

UC Merced

UC Merced Previously Published Works

Title

Coupled cluster theory on modern heterogeneous supercomputers.

Permalink

<https://escholarship.org/uc/item/67j0m9zm>

Authors

Corzo, Hector

Hillers-Bendtsen, Andreas

Barnes, Ashleigh

et al.

Publication Date

2023

DOI

10.3389/fchem.2023.1154526

Peer reviewed



OPEN ACCESS

EDITED BY

Honghui Shang,
Institute of Computing Technology
(CAS), China

REVIEWED BY

Igor Ying Zhang,
Fudan University, China
Zhendong Li,
Beijing Normal University, China

*CORRESPONDENCE

Dmytro Bykov,
✉ bykovd@ornl.gov

RECEIVED 30 January 2023

ACCEPTED 11 May 2023

PUBLISHED 14 June 2023

CITATION

Corzo HH, Hillers-Bendtsen AE, Barnes A,
Zamani AY, Pawłowski F, Olsen J,
Jørgensen P, Mikkelsen KV and Bykov D
(2023), Coupled cluster theory on
modern
heterogeneous supercomputers.
Front. Chem. 11:1154526.
doi: 10.3389/fchem.2023.1154526

COPYRIGHT

© 2023 Corzo, Hillers-Bendtsen, Barnes,
Zamani, Pawłowski, Olsen, Jørgensen,
Mikkelsen and Bykov. This is an open-
access article distributed under the terms
of the [Creative Commons Attribution
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is
permitted, provided the original author(s)
and the copyright owner(s) are credited
and that the original publication in this
journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

Coupled cluster theory on modern heterogeneous supercomputers

Hector H. Corzo¹, Andreas Erbs Hillers-Bendtsen²,
Ashleigh Barnes¹, Abdulrahman Y. Zamani³, Filip Pawłowski⁴,
Jeppe Olsen⁵, Poul Jørgensen⁵, Kurt V. Mikkelsen² and
Dmytro Bykov^{1*}

¹Oak Ridge National Laboratory, Oak Ridge, TN, United States, ²Department of Chemistry, University of Copenhagen, Copenhagen, Denmark, ³Department of Chemistry and Biochemistry and Center for Chemical Computation and Theory, University of California, Merced, CA, United States, ⁴Department of Chemistry and Biochemistry, Auburn University, Auburn, AL, United States, ⁵Department of Chemistry, Aarhus University, Aarhus, Denmark

This study examines the computational challenges in elucidating intricate chemical systems, particularly through *ab-initio* methodologies. This work highlights the Divide-Expand-Consolidate (DEC) approach for coupled cluster (CC) theory—a linear-scaling, massively parallel framework—as a viable solution. Detailed scrutiny of the DEC framework reveals its extensive applicability for large chemical systems, yet it also acknowledges inherent limitations. To mitigate these constraints, the cluster perturbation theory is presented as an effective remedy. Attention is then directed towards the CPS (D-3) model, explicitly derived from a CC singles parent and a doubles auxiliary excitation space, for computing excitation energies. The reviewed new algorithms for the CPS (D-3) method efficiently capitalize on multiple nodes and graphical processing units, expediting heavy tensor contractions. As a result, CPS (D-3) emerges as a scalable, rapid, and precise solution for computing molecular properties in large molecular systems, marking it an efficient contender to conventional CC models.

KEYWORDS

coupled cluster theory, divide-expand-consolidate coupled cluster framework, cluster perturbation theory, excitation energies, tetrahydrocannabinol, deoxyribonucleic acid

1 Introduction

Over the past 6 decades, the field of computational chemistry and molecular modeling has aimed to solve for the energy and expectation values of wave functions for atomic and molecular systems. In the exact limit, the non-relativistic electronic contribution to the total energy of a many-body system can be obtained by finding the exact solution to the N -electron Schrödinger equation. However, due to the challenges in solving atomic and molecular systems composed of more than few electrons in orbitals with angular momentum $l \geq 1$, many numerical approximations have been rather crude. To reduce the mathematical complexity associated with solving multi-electron molecular systems, which often requires modeling a $3N$ dimensional space, many computational chemistry approximations have opted for partial or total neglect of electron-electron correlation (Tew et al., 2007) and relativistic effects (Pyykkö, 2012; Liu, 2020). Furthermore, in many cases, the neglect of inner-core electrons, the acceptance of insufficient Born-Oppenheimer approximation, and

the disregard of chemically unique tailored basis sets in favor of a one-size-fits-all approach have become routine in many quantum chemistry calculations (Woolley and Sutcliffe, 1977; Combes et al., 1981; Davidson and Feller, 1986; Cederbaum, 2008; Nagy and Jensen, 2017). In general, the larger the system with respect to the number of electrons, the cruder the approximations become. It can be argued that, in a way, the current state of many computational chemistry methodologies is as paradoxical as it was 60 years ago. For small systems, where very accurate experiments are often readily accessible, we find increasingly powerful and reliable quantum chemical computational methods and techniques being developed. However, for complex chemical systems like those typically found in biological applications where accurate experimental data requires unambiguous interpretation and quantitative predictions from theoretical models, one often finds an extensive application of low-accuracy quantum chemical methods. Only a few decades ago, modeling of biochemical systems was limited to empirical and semiempirical methods, where approximations such as the Hückel model and methods based on the neglect of diatomic differential overlap were either presented or used as reliable, despite strong evidence to the contrary obtained from studies on small systems (Eltner et al., 2000; Hagebaum-Reignier et al., 2007; Kutzelnigg, 2007; Yates, 2012; Eltner and Seifert, 2014; Thiel, 2014; Christensen et al., 2016; Dral et al., 2019).

Recently, the development of many branches of sciences has been accelerated by the use of machine learning (ML) models that contain a large number of parameters, which are weighted and tuned during the training process. As a result, ML models have had a transformative impact on the chemical sciences. For molecular applications, the design and assembly of a compelling ML model often requires a significant investment of computational resources not only for algorithm processing but also for generating accurately labeled data and ground truths for analysis and pattern inferences necessary for the training. Indeed, once attained, ML models can drastically reduce the computational time for routine tasks in molecular modeling, thereby amplifying the amount of data that can be generated for a given dataset. However, it is important to note that ML models usually fail to generate the insights necessary for explaining the electronic structure of molecules. Many of the reported works on molecular applications of ML models propose that these models are not just proxies but computational ends for theoretical molecular quantum chemistry methods. Nonetheless, many of these methods rely on correlations between families of molecules, where basic parameters such as the molecular topology, molecular local descriptors, and narrow electronic property classification of relationships are emphasized (Haghighatlari and Hachmann, 2019; Corzo et al., 2021; Fung et al., 2021; Keith et al., 2021; Unke et al., 2021). As a consequence, the use of current ML models as an alternative for *ab-initio* quantum-chemical techniques may be less relevant when elucidating molecular phenomena for which unprecedented electron correlation effects and possibly the interactions between electronic and nuclear degrees of freedom play a fundamental role, e.g., redox processes, chemical reactions, photochemical processes, organometallic catalysis, etc.

Computational practitioners working in molecular modeling should be mindful that a given model is only useful within a given range of applications; outside that range, although the theory or

data sets behind the computational model may be correct, the model might not be necessarily useful. Note that there is a tendency to use techniques that were developed for small chemical systems to study large systems, rather than to search for new methods devised specifically to deal with large systems. It follows that modeling biological molecules should rely on direct calculations or ML models trained from accurate *ab-initio* quantum chemical approximations. An understanding of the proper applicability range of ML and computational quantum chemistry models is important, since the traditional neglect of quantities such as time, temperature, entropy, spin angular momentum, and correlation, namely, those physical quantities that have been considered basic since the last century to describe complex chemical systems, can lead to erroneous predictions. One-size-fits-all approaches may present an additional complication, as the only tool available may not be suitable for every problem. With this in mind, the cognitive bias of Maslow's law of the hammer (Maslow, 1966) should be avoided, as reluctance and clemency of this may stifle the advancement of fields such as quantum chemistry, where great efforts to develop new methods and theories are still needed to achieve the deductive and inductive interpretability goals required by experimental molecular sciences.

2 Challenges in large systems

The definition of a large molecular system can be approached from two angles: the number of electrons and the size of the molecule, defined by the ensemble of atoms. For instance, Br₂ contains 70 electrons but only two atoms, making it a small system despite its electron count, while H₁₀, H₅₀, and H₆₄ (Lin et al., 2011; Li et al., 2021; Mitxelena and Piris, 2022) are larger in size yet smaller in electron count than Br₂. In this review, we adopt a definition of large molecules as a compromise between electron count and molecular size. We focus on medium-to large-sized molecules containing dozens of electrons, with potential applications in biological, biochemical, catalytic, photochemical, and technological materials.

To describe large molecular systems accurately and advance scientific and technological endeavors of quantum chemistry, one must develop and implement computational methodologies always guided by the following question: what is the largest number of atoms for which it is still reasonable and, at times, necessary to request electronic energy contributions from solutions of the Schrödinger equation? Although this question may seem simple and perhaps trivial at first, its answer requires not only an understanding of quantum chemistry models, electron-electron correlation, and scalability but also an understanding of the implications associated with the intrinsic phenomenological nature of chemical processes, error propagation, and the ideas related to chemical accuracy first introduced by S.F. Boys (Boys and Rajagopal, 1966) and later popularized by J.A. Pople (Pople, 1999).

In standard *ab-initio* quantum chemistry simulations, the computational bottleneck primarily stems from the number of basis functions used to represent electrons in the system under study, rather than the number of electrons themselves. Accurate

calculations are often obtained by employing a combination of computational approximations designed to yield the exact solution to the electronic Schrödinger equation. In this respect, the exact numerical solution for stationary states of a given Schrödinger equation can be obtained through the configuration interaction (CI) matrix-eigenvalue equation. This matrix-eigenvalue equation and its Hamiltonian representation, expressed in terms of Slater determinants and a sufficiently large orbital basis (where $l \rightarrow \infty$), defines what is commonly known as the Full CI (FCI) method, which is utilized to determine FCI energies. In practice, however, CI is an essentially intractable problem, due to its computational demands. Solutions for CI up to a n number of excitations (CI-nx) and Full CI (FCI) for $n \leq 4$ electrons are typically challenging to obtain. Therefore, in many cases, compact basis sets and computationally feasible approximations to the CI solution, such as the Coupled Cluster approach (CC), are often preferred. In general, molecular systems at their equilibrium geometry are known to exhibit rapid convergence to the Full CI solution through the use of the CC hierarchy of approximations.

Quantum chemistry, in pursuit of high accuracy, employs (CC) family of methods to construct multi-electron wavefunctions using the exponential cluster operator and a molecular orbital basis. The CC method is particularly appealing not only for its remarkable accuracy and rapid convergence, but also for its size extensivity and size consistency properties (Bartlett, 2012). In contrast to other computational methods, CC energy remains unitarily invariant with respect to the rotation within occupied and virtual spin-orbital space, respectively. Hence, it is the ideal computational method for accounting for electron correlation and making accurate determinations for medium-sized molecules. With efficient implementation and approximations correlated *ab-initio* quantum chemistry methods like CC can be performed for an upper limit of 200–500 atoms with paired electrons within orbitals with angular momentum $l \leq 2$. However, despite advances in computational hardware, routine calculations may still be limited to 30–50 atoms. Therefore, large and complex molecular systems, especially those found in biological or biochemical processes, continue to present challenges in computational chemistry. Accurately estimating the molecular electronic energy in these systems requires consideration of both the number of electrons in the atoms and the size of the system, as well as the relationships between internuclear distances and electronic density, decay of overlap integrals, the long-range nature of Coulomb forces, and the conformational flexibility of molecules, among other factors.

Thus, it is of practical importance to develop reliable computational methods that can be used for molecules approaching the system size limit, which is yet to be systematically defined. Possible approaches may include: (1) Employ reasonable approximations to matrix elements at the self-consistent-field (SCF) level. These approximations often require a distinction between the Roothaan Hartree-Fock method and the canonical nature of the SCF method (Löwdin, 1955a; Löwdin, 1955b; Nesbet, 1955). (2) Generate new techniques for integral evaluation and their applicability to new approximations at the SCF and post-SCF level. (3) Formulate post-SCF approximations for obtaining a better description of electron-electron correlation energy. (4) The development of algorithms and

techniques that can leverage the inherent parallelism in computational tasks to achieve optimal performance on modern high-performance computing (HPC) architectures.

The field of quantum chemistry and general electronic structure theory has witnessed productive research efforts in these directions, with researchers employing clever computational simplifications within new theoretical frameworks and developing new approaches that can effectively scale computations to larger problem sizes, improving the accuracy and efficiency of simulations. These developments have been well-documented in the literature (Häser, 1993; Ishikawa and Kuwata, 2012; Monari et al., 2013; Helmich and Hättig, 2014; Díaz-Tinoco et al., 2016; Gyevi-Nagy et al., 2019; Mester et al., 2019; Nagy and Kállay, 2019; Ballesteros et al., 2021; Datta and Gordon, 2021; Gyevi-Nagy et al., 2021; Szabó et al., 2021; Abyar and Novak, 2022; Paudics et al., 2022; Semidalas and Martin, 2022).

3 Scaling and parallelization of quantum chemistry computations

As both molecular system sizes and computer resources grow larger, efficient computation scaling becomes critical. On commodity computer hardware, which is the basis for the majority of available clusters and HPC resources, an effective strategy would be combining computing power of multiple units, allowing for a lower time-to-solution. This implies effective management of computational tasks in hand. Quantum chemistry calculations often involve tasks with variety of computational costs and dependencies. For example, in the SCF procedure, constructing and diagonalizing the Fock matrix depends on computing and summing up many integrals over basis functions that can vary in size and complexity. Finding the way to expose the parallelism in the SCF procedure and other quantum chemistry methods can yield improved performance on many different kinds of computers, especially modern HPC architectures. However, achieving optimal speedup is challenging, and only very few parallel implementations of quantum chemistry methods can demonstrate it. Domain decomposition and speculative parallelization are general techniques that have proven useful in identifying and designing parallel algorithms for large and complex quantum chemistry calculation tasks (Werner, 1995; González-Escribano et al., 2006; Su et al., 2007; Lipparini et al., 2013; Qiu et al., 2017; Nottoli et al., 2019; Sho and Odanaka, 2019; Jha et al., 2022; Shang et al., 2022; Fedorov and Pham, 2023).

3.1 Parallelization strategies

The main parallelization strategies in quantum chemistry computations can generally be categorized into two approaches: fine-grained and coarse-grained parallelism. Fine-grained parallelism focuses on executing numerous small, independent tasks simultaneously across multiple processing elements. This approach is well-suited for tasks with a high degree of data locality, such as dense matrix operations. The matrix and tensor operations are prevalent in quantum chemistry and thus are most often the first target of performance optimization.

Conversely, coarse-grained parallelism divides the computation into larger, independent tasks that can be executed on separate processing elements. Unlike fine-grained parallelism, this approach is more appropriate for situations with fewer, larger tasks that can be executed simultaneously, and that require infrequent communication between them to deliver results. In addition, coarse-grained parallelism is quintessential for tasks that use data that are far apart (i.e., tasks with less data locality). The coarse-grained parallelism is most often achieved through the reformulation of the underlying theory of a particular quantum chemistry method to expose independent work packets. Examples could be as simple as numerical Hessian evaluation where each nuclear displacement represent independent work packet and all the way to very elaborate CC theory reformulations exposing data locality through physical nature of the quantities to be evaluated.

Many scientific applications achieve both fine-grained and coarse-grained parallelism through shared-memory and distributed-memory parallel programming models. Shared-memory parallelism allows multiple threads of execution to access the same memory space, making it suitable for tasks with a high degree of data locality. In contrast, distributed-memory parallelism uses message-passing techniques for communication between multiple processing elements, each with its own private memory space. Distributed memory is better suited for coarse-grained parallelism, with tightly coupled tasks that communicate frequently. These parallel memory programming models are often supported in quantum chemistry codes by two application programming interfaces: the Message Passing Interface (MPI) and the Open Multi-Processing (OpenMP) application programming interface.

MPI is a standardized, portable message-passing model designed for parallel computing architectures. It is a widely-used distributed-memory parallel programming model that efficiently parallelizes tasks with less data locality, such as the distribution of integrals and the communication of partial results between processing elements. While MPI supports both point-to-point and collective communication, it is generally better suited for coarse-grained parallelism. On the other hand, OpenMP offers the capability to incrementally parallelize a serial program, unlike message-passing models like MPI which typically require an all-or-nothing approach. OpenMP can implement both coarse-grain and fine-grain parallelism. However, many chemistry codes find a hybrid OpenMP and MPI approach most appropriate as it allows for the clear treatment of the two separate levels of parallelism that are often found, coarse-grained and fine-grained, nested within each coarse subdomain.

Quantum chemistry codes often benefit from fine-grained parallelism for tasks such as complete SCF calculations, dense matrix multiplications, DMRG calculations, and other fundamental numerical operations. Graphics Processing Units (GPUs) and OpenMP are the commonly used hardware and programming model, respectively, for implementing this type of parallelism. Coarse-grained parallelism, on the other hand, may be better for partitioning entire component grids onto separate processors. Thus, calculations based on theories such as Density Functional Theory can benefit from this type of parallelization. Recently, fine-grained parallelism has been used to accelerate simulations of quantum circuits on Field Programmable Gate Arrays (FPGAs) (Moawad et al., 2022).

3.2 Load balancing

Another crucial aspect for achieving efficient parallelism in large scale supercomputers is load balancing. Load balancing aims to distribute the computational workload evenly among multiple processors or nodes to reduce idle time and communication overhead. This distribution of the workload enhances the performance, efficiency, and scalability of parallel quantum chemistry applications (Nikodem et al., 2014; Ma et al., 2023). Thus, the proper use of the load balancing technique is essential for many production codes and calculations. Load balancing becomes specially important when coarse-grained parallelism is employed. Static Load Balancing (SLB) and Dynamic Load Balancing (DLB) are two types of load balancing techniques that vary depending on whether the workload distribution is fixed or adaptive during the execution of the computation (Ali and Khan, 2012; Patil and Shedge, 2013; Nikodem et al., 2014).

The SLB approach divides the computational work evenly among processing elements before the execution. Thus, in many quantum chemistry codes, SLB is often used to distribute the number of basis functions or integrals equally across processing units. However, although SLB may be effective for many parts of a production code, it may not always lead to optimal load balancing as workloads may vary across tasks. On the other hand, the DLB approach continuously monitors the computational workload of each processing element during execution and redistributes it to balance the load across all processing elements.

On the other hand, the DLB approach involves continuously monitoring the computational workload of each processing element during execution and redistributing tasks as needed to maintain an even distribution of work. For instance, the work-stealing algorithm allows idle processors to steal tasks from busy processors, which can reduce idle time and communication overhead in parallel systems. This algorithm has been shown to be favorable for DFT calculations (Nikodem et al., 2014).

