step of determining and treating the special indices from the remaining reordering procedure.

a) Remaining indices. Any index from the sets R_A, R_B, or R_C must be treated in this fashion.

b) Repeated positional indices, $P_A \cap P_B \neq \varnothing$.

c) Repeated positional indices on a single source or the target tensor, e.g., $n \neq |A|$,

$$Z_{ij} = \sum_k A_{iik} B_{kj}. \tag{A.15}$$

d) Extra summation indices, $S \setminus (S_A \cup S_B)$. These are, however, best optimized away as a simple multiplication. Furthermore, such indices cannot be specified in the ORCA-AGE contraction format.

e) Trace operations, $S_A \triangle S_B = (S_A \setminus S_B) \cup (S_B \setminus S_A)$. An example for this case is

$$Z_{ij} = \sum_{kl} A_{ikl} B_{kj}, \tag{A.16}$$

with $\{ l \} = \{ kl \} \triangle \{ k \}$ defining a trace operation. In this case, we would transpose $A_{ikl} \rightarrow A'_{lik}$, which can be "sliced" to give a matrix $(A'_l)_{ik}$. This must be done for every value of l, hence requiring a `for`-loop. Then, we perform the remaining computations with a normal DGEMM operation,

$$Z_{ij} = \sum_k (A'_l)_{ik} B_{kj} \quad \forall l. \tag{A.17}$$

The other special cases can be treated in complete analogy.

3. Once we have taken care of the special cases in step 2, we have the following guarantees:[1]

$$P_A \cap P_B = \varnothing, \tag{A.18}$$

$$P_C = P_A \cup P_B, \tag{A.19}$$

$$S = S_A = S_B. \tag{A.20}$$

[1] This, of course, refers to the "new" contraction that we get by removing the special indices and using the determined tensor slices.

These guarantees now allow using the DGEMM operation, which can be symbolized by

$$Z_{P_A P_B} = \sum_S X_{P_A S} Y_{P_B S}, \tag{A.21}$$

barring necessary tensor transpositions, which will be determined in the following steps.

4. Now, let us consider the *order* of the indices on the tensors. For a DGEMM operation to be feasible, the summed indices on X and Y must be

   a) in the same order, $s_1, s_2, \ldots, s_k$,

   b) and be situated either entirely at the left or right edge of the tensor, i.e., X_{SP_A} or $X_{P_A S}$.

   We can achieve this by explicitly transposing the tensor axes until the aforementioned conditions are met. This poses an optimization problem since we only want to do the minimum amount of work possible when transposing the tensors.

   In this step, we first determine which tensors *must* be transposed. This is the case for the following cases:

   a) Condition 4b is not met, i.e., the summed and positional indices are *interleaved* on either source tensor.

   b) The positional indices on the target tensor are interleaved, i.e., the target does not fulfill $Z_{P_A P_B}$ or $Z_{P_B P_A}$.

5. In the next step, we determine which tensors *may* need to be transposed. Even if the conditions from item 4 are met, the indices in the sets may still be out of order. Since either tensor with out-of-order indices may be transposed, this can be used to define a graph with edges between out-of-order tensors, e.g., $Z \longleftrightarrow X$ when

$$Z_{ijk} = \sum_l X_{jil} Y_{lk}, \tag{A.22}$$

   as we need to transpose axes $i \leftrightarrow j$ on either Z or X.

   There are only a two possible edge types:

   a) Positional indices on X or Y are not in the same order as on Z, of which Eq. (A.22) is an example (notwithstanding that they may also be interleaved as described in 4b).

   b) The summed indices on X and Y are not in the same order, for example, indices $k \leftrightarrow l$ in

$$Z_{ij} = \sum_{kl} A_{ikl} B_{lkj}. \tag{A.23}$$

We can find which tensors need (or can) be permuted by any *vertex cover* on (the connected parts of) this graph.[2] Note that all permissible vertex covers *must* include the tensors determined in step 4. Once we have found a permissible vertex cover, we estimate its cost by adding up the cost of transposing the tensors, which itself is defined as

$$\mathrm{cost}(X_{a_1 a_2 \ldots a_n}) = \prod_i \dim(a_i). \tag{A.24}$$

After this step, we have determined which tensors are to be transposed to allow a DGEMM call while having the minimal possible transposition cost.

6. After determining which tensors are to be transposed, we need to determine *how* the indices on those tensors need to be rearranged. To this end, we loop over all tensors that need transposing and permute their indices according to the graph edges determined in step 5. After processing a tensor, we delete the graph edges corresponding to the resolved index transpositions to avoid erroneously shuffling the indices on the other tensor.

7. The final step is to generate the actual DGEMM call. This is straightforward as all indices are now in the correct order, and thus we can exactly match the general DGEMM contraction prototype (Eq. A.21).

A.3.3 Details on Repeated "Special" Indices

Repeated "special" indices must, naturally, also be repeated in the transposed contraction. To this end, simply counting them in the original contraction suffices, and we do not need to uniquely identify any axes with "dummy" indices.

For example, let us assume we are given the contraction

$$C_{ab} = \sum_{ik} A_{iaki} B_{kb}. \tag{A.25}$$

After moving the last axis of A to the second position, the contraction reads

$$C_{ab} = \sum_{ik} A'_{iiak} B_{kb} \tag{A.26}$$

$$= \sum_i \left[\sum_k (A'_{ii})_{ak} B_{kb} \right], \tag{A.27}$$

where the outer loop over i must be done separately, and the inner loop over k can be accomplished through a BLAS DGEMM call.

[2] In the highly constrained graphs possible here, it may be faster to just enumerate the distinct cases. However, this will lead to code bloat and will not be as extensible as the graph-based approach.

Note that, despite transposing two axes $A \leftrightarrow A'$, we are not required to *fully* transpose the tensor. The above equations holds when the condition

$$A'_{iiak} = A_{iaki} \tag{A.28}$$

is fulfilled, which is the case for code such as

```
for (int i = 0; i < i_max; ++i)
  for (int a = 0; a < a_max; ++a)
    for (int k = 0; k < k_max; ++k)
      A_p[i][i][a][k] = A[i][a][k][i];
```

Hence, there is no need to use a "dummy" index to uniquely identify *all* axes if we have repeated indices, i.e., we only operate on *parts* of the full tensor. Instead, it is even beneficial not to transpose the entire tensor since that results in less IO overhead.

Had we introduced a dummy index $j \leftrightarrow i$, our condition for tensor equality would have been

$$A'_{ijak} = A_{iakj}, \tag{A.29}$$

which makes obvious that we transposed the last axis to the second position. However, the code required to transpose the tensor accordingly now looks like

```
for (int i = 0; i < i_max; ++i)
  for (int j = 0; j < i_max; ++j)
    for (int a = 0; a < a_max; ++a)
      for (int k = 0; k < k_max; ++k)
        A_p[i][j][a][k] = A[i][a][k][j];
```

which is more expensive than the first variant by a factor of $i_{\max}$!

B Supplementary Information

B.1 Study on the Excited States of Indigo Dyes

B.1.1 Solvation Parameters

In Tab. B.1, we gather the settings that were used for the CPCM solvation model used in Ch. 5. Most of the values were set to the defaults used in the ORCA program package [287], except for the dielectric constants of xylene and TCE (sourced from Jacquemin *et al.* [485] and the refractive index of TCE [536]. The refractive index of xylene was set to ORCA's default for toluene.

Table B.1 Settings used for the dielectric constant ε and refractive index η.

Solvent	$\boldsymbol{\varepsilon}$	$\boldsymbol{\eta}$
Xylene	2.27	1.497
$CHCl_3$	4.9	1.45
TCE	8.2	1.4922
EtOH	24.3	1.361
Benzene	2.28	1.501
CCl_4	2.24	1.466
DMSO	47.2	1.479

B.1.2 Herzberg-Teller Effects from DFT

Computing Herzberg-Teller effects is significantly more expensive than the zeroth-order Franck-Condon results. To assess their magnitude on indigoid systems, we computed their effect on 6,6'-dibromoindigo in DMSO (Fig. B.1).

We note that the spectrum in Fig. B.1 does not contain the energetic shift computed from STEOM as do all other plotted spectra in Ch. 5, but is exclusively computed at the B3LYP-D3(BJ)-CPCM/def2-TZVP level of theory.

Fig. B.1 shows that the first-order Herzberg-Teller effects are minor in that they are more than one order of magnitude smaller than the Franck-Condon results. Additionally, they also do not change the overall shape of the spectrum. Consequently, we did not include them in our results and other computed spectra.

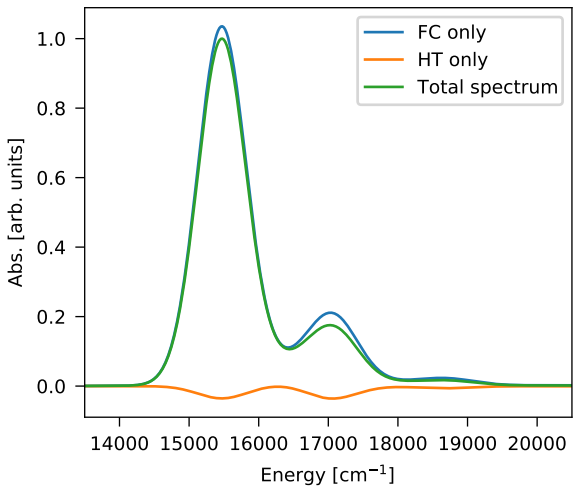

Figure B.1 Contributions of the Franck-Condon and Herzberg-Teller effects to the total spectrum of 6,6'-dibromoindigo in DMSO.

B.1.3 Solvation Effects

In the main body of the text (Tab. 5.3), we presented the effects of a CPCM implicit solvation model on the STEOM transition energies, where CPCM was taken into account for both the geometry optimization as well as in the STEOM calculation and compared that data to all-gas-phase STEOM transition energies.

For completeness' sake, we here also give the corresponding shifts vs. the gas-phase results if we only include the indirect solvation effects on the STEOM calculation (Tab. B.2). The solvated, optimized geometries are identical in these two cases. The direct solvation effects do not have a large impact on the transition energies: on average, the direct effects lead to an additional shift of ≈ -0.03 eV.

B.1.4 Effect of Temperature on the Spectra

In all spectra in Ch. 5, we include the electronic ground state and the first excited electronic state. Since we also include vibrational effects to be able to compare to the experimental 0-0 transition energies, the computed spectra will show several discrete "lines" if the linewidth is artificially set to low values, e.g., 1 cm^{-1}. In case of the spectrum being computed at absolute zero (0 K), the transition lowest in energy corresponds to the 0-0 transition, and all others to "0-n" transitions to a vibrationally excited mode n of the first excited electronic state (Fig. B.2). Note that the following discussion is limited to indigo in DMSO, without loss of generality.

In case we compute the spectrum at a temperature higher than absolute zero, e.g., room temperature (298.15 K), we obtain additional transitions. These are visible in the right spectrum of

Table B.2 Solvation effects compared to gas-phase STEOM calculations with the "fast" CPCM term omitted. Solvents sorted with increasing dielectric constant ε from left to right. All values given in eV.

	CCl4	**Xylene**	**Benzene**	**$CHCl_3$**	**TCE**	**EtOH**	**DMSO**
ε	2.24	2.27	2.28	4.9	8.2	24.3	47.2
Indigo	−0.067	−0.068	−0.068	−0.105	−0.120	−0.128	−0.131
4,4'-Bromoindigo	−0.069	−0.070	−0.070	−0.107	−0.125	−0.139	−0.143
4,4'-Chloroindigo	−0.071	−0.072	−0.072	−0.109	−0.124	−0.139	−0.144
5,5'-Bromoindigo	−0.063	−0.063	−0.063	−0.095	−0.107	−0.119	−0.122
5,5'-Chloroindigo	−0.065	−0.066	−0.066	−0.092	−0.104	−0.116	−0.120
5,5'-Nitroindigo	−0.056	−0.056	−0.048	−0.075	−0.097	−0.111	−0.113
6,6'-Bromoindigo	−0.068	−0.069	−0.069	−0.103	−0.118	−0.128	−0.135
6,6'-Chloroindigo	−0.070	−0.071	−0.071	−0.106	−0.120	−0.135	−0.138
6,6'-Nitroindigo	−0.088	−0.089	−0.089	−0.121	−0.137	−0.153	−0.157
7,7'-Bromoindigo	−0.046	−0.047	−0.047	−0.077	−0.087	−0.098	−0.101
7,7'-Chloroindigo	−0.054	−0.055	−0.055	−0.083	−0.098	−0.106	−0.108
Average	−0.065	−0.066	−0.065	−0.098	−0.112	−0.125	−0.128

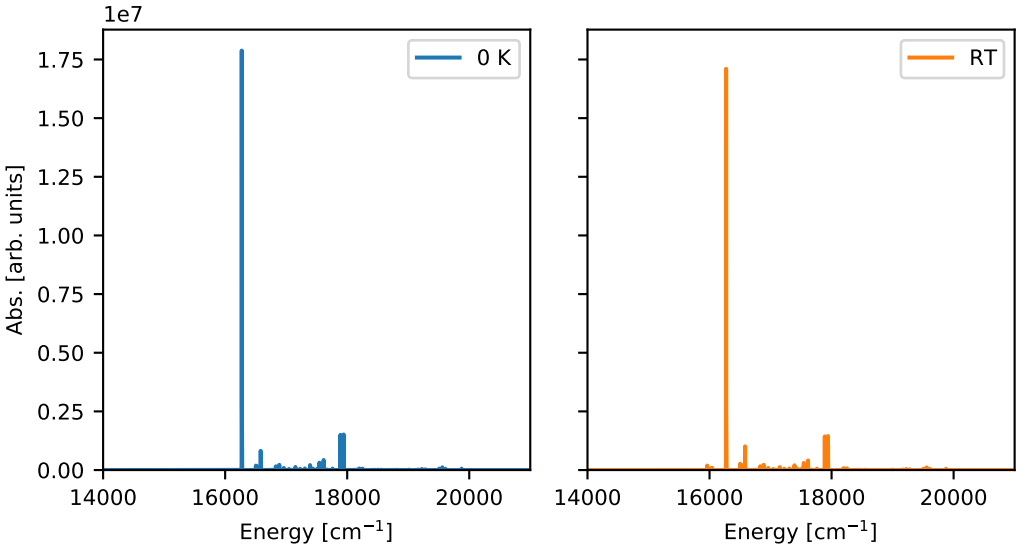

Figure B.2 Computed spectra of indigo in DMSO. The left plot has been computed at absolute zero (0 K), whereas the right spectrum was computed at room temperature (298.15 K).

Fig. B.2 and correspond to n-0 (and possibly n-m) transitions where the molecule is vibrationally excited in its electronic ground state. These transitions are only available at higher-than-zero temperatures since the vibrational modes of the ground state are populated according to the

Boltzmann distribution (see Sec. 2.7). In contrast, we always see the vibrational progression of the final electronic state since the molecule can be electronically and vibrationally excited at the same time (cf. Eq. 2.127).

Notwithstanding, the 0-0 transition remains the by far the strongest and the n-0 transitions are still fairly weak at room temperature since the vibrational ground state of the electronic ground state will still be the most populated one. Moreover, the geometric displacement to the first excited electronic state is small in this case, which also leads to a strong 0-0 transition.

B.1.5 Spectrum of Indigo in $CHCl_3$

In addition to the spectra shown in Ch. 5, we also measured the absorption spectrum of indigo in $CHCl_3$ (Fig. B.3). This commercially available sample is not as pure as the dyes measured in Fig. 5.4, as there are two 'bumps' towards lower energies and a much higher background towards higher energies. Nevertheless, the main peak of indigo is well resolved and coincides almost perfectly with our computed 0-0 transition at 16516.0 cm^{-1}. Furthermore, the left shoulder is very well reproduced after choosing an inhomogeneous linewidth of 700 cm^{-1} for the computed vibrational transitions.

