FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

DOCTORAL THESIS

Jan Brandejs

Tensor Network-based Computational Methods for Strongly Correlated Molecular Quantum Mechanics

Institute of Theoretical Physics

Supervisor of the doctoral thesis: RNDr. Libor Veis, Ph.D.
Study programme: Theoretical Physics, Astronomy and Astrophysics
Study branch: P4F1

Prague 2022

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources. It has not been used to obtain another or the same degree.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In .... date ....

Author's signature

To Hana.

Title: Tensor Network-based Computational Methods for Strongly Correlated Molecular Quantum Mechanics

Author: Jan Brandejs

Institute: Institute of Theoretical Physics

Supervisor: RNDr. Libor Veis, Ph.D., J. Heyrovský Institute of Physical Chemistry of the Czech Academy of Sciences

Abstract: Despite an immense progress in recent decades, a precise treatment of strongly correlated molecular systems still remains a challenge as of today. To help solve this problem, we have developed a massively parallel implementation of DMRG, called MOLMPS. For nonrelativistic systems requiring accurate treatment, we have extended MOLMPS by the means of the almost linear scaling DLPNO-TCCSD method. In relativistic domain, we are the first who employed the 4c-CCSD to add dynamical correlation on top of DMRG, yielding the 4c-TCCSD method. When tested on benchmarks like the π-conjugated anthracene tetramer with CAS(63,63) and the FeMoco cofactor with CAS(113,76). We showed a good parallel performance on up to about 2000 CPUs. On the example of Iron(II)-Porphyrin model, we showed that the DLPNO-TCCSD captures 99.9% of TCCSD correlation energy. Our spectroscopic study on heavy diatomics showed that the 4c-TCCSD approach increases the precision of underlying CCSD to the order of CCSD(T) and that it is a promising approach. The thesis discusses three different implementations of quantum chemical methods based on QC-DMRG.

Keywords: matrix product state tensor networks DMRG quantum chemistry strong correlation parallelization coupled clusters

Contents

Introduction

Defining quantum chemistry

Molecular quantum mechanics, also known as quantum chemistry, is a field of physical chemistry focused on application of quantum mechanics to chemical systems. The central problem of quantum chemistry is finding the electronic structure of molecules and understanding the molecular dynamics by solving the Schrödinger equation.

Quantum chemistry studies ground state of individual molecules, their excited states and the transition states occurring in chemical processes. It is related to field of spectroscopy, which measures energy spectra of molecular systems by analyzing absorption and emission of light and other radiation. Hence, the two fields complement each other. Theoretical results of quantum chemistry can be compared for example with measurements from infrared spectroscopy, nuclear magnetic resonance (NMR) spectroscopy and scanning probe microscopy, see Engel [2006].

This work is focused on the electronic structure problem. Electronic structure represents the states of electrons in an electrostatic field created by stationary nuclei. The nuclei are considered stationary within the frame of Born-Oppenheimer approximation, which is based on the observation that electrons in molecules are much lighter and faster than the relatively cumbersome nuclei. Hence, the wave function of electrons can be treated separately, see Combes [1977].

Ground state of a quantum-mechanical system is defined as a state with the lowest energy. Finding the lowest-lying states is important because electrons naturally fill the lowest available energy levels first (aufbau principle, see Park and Stetten [2001]). Finding the low-lying states is important, because they have the strongest influence on the behaviour of the system, due to the elementary principle of extremal action. It follows that at low temperatures, these are the most frequently populated states, i.e. frequently occuring in the Nature.

Moreover, quantum effects are usually strongest at low temperatures, when the interaction with environment is weak and the system is coherent. When the system interacts with environment, for example due to temperature effects or during a measurement, quantum coherence is shared with the environment and appears to be lost with time in a process called decoherence. Consistently, with increasing temperature, quantum behavior is apparently lost, similarly to how energy appears to be lost by friction in classical mechanics. Notice that here we do not assume equivalence between quantum wave function collapse during measurement and quantum decoherence, as this is a subject of quite a heated debate, see Nielsen and Chuang [2009], Schlosshauer [2019], Sergi and Kapral [2004].

The ability to accurately describe quantum molecular systems is crucial. Finding effective solutions to electronic structure problem allows us to accurately predict the structure of molecules, their spectroscopic properties, or model their reactions. Considering larger systems, quantum chemistry can model properties of nanomaterials, and reactions and processes in biological systems.

Over the time, electronic structure calculations have found numerous applications in chemistry, including the study of catalysis, magnetic systems, nanoparticles, quantum dynamics, excited states, in transition metal chemistry and in organic chemistry, see Taylor-Smith [2019], hong Lam et al. [2020]. Through applications in the aforementioned fields, quantum chemistry also contributes to chemical industry by deepening the understanding of underlying processes, allowing further improvement, see Slater and Johnson [1974], Ess et al. [2019].

Electronic structure theory is used in pharmaceutical research and drug design. Considering the biological systems, for instance, through study of interactions in nucleic acid bases of DNA, quantum chemistry enabled understanding of processes occuring during genetic modifications (see Park et al. [2015]), see also van Mourik [2004] and the debated topics Turin [1996], Higgins et al. [2021].

Moreover, quantum chemistry continues to show value for experimental spectroscopy itself by supporting and interpreting experimental data, see Engel [2006]. It may come as a surprise, but field of spectroscopy saw important contributions even to theoretical physics. In cosmology, the study of observed spectral emission lines of distant galaxies led to the discovery that the universe is expanding rapidly and isotropically (independent of direction) in 1920s, confirmed eventually by the discovery of microwave background radiation in 1965, see Hurst et al. [2021]. Of course, in 1920s, there was no contribution of computational quantum chemistry, we only mention this example to show the connection between physics and spectroscopy.

Solving the electronic structure problem

The electronic structure of molecular systems is described by the time-independent Schrödinger equation

$$H\left|\boldsymbol{\Psi}\right\rangle = E\left|\boldsymbol{\Psi}\right\rangle \tag{1}$$

where E is the energy of eigenstate $\left|\boldsymbol{\Psi}\right\rangle$ of the Hamiltonian H corresponding to the wave function $\boldsymbol{\Psi}$. Note that 1 is an eigenvalue equation. Within the Born-Oppenheimer approximation, the electronic Hamiltonian in its nonrelativistic form reads

$$\begin{aligned} H &= T + U \\ &= -\sum_i \frac{\hbar^2}{2m_e}\nabla^2_{r_i} - \frac{1}{2}\sum_i\sum_A \frac{eZ_A}{4\pi\varepsilon_0}\frac{1}{r_{iA}} + \frac{1}{2}\sum_i\sum_{j>i}\frac{e^2}{4\pi\varepsilon_0}\frac{1}{r_{ij}}. \end{aligned} \tag{2}$$

The Hamiltonian operator is composed of kinetic term of electrons followed by a term corresponding to Coulomb attraction between electrons and nuclei and a term for Coulomb repulsion between electrons themselves. Electrons are labelled with indices i and j, nuclei with A. Z_A represents the nuclear charge, r_{iA} is the distance between electron i and nucleus A and r_{ij} is the distance betweem electrons i and j. Here, $\hbar \doteq 1.054\,572 \times 10^{-34}$ J $\cdot$ s is the reduced plank constant, $\varepsilon_0 \doteq 8.854\,188 \times 10^{-12}$ Fm^{-1} the vacuum permitivity, $m_e \doteq 9.109\,384 \times 10^{-31}$ kg the rest mass of electron and $e \doteq 1.602\,177 \times 10^{-19}$ C its charge. In the rest of this work, we work in atomic units.

The basic method for solving equation 1, upon which other, more advanced methods build, is the *Hartee-Fock method* (HF), also known as the self-consistent field method (SCF). Hartree–Fock method assumes that the wave function can be approximated by a single Slater determinant. By employing the variational method, one can derive k coupled equations for k spin-orbitals. Their solution yields the HF state and HF energy of the system. For a self-contained introduction to the HF method, see Szabo and Ostlund [1996].

The essence of the HF method is to transform the original n-electron problem to one-electron problems by means of the mean-field approximation. We distinguish two types of one-electron wave functions, so called orbitals.

Spatial orbital, $\psi(\mathbf{r})$, where $\mathbf{r}$ is a plain location in space, determines the distribution of an electron within a molecule, such that $|\psi(\mathbf{r})|^2 d\mathbf{r}$ is the probability of finding the electron in the small volume element $d\mathbf{r}$ in around $\mathbf{r}$. A requirement of orthonormality is usually imposed upon the set of molecular orbitals

$$\int \psi_i^\dagger(\mathbf{r})\psi_j(\mathbf{r})d\mathbf{r} = \delta_{ij}. \tag{3}$$

Were the set $\{\psi_i\}$ complete, it would form a basis of all scalar functions $f(\mathbf{r})$. In practice, we work with a finite set of molecular orbitals (MOs), which only span a certain region of complete space. We can, however, describe results as being "exact" within the subspace spanned by $\{\psi_i\}$. In this work, when we talk about the basis, we usually refer to a set of molecular orbitals, which form the one-electron basis for Slater determinants spanning the Hilbert space, within which we search for eigenstates of the electronic Hamiltonian.

In order to completely describe an electron in a molecule, we need to specify its spin. But spatial orbitals, $\psi_i(\mathbf{r})$ do not include the information about spin. *Spin-orbital*, $\chi(\mathbf{x})$, whose location and spin is described by a composite spatial-spin variable $\mathbf{x}$, is a type of molecular orbital which includes this information by incorporating one of the two orthonormal spin-functions, $\alpha(\omega)$ or $\beta(\omega)$,

$$\chi(\mathbf{x}) \;=\; \chi(\mathbf{r},\omega) \;=\; \begin{cases} \psi_i(\mathbf{r})\,\alpha(\omega) \\ \quad\text{or} \\ \psi_i(\mathbf{r})\,\beta(\omega) \end{cases} \tag{4}$$

Here, the unspecified variables of ω serve only for defining α and β as functions, e.g. to be able to use standard integral definition of orthonormality (equation 3) of α and β. Thus, each spin-orbital corresponds either to ↑ or ↓ projection of spin.

Now consider a wave function Ψ, describing a state of n electrons. Wave functions obey the Pauli exclusion principle, which leads to a requirement:

> *A many-electron wave function must be antisymmetric with respect to interchange of coordinates (both space and spin) of any two electrons.*

For n electrons

$$\Psi(\mathbf{x}_1,\ldots,\mathbf{x}_i,\ldots,\mathbf{x}_j,\ldots,\mathbf{x}_n) \;=\; -\Psi(\mathbf{x}_1,\ldots,\mathbf{x}_j,\ldots,\mathbf{x}_i,\ldots,\mathbf{x}_n). \tag{5}$$

It follows that the wave function of HF state must be antisymmetric with respect to electron interchange.

Let use recall that the mathematical concept of a matrix determinant is totally antisymmetric with respect to permutations of rows and columns. Hence, we may use the notion of determinant for a formal description of many-electron wave function in terms of states of individual electrons, if we map them to the rows and columns of a matrix. This is precisely what *Slater determinant* does.

For an n-electron system, *the Slater determinant* is a wave function defined in terms of a determinant of spin-orbitals $\chi_i(\mathbf{x}_i)$

$$\boldsymbol{\Psi}_n(\mathbf{x}_1,\, \mathbf{x}_2,\, \ldots,\, \mathbf{x}_n) \quad = \quad \frac{1}{\sqrt{n!}} \begin{vmatrix} \chi_1(\mathbf{x}_1) & \chi_2(\mathbf{x}_1) & \cdots & \chi_n(\mathbf{x}_1) \\ \chi_1(\mathbf{x}_2) & \chi_2(\mathbf{x}_2) & \cdots & \chi_n(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \chi_1(\mathbf{x}_n) & \chi_2(\mathbf{x}_n) & \cdots & \chi_n(\mathbf{x}_n) \end{vmatrix}. \tag{6}$$

This determinant has exactly the structure that we were looking for. As mentioned above, the HF method variationally optimizes the wave function in the form of a single Slater determinant.

The formalism of second quantization is an efficient and in quantum theory frequently used approach, which found its use also in quantum chemistry, especially in post-HF methods (see below). It brings a description of a many-electron problem in an occupation basis representation. With this formalism, one can transform properties of determinants to algebraic properties of creation and annihilation operators. Consider a state of a system with k spin-orbitals with $n^{\uparrow}$ spin-$\uparrow$ electrons and $n^{\downarrow}$ spin-$\downarrow$ electrons, then the following occupation basis notation is adopted

$$\{|n_1\, n_2\, \ldots\, n_k\rangle\ \} \qquad \forall\, n_i \in \{0,\, 1\},\ \text{ such that } n^{\uparrow},\ n^{\downarrow} \text{ is preserved,} \tag{7}$$

where n_j is the number of electrons occupying spin-orbital j, i.e. n_j determines whether the spin-orbital *is occupied or empty* in the given configuration.

Imagine modifying 7 by lifting the restriction on the total number of electrons. Such set would span the so-called Fock space, which describes scenarios with different number of particles in the system, see Kutzelnigg [1982].

Now let us introduce the notion of creation (annihilation) operator acting on spin-orbital j. These operators are defined by their action

$$\begin{aligned} a_j^{\dagger}\,|0\rangle &= |1\rangle\,, \\ a_j^{\dagger}\,|1\rangle &= 0, \\ a_j\,|1\rangle &= |0\rangle\,, \\ a_j\,|0\rangle &= 0. \end{aligned} \tag{8}$$

Here, † denotes the Hermitian adjoint (conjugate transpose in case of matrices and complex conjugate in the special case of scalars). The second quantization incorporates Pauli exclusion principle by postulating the following anticommutation relations

$$\begin{aligned} \{a_i^{\mu\,\dagger},\, a_j^{\nu}\} &= \delta_{ij}\delta_{\mu\nu}, \\ \{a_i^{\mu\,\dagger},\, a_j^{\nu\,\dagger}\} &= 0, \\ \{a_i^{\mu},\, a_j^{\nu}\} &= 0, \end{aligned} \tag{9}$$

where $\mu,\ \nu\ \in \{\uparrow,\, \downarrow\}$, where we included spin indices to show the anticommutation relation for the case of different spin projections.

The Slater determinant of the HF state (reference) represents the Fermi vacuum for a given system and excited determinants are formed by action of the second-quantized operators on this state.

Notice that in the second-quantized formalism, we can expand any operators in terms of creation and annihilation operators. In the rest of this text, such expansion is used for Hamiltonian.

Note that throughout this work, indices in subscript or superscript do not denote that these are the indices of covariant or contravariant tensors. Also, in this text, the word "tensor" may sometimes be used without the proper meaning from physics, since this term was adopted in computer science for plain multidimensional arrays, without elaborating on how they transform.

In the rest of this thesis, when we mention a molecular orbital, we mean a pair of spin-orbitals with identical spatial part and with the combined basis of $\{|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle\}$, c.f. a Kramers pair, in the relativistic case.

post-Hartee-Fock methods

Let E_{HF} be the Hartree-Fock energy for a system with the exact ground state energy E. Since HF state is an approximation, part of the energy which remained, unaccounted for, is called the *correlation energy*

$$E \quad = \quad E_{HF} + E_{\text{corr}}. \tag{10}$$

The name arises from the fact, that the HF method, a mean-field approach, does not properly take into account the correlations in the motion of electrons with different spin. Correlations of electrons with the same spin are to some extent incorporated through the Pauli exclusion principle and the Slater determinant form of the wave function.

The quantum chemical methods which calculate the correlation energy on top of the HF energy are called the post-HF methods. One of the variational post-HF methods, the configuration interaction (CI) method, employs the following wave function expansion

$$\Psi = c_{HF}\Phi_{HF} + \sum_s c_s \Phi^S_s + \sum_d c_d \Phi^D_d + \sum_t c_t \Phi^T_t + ..., \tag{11}$$

where Φ_{HF} is the Hartree-Fock state and Φ^S_s, Φ^D_d and Φ^T_t represent all possible Slater determinants created from HF state as single, double and triple excitations respectively. In practice, we terminate the expansion after a finite number of terms. The expansion coefficients can be determined in a way which minimizes energy by solving a general matrix eigenvalue equation, whose derivation is shown e.g. in Szabo and Ostlund [1996].

As the level of excitation increases, the description reaches higher accuracy and it gradually approaches the limit called the full configuration interaction (FCI) limit, where all excitations possible within a given set of orbitals are included. However, notice that the number of determinants required in the FCI expansion grows *factorially* with the number of electrons and orbitals. Alternatively, we may say that the associated Hilbert space grows exponentially. In practice, this typically yields FCI calculations with more than 18 orbitals intractable. Even though the FCI energy cannot be calculated for practical systems of moderate size, it represents the "exact" energy within the given finite MO basis and therefore sets a target for other post-HF methods. It follows that the quantum chemistry works almost exclusively with approximate methods. Even then, the scaling of ab-initio methods typically represents a serious limitation. This is the motivation behind the development of tensor-network based methods, a topic central to this work.

An important concept developed for capturing strong correlations within a subset of molecular orbitals is the complete active space (CAS) approach. Within this method, the MOs are classified as:

1. core orbitals, which are always doubly occupied
2. active, partially occupied orbitals
3. virtual orbitals, which are always empty

For n electrons in k active orbitals, we use the notation CAS(n, k). In case of the CASCI method, the configuration interaction wave function expansion contains all Slater determinants corresponding to FCI within the active orbitals. The CI coefficients are then variationally optimized. When also the one-electron basis is optimized simultaneously, we talk about the complete active space self consistent field (CASSCF) method. Notice that the selection of orbitals, which should be included in the active space is a complex process and in practice it often relies on chemical intuition of the researcher, see Golub et al. [2021].

As is apparent from the scaling of the number of terms of sum 11 with the number of orbitals, the problem of post-HF methods is often the huge computational effort, which makes it difficult to apply them efficiently to larger systems, even when employing the CAS approach.

In quantum chemistry, we differentiate between various types of correlation. A certain amount of electron correlation is already considered within the HF approximation, found in the electron exchange term describing the correlation between electrons with parallel spin, i.e. *the exchange correlation.* The rest of the electron correlation is commonly classified either as dynamical or static. The *dynamical correlation* due to the correlated movement of electrons, as they tend to avoid each other. *Static correlation*, on the other hand, arises from the proximity of energy levels in the spectrum of allowed states. Let us quote the textbook of Jensen [2006] on static correlation:

> *“The non-dynamical part is associated with electrons avoiding each other on a more ”permanent” basis, such as those occupying different spatial orbitals. The latter is also sometimes called a near-degeneracy effect, as it becomes important for systems where different orbitals (configurations) have similar energies.”*

Quantum chemical methods perform differently when applied on systems whose electronic structure is dominated by different types of correlation. Let us mention few examples from the wide class of ab-initio methods.

There are methods which typically perform well for systems dominated by static correlation, where many Slater determinants are necessary for even a qualitative description. These include, for instance the methods based on CASCI, or CASSCF. However, these methods usually fail to correctly account for dynamical correlation due to the limited size of CAS and as we mentioned above, the selection of CAS is a complex procedure, see Levine et al. [2021]. An example of approximated CASCI method is the density matrix renormalization group (DMRG), see Schollwöck [2005].

In contrast, the methods based on Møller–Plesset perturbation theory (see Band and Avishai [2013]), or Green's function theory usually perform better when treating systems dominated by dynamical correlation. A notable example is the coupled clusters technique, see Bartlett and Musiał [2007]. Nonetheless, these methods may struggle when multiple Slater determinants are needed for reference, for the case of nearly degenerate problems.

There are also methods which attempt to combine the approaches from both groups (Park et al. [2020]). For instance, the n-electron valence state perturbation theory (NEVPT) is a perturbative treatment, which applies Møller–Plesset perturbation theory to multireference CASSCF wave functions, see Kollmar et al. [2021].

Our focus

Despite an immense progress in recent decades due to novel methods and formidable advance in computational technology, the calculation of correlation energy still remains a challenge as of today. Moreover, there are classes of chemical systems which are still awaiting for a reliable description.

Above, we have defined the two notable classes of electron correlation. Considering the dynamical correlation, presently available methods, like coupled clusters, are able to treat this problem reasonably well. In particular the CCSD(T) method which includes single, double and triple excitations (treated perturbatively), is known as a "golden-standard" of quantum chemisty, which can provide spectroscopic accuracy for systems of a medium size, whose electronic structure is dominated by dynamical correlation. If, for the moment, we put aside the higher order polynomial computational cost of CCSD(T), which prevents its application on larger systems, we can, for the purposes of this text, consider the problem of dynamical correlation to be solved. Although several multireference generalizations of coupled cluster methods have been introduced Lyakh et al. [2011b], none of them is an obvious choice, they are usually applicable only to very small CASes and their development is still an active area of research.

The situation is indeed very different for systems suffering from quasidegeneracy. In fact, there is no "golden standard" method for static correlation except for CASCI/CASSCF, which are however limited by the active space size, due to their factorial scaling mentioned above. The quasidegeneracy (or near-degeneracy) problem appears with small energy gaps between the boundary orbitals. The more orbitals are close in energy, the more the electronic structure is dominated by static correlation. We call these systems strongly correlated. Such class of systems typically require a large active space even for a qualitative description. Notable multireference systems include transition-metal complexes, open-shell and excited electronic states, homolytic bond breaking/formation and transition states of chemical reactions, see Kurashige et al. [2013], Beinert et al. [1997], Fontecave [2006]. As an example of very important molecules which require large active spaces, much larger than CASCI/CASSCF can afford, we mention transition-metal complexes with multiple transition metal atoms, namely the Iron sulfur [Fe-S] clusters, e.g. [4Fe-4S]. They represent an ubiquitous motif in active sites of wide class of metaloenzymes which are responsible for many chemical processes in living organisms (e.g. photosynthesis or respiration). The need for large active spaces stems from open d-shells of multiple Fe atoms. The quantum chemical density matrix renormalization group method (QC-DMRG, see Wouters and Neck [2014]), as an example of tensor-network based approaches, has already proved itself as one of the computational methods capable of an accurate electronic structure description of such systems. Sharma et al. [2014]

In this we work, we present an important step towards a treatment of systems with even larger active spaces than those treated so far, in particular the massively parallel QC-DMRG method. As is shown below, our state-of-the-art implementation named MOLMPS (Figure 1) in fact opens the way for treatment of previously intractable systems, like the so-called “holy grail of quantum chemisty”, the FeMoco cofactor, see Lancaster et al. [2011]. This implementation is designed for supercomputer platforms in which it reduces the computational time significantly by simultaneous execution on thousands of CPU cores. Also, the distribution of memory among individual nodes allows for more accurate treatment of large systems, because DMRG suffers from substantial requirements on RAM memory size. To the best of our knowledge, this is the first implementation of DMRG which is truly massively parallel. On the examples of relatively large benchmark systems (including FeMoco cluster), we show that MOLMPS has a favourable scaling of runtime with respect to the number of CPU cores used.

MOLMPS

Figure 1: The logo of MOLMPS: our massively parallel DMRG program.

The DMRG wave function parametrization brings many advantages, and among others serves as an efficient tool for calculation of useful quantities, for instance various correlation measures, which can be used for a detailed analysis of interactions between parts of the system (orbital clusters). We use it to describe the chemical bonding structure, and we apply it on newly synthesized molecules with complex electron-deficient bonds.

Despite very large active spaces which our massively parallel DMRG implementation offers, except for the smallest molecules, CAS cannot comprise all the virtual orbitals and the missing dynamical electron correlation might play an important role. For this reason, we have developed the tailored coupled cluster method in our group, which combines DMRG with the coupled clusters Kinoshita et al. [2005], Veis et al. [2016]. The former is responsible for proper description of static correlation whereas the later captures the missing dynamical correlation.

In this work, we have extended this approach to the fully relativistic (i.e. four-component) domain. We have developed the new implementation and showed its performance on heavy diatomics. For compounds with heavy atoms, like thallium, arsenic and antimony, which are studied in this work, the relativistic description is crucial, since the movement of their electrons reaches speeds in tens of percents of the speed of light in vacuum. Hence, their masses and the resulting spectra are heavily altered. Let us mention a notorious example of an atom of gold. There, the 1s electron travels on average at 58% of the speed of light. The influence of such speeds on spectrum is so profound, that gold would be of silver color if our description did not include relativistic effects (see Pyykko and Desclaux [1979]). Moreover, because of spin-orbit coupling and the splitting of energy levels due to Zeeman effect, the relativistic systems typically exhibit even stronger multireference behaviour than their nonrelativic counterparts, see Schwarz [2010].

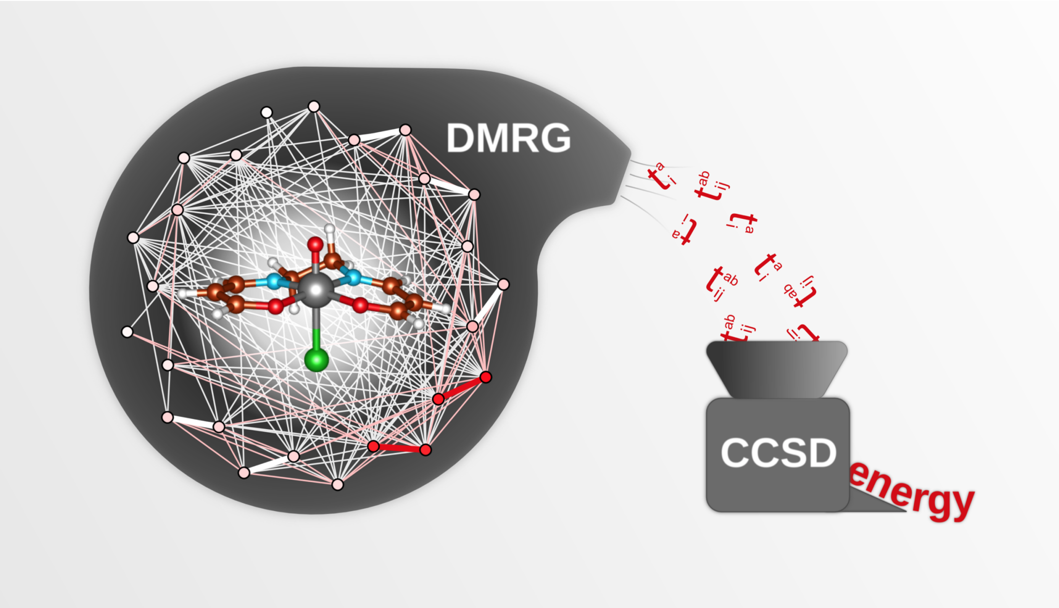

Figure 2: Artist's impression of the TCCSD method, featured frequently in posters and graphical abstracts of our group. Graphical design: A. Antalík 2018.

Let us now turn back to the nonrelativistic domain. There, in real-life applications (molecules with tens to hundreds of atoms), we meet the limits of high order polynomial scaling of coupled clusters (which we put aside in the beginning of this subsection), e.g. $\mathcal{O}(n^6)$, with n being the basis size for CCSD. Therefore, we took the path of *domain-based local pair natural orbital* (DLPNO) approach to the optimization of CCSD. We employ it in the CC part of the TCCSD calculation. DLPNO technique employs the concept of locality and compression of the virtual space and presents a near linear scaling approach. This leads to a new version of tailored coupled cluster method, the DLPNO-TCC, capable of treating systems with thousands of basis functions. In this work, we limit ourselves to showing that our implementation can capture more than 99.9% of correlation energy relative to standard TCCSD with much smaller computational costs. Application on previously intractable systems is beyond the scope of this work, but represents an important outlook.