Another algorithm that has been shown to be advantageous in quantum chemistry calculations is the inspector/executor load balancing algorithm. This approach involves two phases: an inspector phase that analyzes the task dependencies and assigns them to the compute units, and an executor phase that performs the tasks in parallel. By reducing the synchronization and contention costs of the parallel system, this algorithm is applicable to any application requiring load balance where reasonable estimations of computational kernel execution times are available. Furthermore, this algorithm reduces the overhead from centralized dynamic load balancing in codes such as NWChem's Tensor Contraction Engine (TCE) (Ozog et al., 2013).

Load balancing plays a crucial role in efficient parallelization of quantum chemistry computations across multiple computing units, ensuring that no single unit bears too much demand. DLB algorithms are well-suited for fine-grained parallelism in quantum chemistry codes and are designed to adapt to changing workloads by distributing traffic based on real-time conditions. However, they can add communication overhead and slow down the system. SLB algorithms, on the other hand, use fixed rules and are better suited for coarse-grained parallelism in which larger, independent tasks are executed on separate computing units. Therefore, when implementing codes in distributed computing

systems for quantum chemistry applications, it is important to carefully consider the trade-offs between DLB and SLB algorithms.

3.3 Tensor decompositions in quantum chemistry

The use of parallelization techniques in combination with tensor decomposition methods has proven to be highly effective for computational codes (Boudehane et al., 2022). Tensor decomposition techniques, for example, have proven to be successful in solving various quantum chemistry problems, such as calculating molecular wavefunctions (Altman et al., 2021), density matrices (Hoy and Mazziotti, 2015; Khoromskaia and Khoromskij, 2015), and electronic excitation energies (Xie et al., 2009). In CC theory the electronic wave function of molecules is represented as high-dimensional tensor. Efficient manipulation of the tensors is, thus, of crucial importance. One popular technique commonly used for this purpose is tensor decompositions. Tensor decompositions enable higher-order tensors (multidimensional arrays) to be represented as a combination of lower-order tensors, such as vectors or matrices, which can simplify the complexity and dimensionality of the tensor, uncovering patterns and structures, that can facilitate efficient parallelizations. Third-order tensors, which arise naturally as the outer product of a matrix and a vector during the calculation of spectroscopic properties, offer a practical example. The polarizability tensor is a third-order tensor that describes how a molecule responds to an electric field. Representing the fact that molecules can be polarized to varying extents in different directions, the polarizability tensor is usually represented by a 3×3 tensor, denoted as α_{ij} , where $i, j = x, y, z$. Standard manipulations are often challenging to perform on this tensor. Thus, the decomposition of this tensor into vector or matrix components can often reveal critical information about the molecule's electronic structure, and facilitate efficient parallelization of codes.

One common way to decompose a third-order tensor is to use the CANDECOMP/PARAFAC (CP) decomposition, which expresses the tensor as a sum of rank-one tensors (outer products of three vectors) (Vannieuwenhoven et al., 2015). Although this decomposition has some advantages that may be exploited in quantum chemistry applications such as uniqueness, interpretability, and sparsity, it also presents some challenges such as finding the optimal number of rank-one tensors and solving the nonlinear optimization problem (Stegeman, 2006; Favier and de Almeida, 2014; Battaglino et al., 2018). Canonical polyadic decomposition (CPD) (Hitchcock, 1927; Carroll and Chang, 1970; Qiu et al., 2021), Tucker decomposition (Tucker, 1966), and tensor-train decomposition (TTD) (Oseledets, 2011) are additional tensor decomposition tools that have shown potential for computation on molecular systems. For example, the tensor-train decomposition has been used to compress the wave-function tensors in the quantum chemical calculations of large molecules. By representing the wavefunction in a compressed tensor-train format, it becomes possible to perform calculations for larger systems with reduced computational cost. Similarly, the canonical polyadic decomposition has been employed in the study of molecular properties, such as dipole moments and polarizabilities, by

compressing the high-dimensional tensors associated with these properties (Phan et al., 2013).

Nowadays, tensor libraries have been integrated into popular chemistry computational software packages, such as Linear-Scaling Dalton (LSDalton) and North West Computational Chemistry (NWChem), to enhance their performance in CC calculations. NWChem's TCE (tensor contraction engine) module enables the implementation of TCE-generated code for efficient coupled cluster calculations, including CCSD (coupled cluster singles and doubles) and other methods. LS-Dalton has Scalable Tensor Library, developed to execute tensor contractions in the CC parts of the code. There are also a number of stand alone libraries available to manipulate tensors in quantum chemistry software.

Despite the advancements made using tensor contractions and decompositions in quantum chemistry, there are still challenges to overcome. The choice of decomposition technique, as well as the determination of appropriate rank and truncation parameters, can have a significant impact on the accuracy and efficiency of the calculations. Additionally, the development of robust and efficient algorithms for tensor decompositions is an ongoing area of research (Kolda and Bader, 2009). As these challenges are addressed, tensor decompositions are expected to play an increasingly important role in the development of molecular codes.

3.4 Embarrassingly parallel quantum chemistry tasks

In the realm of quantum chemistry, some theoretical methods are considered *embarrassingly parallel* or *trivially parallelizable*. This means that tasks within a computation can be easily divided into independent sub-tasks, that require minimal to no communication or data exchange and coordination between them during execution (Foster, 1995). As a result, these tasks can be seamlessly parallelized, and executed concurrently, without significant alterations to the communication between processing units.

Embarrassingly parallel tasks often arise in the context of evaluating integrals, such as the computation of Hartree-Fock exchange integrals for large basis sets (Titov et al., 2013; Pinski and Neese, 2018), or when solving for quantities like single-point energy calculations, geometry optimizations, and molecular dynamics simulations. In these cases, calculations for each task can be performed independently, with the final result obtained by aggregating or combining individual outcomes. These frameworks allow for parallelization across multiple processing elements, including CPUs or GPUs, enabling larger and more complex problems to be tackled more efficiently.

Unfortunately, the majority of quantum chemistry methods are not easily parallelizable. Certain methods involve complex dependencies between sub-tasks or require frequent communication between processes, making them challenging to parallelize efficiently (Valiev et al., 2010). Examples include wave-function-based *ab-initio* methods, such as CI, CC, and Multi-Configuration Self-Consistent Field (MCSCF) calculations (Hasanein and Evans, 1996). These methods often involve the manipulation of large matrices and the evaluation of high-

dimensional arrays, making them computationally demanding and requiring sophisticated parallelization techniques to achieve optimal performance (Baumgartner et al., 2005).

Although frequently used quantum methods, such as Density Functional Theory (DFT) and semi-empirical methods, are more amenable to parallelization due to their relatively simpler mathematical formulations and reduced computational requirements (Andrade et al., 2015), they may not be optimal for calculations requiring a precise understanding of the electronic structure of the molecular system for elucidating the chemical phenomena at hand. Even these methods can face challenges in parallelizing specific aspects, such as evaluating long-range electron correlation or treating van der Waals interactions (Grimme et al., 2010).

The contrast between embarrassingly parallel tasks and challenging-to-parallelize quantum methods underscores the importance of developing advanced parallelization techniques and algorithms for various quantum chemistry applications. Researchers continue to explore approaches such as tensor decompositions, load balancing protocols, new hybrid MPI-OpenMP strategies, and computational hardware to overcome the limitations of current parallelization techniques for accelerating calculations (Götz et al., 2012). With ongoing developments in parallel computing, the potential to address computational challenges in quantum chemistry continues to expand, paving the way for more efficient and accurate calculations of complex molecular systems (Shao et al., 2015). Bearing this in mind, the following section provides some alternatives for calculating large molecules using CC methods.

4 Subdividing the correlation problem

In molecular orbital theory, the canonical set of molecular orbitals (MO) is typically generated through the linear combination of atom-centered basis functions, commonly Gaussian-type functions. These MOs are called canonical because they are the default MOs for a wave function obtained by diagonalizing the Fock matrix used in the SCF calculation of the HF method (Löwdin, 1955a; Löwdin, 1955b; Nesbet, 1955). Some coefficients in the linear combination can be zero or very small, which reduces computational complexity. This reduction occurs for several reasons.

- **Symmetry:** In many molecular systems, certain basis functions may not contribute to specific molecular orbitals due to symmetry constraints. The symmetry of the molecule can impose conditions on the coefficients of the linear combination, resulting in some of them being zero or negligible.
- **Negligible overlap:** When basis functions are centered on distant atoms, their overlap might be very small. In such cases, the contribution of these basis functions to the molecular orbitals can be negligible.
- **Weak interactions in excited states:** In the case of excited states, some molecular orbitals may have only weak interactions with each other. The contributions of these weakly interacting


orbitals to the overall wave function can be minimal, leading to small coefficients in the linear combination.

These factors can simplify the representation of molecular orbitals and improve the efficiency of calculations by reducing the number of significant coefficients and matrix elements to consider (Boys, 1960; Simons and Nichols, 1997; Helgaker et al., 2000; Levine et al., 2009; Szabo and Ostlund, 2012).

Despite the possible reductions in canonical MOs, they can still lead to computational inefficiencies in large molecular systems. To address this issue, the set of canonical MOs can be transformed into an equally valid set of localized HF molecular orbitals through a unitary transformation that preserves orthonormality. These localized molecular orbitals (LMOs) not only correspond to chemically familiar concepts such as core orbitals on heavy atoms, bond orbitals, and lone pair orbitals, but can also be used to reduce the computational overhead in large molecular systems (Edmiston and Ruedenberg, 1963; Kleier et al., 1974; Pipek and Mezey, 1989; Subotnik and Head-Gordon, 2005; Herbert, 2019).

Transforming canonical MOs to LMOs suggest an intrinsic approach for studying large molecules. By studying a molecule as an assembly of smaller sub-molecules rather than as a whole, the conventional up-front computational overhead can be effectively reduced.

Following this molecular splitting idea, by dividing a large molecular system into smaller fragments and characterizing each fragment with wave-functions localized on specific molecular substructures, the electron correlation contributions for each fragment can be obtained through standard computation of matrix elements, allowing for the description of the entire molecular system by combining the wave-functions of all the fragments to form the final wave-function for the complete molecule. For instance, the [*n*]helicene molecule can serve as a practical example to illustrate the partitioning of a large molecular system into smaller fragments. This molecule can be partitioned into three primary molecular fragments, one with molecular formula C_4H_4 , another with molecular formula C_6H_4 , and $n - 2$ fragments with molecular formula C_4H_2 . Consequently, the [6]Helicene ($C_{38}H_{22}$) molecule can be partitioned into one molecular segment with molecular formula C_4H_4 , another with C_6H_4 , and four with molecular formula C_4H_2 . Figure 1 illustrates this partitioning. Under this partitioning scheme, the orbitals of each of the six fragments of the [6]Helicene molecule are confined to a linear combination of the basis set, possessing atomic functions centered on the C_4H_4 , C_6H_4 , and C_4H_2 nuclei, resulting in the comprehensive and final wave-function for the system. An efficient program design that takes advantage of this partitioning would distribute the workload for each fragment, establishing the orbitals for each fragment and computing the wave-functions and expectation values of each fragment in the initial phase. In the second phase, each fragment's contribution would be added to obtain the complete description and total electronic energy of the entire molecule. This straightforward yet effective subdividing technique for computing the total energy of large systems embodies the core concept behind the Divide-Expand-Consolidate (DEC) scheme for correlated electron methods (Eriksen et al., 2015a; Ettenhuber et al., 2016; Kjærgaard et al., 2017a; Kjærgaard et al., 2017b; Barnes et al., 2019).



The DEC framework for correlated electrons involves dividing a large system into smaller fragments and obtaining the electron correlation contributions for each fragment through standard computation of matrix elements, similar to how it is done for small and intermediate systems. The following section outlines the DEC formalism for correlated wave-function methods, which enables the routine handling of molecular systems composed of over 200 atoms and more with a non-trivial number of electrons.

4.1 Localization of orbitals

A vital aspect of the DEC methodology is the localization of orbitals. Orbital localization forms the foundation for dividing large molecules into smaller fragments. Although the concept of localized orbitals is not new (Edmiston and Krauss, 1965), the literature contains a wide array of procedures aimed at localizing occupied and virtual orbitals as a means of electron-electron correlation calculations. Both wave function-based methods and fragmentation-based correlation approaches often generate localized orbitals through various localization procedures, such as those proposed by S.F. Boys (Boys, 1960) and J. Pipek and P.G. Mezey (Pipek and Mezey, 1989). However, these methods may be sensitive to the nature of molecular systems, which could limit their applicability. To address these limitations, advanced orbital localization techniques based on the central moment have been introduced (Jansík et al., 2011; Høyvik et al., 2012a; 2014). Combined with robust optimization strategies, these advanced techniques can provide more spatially local virtual orbitals even for traditionally delocalized systems (Høyvik et al., 2014; 2012a).

Localized orbitals that account for pair correlation effects have been pivotal in the development of electron-correlated methods (Löwdin, 1955a; Löwdin, 1955b; Edmiston and Ruedenberg, 1963; Davidson, 1972a; Davidson, 1972b; Pople et al., 1976; Surján, 1999). Pair natural orbitals (PNOs) are particularly useful for further compressing the virtual parameter space (Kapuy et al., 1983; Pulay, 1983; Pipek and Mezey, 1989; Saebo and Pulay, 1993; Neese et al., 2009; Yang et al., 2011; Yang et al., 2012).

Constructed from (an approximation to) the MP2 correlation density matrix for each electron pair, PNOs have gained popularity in localized orbital CC methods (Pulay and Saebo, 1986; Hättig et al., 2006; Riplinger and Neese, 2013), such as the domain-based local PNO (DLPNO) CCSD(T) method (Neese et al., 2009; Wang et al., 2013). This method has been successfully applied to various computational chemistry applications (Riplinger et al., 2013; Riplinger and Neese, 2013; Sparta et al., 2017; Sharapa et al., 2019; Stoychev et al., 2021; Wang et al., 2022). Local CC methods relying on orbital-specific virtuals (OSVs) have also been developed, closely relating to PNOs (Kurashige et al., 2012; Yang et al., 2012; Schütz et al., 2013; Tew, 2019). However, PAOs can be non-orthogonal and redundant, complicating algorithmic expressions and posing conceptual challenges for molecular systems with degenerate states arising from symmetry and angular momentum coupling (Krause and Werner, 2012). Having a set of orthogonal and non-redundant localized virtual orbitals is beneficial for CC implementations and, in many cases, necessary for obtaining unambiguous results in molecular applications. Although, localized virtual orbitals obtained using advanced localization functions are more spatially local than PAOs (Høyvik et al., 2014; 2012a), despite their advantages, they might not always be the best choice for local correlated methods. The representation and selection of the optimal set of orbitals for local correlated methods remain active areas of research and continue to evolve.

4.2 Localized orbital-based correlation methods

Orbital localization is commonly utilized to express correlation calculations in a local basis, introducing approximations that reduce the computational complexity of a method in comparison to conventional implementations. Local correlation methods can be broadly classified into two categories, wave-function-based and fragment-based approximations.

Wave-function-based approximations focus on expressing the standard wave-function using a reduced parameter set. Such approximations commonly involve constraining the virtual

excitation domain for each pair of occupied LMOs, while neglecting or approximating pair correlation contributions between well-separated occupied LMOs. For each LMO pair, a local correlation domain containing a subset of PAOs is assigned. PAOs, which form a non-orthogonal and redundant basis for the virtual orbital space, require adaptations to standard (canonical) algorithms to accommodate their unique properties. Notable examples of wave-function-based local approximations can be traced back to the work of Pulay and Sæbø (Pulay, 1983; Sæbø and Pulay, 1985; Pulay and Hamilton, 1988; Sæbø and Pulay, 1988; Sæbø and Pulay, 1993). These methods utilize occupied LMOs and PAOs for the virtual space, and have been further developed and expanded to gradients (Amos and Rice, 1989; Sæbø and Almlöf, 1989; Kirtman, 1995; Hampel and Werner, 1996; Russ and Crawford, 2004; Adler et al., 2009; Helgaker et al., 2012; Riplinger et al., 2013; Rolik et al., 2013; Menezes et al., 2016; Bistoni et al., 2017; Schwilk et al., 2017; Sitkiewicz et al., 2019; Saitow et al., 2022; Wang et al., 2022). In contrast, fragment-based approximations express the correlated method in amplitude equations and partition these amplitude equations into numerous small, typically independent, fragment calculations. Consequently, the energy is divided into fragment energy contributions, and the fragment energies are summed to yield the total energy. Examples of fragment-based approximations include the incremental CCSD(T) method (Friedrich and Dolg, 2009; Friedrich and Hänchen, 2013) and the local energy CCSD(T) method (Zhang and Grüneis, 2019; Altun et al., 2021).

Both wave function-based and fragment-based approximations have their advantages and drawbacks. Fragment-based approximations are better suited for modern multi-core architectures and have storage requirements independent of system size. In contrast, wave function-based approximations have storage requirements that grow with system size, which limits the size of the systems that can be treated. The choice of method depends on the specific application and the balance between computational cost and accuracy. As research in this area continues, further improvements and refinements of these methods are expected, enabling the treatment of larger and more complex molecular systems.

4.3 Breaking down correlation energy

In quantum chemistry theories, the concept of energy correlation, popularized by Löwdin (Löwdin, 1958), is the most prevalent perspective on the electron correlation problem. This description divides the exact total energy of a molecular system into the sum of the HF energy and a correlation contribution:

$$E_{\text{Total}} = E_{\text{HF}} + E_{\text{corr}}. \quad (1)$$

In general, for any correlated method, the relationship between the method's total energy and the correlation energy can be expressed as:

$$E_{\text{Method}} = E_{\text{HF}} + E_{\text{corr}}. \quad (2)$$

In the case of the Møller-Plesset second-order perturbation theory (MP2) method, the energy expression in Eq. 2 becomes:

$$E_{\text{MP2}} = E_{\text{HF}} + E_{\text{corr}}^{\text{MP2}}. \quad (3)$$

For molecular systems, the MP2 correlation energy $E_{\text{corr}}^{\text{MP2}}$ can be expressed as:

$$E_{\text{corr}}^{\text{MP2}} = \sum_{ij} \sum_{ab} t_{ij}^{ab} (2g_{aibj} - g_{biaj}), \quad (4)$$

where g_{aibj} are the electron repulsion integrals (ERIs) using the Mulliken notation, and t_{ij}^{ab} are the MP2 amplitudes.

Similarly, for the coupled cluster (CC) total energy (E_{CC}), the correlation energy for a closed-shell molecular system can be obtained by:

$$E_{\text{corr}}^{\text{CC}} = \sum_{ij} \sum_{ab} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}), \quad (5)$$

where t_i^a and t_{ij}^{ab} represent singles and doubles CC amplitudes. Indices i, j, \dots refer to occupied orbitals, while a, b, \dots refer to virtual orbitals. In the DEC coupled cluster (DEC-CC) framework, this correlation energy can be represented by a set of localized occupied and virtual HF MOs, which are determined and assigned to the atomic site nearest to the MO's center of charge. Thus, the correlation energy for a given correlated method, E_{corr} , becomes:

$$E_{\text{corr}} = \sum_P^{N_{\text{frag}}} \left[E_P + \sum_{Q < P}^{N_{\text{frag}}} \Delta E_{PQ} \right], \quad (6)$$

where N_{frag} is the number of atomic fragments the molecular system was divided into. For the CC method, the atomic fragment energy E_P and the pair fragment interaction energy ΔE_{PQ} are defined as follows:

$$E_P^{\text{CC}} = \sum_{ij \in \underline{P}} \sum_{ab} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}), \quad (7)$$

$$\Delta E_{PQ}^{\text{CC}} = \sum_{i \in \underline{P}} \sum_{j \in \underline{Q}} \sum_{ab} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}) + \sum_{i \in \underline{Q}} \sum_{j \in \underline{P}} \sum_{ab} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}), \quad (8)$$

where the MOs are now assumed to be local, and \underline{P} denotes the set of local occupied orbitals assigned to atomic site P .