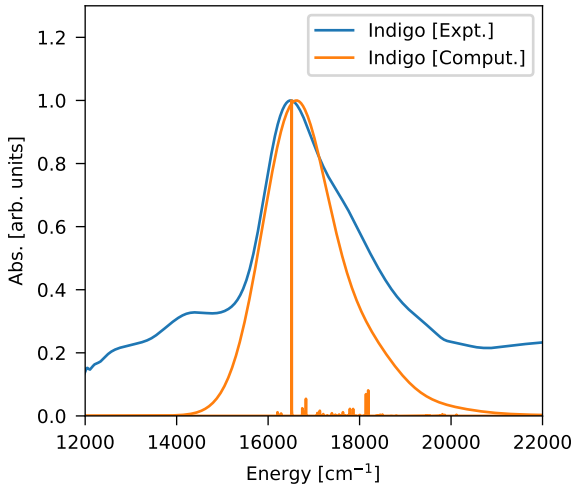

Figure B.3 Normalized experimental and computed spectra of indigo in $CHCl_3$. The unshifted computed spectrum is presented once with an inhomogeneous line broadening of 700 cm^{-1} and once with sticks to indicate the vibrational transitions.

B.1.6 Optimized Geometries and Vibrational Trajectories

The optimized geometries for all substituted indigo dyes can be obtained from `https://doi.org/10.6084/m9.figshare.15183383.v3`. Among these files are also the `xyz` vibrational

trajectories containing the dominant fundamental modes for the spectrum of indigo in DMSO for the spectra discussed in Sec. 5.8.

B.1.7 Experimental Data

The raw experimental data for indigo and genuine Tyrian purple (both measured in $CHCl_3$) is available in plain text format under `https://doi.org/10.6084/m9.figshare.15183383.v3`.

The experimental data has been smoothed slightly with cubic splines prior to plotting using the `scipy.interpolate.splrep` function of the SciPy [537] software package.

B.2 Multireference Equation-of-Motion Perturbation Theory

B.2.1 Explicit Equations for all Perturbative Amplitudes

In this section, we give explicit expressions for the amplitudes used in MR-EOMPT theory that were not fully described in Ch. 6.

In the first similarity transformation of with the cluster operator $\hat{T}$, the amplitudes are computed with the results of the EKT-IP calculation,

$$t_{ab}^{it} = \sum_{\tilde{v}} c_{t\tilde{v}}^{\mathrm{IP}} \frac{(\tilde{v}a|ib)}{\epsilon_{\tilde{v}}^{\mathrm{IP}} + \epsilon_i - (\epsilon_a + \epsilon_b)}, \tag{B.1}$$

and the transformed integral being defined as

$$(\tilde{v}a|ib) = \sum_{t} c_{t\tilde{v}}^{\mathrm{IP}} (ta|ib). \tag{B.2}$$

The $\hat{X}$ and $\hat{D}$ operators describe the spin-flip and spectator excitations, respectively. Hence, their estimates are effectively identical. They require both the EKT-EA and -IP results,

$$x_{au}^{ti} = d_{au}^{it} = \sum_{\tilde{v}\tilde{w}} c_{t\tilde{v}}^{\mathrm{IP}} c_{u\tilde{w}}^{\mathrm{EA}} \frac{(\tilde{v}a|i\tilde{w})}{\epsilon_{\tilde{v}}^{\mathrm{IP}} + \epsilon_i - (\epsilon_a + \epsilon_{\tilde{w}}^{\mathrm{EA}})}, \tag{B.3}$$

and the transformed integrals are defined as

$$(\tilde{v}a|i\tilde{w}) = \sum_{\tilde{v}\tilde{w}} c_{t\tilde{v}}^{\mathrm{IP}} c_{u\tilde{w}}^{\mathrm{EA}} (ta|iu) \tag{B.4}$$

B.2.2 Total Energies for the Diatomic Systems

The total energies for the diatomic systems can be found in Tab. B.3.

B.2.3 Summary of Active Spaces

The active spaces for the systems in Ch. 6 are summarized in Tab. B.4.

Table B.3 Total energies for all the states of the diatomic test systems (E_h).

					MR-EOM	
	State	**ICE-FCI**	**CASSCF**	**NEVPT2**	**-PT**	**-CC**
CH	$^2\Pi$	−38.341204	−38.262875	−38.324470	−38.333238	−38.340435
	$^2\Delta$	−38.229473	−38.148534	−38.213131	−38.222298	−38.228601
	$^2\Sigma^-$	−38.216293	−38.141776	−38.200474	−38.209134	−38.216941
	$^2\Sigma^+$	−38.188757	−38.111817	−38.170641	−38.177167	−38.182974
CN	$^2\Sigma^+$	−92.411165	−92.256484	−92.387161	−92.397334	−92.406678
	$^2\Pi$	−92.357545	−92.199593	−92.332711	−92.341581	−92.352792
	$2\,^2\Sigma^+$	−92.299721	−92.142378	−92.274137	−92.283921	−92.294449
CO	$^1\Sigma^+$	−112.956081	−112.757326	−112.928519	−112.938722	−112.948687
	$^1\Pi$	−112.633516	−112.417543	−112.603116	−112.615749	−112.630480
	$^1\Sigma^-$	−112.577475	−112.373032	−112.547200	−112.557868	−112.570080
	$^1\Delta$	−112.573236	−112.371343	−112.543189	−112.553937	−112.566753
	$^3\Pi$	−112.723349	−112.517232	−112.694774	−112.706374	−112.718998
	$^3\Sigma^+$	−112.639166	−112.436013	−112.608350	−112.616281	−112.627272
	$^3\Delta$	−112.605164	−112.401214	−112.574667	−112.583921	−112.595651
	$^3\Sigma^-$	−112.587212	−112.386993	−112.556752	−112.566144	−112.578439
CO^+	$^2\Sigma^+$	−112.457861	−112.306248	−112.435987	−112.454541	−112.454541
	$^2\Pi$	−112.326110	−112.165787	−112.302484	−112.322395	−112.322395
	$2\,^2\Sigma^+$	−112.243498	−112.080508	−112.218682	−112.239518	−112.239518
N_2	$^1\Sigma^+$	−109.180593	−108.987865	−109.152424	−109.164935	−109.178091
	$^1\Pi$	−108.826143	−108.617307	−108.795425	−108.808434	−108.821585
	$^1\Sigma^-$	−108.797986	−108.588103	−108.766897	−108.781815	−108.795314
	$^1\Delta$	−108.783765	−108.576669	−108.752244	−108.765952	−108.779008
	$^3\Sigma^+$	−108.888730	−108.690236	−108.858394	−108.869729	−108.883982
	$^3\Pi$	−108.877161	−108.666838	−108.847484	−108.858646	−108.870944
	$^3\Delta$	−108.840258	−108.635921	−108.809593	−108.822670	−108.836396
	$^3\Sigma^-$	−108.811041	−108.610542	−108.779487	−108.791419	−108.804613
O_2	$^1\Delta$	−149.813573	−149.540200	−149.776548	−149.791507	−149.809050
	$^1\Sigma^+$	−149.789449	−149.520996	−149.752842	−149.766966	−149.785038
	$^3\Sigma^+$	−149.852988	−149.575157	−149.814595	−149.829455	−149.846299
	$^3\Delta$	−149.619889	−149.349011	−149.581360	−149.596704	−149.615225
OH	$^2\Pi$	−75.492499	−75.344410	−75.471198	−75.479295	−75.490188
	$^2\Sigma^+$	−75.334985	−75.174579	−75.312906	−75.322854	−75.335956

Table B.4 Summary of the active spaces (all calculations) and reference states for the MR-EOMCC and MR-EOMPT calculations. In this condensed notation, "Mult. 4,2 Roots 1,2" would be interpreted as a single quartet and two doublet roots, for example.

System	Active El.	Active Orb.	Mult.	Roots	Project singular?	Notes
Diatomic Systems						
CH	5	5	4,2	1,2	Yes	Full valence
CN	9	8	2	1	Yes	Full valence
CO	10	8	3	2	Yes	Full valence
CO^+	9	8	2	4	Yes	Full valence
N_2	10	8	3	1	Yes	Full valence
O_2	12	8	3	3	No	Full valence
OH	7	5	2	2	Yes	Full valence
Selected Systems from Thiel Benchmark						
Cyclopropene	4	3	3	2	No	
Ethene	2	2	3,1	1,2	No	π-system
Butadiene	4	4	3	2	No	π-system
Benzene	6	6	3,1	3,2	Yes	π-system, `Stol 8e-5`
Pyrrole	6	5	3,1	2,1	Yes	π-system, `Stol 3e-5`
Example Systems						
LiF	2	2	1	2	Yes	Li $2s$ + F p_z
CH_2 S/T Splitting	6	6	3	1	Yes	SA-CASSCF not possible b/c different geometries
CH_2 S/T Splitting	6	6	1	1	Yes	
SiH_2 S/T Splitting	6	6	3	1	Yes	SA-CASSCF not possible b/c different geometries
SiH_2 S/T Splitting	6	6	1	1	Yes	
Co	9	6	4	14	No	
Cr	6	6	7,5	1,6	Yes	50% weight on S states, 50% weight on five roots for D state
Comparison with LR-ic-MRCCSD						
CH_2	6	6	1	7	Yes	Full valence
Hexatriene	6	6	1	3	Yes	π-system

B.2.4 Methylene

The geometry of this test system for comparison to the LR-ic-MRCC method has already been published in Refs. [113, 202]. Note that the distances are given in Bohr radii, a_0, and not in Ångstrom. The basis set is a slightly augmented variant of the cc-pVDZ [538] basis, with an *s*-function added each to carbon and hydrogen with exponents of 0.015 and 0.025, respectively. For ORCA, the input is:

```
! cc-pVDZ Bohrs [ ... ]

%basis
  AddGTO C
    s 1
    1 0.015 1.0
  end
  AddGTO H
    s 1
    1 0.025 1.0
  end
end

[ ... method settings ... ]

* xyz 0 1
  C 0.0  0.0      0.0
  H 0.0  1.644403 1.32213
  H 0.0 -1.644403 1.32213
*
```

In total, this system has 8 electrons and 27 orbitals, which generally allows FCI calculations to be performed on single servers given a recent CPU and sufficient RAM.

B.3 Reduced Scaling for Internally Contracted Multireference Theories

B.3.1 Linear Alk-1-ene Chains

In this section, we give the full system descriptions (Tab. B.5) for the linear alk-1-enes benchmarked in Sec. 7.3.1. We further present the raw, numerical data (Tab. B.6) that was used to discuss the scaling of the reduced-scaling implementations with system size.

Table B.5 System descriptions for linear alk-1-enes. We show the total number of electrons, n_{el}, the number of active electrons, $n_{\mathrm{el,act}}$, as well as the basis set, including the total number of basis functions and the size of the individual orbital ranges.

Geometry	n_{el}	$n_{\mathrm{el,act}}$	Basis	Basis func.	Core	Inactive	Active	Virtual
Ethene	16	2	def2-SVP	48	2	5	2	39
		2	def2-TZVP	86	2	5	2	77
Propene	24	2	def2-SVP	72	3	8	2	59
		2	def2-TZVP	129	3	8	2	116
Butene	32	2	def2-SVP	96	4	11	2	79
		2	def2-TZVP	172	4	11	2	155
Pentene	40	2	def2-SVP	120	5	14	2	99
		2	def2-TZVP	215	5	14	2	194
Hexene	48	2	def2-SVP	144	6	17	2	119
		2	def2-TZVP	258	6	17	2	233
Heptene	56	2	def2-SVP	168	7	20	2	139
		2	def2-TZVP	301	7	20	2	272
Octene	64	2	def2-SVP	192	8	23	2	159
		2	def2-TZVP	344	8	23	2	311
Nonene	72	2	def2-SVP	216	9	26	2	179
		2	def2-TZVP	387	9	26	2	350
Decene	80	2	def2-SVP	240	10	29	2	199
		2	def2-TZVP	430	10	29	2	389
Undecene	88	2	def2-SVP	264	11	32	2	219
		2	def2-TZVP	473	11	32	2	428
Dodecene	96	2	def2-SVP	288	12	35	2	239
		2	def2-TZVP	516	12	35	2	467
Tridecene	104	2	def2-SVP	312	13	38	2	259
		2	def2-TZVP	559	13	38	2	506
Tetradecene	112	2	def2-SVP	336	14	41	2	279
		2	def2-TZVP	602	14	41	2	545
Pentadecene	120	2	def2-SVP	360	15	44	2	299
		2	def2-TZVP	645	15	44	2	584
Hexadecene	128	2	def2-SVP	384	16	47	2	319
		2	def2-TZVP	688	16	47	2	623
Heptadecene	136	2	def2-SVP	408	17	50	2	339
		2	def2-TZVP	731	17	50	2	662
Octadecene	144	2	def2-SVP	432	18	53	2	359
		2	def2-TZVP	774	18	53	2	701

Table B.6 Scaling of standard and reduced-scaling fic-MRCI and fic-MRCC theory on linear alk-1-enes. `CIType` denotes the method being run, `D4Step` indicates the standard (disk) and reduced-scaling versions (efficient), followed by the timings for the constant intermediates, the average time per iteration, the time for all pre-iterative steps, the total time for all iterations as well as the total, overall time for the calculation. All timings given in seconds.

Geometry	Basis	n_{bas}	`CIType`	`D4Step`	Const. inter.	Avg.	Pre-iter.	Total Iter.	Total
ethene	def2-SVP	48	ficmrci	disk	0.1	0.2	0.1	2.0	2.1
	def2-SVP	48	ficmrci	efficient	0.1	0.2	0.1	2.4	2.5
	def2-SVP	48	ficmrcc	disk		5.7	0.0	62.3	62.3
	def2-SVP	48	ficmrcc	efficient		6.2	0.0	68.3	68.3
	def2-TZVP	86	ficmrci	disk	0.4	1.1	0.4	11.0	11.4
	def2-TZVP	86	ficmrci	efficient	0.4	1.1	0.4	11.0	11.4
	def2-TZVP	86	ficmrcc	disk		21.3	0.0	255.7	255.7
	def2-TZVP	86	ficmrcc	efficient		22.4	0.0	268.3	268.3
propene	def2-SVP	72	ficmrci	disk	0.2	0.9	0.2	9.0	9.2
	def2-SVP	72	ficmrci	efficient	0.2	0.9	0.2	9.0	9.2
	def2-SVP	72	ficmrcc	disk		18.6	0.0	205.1	205.1
	def2-SVP	72	ficmrcc	efficient		17.7	0.0	194.6	194.6
	def2-TZVP	129	ficmrci	disk	1.3	6.0	1.3	60.0	61.3
	def2-TZVP	129	ficmrci	efficient	1.4	6.2	1.4	62.0	63.4
	def2-TZVP	129	ficmrcc	disk		104.4	0.0	1148.8	1148.8
	def2-TZVP	129	ficmrcc	efficient		107.2	0.0	1179.5	1179.5
butene	def2-SVP	96	ficmrci	disk	0.5	2.9	0.5	28.7	29.2
	def2-SVP	96	ficmrci	efficient	0.5	2.9	0.5	29.0	29.5
	def2-SVP	96	ficmrcc	disk		48.6	0.0	534.8	534.8
	def2-SVP	96	ficmrcc	efficient		50.8	0.0	559.0	559.0

continued

Table B.6 Scaling results on alkenes (*continued*).