Structure

This thesis is focused on developments of the aforementioned methods, their implementations, and demonstration of their performance on benchmark systems. In the following paragraphs, we specify the topics covered in this thesis.

First part of this work is concerned with the summary of the concept of MPS and the description of the DMRG algorithm (sections 1.1 and 1.2).

In the practical part of this work, we present our "MOLPMS" program for high performance computing (HPC) platform (section 1.3).

In section 1.4, we show how the correlation quantities obtained from matrix product states can be used to analyze complex chemical bonds.

Then, after summarizing the equations behind the coupled cluster method (sections 2.1), we turn to the problem of dynamical correlation and how it can be calculated on top of DMRG by tailored coupled cluster method (section 2.2).

After that, we present the first-ever application of DMRG-TCC to relativistic domain, by benchmarking it on heavy diatomics (section 2.3).

The last part of this work concerns an implementation of the highly optimized version of tailored coupled cluster method, the DLPNO-TCC. (section 2.5).

My personal contribution to the aforementioned projects is disclosed below in the section *List of publications* in the description under each publication. The list of publications is followed by attached fulltext of these publications, distributed as a separate document.

1. Density matrix renormalization group method

1.1 Matrix product states - MPS

Wave function of a general n-orbital state of a molecular system can be written as

$$|\Psi\rangle \;=\; \sum_{\{a_i\}} C_{a_1 a_2 a_3 \ldots a_n} |a_1 a_2 a_3 \ldots a_n\rangle , \tag{1.1}$$

where $C_{a_1 a_2 a_3 \ldots a_n}$ is in general a tensor of complex coefficients. Without imposing any symmetry restrictions (i.e. n_α, n_β, spin) there are in total 4^n coefficients, since vector subspace corresponding to one molecular orbital has a dimension 4, i.e. $a_i \in \{|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle\}$. Obviously, the memory requirements for storage of such a wave function parametrization scale exponentially.

Matrix product state (MPS) is a wave function parametrization which factorizes the coefficient tensor into a product of matrices

$$C_{a_1 a_2 a_3 \ldots a_n} \;=\; \mathbf{A}_1^{a_1} \mathbf{A}_2^{a_2} \ldots \mathbf{A}_{n-1}^{a_{n-1}} \mathbf{A}_n^{a_n}, \tag{1.2}$$

where $\mathbf{A}_1^{a_1}$ and $\mathbf{A}_n^{a_n}$ are row, column vectors respectively and between them are either square or rectangular matrices. The single-orbital basis index a_i in $\mathbf{A}_1^{a_1}$ means that we need to store 4 matrices for each orbital, i.e. $4n$ matrices in total. These indices are called physical indices. When the individual matrix product is carried out, the matrices are contracted over their row/column indices, termed virtual indices, as depicted on the scheme in Figure 1.1. There are multiple approaches to factorization of a wave function into form of contraction of tensors or matrices, called tensor networks. The most common schemes are depicted in Figure 1.1. These types differ by the number of virtual indices and by the topology of virtual indices, which connect the individual tensors. The dimension of virtual indices (bond dimension) is up to some extent arbitrary. With increasing bond dimension, the MPS representation of a wave function converges to the exact form of 1.1, see Schollwöck [2011]. Indeed, in practical applications, we typically seek to balance the computational requirements connected with the contraction of virtual indices with accuracy of the wave function representation given by its bond dimension.

Not just in quantum chemistry and solid state physics, but also in other fields Wang et al. [2021], matrix product states (also called tensor trains) have been used as a data structure capable of capturing various forms of correlations between subsystems. Matrix product states are particularly efficient in representation of correlations between subsystems corresponding to neighbouring sites/orbitals Schollwöck [2011]. Hence, in systems with nearest-neighbour interactions, i.e. 1D structure of correlations, MPS is an optimal choice of a wave function parametrization. Let us cite the Hubbard model and the Kondo problem as notable applications, see White [1992], Nakatani [2018]. Therefore, MPS is an appropriate representation for quantum chemical wave functions in which the correlation between orbitals can be localized to a certain degree, see Barcza et al.

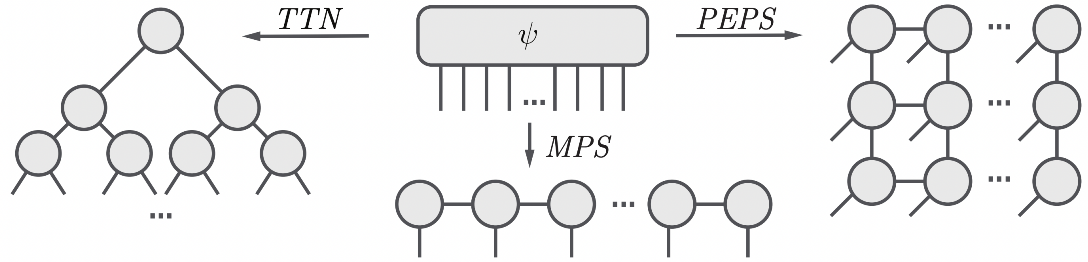

Figure 1.1: Different types of factorization of a wave function into form of contraction of tensors or matrices, called tensor networks. In this figure, the open-ended lines correspond to the physical indices a_i, while the lines which connect two circles (tensors) represent contractions over virtual indices, whose dimension can be chosen to reach desired accuracy of our approximation of Ψ. The factorization in form of matrix product is MPS and corresponds to the equation 1.2. The scheme including three-legged tensors (three virtual indices) is called the tree tensor network (TTN) and the scheme with up to four-legged tensors connected in a two-dimensional lattice is used in the projected entangled pair states (PEPS) technique. The general trend is that the order of computational complexity increases with the maximal number of virtual legs per tensor used in the scheme. Scheme from Biamonte and Bergholm [2017].

[2011]. It follows that a localized basis is preferable. The manual for the MPS-based DMRG method implementation “Block” in Orca program Neese [2017] advises users to use split-localized orbital basis (localization around bonds): *“It is highly recommended to start the calculation with split-localized orbitals”*. These orbitals localize the basis-dependent correlation to between the “neighbours” in the sense of atoms connected with substantial chemical bonds, see Bytautas et al. [2003] for more about split-localized orbitals and Olivares-Amaya et al. [2015] for their performance assessment in DMRG.

Indeed the question arises of how to find an optimal form of MPS 1.2 for a specicif problem. For this purpose, let us introduce the numerical technique of singular value decomposition (SVD) Golub and Reinsch [1970], which generalizes the concept of eigenvalue decomposition. SVD decomposes matrix M into a product form

$$M = U\Sigma V^\dagger, \tag{1.3}$$

where U, V are complex unitary matrices and Σ is rectangular diagonal matrix including only non-negative real numbers on its diagonal. If the original matrix M is real as well, then U, V are guaranteed to be real and orthogonal. The diagonal elements of Σ are called “singular values”. The number of positive singular values is equal to the rank of M. SVD is a powerful tool which found numerous applications, e.g. in signal processing and image compression.

In principle, MPS can be generated step by step by employing SVD on 1.1

$$\begin{aligned}
|\Psi\rangle = & \sum_{\{a_i\}} C_{a_1 a_2 a_3 \ldots a_n} \left|a_1 a_2 a_3 \ldots a_n\right\rangle = \sum_{\{a_i\}} C_{a_1; a_2 a_3 \ldots a_n} \left|a_1 a_2 a_3 \ldots a_n\right\rangle = \\
= & \sum_{\{a_i\}} \sum_{\nu_1} \underbrace{(\mathbf{A}_1^{a_1})_{\nu_1}}_{U} \underbrace{\sigma_{\nu_1}}_{\Sigma} \underbrace{B_{\nu_1; a_2 a_3 \ldots a_n}}_{V^\dagger} \left|a_1 a_2 a_3 \ldots a_n\right\rangle = \\
= & \sum_{\{a_i\}} \sum_{\nu_1} (\mathbf{A}_1^{a_1})_{\nu_1} C'_{\nu_1; a_2 a_3 \ldots a_n} \left|a_1 a_2 a_3 \ldots a_n\right\rangle = \sum_{\{a_i\}} C'_{\nu_1 a_2; a_3 \ldots a_n} \left|a_1 a_2 a_3 \ldots a_n\right\rangle = \\
= & \sum_{\{a_i\}} \sum_{\nu_1, \nu_2} (\mathbf{A}_1^{a_1})_{\nu_1} (\mathbf{A}_2^{a_2})^{\nu_1}_{\nu_2} C''_{\nu_2; a_3 a_4 \ldots a_n} \left|a_1 a_2 a_3 \ldots a_n\right\rangle . \qquad (1.4)
\end{aligned}$$

Here, ν_1 is the first virtual index. In the third row, we have absorbed σ_{ν_1} into B thus creating C'. If we again use SVD, this time on C', we obtain the last row. After repeated analogous applications of SVD, we arrive to MPS

$$|\Psi\rangle = \sum_{\{a_i\}} \sum_{\{\nu_j\}} (\mathbf{A}_1^{a_1})_{\nu_1} (\mathbf{A}_2^{a_2})^{\nu_1}_{\nu_2} (\mathbf{A}_3^{a_3})^{\nu_2}_{\nu_3} \ldots (\mathbf{A}_n^{a_n})^{\nu_{n-1}} \left|a_1 a_2 \ldots a_n\right\rangle , \qquad (1.5)$$

where the virtual indices ν_j are fully contracted. The form of MPS obtained in 1.5 is equivalent to the MPS factorization in equation 1.2 depicted also in Figure 1.1.

Notice that alternatively to equation 1.4, SVD could be applied to the last index on the right of tensor C, i.e. a_n and composing the remaining indices on the left to the B tensor. This SVD would yield matrices $B_{a_1 a_2 \ldots a_{n-1}; \nu_{n-1}} \sigma_{\nu_{n-1}} (\mathbf{A}_n^{a_n})_{\nu_{n-1}}$ and we could continue to generate the $\mathbf{A}$ matrices from the right.

It follows from the properties of the SVD decomposition, that it is not unique. There is some gauge freedom in the basis of vector spaces of U and V. This can be fixed along the MPS “matrix chain” to form a so-called left-canical form of MPS (here shown for the A matrices generated on the left) by demanding

$$\sum_{a_i, \nu_{i+1}} (\mathbf{A}_i^{a_i})^{\dagger\, \nu_i}_{\nu_{i+1}} (\mathbf{A}_i^{a_i})^{\nu_{i+1}}_{\nu_j} = \delta_{\nu_i \nu_j} . \qquad (1.6)$$

Analogous relation may be imposed on the $\mathbf{A}$ matrices generated on the right, leading to right-canonical MPS. For the DMRG algorithm, a relevant form of MPS is the so-called “mixed canonical” MPS (see Schollwöck [2011]), which is a combination of both left-canonical and right-canonical MPS, and in which we generate part of the matrices from the left and part from the right

$$|\Psi\rangle = \sum_{\{a_i\}} \mathbf{A}_1^{a_1} \mathbf{A}_2^{a_2} \ldots \mathbf{A}_l^{a_l} \, \mathbf{W}^{a_{l+1}, a_{l+2}} \mathbf{A}_{l+3}^{a_{l+3}} \ldots \mathbf{A}_n^{a_n} \left|a_1 a_2 \ldots a_n\right\rangle . \qquad (1.7)$$

The special tensor $\mathbf{W}^{a_{l+1}, a_{l+2}}$ in the middle is left with two physical indices. The reason why this is practical will become apparent in section 1.2.

The presented method of finding the exact MPS factorization is indeed useful only when the coefficient tensor C is representable in memory, which typically is not the case due to exponential character of its dimension with growing number of orbitals n. In practical approaches we do not start with the whole $C_{a_1 a_2 \ldots a_n}$, but we instead fill the matrices of 1.2 with some appropriate initial guess and then find optimal form of MPS iteratively, for instance, using the sweeping procedure of the DMRG method, which is presented in section 1.2. Also, it is useful to use sparse representation of tensors and matrices by exploiting various symmetries specific for a given Hamiltonian. This can lead to a significant reduction of storage requirements.

1.2 DMRG algorithm

Density matrix renormalization group method is a practical way of generating optimal matrix product state representation for lowest eigenstates of a certain class of Hamiltonians. It was introduced by White [1992] in solid-state physics for 1D lattice problems. For such problems, DMRG is still one of the most efficient approaches as of today. Since 1992, DMRG has been generalized for use on many different Hamiltonians including the quantum-chemical one, see the extensive review of Schollwöck [2005]. The method is particularly efficient for systems with somewhat local interactions, ideally with 1D topology. However, since its introduction in quantum chemistry by White and Martin [1999], it has been realized that DMRG is a very robust approach also for generic (non 1D) molecules, especially suitable for description of strong correlations in large active spaces.

The nonrelativistic quantum chemical Hamiltonian has the following second quantized form

$$H = \sum_{pq} h_p^q \, a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} \langle pq | \, rs \rangle \; a_p^\dagger a_q^\dagger a_s a_r, \tag{1.8}$$

where h_p^q and $\langle pq| \, rs\rangle$ represent one and two-electron integrals in a spin-orbital basis. Here, a_p is an annihilation operator acting on spin-orbital indexed p. In what follows, when we explicitely work with molecular orbitals (not spin-orbitals), we for simplicity assume the restricted basis (restricted HF). One-electron and two-electron integrals in (1.8) are defined as

$$\begin{aligned} \langle pq| \, rs\rangle &= \int d\mathbf{x}_1 \, d\mathbf{x}_2 \;\; \chi_p^\dagger(\mathbf{x}_1)\chi_q^\dagger(\mathbf{x}_2) \, \frac{1}{r_{12}} \, \chi_r(\mathbf{x}_1)\chi_s(\mathbf{x}_2), \\ h_p^q &= \int d\mathbf{x}_1 \;\; \chi_p^\dagger(\mathbf{x}_1) \; h(\mathbf{r}_1) \; \chi_q(\mathbf{x}_1), \end{aligned} \tag{1.9}$$

where χ_p denotes p-th spin orbital, r_{12} represents the distance between two points in space, and $\mathbf{x}$ is the composite spatial-spin variable. Here, r_{12} originates from Coulomb potential term in the first-quantized form of the electronic Hamiltonian. $h(\mathbf{r}_1)$ is the "core-Hamiltonian" describing kinetic energy contribution of individual electrons and their interaction with nuclei, whose spatial coordinates are frozen within the frame of the Born-Oppenheimer approximation, see Szabo and Ostlund [1996].

For later reference, we introduce the notion of reduced density matrix. For a system characterized by a pure state $|\boldsymbol{\Psi}\rangle$, the state of its subsystem X is encoded in its reduced density matrix ρ_X defined via a partial trace

$$\rho_X \;=\; \mathrm{Tr}_Y \, |\boldsymbol{\Psi}\rangle \, \langle\boldsymbol{\Psi}| \, , \tag{1.10}$$

where Y is the other subsystem, i.e. the rest of the system. The operation Tr_Y means carrying out the trace over Y.

In the rest of this section, we describe the consecutive steps of the QC-DMRG algorithm. Compared with other iterative methods, these steps are, however, relatively involved. The extent of truly self-consistent description of DMRG up to implementation details is beyond the scope of this work, see also the review of Schollwöck [2005] for an alternative description of DMRG and an implementation oriented text Battaglia et al. [2018]. In this section, we attempt to describe the basic workings of the method, assuming a nonrelativistic domain, for simplicity. Let us begin by listing the four main steps of each DMRG iteration, which are further decribed below

1. Formation of second-quantized pre-summed operators
2. Diagonalization of the superblock Hamiltonian using the Davidson method
3. Formation of renormalization matrix via SVD of wave function matrix
4. Blocking and renormalization of operators for the next iteration

Orbital ordering

At the beginning of QC-DMRG procedure, we have to map the molecular orbitals onto a 1D lattice. Orbitals should be ordered in a way which keeps mutually strongly interacting orbitals near each other. A sufficiently good orbital ordering can be determined using techniques based on quantum-information theoretical concept of mutual information as shown in Legeza and Sólyom [2003], defined also in equation 1.30 below. Such optimizations are typically carried out using preliminary low precision DMRG calculations preceding the production run and modern implementations, including our implementation, perform these automatically.

Basis and blocks

Even though other variants exist, in practical QC-DMRG, the Hamiltonian diagonalization is performed within a vector space composed as a tensor product of four blocks (left block, left site, right site, right block).

The DMRG sweeping procedure starts with just a single orbital in the left block. In each DMRG iteration, left block is enlarged by one orbital, as depicted schematically in Figure 1.2 At the point when the maximal size of left block is reached and right block size is just 1 (one orbital), the direction is reversed and the right block is enlarged. The sweeping is then repeated until the energy converges. Notice that the two site-blocks in the middle are always composed of just one orbital each.

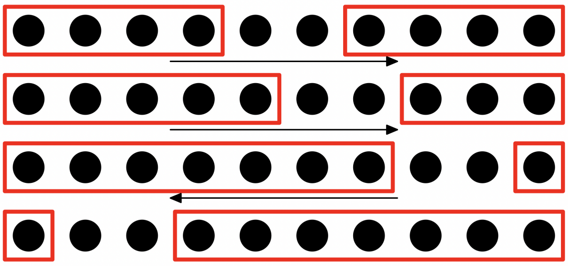

Figure 1.2: The sweeping procedure of two-site DMRG illustrated on the example of 10 orbitals in active space. Each black dot corresponds to a molecular orbital. Red rectangles delimit the subsystems (blocks) L and R, with two single-orbital blocks in the middle. The arrows represent direction of a given "sweep". This scheme shows only frames corresponding to configurations in 4 selected iterations from the total of 13 possible configurations of the procedure for this number of molecular orbitals involved. (The direction is taken into account). Image from Szalay et al. [2015].

As we mentioned in section 1.1, a vector space of a single molecular orbital is spanned by four basis states $\{|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle\}$. In this basis, a full matrix representation of creation operators takes form

$$a_{\downarrow}^{\dagger} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \end{pmatrix}, \qquad a_{\uparrow}^{\dagger} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \tag{1.11}$$

When we start in the first iteration with such operators acting on the left block and then enlarge it via a full Kronecker product with identity $a_{\uparrow}^{\dagger} \otimes I$ to obtain its representation for the enlarged block, its matrix grow from 4×4 to 16×16. Obviously, if we keep enlarging the matrix in this fashion, we might soon run out of storage space.

Here the elementary principle of DMRG renormalization comes into play. This is to determine optimal left and right block many-electron basis with dimension bounded by a threshold M, called bond dimension (corresponds to the same term defined in section 1.1).

Consider having the left block of dimension M and of L orbitals, which is then increased by adding one orbital (site-block). The full vector space of the new enlarged left block is spanned by $\{|l_L\rangle\} \otimes \{|q\rangle\}$, where $\{|l_L\rangle\}$ and $\{|q\rangle\}$ are the bases of the former left block and the added site-block respectively. To retain the bond dimension of M, the enlarged basis is truncated by the action of renormalization matrix $\mathbf{O}$

$$|l_{L+1}\rangle = \sum_{(l_L\, q)} \mathbf{O}^{l_{L+1}}_{(l_L\, q)} \; |l_L\rangle \otimes |q\rangle\,, \tag{1.12}$$

where $(l_L\, q)$ represents composite index of former basis indices l_L and q. For later convenience, we denote the "input" index (column index) in subscript and the "output" index (row index) of $\mathbf{O}$ in superscript. Here $\mathbf{O}$ is a rectangular $M \times 4M$ matrix, which we define further in this section.

Instead of storing the explicit form of many-electron basis determinants in each iteration, we store the representation of second-quantized operators and the renormalization matrices. These operators are needed for performing the action of Hamiltonian on the wave function, while the explicit basis determinants are not required for the process of finding the ground state.

All the operator matrices which we need for an action of the Hamiltonian on a block with single orbital can be generated from creation operator matrices in equation 1.11 by the means of matrix multiplication, Hermitian conjugate and scaling by one-electron or two-electron integrals.

When we perform the action of Hamiltonian on a block of multiple orbitals, we need to take into account that the renormalized (truncated) basis of the block is incomplete. Consider an operator P represented by a product of creation and annihilation operators. As we work in truncated basis, typically, P cannot be generated as a product of matrices of individual creation and annihilation operators. Instead, starting from a single-orbital in the first iteration and then in every further iteration, we generate a renormalized form of P for the enlarged block and we save it. Therefore, before moving to the next iteration, we have to prepare (by "blocking", explained below) and renormalize all the intermediaries necessary for forming P and other operators required for the action of Hamiltonian in the next iteration.

Operators and pre-summation

Projecting the Schrödinger equation onto the product space of bases $\{|l_L\rangle\} \otimes \{|q_1\rangle\} \otimes \{|q_2\rangle\} \otimes \{|r_R\rangle\}$ corresponding to the left block, two site-blocks and the right block, we obtain the effective equation

$$\begin{aligned} H\,\mathbf{\Psi} &= E\,\mathbf{\Psi}, \\ \text{where} \quad |\mathbf{\Psi}\rangle &= \sum_{l_L,\,q_1,\,q_2,\,r_R} (\mathbf{\Psi}\,)_{l_L,\,q_1,\,q_2,\,r_R}\; |l_L\rangle \otimes |q_1\rangle \otimes |q_2\rangle \otimes |r_R\rangle\,, \end{aligned} \tag{1.13}$$

and where $(\mathbf{\Psi}\,)_{l_L,\,q_1,\,q_2,\,r_R}$ are the expansion coefficients of the wave function.

The most computationally demanding part of the DMRG procedure is the application of Hamiltonian on the wave function, due to numerous matrix-matrix multiplications. To decrease the computational complexity, most DMRG implementations work with the so called presummed (or partially summed) operators, see Xiang [1996]. These are intermediares formed by summing operator matrices scaled by MO integrals. Consider, for instance, an operator P corresponding to the action of any two creation operators of ↑ electrons within the left block

$$P_{rs} \;=\; \sum_{p,\,q\,\in\,\text{left}} \langle pq|rs\rangle\; a^\dagger_{p\uparrow} a^\dagger_{q\uparrow}, \qquad r,\, s \notin \text{left}. \tag{1.14}$$

Here P_{rs} is a pre-summed operator with two free indices r, s pointing outside the left block (i.e. $\{P_{rs}\}$ is a set of pre-summed operators).

In order to represent a Hamiltonian term (terms) with four creation/annihilation operators, an application of operator P_{rs} on the left block has to be accompanied by a concurrent application of two other creation/annihilation operators acting on other blocks, in particular $a^\dagger_{r\uparrow} a^\dagger_{s\uparrow}$

$$H^{\uparrow\uparrow\uparrow\uparrow}_{\text{two-el.}} \;\ni \sum_{p,\,q\,\in\,\text{left},\,r,\,s\,\notin\,\text{left}} \langle pq|rs\rangle\; a^\dagger_{p\uparrow} a^\dagger_{q\uparrow} a^\dagger_{r\uparrow} a^\dagger_{s\uparrow} = \sum_{r,\,s\,\notin\,\text{left}} P_{rs}\, a^\dagger_{r\uparrow} a^\dagger_{s\uparrow}. \tag{1.15}$$

Notice the correspondence of the physical indices of $a^\dagger_{r\uparrow}a^\dagger_{s\uparrow}$ with the free indices of P_{rs}. Once a loop over all possible r, s is performed, then all the two-electron terms with two creation operators in the left block have been applied to the wave function. (For a given $\uparrow\uparrow\uparrow\uparrow$ spin combination.)

Notice that the use of partially summed (pre-summed) operators saves the number of matrix-matrix multiplications needed for the $H\left|\mathbf{\Psi}\right\rangle$ action (see below).

The whole Hamiltonian application ($H\left|\mathbf{\Psi}\right\rangle$) is finished once all similar loops over free indices are performed for one, two, three, and four creation/annihilation operators in the left block, and the action corresponding to their respective free indices is applied to the other blocks, including all contributing combinations of spin indices.

The pre-sumation can be also performed for the right block instead of the left block. This option is used when right block is larger than the left block, to contract the maximal possible number of terms in the partial sum.

Notice that the operators that build up the Hamiltonian always have at most two free indices. This holds both for pre-summed and for standard operators. For this to hold, when more than two creation/annihilation operators act on one block, the pre-summation then has to be performed on this block. Otherwise the largest block is chosen for pre-summation

$H\left|\mathbf{\Psi}\right\rangle$ and the Davidson algorithm

When applying the Hamiltonian on the wave function, we do not form the matrix of the Hamiltonian from equation 1.13. Instead, we employ the tensor product form of vector space. Consider an example of the operator combination from equation 1.15, where P_{rs} acts on the left block, and the case when operators $a^\dagger_{r\uparrow}a^\dagger_{s\uparrow}$ act in the right block, i.e. $P_{rs}\otimes I\otimes I\otimes a^\dagger_{r\uparrow}a^\dagger_{s\uparrow}$, with identities in the middle. Then, the product form of the Hamiltonian matrix reads

$$\begin{aligned} \left\langle l'_L\right|\left\langle q'_1\right|\left\langle q'_2\right|\left\langle r'_R\right|\left(P_{rs}\otimes I\otimes I\otimes a^\dagger_{r\uparrow}a^\dagger_{s\uparrow}\right)\left|l_L\right\rangle\left|q_1\right\rangle\left|q_2\right\rangle\left|r_R\right\rangle &= \\ = \left\langle l'_L\right|P_{rs}\left|l_L\right\rangle\,\left\langle r'_R\right|a^\dagger_{r\uparrow}a^\dagger_{s\uparrow}\left|r_R\right\rangle\,\delta^{q'_1}_{q_1}\,\delta^{q'_2}_{q_2}, & \qquad (1.16) \end{aligned}$$

where we have omitted “$\otimes$” between the basis states for better readability. Using the approach in equation 1.16, the application of the whole Hamiltonian term from equation 1.15 on the wave function can be compiled by composing only operator matrices acting in the bases of individual blocks.

Most DMRG implementations use either Davidson algorithm or Lanczos algorithm to find low-lying eigenstates of the superblock Hamiltonian. These numeric diagonalization methods are based on iterative application of the matrix (Hamiltonian) on a trial vector (trial wave function), i.e. performing $H\left|\Psi_t\right\rangle$. By repeated application of Hamiltonian on the trial wave function, these iterative methods build the so-called Krylov space

$$\mathcal{K}_r\left(H,\,\Psi_t\right) \;=\; \mathrm{span}\left\{\Psi_t,\;H\Psi_t,\;H^2\Psi_t,\;\ldots,\;H^{r-1}\Psi_t\right\}. \qquad (1.17)$$

The approximations to the low-lying eigenstates are then formed by minimizing the residual (error in the result) over $\mathcal{K}_r$. Since $\mathcal{K}_r$ spans a basis of the space where H acts (its domain), it follows that the method converges in at most M iterations (in exact arithmetic), where M is the bond dimension. In practice, M is large, but the iterative process reaches sufficient accuracy already far earlier. For more information about iterative diagonalization methods based on matrix-vector multiplication, see Crouzeix et al. [1994].