It is important to note that the correlation energy expressions in Eqs 6–8 do not contain any approximations. Therefore, these equations, in principle, yield the same correlation energy corrections as the original expressions of the CC correlated methods. A crucial aspect of the DEC-CC approximation is dividing the calculation of the correlation energy of the entire molecular system into $N_{\text{frag}} + 1/2 \cdot N_{\text{frag}}(N_{\text{frag}} - 1)$ independent molecular fragments. Computational savings arise when screening techniques are employed for each fragment calculation. In many instances, the locality of the MOs reduces the computational effort in the molecular calculation. Moreover, the integral g_{aibj} becomes negligible when the molecular orbital ϕ_a is spatially distant from ϕ_i , allowing the summation over virtual orbitals in (7)–(8) to be limited. Consequently, only a subset of virtual orbitals, $[\bar{P}]$, is significant for each fragment from the complete set of atomic site orbitals, P . A key advantage of the DEC framework is that these summation constraints in fragment energy calculations are determined in a black-box manner, enabling the definition of atomic fragment and pair fragment interaction CC energies as.

$$E_P = \sum_{ij \in \underline{P}} \sum_{ab \in [\bar{P}]} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}), \quad (9)$$

$$\Delta E_{PQ} = \sum_{\substack{i \in \underline{P} \\ j \in \underline{Q}}} \sum_{ab \in [\bar{P}] \cup [\bar{Q}]} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}) \\ + \sum_{\substack{i \in \underline{Q} \\ j \in \underline{P}}} \sum_{ab \in [\bar{P}] \cup [\bar{Q}]} (t_{ij}^{ab} + t_i^a t_j^b) (2g_{aibj} - g_{biaj}), \quad (10)$$

where, for the pair fragment interaction energies, the set of virtual orbitals is chosen as the union of the atomic fragment spaces, which can be justified by a locality analysis of the results (Kristensen et al., 2011; Ettenhuber et al., 2016).

5 The Divide–Expand–Consolidate CCSD(T) framework

In the realm of computational quantum chemistry, the CCSD(T) model is often referred to as the *gold standard* for molecular calculations. Within the DEC framework, the CCSD(T) method is implemented as follows:

$$E^{(T)} = \sum_P^{N_{\text{frag}}} \left[E_P^{(T)} + \sum_{Q < P}^{N_{\text{frag}}} \Delta E_{PQ}^{(T)} \right] \quad (11)$$

with the corresponding equations.

$$E_P^{(T)} = 2 \sum_{ij \in \underline{P}} \sum_{ab \in [\bar{P}]} (2t_{ij}^{ab} - t_{ji}^{ab}) T_{ij}^{ab} + 2 \sum_{i \in \underline{P}} \sum_{a \in \bar{P}} t_i^a T_i^a \quad (12)$$

$$\Delta E_{PQ}^{(T)} = 2 \left(\sum_{\substack{i \in \underline{P} \\ j \in \underline{Q}}} + \sum_{\substack{i \in \underline{Q} \\ j \in \underline{P}}} \right) \sum_{ab \in [\bar{P}] \cup [\bar{Q}]} (2t_{ij}^{ab} - t_{ji}^{ab}) T_{ij}^{ab} \\ + 2 \left(\sum_{\substack{a \in \bar{P} \\ i \in \underline{Q}}} + \sum_{\substack{a \in \bar{Q} \\ i \in \underline{P}}} \right) t_i^a T_i^a. \quad (13)$$

In these expressions, the intermediate terms T_{ij}^{ab} and T_j^a can be defined as follows.

$$T_{ij}^{ab} = \sum_{cd \in [\bar{P}]} \sum_{k \in [\underline{P}]} (t_{ijk}^{acd} L_{bckd} - t_{kji}^{acd} g_{kldc}) \\ - \sum_{c \in [\bar{P}]} \sum_{kl \in [\underline{P}]} (t_{ikl}^{abc} L_{kjl c} - t_{lki}^{abc} g_{kjl c}) \quad (14a)$$

$$T_j^a = \sum_{cd \in [\bar{P}]} \sum_{kl \in [\underline{P}]} (t_{ikl}^{acd} - t_{lki}^{acd}) L_{kcl d}. \quad (14b)$$

Here, $[\underline{P}]$ represents the set of occupied orbitals assigned to atomic sites near center P , analogous to the virtual spaces. The triples amplitudes t_{ijk}^{abc} are derived from the CCSD doubles amplitudes, and $L_{aibj} = 2g_{aibj} - g_{biaj}$. This implementation of CCSD(T) necessitates an additional o^3v^4 scaling step compared to a conventional CCSD(T) implementation, as the standard CCSD(T) method cannot be easily partitioned into atomic fragment energy contributions due to the (T) corrections (Raghavachari et al., 1989). Nevertheless, with this implementation, the CCSD(T) method can be partitioned analogously to the standard CC correlation energy (Eriksen et al., 2015a).

The energy partitioning presented in Eqs 9, 10 defines what is known as the occupied partitioning scheme. However, DEC utilizes both local occupied and local virtual orbitals, and as a result, virtual and Lagrangian partitioning schemes are also formulated (Høyvik et al., 2012a). These partitioning schemes not only provide independent paths for evaluating correlation energy, maintain comparable error control, and yield reliable results, but they also exhibit distinct characteristics. The virtual and Lagrangian partitioning schemes tend to generate larger fragments in practice, yet they still enable error estimation in DEC calculations. In contrast, the Lagrangian scheme offers some advantages over the occupied and virtual schemes due to its variational nature, which leads to errors in amplitudes and multipliers being roughly proportional to the square root of the fragment optimization threshold (FOT). Additionally, the Lagrangian scheme delivers a more balanced treatment of both occupied and virtual spaces. Although virtual orbitals are generally less localized than occupied orbitals (Høyvik and Jørgensen, 2016), resulting in larger fragments within the virtual and Lagrangian schemes compared to the preferred occupied scheme, DEC calculations of molecular gradients necessitate the use of the virtual partitioning scheme. (Kristensen et al., 2012a; Bykov et al., 2016).

5.1 Atomic fragment optimization

As a continuation of the discussion on partitioning the correlation energy into atomic fragment and pair fragment interaction energies, this section focuses on optimizing the occupied and virtual orbital spaces $[\underline{P}]$ and $[\bar{P}]$ for atomic fragment P . The error associated with this optimization is dictated by the FOT used to obtain the fragment energy E_P .

The atomic fragment energy E_P in Eq. 9 is determined from the energy orbital space (EOS), $P_{\text{EOS}} \equiv \underline{P} \cup [\bar{P}]$. The EOS represents the orbital space which ensures accurate correlation energy. However, due to the coupling between the CC amplitudes, solving the CC amplitude equation in P_{EOS} is not feasible (Eriksen et al., 2015a; Ettenhuber et al., 2016). Instead, the coupling can be accounted for by solving the CC amplitude equations in an extended orbital space, the amplitude orbital space (AOS), $P_{\text{AOS}} \equiv [\underline{P}] \cup [\bar{P}]$ (Eriksen et al., 2015a). It is important to note that the occupied orbitals in the EOS ($i \in \underline{P}$) are fixed by the orbital assignment, and the virtual orbital space is identical for both EOS and AOS. Assuming $[\underline{P}]$ and $[\bar{P}]$ are known, the atomic fragment energy E_P can be calculated as follows.

1. Solve the CC amplitude equations in P_{AOS} .
2. Extract the CC amplitudes from P_{AOS} to P_{EOS} .
3. Calculate the two-electron integrals in P_{EOS} .
4. Use the CC amplitudes and integrals in P_{EOS} to calculate the atomic fragment energy as in Eq 9.

In DEC, the strategy to determine the spaces $[\underline{P}]$ and $[\bar{P}]$ that yield atomic fragment energies with FOT accuracy consists of two steps: fragment expansion followed by fragment reduction (Ettenhuber et al., 2016). In the fragment expansion step, a priority list l_r^P is generated to describe the importance of each local orbital for the fragment energy E_P ,

utilizing the distance between the center of charge of a given orbital and the atomic site P due to the locality of correlation effects (Høyvik and Jørgensen, 2016). However, alternative lists based on numerical overlap of orbitals or Fock matrix elements have also been tested with similar results (Ettenhuber et al., 2016). The process begins by selecting an initial space ($[P]$ and $[\bar{P}]$) from the priority list, calculating the fragment energy as previously described, then expanding the orbital spaces based on the priority list, and subsequently obtaining an improved fragment energy. This procedure is repeated until the difference between the last two fragment energies falls below the FOT. The fragment reduction step involves a binary search to remove orbitals without introducing errors larger than the FOT in the atomic fragment energy (Ettenhuber et al., 2016). This reduces the size of the AOS for atomic fragments, leading to significant computational savings for pair fragments.

The error control of the atomic fragment optimization comes with overhead, and improving the fragment optimization procedure is an ongoing research direction in optimizing the DEC scheme (Riplinger and Neese, 2013; Ettenhuber et al., 2016). One possible approach to reducing the overhead is to explore more efficient algorithms or heuristics that can guide the optimization process and reduce the number of calculations needed to reach the desired FOT accuracy.

Additionally, for DEC-CCSD or DEC-CCSD(T) calculations, fragment optimization can be performed at a lower level of theory, such as DEC-MP2 (Riplinger and Neese, 2013a; Ettenhuber et al., 2016). This approach can lead to considerable computational savings without significantly compromising the accuracy of the final results. However, it is crucial to validate the appropriateness of using a lower level of theory for the specific system being studied, as some systems might require higher levels of theory for accurate predictions.

The locality of electron correlation is system-dependent, and the goal of the fragment optimization procedure is to obtain a method that provides the same recovery of the correlation energy for all systems, independently of the complexity of the electronic structure (Høyvik and Jørgensen, 2016). The fragment spaces tend to be larger for systems characterized by a delocalized electronic-structure, such as graphene, than for systems containing only non-conjugated covalent bonds. In particular, for systems with a delocalized electronic-structure, it is important to use the most advanced orbital localization functions, such as the squared fourth moment localization function (Høyvik et al., 2012a), as these localization functions can generate localized sets of orbitals that are minimally system-dependent (Høyvik and Jørgensen, 2016).

Furthermore, incorporating machine learning techniques into the fragment optimization procedure may offer a promising avenue for future research (Westermayr et al., 2021). By training machine learning models on existing datasets of molecular systems, it may be possible to predict optimal fragment spaces or guide the optimization process more efficiently. This approach could potentially reduce the computational cost associated with the optimization procedure while maintaining the desired level of accuracy.

5.2 DEC amplitudes and the transformation of the basis

In the fragment energy calculations of the DEC scheme, the CC amplitude equations must be solved in the AOS. To achieve this, the

set of local orbitals is transformed into a pseudo-canonical basis by diagonalizing the local Fock matrix blocks F_{ij} ($ij \in [P]$) and F_{ab} ($ab \in [\bar{P}]$). The pseudo-canonical basis is traditionally denoted using capital letters I, J, A, B .

The CC amplitude equations are better conditioned in the pseudo-canonical basis, and the MP2 amplitudes (Kristensen et al., 2011; Høyvik et al., 2012b) and (T) intermediates (Ziółkowski et al., 2010; Eriksen et al., 2015a) can be obtained non-iteratively using standard canonical CC algorithms. After solving the amplitude equations in the AOS, the amplitudes must be transformed back to the local basis ($t_{IJ}^{AB} \rightarrow t_{ij}^{ab}$) in order to extract the EOS amplitudes and calculate the fragment energy. A similar operation is performed for the (T) intermediates (T_{ij}^{ab} and T_i^a in equation (14)).

While the pseudo-canonical basis is utilized throughout the code implementation, the transformation to the local basis is model-dependent. It is worth noting that when the energy is evaluated using the occupied partitioning scheme in Eqs. 9, 10, transforming the virtual orbitals to the local basis is not necessary. Moreover, in the case of the Laplace-transformed variation of the method DEC-RIMP2 (i.e., the DEC-LT-RIMP2 method), the amplitudes are directly obtained in the local basis (Bykov et al., 2016; Bykov and Kjaergaard, 2017).

5.3 Pair fragment calculations

The linear scaling of the DEC scheme is achieved only when the number of pair fragments scales linearly with the system size. Dispersion interactions cause the pair fragment interaction energy ΔE_{PQ} to decay following the R_{PQ}^{-6} pattern with the distance R_{PQ} between atomic sites P and Q . By employing a real-space cutoff R_{screen} , pairs with a distance exceeding the distance screening threshold (e.g., $R_{PQ} > R_{\text{screen}}$) can be screened, leading to a linear-scaling algorithm (Ziółkowski et al., 2010; Kristensen et al., 2011; Kristensen et al., 2012a; Høyvik et al., 2012b; Kristensen et al., 2012b; Jakobsen et al., 2013; Olsen et al., 2015). However, this strategy presents two challenges: (i) the calculated number of pair fragment interaction energies becomes independent of the FOT, requiring adjustments to both the FOT and R_{screen} in order to converge to the standard CC correlation energy, and (ii) the pair fragment interaction energies for a specific pair distance can span several orders of magnitude indicating that numerous pairs included within the distance screening threshold are less significant than some of the pairs excluded. To address these challenges, a pair screening strategy based on estimating the pair fragment interaction energies can be considered. This strategy is introduced by rewriting the correlation energy (Eq. 6) as a sum of effective atomic fragment energies ϵ_p ,

$$E_{\text{corr}} = \sum_p^{N_{\text{frag}}} \epsilon_p. \quad (15)$$

The effective atomic fragment energy for fragment P is the sum of the atomic fragment energy E_P and the average pair fragment interaction energy E_p^{av} . The latter term describes the interaction between the atomic site P and the other atomic sites,

$$\epsilon_p = E_P + E_p^{\text{av}}, \quad (16)$$

with

$$E_p^{av} = \frac{1}{2} \sum_{Q \neq P} \Delta E_{PQ}. \quad (17)$$

This process of partitioning a system into fragments is not exclusive to DEC. Since it is also employed in other fragment-based quantum chemistry methods (Wang et al., 2014; Fedorov, 2017).

In order to accurately determine the average pair fragment interaction energy to the FOT accuracy, a pair fragment screening technique is employed. This technique involves using minimal orbital spaces to calculate pair energy estimates at the MP2 level, based on a priority list I_p^P , such that the estimates recover a significant portion (usually 80–95%) of the *exact* MP2 pair fragment interaction energies, while requiring much less computational resources. To ensure efficiency, a conservative real-space cutoff of $R_{\text{screen}} = 30\text{\AA}$ ensures a linear-scaling number of pair energy estimates.

Once the pair energy estimates have been obtained, the screening proceeds by following this specific strategy.

1. Order all pair energy estimates associated with a N_{frag} number of atomic sites in the molecule within a given P fragment,

$$|\Delta E_{P1}^{\text{esti}}| \leq |\Delta E_{P2}^{\text{esti}}| \dots \leq |\Delta E_{PN_{\text{frag}}}^{\text{esti}}|, \quad (18)$$

2. Sum up the estimated contributions in the list, starting with the smallest values until it adds up to the FOT,

$$\max_{I_p} \left(\frac{1}{2} \sum_{Q=1}^{I_p} |\Delta E_{PQ}^{\text{esti}}| \right) \leq \text{FOT} \quad (19)$$

3. All pairs $\Delta E_{PQ}^{\text{esti}}$ in the ordered list for which $Q \leq I_p$ are then screened away and not calculated at the target CC level.
4. Repeat Steps 1-3 for all atomic sites.

This procedure, combined with fragment optimization and basis transformation, serves as the core of the DEC framework. It facilitates the breakdown of large molecular systems into smaller, more manageable fragments that can be treated independently. Consequently, it enables a series of single fragment and pair fragment calculations, which concentrate on specific sections of the system, and offers the potential to take advantage of tensor hypercontraction concepts and linear-scaling approaches.

5.4 Error estimates

As emphasized earlier, a crucial aspect of practicing computational quantum chemistry is to recognize and appreciate the accuracy limitations inherent to the different approaches that make up the computational molecular method.

In connection with the approximations found in the DEC framework, it is important to recognize the three primary assumptions that greatly impact the calculation's accuracy: (i) the fragment expansion procedure, which is presumed to converge with an error deemed negligible relative to the FOT,

(ii) the pair energy estimates and their ability to deliver a qualitative approximation of the actual pair fragment interaction energies, confirming that the relevant pair fragment interaction energies are assessed, and (iii) the errors in the ultimate pair fragment interaction energies E_{PQ} being minimal in comparison to the FOT.

The FOT is a crucial aspect of the DEC framework, as it directly governs the correlation energy error and implicitly manages errors in the correlated density matrix and molecular gradient. The DEC framework has been purposefully designed to control the error in atomic fragment energy E_p and average pair fragment interaction energy E_p^{av} according to the FOT. This is achieved by determining the AOS for each atomic fragment in a DEC calculation, ensuring precision dictated by the FOT through a black-box like algorithm. The AOS is initially defined with only the nearest neighbor atoms to P and expands layer by layer, assessing the energy contribution from each individual orbital in the AOS. It is then examined whether single orbitals can be removed from the AOS without affecting the calculation's precision, allowing the atomic fragment energies to be determined within the FOT tolerance.

DEC calculations derive properties from the sum of contributions from atomic fragment and pair fragment calculations, determined up to the FOT tolerance. An analysis of numerical experiments has revealed that the accuracy of a DEC calculation is influenced by a correlation error, which scales with $N_{\text{fragments}}$ and is dependent on the number of non-Hydrogen atoms in the system, such that

$$\delta E_{\text{corr}} \approx 2N_{\text{frag}} \cdot \text{FOT} \quad (20)$$


The error of a DEC calculation compared to a conventional calculation depends on the FOT for size-intensive properties but relies on both the FOT and system size for size-extensive properties. For example, the percentage of correlation energy recovered for a given FOT is independent of the system size, while the absolute error in the correlation energy increases linearly with the system size. Similarly, the error at a specific point in space of the correlated density or electrostatic potential depends solely on the FOT, irrespective of the system size.

On the other hand, the error control of the DEC method has been validated through theoretical analyses and numerical results for small and medium-sized molecules. For larger molecules, internal consistency checks can be used to estimate errors in the calculated DEC correlation energy and electron density, revealing that DEC errors for size-intensive properties remain consistent across different system sizes, while errors for size-extensive properties increase with system size.

DEC calculations are capable of recovering over 99% of the correlation energy if a sufficiently stringent FOT is employed, **Table 1**.

Although FOT can assure a reduction in error compared to the method's conventional counterpart, it does not guarantee the validity of the calculation's accuracy for the system being studied. This point is related to a minor alteration of the earlier reflective question:

What is the maximum number of atoms for which electronic energy calculations and quantum mechanical methods remain meaningful and do not lose their significance?