Geometry	Basis	n_{bas}	CIType	D4Step	Const. inter.	Avg.	Pre-iter.	Total Iter.	Total
	def2-TZVP	172	ficmrci	disk	4.1	22.9	4.1	228.6	232.7
	def2-TZVP	172	ficmrci	efficient	4.2	23.1	4.2	231.1	235.3
	def2-TZVP	172	ficmrcc	disk		366.1	0.0	4026.6	4026.6
	def2-TZVP	172	ficmrcc	efficient		372.5	0.0	4097.2	4097.2
pentene	def2-SVP	120	ficmrci	disk	1.0	8.2	1.0	90.2	91.2
	def2-SVP	120	ficmrci	efficient	1.1	8.2	1.1	90.6	91.7
	def2-SVP	120	ficmrcc	disk		118.9	0.0	1426.2	1426.2
	def2-SVP	120	ficmrcc	efficient		123.4	0.0	1480.4	1480.4
	def2-TZVP	215	ficmrci	disk	10.4	70.8	10.4	778.7	789.1
	def2-TZVP	215	ficmrci	efficient	10.7	71.1	10.7	782.0	792.7
	def2-TZVP	215	ficmrcc	disk		1074.8	0.0	12897.7	12897.7
	def2-TZVP	215	ficmrcc	efficient		1092.8	0.0	13113.8	13113.8
hexene	def2-SVP	144	ficmrci	disk	1.9	19.8	1.9	217.9	219.8
	def2-SVP	144	ficmrci	efficient	2.0	19.9	2.0	218.8	220.8
	def2-SVP	144	ficmrcc	disk		267.8	0.0	3213.5	3213.5
	def2-SVP	144	ficmrcc	efficient		261.9	0.0	3143.2	3143.2
	def2-TZVP	258	ficmrci	disk	21.6	184.0	21.7	2024.1	2045.8
	def2-TZVP	258	ficmrci	efficient	21.1	183.6	21.2	2019.6	2040.8
	def2-TZVP	258	ficmrcc	disk		2588.9	0.1	31066.8	31066.9
	def2-TZVP	258	ficmrcc	efficient		2619.1	0.1	31429.6	31429.7
heptene	def2-SVP	168	ficmrci	disk	3.5	41.7	3.5	458.2	461.7
	def2-SVP	168	ficmrci	efficient	3.6	42.2	3.6	464.1	467.7
	def2-SVP	168	ficmrcc	disk		524.1	0.0	6289.0	6289.0

continued

Table B.6 Scaling results on alkenes (*continued*).

Geometry	Basis	n_{bas}	`CIType`	`D4Step`	Const. inter.	Avg.	Pre-iter.	Total Iter.	Total
	def2-SVP	168	ficmrcc	efficient		513.8	0.0	6166.0	6166.0
	def2-TZVP	301	ficmrci	disk	39.4	423.6	39.5	5083.4	5122.9
	def2-TZVP	301	ficmrci	efficient	39.9	425.6	40.0	5107.3	5147.3
	def2-TZVP	301	ficmrcc	disk		5582.1	0.1	66985.7	66985.8
	def2-TZVP	301	ficmrcc	efficient		5624.0	0.1	67488.4	67488.5
octene	def2-SVP	192	ficmrci	disk	8.2	85.5	8.3	1025.9	1034.2
	def2-SVP	192	ficmrci	efficient	8.4	86.4	8.5	1037.0	1045.5
	def2-SVP	192	ficmrcc	disk		1168.9	0.1	14026.7	14026.8
	def2-SVP	192	ficmrcc	efficient		1135.2	0.1	13622.9	13623.0
	def2-TZVP	344	ficmrci	disk	92.2	888.8	92.5	12443.4	12535.9
	def2-TZVP	344	ficmrci	efficient	90.7	888.1	91.0	12433.6	12524.6
	def2-TZVP	344	ficmrcc	disk		11850.9	0.3	142210.5	142210.8
	def2-TZVP	344	ficmrcc	efficient		15064.3	0.4	180771.8	180772.2
nonene	def2-SVP	216	ficmrci	disk	13.3	165.5	13.4	1985.5	1998.9
	def2-SVP	216	ficmrci	efficient	13.5	165.5	13.6	1986.2	1999.8
	def2-SVP	216	ficmrcc	disk		2080.9	0.1	24970.9	24971.0
	def2-SVP	216	ficmrcc	efficient		2076.2	0.1	24914.4	24914.5
	def2-TZVP	387	ficmrci	disk	149.8	1741.9	150.3	26128.3	26278.6
	def2-TZVP	387	ficmrci	efficient	148.8	1740.3	149.3	26104.8	26254.1
	def2-TZVP	387	ficmrcc	disk		25837.6	0.5	310050.6	310051.1
	def2-TZVP	387	ficmrcc	efficient		21316.7	0.6	255799.9	255800.5
decene	def2-SVP	240	ficmrci	disk	20.5	292.5	20.7	3802.8	3823.5
	def2-SVP	240	ficmrci	efficient	20.5	292.6	20.7	3804.4	3825.1

continued

Table B.6 Scaling results on alkenes (*continued*).

Geometry	Basis	n_{bas}	CIType	D4Step	Const. inter.	Avg.	Pre-iter.	Total Iter.	Total
	def2-SVP	240	ficmrcc	disk		3617.7	0.2	43412.0	43412.2
	def2-SVP	240	ficmrcc	efficient		3511.0	0.2	42131.7	42131.9
	def2-TZVP	430	ficmrci	disk	221.8	3285.9	222.5	49287.9	49510.4
	def2-TZVP	430	ficmrci	efficient	229.6	3254.7	230.3	48820.6	49050.9
	def2-TZVP	430	ficmrcc	disk		38569.9	1.7	462838.8	462840.5
	def2-TZVP	430	ficmrcc	efficient		39143.0	1.3	469715.7	469717.0
undecene	def2-SVP	264	ficmrci	disk	29.9	495.1	30.1	6436.0	6466.1
	def2-SVP	264	ficmrci	efficient	41.9	499.7	42.2	6496.1	6538.3
	def2-SVP	264	ficmrcc	disk		5565.5	0.2	66785.6	66785.8
	def2-SVP	264	ficmrcc	efficient		5468.2	0.2	65618.3	65618.5
	def2-TZVP	473	ficmrci	disk	333.4	5759.6	334.4	92154.3	92488.7
	def2-TZVP	473	ficmrci	efficient	335.6	5718.4	336.7	91495.1	91831.8
	def2-TZVP	473	ficmrcc	disk		60499.5	1.1	786493.4	786494.5
	def2-TZVP	473	ficmrcc	efficient		58321.3	1.0	758177.0	758178.0
dodecene	def2-SVP	288	ficmrci	disk	44.5	801.7	44.9	12827.8	12872.7
	def2-SVP	288	ficmrci	efficient	42.4	832.4	42.8	13319.0	13361.8
	def2-SVP	288	ficmrcc	disk		10141.5	0.5	121698.0	121698.5
	def2-SVP	288	ficmrcc	efficient		10368.4	0.5	124420.6	124421.1
	def2-TZVP	516	ficmrcc	disk		99812.4	1.7	1297561.2	1297562.9
	def2-TZVP	516	ficmrcc	efficient		99722.4	1.6	1296391.7	1296393.3
tridecene	def2-SVP	312	ficmrci	disk	57.8	1257.9	58.3	20127.0	20185.3
	def2-SVP	312	ficmrci	efficient	58.1	1261.0	58.6	20175.2	20233.8
	def2-SVP	312	ficmrcc	disk		14576.9	0.6	174923.0	174923.6

continued

Table B.6 Scaling results on alkenes (*continued*).

Geometry	Basis	n_{bas}	CIType	D4Step	Const. inter.	Avg.	Pre-iter.	Total Iter.	Total
	def2-SVP	312	ficmrcc	efficient		12315.4	0.5	147784.8	147785.3
tetradecene	def2-SVP	336	ficmrci	disk	77.0	1963.4	77.7	35340.3	35418.0
	def2-SVP	336	ficmrci	efficient	77.8	1908.2	78.5	34348.0	34426.5
	def2-SVP	336	ficmrcc	disk		17953.5	0.7	215442.0	215442.7
	def2-SVP	336	ficmrcc	efficient		21022.6	0.9	252271.2	252272.1
pentadecene	def2-SVP	360	ficmrci	disk	124.5	2818.4	125.6	53550.5	53676.1
	def2-SVP	360	ficmrci	efficient	110.5	2891.9	111.5	54945.4	55056.9
	def2-SVP	360	ficmrcc	disk		25825.1	1.0	309901.4	309902.4
	def2-SVP	360	ficmrcc	efficient		26147.7	1.0	313772.4	313773.4
hexadecene	def2-SVP	384	ficmrci	disk	133.0	4065.3	134.2	77240.5	77374.7
	def2-SVP	384	ficmrci	efficient	168.7	4017.7	170.3	76335.4	76505.7
	def2-SVP	384	ficmrcc	disk		44927.2	1.2	539126.3	539127.5
	def2-SVP	384	ficmrcc	efficient		43458.4	1.4	521501.1	521502.5
heptadecene	def2-SVP	408	ficmrci	disk	251.3	5742.2	253.3	114844.2	115097.5
	def2-SVP	408	ficmrci	efficient	170.1	5704.3	171.7	114085.8	114257.5
	def2-SVP	408	ficmrcc	disk		48698.0	1.6	584375.9	584377.5
	def2-SVP	408	ficmrcc	efficient		49284.2	1.6	591410.7	591412.3
octadecene	def2-SVP	432	ficmrci	disk	214.4	7957.6	216.3	159152.9	159369.2
	def2-SVP	432	ficmrci	efficient	214.7	7909.8	216.6	158195.5	158412.1
	def2-SVP	432	ficmrcc	disk		67774.5	1.9	813294.5	813296.4
	def2-SVP	432	ficmrcc	efficient		76574.2	2.2	918890.7	918892.9

B.3.2 Polyunsaturated Alkenes

In this section, we give the full system descriptions (Tab. B.7) for the polyunsaturated alkenes benchmarked in Sec. 7.3.2. We further present the raw, numerical data (Tab. B.8) that was used to discuss the scaling of the reduced-scaling implementations with size of the active space.

Table B.7 System descriptions for polyunsaturated alkenes. We show the total number of electrons, n_{el}, the number of active electrons, $n_{\mathrm{el,act}}$, as well as the basis dimension (def2-SVP basis set [463]) and the sizes of the individual orbital ranges.

Geometry	n_{el}	$n_{\mathrm{el,act}}$	**Basis**	**Core**	**Inactive**	**Active**	**Virtual**
Ethene	16	2	48	2	5	2	39
Butadiene	30	2	86	4	10	2	70
	30	4	86	4	9	4	69
Hexatriene	44	2	124	6	15	2	101
	44	4	124	6	14	4	100
	44	6	124	6	13	6	99
Octatetraene	58	2	162	8	20	2	132
	58	4	162	8	19	4	131
	58	6	162	8	18	6	130
	58	8	162	8	17	8	129
Decapentaene	72	2	200	10	25	2	163
	72	4	200	10	24	4	162
	72	6	200	10	23	6	161
	72	8	200	10	22	8	160
	72	10	200	10	21	10	159
Dodecahexaene	86	2	238	12	30	2	194
	86	4	238	12	29	4	193
	86	6	238	12	28	6	192
	86	8	238	12	27	8	191
	86	10	238	12	26	10	190
	86	12	238	12	25	12	189
Tetradecaheptaene	100	2	276	14	35	2	225
	100	4	276	14	34	4	224
	100	6	276	14	33	6	223
	100	8	276	14	32	8	222
	100	10	276	14	31	10	221
	100	12	276	14	30	12	220
	100	14	276	14	29	14	219

Table B.8 Scaling of standard and reduced-scaling fic-MRCI and fic-MRCC theory on polyunsaturated alkenes. `CIType` denotes the method being run, `D4Step` indicates the standard (disk) and reduced-scaling versions (efficient), followed by the timings for the computation of γ_4, γ_5, for the constant intermediates, the average time per iteration, the time for all pre-iterative steps, the total time for all iterations as well as the total, overall time for the calculation. All timings given in seconds.

Geometry	n_{el}	n_{orb}	`CIType`	`D4Step`	γ_4	γ_5	Const. inter.	Avg.	Pre-iter.	Total iter.	Total
ethene	2	2	ficmrci	disk	0.0		0.1	0.2	0.1	2.0	2.1
ethene	2	2	ficmrci	efficient	0.0		0.1	0.2	0.1	2.2	2.3
ethene	2	2	ficmrcc	disk	0.0	0.0		5.6	0.0	61.6	61.6
ethene	2	2	ficmrcc	efficient	0.0			5.7	0.0	62.4	62.4
butadiene	2	2	ficmrci	disk	0.0		0.4	1.8	0.4	21.6	22.0
butadiene	2	2	ficmrci	efficient	0.0		0.4	1.8	0.4	21.6	22.0
butadiene	2	2	ficmrcc	disk	0.0	0.0		32.9	0.0	460.9	460.9
butadiene	2	2	ficmrcc	efficient	0.0			34.8	0.0	487.4	487.4
butadiene	4	4	ficmrci	disk	0.0		1.0	3.2	1.0	32.0	33.0
butadiene	4	4	ficmrci	efficient	0.0		1.0	3.3	1.0	33.0	34.0
butadiene	4	4	ficmrcc	disk	0.0	0.0		75.4	0.0	1131.4	1131.4
butadiene	4	4	ficmrcc	efficient	0.0			78.4	0.0	1175.4	1175.4
hexatriene	2	2	ficmrci	disk	0.0		1.2	9.9	1.2	118.8	120.0
hexatriene	2	2	ficmrci	efficient	0.0		1.2	10.1	1.2	121.0	122.2
hexatriene	2	2	ficmrcc	disk	0.0	0.0		139.7	0.0	2235.7	2235.7
hexatriene	2	2	ficmrcc	efficient	0.0			144.5	0.0	2311.4	2311.4
hexatriene	4	4	ficmrci	disk	0.0		3.0	15.0	3.0	179.7	182.7
hexatriene	4	4	ficmrci	efficient	0.0		3.2	15.3	3.2	183.2	186.4
hexatriene	4	4	ficmrcc	disk	0.0	0.0		294.5	0.0	4711.8	4711.8
hexatriene	4	4	ficmrcc	efficient	0.0			303.8	0.0	4861.3	4861.3
hexatriene	6	6	ficmrci	disk	0.1		8.2	24.6	8.3	246.4	254.7
hexatriene	6	6	ficmrci	efficient	0.0		7.8	25.3	7.9	252.8	260.7
hexatriene	6	6	ficmrcc	disk	0.1	6.7		711.5	6.8	7826.6	7833.4
hexatriene	6	6	ficmrcc	efficient	0.1			711.1	0.1	7822.3	7822.4

continued

Table B.8 Scaling results on polyenes (*continued*).