For instance, in QC-DMRG with chemical accuracy requirements, the Davidson algorithm typically converges to the ground state in dozens of $H\left|\mathbf{\Psi}\right\rangle$ steps. However, as mentioned above, $H\left|\mathbf{\Psi}\right\rangle$ is the most demanding step of the DMRG procedure. Hence, a lot of attention turns to optimization of this step. As described below, in our in-house implementation, we employ several optimization techniques accompanied by parallelization to reduce the scaling. In fact, the computational demands of $H\left|\mathbf{\Psi}\right\rangle$ step are so large, that a QC-DMRG calculation without the most important optimizations easily becomes intractable, see Legeza and Sólyom [2003].

Blocking

Blocking is a term used for enlarging the basis by a tensor product with a single orbital site-block q, i.e. $\{|l_L\rangle\} \otimes \{|q\rangle\}$. With the enlargement of the basis, we form also the operators and operator intermediaries for the next iteration. Consider an example of blocking of the pre-summed operator P_{rs} from equation 1.15. To build an analogous operator $\tilde{P}_{rs}$ for the enlarged block from available intermediaries, we have to sum over following Kronecker product terms corresponding to three possible ways of composing $\tilde{P}_{rs}$ from two operators (intermediaries)

$$\begin{aligned} \tilde{P}_{rs} = & \quad P_{rs} \otimes I \quad + \\ + & \quad \langle qq|rs\rangle \; I \otimes \underbrace{a^\dagger_{q\uparrow} a^\dagger_{q\uparrow}}_{=0} \quad + \\ + & \quad Q_{rs}\sigma_z \otimes a^\dagger_{q\uparrow}, \qquad r,\, s \notin \text{enlarged}, \end{aligned} \tag{1.18}$$

where σ_z is a fermionic phase operator described further in this section. Here the Q_{rs} intermediary is an analogy to P_{rs} with a single creation operator, i.e. a pre-summed operator with one $a^\dagger_{p\uparrow}$ partially contracted over left block according to

$$Q_{rs} \;=\; \sum_{\substack{p\,\in\,\text{left}\\ q\,\in\,\text{left site}}} \langle pq|rs\rangle \; a^\dagger_{p\uparrow}. \qquad r,\, s \notin \text{enlarged}. \tag{1.19}$$

Notice that when composing $\tilde{P}_{rs}$ for enlarged left block, we only need P_{rs}, Q_{rs} intermediaries available from previous iteration along with the simple single-site operators. More importantly, we have completely avoided performing sums over any of p, q indices while still attaining the advantage of having the pre-summed operators available for the new iteration.

In blocking, the dimension of operator increases from $M \times M$ to $4M \times 4M$. (Values without using any sparsity optimizations.)

Formation of renormalization matrix

Our description of DMRG algorithm would be incomplete without explaining how the renormalization matrix $\mathbf{O}$ is obtained. We already mentioned it above in equation 1.12, where we used $\mathbf{O}$ for basis truncation.

In DMRG, the renormalization matrix for enlarged block can be formed using two possible approaches, whose result is equivalent. In both cases, we start by a reshape of the wave function coefficient tensor $(\boldsymbol{\Psi})_{l_L,q_1,q_2,r_R}$ from equation 1.13 into a matrix $(\boldsymbol{\Psi})_{(l_L,q_1),(q_2,r_R)}$ using two composite indices.

Then, we choose one of the possible approaches. The first approach performs singular value decomposition (equation 1.3) on the matrix $(\boldsymbol{\Psi})_{(l_L,q_1),(q_2,r_R)}$. Along with the singular values of this matrix, we obtain its left (right) singular vectors as rows (columns) of matrix U (or $V^\dagger$) from equation 1.3. These singular vectors are then used to form the renormalization matrix $\mathbf{O}$, as described below.

The second approach forms the reduced density matrix of the left (right) enlarged subsystem, which reads

$$\begin{aligned} \rho^L &= \boldsymbol{\Psi}\boldsymbol{\Psi}^\dagger, \\ \rho^R &= \boldsymbol{\Psi}^\dagger\boldsymbol{\Psi}. \end{aligned} \tag{1.20}$$

We then diagonalize ρ^L (or ρ^R). Along with the eigenvalues of the reduced density matrix, we obtain its left (right) eigenvectors.

From this point onward, the two approaches proceed with common steps.

Now we select a set of M of the left (right) singular vectors (eigenvectors) and from them we form the renormalization matrix $\mathbf{O}$. We choose the vectors corresponding to M largest singular values (eigenvalues). For the case of left subsystem, $\mathbf{O}$ is now a $M \times 4M$ matrix.

Notice that in this step, the new bond dimension M is arbitrary. Thus, we may e.g. gradually increase the accuracy in the course of the calculation, as we approach convergence. Moreover, consider defining the "truncation error" as the sum of discarded singular values (eigenvalues). Then, we may set the threshold for maximal truncation error as an input parameter of the DMRG calculation and in each iteration, update the bond dimension automatically, such that we meet this error restriction. This is the most systematic approach to DMRG accuracy, the Dynamical Block State Selection approach (DBSS) by Legeza et al. [2003].

It has been shown by White [1992] that the form of truncation matrix defined above minimizes the norm of error of the wave function, i.e. $\| |\Psi_{\text{exact}}\rangle - |\Psi_{\text{DMRG}}\rangle \|$. The relation between the wave function error (proportional to truncation error) and the error in energy has been studied empirically by Keller and Reiher [2014], Hubig et al. [2018]. Figure 1.3b shows a typical empirical relation of the two measures of error, showing a nearly straight line in a log-scale chart. Consistently with these findings, in practical QC-DMRG calculations, we expect the exponents of truncation error and of the acceptable energy error in hartrees to be directly proportional. Consequently, we may extrapolate the DMRG energies towards the zero-truncation error case representing the exact energy, see Hubig et al. [2018].

Renormalization (truncation)

Before transition to next iteration, the operators created in blocking need to be renormalized, in order to retain the required bond dimension M. Above, we have shown how the renormalization matrix $\mathbf{O}$ is formed. Now we can use this matrix to renormalize the operators.

Without employing sparse representation of operators, after blocking, we consider an operator P which acts on the enlarged left subsystem, and whose matrix has the dimension of $4M \times 4M$. Now we transform P to the truncated basis using the $M \times 4M$ renormalization matrix $\mathbf{O}_L$ for the left subsystem

$$\tilde{\mathbf{P}} \;=\; \mathbf{O}_L \, \mathbf{P} \, \mathbf{O}_L^{\dagger}. \tag{1.21}$$

The resulting operator $\tilde{\mathbf{P}}$ has the dimensions of $M \times M$ in its full form.

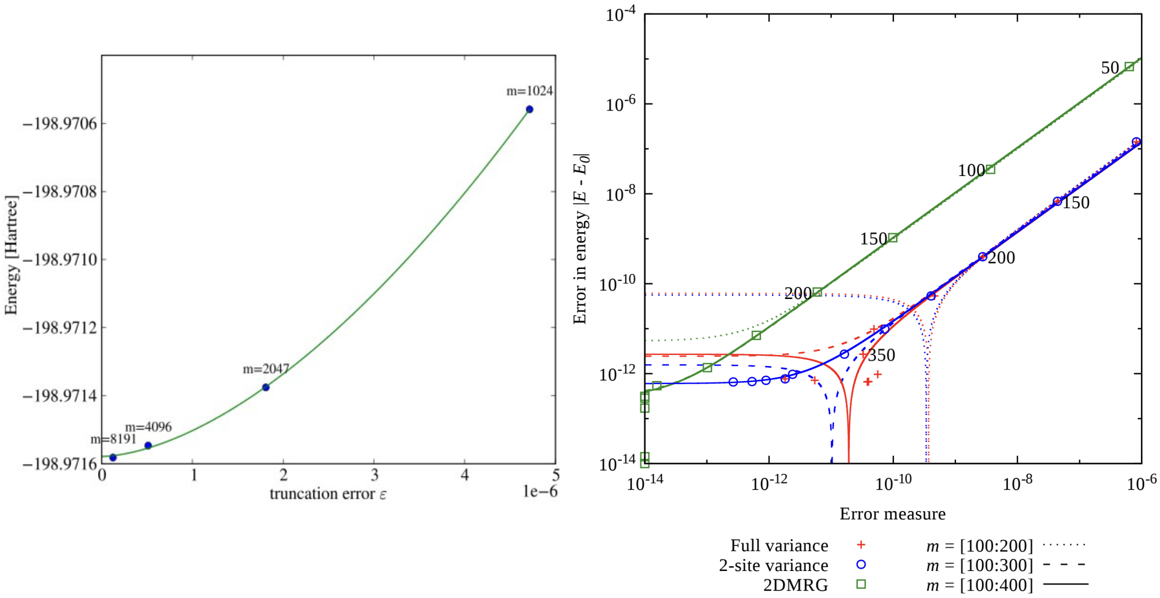

(a) The convergence of DMRG energy with respect to increasing bond dimension m tested on the ground state of F_2 molecule. Chart from Keller and Reiher [2014].

(b) The error in energy with respect to truncation error of two-site DMRG (2DMRG). The chart features logarithmic axes and m is for bond dimension. Presented dependency is for a model with long-range Coulomb-like interactions. Reprinted figure with permission from Hubig et al. [2018]. Copyright 2018 by the American Physical Society.

Figure 1.3: The dependency of DMRG energy error on truncation error, which is calculated as a sum of discarded singular values. Further error measures are also included in the chart on the right. Such empirical dependencies can be used to estimate the truncation error necessary for achieving the required accuracy in energy.

Symmetry sectors

As we mentioned above, a proper handling of sparsity due to quantum symmetries is essential for an efficient QC-DMRG implementation. Our current implementation described in section 1.3 employs U(1) symmetry, however, SU(2) symmetry (spin-adapted version) is under development.

The aforementioned U(1) symmetry corresponds to the conservation of total number of spin-↑ and spin-↓ electrons. All left and right block basis states are grouped into symmetry sectors by the number of ↑ ($n_\uparrow$) and ↓ ($n_\downarrow$) electrons. Then, only the nonzero symmetry blocks of operator matrices are stored, leading to a sparse representation.

The dense sectors of the matrix of operator are labelled by the quantum numbers of the corresponding symmetry sectors. These matrices along with the rules for composing the symmetry sectors together represent the operator, yet with much smaller computational requirements for their storage and application, compared with the full form. Our operator storage scheme is described further in section 1.3 and in Figure 1.4.

It is important to notice that in case of the nonrelativistic QC-DMRG, $n_\uparrow$ and $n_\downarrow$ are good quantum numbers and all the symmetry sectors of a single site (molecular orbital) are one-dimensional. This simplifies the work with operators and wave function considerably. For instance, a sector of $(\mathbf{\Psi})_{l_L,q_1,q_2,r_R}$ can be treated as a matrix instead of a four-index tensor.

Jordan-Wigner mapping

In the first step of QC-DMRG, we map molecular orbitals onto a 1D lattice. The particular choice of ordering has another consequence. Aside from a huge influence on DMRG accuracy, our choice fixes the Jordan-Wigner mapping from fermionic creation and annihilation operators in MO basis onto their representation on 1D lattice. In principle, the Jordan-Wigner mapping effectively maps these fermionic operators onto “bosonic” spin-operators acting on 1D lattice, while retaining the correct anticommutation relations of the mapped operators. By the term “bosonic” here we mean that an operator in our implementation is “bosonic” in the sense that the C++ operator class op_class which represents it does not automatically incorporate fermionic anticommutation rules. For more information on Jordan-Wigner mapping and its application on tensor networks, see Jordan and Wigner [1928], Stoudenmire [2016], Hauschild and Pollmann [2020], Nakatani and Chan [2013].

In QC-DMRG, the Jordan-Wigner mapping is implemented using phase operators σ_z, which appeared in equation 1.18 above, in blocking. The notation stems from the fact that the matrix of phase operators can be built from Pauli matrices. Its action corresponds to a simple switch of signs for sectors with odd number of electrons. For the Jordan-Wigner mapping, we adopt the convention of numbering the lattice sites with numbers increasing from the left to the right.

When creation/annihilation operator acts on site j on the lattice, the Jordan-Wigner mapping requires an application of a phase operator σ_z acting on each orbital i such that $i < j$. Even though this brings certain non-locality, the application of phase operators is still fully compatible with the aformentioned concept of pre-summation and other optimizations.

Warm-up procedure

During the first sweep of QC-DMRG, the renormalized representation of basis and operators is not yet available for the right block. Therefore, we employ a warm-up procedure to generate an initial guess, whose quality is reflected in the rate of convergence of the following DMRG. In our implementation, we use the so-called CI-DEAS procedure, described in Szalay et al. [2015]. First, we describe the basic version thereof, the DEAS procedure (Dynamically Extended Active Space).

For molecular systems treated by DMRG, it is typical that the state of their subsystems are quite strongly correlated with the environment (the rest of the system). To recreate such conditions, we need to find representation of the right block (the environment), which maximizes the von Neumann entropy of the left block. For a definition of von Neumann entropy, see equation 1.27 in section 1.4. At this point it suffices to mention that for a single orbital, its correlation with the rest of the system is proportional to its von Neumann entropy.

When composing the environment basis, we are limited by the bond dimension M. In practice, the maximization of the left block entropy is achieved by starting from a HF determinant and then creating an environment basis by generating all possible excitations of such molecular orbital of the right block, whose von Neumann entropy is the largest. Further combinations are generated by increasing the set of orbitals excited from HF reference, always by the most correlated orbital, until M states are obtained.

Here we assume that we know the von Neumann entropy for individual orbitals. It can be obtained e.g. by calculating it beforehand in a low-precision preliminary DMRG calculation. This can be performed automatically before the production run. Section 1.4 describes how the von Neumann entropy is calculated from MPS.

The version of the DEAS warm-up improved for the use in chemistry is the CI-DEAS, based on the CI expansion in equation 11. Here, we generate the states in a different order. First, we use DEAS, but we retain only states corresponding to slater determinants singly excited from the HF state. Then, we use DEAS again while keeping only the doubly excited states. This way, we increase the excitation level of generated states systematically until the desired amount of M states is obtained.

A representation of both pre-summed and standard operators is then built within this basis by employing the knowledge of how the creation/annihilation operators act on Slater determinants.

Further properies of QC-DMRG

In the following, we cite several notable properties of DMRG method from Schollwöck [2005] and from Chan [2004].

Considering the scaling of DMRG, the dominant term in computational complexity is that of H diagonalization $\mathcal{O}\left(k^3M^3\right)$, where k is the number of orbitals in the active space. The same computational complexity holds also for renormalization. The pre-summation scales as $\mathcal{O}\left(k^4M^2\right)$. The scaling of total RAM memory requirements is $\mathcal{O}\left(k^2M^2\right)$, see Chan [2004].

The QC-DMRG method is variational and size-consistent. As we mentioned, it is capable of capturing static correlation with high accuracy. In contrast to some other available methods like coupled clusters, DMRG is not suitable for systems dominated by dynamical correlation. In practical calculations, DMRG is typically useful for active space size up to ≈ 60 orbitals. For larger active spaces, the computational requirements may be too high, but this of course depends on available hardware and multiple other factors.

Last but not least, we notice the correspondence between DMRG renormalization matrices obtained from consecutive iterations and the MPS matrices. The

MPS matrices may be obtained by a simple reshape of renormaliation matrices. For those **O** matrices used for truncation of left subsystem, the reshape reads

$$\mathbf{O}^{l_{i+1}}_{(l_i\, a_i)} \quad \longmapsto \quad (\mathbf{A}^{a_i}_{i})^{l_{i+1}}_{l_i}\,, \tag{1.22}$$

and analogous relation can be written for the right subsystem. After reshaping matrices for both sides, we arrive to the "mixed canonical" form of matrix product state defined above in equation 1.7 of section 1.2.

1.3 Implementation: Massively parallel program MOLMPS

In the DMRG algorithm, there are several steps which can be performed in parallel to speed up the computation. In this section, we identify such steps and present benchmarks of our massively parallel implementation called MOLMPS, which is designed to exploit modern high performance computing (HPC) platforms.

Several approaches to DMRG parallelization have been developed over the time. Some of them use shared memory, e.g. Hager et al. [2004], and some use distributed memory architectures, e.g. Kurashige and Yanai [2009], Chan [2004].

The distributed approach of Chan [2004] employs parallelization over different terms of the sum in the Hamiltonian. Tasks corresponding to orbital indices of renormalized operators from these terms are assigned to individual CPU cores.

Alternative approach of Kurashige and Yanai [2009] parallelizes over symmetry sectors of the basis. A novel approach has been proposed by Chan et al. [2016], which uses the concept of matrix product operators (MPO) and leads to a different contraction scheme which is shown to minimize the communication between computational nodes, which is often the bottleneck of parallelization techniques.

Presently, there are several notable large DMRG codes well known to the community, see them described in Wouters and Neck [2014]. Most of them are open-source and each is useful for a different class of applications ranging from solid state physics to quantum chemistry and nuclear theory. However, to the best of our knowledge, no previous DMRG implementation is truly massively parallel, so that it would be able to use thousands of computational CPU cores efficiently, i.e. with reasonable scaling of computational time with respect to the number of cores. Our C++ implementation, called MOLMPS Brabec et al. [2020], is the first attempt to truly port the quantum chemical DMRG to the supercomputer platform.

Our approach to DMRG parallelization is, similarly to shared memory scheme Hager et al. [2004], based on merging of symmetry and operator sector loops. We use the global memory model (GM) (described below), which relies on fast internode connection. Such connection is typically, at least for most supercomputers from TOP500 list, provided by high-throughput hardware using InfiniBand standard, an alternative to the Ethernet protocol, see Gara et al. [2011].

In MOLMPS, the storage of operators is based on sector decomposition, which we already mentioned in section 1.2. The storage scheme is depicted in Figure

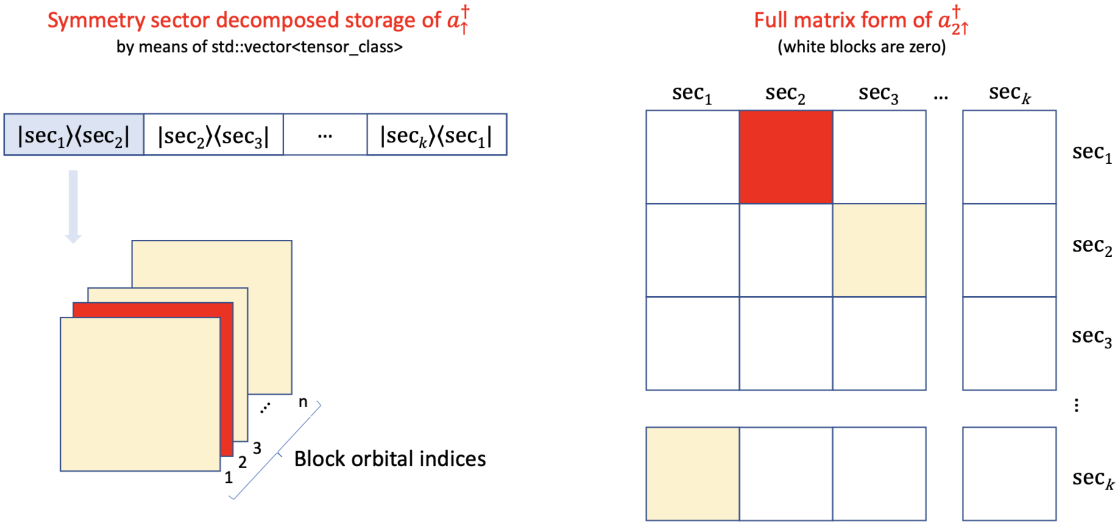

Figure 1.4: Memory representation of individual symmetry sectors for a creation operator $a_{2\uparrow}^\dagger$ acting on orbital 2. One sector of $a_{2\uparrow}^\dagger$ is highlighted in red. On the left, we see how the matrices corresponding to the same symmetry sector, but to different orbital ($a_{i\uparrow}^\dagger$) are kept in memory as slices of a 3D array (tensor_class in our implementation). On the right, we see how these dense sector blocks compose the full matrix form of $a_\uparrow^\dagger$, which is in fact never formed in full. Only the nonzero sectors are saved and all the treatment thereof is realized sector-wise. The matrices are pictured as square for convenience, but in general are rectangular with dimensions equal to the number of basis states contributing to input sector (number of columns) and output sector (rows). Table from our work Brabec et al. [2020].

1.4. The wave function is also stored in sector-decomposed form, which in the case of nonrelativistic two-site DMRG (see Figure 1.2) reads

$$|\Psi\rangle = \sum_{abcd} \sum_{l\in a, r\in d} \left(\mathbf{\Psi}^{abcd}\right)_{l_a, r_d} |l_a\rangle \otimes \left|q_1^b\right\rangle \otimes |q_2^c\rangle \otimes |r_d\rangle \,, \tag{1.23}$$

where a, b, c, d are indices of symmetry sectors of the left block, two middle sites and the right block. $\{|l_a\rangle\} \subset \{|l\rangle\}$ is the basis of left block sector a, $\{|r_d\rangle\} \subset \{|r\rangle\}$ is the basis of sector d of the right block basis and $\left|q_1^b\right\rangle \otimes |q_2^c\rangle$ is a Kronecker product of bases of two orbitals in the middle.

Notice that since $\left|q_1^b\right\rangle$ is a basis of single sector of single-orbital basis, in nonrelativistic case, its dimension is equal to one. Therefore each $\mathbf{\Psi}^{abcd}$ is just a matrix (2D array). This property simplifies the work with $\mathbf{\Psi}^{abcd}$ in the program.

The sparsity among sectors depicted in Figure 1.4 stems from the conservation of total number of electrons and the spin projection (or equivalently $n_\uparrow^{\mathrm{tot}}$, $n_\downarrow^{\mathrm{tot}}$) in sectors a, b, c, d

$$\begin{aligned} n_\uparrow^{\mathrm{tot}} &= n_\uparrow^a + n_\uparrow^b + n_\uparrow^c + n_\uparrow^d, \\ n_\downarrow^{\mathrm{tot}} &= n_\downarrow^a + n_\downarrow^b + n_\downarrow^c + n_\downarrow^d, \end{aligned} \tag{1.24}$$

but also from the way how operators act. For instance, the sector blocks of spin-up creation operator $a_\uparrow^\dagger$ acting on sector a of left block are non-zero only when the equality ${n_{\uparrow}}^{a}_{\mathrm{output}} = {n_{\uparrow}}^{a}_{\mathrm{input}} + 1$ holds, i.e. particle is added. Moreover, some slices of 3D arrays (see Figure 1.4) corresponding to how operator acts on individual orbitals can be zero, for example when multiplied by the corresponding integral from the equation 1.8, which can be zero e.g. due to point group symmetry.

When two or more creation/annihilation operators act on orbitals in the same block, we can sum over such contributions to the given block to form a so-called pre-summed operator. This way we avoid repeated summations of the same terms. For more information about the definition and the formation of pre-summed operators, see equations 1.14, 1.15 section 1.2. Pre-summed operator has a corresponding number of free indices, i.e. those which were not summed over. The free index corresponds to the third index of the matrix of a given dense sector of the operator, as depicted on the left side of Figure 1.4.

As we mentioned above, the parallelization in our implementation is based on merging the loops over orbital indices with the loops over sectors and thus creating a huge large task list over all Hamiltonian terms, when performing the demanding $\left|\tilde{\Psi}\right\rangle = H\left|\Psi\right\rangle$ in the Davidson procedure. The parallelization and the internode communication is implemented using the Message passing interface (MPI) library. During the Hamiltonian application, each MPI process has its own copy of $\tilde{\Psi}$. Then, we use the MPI *reduce* feature to collect the resulting wave function sector matrices $\tilde{\mathbf{\Psi}}^{efgh}$ and perform their addition. Consider the situation from equation 1.15, where we have a term with two free indices in left block and two physical indices in right block. In our implementation, we merge the loop over these indices and perform it sequentially in a single MPI process, in order to avoid fetching of small memory chunks. Aside from decreasing internode messaging due to small packets, here we look to a potential GPU acceleration in future.

A single task of $H\left|\Psi\right\rangle$ can be written as

$$\tilde{\mathbf{\Psi}}^{efgh} \quad += \quad {q_1}^{f}_{\ b}\ {q_2}^{g}_{\ c} \times \mathbf{L}_{a}^{\alpha\ e}\ \mathbf{\Psi}^{abcd}\left(\mathbf{R}_{d}^{\beta\ h}\right)^{\dagger}, \tag{1.25}$$

where $efgh$ are indices of output sectors, $abcd$ are indices of input sectors and ${q_1}^{f}_{\ b}$ are just scalars corresponding to operators acting on the two middle orbitals in the given term of the sum in Hamiltonian. $\mathbf{L}_{a}^{\alpha\ e}$ is a matrix corresponding to a operator acting on a left block, with a free index α (possibly merged from two free indices). To be more specific, $\mathbf{L}_{a}^{\alpha\ e}$ can be for instance the pre-summed operator from equation 1.15 with two free indices, or a single-index variant thereof. $\mathbf{R}_{d}^{\beta\ h}$ is the matrix of its counterpart acting on the right block. Here we use the += operator, which is a C++ notation for assignment with addition to the former values of the LHS. In our implementation, this addition is carried out by the means of MPI reduction.

Most of parallelization techniques listed above are useful also for the case of parallelization within a single computational node with multiple CPU cores. When using multiple computational nodes, one has to take into account various other factors, particularly the overhead due to internode communication. In order to perform its task, the node needs to access the matrices of sectors of given operators acting on LqqR blocks (for reading). If the RAM memory of a single node is sufficient, it is profitable to store copies of all operator matrices on each node and update them in each DMRG iteration, so called local memory model. However, this is not feasible for larger systems. For them, the operators have to be distributed among the nodes preferably in a way which minimizes internode queries. In the global memory (GM) model, which we have developed, the distribution is performed over all available nodes with balanced load per node. The global model also implements internode data requests using the Message passing interface (MPI).

In our implementation, any intermediate of the DMRG calculation may be treated in a different data model (local or global), purely based on the amount of free computer memory, just to maximize data locality. Moreover, our code features allows setting a threshold for the maximal size of matrix to be stored locally. Any matrices larger than a threshold are distributed globally. Notice that different values of this threshold correspond to individual rows in Table 1.1 below, whose content is, nonetheless introduced only later in this section.

For the representation of matrices and tensors, we use tensor_class from our in-house tensor library, which allows for external storage of the data. Thus we can work with a matrix as an abstract object until its physical data are needed. These can then be fetched from the global memory on demand, or stored locally if this is affordable. In order to prevent multiple identical fetch requests, the matrix data can be pre-fetched and used repeatedly, and then discarded automatically with the last usage of the virtual matrix class in a given task.

For performing the $H\left|\Psi\right\rangle$ task list, we have developed a semi-dynamic scheduler, which estimates the amount of internode communication required for executing the task for each different node based on the operator storage distribution and then it assigns the task to the node which minimizes this cost function. The execution of the task is then done in parallel within the given node, using the MPI dynamical task scheduling over the local CPU cores.