5.5 The Divide–Expand–Consolidate parallelization

The DEC algorithm's attractiveness for high-performance computing is due in part to its core procedures: fragment optimization, local energy calculation, and pair fragment screening. These processes enable the DEC framework to divide large molecular systems into independent fragments and taking advantage of tensor hypercontraction concepts (like resolution-of-identity) and other linear-scaling procedures and optimization on the fragment level. The method has been applied to study sizable systems, such as supramolecular wires with up to 40 monomers of 1-aza-adamantane-trione (AAT) molecules, encompassing 2,440 atoms and 24,440 basis functions (Kjærgaard et al., 2017a; Kjærgaard et al., 2017b).

Three levels of parallelization using a multi-threaded OpenMP and MPI implementation exist within the DEC algorithm: coarse, medium, and fine-grained. Coarse-grained parallelization leverages the independent nature of fragment energy calculations, allowing them to be executed in parallel on separate computing units (usually a user defined team of compute nodes). Meanwhile, medium and fine-grained parallelization focus on distributing the solution of the CC amplitude equations for each molecular fragment across multiple compute nodes within computing unit and individual threads on a particular node.


At the coarse-grained level, available compute units are divided into groups under a global work manager that directs the DEC calculation through a set of local managers. These local managers each command a group of workers to execute individual fragment calculations. The global manager generates an ordered job list based on the workload of each job, which represents an atomic or pair fragment interaction energy calculation. The global manager communicates the required information for each fragment energy calculation to local groups, starting with the largest fragments. As local groups become available, the global manager dynamically distributes the remaining fragment calculations in the job list. The size of the local group is also dynamic, splitting if the workload is insufficient to span across the local group, with two new jobs

assigned to the new groups. This combination of dynamic job distribution and dynamic group size adjustment minimizes global and local communication losses, ensuring optimal performance on large parallel machines. Finally, the global manager calculates the correlation energy based on Eq 6.

The DEC scheme has been used on large compute clusters and supercomputers like Titan (Kjærgaard et al., 2017a; Kjærgaard et al., 2017b) and Summit (Luo et al., 2020) at Oak Ridge National Laboratory, as well as on GPUs (Bykov and Kjærgaard, 2017). Several methods have been implemented using the DEC strategy, including DEC resolution-of-identity MP2 (Bykov and Kjærgaard, 2017) with gradients (Bykov et al., 2016), densities (Kristensen et al., 2012a), Laplace-transformed (Kjærgaard, 2017), and Explicitly-correlated (Wang et al., 2016) excitation energies using local framework (LoFEx) (Baudin and Kristensen, 2016), CCSD (Baudin et al., 2017); CCSD(T) (Eriksen et al., 2015a); and multi-layer DEC (Barnes et al., 2019).

5.6 Demonstrative calculations

To demonstrate the capability of the DEC methodology for incorporating correlated techniques, we applied it to simulate 11 biologically and pharmacologically significant molecules at the MP2 level, utilizing the resolution of identity (RI) approximation in conjunction with the cc-pVDZ basis set (Dunning Jr, 1989). The optimized molecular geometries were obtained with the Berny algorithm available in the Gaussian 16 suite of programs (Frisch et al., 2021). Initial guess geometries for codeine, remdesivir, ampicillin, tetrahydrocannabinol (THC), and fentanyl were computed in the gas phase using Hartree-Fock (HF) and the 6-311++G (d,p) basis set (Krishnan et al., 1980; Frisch et al., 1984). These structures were subsequently optimized in water and pentylethanoate with the Polarizable Continuum Model (PCM) to simulate the aqueous and lipidic effects within a biological cellular system. Initial structures for protonated deoxyribonucleic acid (DNA) fragments were optimized with HF/STO-3G (Hehre et al., 1969; Collins et al., 1976) in both gas and

**TABLE 1** Percentage of correlation recovered as a function of the FOT.

FOT	Δ_{DEC} (%)
10^{-3}	98.2
10^{-4}	99.8
10^{-5}	99.985
10^{-6}	99.998

TABLE 2 DEC-RIMP2 total energy for 11 biologically relevant molecules, all values are in a.u.

Molecule	Hartree-Fock energy	RIMP2 total energy
Codeine	-972.862701	-975.992852
Remdesivir	-2,309.345791	-2,315.468663
Ampicillin	-1477.391759	-1480.860488
THC	-962.501681	-965.745358
Fentanyl	-1033.262311	-1036.726093
B-DNA (2)	-2,901.318601	-2,907.974428
B-DNA (4)	-5650.439019	-5663.251418
B-DNA (8)	-11148.833671	-11174.064299
B-DNA (10)	-13882.030042	-13913.440352
B-DNA (14)	-19348.404412	-19392.241663
FDNA	-32403.080000	-32480.965782

aqueous phases. These initial geometries were later refined using the BP86 functional with Orca (Neese et al., 2020) and the cc-pVDZ basis set. Figure 2 illustrates the molecular structures of five

TABLE 3 DEC-CCSD total energies for three representative molecules, all values in a.u.

Molecule	Hartree-Fock energy	CCSD total energy
Codeine	-972.862700	-976.107748
Fentanyl	-1033.262311	-1036.877810
THC	-962.501681	-965.914124

protonated canonical right-handed DNA helix (B-DNA) fragments with two (B-DNA (2)), four (B-DNA (4)), eight (B-DNA (8)), and ten (B-DNA (10)) nucleic acids, as well as five common drugs (pain relievers, an antiviral, an antibiotic, and a cannabinoid). Additionally, the molecular structure reported by Q. Zhang et al. for a protonated Cyclic Polyamide-DNA Complex, FDNA (Zhang et al., 2004), was considered in this study and computed with the cc-pVDZ basis set (see Figure 3).

DEC framework offers a significant advantage in its adaptability to established techniques for accelerating molecular calculations. One such technique is RI for reducing four-index integrals, which offers substantial computational savings for DEC calculations. The RI approximation can be straightforwardly applied on the fragment level for any CC level of excitation.

Table 2 displays the HF and RIMP2 total energy values for the 11 computed molecules. The results demonstrate the accuracy of the DEC framework for modeling large and complex systems, as well as its effectiveness in reducing computational cost. Table 3 presents the CCSD total energies for three representative molecules. These results further showcase the adaptability of the DEC framework to more advanced correlated methods and its ability to provide accurate results for large and complex systems. The flexible nature of the DEC framework makes it amenable to the integration of various correlated methods for the computation of total energies and molecular properties.

The DEC approximation is particularly useful if time-to-solution is to be optimized. Due to embarrassingly parallel nature of the DEC the time-to-solution can be brought down practically to any number provided a computing resources are available. Obviously, this comes at the cost of much higher count of floating point operations (FLOPs). Thus for small and medium size systems it would be more beneficial to use canonical implementation or approximations able to be optimized for the optimal FLOPs count.

The DEC framework presents a flexible and adaptable strategy for modeling intricate and large molecular systems without compromising precision and computational efficiency. By leveraging well established techniques such as the RI approximation and integrating advanced correlated methods, the DEC framework empowers computational chemistry in HPC with a powerful tool. Although, the DEC framework provides a versatile and robust black-box like approach to addressing the challenges of contemporary molecular modeling, there are still several obstacles that must be overcome to establish it as an essential tool in the realm of computational chemistry.

5.7 Challenges and opportunities

Although the DEC framework combined with electron correlation methods may appear attractive for computing large

systems, practitioners in computational chemistry should be mindful of the applicability and limitations of each method. It is crucial to remember that a given model is only useful within a specific range of applications, and its validity may not extend to all types of molecular systems or sizes. In this regard, the DEC framework with MP2 and CC methods presents a vulnerability; while MP2 and CC methods have demonstrated their accuracy and effectiveness for small to medium molecular systems, their systematic validation for large systems remains to be conducted.

Moreover, many atomic and molecular processes necessary for experimental applications are described by measurements inferred from energy differences. A single total energy of a molecular system is not chemically relevant unless the energies for the initial and final states, accounting for the change in energy during the chemical process, are available. Additionally, it has been established that an accuracy of 1 kcal/mol is necessary for quantum mechanical simulations to make reliable predictions for thermochemical properties, chemical kinetics, reactivity profiling, and chemical transformations. The use of DEC or other fragmentation approaches for molecular fragments can result in the recovery of correlation energy exceeding 1 kcal/mol for large molecules. This raises concerns about the validity of the electronic energy error propagation in large systems. When the fragments computed have a contribution to the total energy larger than the accepted error of 1 kcal/mol, it becomes challenging to discuss the accuracy of the correlated total energy.

Moreover, for large molecular systems, electron-electron effects and electronic energy become less important due to the contribution of vibrational effects. Therefore, practitioners should consider the largest number of atoms for which it is still reasonable to compute solutions of the Schrödinger equation. Although a systematic evaluation of the upper limit size for quantum chemistry methods is yet to be conducted, it is well-known that for large molecular systems, the transition from a quantum mechanical domain to a statistical mechanical one occurs. Depending on the system size, one should evaluate whether to use approaches from quantum mechanics, statistical mechanics, fluid mechanics, or thermodynamics.

Nevertheless, correlated electron methods such as CI, CC, and DEC CC can play a crucial role in the development and training of ML and AI chemical applications, where electronic structure elucidation becomes secondary and pragmatic insights are paramount for task prioritization. The DEC framework, in particular, could be a natural choice for ML chemistry models with architectures based on diffusion, active learning, generative adversarial models, normalizing flows, and transformers.

For large systems, properties depending on a correct description of electronic behavior are often localized to a specific site in the molecule or molecular system. Accounting for all electron-electron correlations may be computationally wasteful due to the small contributions from electrons far from the molecular site where the electronic process occurs. Furthermore, accessing many final states needed to describe chemical transformations can be challenging due to variational collapse, symmetry breaking, spin contamination, and the consideration of multiple configurations. Thus, direct methods for computing localized electron-electron properties should be a more effective way to tackle electron-driven phenomena in large molecules. Quantum chemistry methods that directly compute the desired property of the

molecule, such as those based on Green's functions, equation of motion and response theories, are particularly important for large molecules. In the following section, the application of one such method in the framework of the CC theory will be explored.

6 Cluster perturbation theory

An alternative method to the DEC-CC model and other localized *ab-initio* approaches that can be efficiently adapted for use on modern supercomputers to enable CC quality calculations on large molecular systems is the recently developed Cluster Perturbation theory (CP) (Baudin et al., 2019; Pawłowski et al., 2019a; Pawłowski et al., 2019b; Pawłowski et al., 2019c; Pawłowski et al., 2019d; Hillers-Bendtsen et al., 2022; Høyer et al., 2022; Olsen et al., 2022). CP theory offers a hybrid approach that combines CC theory and the Møller-Plesset (Møller and Plesset, 1934) partitioning of the wave-function. This unique combination helps overcome the limitations of the individual models and offers a more comprehensive and accurate approach to molecular calculations (Bartlett and Silver, 1974; Binkley and Pople, 1975; Pople et al., 1976; Bartlett and Shavitt, 1977; Krishnan and Pople, 1978; Raghavachari et al., 1989; Eriksen et al., 2014; 2015b; Eriksen et al., 2015c; Kristensen et al., 2016).

CP theory introduces a new class of perturbation models that rely on a correlated zeroth-order state, which can be truncated at any excitation level, thereby avoiding high excitation levels that have a negligible effect on molecular properties. Perturbation series for both energy and molecular properties, including excitation energies, can be determined by calculating small perturbation corrections, which are obtained by taking the difference between the energy and molecular property for the CC parent and CC target states.

Compared to coupled cluster perturbation theory (CCPT) energy and CCPT Lagrangian series, CP theory exhibits a faster convergence and yields superior results. It allows for the calculation of perturbation series for both ground-state energy and excitation energies, treating the CC parent state Jacobian as a zeroth-order contribution. Excitation energies in CP theory can be determined using either response function theory or equation-of-motion coupled cluster (EOM-CC) theory.

CP theory's unique pathway that connects the determination of the CC parent state and the CC target state makes it possible to determine CP series for both energy and molecular properties. Moreover, its incorporation of only one additional excitation level typically results in a more robust perturbation series and, hence, better convergence. Furthermore, non-iterative and easily parallelizable formulations can be derived within CP theory and some of the targeted model is a higher-level CC model rather than a full configuration interaction (FCI) solution, the derived corrections remain reasonably small.

A brief introduction to CP theory for electronic structure calculations, specifically for excitation energies, is presented in the following section. Additionally, the development of massively parallel implementations of CP theory and some numerical illustrations are discussed. For a comprehensive understanding of CP theory, interested readers may refer to (Pawłowski et al., 2019a,b; Baudin et al., 2019; Pawłowski et al., 2019c,d; Høyer et al., 2022; Olsen et al., 2022; Hillers-Bendtsen et al., 2022).

6.1 Standard CC theory for excitation energies

In the standard CC theory, the coupled cluster (CC) wave function (Čížek, 1966) is written as

$$|CC\rangle = e^T|HF\rangle, \quad (21)$$

where $|HF\rangle$ is the Hartree-Fock reference state. Furthermore, in the above expression the cluster operator reads

$$T = \sum_i \sum_{\mu_i} t_{\mu_i} \theta_{\mu_i}, \quad (22)$$

where the cluster amplitude is denoted t_{μ_i} , and it is associated with the many-body excitation operator θ_{μ_i} which, when applied to the Hartree-Fock reference state, generates an excited determinant: $|\mu_i\rangle = \theta_{\mu_i}|HF\rangle$. Thus, the energy may be determined as

$$E_0 = \langle HF|e^{-T}H_0e^T|HF\rangle = \langle HF|H_0^T|HF\rangle. \quad (23)$$

In this energy expression, the electronic Schrödinger equation has been multiplied from the left with e^{-T} and then projected against the Hartree-Fock state. Consequently, the corresponding amplitude equations are given as

$$\langle \mu_i|e^{-T}H_0e^T|HF\rangle = \langle \mu_i|H_0^T|HF\rangle = 0. \quad (24)$$

The similarity transformed Hamiltonian has been introduced,

$$H_0^T = e^{-T}H_0e^T. \quad (25)$$

Based on linear response theory, the CC excitation energies may be obtained as eigenvalues of the CC Jacobian, J (Harris, 1977; Dalgaard and Monkhorst, 1983), such that,

$$J\mathbf{R}_x = \omega_x\mathbf{R}_x, \quad (26)$$

$$\mathbf{L}_xJ = \mathbf{L}_x\omega_x, \quad (27)$$

$$\mathbf{L}_x\mathbf{R}_y = \delta_{xy}, \quad (28)$$

The CC Jacobian,

$$J_{\mu_i\nu_j} = \langle \mu_i|[H_0^T, \theta_{\nu_j}]|HF\rangle, \quad (29)$$

is non-Hermitian and therefore the left and right eigenvectors, \mathbf{L}_x and \mathbf{R}_x (which represent the excited-state wave function) are not Hermitian conjugates of each other. The Jacobian formulation of CP theory defines a numerical problem that can be solved using a block Davidson eigensolver, where the eigenvalues are determined by an iterative procedure starting from the lowest eigenvalue.

6.2 Cluster perturbation theory for excitation energies

For excitation energies, the CC target excitation space, $1 \leq i \leq t$, is divided into two, the parent excitation space and the auxiliary excitation space. The parent excitation space comprises the excitations from $1 \leq i \leq p$, whereas the auxiliary space includes the excitations from $p < i \leq t$.

The CC parent state wave-function is defined as¹

$$|CC^*\rangle = e^{*T}|HF\rangle. \quad (30)$$

This wave-function is defined by a truncated cluster operator which covers the excitations in the parent space.

$$*T = \sum_{i=1}^p *T_i, \quad (31)$$

$$*T_i = \sum_{\mu_i} *t_{\mu_i} \theta_{\mu_i}. \quad (32)$$

Then, the corresponding energy and amplitude equations for the parent state read

$$*E_0 = \langle HF|H_0^{*T}|HF\rangle, \quad (33)$$

$$\langle \mu_i|H_0^{*T}|HF\rangle = 0, \quad 1 \leq i \leq p. \quad (34)$$

Subsequently, the excitation energies are obtained from the eigenvalue problem involving the Jacobian of the parent space which is denoted by the left superscript P ,

$${}^P J \mathbf{R}^x = * \omega_x \mathbf{R}^x, \quad (35)$$

$$*\mathbf{L}^{xP} J = *\mathbf{L}^x \omega_x, \quad (36)$$

$${}^P J_{\mu_i\nu_j} = \langle \mu_i|[H_0^{*T}, \theta_{\nu_j}]|HF\rangle, \quad 1 \leq i, j \leq p. \quad (37)$$

Now, let's assume that the left and right eigenvectors are orthonormal,

$$*\mathbf{L}^x \mathbf{R}^y = \delta_{xy}. \quad (38)$$

Then, the CC target wave function can be parameterized according to

$$|CC\rangle = e^T|HF\rangle = e^{*T+\delta T}|HF\rangle = e^{\delta T}|CC^*\rangle, \quad (39)$$

where the cluster operator is split into the parent-space component and the correction that arises due to the presence of the auxiliary space,

$$T = *T + \delta T, \quad (40)$$

with the correction term affecting both the parent and the auxiliary space,

$$\delta T = \sum_{i=1}^t \sum_{\mu_i} \delta t_{\mu_i} \theta_{\mu_i}. \quad (41)$$

The target state amplitudes are split in an analogous manner,

$$t_{\mu_i} = *t_{\mu_i} + \delta t_{\mu_i}, \quad 1 \leq i \leq t, \quad (42)$$

where the parent space amplitudes have a vanishing auxiliary excitation space component

$$*t_{\mu_i} = 0, \quad p < i \leq t. \quad (43)$$

Thus, the amplitude equations become

$$\Omega_{\mu_i} = \langle \mu_i|e^{-\delta T}H_0^{*T}e^{\delta T}|HF\rangle = 0, \quad 1 \leq i \leq t. \quad (44)$$

Finally, the expression for the Jacobian of the target state reads

$$J_{\mu_i\nu_j} = \langle \mu_i|[e^{-\delta T}H_0^{*T}e^{\delta T}, \theta_{\nu_j}]|HF\rangle, \quad 1 \leq i, j \leq t. \quad (45)$$

The corrections to the excitation energy of the CC parent state are derived by an expansion of the different components that enter the CC target state eigenvalue problem in Eq. 2) in orders of the similarity-transformed fluctuation potential, Φ^{*T} .

¹ The asterisk in the equation should not be confused with complex conjugation.