Geometry	n_{el}	n_{orb}	CIType	D4Step	γ_4	γ_5	Const. inter.	Avg.	Pre-iter.	Total iter.	Total
octatetraene	2	2	ficmrci	disk	0.0		3.6	36.2	3.6	470.0	473.6
octatetraene	2	2	ficmrci	efficient	0.0		3.4	36.3	3.4	471.3	474.7
octatetraene	2	2	ficmrcc	disk	0.0	0.0		440.4	0.0	6165.5	6165.5
octatetraene	2	2	ficmrcc	efficient	0.0			449.7	0.0	6295.6	6295.6
octatetraene	4	4	ficmrci	disk	0.0		10.0	53.1	10.0	690.7	700.7
octatetraene	4	4	ficmrci	efficient	0.0		8.1	52.6	8.1	684.2	692.3
octatetraene	4	4	ficmrcc	disk	0.0	0.0			0.0		0.0
octatetraene	4	4	ficmrcc	efficient	0.0				0.0		0.0
octatetraene	6	6	ficmrci	disk	0.2		19.5	77.9	19.7	1012.8	1032.5
octatetraene	6	6	ficmrci	efficient	0.0		18.5	78.3	18.6	1018.3	1036.9
octatetraene	6	6	ficmrcc	disk	0.1	6.6		1897.0	6.7	26557.8	26564.5
octatetraene	6	6	ficmrcc	efficient	0.2			1923.6	0.2	26930.1	26930.3
octatetraene	8	8	ficmrci	disk	4.6		50.1	120.3	55.6	1323.0	1378.6
octatetraene	8	8	ficmrci	efficient	1.0		45.7	123.8	48.2	1361.3	1409.5
octatetraene	8	8	ficmrcc	disk	4.6	884.5		5469.7	889.9	60166.4	61056.3
octatetraene	8	8	ficmrcc	efficient	4.6			5123.4	5.4	56357.4	56362.8
decapentaene	2	2	ficmrci	disk	0.0		9.8	110.0	9.9	1760.5	1770.4
decapentaene	2	2	ficmrci	efficient	0.0		9.4	109.6	9.5	1753.0	1762.5
decapentaene	2	2	ficmrcc	disk	0.0	0.0		1419.8	0.1	22716.6	22716.7
decapentaene	2	2	ficmrcc	efficient	0.0			1375.2	0.1	22003.5	22003.6
decapentaene	4	4	ficmrci	disk	0.0		23.3	149.4	23.4	2241.4	2264.8
decapentaene	4	4	ficmrci	efficient	0.0		23.5	148.5	23.6	2227.7	2251.3
decapentaene	4	4	ficmrcc	disk	0.0	0.0			0.1		0.1
decapentaene	4	4	ficmrcc	efficient	0.0				0.1		0.1
decapentaene	6	6	ficmrci	disk	0.1		53.7	213.2	54.1	3197.4	3251.5
decapentaene	6	6	ficmrci	efficient	0.0		50.8	204.9	51.2	3073.9	3125.1
decapentaene	6	6	ficmrcc	disk	0.1	6.8		4922.7	7.2	68917.6	68924.8
decapentaene	6	6	ficmrcc	efficient	0.1			4903.4	0.4	68648.2	68648.6

continued

Table B.8 Scaling results on polyenes (*continued*).

Geometry	n_{el}	n_{orb}	`CIType`	`D4Step`	γ_4	γ_5	Const. inter.	Avg.	Pre-iter.	Total iter.	Total
decapentaene	8	8	ficmrci	disk	4.9		110.5	298.5	116.7	3582.3	3699.0
decapentaene	8	8	ficmrci	efficient	0.8		95.2	309.1	97.9	3709.4	3807.3
decapentaene	8	8	ficmrcc	disk	5.0	930.6		11036.8	936.8	165551.8	166488.6
decapentaene	8	8	ficmrcc	efficient	4.9			11075.7	6.1	166135.3	166141.4
decapentaene	10	10	ficmrci	disk	172.0		243.7	429.6	424.3	5584.4	6008.7
decapentaene	10	10	ficmrci	efficient	13.8		285.7	718.5	312.4	9339.9	9652.3
decapentaene	10	10	ficmrcc	disk	175.7	80592.2		39678.7	80776.6	436466.2	517242.8
decapentaene	10	10	ficmrcc	efficient	171.0			36492.8	179.6	401420.5	401600.1
dodecahexaene	6	6	ficmrcc	disk	0.2	6.7		12249.5	7.3	183742.9	183750.2
dodecahexaene	6	6	ficmrcc	efficient	0.2			12191.3	0.6	182868.9	182869.5
dodecahexaene	8	8	ficmrci	disk	4.8		186.1	654.2	192.3	9813.6	10005.9
dodecahexaene	8	8	ficmrci	efficient	0.8		181.1	664.0	184.0	9959.3	10143.3
dodecahexaene	8	8	ficmrcc	disk	4.8	1358.9		23957.6	1365.2	383321.2	384686.4
dodecahexaene	8	8	ficmrcc	efficient	5.1				6.7		6.7
dodecahexaene	10	10	ficmrci	disk	178.1		476.9	1011.5	664.3	16183.7	16848.0
dodecahexaene	10	10	ficmrci	efficient	15.1		510.3	1435.7	538.9	22971.0	23509.9
dodecahexaene	10	10	ficmrcc	disk	176.8	65231.5			65417.4		65417.4
dodecahexaene	10	10	ficmrcc	efficient	177.5				186.7		186.7
dodecahexaene	12	12	ficmrci	disk	6699.9		853.8	1423.3	7685.2	18503.0	26188.2
dodecahexaene	12	12	ficmrci	efficient	463.2		3572.4	10034.8	4187.1	130451.8	134638.9
dodecahexaene	12	12	ficmrcc	disk					0.0		0.0
dodecahexaene	12	12	ficmrcc	efficient	6692.1				6819.2		6819.2
tetradecaheptaene	2	2	ficmrci	disk	0.0				0.4		0.4
tetradecaheptaene	2	2	ficmrci	efficient	0.0				0.3		0.3
tetradecaheptaene	4	4	ficmrci	disk	0.0				0.4		0.4
tetradecaheptaene	4	4	ficmrci	efficient	0.0				0.4		0.4
tetradecaheptaene	6	6	ficmrci	disk	0.1		229.9	1108.7	230.9	22174.0	22404.9
tetradecaheptaene	6	6	ficmrci	efficient	0.0		231.9	1122.2	232.8	22444.9	22677.7

continued

Table B.8 Scaling results on polyenes (*continued*).

Geometry	n_{el}	n_{orb}	CIType	D4Step	γ_4	γ_5	Const. inter.	Avg.	Pre-iter.	Total iter.	Total
tetradecaheptaene	8	8	ficmrci	disk	5.9		309.9	1310.1	317.8	26202.7	26520.5
tetradecaheptaene	8	8	ficmrci	efficient	0.8		303.8	1333.4	307.2	26667.9	26975.1
tetradecaheptaene	10	10	ficmrci	disk	175.0		763.7	1969.6	948.9	39392.0	40340.9
tetradecaheptaene	10	10	ficmrci	efficient	13.9		800.0	2514.3	828.6	50285.3	51113.9
tetradecaheptaene	12	12	ficmrci	disk	6597.8		1188.9	2379.5	7919.4	38071.4	45990.8
tetradecaheptaene	12	12	ficmrci	efficient	463.1		4314.2	13038.5	4929.8	208615.5	213545.3
tetradecaheptaene	14	14	ficmrci	disk	241931.1		2582.4	3519.4	247659.1	56310.2	303969.3
tetradecaheptaene	14	14	ficmrci	efficient	12088.2				15286.3		15286.3

Bibliography

[1] M. Krupička, K. Sivalingam, L. Huntington, A. A. Auer, and F. Neese, J. Comput. Chem. **38**, 1853–1868 (2017).

[2] M. H. Lechner, R. Izsák, M. Nooijen, and F. Neese, Mol. Phys. **119**, e1939185 (2021).

[3] *Web of Science*, (2021) `https://www.webofscience.com/wos/woscc/summary/34b62a42-b61c-4828-929e-e3504a7f0acc-471ca058/relevance/1` (visited on 08/04/2022).

[4] J. A. Pople, J. S. Binkley, and R. Seeger, Int. J. Quantum Chem. **10**, 1–19 (1976).

[5] J. A. Pople, R. Seeger, and R. Krishnan, Int. J. Quantum Chem. **12**, 149–163 (1977).

[6] J. A. Pople, R. Krishnan, H. B. Schlegel, and J. S. Binkley, Int. J. Quantum Chem. **14**, 545–560 (1978).

[7] R. J. Bartlett and G. D. Purvis, Int. J. Quantum Chem. **14**, 561–581 (1978).

[8] R. J. Bartlett and G. D. Purvis, Phys. Scr. **21**, 255–265 (1980).

[9] R. J. Bartlett and G. D. Purvis III, Ann. N. Y. Acad. Sci. **367**, 62–82 (1981).

[10] M. Nooijen and R. J. Bartlett, J. Chem. Phys. **106**, 6441–6448 (1997).

[11] M. Nooijen and R. J. Bartlett, J. Chem. Phys. **106**, 6449–6455 (1997).

[12] M. Nooijen and R. J. Bartlett, J. Chem. Phys. **107**, 6812–6830 (1997).

[13] R. J. Bartlett and J. F. Stanton, "Applications of Post-Hartree-Fock Methods: A Tutorial", in *Reviews in Computational Chemistry*, Vol. 5, edited by K. B. Lipkowitz and D. B. Boyd (John Wiley & Sons, Inc., Hoboken, NJ, USA, 1994), pp. 65–169.

[14] T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (John Wiley & Sons, Ltd, Chichester, 2000).

[15] T. D. Crawford and H. F. Schaefer III, "An Introduction to Coupled Cluster Theory for Computational Chemists", in *Reviews in Computational Chemistry*, Vol. 14 (John Wiley & Sons, Inc., 2000), pp. 33–136.

[16] R. J. Bartlett and M. Musiał, Rev. Mod. Phys. **79**, 291–352 (2007).

[17] I. Shavitt and R. J. Bartlett, *Many-Body Methods in Chemistry and Physics: MBPT and Coupled-Cluster Theory* (Cambridge University Press, Cambridge, 2009).

[18] F. Coester, Nucl. Phys. **7**, 421–424 (1958).

[19] F. Coester and H. Kümmel, Nucl. Phys. **17**, 477–485 (1960).

[20] J. Čížek, J. Chem. Phys. **45**, 4256–4266 (1966).

[21] J. Paldus, J. Čížek, and I. Shavitt, Phys. Rev. A **5**, 50–67 (1972).

[22] J. Čížek, Theor. Chim. Acta **80**, 91–94 (1991).

[23] H. Kümmel, Theor. Chim. Acta **80**, 81–89 (1991).

[24] O. Sinanoğlu, J. Chem. Phys. **36**, 706–717 (1962).

[25] O. Sinanoğlu, J. Chem. Phys. **36**, 3198–3208 (1962).

[26] O. Sinanoğlu and D. F.-t. Tuan, J. Chem. Phys. **38**, 1740–1748 (1963).

[27] G. D. Purvis and R. J. Bartlett, J. Chem. Phys. **76**, 1910–1918 (1982).

[28] J. D. Watts, J. Gauss, and R. J. Bartlett, J. Chem. Phys. **98**, 8718–8733 (1993).

[29] J. Gauss and J. F. Stanton, J. Chem. Phys. **104**, 2574–2583 (1996).

[30] J. Gauss and J. F. Stanton, Chem. Phys. Lett. **276**, 70–77 (1997).

[31] G. E. Scuseria, C. L. Janssen, and H. F. Schaefer III, J. Chem. Phys. **89**, 7382–7387 (1988).

[32] J. F. Stanton, J. Gauss, J. D. Watts, and R. J. Bartlett, J. Chem. Phys. **94**, 4334–4345 (1991).

[33] P. Pulay, S. Saebo, and W. Meyer, J. Chem. Phys. **81**, 1901–1905 (1984).

[34] H.-J. Werner, "Matrix-Formulated Direct Multiconfiguration Self-Consistent Field and Multiconfiguration Reference Configuration-Interaction Methods", in *Advances in Chemical Physics*, Vol. 69 (John Wiley & Sons, Inc., 1987), pp. 1–62.

[35] I. Nebot-Gil, J. Sanchez-Marin, J. L. Heully, J. P. Malrieu, and D. Maynau, Chem. Phys. Lett. **234**, 45–49 (1995).

[36] F. Neese, A. Hansen, and D. G. Liakos, J. Chem. Phys. **131**, 064103 (2009).

[37] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon, Chem. Phys. Lett. **157**, 479–483 (1989).

[38] R. J. Bartlett, J. D. Watts, S. A. Kucharski, and J. Noga, Chemical Physics Letters **165**, 513–522 (1990).

[39] J. F. Stanton, Chem. Phys. Lett. **281**, 130–134 (1997).

[40] J. Řezáč and P. Hobza, J. Chem. Theory Comput. **9**, 2151–2155 (2013).

[41] E. Paulechka and A. Kazakov, J. Phys. Chem. A **121**, 4379–4387 (2017).

[42] T. Kato, Commun. Pure Appl. Math. **10**, 151–177 (1957).

[43] D. Prendergast, M. Nolan, C. Filippi, S. Fahy, and J. C. Greer, J. Chem. Phys. **115**, 1626–1634 (2001).

[44] W. Kutzelnigg, Theor. Chim. Acta **1**, 327–342 (1963).

[45] C. Edmiston and M. Krauss, J. Chem. Phys. **42**, 1119–1120 (1965).

[46] C. Edmiston and M. Krauss, J. Chem. Phys. **45**, 1833–1839 (1966).

[47] H.-J. Werner, C. Köppl, Q. Ma, and M. Schwilk, "Explicitly Correlated Local Electron Correlation Methods", in *Fragmentation*, edited by M. S. Gordon (John Wiley & Sons, Ltd, Chichester, UK, 2017), pp. 1–79.

[48] C. Hampel and H.-J. Werner, J. Chem. Phys. **104**, 6286–6297 (1996).

[49] M. Schütz and H.-J. Werner, J. Chem. Phys. **114**, 661–681 (2001).

[50] M. Schütz, Phys. Chem. Chem. Phys. **4**, 3941–3947 (2002).

[51] M. Schütz and F. R. Manby, Phys. Chem. Chem. Phys. **5**, 3349–3358 (2003).

[52] F. Neese, F. Wennmohs, and A. Hansen, J. Chem. Phys. **130**, 114108 (2009).

[53] P. Pinski, C. Riplinger, E. F. Valeev, and F. Neese, J. Chem. Phys. **143**, 034108 (2015).

[54] C. Riplinger, P. Pinski, U. Becker, E. F. Valeev, and F. Neese, J. Chem. Phys. **144**, 024109 (2016).

[55] Y. Guo, K. Sivalingam, E. F. Valeev, and F. Neese, J. Chem. Phys. **144**, 094111 (2016).

[56] F. Pavošević, P. Pinski, C. Riplinger, F. Neese, and E. F. Valeev, J. Chem. Phys. **144**, 144109 (2016).

[57] W. Kutzelnigg, Theor. Chim. Acta **68**, 445–469 (1985).

[58] J. Noga, W. Kutzelnigg, and W. Klopper, Chem. Phys. Lett. **199**, 497–504 (1992).

[59] J. Noga and W. Kutzelnigg, J. Chem. Phys. **101**, 7738–7762 (1994).

[60] T. Shiozaki, M. Kamiya, S. Hirata, and E. F. Valeev, Phys. Chem. Chem. Phys. **10**, 3358–3370 (2008).

[61] A. J. May and F. R. Manby, J. Chem. Phys. **121**, 4479–4485 (2004).

[62] S. Ten-no, Chem. Phys. Lett. **398**, 56–61 (2004).

[63] S. Ten-no, J. Chem. Phys. **121**, 117–129 (2004).

[64] A. Köhn, G. W. Richings, and D. P. Tew, J. Chem. Phys. **129**, 201103 (2008).

[65] J. Noga, S. Kedžuch, J. Šimunek, and S. Ten-no, J. Chem. Phys. **128**, 174103 (2008).

[66] D. G. Liakos, R. Izsák, E. F. Valeev, and F. Neese, Mol. Phys. **111**, 2653–2662 (2013).

[67] L. Kong, F. A. Bischoff, and E. F. Valeev, Chem. Rev. **112**, 75–107 (2012).

[68] Q. Ma and H.-J. Werner, Wiley Interdiscip. Rev. Comput. Mol. Sci. **8**, e1371 (2018).

[69] K. Sneskov and O. Christiansen, Wiley Interdiscip. Rev. Comput. Mol. Sci. **2**, 566–584 (2012).

[70] R. Izsák, Wiley Interdiscip. Rev. Comput. Mol. Sci. **10**, e1445 (2020).

[71] R. Izsák, Int. J. Quantum Chem. **121**, e26327 (2021).

[72] A. I. Krylov, Annu. Rev. Phys. Chem. **59**, 433–462 (2008).

[73] R. J. Bartlett, Wiley Interdiscip. Rev. Comput. Mol. Sci. **2**, 126–138 (2011).

[74] K. Emrich, Nucl. Phys. A **351**, 379–396 (1981).

[75] J. F. Stanton and R. J. Bartlett, J. Chem. Phys. **98**, 7029–7039 (1993).

