

Orbital Optimization of Large Active Spaces via AI-Accelerators

Örs Legeza,* Andor Menczer, Ádám Ganyecz, Miklós Antal Werner, Kornél Kapás, Jeff Hammond,* Sotiris S. Xantheas,* Martin Ganahl,* and Frank Neese*

Cite This: *J. Chem. Theory Comput.* 2025, 21, 6545–6558

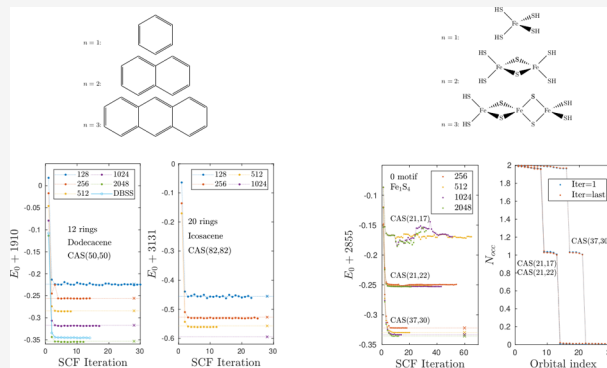
Read Online

ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: We present an efficient orbital optimization procedure that combines the highly GPU accelerated, spin-adapted density matrix renormalization group (DMRG) method with the complete active space self-consistent field (CAS-SCF) approach for quantum chemistry implemented in the ORCA program package. Leveraging the computational power of the latest generation of Nvidia GPU hardware, we perform CAS-SCF based orbital optimizations for unprecedented CAS sizes of up to 82 electrons in 82 orbitals [CAS(82,82)] in molecular systems comprising active space sizes of hundreds of electrons in thousands of orbitals. For both the NVIDIA DGX-A100 and DGX-H100 hardware, we provide a detailed scaling and error analysis of our DMRG-SCF approach for benchmark systems consisting of polycyclic aromatic hydrocarbons and iron–sulfur complexes of varying sizes. Our efforts demonstrate for the first time that highly accurate DMRG calculations at large bond dimensions are critical for obtaining reliably converged CAS-SCF energies. For the more challenging iron–sulfur benchmark systems, we furthermore find the optimized orbitals of a converged CAS-SCF calculation to depend more sensitively on the DMRG parameters than those for the polycyclic aromatic hydrocarbons. The ability to obtain converged CAS-SCF energies and orbitals for active spaces of such large sizes within days reduces the challenges of including the appropriate orbitals into the CAS or selecting the correct minimal CAS, and may open up entirely new avenues for tackling strongly correlated molecular systems.



1. INTRODUCTION

Since their early days more than 50 years ago, computational quantum chemistry methods have nowadays become a powerful and indispensable tool assisting scientists unraveling the mysteries of the microscopic world of atoms and molecules. Through substantial advances in algorithmic and hardware development, today's quantum chemistry algorithms are able to accurately model molecules and materials with hundreds or thousands of electrons and orbitals, often requiring only moderate computational resources. A key element of success in this context has been the combination of groundbreaking, innovative new algorithms coupled with novel hardware advances. Today, strongly correlated systems i.e. multireference systems, have become the main frontier in computational quantum chemistry. While modern quantum chemistry methods are nowadays routinely applied to compute properties of large, single-reference molecules,^{1–13} multireference quantum systems still pose significant challenges for computational methods, often to the point that a reliable description of molecular properties remains beyond reach.^{14–17} Reliable ab initio computational prediction of properties of such materials and molecules is considered the holy grail of quantum chemistry, and lies at the center of contemporary

research in classical and quantum approaches to computational quantum chemistry.^{16,18–20}

In the context of multireference quantum chemistry, tensor networks methods^{21,22} are a particularly successful class of classical algorithms^{23–25} able to accurately model quantum systems with strong, local correlations.^{20,26–34} Over the past decades they have emerged as the gold-standard method for strongly correlated quantum systems in low dimensions,^{35–40} and today often serve as benchmark tools against which novel computational approaches need to be tested.^{20,33} A key role is played in this context by the so-called density matrix renormalization group (DMRG) algorithm,²³ a variational optimization algorithm over the space of so-called matrix product states, a simple but very powerful tensor network. In this work we employ a highly GPU-efficient implementation of the DMRG algorithm as a complete active space (CAS) solver