For the case of blocking and renormalization we consider an example of enlarging a left block operator by one orbital. To be specific, in this example, we choose to block and renormalize the operator pairs from the rows of equation 1.18. However, the form of following relation can be used for other operator pairs. Let us write down a single task for an input sector a and output sector e of left block operator $\mathbf{L}$

$$\tilde{\mathbf{L}}^{\tilde{\alpha}\,(ef)}_{(ab)} \quad \mathrel{+}= \quad \mathbf{O}^{(ef)}_{(ef)} \, \left(\mathbf{L}^{\alpha\;e}_{a} \otimes q_{1\;b}^{\;f}\right)^{(ef)}_{(ab)} \, \mathbf{O}^{(ab)\dagger}_{(ab)} \quad , \tag{1.26}$$

where $\mathbf{O}^{(ef)}_{(ef)}$ is the matrix of renormalization operator (composed of chosen RDM eigenvectors) acting on sector (ef) in the basis of enlarged block $L \otimes q_1$. Here (ef) refers to index of sector in this enlarged basis composed of sectors e and f from the bases of left block L and the orbital q_1 respectively. In the equation 1.26, α and $\tilde{\alpha}$ correspond to the free index (or a merged pair of free indices) of the left block operator from the equation 1.18 and its enlarged version respectively.

For the sake of readability, a simplification was done in equation 1.26 by assuming that the sector (ef) is composed only of sectors e and f. However, in most cases, multiple symmetry sectors contribute to specific parts of sector (ef) and when performing the addition $\mathrel{+}=$, one needs to take this into account and add only to the corresponding subsector. Naturally, the same holds for input sector (ab). Therefore, the addition is only to a submatrix of the matrix $\mathbf{L}^{\tilde{\alpha}\,(ef)}_{(ab)}$.

Notice that renormalization operator $\mathbf{O}^{(ef)}_{(ef)}$ is block diagonal in its sector-decomposed form as it performs only the truncation of the basis. Hence, it maps a sector of any enlarged operator to the same sector, while reducing its matrix dimensions. Except for the trivial case of no truncation, the matrices of operator **O** are rectangular ($n_{\text{cols}} > n_{\text{rows}}$), since they, via matrix multiplication, decrease the dimensions of matrix of the operator, which has just been enlarged by blocking.

When executing the blocking and renormalization tasks, the demands for processing power are usually lower than in case of performing $H\left|\Psi\right\rangle$ during Hamiltonian diagonalization, since matrix multiplication is not used as many times. On the other hand, the amount of required internode communication is much higher, since Kronecker product of pairs of operators is being performed, accessing many different sectors in memory to build the composed operator. Therefore we once again merge the free index loops in larger chunks, similarly as we did for $H\left|\Psi\right\rangle$ in order to reduce the required internode MPI messaging. We also reorder the task list for renormalization in the way which maximizes the reuse of the fetched dense sector matrices in consecutive tasks. However, note that the biggest optimization in blocking and renormalization is already intrinsic to our approach to blocking procedure, by employing the locality of newly formed operators, i.e. when forming the operators, we need only the intermediaries available from the previous iteration. This is explained in section 1.2 in the paragraph below equation 1.18.

Let us now turn to results on benchmark systems. The major indicator of successful parallelization is the dependency of processing time on the number of CPU cores used in the calculation, which would ideally be linear, as when running completely independent tasks. Of course, in DMRG, certain level of internode communication cannot be avoided and hence we expect the curves deviate from linear to some degree, even though our implementation, MOLMPS, attempts to minimize this deviation.

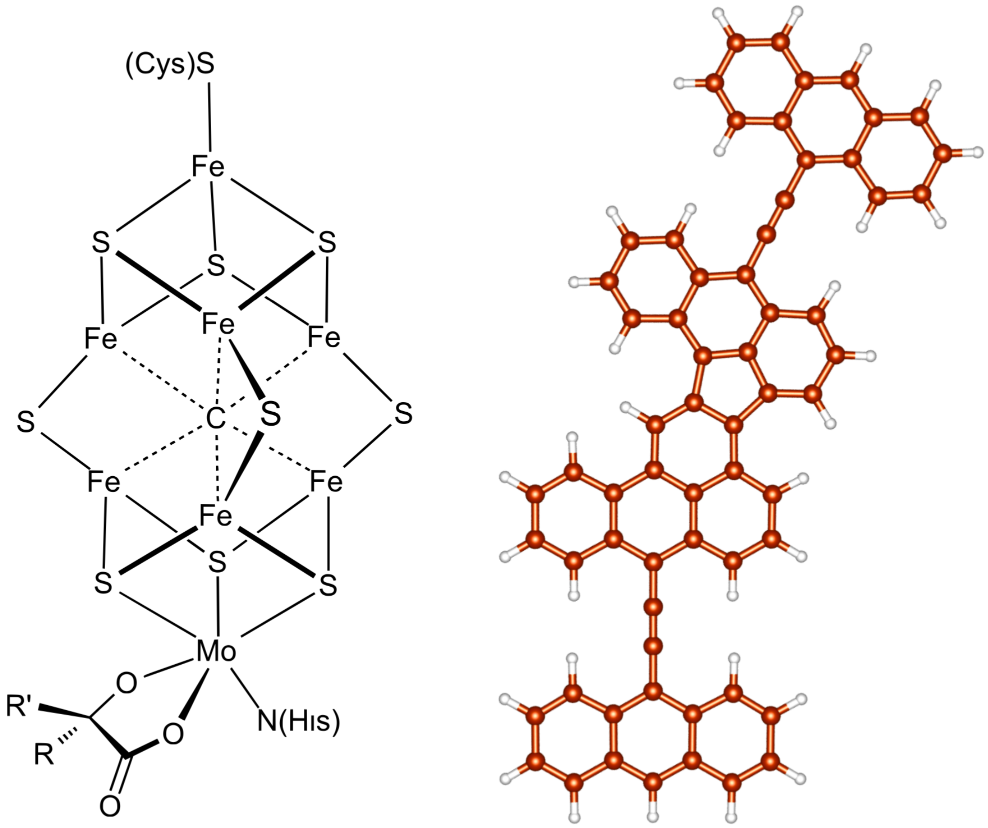

(a) FeMoco cluster scheme.
Figure in public domain.

(b) Defected π-conjugated
anthracene tetramer.

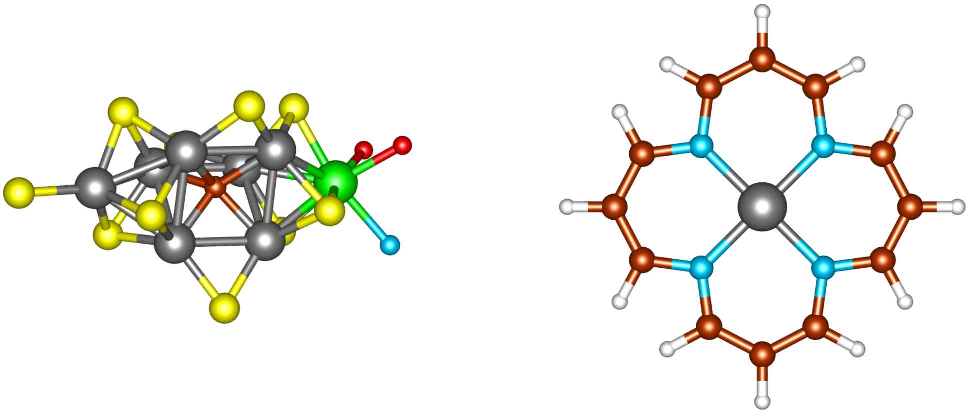

(c) FeMoco cluster.

(d) Iron(II)-Porphyrin model.

Figure 1.5: Structures of the molecules used as benchmarks for the massively parallel MOLMPS implementation. Notice that for the sake of clariry, the FeMoco cluster is pictured only partially, i.e. only the atoms of the ligands, which bond directly to Fe and Mo atoms. Colors of atoms: nitrogen—blue, sulfur—yellow, oxygen—red, carbon—brown, hydrogen—white, iron—gray, molybdenum—green. Images b), c) and d) from our work Brabec et al. [2020].

The Figure 1.5 shows the systems which we chose as benchmarks. These are the typical representatives of molecules calculated with QC-DMRG, characteristic with a necessity of large CAS for a correct description of their ground state.

The first system, Fe(II)-porphyrin, was selected since it was shown by Manni et al. [2019], that a description of complete π space, CAS(32,34), is required for a quantitative determination of its ground state.

The second system chosen is the defected π-conjugated anthracene tetramer, which we selected because our group recently studied its ground state with DMRG and since this and similar molecules are produced as unwanted by-products during on-surface synthesis of ethynylene-bridged anthracene polymers, see Sánchez-Grande et al. [2019].

The largest and somewhat symbolic benchmark is the nitrogenase FeMo cofactor (FeMoco) cluster. This molecule is responsible for nitrogen reduction during the energetically-efficient process of nitrogen fixation under ambient conditions in certain types of bacteria. The corresponding industrial Haber-Bosch process to produce ammonia for fertilizers is, in contrast, very energetically demanding. Thus, numerous efforts have been developed to understand the reaction in bacteria with the potential industrial application in mind. However, as shown in Li et al. [2019], the electronic structure of FeMoco still remains poorly understood.

The dependency of processing time on the number of CPU cores, which we obtained for Fe(II)-porphyrin, is shown in Figures 1.6a and 1.6b. We plot separately the figures for time spent on diagonalization, where $H\left|\Psi\right\rangle$ is performed and the figures for renormalization with blocking. For this case the bond dimension is $M = 2000$, here we plot also the time spent on pre-summation of Hamiltonian sum from equation 1.8. In this case the curve for pre-summation is slightly higher than the curve for renormalization, but for larger systems the remormalization is more costly (as apparent in Table 1.1) and hence in other plots we ommit the pre-summation timings.

We see that nearly optimal scaling is achieved for up to about 512 cores, but then the curve in Figure 1.6a takes an upturn with the growing number of cores. This subfigure corresponds to the case when all data fit into the memory of a single node (local memory model) and the upturn is due to the lack of tasks to be performed.

Figures 1.6c and 1.6d show similar dependency for the same system of Fe(II)-porphyrin, but this time with higher accuracy with bond dimension of $M = 8000$ and the global memory model. Even though not linear anymore, the curve still shows quite favourable scaling. However, once we use 2048 or more CPU cores, the overhead becomes so large that it is not worthy to increase the number of CPU cores anymore. When this occurs, we say that a *saturation* is reached. Notice that the curve in Figure 1.6d for renormalization has even worse scaling, since this step is even more dependent on internode messaging. The saturation is reached at the same level of 2048 cores. However, notice that the absolute renormalization/pre-summation timings are generally order of magnitude lower than those of H diagonalization, and thus the impact of renormalization costs is comparably low.

The Figure 1.7a shows both diagonalization and renormalization costs for the anthracene tetramer, since now we have smaller amount of data points. Notice the increased size of active space, CAS(63,63), and the global memory model.

The computational time is again reduced significantly with the number of core despite overhead due to internode communication. Since the overhead can be significant, this improvement is already a positive result.

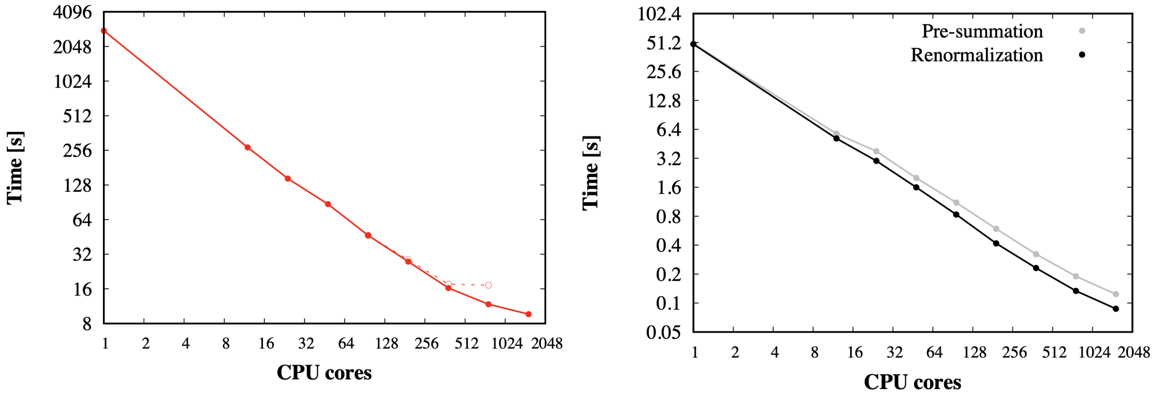

(a) H diagonalization, $M = 2000$.

(b) Renormalization, $M = 2000$.

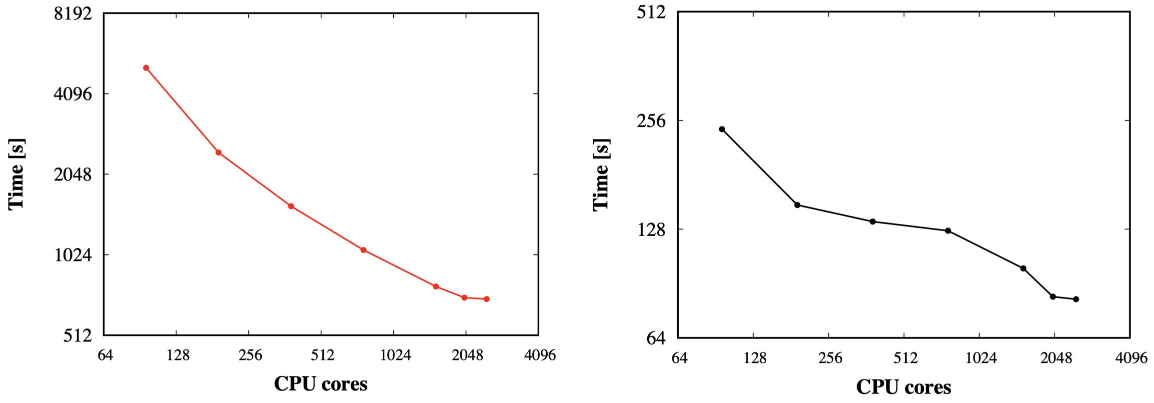

(c) H diagonalization, $M = 8000$.

(d) Renormalization, $M = 8000$.

Figure 1.6: Iron(II)-Porphyrin model with CAS(32,34). The dependency of processing time on the number of CPU cores for Davidson diagonalization of Hamiltonian and for the renormalization with blocking. Charts from our work Brabec et al. [2020].

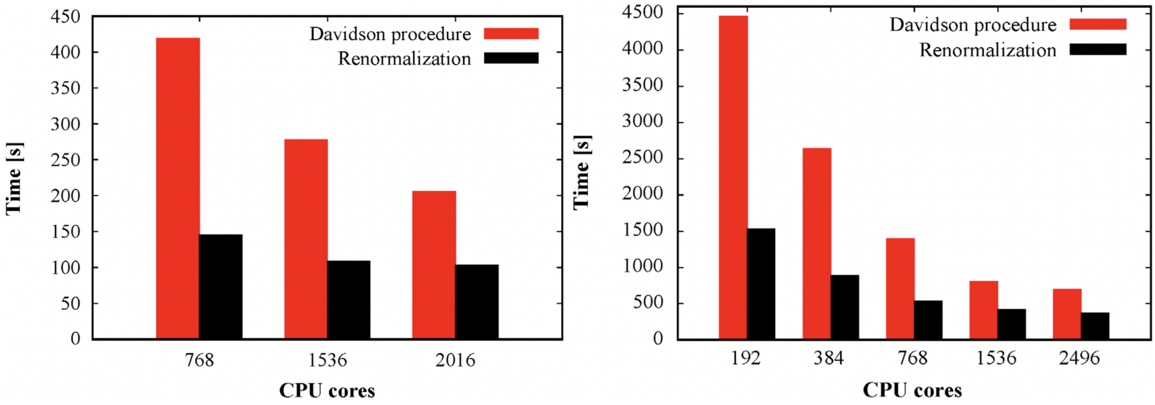

(a) π-conjugated anthracene tetramer with CAS(63,63), $M = 4096$.

(b) FemoCo cluster with CAS(113,76) and bond dimension of $M = 6000$.

Figure 1.7: The dependency of processing time on the number of CPU cores for Davidson diagonalization of Hamiltonian and for the renormalization with blocking for the π-conjugated anthracene tetramer and the FeMoco cluster. Charts from our work Brabec et al. [2020].

The Figure 1.7b shows the timings obtained for the FemoCo cluster with a very large active space, CAS(113,76), the bond dimension of $M = 6000$, and the global memory model. Even for this demanding system, the massive parallelization brings the timing down significantly, but this time the saturation is reached much earlier, near 1536 CPU cores.

To analyze the slowdown due to internode communication we present the timings in Table 1.1 for different values of threshold of maximal matrix size stored in local memory. The rest of matrices is then automatically distributed and stored on different nodes. We can see that for this representative example, the local memory storage is about twice as fast for both the H diagonalization and the renormalization steps. Thus, as we mentioned above, our code copies matrices to local nodes when their RAM memory is sufficient. However, for large calculations like the one with FeMoco cluster, the RAM of a single node is usually insufficient and the matrices have to be distributed.

The slowdown due to internode messaging is the major limiting factor when we look to use more CPU cores. The typical behaviour observed in the benchmarks is that the scaling of time with respect to the number of CPU cores is almost linear while local storage is used. The less favourable scaling for the larger number of CPU cores can in such cases occur due to insufficient number of tasks for such large number of MPI processes.

For the larger systems among our benchmarks, the program has to distribute part of the matrices and more data are interchanged with the increasing number of CPU cores. Then, the scaling starts again as almost linear, but then takes a less favorable tendency and eventually approaches saturation.

The benchmark calculations were performed on the Salomon supercomputer of the Czech national supercomputing center in Ostrava. We have used up to 2480 CPU cores simultaneously for a single DMRG calculation. The hardware specifications of individual node is 24 cores (2 × Intel Xeon E5-2680v3, 2.5 GHz), 128 GB RAM with InfiniBand FDR56 interconnect.

Max. size	% local	pre-summation	diagonalization	renormalization
0.0 MB	0.00	5.26	344.85	38.36
0.8 MB	0.08	5.23	339.49	33.43
1.6 MB	0.20	5.23	305.83	27.90
3.2 MB	0.30	5.86	261.84	22.08
40.0 MB	0.56	5.27	204.84	21.59
all local	100.0	5.90	159.35	20.14

Table 1.1: Fe(II)-Porphyrin model with CAS(32,34), $M = 4096$ on 196 CPU cores. Timings (in seconds) of H diagonalization, renormalization and pre-summation with respect to the distribution of operator matrices to computational nodes. The percentage of matrix elements stored locally within RAM of the node is shown in the second column, while the threshold imposed on the maximal size of a matrix to be saved locally is listed in the first column. Charts from our work Brabec et al. [2020].

For more results, computational details and broader discussion, see our peer-reviewed study "Massively parallel quantum chemical density matrix renormalization group method" by Jiří Brabec, Jan Brandejs, Karol Kowalski, Sotiris Xantheas, Örs Legeza and Libor Veis included in the attachment below, Brabec et al. [2020].

1.4 Analysis of chemical bonding

As we mentioned in section 1.1, the distribution of correlation in the system has strong influence on the performance of DMRG method. Typically, we restrict ourselves to a basis which is localized in a way that it respects the division of the system to groups of mutually correlated orbitals, i.e. such basis where correlations between these clusters is reduced. It has been shown that the more successful the reduction of entanglement between parts of the system, the more the DMRG error decreases, the faster the method converges and the smaller the bond dimension is required, see Barcza et al. [2011].

Now let us look at this another way. Since DMRG is strong in capturing static correlation and from the MPS state, we can efficiently calculate reduced density matrices, which contain the information on correlations between parts of the system, why not use this for an analysing of the system itself?

In the work of Szalay et al. [2017], a correlation theory of the chemical bond has been proposed, which connects the concepts of chemical bond with correlation among orbitals localized on individual atoms. As an example, consider the case of a simple covalent bond. When we localize the bonding and antibonding molecular orbitals into their atomic contributions, these localized orbitals will be highly correlated. Moreover, the strength of the correlation corresponds with the strength of the bond from the quantum information theoretical point of view. We chose to employ this correspondence in a quantitative study of the often counter-intuitive bonding in systems with electron-deficiency.

To work with the concept of correlation quantitatively, we employ the correlation measures introduced in Szalay et al. [2017]. These measures are based on von Neumann entropy. Von Neumann entropy of system partition X of can be defined as

$$S(X) = -\mathrm{Tr}\left[\, \varrho_X \, \log_4 \varrho_X \right], \tag{1.27}$$

where ϱ_X is the reduced density matrix of X. The choice of base 4 matrix logarithm serves for convenient work with numerical values. 4 is given by the dimension of single orbital basis. The resulting values are the same as of the original measures with natural logarithm given in the units of $\ln 4$. For more about entropy in quantum information theory, see Ohya and Watanabe [2010]. In this work, we always partition the active space of a molecular system in parts corresponding to groups of localized atomic orbitals (clusters). When X is composed of n orbitals, then it holds for $S(X)$ that $S(X) \leq n$.

The reduced density matrix of a particular set of orbitals X within system L (e.g. L=CAS) can be efficiently calculated by a partial contraction (partial trace) of a matrix product state, which is, in our case, available from DMRG. The partial trace is used to contract all indices in MPS except for physical indices of X. For instance, when X is composed of two orbitals $X = \{i,\ j\}$, the reduced density matrix of X is calculated as

$$\varrho_{ij} = \sum_{\substack{\alpha_k,\, \alpha'_k \\ \forall k \neq i,\, j}} \mathbf{A}_1^{\dagger\, \alpha'_1} \ldots \mathbf{A}_{|L|}^{\dagger\, \alpha'_{|L|}} \; \mathbf{A}_1^{\alpha_1} \ldots \mathbf{A}_{|L|}^{\alpha_{|L|}}, \tag{1.28}$$

where α_i is a physical index of orbital i and where $|L|$ refers to the number of orbitals in the system L. A graphical scheme of two-orbital reduced density matrix calculation is shown in Figure 1.8. For general subsystem X, the sum is analogous, except for a different set of physical indices that remain uncontracted.

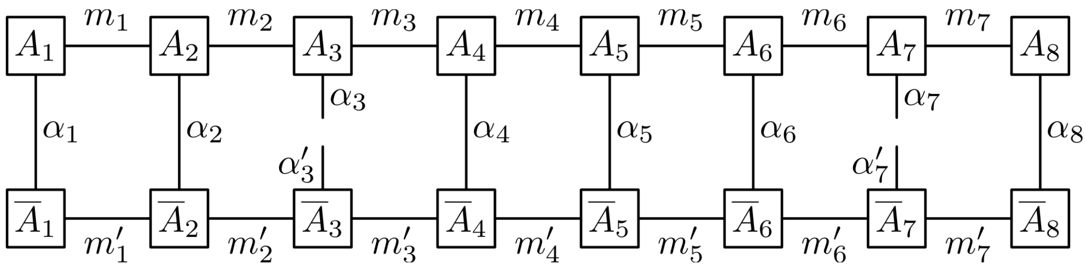

Figure 1.8: Example calculation of reduced density matrix ϱ_X for X composed of two orbitals with physical indices α_3 and α_7, i.e. when $X = \{3,\ 7\}$. Since this scheme represents a partial trace, i.e. contractions of matrix indices "with themselves", a matrix of each orbital is present twice in the scheme, as A_i and $\bar{A}_i = \bar{A}^\dagger$. In the scheme, all the virtual (m_i) and physical (α_i) indices are contracted, except for α_3, α_7 and α'_3, α'_7, which form the indices the reduced density matrix ϱ_X. Therefore, ϱ_X is in our case a 16×16 matrix. When we include multiple orbitals in X, the indices of ϱ_X are then formed by a composition of multiple physical indices α_i and the dimensions of ϱ_X increase considerably. Therefore, our implementation uses a sparse, sector-decomposed representation of reduced density matrix. Image from Szalay et al. [2015].

Consider a partitioning of a group of orbitals, L, into k disjoint clusters of orbitals $L = X_1 \cup X_2 \cup \ldots \cup X_k$. The measure of correlation with respect to this partitioning (correlation among the parts) is the ξ-correlation, defined as

$$C(\{X_1, X_2, \ldots, X_k\}) = \sum_{i=1}^{k} S(X_i) \; - \; S(L). \tag{1.29}$$

Here L can again be the whole active space, or its subset. A special case is when L is composed of just two orbitals, then we denote it as

$$C(i|j) = S(i) + S(j) - S(i,j) = I(i|j) \leq 2, \tag{1.30}$$

and it is termed mutual information between orbitals i and j. This is a well-known concept, which has been considered in chemistry previously by e.g. Boguslawski et al. [2013], Freitag et al. [2015]. The value of $C(\{\{i\},\{j\}\})$ is bounded by 0 and 2. For a general partition, the upper bound is

$$C(\{X_1, X_2, \ldots, X_k\}) \leq 2(|L| - \max_{1 \leq i \leq k} |X_i|). \tag{1.31}$$

Note that C is zero in trivial case $L = X_1$, $k = 1$ and it reaches its maximal value when $k = n$ and each X_i contains one orbital. In the case when cluster L is described by a pure state, then $S(L) = 0$, and the correlation is entirely quantum entanglement, see Szalay [2015].

Since the specific localization of orbitals is inherently basis-dependent, we expect these correlation measures to be basis-dependent as well, however, as we show below, they can still provide meaningful information about the system. Moreover, in certain cases, the discussed correlation measures represent a quantitative measure of entanglement between system partitions. The measures defined above typically account for correlation both due to entanglement and of a different origin. Even though the entanglement itself is a basis-independent concept, many practical entanglement measures are basis dependent, as they are typically based on von Neumann entropy. There is a whole field of study around entanglement, the quantum information theory, which offers numerous different definitions of measures thereof. For more about quantum information, see the textbook of Nielsen and Chuang [2009].