[76] S. V. Levchenko and A. I. Krylov, J. Chem. Phys. **120**, 175–185 (2004).

[77] M. Nooijen and J. G. Snijders, Int. J. Quantum Chem. **48**, 15–48 (1993).

[78] J. F. Stanton and J. Gauss, J. Chem. Phys. **101**, 8938–8944 (1994).

[79] K. W. Sattelmeyer, H. F. Schaefer III, and J. F. Stanton, Chem. Phys. Lett. **378**, 42–46 (2003).

[80] M. Nooijen and R. J. Bartlett, J. Chem. Phys. **102**, 3629–3647 (1995).

[81] M. Nooijen and R. J. Bartlett, J. Chem. Phys. **102**, 6735–6756 (1995).

[82] S. Gulania, T.-C. Jagau, and A. I. Krylov, Faraday Discuss. **217**, 514–532 (2019).

[83] J. Schirmer, Phys. Rev. A **26**, 2395–2416 (1982).

[84] O. Christiansen, H. Koch, and P. Jørgensen, Chem. Phys. Lett. **243**, 409–418 (1995).

[85] A. B. Trofimov and J. Schirmer, J. Phys. B: At. Mol. Opt. Phys. **28**, 2299–2324 (1995).

[86] J. F. Stanton and J. Gauss, J. Chem. Phys. **103**, 1064–1076 (1995).

[87] M. Nooijen and J. G. Snijders, J. Chem. Phys. **102**, 1681–1688 (1995).

[88] M. Nooijen, Spectrochim. Acta A **55**, 539–559 (1999).

[89] B. Helmich and C. Hättig, J. Chem. Phys. **135**, 214106 (2011).

[90] M. S. Frank and C. Hättig, J. Chem. Phys. **148**, 134102 (2018).

[91] C. Peng, M. C. Clement, and E. F. Valeev, J. Chem. Theory Comput. **14**, 5597–5607 (2018).

[92] A. K. Dutta, F. Neese, and R. Izsák, J. Chem. Phys. **145**, 034102 (2016).

[93] A. K. Dutta, M. Saitow, C. Riplinger, F. Neese, and R. Izsák, J. Chem. Phys. **148**, 244101 (2018).

[94] S. Haldar, C. Riplinger, B. Demoulin, F. Neese, R. Izsak, and A. K. Dutta, J. Chem. Theory Comput. **15**, 2265–2277 (2019).

[95] A. K. Dutta, M. Saitow, B. Demoulin, F. Neese, and R. Izsák, J. Chem. Phys. **150**, 164123 (2019).

[96] S. Haldar and A. K. Dutta, J. Phys. Chem. A **124**, 3947–3962 (2020).

[97] A. K. Dutta, M. Nooijen, F. Neese, and R. Izsák, J. Chem. Theory Comput. **14**, 72–91 (2018).

[98] C. E. Schulz, A. K. Dutta, R. Izsák, and D. A. Pantazis, J. Comput. Chem. **39**, 2439–2451 (2018).

[99] A. Dittmer, R. Izsák, F. Neese, and D. Maganas, Inorg. Chem. **58**, 9303–9315 (2019).

[100] R. Berraud-Pache, F. Neese, G. Bistoni, and R. Izsák, J. Phys. Chem. Lett. **10**, 4822–4828 (2019).

[101] A. Sirohiwal, R. Berraud-Pache, F. Neese, R. Izsák, and D. A. Pantazis, J. Phys. Chem. B **124**, 8761–8771 (2020).

[102] R. Berraud-Pache, E. Santamaría-Aranda, B. de Souza, G. Bistoni, F. Neese, D. Sampedro, and R. Izsák, Chem. Sci. **12**, 2916–2924 (2021).

[103] R. Berraud-Pache, F. Neese, G. Bistoni, and R. Izsák, J. Chem. Theory Comput. **16**, 564–575 (2020).

[104] A. Sirohiwal, F. Neese, and D. A. Pantazis, J. Am. Chem. Soc. **142**, 18174–18190 (2020).

[105] A. Sirohiwal, F. Neese, and D. A. Pantazis, J. Chem. Theory Comput. **17**, 1858–1873 (2021).

[106] H. J. Monkhorst, Int. J. Quantum Chem. **12**, 421–432 (1977).

[107] E. Dalgaard and H. J. Monkhorst, Phys. Rev. A **28**, 1217–1222 (1983).

[108] H. Sekino and R. J. Bartlett, Int. J. Quantum Chem. **26**, 255–265 (1984).

[109] H. Koch and P. Jørgensen, J. Chem. Phys. **93**, 3333–3344 (1990).

[110] H. Koch, H. J. A. Jensen, P. Jørgensen, and T. Helgaker, J. Chem. Phys. **93**, 3345–3350 (1990).

[111] O. Christiansen, H. Koch, and P. Jørgensen, J. Chem. Phys. **103**, 7429–7441 (1995).

[112] H. Koch, O. Christiansen, P. Jørgensen, A. M. Sanchez de Méras, and T. Helgaker, J. Chem. Phys. **106**, 1808–1818 (1997).

[113] H. Koch, O. Christiansen, P. Jørgensen, and J. Olsen, Chem. Phys. Lett. **244**, 75–82 (1995).

[114] O. Christiansen, H. Koch, P. Jørgensen, and J. Olsen, Chem. Phys. Lett. **256**, 185–194 (1996).

[115] P. Cronstrand, O. Christiansen, P. Norman, and H. Ågren, Phys. Chem. Chem. Phys. **2**, 5357–5363 (2000).

[116] H. Larsen, J. Olsen, P. Jørgensen, and O. Christiansen, J. Chem. Phys. **113**, 6677–6686 (2000).

[117] P. Norman, Phys. Chem. Chem. Phys. **13**, 20519–20535 (2011).

[118] P. Norman, K. Ruud, and T. Saue, *Principles and Practices of Molecular Properties: Theory, Modeling and Simulations* (John Wiley & Sons, Ltd, Chichester, UK, 2018).

[119] T. D. Crawford, A. Kumar, A. P. Bazanté, and R. Di Remigio, Wiley Interdiscip. Rev. Comput. Mol. Sci. **9**, e1406 (2019).

[120] W. Meyer, "Configuration Expansion by Means of Pseudonatural Orbitals", in *Methods of Electronic Structure Theory*, Vol. 3, edited by H. F. Schaefer III (Springer Science+Business Media, New York, 1977), pp. 413–446.

[121] P. E. M. Siegbahn, Int. J. Quantum Chem. **18**, 1229–1242 (1980).

[122] D. Mukherjee, R. K. Moitra, and A. Mukhopadhyay, Pramana **4**, 247–263 (1975).

[123] C. F. Bender and E. R. Davidson, Phys. Rev. **183**, 23–30 (1969).

[124] C. F. Bender and E. R. Davidson, J. Phys. Chem. **70**, 2675–2685 (1966).

[125] R. Ahlrichs and F. Driessler, Theor. Chim. Acta **36**, 275–287 (1975).

[126] B. O. Roos and P. E. M. Siegbahn, Int. J. Quantum Chem. **17**, 485–500 (1980).

[127] H. Dachsel, R. J. Harrison, and D. A. Dixon, J. Phys. Chem. A **103**, 152–155 (1999).

[128] A. D. Powell, N. S. Dattani, R. F. K. Spada, F. B. C. Machado, H. Lischka, and R. Dawes, J. Chem. Phys. **147**, 094306 (2017).

[129] R. J. Buenker and S. D. Peyerimhoff, Theor. Chim. Acta **35**, 33–58 (1974).

[130] P. E. M. Siegbahn, J. Chem. Phys. **72**, 1647–1656 (1980).

[131] R. J. Gdanitz and R. Ahlrichs, Chem. Phys. Lett. **143**, 413–420 (1988).

[132] R. J. Gdanitz, Int. J. Quantum Chem. **85**, 281–300 (2001).

[133] P. G. Szalay and R. J. Bartlett, Chem. Phys. Lett. **214**, 481–488 (1993).

[134] L. Füsti-Molnár and P. G. Szalay, J. Phys. Chem. **100**, 6288–6297 (1996).

[135] P. G. Szalay and R. J. Bartlett, J. Chem. Phys. **103**, 3600–3612 (1995).

[136] P. G. Szalay, Chem. Phys. **349**, 121–125 (2008).

[137] P. J. A. Ruttink, J. H. van Lenthe, R. Zwaans, and G. C. Groenenboom, J. Chem. Phys. **94**, 7212–7220 (1991).

[138] P. J. A. Ruttink, J. H. Van Lenthe, and P. Todorov, Mol. Phys. **103**, 2497–2506 (2005).

[139] D. I. Lyakh, M. Musiał, V. F. Lotrich, and R. J. Bartlett, Chem. Rev. **112**, 182–243 (2012).

[140] S. R. Langhoff and E. R. Davidson, Int. J. Quantum Chem. **8**, 61–72 (1974).

[141] M. Kállay, P. G. Szalay, and P. R. Surján, J. Chem. Phys. **117**, 980–990 (2002).

[142] T. Müller, M. Dallos, H. Lischka, Z. Dubrovay, and P. G. Szalay, Theor. Chem. Acc. **105**, 227–243 (2001).

[143] H.-J. Werner and P. J. Knowles, J. Chem. Phys. **89**, 5803–5814 (1988).

[144] B. Jeziorski and H. J. Monkhorst, Phys. Rev. A **24**, 1668–1681 (1981).

[145] B. Jeziorski, Mol. Phys. **108**, 3043–3054 (2010).

[146] A. Köhn, M. Hanauer, L. A. Mück, T.-C. Jagau, and J. Gauss, Wiley Interdiscip. Rev. Comput. Mol. Sci. **3**, 176–197 (2012).

[147] U. S. Mahapatra, B. Datta, and D. Mukherjee, Mol. Phys. **94**, 157–171 (1998).

[148] U. S. Mahapatra, B. Datta, B. Bandyopadhyay, and D. Mukherjee, Adv. Quantum Chem. **30**, 163–193 (1998).

[149] U. S. Mahapatra, B. Datta, and D. Mukherjee, J. Chem. Phys. **110**, 6171–6188 (1999).

[150] E. Prochnow, F. A. Evangelista, H. F. Schaefer III, W. D. Allen, and J. Gauss, J. Chem. Phys. **131**, 064109 (2009).

[151] T.-C. Jagau, E. Prochnow, F. A. Evangelista, and J. Gauss, J. Chem. Phys. **132**, 144110 (2010).

[152] T.-C. Jagau and J. Gauss, J. Chem. Phys. **137**, 044116 (2012).

[153] T.-C. Jagau and J. Gauss, J. Chem. Phys. **137**, 044115 (2012).

[154] O. Demel, J. Pittner, and F. Neese, J. Chem. Theory Comput. **11**, 3104–3114 (2015).

[155] J. Brabec, J. Lang, M. Saitow, J. Pittner, F. Neese, and O. Demel, J. Chem. Theory Comput. **14**, 1370–1382 (2018).

[156] L. Kong, Int. J. Quantum Chem. **109**, 441–447 (2009).

[157] J. Mášik and I. Hubač, Adv. Quantum Chem. **31**, 75–104 (1998).

[158] J. Pittner, P. Nachtigall, P. Čársky, J. Mášik, and I. Hubač, J. Chem. Phys. **110**, 10275–10282 (1999).

[159] J. Mášik and I. Hubač, "Multireference Brillouin—Wigner Coupled-Cluster Theory: Hilbert Space Approach", in *Quantum Systems in Chemistry and Physics. Trends in Methods and Applications*, Vol. 16, edited by R. McWeeny, J. Maruani, Y. G. Smeyers, and S. Wilson, red. by W. N. Lipscomb and J. Maruani (Springer Netherlands, Dordrecht, 1997), pp. 283–308.

[160] U. S. Mahapatra and S. Chattopadhyay, J. Chem. Phys. **133**, 074102 (2010).

[161] E. Giner, C. Angeli, Y. Garniron, A. Scemama, and J.-P. Malrieu, J. Chem. Phys. **146**, 224108 (2017).

[162] M. Hanrath, J. Chem. Phys. **123**, 084102 (2005).

[163] M. Hanrath, Theor. Chem. Acc. **121**, 187–195 (2008).

[164] F. A. Evangelista and J. Gauss, J. Chem. Phys. **133**, 044101 (2010).

[165] L. Kong, Int. J. Quantum Chem. **110**, 2603–2613 (2010).

[166] S. Hirata, Theor. Chem. Acc. **116**, 2–17 (2006).

[167] K. Andersson, P.-Å. Malmqvist, B. O. Roos, A. J. Sadlej, and K. Wolinski, J. Phys. Chem. **94**, 5483–5488 (1990).

[168] K. Andersson, P.-Å. Malmqvist, and B. O. Roos, J. Chem. Phys. **96**, 1218–1226 (1992).

[169] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J.-P. Malrieu, J. Chem. Phys. **114**, 10252–10264 (2001).

[170] C. Angeli, R. Cimiraglia, and J.-P. Malrieu, Chem. Phys. Lett. **350**, 297–305 (2001).

[171] C. Angeli, R. Cimiraglia, and J.-P. Malrieu, J. Chem. Phys. **117**, 9138–9153 (2002).

[172] Y. Guo, K. Sivalingam, E. F. Valeev, and F. Neese, J. Chem. Phys. **147**, 064110 (2017).

[173] J. W. Park, J. Chem. Theory Comput. **15**, 5417–5425 (2019).

[174] H.-J. Werner and E.-A. Reinsch, J. Chem. Phys. **76**, 3144–3156 (1982).

[175] C. Angeli, R. Cimiraglia, and M. Pastore, Mol. Phys. **110**, 2963–2968 (2012).

[176] K. Sivalingam, M. Krupička, A. A. Auer, and F. Neese, J. Chem. Phys. **145**, 054104 (2016).

[177] M. Saitow, Y. Kurashige, and T. Yanai, J. Chem. Phys. **139**, 044118 (2013).

[178] M. Saitow, Y. Kurashige, and T. Yanai, J. Chem. Theory Comput. **11**, 5120–5131 (2015).

[179] T. Shiozaki, G. Knizia, and H.-J. Werner, J. Chem. Phys. **134**, 034113 (2011).

[180] A. Berning, M. Schweizer, H.-J. Werner, P. J. Knowles, and P. Palmieri, Mol. Phys. **98**, 1823–1833 (2000).

[181] K. R. Shamasundar, G. Knizia, and H.-J. Werner, J. Chem. Phys. **135**, 054101 (2011).

[182] P. G. Szalay, T. Müller, G. Gidofalvi, H. Lischka, and R. Shepard, Chem. Rev. **112**, 108–181 (2012).

[183] P. Fulde and H. Stoll, J. Chem. Phys. **97**, 4185–4187 (1992).

[184] R. Fink and V. Staemmler, Theor. Chim. Acta **87**, 129–145 (1993).

[185] J. A. Black and A. Köhn, J. Chem. Phys. **150**, 194107 (2019).

[186] M. A. Haque and D. Mukherjee, J. Chem. Phys. **80**, 5058–5069 (1984).

[187] U. Kaldor, Theor. Chim. Acta **80**, 427–439 (1991).

[188] I. Lindgren, Int. J. Quantum Chem. **14**, 33–58 (1978).

[189] L. Z. Stolarczyk and H. J. Monkhorst, Phys. Rev. A **32**, 725–742 (1985).

[190] A. Banerjee and J. Simons, Int. J. Quantum Chem. **19**, 207–216 (1981).

[191] A. Banerjee and J. Simons, J. Chem. Phys. **76**, 4548–4559 (1982).

[192] A. Banerjee and J. Simons, Chem. Phys. **81**, 297–302 (1983).

[193] D. Mukherjee, R. K. Moitra, and A. Mukhopadhyay, Mol. Phys. **30**, 1861–1888 (1975).

[194] F. A. Evangelista and J. Gauss, J. Chem. Phys. **134**, 114102 (2011).