Received: April 10, 2025

Revised: May 23, 2025

Accepted: May 27, 2025

Published: June 13, 2025



within a complete active space self-consistent field (CAS-SCF) approach^{41–43} for both closed and open shell systems. We demonstrate that by replacing the conventional CI solver with the massively parallel hybrid CPU-GPU DMRG method^{44,45} the quality of the resulting DMRG-SCF^{28,46–49} framework can be raised to a significantly higher level, leading to a well controlled and stable convergence of the CAS-SCF method, while simultaneously targeting active spaces of unprecedented size and complexity. This enables us to apply the CAS-SCF method to complex chemical problems far beyond the ones currently possible, completed in a fraction of the time and cost compared to alternative approaches. Moreover, the accessible large bond dimensions and active space sizes make it possible to carry out a detailed error analysis of the DMRG-SCF procedure as a function of CAS size, bond-dimension, accuracy settings of subroutines, and other tunable hyper-parameters.

We utilize a message passing interface (MPI) based massively parallel implementation of the SCF procedure via the ORCA program package,¹ which is coupled efficiently with our hybrid CPU-GPU DMRG implementation.^{44,50,51} We focus on results obtained on a single node supplied with eight NVIDIA A100 or H100 graphics processing units leading to a factor of up to 20–70 speedup compared to a 128 CPU core single node implementation.

2. METHODS

In this section, we briefly overview the underlying theory of the applied DMRG-SCF method focusing on the various error sources, technical aspects related to its massive parallelization, and background on the molecular systems it was applied.

2.1. Complete Active Space SCF. Within the Born–Oppenheimer approximation, the nonrelativistic quantum chemical many-body Hamiltonian is given (in atomic units) by

$$H = \sum_{pq} h_{pq} \hat{e}_{pq} + \sum_{pqrs} V_{pqrs} \hat{e}_{pqrs}, \quad \text{with} \quad (1)$$

$$\hat{e}_{pq} \equiv \sum_{\sigma} a_{p\sigma}^{\dagger} a_{q\sigma} \quad \text{and} \quad (2)$$

$$\hat{e}_{pqrs} \equiv \sum_{\sigma\sigma'} a_{p\sigma}^{\dagger} a_{q\sigma'}^{\dagger} a_{r\sigma} a_{s\sigma'} \quad (3)$$

Here, $h_{pq} \equiv (p|\hat{h}|q)$ and $V_{pqrs} \equiv (pq|rs)$ are the one- and two-electron integrals

$$(p|\hat{h}|q) = \int d\mathbf{r} \phi_p^*(\mathbf{r}) \hat{h}(\mathbf{r}) \phi_q(\mathbf{r}) \quad (4)$$

$$(pq|rs) = \int d\mathbf{r}_1 d\mathbf{r}_2 \frac{\phi_p^*(\mathbf{r}_1) \phi_q(\mathbf{r}_1) \phi_r^*(\mathbf{r}_2) \phi_s(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} \quad (5)$$

with $\hat{h}(r)$ describing both the kinetic energy of the electrons as well as an external one-body potential, including the potential energy of the charged nuclei. $\phi_p(\mathbf{r})$, $p = 1, \dots, N$, are a set of N orthonormal single electron wave functions, and $a_{p\sigma}^{\dagger}$ and $a_{p\sigma}$ are Fermionic creation and annihilation operators of electrons with spin σ in orbital p .

The Multi Configuration SCF (MCSCF) method aims at incorporating static electronic correlations into a Hartree–Fock or DFT calculation both by variationally incorporating multiple Slater determinants into a variationally optimized wave function and by optimizing the single-particle orbitals $\phi_p(\mathbf{r})$ of the Slater determinants at the same time. The orbitals

$p = 1, \dots, N$ are grouped into three distinct, nonoverlapping sets of core (doubly occupied), active and virtual (unoccupied) orbitals. The active space (AS) orbitals that can have occupations of 0, 1, or 2 are chosen with the intent to capture the most prominent many-body effects of the problem at hand. Typically, the AS includes part of or the entirety of valence orbitals. Their corresponding HF single-particle energies are typically located around the Fermi level of the HF solution. In the complete active space SCF (CAS-SCF) approach, the variational optimization is carried out over all possible Slater determinants that can be obtained from redistributing electronic occupations within the AS only, while in the SCF, simultaneously, all single-electron orbitals are also unitarily varied. Here, we remark, that ORCA is using configuration state functions (CSFs) not determinants but in this context, this is an insignificant detail.