Last but not least, we define a correlation measure which for a given orbital cluster L takes into account all possible partitions to at least m parts. Its value is given by such partition with minimal correlation. In other words, among these partitions, this newly defined measure is equal to correlation of a cut with least correlated parts. We call this measure m-partitionability of L and it is defined as

$$C_{m\text{-part}}(L) \quad = \min\nolimits_{|\xi| \geq m} C(\xi), \tag{1.32}$$

where ξ labels a partition of L into orbital clusters. The minimum is among all possible ξ for given L. From the previous thoughts on maximal and minimal values, we expect that minimal partition should have m parts and it should not the finest split into individual orbitals, which takes the maximal value of correlation possible in L, as we mentioned above. In general, the m-partitionability is bounded by

$$C_{m\text{-part}}(L) \quad \leq 2(m-1). \tag{1.33}$$

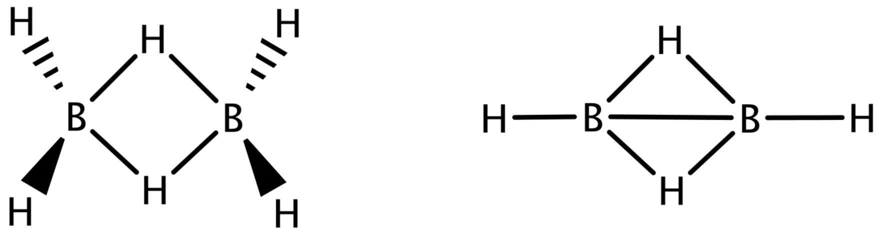

(a) diborane(6), B_2H_6 (b) diborane(4), B_2H_4

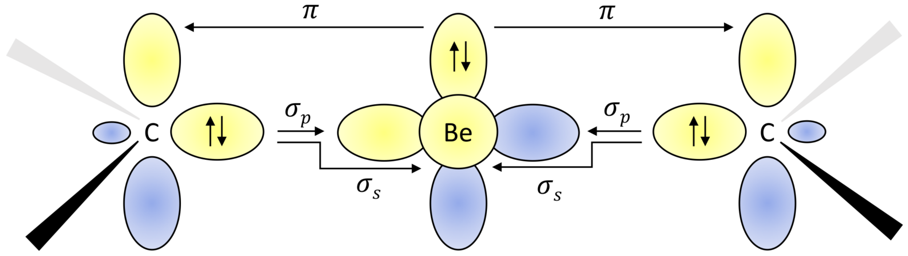

(c) central part of neutral zero-valent beryllium complex

Figure 1.9: Illustration of expected bonding structure of systems under investigation based on available literature. Figures from our work Brandejs et al. [2019].

When using m-partitionability, in most cases, we employ $C_{2\text{-part}}$ and $C_{3\text{-part}}$, as we seek to examine three-center two-electron bonds.

We have chosen three molecular systems as objects of our analysis. The first one, depicted in Figure 1.9a, is diborane(6) (B_2H_6), a well-known textbook example of a molecule with of three-center two-electron bonds. This was to verify that the analysis gives expected results for a well-known system.

Then we turned to two newly synthesized systems featuring electron-deficiency whose more complicated bonding has not yet been studied in detail. The first one is diborane(4) (B_2H_4) shown in Figure 1.9b, characterized recently by Chou et al. [2015] and the second is a neutral compound with zero-valent s-block metal, beryllium in its center. Here CAC corresponds to cyclic amino carbene donors, which stabilize the compound due to their π-acidity. This compound, labelled $Be(CAC)_2$, has been synthesized for the first time by Arrowsmith et al. [2016] and according to their theoretical and spectroscopic results, the molecule adopt a closed-shell singlet configuration with the Be(0) metal centre, and a bonding structure depicted in Figure 1.9c. The surprising stability of the compound has been ascribed to a strong three-center two-electron π-bond in its center.

Considering our results, Figure 1.10a depicts the scheme of correlations within diborane(6). The red circles show how the orbitals group into clusters, i.e. independent bonds. This means that among all possible clusters, we have highlighted those that are strongly correlated within, but weakly with the rest of the system. Notice the two red triangles in the middle. These correspond to the expected three-center two-electron B-H-B bonds.

The correlation plotted in the Figure 1.10a is the mutual information from equation 1.30, a measure of correlation between two-orbitals. From the perspective of orbital pairs, it seems that the bond is three-center. But is the correlation within the triangle truly of inseparable, three-center nature? To answer this question, we employ the measure of 2-partitionability to the corresponding triple of orbitals labelled X_1 in the Figure 1.10a. We obtain $C_{\text{2-part}}(X_1) = 1.5$, which is, as shown in Table 1.2a, 75% of the bounding value. The value is large within this system, but without proper comparison we cannot conclude that it supports the hypothesis of an independent three-center bond. To corroborate further, we look at the correlation (entanglement) of X_1 with the remaining orbitals, $C(X_1 \mid \text{rest})$. This is is very weak, only 8.6% of the theoretical maximum, which indicates that the bond within X_1 is independent. To confirm our hypothesis, we look at the correlation of separate pairs of orbitals from X_1 with the rest of the system. As we can see in Table 1.2a, the correlations of these pair with the rest of the system are mutually similar and with values of 46% and 53%, much stronger than $C(X_1 \mid \text{rest})$. This justifies existence of the three-center bond.

It appears that all these comparisons are inherently contained within a single value, the 2-partitionability $C_{\text{2-part}}(X_1)$. It turns out that we can use this quantity to determine to what extent does a triple of orbitals form a three-center bond. Note that here the correlation between the pairs in X_1 is only around 46%, while the 2-partitionability is already 75%. The reason behind this is that for a multi-center bond, the principle of entanglement monogamy implies that all the three pairs cannot reach their theoretical maxima simultaneously. For a model example, see how entanglement monogamy results in a similar limitation on an entanglement of three qubits in Coffman et al. [2000].

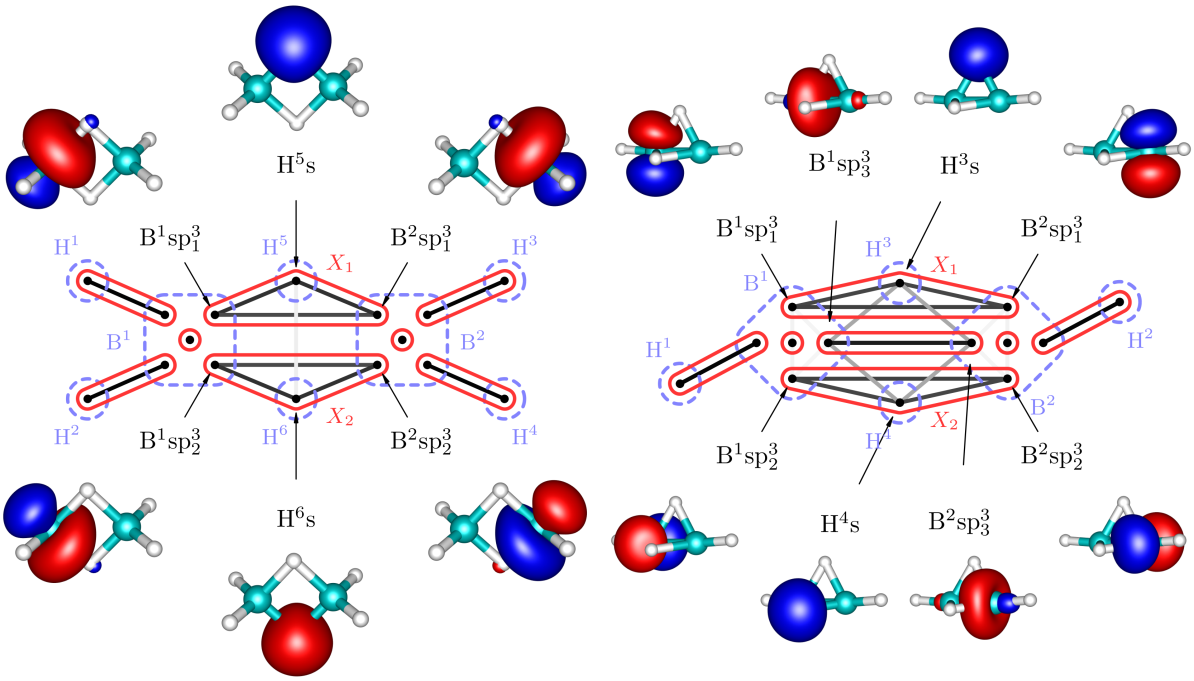

(a) diborane(6) with mutual information (b) diborane(4) with mutual information

Figure 1.10: Schematic view of diborane(6) and diborane(4) with mutual information. Each dot represents a localized orbital, dashed blue line encircles individual atoms, edges correspond to mutual information (plot shaded by a logarithmic scale depending on strength) and red circles show how the orbitals group into clusters, i.e. independent bonds. The scale of mutual information is not plotted, since the logarithmic scaling is difficult to comprehend by colour. This plot serves only for qualitative comparison. For a quantitative description, see the dominant correlation values listed in Table 1.2. Schemes from our work Brandejs et al. [2019].

correlation	abs. value	rel.
$C(X_1 \mid \text{rest})$	0.515	8.6%
$C(X_1 \cup X_2 \mid \text{rest})$	0.412	3.4%
$C(\mathrm{B}^1\mathrm{sp}^3_1, \mathrm{H}^5\mathrm{s} \mid \text{rest})$	1.852	46%
$C(\mathrm{H}^5\mathrm{s}, \mathrm{B}^2\mathrm{sp}^3_1 \mid \text{rest})$	1.852	46%
$C(\mathrm{B}^1\mathrm{sp}^3_2, \mathrm{B}^2\mathrm{sp}^3_1 \mid \text{rest})$	2.114	53%
$C(\mathrm{B}^1\mathrm{sp}^3_1 \mid \mathrm{H}^5\mathrm{s})$	0.894	45%
$C(\mathrm{H}^5\mathrm{s} \mid \mathrm{B}^2\mathrm{sp}^3_1)$	0.894	45%
$C(\mathrm{B}^1\mathrm{sp}^3_1 \mid \mathrm{B}^2\mathrm{sp}^3_1)$	0.605	30%
$C_{2\text{-part}}(X_1)$	1.500	75%
$C_{3\text{-part}}(X_1)$	2.394	60%
$C(X_1 \mid X_2)$	0.309	5.2%
$C(\mathrm{H}^5\mathrm{s} \mid \mathrm{H}^6\mathrm{s})$	0.042	2.1%

(a) diborane(6)

correlation	abs. value	rel.
$C(X_1 \mid \text{rest})$	1.328	22%
$C(\mathrm{B}^1\mathrm{sp}^3_1 \mid \mathrm{H}^3\mathrm{s})$	0.647	32%
$C(\mathrm{B}^1\mathrm{sp}^3_1 \mid \mathrm{B}^2\mathrm{sp}^3_1)$	0.701	35%
$C_{2\text{-part}}(X_1)$	1.388	69%
$C_{3\text{-part}}(X_1)$	2.089	52%
$C(X_3 \mid \text{rest})$	1.438	36%
$C(\mathrm{B}^1\mathrm{sp}^3_3 \mid \mathrm{B}^2\mathrm{sp}^3_3)$	1.245	62%
$C(\mathrm{B}^1\mathrm{sp}^3_3 \mid \mathrm{H}^3\mathrm{s})$	0.130	6.5%
$C(X_1 \cup X_2 \cup X_3 \mid \text{rest})$	0.535	6.7%
$C(X_1 \mid X_2)$	0.066	1.1%
$C(X_1 \mid X_3)$	0.639	16%

(b) diborane(4)

Table 1.2: Correlation measures for diborane(6) and diborane(4). Relative values are related to the upper bounds. Labeling of orbitals corresponds to Figure 1.10. Data from our work Brandejs et al. [2019].

For the case of recently characterized system of diborane(4), the situation of three-center bonds plotted in Figure 1.10b is analogous to diborane(6) and we can reach a similar conclusion about the nature of its three-center bonds, based on values of correlations in Table 1.2b. However, here the value of $C_{2-part}(X_1)$ is a little lower, 69%, as it is weakend by an extra correlation of 6.5% between $\mathrm{B}^1\mathrm{sp}_3^3$ and H^3s. The correlation between these two orbitals together with their symmetric counterparts forms the grey rhombus in the middle of the Figure 1.2b. Had this correlation been stronger, the description with three-center bonding would become implausible. However, this correlation is still relatively weak compared to other correlations within the system and therefore the description proposed by Chou et al. [2015] is reasonable in terms of our analysis.

The last among the analysed systems is the neutral zero-valent beryllium complex $\mathrm{Be(CAC)_2}$, whose correlation scheme is plotted in Figure 1.11a and the relevant correlations are listed in Table 1.3. The cluster X_1 here is expected to correspond to a strong three-center C-Be-C π-bond. The bonding structure of this system as described by Arrowsmith et al. [2016] is depicted in Figure 1.9c. Here the cluster X_2 corresponds to σ-bonding between C-Be and Be-C. To show that there are two independent σ-bonds in X_2, we have rotated the orbitals within X_2 subsystem and the correlation scheme obtained and shown in Figure 1.12 confirms that the σ-bonds are indeed independent.

Now, let us turn to the peculiar cluster labelled X_1. Since the relative value of $C(X_1 \mid \text{rest})$ is quite high with 34% of theoretical maximum, the idea of strong independent π-bond proposed by Arrowsmith et al. [2016] is questionable. We see surprisingly strong values of mutual information between a orbital of nitrogen on the $\mathrm{Be(CAC)_2}$ rings and the carbon from the opposite side of the molecule (with respect to beryllium center). It appears that this is an unexpected direct correlation between the rings, without the participation of Be orbitals. After more careful analysis, we have hypothesized that instead, the two electrons in X_1 are spatially delocalized and reach outside X_1 to the nitrogen orbitals, which then weakly participate in the *pi*-bond over the C-Be-C core.

To support this, we removed two electrons from the system and calculated the mutual information for $[\mathrm{Be(CAC)_2}]^{2+}$, shown in Figure 1.11b. As expected, in this positively charged version of the system, the correlations in have X_1 disappeared, since the electrons necessary to form the *pi*-bond are not present. Moreover, the problematic N-C correlations has disappeared as well. This supports the hypothesized origin of N-C correlation due to the two delocalized π-bond electrons.

We conclude that the structures of studied beryllium complexes suggested by our correlation analysis are those which we depicted schematically in Figure 1.13a and that the nitrogen atoms do contribute to the bonding scheme.

For more detailed discussion of the results, see our peer-reviewed study “Quantum information-based analysis of electron-deficient bonds” by Jan Brandejs, Libor Veis, Szilárd Szalay, Gergely Barcza, Jiří Pittner, Örs Legeza included in the attachment below, Brandejs et al. [2019].

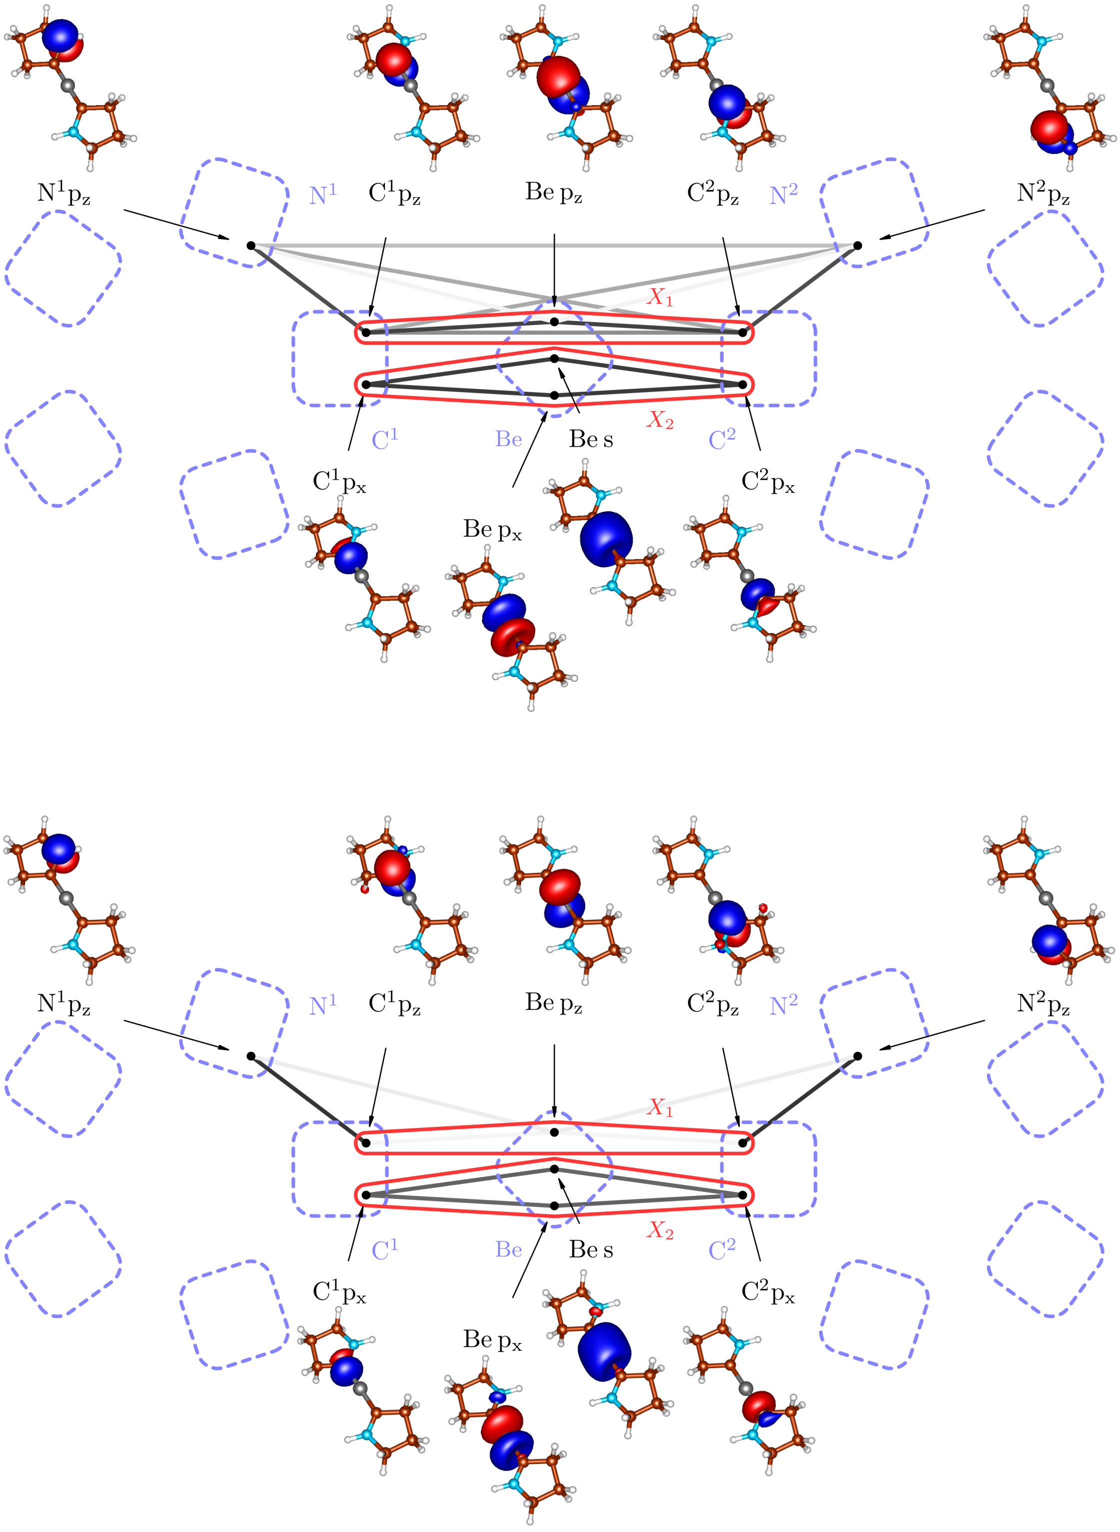

(b) Be(CAC)$_2$ 2+

Figure 1.11: Schematic view of central part of neutral zero-valent beryllium complex Be(CAC)$_2$ and a positively charged version thereof with mutual information. Each dot represents a localized orbital, dashed blue line encircles individual atoms, edges correspond to mutual information (plot shaded by a logarithmic scale depending on strength for a qualitative comparison) and red circles show how the orbitals group into clusters, i.e. independent bonds, see the selected individual values in Table 1.3. Plot from our work Brandejs et al. [2019].

correlation	abs. value	rel.
$C(\mathrm{C}^1\mathrm{p_z}, \mathrm{C}^2\mathrm{p_z} \mid \mathrm{rest})$	3.424	86%
$C(\mathrm{C}^1\mathrm{p_z}, \mathrm{Be\,p_z} \mid \mathrm{rest})$	2.432	61%
$C(X_1 \mid \mathrm{rest})$	2.032	34%
$C(\mathrm{C}^1\mathrm{p_z} \mid \mathrm{C}^2\mathrm{p_z})$	0.194	9.7%
$C(\mathrm{C}^1\mathrm{p_z} \mid \mathrm{Be\,p_z})$	0.681	34%
$C_{2-part}(X_1)$	1.153	58%
$C_{3-part}(X_1)$	1.834	46%
$C(X_1 \mid \mathrm{N}^1\mathrm{p_z})$	0.915	46%

correlation	abs. value	rel.
$C(X_1' \mid \mathrm{rest})$	0.209	2.6%
$C(\mathrm{N}^1\mathrm{p_z}, \mathrm{C}^1\mathrm{p_z} \mid \text{rest of } X_1')$	1.737	43%
$C(\mathrm{N}^1\mathrm{p_z} \mid \mathrm{C}^1\mathrm{p_z})$	0.560	28%
$C(X_2 \mid \mathrm{rest})$	0.209	2.6%
$C(\mathrm{C}^1\mathrm{p_x} \mid \mathrm{Be\,s})$	0.765	38%
$C(\mathrm{C}^1\mathrm{p_x} \mid \mathrm{Be\,p_x})$	0.771	39%
$C_{\text{2-part}}(X_2)$	1.936	97%
$C(\mathrm{C}^1\mathrm{p_x} \mid \mathrm{Be\,sp_2})$	1.912	96%

Table 1.3: Correlation measures for the neutral zero-valent beryllium complex Be(CAC)$_2$. Relative values are related to the upper bounds. Labeling of localized orbitals corresponds to Figure 1.11a. Data from our work Brandejs et al. [2019].

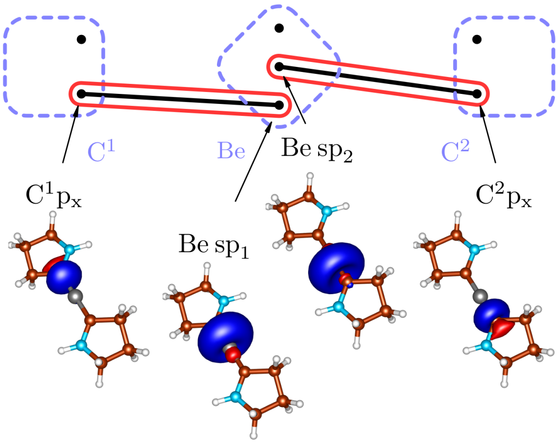

Figure 1.12: Subspace X_2 of neutral zero-valent beryllium complex Be(CAC)$_2$, which is composed of the four orbitals participating on σ-bonding between C-Be and Be-C. Here, the orbitals have been rotated within X_2 in order to show the independence of these bonds. Scheme with mutual information plotted. Scheme from our work Brandejs et al. [2019].

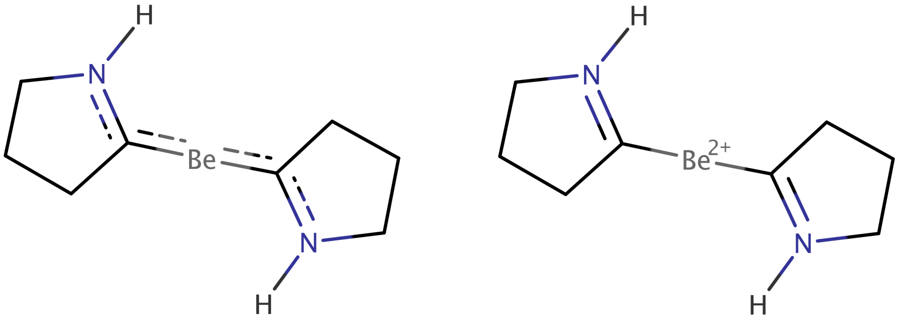

(a) neutral

(b) positively charged 2+

Figure 1.13: Structures of studied beryllium complexes suggested by our correlation analysis. The Images from our work Brandejs et al. [2019].

2. Tailored coupled cluster method - TCC

The coupled cluster approach was first brought to chemistry by Jiří Čížek in 1966, thanks to his excellent multidisciplinary knowledge extending to techniques from nuclear physics Čížek [1966]. The resulting CCSD(T) method is nowadays considered a gold standard in quantum chemistry. Coupled clusters (CC) treatment is particularly suitable for systems dominated by dynamical electron correlation Helgaker et al. [2000].

The central idea behind coupled clusters theory is to use a following ansatz for the parametrization of the wave function Ψ

$$|\Psi\rangle = \exp\left(T\right)|\Phi_0\rangle\,, \tag{2.1}$$

where Φ_0 is the reference determinant (Fermi vacuum). The exponential of the cluster operator T encodes the necessary excitations and their respective amplitudes $t_{ijk\dots}^{abc\cdots}$, which are the parameters of the wave function. T takes the form

$$T = \sum_{ai} t_i^a a_a^\dagger a_i + \frac{1}{4}\sum_{abij} t_{ij}^{ab} a_a^\dagger a_b^\dagger a_j a_i + \frac{1}{36}\sum_{abcijk} t_{ijk}^{abc} a_a^\dagger a_b^\dagger a_c^\dagger a_k a_j a_i + \dots \tag{2.2}$$

Each sum in T corresponds to a certain level of excitation with respect the Fermi vacuum. The indices of amplitudes t denoted by subscript $(i,\, j\, k\, \dots)$ run over occupied orbitals and correspond to excitations from these to virtuals, whose indices are in superscript $(a,\, b\, c\, \dots)$. A simple approximation to T is to truncate excitations above certain level by leaving out the corresponding terms in equation 2.2. As an example of simple, but practical choice is to take into account only the first two sums (CC with single and double excitations), which results in the CCSD method.

The CC ansatz essentially performs an infinite summation of certain classes of many body perturbation theory (MBPT) diagrams Shavitt and Bartlett [2009] using the exponential in a way similar to how the infinite summations of Feynman diagrams are handled in quantum field theory. Peskin [2018]

This approach guarantees rigorous size extensivity even for truncated T, unlike some other quantum chemical approximations such as truncated CI.

2.1 Coupled cluster amplitude equations

The coupled cluster equations are derived in several steps. First we employ the CC ansatz from 2.1 and substitute it into the time-independent Schrodinger equation. Then we can rewrite this as the same equation with Hamiltonian transformed by the exponential operator as

$$\begin{aligned} H_N e^T |\Phi_0\rangle &= E\left(e^T|\Phi_0\rangle\right), \\ e^{-T} H_N e^T |\Phi_0\rangle &= E_{corr}|\Phi_0\rangle\,, \end{aligned} \tag{2.3}$$

where H_N is the normal-ordered form of the electronic Hamiltonian, see Shavitt and Bartlett [2009]. Notice that the energy E_{corr} in 2.3 corresponds to the correlation energy defined by $E = E_{HF} + E_{corr}$, as the difference between the total energy and the Hartree-Fock energy E_{HF}. To solve 2.3 we need to optimize amplitudes $t_{ijk...}^{abc...}$ in the cluster operator T defined by 2.2.

In this work, we limit ourselves to the outline of the rest of the derivation, for more details, we refer the reader to the textbook of Shavitt and Bartlett [2009].