[195] M. Hanauer and A. Köhn, J. Chem. Phys. **134**, 204111 (2011).

[196] I. Lindgren and D. Mukherjee, Phys. Rep. **151**, 93–127 (1987).

[197] W. D. Laidig, P. Saxe, and R. J. Bartlett, J. Chem. Phys. **86**, 887–907 (1987).

[198] D. Mukherjee, "A Coupled Cluster Approach to the Electron Correlation Problem Using a Correlated Reference State", in *Recent Progress in Many-Body Theories*, edited by E. Schachinger, H. Mitter, and H. Sormann (Springer US, Boston, MA, 1995), pp. 127–133.

[199] M. Nooijen and V. Lotrich, J. Mol. Struct.: THEOCHEM **547**, 253–267 (2001).

[200] M. Hanauer and A. Köhn, J. Chem. Phys. **136**, 204107 (2012).

[201] W. Liu, M. Hanauer, and A. Köhn, Chem. Phys. Lett. **565**, 122–127 (2013).

[202] P. K. Samanta, D. Mukherjee, M. Hanauer, and A. Köhn, J. Chem. Phys. **140**, 134108 (2014).

[203] P. K. Samanta and A. Köhn, J. Chem. Phys. **149**, 064101 (2018).

[204] F. Lipparini, T. Kirsch, A. Köhn, and J. Gauss, J. Chem. Theory Comput. **13**, 3171–3184 (2017).

[205] Y. A. Aoto and A. Köhn, J. Chem. Phys. **144**, 074103 (2016).

[206] D. J. Coughtrie, R. Giereth, D. Kats, H.-J. Werner, and A. Köhn, J. Chem. Theory Comput. **14**, 693–709 (2018).

[207] Y. A. Aoto, A. Bargholz, D. Kats, H.-J. Werner, and A. Köhn, J. Chem. Theory Comput. **15**, 2291–2305 (2019).

[208] O. Demel, K. R. Shamasundar, L. Kong, and M. Nooijen, J. Phys. Chem. A **112**, 11895–11902 (2008).

[209] L. Kong, K. R. Shamasundar, O. Demel, and M. Nooijen, J. Chem. Phys. **130**, 114101 (2009).

[210] D. Datta, L. Kong, and M. Nooijen, J. Chem. Phys. **134**, 214116 (2011).

[211] D. Datta and M. Nooijen, J. Chem. Phys. **137**, 204107 (2012).

[212] O. Demel, D. Datta, and M. Nooijen, J. Chem. Phys. **138**, 134108 (2013).

[213] M. Nooijen, O. Demel, D. Datta, L. Kong, K. R. Shamasundar, V. Lotrich, L. M. Huntington, and F. Neese, J. Chem. Phys. **140**, 081102 (2014).

[214] F. A. Evangelista, M. Hanauer, A. Köhn, and J. Gauss, J. Chem. Phys. **136**, 204108 (2012).

[215] L. M. J. Huntington, O. Demel, and M. Nooijen, J. Chem. Theory Comput. **12**, 114–132 (2016).

[216] Z. Liu, O. Demel, and M. Nooijen, J. Mol. Spectrosc. **311**, 54–63 (2015).

[217] L. M. J. Huntington and M. Nooijen, J. Chem. Phys. **142**, 194111 (2015).

[218] D. Maganas, J. K. Kowalska, M. Nooijen, S. DeBeer, and F. Neese, J. Chem. Phys. **150**, 104106 (2019).

[219] W. Kutzelnigg and D. Mukherjee, J. Chem. Phys. **107**, 432–449 (1997).

[220] D. Mukherjee, Chem. Phys. Lett. **274**, 561–566 (1997).

[221] M. Nooijen and R. J. Bartlett, J. Chem. Phys. **104**, 2652–2668 (1996).

[222] W. Kutzelnigg, K. R. Shamasundar, and D. Mukherjee, Mol. Phys. **108**, 433–451 (2010).

[223] K. R. Shamasundar, J. Chem. Phys. **131**, 174109 (2009).

[224] M. R. Hoffmann and J. Simons, Chem. Phys. Lett. **142**, 451–454 (1987).

[225] M. R. Hoffmann and J. Simons, J. Chem. Phys. **88**, 993–1002 (1988).

[226] T. Yanai and G. K.-L. Chan, J. Chem. Phys. **124**, 194106 (2006).

[227] T. Yanai and G. K.-L. Chan, J. Chem. Phys. **127**, 104107 (2007).

[228] E. Neuscamman, T. Yanai, and G. K.-L. Chan, J. Chem. Phys. **130**, 169901 (2009).

[229] E. Neuscamman, T. Yanai, and G. K.-L. Chan, J. Chem. Phys. **130**, 124102 (2009).

[230] E. Neuscamman, T. Yanai, and G. K.-L. Chan, Int. Rev. Phys. Chem. **29**, 231–271 (2010).

[231] F. A. Evangelista and J. Gauss, Chem. Phys. **401**, 27–35 (2012).

[232] F. A. Evangelista, J. Chem. Phys. **141**, 054109 (2014).

[233] C. Li and F. A. Evangelista, J. Chem. Theory Comput. **11**, 2097–2108 (2015).

[234] S. Li, J. Chem. Phys. **120**, 5017–5026 (2004).

[235] T. Fang and S. Li, J. Chem. Phys. **127**, 204108 (2007).

[236] T. Fang, J. Shen, and S. Li, J. Chem. Phys. **128**, 224107 (2008).

[237] J. Shen, T. Fang, S. Li, and Y. Jiang, J. Phys. Chem. A **112**, 12518–12525 (2008).

[238] R. E. Watson, Phys. Rev. **118**, 1036–1045 (1960).

[239] R. E. Watson, Phys. Rev. **119**, 1934–1939 (1960).

[240] M. R. Silva-Junior, S. P. Sauer, M. Schreiber, and W. Thiel, Mol. Phys. **108**, 453–465 (2010).

[241] C. Riplinger and F. Neese, J. Chem. Phys. **138**, 034106 (2013).

[242] R. R. Schaller, IEEE Spectr. **34**, 52–59 (1997).

[243] T. Yanai, Y. Kurashige, M. Saitow, J. Chalupský, R. Lindh, and P.-Å. Malmqvist, Mol. Phys. **115**, 2077–2085 (2017).

[244] S. Hirata, J. Phys. Chem. A **107**, 9887–9897 (2003).

[245] A. A. Auer, G. Baumgartner, D. E. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. Harrison, S. Krishnamoorthy, S. Krishnan, C.-C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, and A. Sibiryakov, Mol. Phys. **104**, 211–228 (2006).

[246] M. K. MacLeod and T. Shiozaki, J. Chem. Phys. **142**, 051103 (2015).

[247] M. Saitow and T. Yanai, J. Chem. Phys. **152**, 114111 (2020).

[248] L. Kong, "Internally Contracted Multireference Coupled Cluster Method and Normal-Order-Based Automatic Code Generator", PhD thesis (University of Waterloo, 2009).

[249] A. Köhn, J. Chem. Phys. **130**, 104104 (2009).

[250] G. C. Wick, Phys. Rev. **80**, 268–272 (1950).

[251] C. L. Janssen and H. F. Schaefer III, Theor. Chim. Acta **79**, 1–42 (1991).

[252] J. Paldus and H. C. Wong, Comput. Phys. Commun. **6**, 1–7 (1973).

[253] H. C. Wong and J. Paldus, Comput. Phys. Commun. **6**, 9–16 (1973).

[254] Z. Csépes and J. Pipek, J. Comput. Phys. **77**, 1–17 (1988).

[255] M. Kállay and P. R. Surján, J. Chem. Phys. **115**, 2945–2954 (2001).

[256] A. Hartono, A. Sibiryakov, M. Nooijen, G. Baumgartner, D. E. Bernholdt, S. Hirata, C.-C. Lam, R. M. Pitzer, J. Ramanujam, and P. Sadayappan, "Automated Operation Minimization of Tensor Contraction Expressions in Electronic Structure Calculations", in *Computational Science – ICCS 2005*, Vol. 3514, edited by V. S. Sunderam, G. D. van Albada, P. M. A. Sloot, and J. J. Dongarra, red. by D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, and G. Weikum, Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2005), pp. 155–164.

[257] A. D. Bochevarov and C. D. Sherrill, J. Chem. Phys. **121**, 3374–3383 (2004).

[258] A. Hartono, Q. Lu, X. Gao, S. Krishnamoorthy, M. Nooijen, G. Baumgartner, D. E. Bernholdt, V. Choppella, R. M. Pitzer, J. Ramanujam, A. Rountev, and P. Sadayappan, "Identifying Cost-Effective Common Subexpressions to Reduce Operation Count in Tensor Contraction Evaluations", in *Computational Science – ICCS 2006*, Vol. 3991, edited by V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, red. by D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, and G. Weikum, Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2006), pp. 267–275.

[259] A. Engels-Putzka and M. Hanrath, J. Chem. Phys. **134**, 124106 (2011).

[260] C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh, ACM Trans. Math. Softw. **5**, 308–323 (1979).

[261] M. Hanrath and A. Engels-Putzka, J. Chem. Phys. **133**, 064108 (2010).

[262] P. Springer and P. Bientinesi, ACM Trans. Math. Softw. **44**, 1–29 (2018).

[263] E. Solomonik, D. Matthews, J. Hammond, and J. Demmel, in 2013 IEEE 27th International Symposium on Parallel and Distributed Processing (2013), pp. 813–824.

[264] E. Solomonik, D. Matthews, J. R. Hammond, J. F. Stanton, and J. Demmel, J. Parallel Distrib. Comput. **74**, 3176–3190 (2014).

[265] E. Epifanovsky, M. Wormit, T. Kuś, A. Landau, D. Zuev, K. Khistyaev, P. Manohar, I. Kaliman, A. Dreuw, and A. I. Krylov, J. Comput. Chem. **34**, 2293–2309 (2013).

[266] D. Kats and F. R. Manby, J. Chem. Phys. **138**, 144101 (2013).

[267] P. Springer, *Tensor Contraction Library (TCL)*, 2019, `https://github.com/springer13/tcl` (visited on 09/20/2021).

[268] J. A. Calvin and E. F. Valeev, *TiledArray: A general-purpose scalable block-sparse tensor framework*, `https://github.com/valeevgroup/tiledarray` (visited on 06/30/2021).

[269] J. A. Calvin, C. A. Lewis, and E. F. Valeev, in Proceedings of the 5th Workshop on Irregular Applications: Architectures and Algorithms (2015), pp. 1–8.

[270] N. Jindal, "Using interpretation to optimize and analyze parallel programs in the super instruction architecture", PhD thesis (University of Florida, Gainesville, Florida, 2015).

[271] E. Deumens, V. F. Lotrich, A. Perera, M. J. Ponton, B. A. Sanders, and R. J. Bartlett, Wiley Interdiscip. Rev. Comput. Mol. Sci. **1**, 895–901 (2011).

[272] E. Deumens, V. F. Lotrich, A. S. Perera, R. J. Bartlett, N. Jindal, and B. A. Sanders, "The Super Instruction Architecture", in *Annual Reports in Computational Chemistry*, Vol. 7 (Elsevier, 2011), pp. 179–191.

[273] H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, and M. Schütz, Wiley Interdiscip. Rev. Comput. Mol. Sci. **2**, 242–253 (2012).

[274] J. M. Herbert and W. C. Ermler, Comput. Chem. **22**, 169–184 (1998).

[275] M. Kállay and P. R. Surján, J. Chem. Phys. **113**, 1359–1365 (2000).

[276] K. Kowalski, J. R. Hammond, W. A. de Jong, P.-D. Fan, M. Valiev, D. Wang, and N. Govind, "Coupled-Cluster Calculations for Large Molecular and Extended Systems", in *Computational Methods for Large Systems*, edited by J. R. Reimers (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2011), pp. 167–200.

[277] H. J. J. van Dam, W. A. de Jong, E. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, and M. Valiev, Wiley Interdiscip. Rev. Comput. Mol. Sci. **1**, 888–894 (2011).

[278] S. Hirata, J. Chem. Phys. **121**, 51–59 (2004).

[279] S. Hirata, P.-D. Fan, A. A. Auer, M. Nooijen, and P. Piecuch, J. Chem. Phys. **121**, 12197–12207 (2004).

[280] T. Shiozaki, M. Kamiya, S. Hirata, and E. F. Valeev, J. Chem. Phys. **129**, 071101 (2008).

[281] B. Vlaisavljevich and T. Shiozaki, J. Chem. Theory Comput. **12**, 3781–3787 (2016).

[282] K. Chatterjee and A. Y. Sokolov, J. Chem. Theory Comput. **15**, 5908–5924 (2019).

[283] I. M. Mazin and A. Y. Sokolov, J. Chem. Theory Comput. **17**, 6152–6165 (2021).

[284] F. A. Evangelista, J. Chem. Phys. **157**, 064111 (2022).

[285] L. M. J. Huntington, M. Krupička, F. Neese, and R. Izsák, J. Chem. Phys. **147**, 174104 (2017).

[286] F. Neese, Wiley Interdiscip. Rev. Comput. Mol. Sci. **2**, 73–78 (2012).

[287] F. Neese, Wiley Interdiscip. Rev. Comput. Mol. Sci. **8**, e1327 (2018).

[288] F. Neese, F. Wennmohs, U. Becker, and C. Riplinger, J. Chem. Phys. **152**, 224108 (2020).

[289] F. Neese, Wiley Interdiscip. Rev. Comput. Mol. Sci., `10.1002/wcms.1606` (2022).

[290] B. O. Roos, "The Complete Active Space Self-Consistent Field Method and its Applications in Electronic Structure Calculations", in *Advances in Chemical Physics*, Vol. 69 (John Wiley & Sons, Inc., 1987), pp. 399–445.

[291] P. J. Knowles, M. Schütz, and H.-J. Werner, "Ab Initio Methods for Electron Correlation in Molecules", in *Modern Methods and Algorithms of Quantum Chemistry*, 2nd ed. (John von Neumann Institute for Computing (NIC), Jülich, 2000), pp. 97–161.

[292] M. Radoń, Phys. Chem. Chem. Phys. **21**, 4854–4870 (2019).

[293] M. Musial and R. J. Bartlett, J. Chem. Phys. **129**, 134105 (2008).

[294] A. Mukhopadhyay, R. K. Moitra, and D. Mukherjee, J. Phys. B: At. Mol. Phys. **12**, 1–18 (1979).

[295] D. Mukherjee, R. K. Moitra, and A. Mukhopadhyay, Mol. Phys. **33**, 955–969 (1977).

[296] D. Mukherjee, Int. J. Quantum Chem. **30**, 409–435 (1986).

[297] L. Meissner, Chem. Phys. Lett. **255**, 244–250 (1996).

[298] S. Koch, Theor. Chim. Acta **81**, 169–183 (1991).

[299] A. Landau, E. Eliav, and U. Kaldor, Chem. Phys. Lett. **313**, 399–403 (1999).

[300] M. Musiał, L. Meissner, S. A. Kucharski, and R. J. Bartlett, J. Chem. Phys. **122**, 224110 (2005).

[301] D. Sinha, S. K. Mukhopadhyay, R. Chaudhuri, and D. Mukherjee, Chem. Phys. Lett. **154**, 544–549 (1989).

[302] R. Chaudhuri, D. Mukhopadhyay, and D. Mukherjee, Chem. Phys. Lett. **162**, 393–398 (1989).

[303] B. Datta, R. Chaudhuri, and D. Mukherjee, J. Mol. Struct.: THEOCHEM **361**, 21–31 (1996).

[304] L. Z. Stolarczyk and H. J. Monkhorst, Phys. Rev. A **32**, 743–747 (1985).

[305] L. Z. Stolarczyk and H. J. Monkhorst, Phys. Rev. A **37**, 1908–1925 (1988).

[306] L. Z. Stolarczyk and H. J. Monkhorst, Phys. Rev. A **37**, 1926–1933 (1988).