Within a given basis set ϕ_p , a general CAS wave function has the form

$$|\Psi\rangle = \sum_{I \in AS} c_I |\Phi_I\rangle \quad (6)$$

where the sum over $I \in AS$ runs over the Slater determinants $|\Phi_I\rangle$ obtained from all possible electron configurations within the active space, and c_I are the expansion coefficients of each Slater determinant. Additionally, in the CAS-SCF approach, all single-particle orbitals ϕ_p are unitarily optimized using a single-body rotation matrix $U \equiv e^{-\hat{\kappa}}$ with $\hat{\kappa} \equiv \sum_{p>q} \kappa_{pq} (\hat{e}_{pq} - \hat{e}_{qp})$ on top of the expansion coefficients c_I in order to minimize the energy

$$E = \langle \Psi | U^{\dagger} H U | \Psi \rangle = \sum_{pq} h_{pq} \langle \Psi | U^{\dagger} \hat{e}_{pq} U | \Psi \rangle + \sum_{pqrs} V_{pqrs} \langle \Psi | U^{\dagger} \hat{e}_{pqrs} U | \Psi \rangle \quad (7)$$

from which the optimal parameters are obtained as

$$U_{\text{opt}}, |\Psi_{\text{opt}}\rangle = \arg \min_{U, c_I} \langle \Psi | U^{\dagger} H U | \Psi \rangle \quad (8)$$

Instead of optimizing both the CI coefficients c_I and the unitary U simultaneously, practical CAS-SCF implementations employ an alternating scheme of a CAS-CI step optimizing the c_I coefficients at fixed U followed by an orbital optimization of the unitary U at fixed c_I .⁵⁹ In this work we employ the density matrix renormalization group (DMRG) method as the CAS-CI solver, and use ORCA to perform the orbital optimization step. The gradient of the energy with respect to the orbital rotation matrix can be obtained from the one and two electron reduced density matrices

$$\tilde{\gamma}_{pq} = \langle \Psi | \hat{e}_{pq} | \Psi \rangle \quad (9)$$

$$\tilde{\Gamma}_{pqrs} = \langle \Psi | \hat{e}_{pqrs} | \Psi \rangle \quad (10)$$

which can be computed efficiently from the CAS-CI wave function obtained via DMRG.⁶⁰ The iteration is stopped once the gradient for the orbital rotations is sufficiently small. In general, the threshold depends on the chosen convergence tightness. In the CASSCF module of the ORCA program for “TightSCF” it is $\times 10^{-3}$ on the gradient and at the same time $\times 10^{-7}$ on the energy change.

Regarding the superiority of CAS-SCF over CAS-CI we recall that variationally optimizing the CI coefficients and orbital rotations at the same time offers a number of

advantages. First of all, derivatives of variational wave functions are far easier to calculate than for nonvariational wave functions. Second, enforcing orbital optimization (usually) leads to a very well-defined wave function that eliminates possible arbitrariness. Therefore, the results of a CAS-SCF calculation are typically of much better quality compared to those of a CAS-CI approach.

An important ingredient, however, for obtaining qualitatively and quantitatively accurate results with CAS-CI is the inclusion of the appropriate orbitals in the AS. This is far from a trivial task and usually requires the intuition of a skilled quantum chemist familiar with the intricacies of the problem at hand. We do, however, mention here recent efforts for both the automatic selection of active spaces^{61,62} and the development of powerful tools that assist with the selection of the initial active space orbitals (e.g., AVAS⁶³). CAS-SCF on the other hand is computationally much more challenging but is more resilient to picking the wrong orbitals in the active space since the orbital update may mix in orbital components that have originally been left out. Most definitely, a “self-repairing” active space is rare. If one chooses a bad starting point, which is really easy, one most certainly ends up with nonsensical results. It is rather an art form to make sure that the initially picked orbitals stay in the active space. ORCA has a lot of infrastructure to try to ensure that, but the variational principle can be merciless. Ultimately, the division of orbitals into an inactive, active, and virtual space is very non-natural, and mathematics “rewards” us for insisting on this choice with convergence problems of all kinds.