To this end, the ground-state parent-state cluster amplitudes are initially expanded in the orders of Φ^{*T} ,

$$t_{\mu_i} = t_{\mu_i}^{(0)} + \delta t_{\mu_i}^{(1)} + \delta t_{\mu_i}^{(2)} + \dots, \quad 1 \leq i \leq t \quad (46)$$

with

$$t_{\mu_i}^{(0)} = {}^*t_{\mu_i}. \quad (47)$$

An analogous expansion is applied to the Jacobian,

$$J_{\mu_i\nu_j} = J_{\mu_i\nu_j}^{(0)} + J_{\mu_i\nu_j}^{(1)} + J_{\mu_i\nu_j}^{(2)} + \dots, \quad 1 \leq i, j \leq t, \quad (48)$$

to the right excited-state vector

$$R^x = R^{x(0)} + R^{x(1)} + R^{x(2)} + \dots \quad (49)$$

with

$$R^{x(0)} = {}^*R^x, \quad (50)$$

and the corresponding excitation energy

$$\omega_x = \omega_x^{(0)} + \omega_x^{(1)} + \omega_x^{(2)} + \dots \quad (51)$$

with

$$\omega^{x(0)} = {}^*\omega^x. \quad (52)$$

One of the key concepts of the CP theory is that the zeroth-order Jacobian, $J^{(0)}$, in Eq. 48 does incorporate a term that is of first order in Φ^{*T} . (This is similar to the “pseudo-perturbation” concept introduced, for example, in the derivation of the non-iterative second-order correction to the random-phase approximation excitation energies. (Christiansen et al., 1998) By this particular choice, the zeroth-order Jacobian, $J^{(0)}$, is guaranteed to correspond to the Jacobian of the parent-state model. To see this, consider the CC extended-parent-state Jacobian,

$${}^*J_{\mu_i\nu_j} = \langle \mu_i | [H_0^{*T}, \theta_{\nu_j}] | \text{HF} \rangle \quad i, \quad j = 1, 2, \dots, t, \quad (53)$$

Notice that ${}^*J_{\mu_i\nu_j}$ has the same structure as the Jacobian of the parent space, ${}^PJ_{\mu_i\nu_j}$ in Eq. 37, but the difference is that the indices i and j run for ${}^*J_{\mu_i\nu_j}$ over the entire target space rather than just the parent space (hence the word “extended” in the name of ${}^*J_{\mu_i\nu_j}$). The extended-parent-state Jacobian may now be split into the zeroth- and first-order terms,

$${}^*J = J^{(0)} + J^{(1)}, \quad (54)$$

where the zeroth-order part does contain a Φ^{*T} term.

$$J^{(0)} = \begin{pmatrix} \langle \mu_P | [H_0^{*T}, \theta_{\nu_P}] | \text{HF} \rangle & 0 \\ 0 & \langle \mu_A | [J^{*T}, \theta_{\nu_A}] | \text{HF} \rangle \end{pmatrix} \quad (55a)$$

$$J^{(1)} = \begin{pmatrix} 0 & \langle \mu_P | [\Phi^{*T}, \theta_{\nu_A}] | \text{HF} \rangle \\ \langle \mu_A | [\Phi^{*T}, \theta_{\nu_P}] | \text{HF} \rangle & \langle \mu_A | [\Phi^{*T}, \theta_{\nu_A}] | \text{HF} \rangle \end{pmatrix}, \quad (55b)$$

where the P and A subscripts denote the parent- and auxiliary-space components, respectively.

We have illustrated (with the excitation-energy example) the conceptual foundations of the CP theory, and showed how these foundations guarantee that the zeroth order wave function and its properties are those of the parent excitation space. At infinite order, the target state and its properties are formally recovered.

In practice, the target-state quality is usually recovered at third order (Baudin et al., 2019). The excitation energy corrections through third order become:

$$\omega_x^{(1)} = 0, \quad (56)$$

$$\omega_x^{(2)} = \sum_{q=p+1}^t \langle {}^*L_x | [[\Phi^{*T}, \delta T_q^{(1)}], {}^*R_x] | \text{HF} \rangle + \langle {}^*L_x | [\Phi^{*T}, R_x^{(1)}] | \text{HF} \rangle, \quad (57)$$

$$\begin{aligned} \omega_x^{(3)} = & \langle {}^*L_x | [[\Phi^{*T}, \delta T_q^{(2)}], {}^*R_x] | \text{HF} \rangle \\ & + \frac{1}{2} \sum_{q,r=p+1}^t \langle {}^*L_x | [[[\Phi^{*T}, \delta T_q^{(1)}], \delta T_r^{(1)}], {}^*R_x] | \text{HF} \rangle \\ & + \sum_{q=p+1}^t \langle {}^*L_x | [[\Phi^{*T}, \delta T_q^{(1)}], R_x^{(1)}] | \text{HF} \rangle \\ & + \langle {}^*L_x | [\Phi^{*T}, R_x^{(2)}] | \text{HF} \rangle. \end{aligned} \quad (58)$$

The quantities needed to determine excitation energies through third order are collected in Table 4.

6.3 A strategy for massively parallel implementations of cluster perturbation theory in singles and doubles excitation space

As discussed in (Baudin et al., 2019), CP theory excitation energies have been explicitly derived up to the third order for a CC singles parent excitation space and a doubles auxiliary excitation space in the CPS(D-3) model. A significant innovation over traditional methods is that the calculations of $\omega_x^{(3)}$ are non-iterative in the doubles excitation space. The same, of course, pertains to $\omega_x^{(2)}$, which turns out to be identical to the well-known CIS(D) model (Head-Gordon et al., 1994). Furthermore, for computing excitation energies, the CP framework allows contributions to $\omega_x^{(2)}$ and $\omega_x^{(3)}$ from batches over a virtual molecular orbital index to be determined independently and subsequently summed, which greatly enhances the efficiency and parallelizability of the method. In the following sections, the most crucial aspects of the CP theory implementation for the excitation energies are discussed. For a more extensive and thorough insight into the algorithm, including performance, wall times, computational cost, and various numerical tests, the reader is encouraged to consult (Baudin et al., 2019; Hillers-Bendtsen et al., 2023).

6.4 The singlet biorthonormal basis

Efficient parallelization of the third-order excitation energy correction, $\omega^{(3)}$, in the CP method requires the use of singlet biorthonormal basis working equations. By expressing the excitation energy corrections to the coupled cluster singles (CCS) excitation energy in singlet biorthonormal basis, the working equations for the second-order correction of the CPS(D-3) excitation energy model (Baudin et al., 2019) can be obtained as

TABLE 4 The cluster amplitudes, the Jacobian and its right eigenvector through third order. ε_{μ_i} denotes the orbital energy difference between orbitals that differ in the $|\text{HF}\rangle$ and $|\mu_i\rangle$ determinants. S_{ip} is a step function, which vanishes for $i \leq p$ and equals 1 otherwise.

$\delta t_{\mu_i}^{(1)}$	$= 0$	$1 \leq i \leq p$
$\varepsilon_{\mu_i} \delta t_{\mu_i}^{(1)}$	$= -\langle \mu_i \Phi^{*T} \text{HF} \rangle$	$p < i \leq t$
$\sum_{j=1}^p \sum_{\nu_j} J_{\mu_i \nu_j}^p \delta t_{\nu_j}^{(2)}$	$= -\sum_{j=p+1}^t \sum_{\nu_j} \langle \mu_i [\Phi^{*T}, \theta_{\nu_j}] \text{HF} \rangle \delta t_{\nu_j}^{(1)}$	$1 \leq i \leq p$
$\varepsilon_{\mu_i} \delta t_{\mu_i}^{(2)}$	$= -\sum_{j=p+1}^t \sum_{\nu_j} \langle \mu_i [\Phi^{*T}, \theta_{\nu_j}] \text{HF} \rangle \delta t_{\nu_j}^{(1)}$	$p < i \leq t$
$J_{\mu_i \nu_j}^{(0)}$	$= \langle \mu_i [H_0^{*T}, \theta_{\nu_j}] \text{HF} \rangle (1 - S_{ip})(1 - S_{jp}) + \varepsilon_{\nu_j} \delta_{\mu_i \nu_j} S_{ip} S_{jp}$	$1 \leq i, j \leq t$
$J_{\mu_i \nu_j}^{(1)}$	$= \langle \mu_i [\Phi^{*T}, \theta_{\nu_j}] \text{HF} \rangle (1 - S_{ip}) S_{jp}$ $+ \langle \mu_i [\Phi^{*T}, \theta_{\nu_j}] \text{HF} \rangle S_{ip} (1 - S_{jp})$ $+ \langle \mu_i [\Phi^{*T}, \theta_{\nu_j}] \text{HF} \rangle S_{ip} S_{jp}$	$1 \leq i, j \leq t$
$J_{\mu_i \nu_j}^{(2)}$	$= \sum_{q=p+1}^t \langle \mu_i [[\Phi^{*T}, \delta T_q^{(1)}], \theta_{\nu_j}] \text{HF} \rangle$	$1 \leq i, j \leq t$
$J_{\mu_i \nu_j}^{(3)}$	$= \sum_{q=1}^t \langle \mu_i [[\Phi^{*T}, \delta T_q^{(2)}], \theta_{\nu_j}] \text{HF} \rangle$ $+ \frac{1}{2} \sum_{q,r=p+1}^t \langle \mu_i [[[\Phi^{*T}, \delta T_q^{(1)}], \delta T_r^{(1)}], \theta_{\nu_j}] \text{HF} \rangle$	$1 \leq i, j \leq t$
$R_{\nu_j}^{x(1)}$	$= 0$	$1 \leq i \leq p$
$(\varepsilon_{\mu_i} - \omega_x) R_{\nu_j}^{x(1)}$	$= -\langle \mu_i [\Phi^{*T}, R^x] \text{HF} \rangle$	$p < i \leq t$
$\sum_{j=1}^p \sum_{\nu_j} (J_{\mu_i \nu_j}^p - \omega_x \delta_{\mu_i \nu_j}) R_{\nu_j}^{x(2)}$	$= \omega_x^{(2)*} R^x - \sum_{q=p+1}^t \langle \mu_i [[\Phi^{*T}, \delta T_q^{(1)}], R^x] \text{HF} \rangle$ $-\langle \mu_i [\Phi^{*T}, R^{x(1)}] \text{HF} \rangle$	$1 \leq i \leq p$
$(\varepsilon_{\mu_i} - \omega_x) R_{\nu_j}^{x(2)}$	$= -(\sum_{q=p+1}^t \langle \mu_i [[\Phi^{*T}, \delta T_q^{(1)}], R^x] \text{HF} \rangle$ $+ \langle \mu_i [\Phi^{*T}, R^{x(1)}] \text{HF} \rangle)$	$p < i \leq t$

$$\omega_x^{(2)} = -\sum_{ai} L_{ai}^{\text{CCS}} \left[\sum_{jkcd} R_{di}^{\text{CCS}t(1)} L_{ckjd} + \sum_{bjdl} R_{dl}^{\text{CCS}t(1)} L_{bjld} - \sum_{bjcl} R_{al}^{\text{CCS}t(1)} L_{bjlc} + \sum_{bck} \tilde{R}_{bick}^{(1)} (kc|ab) - \sum_{jck} \tilde{R}_{ajck}^{(1)} (kc|ji) \right], \quad (59)$$

where $L_{pqrs} = 2(pq|rs) - (ps|rq)$, and the barred integrals may be represented by

$$(pq|rs) = \hat{P}_{qs}^{PT} \sum_{\alpha\beta\gamma\delta} (\bar{X}_{\alpha p} C_{\beta q} + C_{\alpha p} \bar{Y}_{\beta q}) C_{\gamma r} C_{\delta s} (\alpha\beta|\gamma\delta), \quad (60)$$

with

$$\bar{X}_{ai} = 0, \quad \bar{X}_{aa} = -\sum_i C_{ai} R_{ai}^{\text{CCS}}, \quad (61a)$$

$$\bar{Y}_{ai} = \sum_a C_{aa} R_{ai}^{\text{CCS}}, \quad \bar{Y}_{aa} = 0, \quad (61b)$$

where the i,j,k,l and a,b,c,d indices denote molecular orbitals (MOs) that are, respectively, occupied and unoccupied in the reference Hartree-Fock state. The Greek indices pertain to the atomic-orbital (AO) basis and $C_{\alpha p}$ are AO-to-MO transformation coefficients. Adopting the conventional approach for computing the tensor elements (Helgaker et al., 2000), the first-order cluster amplitudes read

$$\tilde{t}_{aibj}^{(1)} = \frac{2(ai|jb) - (aj|ib)}{\varepsilon_i - \varepsilon_a + \varepsilon_j - \varepsilon_b}, \quad (62)$$

whereas the first-order right eigenvector is symmetrized and written in the singlet basis as

$$\tilde{R}_{aibj}^{(1)} = \frac{2(ai|jb) - (aj|ib)}{\varepsilon_i - \varepsilon_a + \varepsilon_j - \varepsilon_b + \omega^{\text{CCS}}}. \quad (63)$$

The third-order excitation energy correction may be formulated via

$$\omega_x^{(3)} = \sum_{aibj} \tilde{t}_{aibj}^{(2)} \left[\tilde{F}_{ij}^{ab} R_{ai}^{\text{CCS}} \bar{F}_{jb} - (ia|jb) \right] + \sum_{aibj} \left[\tilde{R}_{aibj}^{(2)} + 2R_{ai}^{\text{CCS}t(2)} - R_{aj}^{\text{CCS}t(2)} \right] (ai|bj), \quad (64)$$

where the expression involving the two-electron integral is defined by

$$(ia|jb) = \hat{P}_{ab}^{ij} \sum_{\alpha\beta\gamma\delta} (\bar{X}_{ai} C_{\beta a} + C_{ai} \bar{Y}_{\beta a}) C_{\gamma j} C_{\delta b} (\alpha\beta|\gamma\delta), \quad (65)$$

with

$$\bar{X}_{ai} = \sum_j C_{aj} \sum_b L_{bi}^{\text{CCS}} R_{bj}^{\text{CCS}} = \sum_j C_{aj} D_{ij}^{\text{CCS}} \quad (66)$$

$$\bar{Y}_{aa} = \sum_b C_{ab} \sum_j L_{aj}^{\text{CCS}} R_{bj}^{\text{CCS}} = \sum_j C_{ab} D_{ab}^{\text{CCS}}. \quad (67)$$

Finally, the second-order doubles cluster amplitudes and right eigenvectors may be expressed as

TABLE 5 The first 5 roots for the retinal excited states, all values in eV.

Root	CCS	CPS(D-2)	CPS(D-3)
1	3.274	2.370	2.498
2	5.038	3.621	3.930
3	6.105	4.650	4.874
4	6.161	4.443	4.560
5	6.993	5.388	5.590

$$\tilde{t}_{aibj}^{(2)} = \frac{2X_{aibj}^t - X_{ajbi}^t}{\epsilon_i - \epsilon_a + \epsilon_j - \epsilon_b} \quad (68)$$

and

$$\tilde{R}_{aibj}^{(2)} = \frac{2X_{aibj}^R - X_{ajbi}^R}{\epsilon_i - \epsilon_a + \epsilon_j - \epsilon_b}, \quad (69)$$

where the second order intermediates X_{aibj}^t and X_{aibj}^R are given in Tables 5 and VI of Ref. (Baudin et al., 2019).

6.5 The multi-node multi-GPU accelerated CPS(D-3)

The implementation of the CPS(D-3) is based on the RI approximation (Vahtras et al., 1993; Hättig and Weigend, 2000) for two-electron repulsion integrals, which facilitates the efficient calculation of $\omega^{(3)}$ in batches over a virtual index. The RI approximation is achieved by constructing three-center integral fitting coefficients from an auxiliary basis set using the Coulomb metric:

$$B_{pq}^P = \sum_P (pq|P)(P|Q)^{-1/2}. \quad (70)$$

This allows the four-index two-electron repulsion integrals to be represented by

$$(pq|rs) = \sum_{PQ} (pq|P)(P|Q)^{-1}(Q|rs) = \sum_P B_{pq}^P B_{rs}^P. \quad (71)$$

The three-index fitting coefficients obtained from the RI approximation can be written to disk and read into memory in batches over a virtual index. This enables the calculation of the third-order excitation energy correction using a massively parallel algorithm that can efficiently adapt to the available computational resources.

By calculating the incremental contributions to $\omega^{(3)}$ independently and on-demand for each batch, the need to store fourth-order tensors in main memory and recalculate the first-order doubles amplitudes $t_2^{(1)}$ is eliminated. This allows for scalability of the batch size according to the available memory. The main CPS(D-3) algorithm is summarized by Algorithm 1.

Step 1: Initialize transformation matrices, integral fitting coefficients, and Fock matrices

Step 2: Calculate fully virtual three index integral fitting coefficients and distribute in global memory

for A (batch over virtual index a) do

Step 3.1: Calculate second order doubles intermediate X_{Aibj}^t

Step 3.2: Calculate $(iA|jb)$

Step 3.3: Update excitation energy contribution,

$$\omega^{(3)} = \omega^{(3)} + \sum_{a \in A} \sum_{ibj} \frac{(2X_{aibj}^t - X_{ajbi}^t)}{\epsilon_i - \epsilon_a + \epsilon_j - \epsilon_b} (\hat{P}_{ij}^{ab} R_{ai}^{CCS} \bar{F}_{jb} - (iA|jb))$$

end for

Step 4: Deallocate B_{ia}^P and B_{ia}^P from memory

Step 5: Calculate three center integral fitting coefficients B_{ai} , B_{ai} , and B_{ij}

for A (batch over virtual index a) do

Step 6.1: Calculate second order doubles intermediate X_{Aibj}^R

Step 6.2: Calculate $(A|\bar{i}b_j)$

Step 6.3: Update excitation energy contribution,


$$\omega^{(3)} = \omega^{(3)} + \sum_{a \in A} \sum_{ibj} \left(\frac{(2X_{aibj}^R - X_{ajbi}^R)}{\epsilon_i - \epsilon_a + \epsilon_j - \epsilon_b} + 2R_{ai}^{CCS} t_{bj}^{(2)} - R_{aj}^{CCS} t_{bi}^{(2)} \right) (A|\bar{i}b_j)$$

end for

Algorithm 1. Massively Parallel Algorithm for Calculating the Third Order Excitation Energy Correction to CCS via the CPS(D-3) model.

The CPS(D-3) algorithm described in Algorithm 1 is parallelized for distributed memory architectures using a master-worker model and MPI. The algorithm for the $\omega^{(3)}$ calculation is initiated by the master rank, which determines the transformation matrices \bar{X} , \bar{Y} , \bar{X} , and \bar{Y} . Next, the algorithm calculates the three-center integral fitting coefficients, B_{ai}^P , B_{ij}^P , B_{ia}^P , and B_{ia}^P , followed by the one-index transformed Fock matrices, \bar{F}_{ia} , \bar{F}_{ij} , and \bar{F}_{ab} . Finally, the fully virtual three-index integral fitting coefficients, B_{ab}^P and B_{ab}^P , are determined. Once these quantities are determined by the master rank, they are broadcasted to all other ranks. To distribute the fully virtual three index integral fitting coefficients, tiled distributed tensors are utilized through the Scalable Tensor Library (ScaTeLib) (Ettenhuber, 2023). Subsequently, the master rank assigns the calculation of the first contributions to $\omega^{(3)}$ to the other ranks, to be computed in batches over the virtual index a . These ranks execute the computation of the contribution to $\omega^{(3)}$ for a particular batch and return the incremental contribution. The master rank then accumulates these incremental contributions to determine the total $\omega^{(3)}$ correction. After all workers complete the first loop of Algorithm 1, the master node identifies a new set of three center integral fitting coefficients required for the second contribution to $\omega^{(3)}$, which is computed in the second loop of Algorithm 1. The master rank then assigns the computation of the second contribution to $\omega^{(3)}$ to the other ranks in batches over the virtual index a . This process leads to the determination of the full $\omega^{(3)}$ correction.