[307] J. P. Zobel, J. J. Nogueira, and L. González, Chem. Sci. **8**, 1482–1499 (2017).

[308] I. Schapiro, K. Sivalingam, and F. Neese, J. Chem. Theory Comput. **9**, 3567–3580 (2013).

[309] H. Lischka, D. Nachtigallová, A. J. A. Aquino, P. G. Szalay, F. Plasser, F. B. C. Machado, and M. Barbatti, Chem. Rev. **118**, 7293–7361 (2018).

[310] P.-F. Loos, M. Boggio-Pasqua, A. Scemama, M. Caffarel, and D. Jacquemin, J. Chem. Theory Comput. **15**, 1939–1956 (2019).

[311] S. K. Singh, M. Atanasov, and F. Neese, J. Chem. Theory Comput. **14**, 4662–4677 (2018).

[312] S. Vancoillie, P.-Å. Malmqvist, and K. Pierloot, ChemPhysChem **8**, 1803–1815 (2007).

[313] M. Schreiber, M. R. Silva-Junior, S. P. A. Sauer, and W. Thiel, J. Chem. Phys. **128**, 134110 (2008).

[314] M. R. Silva-Junior, M. Schreiber, S. P. A. Sauer, and W. Thiel, J. Chem. Phys. **133**, 174318 (2010).

[315] Y. Kurashige and T. Yanai, J. Chem. Phys. **135**, 094104 (2011).

[316] S. Guo, M. A. Watson, W. Hu, Q. Sun, and G. K.-L. Chan, J. Chem. Theory Comput. **12**, 1583–1591 (2016).

[317] S. Knecht, E. D. Hedegård, S. Keller, A. Kovyrshin, Y. Ma, A. Muolo, C. J. Stein, and M. Reiher, CHIMIA **70**, 244–251 (2016).

[318] Y. Kurashige, J. Chalupský, T. N. Lan, and T. Yanai, J. Chem. Phys. **141**, 174111 (2014).

[319] Q. M. Phung, S. Wouters, and K. Pierloot, J. Chem. Theory Comput. **12**, 4352–4361 (2016).

[320] D. Zgid, D. Ghosh, E. Neuscamman, and G. K.-L. Chan, J. Chem. Phys. **130**, 194107 (2009).

[321] D. Kats and H.-J. Werner, J. Chem. Phys. **150**, 214107 (2019).

[322] F. Menezes, D. Kats, and H.-J. Werner, J. Chem. Phys. **145**, 124115 (2016).

[323] H. Nakano, J. Chem. Phys. **99**, 7983–7992 (1993).

[324] H. Nakano, Chem. Phys. Lett. **207**, 372–378 (1993).

[325] A. A. Granovsky, J. Chem. Phys. **134**, 214113 (2011).

[326] J. Finley, P.-Å. Malmqvist, B. O. Roos, and L. Serrano-Andrés, Chem. Phys. Lett. **288**, 299–306 (1998).

[327] C. Angeli, S. Borini, M. Cestari, and R. Cimiraglia, J. Chem. Phys. **121**, 4043–4049 (2004).

[328] B. G. Levine, J. D. Coe, and T. J. Martínez, J. Phys. Chem. B **112**, 405–413 (2008).

[329] Y. Nishimoto, Chem. Phys. Lett. **744**, 137219 (2020).

[330] J.-P. Malrieu, J.-L. Heully, and A. Zaitsevskii, Theor. Chim. Acta **90**, 167–187 (1995).

[331] I. Shavitt, Int. J. Mol. Sci. **3**, 639–655 (2002).

[332] S. R. Yost, T. Kowalczyk, and T. Van Voorhis, J. Chem. Phys. **139**, 174104 (2013).

[333] S. R. Yost and M. Head-Gordon, J. Chem. Phys. **145**, 054105 (2016).

[334] S. Pathak, L. Lang, and F. Neese, J. Chem. Phys. **147**, 234109 (2017).

[335] B. O. Roos, P. R. Taylor, and P. E. M. Siegbahn, Chem. Phys. **48**, 157–173 (1980).

[336] B. O. Roos, Int. J. Quantum Chem. **18**, 175–189 (1980).

[337] P. Siegbahn, A. Heiberg, B. Roos, and B. Levy, Phys. Scr. **21**, 323–327 (1980).

[338] P. E. M. Siegbahn, J. Almlöf, A. Heiberg, and B. O. Roos, J. Chem. Phys. **74**, 2384–2396 (1981).

[339] P.-Å. Malmqvist, A. Rendell, and B. O. Roos, J. Phys. Chem. **94**, 5477–5482 (1990).

[340] J. Olsen, B. O. Roos, P. Jørgensen, and H. J. A. Jensen, J. Chem. Phys. **89**, 2185–2192 (1988).

[341] G. K.-L. Chan and M. Head-Gordon, J. Chem. Phys. **116**, 4462–4476 (2002).

[342] S. Daul, I. Ciofini, C. Daul, and S. R. White, Int. J. Quantum Chem. **79**, 331–342 (2000).

[343] G. Fano, F. Ortolani, and L. Ziosi, J. Chem. Phys. **108**, 9246–9252 (1998).

[344] A. O. Mitrushenkov, G. Fano, F. Ortolani, R. Linguerri, and P. Palmieri, J. Chem. Phys. **115**, 6815–6821 (2001).

[345] S. R. White and R. L. Martin, J. Chem. Phys. **110**, 4127–4130 (1999).

[346] S. R. White, Phys. Rev. Lett. **69**, 2863–2866 (1992).

[347] S. R. White and R. M. Noack, Phys. Rev. Lett. **68**, 3487–3490 (1992).

[348] C. Angeli, S. Borini, A. Cavallini, M. Cestari, R. Cimiraglia, L. Ferrighi, and M. Sparta, Int. J. Quantum Chem. **106**, 686–691 (2006).

[349] E. Maradzike, M. Hapka, K. Pernal, and A. E. DePrince, J. Chem. Theory Comput. **16**, 4351–4360 (2020).

[350] F. Colmenero, C. Pérez del Valle, and C. Valdemoro, Phys. Rev. A **47**, 971–978 (1993).

[351] F. Colmenero and C. Valdemoro, Phys. Rev. A **47**, 979–985 (1993).

[352] F. Colmenero and C. Valdemoro, Int. J. Quantum Chem. **51**, 369–388 (1994).

[353] F. E. Harris, Int. J. Quantum Chem. **90**, 105–113 (2002).

[354] W. Kutzelnigg and D. Mukherjee, J. Chem. Phys. **110**, 2800–2809 (1999).

[355] D. A. Mazziotti, Chem. Phys. Lett. **326**, 212–218 (2000).

[356] D. A. Mazziotti, Phys. Rev. A **60**, 3618–3626 (1999).

[357] H. Nakatsuji and K. Yasuda, Phys. Rev. Lett. **76**, 1039–1042 (1996).

[358] D. A. Mazziotti, Phys. Rev. A **57**, 4219–4234 (1998).

[359] H. Nakatsuji, Phys. Rev. A **14**, 41–50 (1976).

[360] N. Nakatani and S. Guo, J. Chem. Phys. **146**, 094102 (2017).

[361] Y. Guo, K. Sivalingam, and F. Neese, J. Chem. Phys. **154**, 214111 (2021).

[362] Y. Guo, K. Sivalingam, C. Kollmar, and F. Neese, J. Chem. Phys. **154**, 214113 (2021).

[363] C. Li and F. A. Evangelista, J. Chem. Phys. **148**, 079903 (2018).

[364] C. Li and F. A. Evangelista, Annu. Rev. Phys. Chem. **70**, 245–273 (2019).

[365] C. Li and F. A. Evangelista, J. Chem. Phys. **144**, 164114 (2016).

[366] P. Celani and H.-J. Werner, J. Chem. Phys. **112**, 5546–5557 (2000).

[367] A. Y. Sokolov and G. K.-L. Chan, J. Chem. Phys. **144**, 064102 (2016).

[368] A. Y. Sokolov, S. Guo, E. Ronca, and G. K.-L. Chan, J. Chem. Phys. **146**, 244102 (2017).

[369] S. Sharma, G. Knizia, S. Guo, and A. Alavi, J. Chem. Theory Comput. **13**, 488–498 (2017).

[370] S. Sharma and A. Alavi, J. Chem. Phys. **143**, 102815 (2015).

[371] S. Sharma and G. K.-L. Chan, J. Chem. Phys. **141**, 111101 (2014).

[372] S. Wouters, V. Van Speybroeck, and D. Van Neck, J. Chem. Phys. **145**, 054120 (2016).

[373] M. Roemelt, S. Guo, and G. K.-L. Chan, J. Chem. Phys. **144**, 204113 (2016).

[374] A. Mahajan, N. S. Blunt, I. Sabzevari, and S. Sharma, J. Chem. Phys. **151**, 211102 (2019).

[375] R. J. Anderson, T. Shiozaki, and G. H. Booth, J. Chem. Phys. **152**, 054101 (2020).

[376] K. P. Hannon, C. Li, and F. A. Evangelista, J. Chem. Phys. **144**, 204111 (2016).

[377] C. Li and F. A. Evangelista, J. Chem. Phys. **148**, 124106 (2018).

[378] C. Li and F. A. Evangelista, J. Chem. Phys. **146**, 124132 (2017).

[379] C. Li and F. A. Evangelista, J. Chem. Phys. **148**, 079902 (2018).

[380] K. Pernal, Phys. Rev. Lett. **120**, 013001 (2018).

[381] K. Chatterjee and A. Y. Sokolov, J. Chem. Theory Comput. **16**, 6343–6357 (2020).

[382] C. Kollmar, K. Sivalingam, Y. Guo, and F. Neese, J. Chem. Phys. **155**, 234104 (2021).

[383] R. Pauncz, *Spin Eigenfunctions: Construction and Use* (Plenum Press, New York, 1979).

[384] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, Mineola, N.Y, 1996).

[385] I. Mayer, *Simple Theorems, Proofs, and Derivations in Quantum Chemistry* (Springer US, 2003).

[386] E. R. Davidson, *Reduced Density Matrices in Quantum Chemistry*, Vol. 6, Theoretical Chemistry (Academic Press, New York, 1976).

[387] P.-O. Löwdin, Phys. Rev. **97**, 1474–1489 (1955).

[388] J. E. Mayer, Phys. Rev. **100**, 1579–1586 (1955).

[389] J. Paldus and J. Čížek, "Time-Independent Diagrammatic Approach to Perturbation Theory of Fermion Systems", in *Advances in Quantum Chemistry*, Vol. 9 (Elsevier, 1975), pp. 105–197.

[390] W. Kutzelnigg, J. Chem. Phys. **77**, 3081–3097 (1982).

[391] J. Paldus and B. Jeziorski, Theor. Chim. Acta **73**, 81–103 (1988).

[392] W. Kutzelnigg, J. Chem. Phys. **80**, 822–830 (1984).

[393] W. Kutzelnigg, J. Chem. Phys. **82**, 4166–4186 (1985).

[394] J. Goldstone, Proc. R. Soc. London A **239**, 267–279 (1957).

[395] N. Hugenholtz, Physica **23**, 481–532 (1957).

[396] B. H. Brandow, Rev. Mod. Phys. **39**, 771–828 (1967).

[397] J. Paldus, J. Chem. Phys. **67**, 303–318 (1977).

[398] S. A. Kucharski and R. J. Bartlett, "Fifth-Order Many-Body Perturbation Theory and Its Relationship to Various Coupled-Cluster Approaches", in *Advances in Quantum Chemistry*, Vol. 18 (Elsevier, 1986), pp. 281–344.

[399] R. J. Bartlett and D. M. Silver, Int. J. Quantum Chem. **9**, 183–198 (1975).

[400] F. E. Harris, Int. J. Quantum Chem. **75**, 593–597 (1999).

[401] W. Kutzelnigg and D. Mukherjee, "Generalized Normal Ordering, Irreducible Brillouin Conditions, and Contracted Schrödinger Equations", in *Reduced-Density-Matrix Mechanics: With Application to Many-Electron Atoms and Molecules* (John Wiley & Sons, Inc., Hoboken, New Jersey, 2007), pp. 293–330.

[402] D. Sinha, R. Maitra, and D. Mukherjee, Comput. Theor. Chem. **1003**, 62–70 (2013).

[403] W. Kutzelnigg and D. Mukherjee, J. Chem. Phys. **120**, 7350–7368 (2004).

[404] L. Kong and E. F. Valeev, J. Chem. Phys. **134**, 214109 (2011).

[405] M. Hanauer and A. Köhn, Chem. Phys. **401**, 50–61 (2012).

[406] D. A. Mazziotti, Chem. Phys. Lett. **289**, 419–427 (1998).

[407] D. A. Mazziotti, Phys. Rev. A **60**, 4396–4408 (1999).

[408] J. E. Harriman, Phys. Rev. A **65**, 052507 (2002).

[409] W. Kutzelnigg and D. Mukherjee, J. Chem. Phys. **116**, 4787–4801 (2002).

[410] L. Lain, A. Torre, and R. Bochicchio, J. Chem. Phys. **117**, 5497–5498 (2002).

[411] M. Nooijen and V. Lotrich, J. Chem. Phys. **113**, 4549–4557 (2000).

[412] M. Nooijen and V. Lotrich, J. Chem. Phys. **113**, 494–507 (2000).

[413] M. Nooijen, Int. J. Mol. Sci. **3**, 656–675 (2002).

[414] M. Nooijen, M. Wladyslawski, and A. Hazra, J. Chem. Phys. **118**, 4832–4848 (2003).

[415] J. M. Herbert, Int. J. Quantum Chem. **107**, 703–711 (2007).

[416] D. R. Hartree, Math. Proc. Cambridge Philos. Soc. **24**, 89–110 (1928).

[417] D. R. Hartree, Math. Proc. Cambridge Philos. Soc. **24**, 111–132 (1928).

[418] V. Fock, Z. Physik **61**, 126–148 (1930).

[419] V. Fock, Z. Physik **62**, 795–805 (1930).

[420] J. C. Slater, Phys. Rev. **35**, 210–211 (1930).

[421] D. R. Hartree and W. Hartree, Proc. R. Soc. London A **150**, 9–33 (1935).

[422] W. Kohn and L. J. Sham, Phys. Rev. **140**, A1133–A1138 (1965).

[423] J. C. Slater, Phys. Rev. **34**, 1293–1322 (1929).

[424] E. U. Condon, Phys. Rev. **36**, 1121–1133 (1930).

[425] C. C. J. Roothaan, Rev. Mod. Phys. **23**, 69–89 (1951).

[426] C. Froese Fischer, J. Phys. B: At. Mol. Phys. **10**, 1241–1251 (1977).

[427] C. W. Bauschlicher, P. Siegbahn, and L. G. M. Pettersson, Theor. Chim. Acta **74**, 479–491 (1988).

[428] B. O. Roos, "Multiconfigurational quantum chemistry for ground and excited states", in *Radiation Induced Molecular Phenomena in Nucleic Acids: A Comprehensive Theoretical and Experimental Analysis*, edited by M. K. Shukla and J. Leszczynski, Challenges and Advances In Computational Chemistry and Physics (Springer Netherlands, Dordrecht, 2008), pp. 125–156.

[429] G. H. Booth, A. J. W. Thom, and A. Alavi, J. Chem. Phys. **131**, 054106 (2009).

[430] B. Huron, J. P. Malrieu, and P. Rancurel, J. Chem. Phys. **58**, 5745–5759 (1973).

[431] R. J. Buenker and S. D. Peyerimhoff, Theor. Chim. Acta **39**, 217–228 (1975).

[432] F. A. Evangelista, J. Chem. Phys. **140**, 124114 (2014).

[433] B. S. Fales, S. Seritan, N. F. Settje, B. G. Levine, H. Koch, and T. J. Martínez, J. Chem. Theory Comput. **14**, 4139–4150 (2018).