We would like to obtain more explicit form of 2.3. By applying the Baker-Campbell-Hausdorff formula to the exponential in 2.3, we expand it to a number of commutators. As described in Shavitt and Bartlett [2009], this expansion actually has a finite number of terms, since the electronic Hamiltonian is a two-body operator in normal form with respect to the vacuum $|\Phi_0\rangle$. Hence the series terminates with four-fold nested commutators.

We can then use the generalized Wick theorem to show that the only nonzero terms will be those in the form of $H_N T^k$, $0 \leq k \leq 4$, i.e. with the Hamiltonian on the left and with up to the fourth power of cluster operator on the right. This allows us to replace the commutators by restricting ourselves to connected MBPT diagrams.

We can do this since the cluster operator T contains only excitations with respect to Fermi vacuum $|\Phi_0\rangle$, and hence, only nonzero terms $H_N T^k$ are those where each cluster operator is contracted with H_N in at least one index, i.e. the diagram corresponding to this term is connected.

The operator on the left hand side of 2.3 can now be rewritten as

$$e^{-T} H_N e^T \quad = \quad H_N e^T \Big|_{\text{connected}} \quad , \tag{2.4}$$

where the whole right hand side is restricted to connected diagrams.

Now if we calculate the expectation value of operator 2.4 in the state $|\Phi_0\rangle$, all the excitations from T higher than doubles will not contribute to energy directly, thus we can write

$$\langle\Phi_0| H_N e^{T_S+T_D} |\Phi_0\rangle \Big|_{\text{connected}} = E_{corr} \langle\Phi_0|\Phi_0\rangle = E_{corr} \quad , \tag{2.5}$$

using equation 2.3. Here T_S, T_D correspond to the first, respectively the second sum in 2.2. Higher excitations in T do not contribute since in order to yield non-zero contribution, the Hamiltonian with one and two electron terms can only compensate for up to double excitations from the Fermi vacuum when acting on $\langle\Phi_0|$ from the right. However, they do contribute to the energy indirectly through the amplitude equations.

Let us now limit T to singles and doubles. Similarly as with the expectation value above, the projection onto singly and doubly excited vacuum, i.e. $\langle\Phi_i^a|$ and $\left\langle\Phi_{ij}^{ab}\right|$ can be used to extract the amplitudes t_i^a and t_{ij}^{ab} when we realize that its value has to be zero

$$\begin{aligned} \langle\Phi_i^a| H_N e^{T_S+T_D} |\Phi_0\rangle \Big|_{\text{connected}} &= 0, \\ \left\langle\Phi_{ij}^{ab}\middle|H_N e^{T_S+T_D}\middle|\Phi_0\right\rangle \Big|_{\text{connected}} &= 0. \end{aligned} \tag{2.6}$$

The set of equations 2.6 is called the CCSD amplitude equations, and it is used to find the coupled cluster amplitudes for T approximated by singles and doubles.

Higher order methods (CCSDT, CCSDTQ) are generated by including further excitations in T, and then projecting onto all excitations of Fermi vacuum up to the given order.

Notice that in practical calculations, perturbative corrections are common, especially CCSD(T), where the estimate of triple excitations is calculated from the converged T_D amplitudes from CCSD using MBPT. This is accompanied by one order of magnutide computational cost on top of CCSD with respect to the number of orbitals, compared with two orders in case of full CCSDT, i.e. $\mathcal{O}(n^7)$ and $\mathcal{O}(n^8)$.

Since the CCSD equations 2.6 are not linear in t_i^a and t_{ij}^{ab}, iterative methods are usually employed to solve them. By using some initial guess for the amplitudes, enumerating the left hand side of 2.6 will yield some generally nonzero values σ_i^a and σ_{ij}^{ab}. These are called residuals.

As illustrated in Shavitt and Bartlett [2009], we can use these residuals to improve our initial guess for the amplitudes, in a way inspired by MBPT

$$t_{ij\,new}^{ab} = t_{ij\,old}^{ab} - \frac{\sigma_{ij\,old}^{ab}}{\epsilon_a + \epsilon_b - \epsilon_i - \epsilon_j}, \tag{2.7}$$

where ϵ_k are HF energies of molecular orbitals. These improved amplitudes can then again be substituted to the left hand side of 2.6 in order to find new residuals. Hence the loop is closed and we may continue to iterate until convergence. In practice, the rate of convergence is usually optimized using the direct inversion in the iterative subspace (DIIS) technique, see Pulay [1980].

2.2 Wave function analysis and the TCC ansatz

The concept behind externally corrected coupled cluster method as described in Li and Paldus [1997], is to retrieve information on static correlation from external source, e.g. CASCI, and to use it in the subsequent CC calculation. A simple approach is the tailored coupled cluster method (TCC) originally proposed by Bartlett et al., see Kinoshita et al. [2005], Lyakh et al. [2011a], Melnichuk and Bartlett [2012, 2014], which uses the split-amplitude ansatz for the wave function introduced in Piecuch et al. [1993], Piecuch and Adamowicz [1994].

We have followed the pathway to correct CC with DMRG. In general, the DMRG extensions efforts aiming to add dynamical correlation on top of DMRG are in relatively early stage. Past works focused on this include for example second order perturbation theory in Kurashige and Yanai [2011], canonical transformation method in Yanai and Chan [2006], random phase approximation in Wouters et al. [2013], and the internally contracted MRCI in Saitow et al. [2013].

The split-amplitude ansatz, on which a general tailored coupled cluster wave function takes the form

$$|\Psi\rangle = e^{T_{\text{ext}}+T_{\text{CAS}}}|\Phi_0\rangle = e^{T_{\text{ext}}}e^{T_{\text{CAS}}}|\Phi_0\rangle \tag{2.8}$$

where T_{CAS} contains just the amplitudes with all indices inside the active space (CAS) of the method used to source the static correlation, i.e. CASCI (DMRG in our case). The amplitudes of T_{CAS} are fully determined by its output and they are fixed for the rest of the calculation.

By contrast, the external part of the cluster operator, T_{ext}, is composed of amplitudes with at least one index outside the CAS space (DMRG CAS). Only these amplitudes are optimized in the subsequent CC calculation.

Please note that the last equality in 2.8 holds since we apply the exponential operator only to the single Slater determinant of $|\Phi_0\rangle$ and therefore the cluster operators T_{CAS}, T_{ext} commute. As noted in Veis et al. [2016], this helps to keep the method simple.

This way of combining the two methods seeks to take the information on static correlation from a method particularly strong in treating static correlation (DMRG) and retain it, while using coupled cluster theory to describe the dynamical correlation. In other words, T_{CAS} is used to "tailor" the external space amplitudes T_{ext}.

The implementation of TCC by our group is TCCSD, where only the amplitudes corresponding to single (S) and double (D) excitations are extracted from DMRG and then used in CCSD. Hence we can write the parametrization of wave function as

$$|\Psi\rangle = e^{T^S_{\text{CAS}}+T^D_{\text{CAS}}}\, e^{T^S_{\text{ext}}+T^D_{\text{ext}}}\, |\Phi_0\rangle\,. \tag{2.9}$$

The CAS amplitudes of T^S_{CAS} and T^D_{CAS} are obtained from the coefficients of CI expansion from the initial CASCI (DMRG) calculation as

$$\begin{aligned} T^S_{\text{CAS}} &= C_1, \\ T^D_{\text{CAS}} &= C_2 - \frac{1}{2}(C_1)^2. \end{aligned} \tag{2.10}$$

Before performing the wavefuction analysis defined in 2.10, the CI coeffients are extracted from a converged DMRG wave function here shown in MPS form 1.7, by performing a contraction over the corresponding indices of the matrix product chain. Considering our DMRG implementation, we have extended the MOLMPS program to for calculation of active space T^S_{CAS} and T^D_{CAS} amplitudes.

Once we have the T_{CAS} amplitudes, we continue in a way similar to single reference CC equations explained in section 2.1. Projecting 2.9 onto excited determinants $\left\langle \Phi^{ab}_{ij} \right|$ lead to an iterative method. The equation 2.6 now takes form

$$\begin{aligned} \langle \Phi^a_i |\, H_N e^{T_{\text{ext}}} e^{T_{\text{CAS}}}\, |\Phi_0\rangle \Big|_{\text{connected}} &= 0, \\ \left\langle \Phi^{ab}_{ij} \middle| H e^{T_{\text{ext}}} e^{T_{\text{CAS}}} \middle| \Phi_0 \right\rangle \Big|_{\text{connected}} &= 0, \end{aligned} \tag{2.11}$$

where $\{i, j, a, b\} \not\subset$ CAS. The iterative procedure optimizes T_{ext} while keeping frozen T_{CAS} obtained from 2.10.

As noted in Veis et al. [2016], the TCCSD method has been used for a number of nonrelativistic quantum chemical systems and performs well, even though in some cases it might require either larger active spaces or the use of CASSCF orbitals for high accuracy, see Kinoshita et al. [2005] for details on its accuracy.

One limitation of this method stems from the fact that it uses just a single Slater determinant as a Fermi vacuum, which introduces a certain bias. This may damage the performance of the method for exactly degenerate states, when

TCCSD breaks the spatial symmetry of the degenerate components. Hence, while the method behaves as size-extensive, a rigorous size-consistency is not achieved. In the work of Veis et al. [2016], a limited analysis of size consistency is performed on the N_2 dimer. In theory, the error vanishes with the growing size of the active space, as TCCSD approaches the FCI limit.

2.3 Four component relativistic TCCSD

All the previous DMRG-extension efforts (post-DMRG methods) have so far been restricted to non-relativistic systems. Hence, the relativistc domain has remained unexplored until recent study Brandejs et al. [2020] produced in close cooperation between our group and the group of Prof. Örs Legeza from Wigner Research Centre for Physics, Budapest.

When one employs the relativistic theory in quantum chemistry, the concept of spin is replaced with the Kramers projection. Let us now turn to relativistic four-component CC (4c-CC). The first implementation of the method was in Dirac program in the Kramers-restricted closed shell form, see Visscher et al. [1995], Shee et al. [2016], and this was later extended to Kramers-unrestricted form in Visscher et al. [1996]. Another direction of development has been the string-based general-order 4c-CC by Nataraj et al. [2010]. Moreover, a Fock-space version of 4c-MRCC has been developed by Visscher et al. [2001]. Recently, a relativistic implementatin of the state-specific Mukherjee's MRCC method has been published in Ghosh et al. [2016].

Considering the relativistic implementatins of DMRG, first attemp was at the scalar DKH level and has been developed by Moritz et al. [2005]. Perturbative approach to spin-orbit couplings on top of DMRG has also been presented by Sayfutyarova and Chan [2016]. In Lan et al. [2014], the authors used DKH3 Hamiltonian in DMRG to evaluate coupling constants. The implementation used in our work is the DMRG recently extended to the 4c relativistic version by Knecht, Legeza, and Reiher in Knecht et al. [2014]. This has opened the way for relativistic post-DMRG approaches, like 4c-TCCSD discussed here.

In relativistic formulations of 4c post-Hartree-Fock methods, the Dirac-Coulomb Hamiltonian is usually expressed in its second quantized form employing the no-pair approximation, which projects out the negative energy solutions as defined in Almoukhalalati et al. [2016]. The QED effects are also neglected, resulting in

$$H = \sum_{PQ} h_P^Q \, a_P^\dagger a_Q + \frac{1}{4} \sum_{PQRS} \langle PQ \,||RS\, \rangle \; a_P^\dagger a_Q^\dagger a_S a_R \tag{2.12}$$

where the indices P, Q, R, S run over the positive-energy 4-component spinors, which span the one-electron basis. $\langle PQ \,||RS\, \rangle \in \mathbb{C}$ are the two-integral coefficients and h_P^Q the coefficients corresponding to one-electron terms in the Hamiltonian. These spinors form Kramers pairs related to each other by action of the time-reversal operator K

$$\begin{aligned} K\phi_p &= \phi_{\bar{p}} \\ K\phi_{\bar{p}} &= -\phi_p \end{aligned} \tag{2.13}$$

As mentioned above, this Kramers symmetry replaces the spin symmetry in the nonrelativistic theory. Now the spin projection M_S is not a good quantum number anymore and is replaced by Kramers projection M_K. Its values are 1/2 for unbarred spinors and $-1/2$ for spinors with barred indices. Indices P, Q, R, S in 2.12 run over both spinors of a Kramers pair and, in contrast with the nonrelativistic theory, the integrals corresponding to excitation operators which change the M_K projection do not necessarily vanish. Therefore, the Hamiltonian (2.12) is not necessarily block-diagonal with respect to M_K. Each creation or annihilation operator in (2.12) changes M_K by $\pm 1/2$, thus states with $|\Delta M_K| \leq 2$ are coupled in this Hamiltonian. Moreover, compared with nonrelativistic case, the complex-valued integrals lead to a decrease in index permutation symmetry of two-electron terms. In other aspects, the second quantized form 2.12 is analogous to its nonrelativistic counterpart. The usual CC diagrams apply, similar to 2.6 and 2.11 above.

Quoting from Brandejs et al. [2020] on the implementation of "quaternion groups":

> *"The Dirac program [see Dirac18] employs a quaternion symmetry approach which combines the Kramers and binary double group symmetry (D^*_{2h} and subgroups) Saue and Jensen [1999]. The double groups can be sorted into three classes based on the application of the Frobenius-Schur indicator to their irreducible representations: "real groups" (D^*_{2h}, D^*_2, and C^*_{2v}); "complex groups" (C^*_{2h}, C^*_2, and C^*_s); and "quaternion groups" (C^*_i and C^*_1). Generalization of nonrelativistic methods is simplest in the "real groups" case, where the integrals are real-valued and the ones with odd number of barred (B) indices vanish. In practice, it means that additional "spin cases" of integrals* (AB|AB) *and* (AB|BA) *(in Mulliken notation) have to be included. For the complex groups, the integrals are complex-valued, but still only integrals with even number of barred indices are non-zero. Finally, in the remaining case of "quaternion groups" all the integrals have to be included and are complex-valued, see Thyssen [2001], Dyall and Faegri [2007]."*

In analogy with nonrelativistic chemistry, the single-determinantal Fermi vacuum guarantees that cluster operators T_{ext} and T_{CAS} commute, leading to simplifications, most importantly the factorization of exponential operator in the 2.8, which enables us to construct the ansatz and the wave function parametrization in analogous manner.

In order to implement relativistic TCCSD, we have modified the existing implementations of DMRG and CC and created an interface which extracts the CAS amplitudes from DMRG and passes them into CC in a standardized file format.

The CC implementation in Dirac program Dirac18 has been modified to load T_{CAS} amplitudes for the corresponding subspace and then to keep them frozen during the iterative optimization of T_{ext}.

The existing nonrelativistic DMRG Legeza et al. [2018] has been modified to include Dirac-Coulomb Hamiltonian with terms with nonzero ΔM_K and to generate explicit MPS for the CC interface, see Battaglia et al. [2018]. DMRG input interface has been modified to read a set of complex integrals generated from 4c-SCF calculation in Dirac program prior to the DMRG calculation. Except for the “real groups”, the DMRG procedure has to work with complex matrices in MPS. The generated T_{CAS} amplitudes are also complex-valued in general.

For the calculations in our first publication Brandejs et al. [2020] described in section 2.4, we have selected numerical examples with “real groups” symmetry (diatomics), while the complex generalization of the DMRG code was still in progress. Eventually, the complex generalization has been finished, but only very recently and the preliminary results are currently being analysed and prepared for a journal submission.

2.4 Application to heavy diatomics

This first publication with relativistic TCC is limited to diatomics, namely thallium hydride, arsenic hydride and astimony hydride, consistent with the way how our implementation was developed where the integrals remained real in the first production version and were only generalized to complex version in the summer of 2021.

method	E_{el} [E$_h$]	ΔE_{el} [mE$_h$]
4c-CCSDTQ(14,47)* ref., higher order	-20275.84024	0.00
4c-DMRG(14,47)** ref., large CAS	-20275.83767	2.57
4c-MP2(14,47)	-20275.85372	-13.49
4c-CCSD(14,47)	-20275.82966	10.58
4c-TCCSD, CAS(14,14) $\cup$ {ext} $=$ (14,47)	-20275.83430	5.94
4c-CCSD(T)(14,47)	-20275.84056	-0.32

*From Knecht et al. [2014].

**4c-DMRG(14,47)[4500,1024,2048,10^{-5}], from Knecht et al. [2014].

Table 2.1: Excerpt from table II in Brandejs et al. [2020]. Total electronic energy and energy differences ΔE_{el} (in mE$_h$) for various methods with respect to the 4c-CCSDTQ(14,47) reference energy of -20275.84024233 E$_h$ for TlH at the experimental equilibrium internuclear distance 1.872 Å.

In this text, we will only include most important charts from our publication Brandejs et al. [2020] with a short commentary. First of all, we would like to present Table 2.1, which shows that the 4c-TCCSD was able to capture more correlation energy of ^{205}TlH than 4c-CCSD. However, the result of 4c-CCSD(T) was closer to the reference energy, even though it undershoots the reference energy. Considering the spectroscopic constants listed in Table 2.2, 4c-TCCSD with CAS(14,10) performed better than 4c-CCSD in comparison with experimental values. Of course, it could not outperform 4c-DMRG on the whole space of (14,47) without extending external space beyond this DMRG space. Notice also that this 4c-DMRG (14,47) calculation is much more demanding in terms of computational requirements.

As we mentioned, an optimal choice of CAS is necessary for the right operation of 4c-TCCSD. In Figure 2.1 (b), we can see the procedure which we used for the optimization of CAS size in equilibrium internuclear distance. For larger cal-

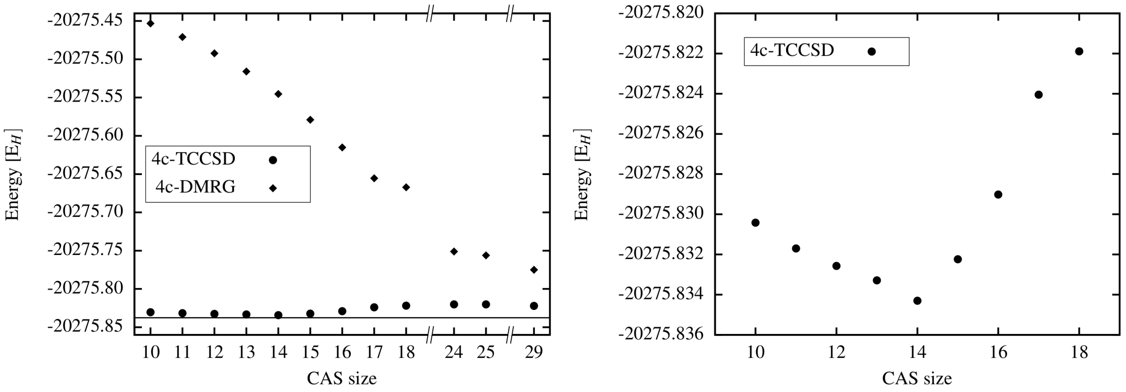

(a) Comparison of TCCSD and DMRG methods. The horizontal solid line represents the "FCI-limit" from the large 4c-DMRG(14,47)Knecht et al. [2014] calculation.

(b) Dependency of 4c-TCCSD energy on CAS size k, with restricted size of external space CAS(14,k) $\cup$ {ext} = (14,47).

Figure 2.1: Equilibrium energy of TlH calculated using the 4c-TCCSD and 4c-DMRG methods with different sizes of DMRG active space, as given in Table 2.1. Charts from our work Brandejs et al. [2020], reproduced with the permission of AIP Publishing.

method	r_e [Å]	ω_e [cm^{-1}]	$\omega_e x_e$ [cm^{-1}]
experiment†	1.872	1391	22.7
4c-DMRG(14,47)*	1.873	1411	26.6
4c-CCSD(14,47)*	1.871	1405	19.4
4c-TCCSD$^{a)}$ CAS(14,10) $\cup$ {ext} = (14,47)	1.874	1403	23.4
4c-TCCSD$^{a)}$ CAS(14,14) $\cup$ {ext} = (14,47)	1.869	1411	22.6

† GRECP spin–orbit MRD-CI from Titov et al. [2001].

*From Knecht et al. [2014].

a) Using TWOFIT module for curve fit in DIRAC, 4th order polynomial.

Table 2.2: Excerpt from table III in Brandejs et al. [2020]. Spectroscopic constants of ^{205}TlH obtained from 4c-TCCSD, compared with calculations and experimental work from the literature. The spectroscopic constants have been evaluated from potential energy curve fit using TWOFIT methodology. The number of points have been selected according to Mean displacement in harmonic ground state criterion. Internuclear separation axis sampling was chosen to be 0.02 Å. Here r_e is the internuclear distance and ω_e, $\omega_e x_e$ is the first, respectively the second term of vibrational constant.

culations, performing such preliminary estimates is essential. Fortunately, these estimates can be based on cheap low-bond dimension DMRG calculations only a neglidible increase in computational requirements compared with the production calculations.

Figure 2.1 (a) shows how the 4c-DMRG energies gradually approach the 4c-TCCSD energies with increasing CAS size. The FCI-limit should be located near the solid horizontal line, which represents the 4c-DMRG (14,47) energy from Knecht et al. [2014].

See our peer-reviewed publication “Toward DMRG-tailored coupled cluster method in the 4c-relativistic domain” by Jan Brandejs, Jakub Višňák, Libor Veis, Mihály Maté, Örs Legeza and Jiří Pittner included in the attachment Brandejs et al. [2020], where potential energy surfaces and spectroscopic constants for further systems are presented, discussed and compared with results from the literature.

2.5 Domain-based local pair natural orbital approach

Let us now turn back to non-relativistic TCC. Here the progress in our group has already gotten much further, as our group produced its first non-relativistic implementation of TCCSD based on Budapest DMRG already in 2016. Later in 2020, the CC interface was also included to MOLMPS - the highly parallel implementation, see Brabec et al. [2020]. The CC part of the method has been implemented into Orca program (see Neese et al. [2020]) in 2016 by members of our group.

However, the problem of CC calculations in general is the computational scaling, which allows us to treat only orbital spaces of limited size (basis size). There are approximations which efficiently reduce size of orbital space, like the frozen core approximation which restricts the correlation treatment to valence orbitals. Unfortunately, limiting the size of virtual space is more difficult, particularly due to to the nature of virtual orbitals, which makes it hard to truncate the space in a way to capture the most of dynamical correlation. Attempts to optimize the virtual space basis turned out to be difficult as well.

We chose to pursue the approach which employs the localizability of dynamical electron correlation, in order to limit the size of virtual space. Efforts related to localizability of dynamical correlation date back to the work of Pulay [1983]. The development continued by the use of projected atomic orbitals (PAO) Saebo/ and Pulay [1987], where the orbital domain-based approach allowed to devise an efficient truncation scheme. Using these concepts, a local CC method was formulated Hampel and Werner [1996], sparking new developments in the area. One of the most successful approaches which we discuss in the next paragraph is based on the concept of pair natural orbitals (PNO) introduced originally by Edmiston and Krauss [1965]. Natural orbitals are known as a basis notable for compact representation of correlated wave functions. In PNO, we seek to combine this with locality to make the best of both worlds.

For some time, the development in this area plummeted, since the computational demands of the newly devised methods were still too large for the options at the time. It wasn’t until modern integral transformation techniques and faster hardware allowed for more relevant applications, see Vahtras et al. [1993]. Later on, PNOs saw their revival in the local pair natural orbital method (LPNO) of Neese et al. [2009]. In their domain-based local PNO approach, they achieved nearly linear scaling for CCSD(T) Pinski et al. [2015] by exploiting sparse matrix representation and detailed prescreening procedures. This allowed for treatment of large systems with hundreds of atoms.

The domain-based local pair natural orbital approach seeks to allow for TCC calculations with large scale virtual space. This approach is based on careful selection of pair natural orbitals comprising the virtual space in CCSD calculation. The CC calculation is restricted to these orbital-pair subspaces, which are obtained by a truncation procedure at the level of pair density matrices (see below).

In all LPNO approaches, the first step is to localize the occupied orbitals. Then the virtual space is transformed to the PNO basis. This process is composed of several steps described below.

We employ the fact that the correlation energy can be written as a sum of orbital pair contributions with coefficients c_{ij}^{ab}, stemming from the equation 2.5

$$E^{corr} = \frac{1}{2}\sum_{abij} c_{ij}^{ab} \left\langle \Phi_{ij}^{ab} \middle| H \middle| \Phi_0 \right\rangle \;=\; \frac{1}{2}\sum_{ij} \varepsilon_{ij} = \frac{1}{2}\sum_{abij} c_{ij}^{ab} \left\langle ab || ij \right\rangle . \tag{2.14}$$

Here, ε_{ij} is the correlation energy corresponding to a pair of occupied indices i, j. This sum can be obtained by a projection of the Schrödinger equation with the exact wave function onto Fermi vacuum $|\Phi_0\rangle$.

The accuracy of the LPNO approach is determined by two cut-off parameters. Firstly, the pair energies ε_{ij} are estimated during pre-screening. Some of the pairs are excluded, some treated at the MP2 level

$$\varepsilon_{ij} = \sum_{ab} \frac{\langle ij|ab\rangle \left[4\langle ij|ab\rangle - 2\langle ij|ba\rangle\right]}{f_{ii} + f_{jj} - \varepsilon_a - \varepsilon_b}, \tag{2.15}$$

and only a fraction of pairs is treated at the CCSD level, based on the threshold T_{cutPairs}. The other cut-off parameter, T_{cutPNO} is defined further in this section. In equation 2.15, f_{ii} represent the Fock matrix elements and ε_a are the orbitals energies. Notice that since $\langle ij|ab\rangle$ reach significant values only locally, only local terms will contribute significantly.

Now the virtual space corresponding to orbital pair ij is to be localized by a transformation to PNO basis. Since we work in virtual space, this is done for inactive pairs ij, i.e. those outside the DMRG CAS, $\{ij\} \not\subset \mathrm{CAS}$. First we construct MP2 amplitude matrix $\mathbf{T}^{ij}$ with elements

$$\left(\mathbf{T}^{ij}\right)_{ab} = \frac{\langle ij|ab\rangle}{f_{ii} + f_{jj} - \varepsilon_a - \varepsilon_b}, \tag{2.16}$$

then we build the pair density matrix as

$$\begin{aligned} \mathbf{D}^{ij} &= \left(\mathbf{T}^{ij}\right)^{\dagger} \tilde{\mathbf{T}}^{ij} + \mathbf{T}^{ij} \left(\tilde{\mathbf{T}}^{ij}\right)^{\dagger}, \\ \text{where} \quad \tilde{\mathbf{T}}^{ij} &= \frac{4\mathbf{T}^{ij} - 2\mathbf{T}^{ij\dagger}}{1 + \delta_{ij}}. \end{aligned} \tag{2.17}$$

The pair density matrices $\mathbf{D}^{ij}$ are then diagonalized. Their eigenvalues correspond to occupations of PNOs and their eigenvectors correspond to the coefficients of PNOs in MO basis. The efficient truncation of the virtual space ij is then done by discarding all PNOs with occupations below certain threshold, T_{cutPNO} for a given orbital pair.