In the primary algorithm of the CPS(D-3) implementation, Algorithm 2 is employed to calculate the X_{aibj}^t intermediates for a given batch A . Minor modifications are made to utilize the algorithm for the computation of X_{aibj}^R ; Algorithm 2 begins by calculating the first-order correction for the doubles excitation amplitudes, $t_{Aidk}^{(1)}$. It then retrieves the tiled distributed tensor B_{cA}^P . The algorithm proceeds with two nested loops, iterating over virtual index D in batches, and occupied index L in batches. In the loop over virtual index D , several intermediate values are computed, such as $t_{Aidk}^{(1)}$ ($bD|kj$) ($cA|bD$), $t_{AkDj}^{(1)}$, and $t_{Djci}^{(1)}$.



These values are used to iteratively update the X_{Aibj}^t . Once this loop is completed, the algorithm proceeds with the loop over the occupied index L . Here, $t_{AkbL}^{(1)}$ and $(ik|jL)$ are calculated and used to further update the X_{Aibj}^t value. After completing the loop over L , the algorithm calculates Y_{Ai}^P outside of the loops and uses it to finalize the update of the X_{Aibj}^t value. This method allows for an efficient and systematic computation of the first-order correction to the doubles excitation amplitudes and the X_{aibj}^t (or X_{aibj}^R) value for a given batch A , accounting for the contributions from both occupied and virtual indices in the process.

Step 1: Calculate first-order correction for doubles excitation amplitudes, $t_{Aidk}^{(1)} = (Aid|dk) / (\epsilon_i - \epsilon_A + \epsilon_k - \epsilon_d)$

Step 2: Retrieve tiled distributed tensor B_{CA}^P

for D (batch over virtual index) **do**

Step 3.1: Extract $t_{Aidk}^{(1)}$ from $t_{Aidk}^{(1)}$

Step 3.2: Get tiled distributed tensor B_{bD}^P

Step 3.3: Calculate $(bd|kj)$

Step 3.4: Update X_{Aibj}^t using $t_{Aidk}^{(1)}$ and $(bd|kj)$

Step 3.5: Update X_{Aibj}^t using $t_{Akdj}^{(1)}$ and $(bd|ki)$

Step 3.6: Calculate $(cA|bD)$

Step 3.7: Calculate first-order correction for $t_{Djci}^{(1)}$

Step 3.8: Update X_{Aibj}^t using $t_{Djci}^{(1)}$ and $(cA|bD)$

end for

for L (batch over occupied index) **do**

Step 4.1: Extract $t_{AkbL}^{(1)}$ from $t_{AkbL}^{(1)}$

Step 4.2: Calculate $(ik|jL)$

Step 4.3: Update X_{Aibj}^t using $t_{AkbL}^{(1)}$ and $(ik|jL)$

end for

Step 5: Calculate Y_{Ai}^P using the given summation formula

Step 6: Update the X_{Aibj}^t value using the calculated values of Y_{Ai}^P and B_{bj}^P

Algorithm 2. Algorithm for calculating the intermediate X_{aibj}^t for a given batch A .

Algorithm 1; Algorithm 2 provide a scheme that enables the calculation of CPS(D-3) excitation energies for system sizes beyond the reach of conventional CCSD calculations. To perform all the costly tensor contractions in each batch, these algorithms use the Tensor Algebra Library for Shared Memory Computers (TALSH) (Lyakh, 2023), which transfers them to GPUs. Open Multiprocessing (OMP) allows each rank to exploit shared memory parallelism locally. Moreover, the computation of $\langle \mu_1 | [\Phi, T_2^{(1)}] | HF \rangle$ is parallelized as the $\mathcal{O}(N^5)$ scaling of this term is a bottleneck for larger systems.

6.6 Numerical illustrations: CPS(D-3) excitation energies

In the case of CPS(D- n) methods, the retinal molecule (see Figure 4) with a cc-pVDZ basis set serves as an excellent example of a molecule with degenerate orbitals and delocalized electrons. This molecule is also of significant importance due to its biological role in photochemical reactions. The retinal system contains 63 atoms, and for the demonstrative calculation, 540 basis functions and 1932 auxiliary basis functions were employed. The calculation was performed using the GPU-accelerated LS-Dalton implementation on the Summit supercomputer at Oak Ridge National Laboratory. The first five roots, corresponding to the first excitation energies, are summarized in the table below.

The total time for the CPS(D-3) calculation on the retinal molecules was less than 20 min using 64 MPI processes. This balance between wall time and computational resources makes approximations such as CPS(D-3) practical tools capable of providing experimental molecular scientists with chemical observables at an accuracy of the CC target model (in this case, CCSD).

7 Discussion and outlook

The DEC scheme presents a divide-and-conquer linear-scaling and massively parallel framework for CC calculations, assuring error control in a black-box manner. These characteristics make DEC an enticing computational basis for molecular modeling on extensive molecular systems. Specifically, the DEC scheme imparts CC methods with computational abilities that are generally unattainable with conventional CC algorithms. Nevertheless, the DEC framework also faces challenges, such as the high DEC prefactor resulting from the need to recalculate integrals and amplitudes due to overlapping orbital spaces across different fragments.

In single compute node calculations, the crossover point in computational effort between DEC and a traditional, canonical calculation arises in large molecular systems. If a canonical calculation is achievable, it is likely more efficient than the DEC calculation. However, when multiple compute nodes are available, the massively parallel nature of the DEC algorithm leads to a reduced time-to-solution compared to a canonical calculation. As floating point operations continue to become more cost-effective and the number of cores on a compute node increases, parallelization will be crucial, and a considerable amount of recalculation will be acceptable if it facilitates a massively parallel computational strategy.

The DEC scheme, in principle, lays the foundation for linear-scaling and massively parallel implementations of any CC model. However, multiple technical challenges must be addressed before DEC can become a mainstream tool for more accurate CC models. For example, although the error control of the DEC-CCSD(T) method has been tackled, the current DEC-CCSD(T) algorithm can only be applied to large molecules for loose FOT values due to the fragment sizes encountered. When analyzing large molecules with a triple- ζ quality basis set, the resulting fragments often contain over 1000 basis functions, making such calculations impractical

even for massively parallel CCSD(T) implementations and hindering high-accuracy DEC-CCSD(T) applications. The large pair fragments, however, are frequently associated with minor energy contributions, and it may be possible to further reduce the pair orbital spaces without sacrificing the accuracy of the final correlation energy. Alternatively, the scaling of high-level CC fragment calculations could be decreased by considering tensor factorization techniques, PNOs, or other fragment-specific orbitals. Combining the DEC scheme with PNO-based local CC methods, for instance, would fully harness the sparsity of correlation effects in each fragment calculation, enabling CC calculations on systems of unprecedented sizes by removing all bottlenecks of wave function-based approximations through the DEC scheme—provided that the HF solver which generates the underlying reference is implemented with comparable efficiency.

Lastly, it is essential to note that the implementation of the DEC scheme ensures performance portability, allowing the DEC scheme to automatically benefit from new hardware developments and expand its application range as computational resources become more accessible. Although this work anticipates a promising future for DEC CC calculations, it is crucial to remember that many chemistry applications require energy differences rather than total energies. In these scenarios, computational chemistry practitioners should concentrate on direct approaches, such as CP theory.

CP theories provide a promising potential in addressing the limitations of total energy calculations with DEC for large molecular systems. The systematic approach offered by CP theories allows for the accurate calculation of excitation energies and other molecular properties at a CCSD level, providing a more robust, efficient and direct way to understand and predict the behavior of complex molecular systems.

The development of the CPS(D-3) model, as discussed in the literature, serves as a significant advancement in this area. By employing perturbation corrections up to the third order within the CP framework, the CPS(D-3) model delivers excitation energies of CCSD quality. The non-iterative nature of the $\omega_x^{(2)}$ and $\omega_x^{(3)}$ correction calculations in the doubles excitation space enables a massively parallel implementation, distributing batches over a virtual molecular orbital index across different compute nodes using the MPI.

The application of modern heterogeneous supercomputers and GPUs for accelerating heavy tensor contractions further enhances the efficiency and scalability of the CPS(D-3) calculations. This approach makes it feasible to calculate high quality excitation energies for molecular systems with several thousand basis functions, given sufficient computational resources.

Moreover, CP theory demonstrates versatility in its potential applications, extending to the calculation of other time-independent and time-dependent properties (Pawłowski et al., 2019c; Hillers-

Bendtsen et al., 2022). This adaptability paves the way for the development of similar massively parallel implementations in other areas of research.

Data availability statement

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

Author contributions

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

Funding

This research used resources of the Oak Ridge Leadership Computing Facility (OLCF), which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. AH-B and KM acknowledge the Danish Council for Independent Research, DFF-0136-00081B, and the European Union's Horizon 2020 Framework Programme under grant agreement number 951801 for financial support. AYZ acknowledges the additional computational resources that were provided by the Pinnacles cluster at UC Merced, which is supported by the National Science Foundation under OAC-2019144.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Abyar, F., and Novak, I. (2022). Electronic structure analysis of riboflavin: OVG and EOM-CCSD study. *Acta A Mol. Biomol. Spectrosc.* 264, 120268. doi:10.1016/j.saa.2021.120268
- Adler, T. B., Werner, H.-J., and Frederick, R. (2009). Local explicitly correlated second-order perturbation theory for the accurate treatment of large molecules. *J. Chem. Phys.* 130, 054106. doi:10.1063/1.3040174
- Ali, M. F., and Khan, R. Z. (2012). The study on load balancing strategies in distributed computing system. *Int. J. Comput. Sci. Eng. Surv.* 3, 19–30. doi:10.5121/ijcses.2012.3203
- Altman, E., Brown, K. R., Carleo, G., Carr, L. D., Demler, E., Chin, C., et al. (2021). *Quantum Simulators: Architectures and Opportunities PRX Quantum*, 2.017003.
- Altun, A., Izsák, R., and Bistoni, G. (2021). Local energy decomposition of coupled-cluster interaction energies: Interpretation, benchmarks, and comparison with symmetry-adapted perturbation theory. *Int. J. Quantum Chem.* 121, e26339. doi:10.1002/qua.26339
- Amos, R. D., and Rice, J. E. (1989). Implementation of analytic derivative methods in quantum chemistry. *Comput. Phys. Rep.* 10 (4), 147–187. doi:10.1016/0167-7977(89)90001-4

- Andrade, X., Strubbe, D., De, G., Larsen, A. H., Oliveira, M. J. T., Alberdi-Rodriguez, Joseba, et al. (2015). Real-space grids and the Octopus code as tools for the development of new simulation approaches for electronic systems. *Phys. Chem. Chem. Phys.* 17, 31371–31396. doi:10.1039/c5cp00351b
- Ballesteros, Francisco, Dunivan, Shelbie, and Lao, Ka Un (2021). Coupled cluster benchmarks of large noncovalent complexes: The L7 dataset as well as DNA–ellipticine and buckycatcher–fullerene. *J. Chem. Phys.* 154, 154104. doi:10.1063/5.0042906
- Barnes, A. L., Bykov, D., Lyakh, D. I., and Tjerk, P. (2019). Multilayer divide-expand-consolidate coupled-cluster method: Demonstrative calculations of the adsorption energy of carbon dioxide in the Mg-MOF-74 metal-organic framework. *J. Phys. Chem. A* 123, 8734–8743. doi:10.1021/acs.jpca.9b08077
- Bartlett, R. J. (2012). Coupled-cluster theory and its equation-of-motion extensions. *WIREs Comput. Mol. Sci.* 2 (1), 126–138. doi:10.1002/wcms.76
- Bartlett, R. J., and Shavitt, I. (1977). Comparison of high-order many-body perturbation theory and configuration interaction for H₂O. *Phys. Lett.* 50, 190–198. doi:10.1016/0009-2614(77)80161-9
- Bartlett, R. J., and Silver, D. M. (1974). Correlation energy in LiH, BH, and HF with many-body perturbation theory using Slater-type atomic orbitals. *Int. J. Quantum Chem.* 8, 271–276. doi:10.1002/qua.560080831
- Battaglino, C., Ballard, G., and Tamara, G. (2018). A practical randomized CP tensor decomposition. *SIAM J. Matrix Analysis Appl.* 39, 876–901. doi:10.1137/17m112303
- Baudin, P., Bykov, D., Liakh, D., Ettenhuber, P., and Kristensen, K. (2017). A local framework for calculating coupled cluster singles and doubles excitation energies (LoFEx-CCSD). *Mol. Phys.* 115, 2135–2144. doi:10.1080/00268976.2017.1290836
- Baudin, P., and Kristensen, K. (2016). LoFEx — a local framework for calculating excitation energies: Illustrations using RI-CC2 linear response theory. *J. Chem. Phys.* 144, 224106. doi:10.1063/1.4953360
- Baudin, P., Pawłowski, F., Bykov, D., Liakh, D., Kristensen, K., Olsen, J., et al. (2019). Cluster perturbation theory. III. Perturbation series for coupled cluster singles and doubles excitation energies. *J. Chem. Phys.* 150, 134110. doi:10.1063/1.5046935
- Baumgartner, G., Auer, A., Bernholdt, D. E., Bibireata, A., Choppella, V., Cociorva, D., et al. (2005). Synthesis of high-performance parallel programs for a class of *ab initio* quantum chemistry models. *Proc. IEEE* 93 (2), 276–292. doi:10.1109/JPROC.2004.840311
- Binkley, J. S., and Pople, J. A. (1975). Møller–Plesset theory for atomic ground state energies. *Int. J. Quantum Chem.* 9, 229–236. doi:10.1002/qua.560090204
- Bistoni, G., Ripplinger, C., Minenkov, Y., Cavallo, L., Auer, A. A., and Neese, F. (2017). Treating subvalence correlation effects in domain based pair natural orbital coupled cluster calculations: An out-of-the-box approach. *J. Chem. Theory Comput.* 13, 3220–3227. doi:10.1021/acs.jctc.7b00352
- Boudehane, A., Albera, L., Tenenhaus, A., Le Brusquet, L., and Boyer, R. (2022). Parallelization scheme for canonical polyadic decomposition of large-scale high-order tensors. *Signal Processing* 199, 108610.
- Boys, S. F. (1960). Construction of some molecular orbitals to be approximately invariant for changes from one molecule to another. *Rev. Mod. Phys.* 32, 296–299. doi:10.1103/revmodphys.32.296
- Boys, S. F., and Rajagopal, P. (1966). Quantum calculations: Which are accumulative in accuracy, unrestricted in expansion functions, and economical in computation. *Adv. Quantum Chem.* 2, 1–24.
- Bykov, D., and Kjaergaard, T. (2017). The GPU-enabled divide-expand-consolidate RI-MP2 method (DEC-RI-MP2). *J. Comput. Chem.* 38, 228–237. doi:10.1002/jcc.24678
- Bykov, D., Kristensen, K., and Kjaergaard, T. (2016). The molecular gradient using the divide-expand-consolidate resolution of the identity second-order Møller–Plesset perturbation theory: The DEC-RI-MP2 gradient. *J. Chem. Phys.* 145, 024106. doi:10.1063/1.4956454
- Carroll, J. D., and Chang, J. J. (1970). Analysis of individual differences in multidimensional scaling via an n-way generalization of Eckart–Young decomposition. *Psychometrika*, 35, 283–319.
- Cederbaum, L. S. (2008). Born–Oppenheimer approximation and beyond for time-dependent electronic processes. *J. Chem. Phys.* 128, 124101. doi:10.1063/1.2895043
- Christensen, A. S., Kubar, T., Cui, Q., and Elstner, M. (2016). Semiempirical quantum mechanical methods for noncovalent interactions for chemical and biochemical applications. *Chem. Rev.* 116 (9), 5301–5337. doi:10.1021/acs.chemrev.5b00584
- Christiansen, O., Bak, K. L., Koch, H. S., and Stephan, P. A. (1998). A second-order doubles correction to excitation energies in the random-phase approximation. *Phys. Lett.* 284, 47–55. doi:10.1016/s0009-2614(97)01285-2
- Čížek, J. (1966). On the correlation problem in atomic and molecular systems. Calculation of wavefunction components in Ursell-type expansion using quantum-field theoretical methods. *J. Chem. Phys.* 45, 4256–4266. doi:10.1063/1.1727484
- Collins, J. B., Schleyer, V. R., Binkley, J. S., and Pople, J. A. (1976). Self-consistent molecular orbital methods. XVII. Geometries and binding energies of second-row molecules. A comparison of three basis sets. *J. Chem. Phys.* 64, 5142–5151. doi:10.1063/1.432189
- Combes, J.-M., Duclos, P., and Seiler, R. (1981). The Born–Oppenheimer approximation. *Rigorous At. Mol. Phys.*, 185–213.
- Corzo, H. H., Sehanobish, A., and Kara, O. (2021). Learning full configuration interaction electron correlations with deep learning. *Mach. Learn. Phys. Sci. Neural Inf. Processing Syst.*, 35. doi:10.48550/ARXIV.2106.08138
- Dalgaard, E., and Monkhorst, H. J. (1983). Some aspects of the time-dependent coupled-cluster approach to dynamic response functions. *Phys. Rev. A* 28, 1217–1222. doi:10.1103/physrev.28.1217
- Datta, D., and Gordon, M. S. (2021). A massively parallel implementation of the CCSD(T) method using the resolution-of-the-identity approximation and a hybrid distributed/shared memory parallelization model. *J. Chem. Theory Comput.* 17, 4799–4822. doi:10.1021/acs.jctc.1c00389
- Davidson, E. R. (1972). Natural orbitals. *Adv. Quantum Chem.* 6, 235–266.
- Davidson, E. R., and Feller, D. (1986). Basis set selection for molecular calculations. *Chem. Rev.* 86, 681–696. doi:10.1021/cr00074a002
- Davidson, E. R. (1972). Properties and uses of natural orbitals. *Rev. Mod. Phys.* 44, 451–464. doi:10.1103/revmodphys.44.451
- Díaz-Tinoco, M., Dolgounitcheva, O., Zakrzewski, V. G., and Ortiz, J. V. (2016). Composite electron propagator methods for calculating ionization energies. *J. Chem. Phys.* 144, 224110. doi:10.1063/1.4953666
- Dral, P. O., Wu, X., and Thiel, W. (2019). Semiempirical quantum-chemical methods with orthogonalization and dispersion corrections. *J. Chem. Theory Comput.* 15, 1743–1760. doi:10.1021/acs.jctc.8b01265
- Dunning Jr., and Thom, H. (1989). Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.* 90, 1007–1023. doi:10.1063/1.456153
- Edmiston, C., and Krauss, M. (1965). Configuration-interaction calculation of H₃ and H₂. *J. Chem. Phys.* 42, 1119–1120. doi:10.1063/1.1696050
- Edmiston, C., and Ruedenberg, K. (1963). Localized atomic and molecular orbitals. *Rev. Mod. Phys.* 35, 457–464. doi:10.1103/revmodphys.35.457
- Elstner, M., Frauenheim, T., Kaxiras, E., Seifert, G., and Suhai, S. (2000). A self-consistent charge density-functional based tight-binding scheme for large biomolecules. *Phys. Status Solidi B* 217, 357–376. doi:10.1002/(sici)1521-3951(200001)217:1<357::aid-psb357>3.0.co;2-j
- Elstner, M., and Seifert, G. (2014). Density functional tight binding. *Philos. Trans. A Math. Phys. Eng. Sci.* 372, 20120483. doi:10.1098/rsta.2012.0483
- Eriksen, J. J., Baudin, P., Ettenhuber, P., Kristensen, K., Kjaergaard, T., and Jørgensen, P. (2015). Linear-scaling coupled cluster with perturbative triple excitations: The divide-expand-consolidate CCSD (T) model. *J. Chem. Theory Comput.* 11, 2984–2993. doi:10.1021/acs.jctc.5b00086
- Eriksen, J. J., Jørgensen, P., and Gauss, J. (2015). On the convergence of perturbative coupled cluster triples expansions: Error cancellations in the CCSD (T) model and the importance of amplitude relaxation. *J. Chem. Phys.* 142, 014102. doi:10.1063/1.4904754
- Eriksen, J. J., Kristensen, K., Kjaergaard, T., Jørgensen, P., and Gauss, J. (2014). A Lagrangian framework for deriving triples and quadruples corrections to the CCSD energy. *J. Chem. Phys.* 140, 064108. doi:10.1063/1.4862501
- Eriksen, J. J., Matthews, Devin A., Jørgensen, Poul, and Gauss, Jürgen (2015). Communication: The performance of non-iterative coupled cluster quadruples models. *J. Chem. Phys.* 143, 041101. doi:10.1063/1.4927247
- Ettenhuber, P., Baudin, P., Kjaergaard, T., Jørgensen, P., and Kristensen, K. (2016). Orbital spaces in the divide-expand-consolidate coupled cluster method. *J. Chem. Phys.* 144 (16), 164116. doi:10.1063/1.4947019
- Ettenhuber, Patrick (2023). *ScaTeLib - a scalable tensor library*.
- Favier, G., and de Almeida, A. L. (2014). Overview of constrained PARAFAC models. *EURASIP, J. Adv. Signal Process.* 142.
- Fedorov, D. G. (2017). The fragment molecular orbital method: Theoretical development, implementation in GAMESS, and applications wiley interdisciplinary reviews. *Comput. Mol. Sci.* 7, e1322.
- Fedorov, D. G., and Pham, B. Q. (2023). Multi-level parallelization of quantum-chemical calculations. *J. Chem. Phys.* 158 (16), 164106. doi:10.1063/5.0144917
- Foster, I. (1995). *Designing and building parallel programs: Concepts and tools for parallel software engineering*.
- Friedrich, J., and Dolg, M. (2009). Fully automated incremental evaluation of MP2 and CCSD (T) energies: Application to water clusters. *J. Chem. Theory Comput.* 5, 287–294. doi:10.1021/ct800355e
- Friedrich, J., and Hänchen, J. (2013). Incremental CCSD(T)(F12*)/MP2: A black box method to obtain highly accurate reaction energies. *J. Chem. Theory Comput.* 9, 5381–5394. doi:10.1021/ct4008074
- Frisch, M. J., Pople, J. A., and Binkley, J. S. (1984). Self-consistent molecular orbital methods 25. Supplementary functions for Gaussian basis sets. *J. Chem. Phys.* 80 (7), 3265–3269. doi:10.1063/1.447079
- Frisch, M. J., Trucks, G. W., Schlegel, H. B., Scuseria, G. E., Robb, M. A., Cheeseman, J. R., et al. (2021). Gaussian development version revision. *J* 15.