[434] Y. Garniron, A. Scemama, E. Giner, M. Caffarel, and P.-F. Loos, J. Chem. Phys. **149**, 064103 (2018).

[435] N. M. Tubman, D. S. Levine, D. Hait, M. Head-Gordon, and K. B. Whaley, "An efficient deterministic perturbation theory for selected configuration interaction methods", 2018.

[436] D. S. Levine, D. Hait, N. M. Tubman, S. Lehtola, K. B. Whaley, and M. Head-Gordon, J. Chem. Theory Comput. **16**, 2340–2354 (2020).

[437] N. Zhang, W. Liu, and M. R. Hoffmann, J. Chem. Theory Comput. **16**, 2296–2316 (2020).

[438] V. G. Chilkuri and F. Neese, J. Comput. Chem. **42**, 982–1005 (2021).

[439] V. G. Chilkuri and F. Neese, J. Chem. Theory Comput. **17**, 2868–2885 (2021).

[440] L. Freitag, S. Knecht, C. Angeli, and M. Reiher, J. Chem. Theory Comput. **13**, 451–459 (2017).

[441] W. Kutzelnigg, Mol. Phys. **94**, 65–71 (1998).

[442] B. Cooper and P. J. Knowles, J. Chem. Phys. **133**, 234102 (2010).

[443] A. K. Dutta, M. Nooijen, F. Neese, and R. Izsák, J. Chem. Phys. **146**, 074103 (2017).

[444] S. F. Boys, Rev. Mod. Phys. **32**, 296–299 (1960).

[445] J. Pipek and P. G. Mezey, J. Chem. Phys. **90**, 4916–4926 (1989).

[446] P. Pulay, Chem. Phys. Lett. **100**, 151–154 (1983).

[447] O. Vahtras, J. Almlöf, and M. Feyereisen, Chem. Phys. Lett. **213**, 514–518 (1993).

[448] K. Eichkorn, O. Treutler, H. Öhm, M. Häser, and R. Ahlrichs, Chem. Phys. Lett. **242**, 652–660 (1995).

[449] D. G. Liakos, M. Sparta, M. K. Kesharwani, J. M. L. Martin, and F. Neese, J. Chem. Theory Comput. **11**, 1525–1539 (2015).

[450] S. Ghosh, A. Kumar Dutta, B. de Souza, R. Berraud-Pache, and R. Izsák, Mol. Phys. **118**, e1818858 (2020).

[451] V. Barone and M. Cossi, J. Phys. Chem. A **102**, 1995–2001 (1998).

[452] C. Hampel, K. A. Peterson, and H.-J. Werner, Chem. Phys. Lett. **190**, 1–12 (1992).

[453] W. D. Laidig and R. J. Bartlett, Chem. Phys. Lett. **104**, 424–430 (1984).

[454] M. Hanauer and A. Köhn, J. Chem. Phys. **137**, 131103 (2012).

[455] A. Köhn, J. A. Black, Y. A. Aoto, and M. Hanauer, Mol. Phys. **118**, e1743889 (2020).

[456] K. K. Irikura, "Appendix B: Essential Statistical Thermodynamics", in *Computational Thermochemistry*, Vol. 677, ACS Symposium Series 677 (American Chemical Society, Washington, DC, 1998), pp. 402–418.

[457] Y. Niu, Q. Peng, C. Deng, X. Gao, and Z. Shuai, J. Phys. Chem. A **114**, 7817–7831 (2010).

[458] A. Baiardi, J. Bloino, and V. Barone, J. Chem. Theory Comput. **9**, 4097–4115 (2013).

[459] B. de Souza, F. Neese, and R. Izsák, J. Chem. Phys. **148**, 034104 (2018).

[460] F. J. Avila Ferrer and F. Santoro, Phys. Chem. Chem. Phys. **14**, 13549–13563 (2012).

[461] V. Barone, M. Biczysko, C. Latouche, and A. Pasti, Theor. Chem. Acc. **134**, 145 (2015).

[462] M. Nooijen, J. Chem. Phys. **104**, 2638–2651 (1996).

[463] F. Weigend and R. Ahlrichs, Phys. Chem. Chem. Phys. **7**, 3297–3305 (2005).

[464] N. Welsch, *Indigo*, in *Römpp online*, edited by F. Böckler, B. Dill, G. Eisenbrand, F. Faupel, B. Fugmann, T. Gamse, R. Matissek, G. Pohnert, A. Rühling, S. Schmidt, and G. Sprenger (Thieme Gruppe, Stuttgart, 2006), `https://roempp.thieme.de/lexicon/RD-09-00480` (visited on 08/03/2022).

[465] T. Maugard, E. Enaud, P. Choisy, and M. D. Legoy, Phytochemistry **58**, 897–904 (2001).

[466] A. Baeyer and A. Emmerling, Ber. Dtsch. Chem. Ges. **3**, 514–517 (1870).

[467] A. Baeyer, Ber. Dtsch. Chem. Ges. **16**, 2188–2204 (1883).

[468] T. M. Hsu, D. H. Welner, Z. N. Russ, B. Cervantes, R. L. Prathuri, P. D. Adams, and J. E. Dueber, Nat. Chem. Biol. **14**, 256–261 (2018).

[469] K. Holzach, Angew. Chem. **60**, 200–204 (1948).

[470] J. Gosteli, Helv. Chim. Acta **60**, 1980–1983 (1977).

[471] L. Wolf, Chem. Eng. News **89**, 44 (2011).

[472] D. Weiß, *Purpur*, in *Römpp online*, edited by F. Böckler, B. Dill, G. Eisenbrand, F. Faupel, B. Fugmann, T. Gamse, R. Matissek, G. Pohnert, A. Rühling, S. Schmidt, and G. Sprenger (Thieme Gruppe, Stuttgart, 2019), `https://roempp.thieme.de/lexicon/RD-16-05002` (visited on 08/03/2022).

[473] F. Sachs and R. Kempf, Ber. Dtsch. Chem. Ges. **36**, 3299–3303 (1903).

[474] P. Friedländer, Ber. Dtsch. Chem. Ges. **42**, 765–770 (1909).

[475] H. Hermann and W. Lüttke, Chem. Ber. **101**, 1708–1714 (1968).

[476] M. Klessinger and W. Lüttke, Tetrahedron **19**, 315–335 (1963).

[477] W. Lüttke, H. Hermann, and M. Klessinger, Angew. Chem. Int. Ed. **5**, 598–599 (1966).

[478] E. Hückel, Z. Physik **70**, 204–286 (1931).

[479] E. Hückel, Z. Physik **72**, 310–337 (1931).

[480] E. Hückel, Z. Physik **76**, 628–648 (1932).

[481] E. Hückel, Z. Physik **83**, 632–668 (1933).

[482] R. Pariser and R. G. Parr, J. Chem. Phys. **21**, 466–471 (1953).

[483] R. Pariser and R. G. Parr, J. Chem. Phys. **21**, 767–776 (1953).

[484] J. A. Pople, Trans. Faraday Soc. **49**, 1375–1385 (1953).

[485] D. Jacquemin, J. Preat, V. Wathelet, and E. A. Perpète, J. Chem. Phys. **124**, 074104 (2006).

[486] M. Moreno, J. M. Ortiz-Sánchez, R. Gelabert, and J. M. Lluch, Phys. Chem. Chem. Phys. **15**, 20236–20246 (2013).

[487] F. Cervantes-Navarro and D. Glossman-Mitnik, J. Photochem. Photobiol. A **255**, 24–26 (2013).

[488] S. Shityakov, N. Roewer, C. Förster, and J.-A. Broscheit, Nanoscale Res. Lett. **12**, 439 (2017).

[489] J. Pina, D. Sarmento, M. Accoto, P. L. Gentili, L. Vaccaro, A. Galvão, and J. S. Seixas de Melo, J. Phys. Chem. B **121**, 2308–2318 (2017).

[490] Z. Ju, J. Sun, and Y. Liu, Molecules **24**, 3831 (2019).

[491] L. Serrano-Andrés and B. O. Roos, Chem. Eur. J. **3**, 717–725 (1997).

[492] S. Yamazaki, A. L. Sobolewski, and W. Domcke, Phys. Chem. Chem. Phys. **13**, 1618–1628 (2011).

[493] J. F. Stanton, Faraday Discuss. **150**, 331–343 (2011).

[494] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, J. Phys. Chem. **98**, 11623–11627 (1994).

[495] F. Weigend, Phys. Chem. Chem. Phys. **8**, 1057–1065 (2006).

[496] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, J. Chem. Phys. **132**, 154104 (2010).

[497] F. Neese, F. Wennmohs, A. Hansen, and U. Becker, Chem. Phys. **356**, 98–109 (2009).

[498] R. Izsák and F. Neese, J. Chem. Phys. **135**, 144105 (2011).

[499] S. Dähne and D. Leupold, Angew. Chem. Int. Ed. **5**, 984–993 (1966).

[500] Z. C. Koren, "HPLC-PDA analysis of brominated indirubinoid, indigoid, and isatinoid dyes", in *Indirubin, the Red Shade of Indigo*, edited by L. Meijer, N. Guyard, L. Skaltsounis, and G. Eisenbrand (Ed. "Life in Progress", Roscoff, France, 2006), pp. 45–53.

[501] O. W. Day, D. W. Smith, and C. Garrod, Int. J. Quantum Chem. **8**, 501–509 (1974).

[502] D. W. Smith and O. W. Day, J. Chem. Phys. **62**, 113–114 (1975).

[503] M. M. Morrell, R. G. Parr, and M. Levy, J. Chem. Phys. **62**, 549–554 (1975).

[504] R. C. Morrison, J. Chem. Phys. **96**, 3718–3722 (1992).

[505] D. Sundholm and J. Olsen, J. Chem. Phys. **98**, 3999–4002 (1993).

[506] M. R. Silva-Junior, M. Schreiber, S. P. A. Sauer, and W. Thiel, J. Chem. Phys. **129**, 104103 (2008).

[507] A. Schäfer, C. Huber, and R. Ahlrichs, J. Chem. Phys. **100**, 5829–5835 (1994).

[508] R. A. Kendall, T. H. Dunning, and R. J. Harrison, J. Chem. Phys. **96**, 6796–6806 (1992).

[509] X. Li and J. Paldus, J. Chem. Phys. **103**, 1024–1034 (1995).

[510] D. Gerbig and D. Ley, Wiley Interdiscip. Rev. Comput. Mol. Sci. **3**, 242–272 (2013).

[511] W. C. Lineberger and W. T. Borden, Phys. Chem. Chem. Phys. **13**, 11792–11813 (2011).

[512] R. W. A. Havenith, P. R. Taylor, C. Angeli, R. Cimiraglia, and K. Ruud, J. Chem. Phys. **120**, 4619–4625 (2004).

[513] A. Kramida, Yu. Ralchenko, J. Reader, and NIST ASD Team, *NIST Atomic Spectra Database (ver. 5.7.1)* (National Institute of Standards and Technology, Gaithersburg, MD, 2019).

[514] J. F. Stanton, J. Gauss, J. D. Watts, W. J. Lauderdale, and R. J. Bartlett, Int. J. Quantum Chem. **44**, 879–894 (1992).

[515] M. Douglas and N. M. Kroll, Ann. Phys. **82**, 89–155 (1974).

[516] B. A. Hess, Phys. Rev. A **33**, 3742–3748 (1986).

[517] T. B. Pedersen, "Introduction to Response Theory", in *Handbook of Computational Chemistry*, edited by J. Leszczynski (Springer Netherlands, Dordrecht, 2015), pp. 1–26.

[518] P. J. Knowles and H.-J. Werner, Chem. Phys. Lett. **145**, 514–522 (1988).

[519] P. E. Siegbahn, Chem. Phys. Lett. **109**, 417–423 (1984).

[520] I. L. Cooper and R. McWeeny, J. Chem. Phys. **45**, 3484–3484 (1966).

[521] I. L. Cooper and R. McWeeny, J. Chem. Phys. **45**, 226–234 (1966).

[522] J. Paldus, J. Chem. Phys. **61**, 5321–5330 (1974).

[523] W. I. Salmon and K. Ruedenberg, J. Chem. Phys. **57**, 2776–2786 (1972).

[524] R. W. Wetmore and G. A. Segal, Chem. Phys. Lett. **36**, 478–483 (1975).

[525] K. G. Dyall, J. Chem. Phys. **102**, 4909–4918 (1995).

[526] F. Cardone and J. R. Hindley, "Lambda-Calculus and Combinators in the 20th Century", in *Handbook of the History of Logic*, Vol. 5, edited by D. M. Gabbay and J. Woods, Logic from Russell to Church (North-Holland, Amsterdam, 2009), pp. 723–817.

[527] A. D. Becke, J. Chem. Phys. **98**, 5648–5652 (1993).

[528] C. Lee, W. Yang, and R. G. Parr, Phys. Rev. B **37**, 785–789 (1988).

[529] S. H. Vosko, L. Wilk, and M. Nusair, Can. J. Phys. **58**, 1200–1211 (1980).

[530] S. Grimme, S. Ehrlich, and L. Goerigk, J. Comput. Chem. **32**, 1456–1465 (2011).

[531] D. A. Matthews, SIAM J. Sci. Comput. **40**, C1–C24 (2018).

[532] E. F. Valeev and J. T. Fermann, *Libint*, 2021, `https://github.com/evaleev/libint` (visited on 09/16/2021).

[533] F. Neese, J. Comput. Chem., `10.1002/jcc.26942` (2022).

[534] V. Lotrich, N. Jindal, E. Deumens, R. J. Bartlett, and B. A. Sanders, *Final Report: The Super Instruction Architecture*, Online Report (2013), `https://bluewaters.ncsa.illinois.edu/liferay-content/document-library/NEIS-P2-C1-Reports/Bartlett_Final_Report_v2.pdf` (visited on 09/08/2021).

[535] V. F. Lotrich, J. M. Ponton, A. S. Perera, E. Deumens, R. J. Bartlett, and B. A. Sanders, Mol. Phys. **108**, 3323–3330 (2010).

[536] J. G. Baragi, M. I. Aralaguppi, T. M. Aminabhavi, M. Y. Kariduraganavar, and A. S. Kittur, J. Chem. Eng. Data **50**, 910–916 (2005).

[537] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, A. Vijaykumar, A. P. Bardelli, A. Rothberg, A. Hilboll, A. Kloeckner, A. Scopatz, A. Lee, A. Rokem, C. N. Woods, C. Fulton, C. Masson, C. Häggström, C. Fitzgerald, D. A. Nicholson, D. R. Hagen, D. V. Pasechnik, E. Olivetti, E. Martin, E. Wieser, F. Silva, F. Lenders, F. Wilhelm, G. Young, G. A. Price, G.-L. Ingold, G. E. Allen, G. R. Lee, H. Audren, I. Probst, J. P. Dietrich, J. Silterra, J. T. Webber, J. Slavič, J. Nothman, J. Buchner, J. Kulick, J. L. Schönberger, J. V. de Miranda Cardoso, J. Reimer, J. Harrington, J. L. C. Rodríguez, J. Nunez-Iglesias, J. Kuczynski, K. Tritz, M. Thoma, M. Newville, M. Kümmerer, M. Bolingbroke, M. Tartre, M. Pak, N. J. Smith, N. Nowaczyk, N. Shebanov, O. Pavlyk, P. A. Brodtkorb, P. Lee, R. T. McGibbon, R. Feldbauer, S. Lewis, S. Tygier, S. Sievert, S. Vigna, S. Peterson, S. More, T. Pudlik, T. Oshima, T. J. Pingel, T. P. Robitaille, T. Spura, T. R. Jones, T. Cera, T. Leslie, T. Zito, T. Krauss, U. Upadhyay, Y. O. Halchenko, and Y. Vázquez-Baeza, Nat. Methods **17**, 261–272 (2020).

[538] T. H. Dunning, J. Chem. Phys. **90**, 1007–1023 (1989).