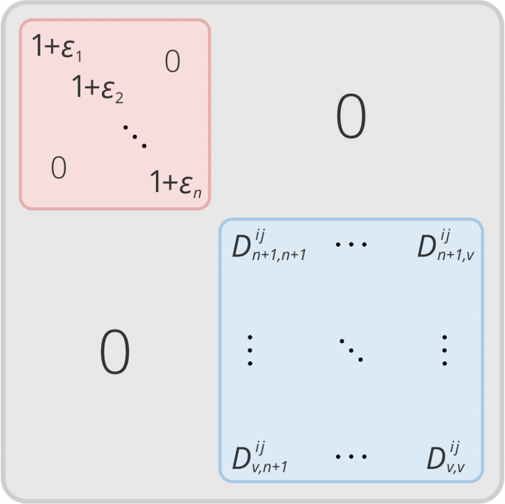

Figure 2.2: Pair density matrix (transformation matrix) of composed active and virtual space for orbital pair ij. The submatrix highlighted in red corresponds to the active-active (CAS-CAS) elements, while the blue part corresponds to virtual-virtual elements defined in equation 2.17. The mixing CAS-virtual terms are set to zero. Scheme from Antalík et al. [2019].

To preserve the alignment with original active orbitals in MO basis, we need to keep their coefficients untouched during PNO transformation. Therefore when we form the pair density matrix for the whole system, we keep nonzero only the diagonal elements of its CAS-CAS part. The CAS-virtual elements are set to zero, and the virtual-virtual part is composed from $\mathbf{D}^{ij}$ matrices, as shown in Figure 2.2. Please note that we set CAS-CAS part of the matrix to identity with very small $\varepsilon_i > \varepsilon_{i+1} > 0$ added to the diagonal to preserve the original order of CAS orbitals in MO basis.

The CC equation 2.11 with amplitudes corresponding to double excitations now takes form

$$\left\langle \Phi^{\bar{a}\bar{b}}_{ij} \middle| H\, e^{\mathbf{T}^S_{\text{ext}}+\bar{\mathbf{T}}^D_{\text{ext}}}\, e^{\mathbf{T}_{\text{CAS}}} \middle| \Phi_0 \right\rangle \Bigg|_{\text{connected}} = 0 \;\; ; \qquad \{a,b,i,j\} \not\subset \text{CAS}, \tag{2.18}$$

where barred indices refer to the PNO basis.

Now let us turn to domain based LPNO approach (DLPNO) for TCCSD. There the procedure is similar to LPNO described above and we again take advantage of PNO basis, but this time several steps differ, in order to keep the scaling of computational costs almost linear.

Again, we start with localized orbitals. The next step is to construct orbital domains from Adler and Werner [2011]. First we project out the localized active and occupied orbitals $|i\rangle$ from the former AO basis $|\mu\rangle$

$$|\tilde{\mu}\rangle = \left(1 - \sum_i |i\rangle\langle i|\right)|\mu\rangle\,, \tag{2.19}$$

hence we constructed the projected atomic orbitals (PAO). PAOs span local virtual space of a given orbital pair. However, they are not orthogonal. Therefore, PAOs can be truncated further based on their mutual differential overlap integrals to remove redundancy in PAO basis

$$(\mathrm{DOI})_{\mu i} = \sqrt{\int_{\mathbb{R}^3} |\phi_\mu|^2 \, |\phi_i|^2 \, \mathrm{d}^3\mathbf{r}}, \tag{2.20}$$

when $(\mathrm{DOI})_{\mu i}$ is below certain threshold.

Then, for each orbital pair, we continue by transforming its virtual space in PAO basis to PNO basis by employing equations 2.17 and following steps analogous to LPNO approach described above, including the PNO truncation. Then, the final PAO/PNO transformation matrix is constructed as

$$\tilde{\mathbf{d}}_{ij} = \mathbf{I}_{n\times n} \oplus \mathbf{d}_{ij} \tag{2.21}$$

for a given pair by enlarging the diagonalized form of pair density matrix $\mathbf{D}_{ij}$, i.e. $\mathbf{d}_{ij}$, by a unit matrix $\mathbf{I}_{n\times n}$, with n, the number of singly occupied and virtual active orbitals.

The CC equations now remain the same as in equation 2.18, but now the singles cluster operator $\bar{\mathbf{T}}^{S}_{\mathrm{ext}}$ is barred as well, allowing the linear scaling of computational costs.

Let us now turn to results from application of our DLPNO-TCCSD implementation on the simplified model of Iron(II)-Porphyrin depicted in Figure 1.5d. Iron(II)-Porphyrin participates in processes in biology and material science and it features close-lying spin states. It has been studied previously with large CASSCF calculations by Manni et al. [2019], which makes it a useful benchmark for DLPNO-TCCSD.

In Figure 2.3a, we can see retrieved percentage of correlation energy with respect to T_{cutPairs} threshold, as compared with the results from the full canonical version of TCCSD. The energies are near convergence at cutoff of about 10^{-5}, with the difference of about 0.005% from the canonical version, which takes much larger computational effort.

A similar dependency on T_{cutPNO} threshold is shown in Figure 2.3b. A discrepancy between the accuracy for different spins states is apparent for larger CAS. After an overestimate of energy, the method converges to nearly 100% for T_{cutPNO} of 10^{-5}.

Overall, the DLPNO-TCCSD was able retrieve more than 99.8% of the canonical correlation energy with default values of the aforementioned tresholds, reaching its target to allow for almost linear scaling of the CC-part of the calculation with comparable results.

In this dissertation text, we do not reiterate all the results we obtained from the DLPNO-TCCSD method, which we instead present in our peer-reviewed study “Near-Linear Scaling in DMRG-Based Tailored Coupled Clusters: An Implementation of DLPNO-TCCSD and DLPNO-TCCSD(T)” by Jakub Lang, Andrej Antalík, Libor Veis, Jan Brandejs, Jiří Brabec, Örs Legeza and Jiří Pittner included in the attachment below, Lang et al. [2020].

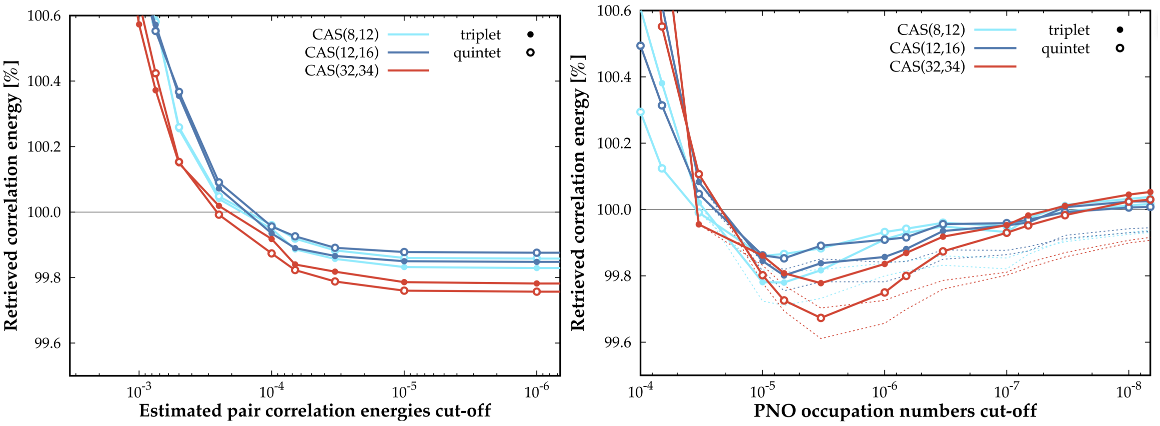

(a) The percentage of retrieved correlation energy with respect to the T_{cutPairs} threshold. The convergence is reached starting from $T_{\text{cutPairs}} = 10^{-5}$.

(b) Retrieved correlation energy with respect to T_{cutPNO}. The dashed line corresponds to fixed value of $T_{\text{cutPairs}} = 10^{-5}$, while the solid line to a dynamically adjusted value.

Figure 2.3: DLPNO-TCCSD method applied on Iron(II)-Porphyrin model in def2-SVP basis as compared with canonocal TCCSD energies. Chart from our work Lang et al. [2020].

Conclusion

In this work, I present three implementations of advanced quantum chemical methods, to the development of which I contributed during my PhD study. The central topic of my work has been the density matrix renormalization group method (DMRG) and its generalizations for special tasks.

The first goal has been to push the boundaries for the active space size even further than the current state-of-the art. We employed the high performance computing (HPC) techniques to achieve this through parallelization to thousands of CPU cores and by distributing tensors to RAM memory of hundreds of computers. This resulted in the MOLMPS program, which is currently being prepared for a source code release for the community, as the first massively parallel DMRG program. During the developments, MOMLPS presented a real programming challenge, since it implements the DMRG algorithm with numerous optimizations described in the method summary above. To give some perspective, our C++ implementation feature more than 40 000 lines of code. This is because we have written, from scratch, an implementation as efficient as the few great DMRG codes well-known in the community.

To test the scaling of MOLMPS, we have performed benchmark calculations of the model of Iron(II)-Porphyrin, of π-conjugated anthracene tetramer and of the FeMoco cofactor, using large active spaces up to CAS(113,76). We have showed that while the local memory is used, the runtime scales linearly with the number of CPU cores. Even when using the shared memory approach of our global memory model the parallelization still exhibits quite favourable scaling, but of course the effects of internode communication are significant. It may come as a surprise that in order to reach the parallel speedup, the system has to be sufficiently large, such that the task-list, generated for the $H \left|\mathbf{\Psi}\right\rangle$ application, provides enough tasks to properly saturate so many CPU cores.

The second goal has been to design a method which would address both the problem of static and dynamical correlation. These two phenomena are of different character and therefore we employed two different methods to treat them. Moreover, here, we chose to study the previously unexplored relativistic domain. We have built upon history of our wider group of in-house DMRG implementations and we have developed a relativistic 4c-DMRG to treat the static correlation. We have interfaced this method to the open-source program DIRAC and we have modified the relativistic coupled cluster equations in DIRAC to conform to 4c-TCCSD method. Our program has become the first ever implementation of relativistic 4c-TCCSD.

This method has been applied to three heavy diatomics: TlH, AsH, SbH and we successfully improved upon both CCSD and DMRG spectroscopic results for a corresponding CAS size. However, we were unable to “outperform” CCSD(T), which is, on the other hand, a method of higher order.

The third goal has been to push the boundaries for the basis size of the CCSD method within the framework of nonrelativistic TCCSD. Here, for the first time, domain-based local pair natural orbital (DLPNO) approach has been applied in the TCCSD method, by modifying the CC implementation in the Orca program and integrating it to its production version. A scalable behaviour has been demonstrated and we have shown that DLPNO-TCCSD accounts for more than 99% of correlation energy compared with the results from the standard TCCSD.

Last but not least, we have used the MPS state from DMRG to analyze the bonding structure. We implemented an efficient MPS contraction scheme to generate the von Neumann entropy and mutual information between various subsystems. in newly synthesised systems of diborane(4), i.e. B_2H_4 and a neutral zero-valent beryllium complex $Be(CAC)_2$. For the former, we confirmed the hypothesized structure with three-center two-electron bonds. In the later case, we used the correlation measures to hypothesize 5-center scheme, due to a delocalization of the two electrons from the electron-deficient bond stretching over the C-Be-C core, towards the nitrogen atoms on the rings.

Even though there are several appealing “firsts” here, the real added value for the community is the software, whose source code has been, or will be made freely available online.

As an outlook to the future, we consider implementation of GPU, or even TPU (dedicated tensor contraction processing unit) version of MOLMPS, promising a huge speedup on hardware recently developed for (neural-) tensor-network techniques for machine learning applications. For the case of relativistic 4c-TCCSD, we are already preparing a manuscript with the recently implemented generalization to “complex” groups, a domain accessible to only a handful implementations of quantum chemical methods.

Bibliography

Thomas B. Adler and Hans-Joachim Werner. An explicitly correlated local coupled cluster method for calculations of large molecules close to the basis set limit. *The Journal of Chemical Physics*, 135(14):144117, October 2011. doi: 10.1063/1.3647565. URL `https://doi.org/10.1063/1.3647565`.

Adel Almoukhalalati, Stefan Knecht, Hans Jørgen Aa. Jensen, Kenneth G. Dyall, and Trond Saue. Electron correlation within the relativistic no-pair approximation. *The Journal of Chemical Physics*, 145(7):074104, August 2016. doi: 10.1063/1.4959452. URL `https://doi.org/10.1063/1.4959452`.

Andrej Antalík, Libor Veis, Jiří Brabec, Ondřej Demel, Örs Legeza, and Jiří Pittner. Toward the efficient local tailored coupled cluster approximation and the peculiar case of oxo-mn(salen). *The Journal of Chemical Physics*, 151(8): 084112, August 2019. doi: 10.1063/1.5110477. URL `https://doi.org/10.1063/1.5110477`.

Merle Arrowsmith, Holger Braunschweig, Mehmet Ali Celik, Theresa Dellermann, Rian D. Dewhurst, William C. Ewing, Kai Hammond, Thomas Kramer, Ivo Krummenacher, Jan Mies, Krzysztof Radacki, and Julia K. Schuster. Neutral zero-valent s-block complexes with strong multiple bonding. *Nature Chemistry*, 8(9):890–894, June 2016. doi: 10.1038/nchem.2542. URL `https://doi.org/10.1038/nchem.2542`.

Yehuda B. Band and Yshai Avishai. Electronic structure of multielectron systems. In *Quantum Mechanics with Applications to Nanotechnology and Information Science*, pages 545–578. Elsevier, 2013. doi: 10.1016/b978-0-444-53786-7.00010-1. URL `https://doi.org/10.1016/b978-0-444-53786-7.00010-1`.

G. Barcza, Ö. Legeza, K. H. Marti, and M. Reiher. Quantum-information analysis of electronic states of different molecular structures. *Physical Review A*, 83(1), January 2011. doi: 10.1103/physreva.83.012508. URL `https://doi.org/10.1103/physreva.83.012508`.

Rodney J. Bartlett and Monika Musiał. Coupled-cluster theory in quantum chemistry. *Reviews of Modern Physics*, 79(1):291–352, February 2007. doi: 10.1103/revmodphys.79.291. URL `https://doi.org/10.1103/revmodphys.79.291`.

Stefano Battaglia, Sebastian Keller, and Stefan Knecht. Efficient relativistic density-matrix renormalization group implementation in a matrix-product formulation. *Journal of Chemical Theory and Computation*, 14(5):2353–2369, March 2018. doi: 10.1021/acs.jctc.7b01065. URL `https://doi.org/10.1021/acs.jctc.7b01065`.

Helmut Beinert, Richard H. Holm, and Eckard Munck. Iron-sulfur clusters: Nature's modular, multipurpose structures. *Science*, 277(5326):653–659, August 1997. doi: 10.1126/science.277.5326.653. URL `https://doi.org/10.1126/science.277.5326.653`.

Jacob Biamonte and Ville Bergholm. Tensor networks in a nutshell, 2017. URL `https://arxiv.org/abs/1708.00006`.

Katharina Boguslawski, Paweł Tecmer, Gergely Barcza, Örs Legeza, and Markus Reiher. Orbital entanglement in bond-formation processes. *Journal of Chemical Theory and Computation*, 9(7):2959–2973, June 2013. doi: 10.1021/ct400247p. URL `https://doi.org/10.1021/ct400247p`.

Jiri Brabec, Jan Brandejs, Karol Kowalski, Sotiris Xantheas, Örs Legeza, and Libor Veis. Massively parallel quantum chemical density matrix renormalization group method. *Journal of Computational Chemistry*, 42(8):534–544, December 2020. doi: 10.1002/jcc.26476. URL `https://doi.org/10.1002/jcc.26476`.

Jan Brandejs, Libor Veis, Szilárd Szalay, Gergely Barcza, Jiří Pittner, and Örs Legeza. Quantum information-based analysis of electron-deficient bonds. *The Journal of Chemical Physics*, 150(20):204117, May 2019. doi: 10.1063/1.5093497. URL `https://doi.org/10.1063/1.5093497`.

Jan Brandejs, Jakub Višňák, Libor Veis, Mihály Maté, Örs Legeza, and Jiří Pittner. Toward DMRG-tailored coupled cluster method in the 4c-relativistic domain. *The Journal of Chemical Physics*, 152(17):174107, May 2020. doi: 10.1063/1.5144974. URL `https://doi.org/10.1063/1.5144974`.

Laimutis Bytautas, Joseph Ivanic, and Klaus Ruedenberg. Split-localized orbitals can yield stronger configuration interaction convergence than natural orbitals. *The Journal of Chemical Physics*, 119(16):8217–8224, October 2003. doi: 10.1063/1.1610434. URL `https://doi.org/10.1063/1.1610434`.

Garnet Kin-Lic Chan. An algorithm for large scale density matrix renormalization group calculations. *The Journal of Chemical Physics*, 120(7):3172–3178, February 2004. doi: 10.1063/1.1638734. URL `https://doi.org/10.1063/1.1638734`.

Garnet Kin-Lic Chan, Anna Keselman, Naoki Nakatani, Zhendong Li, and Steven R. White. Matrix product operators, matrix product states, and ab initio density matrix renormalization group algorithms. *The Journal of Chemical Physics*, 145(1):014102, July 2016. doi: 10.1063/1.4955108. URL `https://doi.org/10.1063/1.4955108`.

Sheng-Lung Chou, Jen-Iu Lo, Yu-Chain Peng, Meng-Yeh Lin, Hsiao-Chi Lu, Bing-Ming Cheng, and J. F. Ogilvie. Identification of diborane(4) with bridging b–h–b bonds. *Chemical Science*, 6(12):6872–6877, 2015. doi: 10.1039/c5sc02586a. URL `https://doi.org/10.1039/c5sc02586a`.

Jiří Čížek. On the correlation problem in atomic and molecular systems. calculation of wavefunction components in ursell-type expansion using quantum-field theoretical methods. *The Journal of Chemical Physics*, 45(11):4256–4266, December 1966. doi: 10.1063/1.1727484. URL `https://doi.org/10.1063/1.1727484`.

Valerie Coffman, Joydip Kundu, and William K. Wootters. Distributed entanglement. *Physical Review A*, 61(5), April 2000. doi: 10.1103/physreva.61.052306. URL `https://doi.org/10.1103/physreva.61.052306`.

J.-M. Combes. The born-oppenheimer approximation. In *The Schrödinger Equation*, pages 139–159. Springer Vienna, 1977. doi: 10.1007/978-3-7091-7673-3_7. URL `https://doi.org/10.1007/978-3-7091-7673-3_7`.

M. Crouzeix, B. Philippe, and M. Sadkane. The davidson method. *SIAM Journal on Scientific Computing*, 15(1):62–76, January 1994. doi: 10.1137/0915004. URL `https://doi.org/10.1137/0915004`.

Dirac18, 2018. DIRAC, a relativistic ab initio electronic structure program, Release DIRAC18, written by T. Saue, L. Visscher, H. J. Aa. Jensen, and R. Bast, with contributions from V. Bakken, K. G. Dyall, S. Dubillard, U. Ekström, E. Eliav, T. Enevoldsen, E. Faßhauer, T. Fleig, O. Fossgaard, A. S. P. Gomes, E. D. Hedegård, T. Helgaker, J. Henriksson, M. Iliaš, Ch. R. Jacob, S. Knecht, S. Komorovský, O. Kullie, J. K. Lærdahl, C. V. Larsen, Y. S. Lee, H. S. Nataraj, M. K. Nayak, P. Norman, G. Olejniczak, J. Olsen, J. M. H. Olsen, Y. C. Park, J. K. Pedersen, M. Pernpointner, R. di Remigio, K. Ruud, P. Sałek, B. Schimmelpfennig, A. Shee, J. Sikkema, A. J. Thorvaldsen, J. Thyssen, J. van Stralen, S. Villaume, O. Visser, T. Winther, and S. Yamamoto (available at `https://doi.org/10.5281/zenodo.2253986`, see also `http://www.diracprogram.org`).

Kenneth G. Dyall and Knut Faegri. *Introduction to Relativistic Quantum Chemistry*. Oxford University Press, July 2007. doi: 10.1093/oso/9780195140866.001.0001. URL `https://doi.org/10.1093/oso/9780195140866.001.0001`.

C. Edmiston and M. Krauss. Configuration-interaction calculation of h3 and h2. *The Journal of Chemical Physics*, 42(3):1119–1120, February 1965. doi: 10.1063/1.1696050. URL `https://doi.org/10.1063/1.1696050`.

T. Engel. *Quantum Chemistry and Spectroscopy*. Pearson/Dorling Kindersley, 2006. ISBN 9788131712856.

Daniel Ess, Laura Gagliardi, and Sharon Hammes-Schiffer. Introduction: Computational design of catalysts from molecules to materials. *Chemical Reviews*, 119(11):6507–6508, June 2019. doi: 10.1021/acs.chemrev.9b00296. URL `https://doi.org/10.1021/acs.chemrev.9b00296`.

Marc Fontecave. Iron-sulfur clusters: ever-expanding roles. *Nature Chemical Biology*, 2(4):171–174, April 2006. doi: 10.1038/nchembio0406-171. URL `https://doi.org/10.1038/nchembio0406-171`.

Leon Freitag, Stefan Knecht, Sebastian F. Keller, Mickaël G. Delcey, Francesco Aquilante, Thomas Bondo Pedersen, Roland Lindh, Markus Reiher, and Leticia González. Orbital entanglement and CASSCF analysis of the ru–NO bond in a ruthenium nitrosyl complex. *Physical Chemistry Chemical Physics*, 17(22):14383–14392, 2015. doi: 10.1039/c4cp05278a. URL `https://doi.org/10.1039/c4cp05278a`.

Alan Gara, José E. Moreira, Tejas S. Karkhanis, José E. Moreira, José E. Moreira, Michael Flynn, Yoichi Muraoka, Matthias Bollhöfer, José I. Aliaga, Alberto F. Martin, Enrique S. Quintana-Ortí, Dhabaleswar K. Panda, Sayantan Sur, Josef Weidendorfer, Paul Petersen, Paul Petersen, Mark Dewing, Douglas Armstrong, Arch D. Robison, Sudhakar Yalamanchili, Xiaosong Ma, Xiaosong Ma, and David Padua. InfiniBand. In *Encyclopedia of Parallel Computing*, pages 927–935. Springer US, 2011. doi: 10.1007/978-0-387-09766-4_21. URL `https://doi.org/10.1007/978-0-387-09766-4_21`.

Anirban Ghosh, Rajat K. Chaudhuri, and Sudip Chattopadhyay. Relativistic state-specific multireference coupled cluster theory description for bond-breaking energy surfaces. *The Journal of Chemical Physics*, 145(12):124303, September 2016. doi: 10.1063/1.4962911. URL `https://doi.org/10.1063/1.4962911`.

G. H. Golub and C. Reinsch. Singular value decomposition and least squares solutions. *Numerische Mathematik*, 14(5):403–420, April 1970. doi: 10.1007/bf02163027. URL `https://doi.org/10.1007/bf02163027`.

Pavlo Golub, Andrej Antalik, Libor Veis, and Jiri Brabec. Machine learning-assisted selection of active spaces for strongly correlated transition metal systems. *Journal of Chemical Theory and Computation*, 17(10):6053–6072, September 2021. doi: 10.1021/acs.jctc.1c00235. URL `https://doi.org/10.1021/acs.jctc.1c00235`.

G. Hager, E. Jeckelmann, H. Fehske, and G. Wellein. Parallelization strategies for density matrix renormalization group algorithms on shared-memory systems. *Journal of Computational Physics*, 194(2):795–808, March 2004. doi: 10.1016/j.jcp.2003.09.018. URL `https://doi.org/10.1016/j.jcp.2003.09.018`.

Claudia Hampel and Hans-Joachim Werner. Local treatment of electron correlation in coupled cluster theory. *The Journal of Chemical Physics*, 104(16):6286–6297, April 1996. doi: 10.1063/1.471289. URL `https://doi.org/10.1063/1.471289`.

Johannes Hauschild and Frank Pollmann. Fermions and the Jordan-Wigner transformation: TeNPy documentation. `https://tenpy.github.io/intro/JordanWigner.html`, 2020. Webpage accessed: 2022-01-15.

Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen. *Molecular Electronic-Structure Theory*. John Wiley & Sons, Ltd, August 2000. doi: 10.1002/9781119019572. URL `https://doi.org/10.1002/9781119019572`.

Jacob S. Higgins, Lawson T. Lloyd, Sara H. Sohail, Marco A. Allodi, John P. Otto, Rafael G. Saer, Ryan E. Wood, Sara C. Massey, Po-Chieh Ting, Robert E. Blankenship, and Gregory S. Engel. Photosynthesis tunes quantum-mechanical mixing of electronic and vibrational states to steer exciton energy transfer. *Proceedings of the National Academy of Sciences*, 118(11):e2018240118, March 2021. doi: 10.1073/pnas.2018240118. URL `https://doi.org/10.1073/pnas.2018240118`.

Yu hong Lam, Yuriy Abramov, Ravi S. Ananthula, Jennifer M. Elward, Lori R. Hilden, Sten O. Nilsson Lill, Per-Ola Norrby, Antonio Ramirez, Edward C. Sherer, Jason Mustakis, and Gerald J. Tanoury. Applications of quantum chemistry in pharmaceutical process development: Current state and opportunities. *Organic Process Research & Development*, 24(8):1496–1507, July 2020. doi: 10.1021/acs.oprd.0c00222. URL `https://doi.org/10.1021/acs.oprd.0c00222`.

C. Hubig, J. Haegeman, and U. Schollwöck. Error estimates for extrapolations with matrix-product states. *Physical Review B*, 97(4), January 2018. doi: 10.1103/physrevb.97.045125. URL `https://doi.org/10.1103/physrevb.97.045125`.

George Samuel Hurst, Jack D. Graybeal, Stoner, John Oliver, and Steven Chu. Spectroscopy. encyclopedia britannica, 2021. URL `https://www.britannica.com/science/spectroscopy`. accessed January 10th 2022.

Frank Jensen. *Introduction to Computational Chemistry*. John Wiley and Sons, Inc., Hoboken, NJ, USA, 2006. ISBN 0470011874.

P. Jordan and E. Wigner. Über das paulische Äquivalenzverbot. *Zeitschrift für Physik*, 47(9-10):631–651, September 1928. doi: 10.1007/bf01331938. URL `https://doi.org/10.1007/bf01331938`.

Sebastian F. Keller and Markus Reiher. Determining factors for the accuracy of DMRG in chemistry. *CHIMIA International Journal for Chemistry*, 68(4):200–203, April 2014. doi: 10.2533/chimia.2014.200. URL `https://doi.org/10.2533/chimia.2014.200`.

Tomoko Kinoshita, Osamu Hino, and Rodney J. Bartlett. Coupled-cluster method tailored by configuration interaction. *The Journal of Chemical Physics*, 123(7):074106, August 2005. doi: 10.1063/1.2000251. URL `https://doi.org/10.1063/1.2000251`.

Stefan Knecht, Örs Legeza, and Markus Reiher. Communication: Four-component density matrix renormalization group. *The Journal of Chemical Physics*, 140(4):041101, January 2014. doi: 10.1063/1.4862495. URL `https://doi.org/10.1063/1.4862495`.