- Fung, V., Zhang, J., Juarez, E., and Sumpter, B. G. (2021). Benchmarking graph neural networks for materials chemistry. *Npj Comput. Mater.* 7 (1), 1–8.
- Gonzalez-Escribano, A., Llanos, D. R., Orden, D., and Palop, B. (2006). Parallelization alternatives and their performance for the convex hull problem. *Appl. Math. Model.* 30, 563–577. doi:10.1016/j.apm.2005.05.022
- Götz, A. W., Williamson, M. J., Xu, D., Poole, D., Le GrandScott, et al. (2012). Routine microsecond molecular dynamics simulations with AMBER on GPUs. I. Generalized born. *Gen. born J. Chem. Theory Comput.* 8, 1542–1555. doi:10.1021/ct200909j
- Grimme, S., Antony, J., Ehrlich, S., and Krieg, H. (2010). A consistent and accurate *ab initio* parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *J. Chem. Phys.* 132, 154104. doi:10.1063/1.3382344
- Gyevi-Nagy, L., Kállay, M., and Nagy, P. R. (2021). Accurate reduced-cost CCSD(T) energies: Parallel implementation, benchmarks, and large-scale applications. *J. Chem. Theory Comput.* 17, 860–878. doi:10.1021/acs.jctc.0c01077
- Gyevi-Nagy, L., Kállay, M., and Nagy, P. R. (2019). Integral-direct and parallel implementation of the CCSD(T) method: Algorithmic developments and large-scale applications. *J. Chem. Theory Comput.* 16, 366–384. doi:10.1021/acs.jctc.9b00957
- Hagebaum-Reignier, D., Girardi, R., Carissan, Y., and Humbel, S. (2007). Hückel theory for Lewis structures: Hückel–Lewis configuration interaction (HL-CI). *J. Mol. Struct. THEOCHEM.* 817, 99–109. doi:10.1016/j.theochem.2007.04.026
- Haghighatdari, M., and Hachmann, J. (2019). Advances of machine learning in molecular modeling and simulation. *Curr. Opin. Chem. Eng.* 23, 51–57. doi:10.1016/j.coche.2019.02.009
- Hampel, C., and Werner, H-J (1996). Local treatment of electron correlation in coupled cluster theory. *J. Chem. Phys.* 104, 6286–6297. doi:10.1063/1.471289
- Harris, F. E. (1977). Coupled-cluster method for excitation energies. *Int. J. Quantum Chem.* 12, 403–411. doi:10.1002/qua.560120848
- Hasanein, A. A., and Evans, M. W. (1996). *Computational methods in quantum chemistry*, 2.
- Häser, M. (1993). Møller-Plesset (MP2) perturbation theory for large molecules. *Theor. Chim. Acta* 87, 147–173.
- Hättig, C., Hellweg, A., and Köhn, A. (2006). Distributed memory parallel implementation of energies and gradients for second-order Møller–Plesset perturbation theory with the resolution-of-the-identity approximation. *Phys. Chem. Chem. Phys.* 8, 1159. doi:10.1039/b515355g
- Hättig, C., and Weigend, F. (2000). CC2 excitation energy calculations on large molecules using the resolution of the identity approximation. *J. Chem. Phys.* 113, 5154–5161. doi:10.1063/1.1290013
- Head-Gordon, M. R., Rudolph, J., Oumi, M., and Lee, T. J. (1994). A doubles correction to electronic excited states from configuration interaction in the space of single substitutions. *Chem. Phys. Lett.* 219, 21–29. doi:10.1016/0009-2614(94)00070-0
- Hehre, W. J., Stewart, R. F., and Pople, J. A. (1969). Self-consistent molecular-orbital methods. I. Use of Gaussian expansions of Slater-type atomic orbitals. *J. Chem. Phys.* 51, 2657–2664. doi:10.1063/1.1672392
- Helgaker, T., Coriani, S., Jørgensen, P., Kristensen, K., Olsen, J., and Ruud, K. (2012). Recent advances in wave function-based methods of molecular-property calculations. *Chem. Rev.* 112 (1), 543–631. doi:10.1021/cr2002239
- Helgaker, T., Jørgensen, P., and Olsen, J. (2000). *Molecular electronic-structure theory*.
- Helmich, B., and Hättig, C. (2014). A pair natural orbital based implementation of ADC(2)-x: Perspectives and challenges for response methods for singly and doubly excited states in large molecules. *Comput. Theor. Chem.* 1040, 1041 35–44.
- Herbert, J. M. (2019). Fantasy versus reality in fragment-based quantum chemistry. *J. Chem. Phys.* 151, 170901. doi:10.1063/1.5126216
- Hillers-Bendtsen, A. E., Bykov, D., Barnes, A., Liakh, D., Corzo, H. H., Olsen, J., et al. (2023). Massively parallel GPU enabled third-order cluster perturbation excitation energies for cost-effective large scale excitation energy calculations. *J. Chem. Phys.* 158 (14), 144111. doi:10.1063/5.0142780
- Hillers-Bendtsen, A. E., Hoyer, N. M., Kjeldal, F. Ø., Mikkelsen, K. V., Olsen, J., and Jørgensen, P. (2022). Cluster perturbation theory. VIII. First order properties for a coupled cluster state. *J. Chem. Phys.* 157, 024108. doi:10.1063/5.0082585
- Hitchcock, F. L. (1927). The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys.* 6, 164–189. doi:10.1002/sapm192761164
- Hoy, E. P., and Mazziotti, D. A. (2015). Positive semidefinite tensor factorizations of the two-electron integral matrix for low-scaling *ab initio* electronic structure. *J. Chem. Phys.* 143, 064103. doi:10.1063/1.4928064
- Høyer, N. M., Kjeldal, F. Ø., Hillers, B., Erbs, A., Mikkelsen, K. V., Olsen, J., et al. (2022). Cluster perturbation theory. VI. Ground-state energy series using the Lagrangian. *J. Chem. Phys.* 157, 024106. doi:10.1063/5.0082583
- Høyvik, I.-M., Jansik, B., and Jørgensen, P. (2012). Orbital localization using fourth central moment minimization. *J. Chem. Phys.* 137, 224114. doi:10.1063/1.4769866
- Høyvik, I.-M., and Jørgensen, P. (2016). Characterization and generation of local occupied and virtual Hartree–Fock orbitals. *Chem. Rev.* 116, 3306–3327. doi:10.1021/acs.chemrev.5b00492
- Høyvik, I.-M., Kristensen, K., Jansik, B., and Jørgensen, P. (2012). The divide-expand-consolidate family of coupled cluster methods: Numerical illustrations using second order møller-Plesset perturbation theory. *J. Chem. Phys.* 136, 014105. doi:10.1063/1.3667266
- Høyvik, I. M., Kristensen, K., Kjærgaard, T., and Jørgensen, P. (2014). A perspective on the localizability of Hartree–Fock orbitals. *Theor. Chem. Acc.* 133, 1417. doi:10.1007/s00214-013-1417-x
- Ishikawa, T., and Kuwata, K. (2012). RI-MP2 gradient calculation of large molecules using the fragment molecular orbital method. *J. Phys. Chem. Lett.* 3, 375–379. doi:10.1021/jz201697x
- Jakobsen, S., Kristensen, K., and Jensen, F. (2013). Electrostatic potential of insulin: Exploring the limitations of density functional theory and force field methods. *J. Chem. Theory Comput.* 9, 3978–3985. doi:10.1021/ct400452f
- Jansik, B., Høst, S., Kristensen, K., and Jørgensen, P. (2011). Local orbitals by minimizing powers of the orbital variance. *J. Chem. Phys.* 134, 194104. doi:10.1063/1.3590361
- Jha, A., Nottoli, M., Mikhalev, A., Quan, C., and Stamm, B. (2022). Linear scaling computation of forces for the domain-decomposition linear Poisson–Boltzmann method. *J. Chem. Phys.* 158 (10), 104105. doi:10.1063/5.0141025
- Kapuy, E., Csépes, Z., and Kozmutza, C. (1983). Application of the many-body perturbation theory by using localized orbitals. *Int. J. Quantum Chem.* 23, 981–990. doi:10.1002/qua.560230321
- Keith, J. A., Vassilev-Galindo, V., Cheng, B., Chmiela, S., Gastegger, M., Müller, K.-R., et al. (2021). Combining machine learning and computational chemistry for predictive insights into chemical systems. *Chem. Rev.* 121, 9816–9872. doi:10.1021/acs.chemrev.1c00107
- Khoromskaja, V., and Khoromskij, B. N. (2015). Tensor numerical methods in quantum chemistry: From Hartree–Fock to excitation energies. *Phys. Chem. Chem. Phys.* 17, 31491–31509. doi:10.1039/c5cp01215e
- Kirtman, B. (1995). Local quantum chemistry: The local space approximation for Møller–Plesset perturbation theory. *Int. J. Quantum Chem.* 55 (2), 103–108. doi:10.1002/qua.560550204
- Kjærgaard, T., Baudin, P., Bykov, D., Eriksen, J. J., Eitenhuber, P., Kristensen, K., et al. (2017). Massively parallel and linear-scaling algorithm for second-order Møller–Plesset perturbation theory applied to the study of supramolecular wires. *Comput. Phys. Commun.* 212, 152–160. doi:10.1016/j.cpc.2016.11.002
- Kjærgaard, T., Baudin, P., Bykov, D., Kristensen, K., and Jørgensen, P. (2017). The divide-expand-consolidate coupled cluster scheme. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* 7, e1319.
- Kjærgaard, T. (2017). The Laplace transformed divide-expand-consolidate resolution of the identity second-order Møller–Plesset perturbation (DEC-LT-RIMP2) theory method. *J. Chem. Phys.* 146, 044103. doi:10.1063/1.4973710
- Kleier, D. A., Halgren, T. A., Hall Jr, John, H., and Lipscomb, W. N. (1974). Localized molecular orbitals for polyatomic molecules. I. A comparison of the Edmiston-Ruedenberg and Boys localization methods. *J. Chem. Phys.* 61, 3905–3919. doi:10.1063/1.1681683
- Kolda, T. G., and Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Rev.* 51 (3), 455–500. doi:10.1137/07070111x
- Krause, C., and Werner, H.-J. (2012). Comparison of explicitly correlated local coupled-cluster methods with various choices of virtual orbitals. *Phys. Chem. Chem. Phys.* 14, 7591–7604. doi:10.1039/c2cp40231a
- Krishnan, R. B. J. S., Binkley, J. S., Seeger, R., and Pople, J. A. (1980). Self-consistent molecular orbital methods. XX. A basis set for correlated wave functions. *J. Chem. Phys.* 72, 650–654. doi:10.1063/1.438955
- Krishnan, R., and Pople, J. A. (1978). Approximate fourth-order perturbation theory of the electron correlation energy. *Int. J. Quantum Chem.* 14, 91–100. doi:10.1002/qua.560140109
- Kristensen, K., Eriksen, J. J., Matthews, D. A., Olsen, J., and Jørgensen, P. (2016). A view on coupled cluster perturbation theory using a bivariational Lagrangian formulation. *J. Chem. Phys.* 144, 064103. doi:10.1063/1.4941605
- Kristensen, K., Høyvik, I.-M., Jansik, B., Jørgensen, P., Kjærgaard, T., Reine, S., et al. (2012). MP2 energy and density for large molecular systems with internal error control using the Divide-Expand-Consolidate scheme. *Phys. Chem. Chem. Phys.* 14, 15706–15714. doi:10.1039/c2cp41958k
- Kristensen, K., Jørgensen, P., Jansik, B., Kjærgaard, T., and Reine, S. (2012). Molecular gradient for second-order Møller–Plesset perturbation theory using the divide-expand-consolidate (DEC) scheme. *J. Chem. Phys.* 137, 114102. doi:10.1063/1.4752432
- Kristensen, K., Ziolkowski, M., Jansik, B., Kjærgaard, T., and Jørgensen, P. (2011). A locality analysis of the divide-expand-consolidate coupled cluster amplitude equations. *J. Chem. Theory Comput.* 7, 1677–1694. doi:10.1021/ct200114k
- Kurashige, Y., Yang, J., Chan, G. K.-L., and Manby, F. R. (2012). Optimization of orbital-specific virtuals in local Møller–Plesset perturbation theory. *J. Chem. Phys.* 136 (12), 124106. doi:10.1063/1.3696962

- Kutzelnigg, W (2007). What I like about Hückel theory. *J. Comput. Chem.* 28, 25–34. doi:10.1002/jcc.20470
- Levine, I. N., Busch, D. H., and Shull, H. (2009). *Quantum chemistry*. 6th edition.
- Li, R R., Liebenthal, M D., De Prince III, and Eugene, A. (2021). Challenges for variational reduced-density-matrix theory with three-particle N-representability conditions. *J. Chem. Phys.* 155, 174110. doi:10.1063/5.0066404
- Liakos, D G., Sparta, M, Kesharwani, M K., Martin, J M. L., and Neese, F (2015). Exploring the accuracy limits of local pair natural orbital coupled-cluster theory. *J. Chem. Theory Comput.* 11, 1525–1539. doi:10.1021/ct501129s
- Lin, N, Marianetti, C. A., Millis, A J., and Reichman, D R. (2011). Dynamical mean-field theory for quantum chemistry. *Phys. Rev. Lett.* 106, 096402. doi:10.1103/physrevlett.106.096402
- Lipparini, F, Stamm, B, Canc'és, E, Maday, Y, and Mennucci, B (2013). Fast domain decomposition algorithm for continuum solvation models: Energy and first derivatives. *J. Chem. Theory Comput.* 9, 3637–3648. doi:10.1021/ct400280b
- Liu, W (2020). Essentials of relativistic quantum chemistry. *J. Chem. Phys.* 152, 180901. doi:10.1063/5.0008432
- Löwdin, P O (1958). Correlation problem in many-electron quantum mechanics I. Review of different approaches and discussion of some current ideas. *Adv. Chem. Phys.* 207–322.
- Löwdin, P O (1955). Quantum theory of many-particle systems. II. Study of the ordinary Hartree-Fock approximation. *Phys. Rev.* 97, 1490–1508. doi:10.1103/physrev.97.1490
- Löwdin, P-O (1955). Quantum theory of many-particle systems. I. Physical interpretations by means of density matrices, natural spin-orbitals, and convergence problems in the method of configurational interaction. *Phys. Rev.* 97, 1474–1489. doi:10.1103/physrev.97.1474
- Luo, L., Straatsma, T., Suarez, L. E. A., Broer, R., Bykov, D., et al. (2020). Pre-exascale accelerated application development: The ORNL Summit experience. *IBM J. Res. Dev.* 64, 1–11. doi:10.1147/jrd.2020.2965881
- Lyakh, D I. (2023). *TAL-SH: Tensor algebra library for shared memory computers*.
- Ma, Y, Li, Z Y, Chen, X, Ding, B, Li, N, Lu, T, et al. (2023). Machine-learning assisted scheduling optimization and its application in quantum chemical calculations. *J. Comput. Chem.* 44, 1174–1188. doi:10.1002/jcc.27075
- Maslow, A. H. (1966). *The psychology of science: A reconnaissance*.
- Menezes, F, Kats, D, and Werner, H-J (2016). Local complete active space second-order perturbation theory using pair natural orbitals (PNO-CASPT2). *J. Chem. Phys.* 145, 124115. doi:10.1063/1.4963019
- Mester, D, Nagy, P R., and Kállay, M (2019). Reduced-Scaling correlation methods for the excited states of large molecules: Implementation and benchmarks for the second-order algebraic-diagrammatic construction approach. *J. Chem. Theory Comput.* 15, 6111–6126. doi:10.1021/acs.jctc.9b00735
- Mitxelena, I, and Piris, M (2022). Benchmarking GNOF against FCI in challenging systems in one, two, and three dimensions. *J. Chem. Phys.* 156, 214102. doi:10.1063/5.0092611
- Moawad, Y, Vanderbauwhede, W, and Steijl, R (2022). Investigating hardware acceleration for simulation of CFD quantum circuits. *Front. Mech. Eng.* 8. doi:10.3389/fmech.2022.925637
- Møller, C., and Plesset, M S. (1934). Note on an approximation treatment for many-electron systems. *Phys. Rev.* 46, 618–622. doi:10.1103/physrev.46.618
- Monari, A., Rivail, J-L., and Assfeld, X. (2013). Theoretical modeling of large molecular systems. Advances in the local self consistent field method for mixed quantum mechanics/molecular mechanics calculations. *Acc. Chem. Res.* 46, 596–603. doi:10.1021/ar300278j
- Nagy, B., and Jensen, F. (2017). “Basis sets in quantum chemistry,” in *Reviews in computational chemistry*, 93–149.
- Nagy, P R., and Kállay, M (2019). Approaching the basis set limit of CCSD(T) energies for large molecules with local natural orbital coupled-cluster methods. *J. Chem. Theory Comput.* 15, 5275–5298. doi:10.1021/acs.jctc.9b00511
- Neese, F, Wennmohs, F., and Hansen, A. (2009). Efficient and accurate local approximations to coupled-electron pair approaches: An attempt to revive the pair natural orbital method. *J. Chem. Phys.* 130, 114108. doi:10.1063/1.3086717
- Neese, F, Wennmohs, F, Becker, U, and Riplinger, C (2020). The ORCA quantum chemistry program package. *J. Chem. Phys.* 152, 224108. doi:10.1063/5.0004608
- Nesbet, R. K. (1955). Configuration interaction in orbital theories. *Proc. R. Soc. Lond. A Math. Phys. Sci.* 230, 312–321.
- Nikodem, A, Matveev, A V., Soini, T M., and Rösch, N (2014). Load balancing by work-stealing in quantum chemistry calculations: Application to hybrid density functional methods. *Int. J. Quantum Chem.* 114, 813–822. doi:10.1002/qua.24677
- Nottoli, M., Stamm, B., Scalmani, G., and Lipparini, F. (2019). Quantum calculations in solution of energies, structures, and properties with a domain decomposition polarizable continuum model. *J. Chem. Theory Comput.* 15, 6061–6073. doi:10.1021/acs.jctc.9b00640
- Olsen, J., Erbs, A., Kjeldal, F Ø, Høyer, N M, Mikkelsen, K V., et al. (2022). Cluster perturbation theory. VII. The convergence of cluster perturbation expansions. *J. Chem. Phys.* 157, 024107. doi:10.1063/5.0082584
- Olsen, J M. H., List, N H., Kristensen, K, and Kongsted, J (2015). Accuracy of protein embedding potentials: An analysis in terms of electrostatic potentials. *J. Chem. Theory Comput.* 11, 1832–1842. doi:10.1021/acs.jctc.5b00078
- Oseledets, I. V. (2011). Tensor-train decomposition. *SIAM J. Sci. Comput.* 33 (5), 2295–2317. doi:10.1137/090752286
- Ozog, D., Hammond, J. R., Dinan, J., Balaji, P., Shende, S., and Malony, A. (2013). “Inspector-executor load balancing algorithms for block-sparse tensor contractions,” in *2013 42nd International conference on parallel processing*, 30–39. doi:10.1109/ICPP.2013.12
- Patil, U, and Shedge, R (2013). Improved hybrid dynamic load balancing algorithm for distributed environment. *Int. J. Sci. Res. Publ.* 3, 1.
- Paudics, A, Hesz, D., Bojtár, M., Bitter, I., Horváth, V., Kállay, M., et al. (2022). A pillararene-based indicator displacement assay for the fluorescence detection of vitamin B1. *Sensors Actuators B Chem.* 369, 132364. doi:10.1016/j.snb.2022.132364
- Pawłowski, F., Olsen, J., and Jørgensen, P. (2019). Cluster perturbation theory. II. Excitation energies for a coupled cluster target state. *J. Chem. Phys.* 150, 134109. doi:10.1063/1.5053167
- Pawłowski, F., Olsen, J., and Jørgensen, P. (2019). Cluster perturbation theory. I. Theoretical foundation for a coupled cluster target state and ground-state energies. *J. Chem. Phys.* 150, 134108. doi:10.1063/1.5004037
- Pawłowski, F., Olsen, J., and Jørgensen, P. (2019). Cluster perturbation theory. IV. Convergence of cluster perturbation series for energies and molecular properties. *J. Chem. Phys.* 150, 134111. doi:10.1063/1.5053622
- Pawłowski, F., Olsen, J., and Jørgensen, P. (2019). Cluster perturbation theory. V. Theoretical foundation for cluster linear target states. *J. Chem. Phys.* 150, 134112. doi:10.1063/1.5053627
- Phan, A. H., Tichavský, P., and Cichocki, A. (2013). Fast alternating LS algorithms for high order CANDECOMP/PARAFAC tensor factorizations. *IEEE Trans. Signal Process.* 61, 4834–4846. doi:10.1109/tsp.2013.2269903
- Pinski, P., and Neese, F. (2018). Communication: Exact analytical derivatives for the domain-based local pair natural orbital MP2 method (DLPNO-MP2). *J. Chem. Phys.* 148, 031101. doi:10.1063/1.5011204
- Pipek, J., and Mezey, P. G. (1989). Pair natural orbitals: A concept for simplifying Hartree-Fock and CI wavefunctions. *J. Chem. Phys.* 90, 4916–4926. doi:10.1063/1.456588
- Pople, J. A. (1999). Nobel Lecture: Quantum chemical models. *Rev. Mod. Phys.* 71, 1267–1274. doi:10.1103/RevModPhys.71.1267
- Pople, J. A., Binkley, J. S., and Seeger, R. (1976). Theoretical models incorporating electron correlation. *Int. J. Quantum Chem.* 10, 1–19. doi:10.1002/qua.560100802
- Pulay, P., and Saebo, S. (1986). Orbital-invariant formulation and second-order gradient evaluation in Miller-Plesset perturbation theory. *Chem. Acc.* 69, 357–368. doi:10.1007/bf00526697
- Pulay, P., and Hamilton, T. P. (1988). UHF natural orbitals for defining and starting MC-SCF calculations. *J. Chem. Phys.* 88, 4926–4933. doi:10.1063/1.454704
- Pulay, P. (1983). Localizability of dynamic electron correlation. *Chem. Phys. Lett.* 100, 151–154. doi:10.1016/0009-2614(83)80703-9
- Pykkö, P. (2012). Relativistic effects in chemistry: More common than you thought. *Annu. Rev. Phys. Chem.* 63, 45–64. doi:10.1146/annurev-physchem-032511-143755
- Qiu, J., Zhao, Z., Wu, B., Vishnu, A., and Song, S. (2017). Enabling scalability-sensitive speculative parallelization for FSM computations. *Proc. Int. Conf. Supercomput.*, 2.
- Qiu, Y., Zhou, G., Zhang, Y., and Cichocki, A. (2021). Canonical polyadic decomposition (CPD) of big tensors with low multilinear rank. *Multimed. Tools Appl.* 80, 22987–23007. doi:10.1007/s11042-020-08711-1
- Raghavachari, K., Trucks, G. W., Pople, J. A., and Head-Gordon, M. (1989). A fifth-order perturbation comparison of electron correlation theories. *Chem. Phys. Lett.* 157 (6), 479–483. doi:10.1016/s0009-2614(89)87395-6
- Riplinger, C., and Neese, F. (2013). An efficient and near linear scaling pair natural orbital based local coupled cluster method. *J. Chem. Phys.* 138, 034106. doi:10.1063/1.4773581
- Riplinger, C., Sandhoefer, B., Hansen, A., and Neese, F. (2013). Natural triple excitations in local coupled cluster calculations with pair natural orbitals. *J. Chem. Phys.* 139, 134101. doi:10.1063/1.4821834
- Rolik, Z., Szegegy, L., Ladjanszki, I., Ladóczki, B., and Kállay, M. (2013). An efficient linear-scaling CCSD (T) method based on local natural orbitals. *J. Chem. Phys.* 139, 094105. doi:10.1063/1.4819401
- Russ, N. J., and Crawford, T. D. (2004). Local correlation in coupled cluster calculations of molecular response properties. *Phys. Lett.* 400, 104–111. doi:10.1016/j.cpllett.2004.10.083

- Sæbo, S., and Almlöf, J. (1989). Avoiding the integral storage bottleneck in LCAO calculations of electron correlation. *Chem. Phys. Lett.* 154, 83–89. doi:10.1016/0009-2614(89)87442-1
- Sæbo, S., and Pulay, P. (1985). Local configuration interaction: An efficient approach for larger molecules. *Chem. Phys. Lett.* 113, 13–18. doi:10.1016/0009-2614(85)85003-x
- Sæbo, S., and Pulay, P. (1993). Local treatment of electron correlation. *Annu. Rev. Phys. Chem.* 44, 213–236. doi:10.1146/annurev.pc.44.100193.001241
- Sæbo, S., and Pulay, P. (1988). The local correlation treatment. II. Implementation and tests. *J. Chem. Phys.* 88, 1884–1890. doi:10.1063/1.454111
- Saitow, M., Uemura, K., and Yanai, T. (2022). A local pair-natural orbital-based complete-active space perturbation theory using orthogonal localized virtual molecular orbitals. *J. Chem. Phys.* 157, 084101. doi:10.1063/5.0094777
- Schriber, J. B., and Evangelista, F. A. (2017). Adaptive configuration interaction for computing challenging electronic excited states with tunable accuracy. *J. Chem. Theory Comput.* 13, 5354–5366. doi:10.1021/acs.jctc.7b00725
- Schütz, M., Yang, J., Frederick, R., and Werner, H. J. (2013). The orbital-specific virtual local triples correction: OSV-L (t). *J. Chem. Phys.* 138, 054109. doi:10.1063/1.4789415
- Schwilk, M., Ma, Q., Köppl, C., and Werner, H. J. (2017). Scalable electron correlation methods. 3. Efficient and accurate parallel local coupled cluster with pair natural orbitals (PNO-LCCSD). *J. Chem. Theory Comput.* 13, 3650–3675. doi:10.1021/acs.jctc.7b00554
- Semidalas, E., and Martin, J. M. L. (2022). The MOBH35 metal-organic barrier heights reconsidered: Performance of local-orbital coupled cluster approaches in different static correlation regimes. *J. Chem. Theory Comput.* 18, 883–898. doi:10.1021/acs.jctc.1c01126
- Shang, H., Shen, L., Fan, Y., Xu, Z., Guo, C., Liu, J., et al. (2022). Large-Scale Simulation of Quantum Computational Chemistry on a New Sunway Supercomputer. *SC22: Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 1–14. doi:10.1109/SC41404.2022.00019
- Shao, Y., Gan, Z., Epifanovsky, E., Gilbert, A. T. B., Wormit, M., Kussmann, J., et al. (2015). Advances in molecular quantum chemistry contained in the Q-Chem 4 program package. *Mol. Phys.* 113, 184–215. doi:10.1080/00268976.2014.952696
- Sharapa, D. I., Genaev, A., Cavallo, L., and Minenkov, Y. (2019). A robust and cost-efficient scheme for accurate conformational energies of organic molecules. *ChemPhysChem* 20, 92–102. doi:10.1002/cphc.201801063
- Sho, S., and Odanaka, S. (2019). Parallel domain decomposition methods for a quantum-corrected drift-diffusion model for MOSFET devices. *Phys. Commun.* 237, 8–16. doi:10.1016/j.cpc.2018.10.029
- Simons, J., and Nichols, J. (1997). *Quantum mechanics in chemistry*.
- Sitkiewicz, S. P., Rodriguez-Mayorga, M., Luis, J. M., and Matito, E. (2019). Partition of optical properties into orbital contributions. *Phys. Chem. Chem. Phys.* 21, 15380–15391. doi:10.1039/c9cp02662b
- Sparta, M., Retegan, M., Pinski, P., Riplinger, C., Becker, U., and Neese, F. (2017). Multilevel approaches within the local pair natural orbital framework. *J. Chem. Theory Comput.* 13, 3198–3207. doi:10.1021/acs.jctc.7b00260
- Stegeman, A. (2006). Degeneracy in Candecomp/Parafac explained for $p \times p \times 2$ arrays of rank $p + 1$ or higher. *Psychometrika* 71, 483–501.
- Stoychev, G. L., Auer, A. A., Gauss, J., and Neese, F. (2021). DLPNO-MP2 second derivatives for the computation of polarizabilities and NMR shieldings. *J. Chem. Phys.* 154, 164110. doi:10.1063/5.0047125
- Su, H. C., Jiang, H., and Zhang, B. (2007). Synchronization on Speculative Parallelization of Many-Particle Collision Simulation. *World Congr. Eng. Comput. Sci.*
- Subotnik, J. E., and Head-Gordon, M. (2005). A local correlation model that yields intrinsically smooth potential-energy surfaces. *J. Chem. Phys.* 123, 064108. doi:10.1063/1.2000252
- Surján, P. R. (1999). “An introduction to the theory of geminals,” in *Correlation and localization*, 63–88.
- Szabo, A., and Ostlund, N. S. (2012). *Modern quantum chemistry: Introduction to advanced electronic structure theory*.
- Szabó, P. B., Csóka, J., Kállay, M., and Nagy, P. R. (2021). Linear-Scaling open-shell MP2 approach Algorithm benchmarks and large-scale applications. *J. Chem. Theory Comput.* 17, 2886–2905. doi:10.1021/acs.jctc.1c00093
- Tew, D. P., Klopper, W., and Helgaker, T. (2007). Electron correlation: The many-body problem at the heart of chemistry. *J. Comput. Chem.* 28, 1307–1320. doi:10.1002/jcc.20581
- Tew, D. P. (2019). Principal domains in local correlation theory. *J. Chem. Theory Comput.* 15, 6597–6606. doi:10.1021/acs.jctc.9b00619
- Thiel, W. (2014). Semiempirical quantum-chemical methods. *Rev. Comput. Mol. Sci.* 4, 145–157. doi:10.1002/wcms.1161
- Titov, A. V., Ufimtsev, I. S., Luehr, N., and Martinez, T. J. (2013). Generating efficient quantum chemistry codes for novel architectures. *J. Chem. Theory Comput.* 9 (1), 213–221. doi:10.1021/ct300321a
- Tucker, L. R. (1966). Some mathematical notes on three-mode factor analysis. *Psychometrika* 31, 279–311. doi:10.1007/BF02289464
- Unke, O., Bogojeski, M., Gastegger, M., Geiger, M., Smidt, T., and Müller, K.-R. (2021). SE(3)-equivariant prediction of molecular wavefunctions and electronic densities. *Adv. Neural Inf. Process. Syst.* 34 14434–14447.
- Vahtas, O., Almlöf, J., and Feyereisen, M. W. (1993). Integral approximations for LCAO-SCF calculations. *Chem. Phys. Lett.* 213, 514–518. doi:10.1016/0009-2614(93)89151-7
- Valiev, M., Bylaska, E. J., Govind, N., Kowalski, K., Straatsma, T. P., Van Dam Hubertus, J. J., et al. (2010). NWChem: A comprehensive and scalable open-source solution for large scale molecular simulations. *Phys. Commun.* 181, 1477–1489. doi:10.1016/j.cpc.2010.04.018
- Vannieuwenhoven, N., Meerbergen, K., and Vandebril, R. (2015). Computing the gradient in optimization algorithms for the CP decomposition in constant memory through tensor blocking. *SIAM J. Sci. Comput.* 37 (3), C415–C438. doi:10.1137/14097968x
- Wang, B., Yang, K. R., Xu, X., Isegawa, M., Leverentz, H. R., and Truhlar, D. G. (2014). Quantum mechanical fragment methods based on partitioning atoms or partitioning coordinates. *Accounts Chem. Res.* 47, 2731–2738. doi:10.1021/ar500068a
- Wang, H., Neese, C. F., Morong, C. P., Kleshcheva, M., and Oka, T. (2013) High-Resolution near-infrared spectroscopy of and its deuterated isotopologues *J. Phys. Chem. A* 117 9908–9918.
- Wang, Y. M., Hättig, C., Reine, S., Valeev, E., Kjærgaard, T., and Kristensen, K. (2016). Explicitly correlated second-order Møller-Plesset perturbation theory in a Divide-Expand-Consolidate (DEC) context. *J. Chem. Phys.* 144, 204112. doi:10.1063/1.4951696
- Wang, Y., Ni, Z., Neese, F., Li, W., Guo, Y., and Li, S. (2022). Cluster-in-Molecule method combined with the domain-based local pair natural orbital approach for electron correlation calculations of periodic systems. *J. Chem. Theory Comput.* 18, 6510–6521. doi:10.1021/acs.jctc.2c00412
- Werner, H.-J. (1995). “Problem decomposition in quantum chemistry,” in *Domainbased parallelism and problem decomposition methods in computational science and engineering* (SIAM), 239–261. doi:10.1137/1.9781611971507.ch14
- Westermayr, J., Gastegger, M., Schütt, K. T., and Reinhard, J. (2021). Perspective on integrating machine learning into computational chemistry and materials science. *J. Chem. Phys.* 154, 230903. doi:10.1063/5.0047760
- Woolley, R. G., and Sutcliffe, B. T. (1977). Molecular structure and the Born-Oppenheimer approximation. *Phys. Lett.* 45, 393–398. doi:10.1016/0009-2614(77)80298-4
- Xie, Z. Y., Jiang, H. C., Chen, Q. N., Weng, Z. Y., and Xiang, T. (2009). Second renormalization of tensor-network states. *Phys. Rev. Lett.* 103, 160601. doi:10.1103/physrevlett.103.160601
- Yang, J., Chan, G., Frederick, R., Schütz, M., and Werner, H.-J. (2012). The orbital-specific-virtual local coupled cluster singles and doubles method. *J. Chem. Phys.* 136, 144105. doi:10.1063/1.3696963
- Yang, J., Kurashige, Y., Manby, F. R., and Chan, G. K. L. (2011). Tensor factorizations of local second-order Møller-Plesset theory. *J. Chem. Phys.* 134, 044123. doi:10.1063/1.3528935
- Yates, K. (2012). *Hückel molecular orbital theory*.
- Zhang, I. Y., and Grüneis, A. (2019). Coupled cluster theory in materials science. *Front. Mater.* 6. doi:10.3389/fmats.2019.00123
- Zhang, Q., Dwyer, T. J., Tsui, V., Case, D. A., Cho, J., Dervan, P. B., et al. (2004). NMR structure of a cyclic polyamide-DNA complex. *J. Am. Chem. Soc.* 126, 7958–7966. doi:10.1021/ja0373622
- Ziólkowski, M., Jansik, B., Kjærgaard, T., and Jørgensen, P. (2010). Linear scaling coupled cluster method with correlation energy based error control. *J. Chem. Phys.* 133, 014107. doi:10.1063/1.3456535