Christian Kollmar, Kantharuban Sivalingam, Yang Guo, and Frank Neese. An efficient implementation of the NEVPT2 and CASPT2 methods avoiding higher-order density matrices. *The Journal of Chemical Physics*, 155(23):234104, December 2021. doi: 10.1063/5.0072129. URL `https://doi.org/10.1063/5.0072129`.

Yuki Kurashige and Takeshi Yanai. High-performance ab initio density matrix renormalization group method: Applicability to large-scale multireference problems for metal compounds. *The Journal of Chemical Physics*, 130(23):234114, June 2009. doi: 10.1063/1.3152576. URL `https://doi.org/10.1063/1.3152576`.

Yuki Kurashige and Takeshi Yanai. Second-order perturbation theory with a density matrix renormalization group self-consistent field reference function: Theory and application to the study of chromium dimer. *The Journal of Chemical Physics*, 135(9):094104, September 2011. doi: 10.1063/1.3629454. URL `https://doi.org/10.1063/1.3629454`.

Yuki Kurashige, Garnet Kin-Lic Chan, and Takeshi Yanai. Entangled quantum electronic wavefunctions of the mn4cao5 cluster in photosystem II. *Nature Chemistry*, 5(8):660–666, June 2013. doi: 10.1038/nchem.1677. URL `https://doi.org/10.1038/nchem.1677`.

Werner Kutzelnigg. Quantum chemistry in fock space. i. the universal wave and energy operators. *The Journal of Chemical Physics*, 77(6):3081–3097, September 1982. doi: 10.1063/1.444231. URL `https://doi.org/10.1063/1.444231`.

Tran Nguyen Lan, Yuki Kurashige, and Takeshi Yanai. Scalar relativistic calculations of hyperfine coupling constants using ab initio density matrix renormalization group method in combination with third-order douglas–kroll–hess transformation: Case studies on 4d transition metals. *Journal of Chemical Theory and Computation*, 11(1):73–81, December 2014. doi: 10.1021/ct5007778. URL `https://doi.org/10.1021/ct5007778`.

Kyle M. Lancaster, Michael Roemelt, Patrick Ettenhuber, Yilin Hu, Markus W. Ribbe, Frank Neese, Uwe Bergmann, and Serena DeBeer. X-ray emission spectroscopy evidences a central carbon in the nitrogenase iron-molybdenum cofactor. *Science*, 334(6058):974–977, November 2011. doi: 10.1126/science.1206445. URL `https://doi.org/10.1126/science.1206445`.

Jakub Lang, Andrej Antalík, Libor Veis, Jan Brandejs, Jiří Brabec, Örs Legeza, and Jiří Pittner. Near-linear scaling in DMRG-based tailored coupled clusters: An implementation of DLPNO-TCCSD and DLPNO-TCCSD(t). *Journal of Chemical Theory and Computation*, 16(5):3028–3040, April 2020. doi: 10.1021/acs.jctc.0c00065. URL `https://doi.org/10.1021/acs.jctc.0c00065`.

Ö. Legeza and J. Sólyom. Optimizing the density-matrix renormalization group method using quantum information entropy. *Physical Review B*, 68(19), November 2003. doi: 10.1103/physrevb.68.195116. URL `https://doi.org/10.1103/physrevb.68.195116`.

Ö. Legeza, J. Röder, and B. A. Hess. Controlling the accuracy of the density-matrix renormalization-group method: The dynamical block state selection approach. *Physical Review B*, 67(12), March 2003. doi: 10.1103/physrevb.67.125114. URL `https://doi.org/10.1103/physrevb.67.125114`.

Örs Legeza, Libor Veis, and Tamas Mosoni. QC-DMRG-Budapest, a program for quantum chemical DMRG calculations, 2018.

Benjamin G. Levine, Andrew S. Durden, Michael P. Esch, Fangchun Liang, and Yinan Shu. CAS without SCF—why to use CASCI and where to get the orbitals. *The Journal of Chemical Physics*, 154(9):090902, March 2021. doi: 10.1063/5.0042147. URL `https://doi.org/10.1063/5.0042147`.

Xiangzhu Li and Josef Paldus. Reduced multireference CCSD method: An effective approach to quasidegenerate states. *The Journal of Chemical Physics*, 107(16):6257–6269, October 1997. doi: 10.1063/1.474289. URL `https://doi.org/10.1063/1.474289`.

Zhendong Li, Junhao Li, Nikesh S. Dattani, C. J. Umrigar, and Garnet Kin-Lic Chan. The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations. *The Journal of Chemical Physics*, 150(2):024302, January 2019. doi: 10.1063/1.5063376. URL `https://doi.org/10.1063/1.5063376`.

Dmitry I. Lyakh, Victor F. Lotrich, and Rodney J. Bartlett. The 'tailored' CCSD(t) description of the automerization of cyclobutadiene. *Chemical Physics Letters*, 501(4-6):166–171, January 2011a. doi: 10.1016/j.cplett.2010.11.058. URL `https://doi.org/10.1016/j.cplett.2010.11.058`.

Dmitry I. Lyakh, Monika Musiał, Victor F. Lotrich, and Rodney J. Bartlett. Multireference nature of chemistry: The coupled-cluster view. *Chemical Reviews*, 112(1):182–243, December 2011b. doi: 10.1021/cr2001417. URL `https://doi.org/10.1021/cr2001417`.

Giovanni Li Manni, Daniel Kats, David P. Tew, and Ali Alavi. Role of valence and semicore electron correlation on spin gaps in fe(II)-porphyrins. *Journal of Chemical Theory and Computation*, 15(3):1492–1497, January 2019. doi: 10.1021/acs.jctc.8b01277. URL `https://doi.org/10.1021/acs.jctc.8b01277`.

Ann Melnichuk and Rodney J. Bartlett. Relaxed active space: Fixing tailored-CC with high order coupled cluster. II. *The Journal of Chemical Physics*, 140(6):064113, February 2014. doi: 10.1063/1.4862676. URL `https://doi.org/10.1063/1.4862676`.

Anna Melnichuk and Rodney J. Bartlett. Relaxed active space: Fixing tailored-CC with high order coupled cluster. i. *The Journal of Chemical Physics*, 137(21):214103, December 2012. doi: 10.1063/1.4767900. URL `https://doi.org/10.1063/1.4767900`.

Gerrit Moritz, Alexander Wolf, and Markus Reiher. Relativistic DMRG calculations on the curve crossing of cesium hydride. *The Journal of Chemical Physics*, 123(18):184105, November 2005. doi: 10.1063/1.2104447. URL `https://doi.org/10.1063/1.2104447`.

Naoki Nakatani. Matrix product states and density matrix renormalization group algorithm. In *Reference Module in Chemistry, Molecular Sciences and Chemical Engineering*. Elsevier, 2018. doi: 10.1016/b978-0-12-409547-2.11473-8. URL `https://doi.org/10.1016/b978-0-12-409547-2.11473-8`.

Naoki Nakatani and Garnet Kin-Lic Chan. Efficient tree tensor network states (TTNS) for quantum chemistry: Generalizations of the density matrix renormalization group algorithm. *The Journal of Chemical Physics*, 138(13):134113, April 2013. doi: 10.1063/1.4798639. URL `https://doi.org/10.1063/1.4798639`.

Huliyar S. Nataraj, Mihály Kállay, and Lucas Visscher. General implementation of the relativistic coupled-cluster method. *The Journal of Chemical Physics*, 133(23):234109, December 2010. doi: 10.1063/1.3518712. URL `https://doi.org/10.1063/1.3518712`.

Frank Neese. Software update: the ORCA program system, version 4.0. *WIREs Computational Molecular Science*, 8(1), July 2017. doi: 10.1002/wcms.1327. URL `https://doi.org/10.1002/wcms.1327`.

Frank Neese, Frank Wennmohs, and Andreas Hansen. Efficient and accurate local approximations to coupled-electron pair approaches: An attempt to revive the pair natural orbital method. *The Journal of Chemical Physics*, 130(11):114108, March 2009. doi: 10.1063/1.3086717. URL `https://doi.org/10.1063/1.3086717`.

Frank Neese, Frank Wennmohs, Ute Becker, and Christoph Riplinger. The ORCA quantum chemistry program package. *The Journal of Chemical Physics*, 152 (22):224108, June 2020. doi: 10.1063/5.0004608. URL `https://doi.org/10.1063/5.0004608`.

Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2009. doi: 10.1017/cbo9780511976667. URL `https://doi.org/10.1017/cbo9780511976667`.

Masanori Ohya and Noboru Watanabe. Quantum entropy and its applications to quantum communication and statistical physics. *Entropy*, 12(5):1194–1245, May 2010. doi: 10.3390/e12051194. URL `https://doi.org/10.3390/e12051194`.

Roberto Olivares-Amaya, Weifeng Hu, Naoki Nakatani, Sandeep Sharma, Jun Yang, and Garnet Kin-Lic Chan. Theab-initiodensity matrix renormalization group in practice. *The Journal of Chemical Physics*, 142(3):034102, January 2015. doi: 10.1063/1.4905329. URL `https://doi.org/10.1063/1.4905329`.

Buhm Soon Park and DeWitt Stetten. A principle written in diagrams: The aufbau principle for molecules and its visual representations, 1927–1932. In *Boston Studies in the Philosophy and History of Science*, pages 179–198. Springer Netherlands, 2001. doi: 10.1007/978-94-015-9737-1_11. URL `https://doi.org/10.1007/978-94-015-9737-1_11`.

Heechul Park, Nimrod Heldman, Patrick Rebentrost, Luigi Abbondanza, Alessandro Iagatti, Andrea Alessi, Barbara Patrizi, Mario Salvalaggio, Laura Bussotti, Masoud Mohseni, Filippo Caruso, Hannah C. Johnsen, Roberto Fusco, Paolo Foggi, Petra F. Scudo, Seth Lloyd, and Angela M. Belcher. Enhanced energy transport in genetically engineered excitonic networks. *Nature Materials*, 15(2):211–216, October 2015. doi: 10.1038/nmat4448. URL `https://doi.org/10.1038/nmat4448`.

Jae Woo Park, Rachael Al-Saadon, Matthew K. MacLeod, Toru Shiozaki, and Bess Vlaisavljevich. Multireference electron correlation methods: Journeys along potential energy surfaces. *Chemical Reviews*, 120(13):5878–5909, April

2020. doi: 10.1021/acs.chemrev.9b00496. URL `https://doi.org/10.1021/acs.chemrev.9b00496`.

Michael E. Peskin. *An Introduction To Quantum Field Theory*. CRC Press, May 2018. doi: 10.1201/9780429503559. URL `https://doi.org/10.1201/9780429503559`.

Piotr Piecuch and Ludwik Adamowicz. State-selective multireference coupled-cluster theory employing the single-reference formalism: Implementation and application to the h8 model system. *The Journal of Chemical Physics*, 100(8): 5792–5809, April 1994. doi: 10.1063/1.467143. URL `https://doi.org/10.1063/1.467143`.

Piotr Piecuch, Nevin Oliphant, and Ludwik Adamowicz. A state-selective multireference coupled-cluster theory employing the single-reference formalism. *The Journal of Chemical Physics*, 99(3):1875–1900, aug 1993. doi: 10.1063/1.466179. URL `https://doi.org/10.1063/1.466179`.

Peter Pinski, Christoph Riplinger, Edward F. Valeev, and Frank Neese. Sparse maps—a systematic infrastructure for reduced-scaling electronic structure methods. i. an efficient and simple linear scaling local MP2 method that uses an intermediate basis of pair natural orbitals. *The Journal of Chemical Physics*, 143(3):034108, jul 2015. doi: 10.1063/1.4926879. URL `https://doi.org/10.1063%2F1.4926879`.

Péter Pulay. Convergence acceleration of iterative sequences. the case of scf iteration. *Chemical Physics Letters*, 73(2):393–398, July 1980. doi: 10.1016/0009-2614(80)80396-4. URL `https://doi.org/10.1016/0009-2614(80)80396-4`.

Peter Pulay. Localizability of dynamic electron correlation. *Chemical Physics Letters*, 100(2):151–154, September 1983. doi: 10.1016/0009-2614(83)80703-9. URL `https://doi.org/10.1016/0009-2614(83)80703-9`.

Pekka Pyykko and Jean Paul Desclaux. Relativity and the periodic system of elements. *Accounts of Chemical Research*, 12(8):276–281, August 1979. doi: 10.1021/ar50140a002. URL `https://doi.org/10.1021/ar50140a002`.

Svein Saebo/ and Peter Pulay. Fourth-order mo/ller–plessett perturbation theory in the local correlation treatment. i. method. *The Journal of Chemical Physics*, 86(2):914–922, January 1987. doi: 10.1063/1.452293. URL `https://doi.org/10.1063/1.452293`.

Masaaki Saitow, Yuki Kurashige, and Takeshi Yanai. Multireference configuration interaction theory using cumulant reconstruction with internal contraction of density matrix renormalization group wave function. *The Journal of Chemical Physics*, 139(4):044118, July 2013. doi: 10.1063/1.4816627. URL `https://doi.org/10.1063/1.4816627`.

Ana Sánchez-Grande, Bruno de la Torre, José Santos, Borja Cirera, Koen Lauwaet, Taras Chutora, Shayan Edalatmanesh, Pingo Mutombo, Johanna Rosen, Radek Zbořil, Rodolfo Miranda, Jonas Björk, Pavel Jelínek, Nazario

Martín, and David Écija. On-surface synthesis of ethynylene-bridged anthracene polymers. *Angewandte Chemie*, 131(20):6631–6635, March 2019. doi: 10.1002/ange.201814154. URL `https://doi.org/10.1002/ange.201814154`.

T. Saue and H. J. Aa Jensen. Quaternion symmetry in relativistic molecular calculations: The dirac–hartree–fock method. *The Journal of Chemical Physics*, 111(14):6211–6222, October 1999. doi: 10.1063/1.479958. URL `https://doi.org/10.1063/1.479958`.

Elvira R. Sayfutyarova and Garnet Kin-Lic Chan. A state interaction spin-orbit coupling density matrix renormalization group method. *The Journal of Chemical Physics*, 144(23):234301, June 2016. doi: 10.1063/1.4953445. URL `https://doi.org/10.1063/1.4953445`.

Maximilian Schlosshauer. Quantum decoherence. *Physics Reports*, 831:1–57, October 2019. doi: 10.1016/j.physrep.2019.10.001. URL `https://doi.org/10.1016/j.physrep.2019.10.001`.

U. Schollwöck. The density-matrix renormalization group. *Reviews of Modern Physics*, 77(1):259–315, April 2005. doi: 10.1103/revmodphys.77.259. URL `https://doi.org/10.1103/revmodphys.77.259`.

Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, January 2011. doi: 10.1016/j.aop.2010.09.012. URL `https://doi.org/10.1016/j.aop.2010.09.012`.

W. H. Eugen Schwarz. An introduction to relativistic quantum chemistry. In *Challenges and Advances in Computational Chemistry and Physics*, pages 1–62. Springer Netherlands, 2010. doi: 10.1007/978-1-4020-9975-5_1. URL `https://doi.org/10.1007/978-1-4020-9975-5_1`.

Alessandro Sergi and Raymond Kapral. Quantum-classical limit of quantum correlation functions. *The Journal of Chemical Physics*, 121(16):7565, 2004. doi: 10.1063/1.1797191. URL `https://doi.org/10.1063/1.1797191`.

Sandeep Sharma, Kantharuban Sivalingam, Frank Neese, and Garnet Kin-Lic Chan. Low-energy spectrum of iron–sulfur clusters directly from many-particle quantum mechanics. *Nature Chemistry*, 6(10):927–933, August 2014. doi: 10.1038/nchem.2041. URL `https://doi.org/10.1038/nchem.2041`.

Isaiah Shavitt and Rodney J. Bartlett. *Many – Body Methods in Chemistry and Physics*. Cambridge University Press, 2009. doi: 10.1017/cbo9780511596834. URL `https://doi.org/10.1017/cbo9780511596834`.

Avijit Shee, Lucas Visscher, and Trond Saue. Analytic one-electron properties at the 4-component relativistic coupled cluster level with inclusion of spin-orbit coupling. *The Journal of Chemical Physics*, 145(18):184107, November 2016. doi: 10.1063/1.4966643. URL `https://doi.org/10.1063/1.4966643`.

John C. Slater and Keith H. Johnson. Quantum chemistry and catalysis. *Physics Today*, 27(10):34–41, October 1974. doi: 10.1063/1.3128915. URL `https://doi.org/10.1063/1.3128915`.

Miles Stoudenmire. Fermions and Jordan-Wigner String. `https://itensor.org/docs.cgi?page=tutorials/fermions`, 2016. Webpage accessed: 2022-01-17.

A. Szabo and N.S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory.* Dover Books on Chemistry. Dover Publications, 1996. ISBN 9780486691862.

Szilárd Szalay. Multipartite entanglement measures. *Physical Review A*, 92(4), October 2015. doi: 10.1103/physreva.92.042329. URL `https://doi.org/10.1103/physreva.92.042329`.

Szilárd Szalay, Max Pfeffer, Valentin Murg, Gergely Barcza, Frank Verstraete, Reinhold Schneider, and Örs Legeza. Tensor product methods and entanglement optimization for ab initio quantum chemistry. *International Journal of Quantum Chemistry*, 115(19):1342–1391, May 2015. doi: 10.1002/qua.24898. URL `https://doi.org/10.1002/qua.24898`.

Szilárd Szalay, Gergely Barcza, Tibor Szilvási, Libor Veis, and Örs Legeza. The correlation theory of the chemical bond. *Scientific Reports*, 7(1), May 2017. doi: 10.1038/s41598-017-02447-z. URL `https://doi.org/10.1038/s41598-017-02447-z`.

Kerry Taylor-Smith. Applications of quantum chemistry, 2019. URL `https://www.azoquantum.com/Article.aspx?ArticleID=121`. accessed January 12th 2022.

Jbrn Thyssen. *Development and Applications of Methods for Correlated Relativistic Calculations of Molecular Properties.* PhD thesis, Univeristy of Southern Denmark, 2001.

A. V. Titov, N. S. Mosyagin, A. B. Alekseyev, and R. J. Buenker. GRECP/MRD-CI calculations of spin-orbit splitting in ground state of tl and of spectroscopic properties of TlH. *International Journal of Quantum Chemistry*, 81(6):409–421, 2001. doi: 10.1002/1097-461x(2001)81:6⟨409::aid-qua1010⟩3.0.co;2-0. URL `https://doi.org/10.1002/1097-461x(2001)81:6<409::aid-qua1010>3.0.co;2-0`.

Luca Turin. A spectroscopic mechanism for primary olfactory reception. *Chemical Senses*, 21(6):773–791, 1996. doi: 10.1093/chemse/21.6.773. URL `https://doi.org/10.1093/chemse/21.6.773`.

O. Vahtras, J. Almlöf, and M.W. Feyereisen. Integral approximations for LCAO-SCF calculations. *Chemical Physics Letters*, 213(5-6):514–518, October 1993. doi: 10.1016/0009-2614(93)89151-7. URL `https://doi.org/10.1016/0009-2614(93)89151-7`.

Tanja van Mourik. First–principles quantum chemistry in the life sciences. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 362(1825):2653–2670, September 2004. doi: 10.1098/rsta.2004.1469. URL `https://doi.org/10.1098/rsta.2004.1469`.

Libor Veis, Andrej Antalík, Jiří Brabec, Frank Neese, Örs Legeza, and Jiří Pittner. Coupled cluster method with single and double excitations tailored by matrix product state wave functions. *The Journal of Physical Chemistry Letters*, 7 (20):4072–4078, October 2016. doi: 10.1021/acs.jpclett.6b01908. URL `https://doi.org/10.1021/acs.jpclett.6b01908`.

Lucas Visscher, Kenneth G. Dyall, and Timothy J. Lee. Kramers-restricted closed-shellCCSD theory. *International Journal of Quantum Chemistry*, 56 (S29):411–419, February 1995. doi: 10.1002/qua.560560844. URL `https://doi.org/10.1002/qua.560560844`.

Lucas Visscher, Timothy J. Lee, and Kenneth G. Dyall. Formulation and implementation of a relativistic unrestricted coupled-cluster method including noniterative connected triples. *The Journal of Chemical Physics*, 105(19):8769–8776, November 1996. doi: 10.1063/1.472655. URL `https://doi.org/10.1063/1.472655`.

Lucas Visscher, Ephraim Eliav, and Uzi Kaldor. Formulation and implementation of the relativistic fock-space coupled cluster method for molecules. *The Journal of Chemical Physics*, 115(21):9720–9726, December 2001. doi: 10.1063/1.1415746. URL `https://doi.org/10.1063/1.1415746`.

Dingheng Wang, Guangshe Zhao, Hengnu Chen, Zhexian Liu, Lei Deng, and Guoqi Li. Nonlinear tensor train format for deep neural network compression. *Neural Networks*, 144:320–333, December 2021. doi: 10.1016/j.neunet.2021.08.028. URL `https://doi.org/10.1016/j.neunet.2021.08.028`.

J. A. White. Numerical exact-diagonalization study of the one-dimensional symmetric kondo lattice. *Physical Review B*, 46(21):13905–13913, December 1992. doi: 10.1103/physrevb.46.13905. URL `https://doi.org/10.1103/physrevb.46.13905`.

Steven R. White and Richard L. Martin. Ab initio quantum chemistry using the density matrix renormalization group. *The Journal of Chemical Physics*, 110 (9):4127–4130, March 1999. doi: 10.1063/1.478295. URL `https://doi.org/10.1063/1.478295`.

Sebastian Wouters and Dimitri Van Neck. The density matrix renormalization group for ab initio quantum chemistry. *The European Physical Journal D*, 68 (9), September 2014. doi: 10.1140/epjd/e2014-50500-1. URL `https://doi.org/10.1140/epjd/e2014-50500-1`.

Sebastian Wouters, Naoki Nakatani, Dimitri Van Neck, and Garnet Kin-Lic Chan. Thouless theorem for matrix product states and subsequent post density matrix renormalization group methods. *Physical Review B*, 88(7), August 2013. doi: 10.1103/physrevb.88.075122. URL `https://doi.org/10.1103/physrevb.88.075122`.

T. Xiang. Density-matrix renormalization-group method in momentum space. *Physical Review B*, 53(16):R10445–R10448, April 1996. doi: 10.1103/physrevb.53.r10445. URL `https://doi.org/10.1103/physrevb.53.r10445`.

Takeshi Yanai and Garnet Kin-Lic Chan. Canonical transformation theory for multireference problems. *The Journal of Chemical Physics*, 124(19):194106, May 2006. doi: 10.1063/1.2196410. URL `https://doi.org/10.1063/1.2196410`.

List of Figures

List of Tables

List of Abbreviations

- DMRG - density matrix renormalization group
- QC - quantum chemistry
- CC - coupled clusters
- TCC - tailored coupled clusters
- CCSD - coupled clusters with single and double excitations
- CCSD(T) - CC with single, double and perturbative triple excitations
- MBPT - many-body perturbation theory
- LHS - left hand side
- HF - Hartree-Fock
- CI - configuration interaction
- CASCI - complete active space configuration interaction
- CASSCF - complete active space self-consistent field
- FCI - full configuration interaction
- MRCI - multi-reference configuration interaction
- MRCC - multi-reference coupled clusters
- DKH - Douglas-Kroll-Hess
- DKH3 - 3rd-order Douglas-Kroll-Hess
- QED - quantum electrodynamics
- PAO - projected atomic orbital
- PNO - pair natural orbital
- GRECP - generalized relativistic effective core potential
- MRD-CI - multireference single- and double-excitation CI
- PEPS - Projected Entangled Pair States
- TTN - tree tensor networks
- SVD - singular value decomposition
- RDM - reduced density matrix
- RAM - random access memory
- AO - atomic orbital
- MO - molecular orbital
- GM - global memory model
- HPC - high performance computing
- CPU - computational processing unit
- NMR - nuclear magnetic resonance

List of publications

All publications are included in the attachment.

1. Jiří Brabec, Jan Brandejs, Karol Kowalski, Sotiris Xantheas, Örs Legeza and Libor Veis. Massively parallel quantum chemical density matrix renormalization group method. *Journal of Computational Chemistry*, 42(8):534–544, December 2020. doi: 10.1002/jcc.26476. 5093497. URL https://doi.org/10.1002/jcc.26476. Brabec et al. [2020].
    - As one of the three core-developers, I participated significantly on the implementation. Notably, I implemented the parallel warm-up procedure CI-DEAS, including the pre-summation of operators. I programmed the Davidson diagonalizer, the SVD procedure, the generation of renormalization matrix and the saving and loading of operators to/from the disk. MOLMPS is a large project with more than 40 000 lines of code and its development took several years.
    - I carried out testing calculations.
    - I participated on the proposal of the parallel scheme.
2. Jan Brandejs, Libor Veis, Szilárd Szalay, Gergely Barcza, Jiří Pittner and Örs Legeza. Quantum information-based analysis of electron-deficient bonds. The *Journal of Chemical Physics*, 150(20):204117, May 2019. doi: 10.1063/1.5093497. URL https://doi.org/10.1063/1.5093497. Brandejs et al. [2019]
    - I programmed a module for the calculation of correlation measures and I parallelized it.
    - I made all the calculations.
    - I participated on the analysis of the results.
    - I wrote the manuscript.
3. Jan Brandejs, Jakub Višňák, Libor Veis, Mihály Maté, Örs Legeza and Jiří Pittner. Toward DMRG-tailored coupled cluster method in the 4c-relativistic domain. *The Journal of Chemical Physics*, 152(17):174107, May 2020. doi: 10.1063/1.5144974. URL https://doi.org/10.1063/1.5144974.
    - I participated on the implementation of the interface between Budapest DMRG and the DIRAC program.
    - I modified DIRAC to output MP2-NO orbitals for 4c-DMRG.
    - I performed all the relativistic DMRG and CCSD (4c-TCCSD) calculations.
    - I wrote the majority of the text of the manuscript.
4. Jakub Lang and Andrej Antalík, Libor Veis, Jan Brandejs, Jiří Brabec, Örs Legeza and Jiří Pittner. Near-linear scaling in DMRG-based tailored coupled clusters: An implementation of DLPNO-TCCSD and DLPNO-TCCSD(t). *Journal of Chemical Theory and Computation*, 16(5):3028–3040,

April 2020. doi: 10.1021/ acs.jctc.0c00065.
URL https://doi.org/10.1021/acs.jctc.0c00065.

  - I participated on the calculations.
  - I prepared a significant part of answers to the questions of referees and prepared important material and charts for these answers.

5. Libor Veis, Jan Brandejs, and Jiří Pittner. Strongly Correlated Systems and the Method of the Renormalization Group Matrix in Quantum Chemistry. *Chemické listy*, 112(10):655-666, October 2018.
   E-print https://arxiv.org/abs/1807.06289.
   URL http://www.chemicke-listy.cz/ojs3/index.php/chemicke-listy/article/view/3189.
   - Note: A review article written in Czech language.
   - Special issue to celebrate 90th birthday of prof. Rudolf Zahradník
   - I participated on writing of the text of this review.

A. Attachments