The two main bottlenecks of CAS-SCF are the computation of the CAS-CI wave function at fixed U , and the computation of one and two-body reduced density matrices $\tilde{\gamma}_{pq}$ and $\tilde{\Gamma}_{pqrs}$. The main contribution of this work is the application of our novel GPU-accelerated approach for computing the reduced density matrices, in combination with a GPU-accelerated DMRG CAS-CI solver. Leveraging the latest Nvidia supercomputing hardware, and ORCA's efficient parallelization techniques for orbital optimization, this enables us to run CAS-SCF calculations for more than 80 electrons in 80 orbitals within a day or so, far beyond any other current state-of-the-art approach. The ability to use such large active spaces with CAS-SCF has the potential to substantially reduce the burden of researchers to pick the right orbitals for their problem, and may help to attack challenging strongly correlated molecular systems that are beyond the scope of conventional methods.

2.2. Tensor Network States and Density Matrix Renormalization Group. In the following, we consider a system of N Hilbert spaces $|i_p\rangle$, $p = 1, \dots, N$, with $i_p = 0, 1, 2, 3$ corresponding to states with no electron, one down electron, one up electron, and one up and one down electron, respectively. In the context of this manuscript, the states $|i_p\rangle$ correspond to the single-electron orbitals ϕ_p obtained from a Hartree–Fock or DFT calculation. For example, a restricted Hartree–Fock wave function is a simple product state obtained by filling all orbitals $|i_p\rangle$ with low (HF orbital) energies ϵ_p , $p < p^*$ with two electrons, and keeping all other orbitals empty:

$$|\Psi\rangle = |3\rangle_{i_1} \cdots |3\rangle_{i_{p^*}} |0\rangle_{i_{p^*+1}} \cdots |0\rangle_{i_N} \quad (11)$$

For a given set of inactive but occupied orbitals $|i_{o_k}\rangle$, $k = 1, \dots, N_{IA}\rangle$,

$$|\Psi\rangle = \left(\bigotimes_{k=1}^{N_{IA}} |3\rangle_{o_k} \right) \sum_{\{i_{p_k}\}} \psi_{i_{p_1} \dots i_{p_{N_A}}} |i_{p_1} \dots i_{p_{N_A}}\rangle \left(\bigotimes_{k=1}^{N_{UO}} |0\rangle_{q_k} \right) \quad (12)$$

For the sake of brevity, we will relabel all orbitals such that the first N_A orbitals $|i_k\rangle$, $k = 1, \dots, N_A$ correspond to the active space orbitals, and omit the explicit reference to the trivial tensor products over inactive and empty spaces, i.e.

$$|\Psi\rangle = \sum_{\{i_1 \dots i_{N_A}\}} \psi_{i_1 \dots i_{N_A}} |i_1 \dots i_{N_A}\rangle \quad (13)$$

Matrix product state (MPS) wave functions are a special class of many body wave functions parametrized by a set of N_A tensors $A_{\alpha_{p-1}\alpha_p}^{i_p}$ of shape $(4, D_{p-1}, D_p)$, such that the many body wave function eq 13 takes on the special form

$$\psi_{i_1 \dots i_{N_A}} = \sum_{\{\alpha_p\}} [A_1]_{1\alpha_1}^{i_1} [A_2]_{\alpha_1\alpha_2}^{i_2} \cdots [A_{N_A}]_{\alpha_{N_A-1}1}^{i_{N_A}} \quad (14)$$

For a fixed set of orbital indices i_1, \dots, i_{N_A} , the amplitude $\psi_{i_1 \dots i_{N_A}}$ reduces to a sequence of matrix products, giving the ansatz its name.³⁶ The largest of the tensor dimension $D \equiv \max_k D_k$ is called the bond dimension of the MPS. The expressiveness of the MPS wave function increases with increasing bond dimension D . MPS wave functions with small values of D are particularly well suited to represent wave functions with low entanglement. Note that any state eq 13 can be cast into an MPS form, though the bond dimension required for a faithful representation might be exponentially large in N_A . In this work we employ the density matrix renormalization group (DMRG) method as a CAS-CI solver.^{22,23} In its modern formulation, the DMRG can be understood as a variational minimization over the space of MPS.⁶⁴ Specifically, we use it to variationally determine an approximation of the ground state of the CAS-CI Hamiltonian

$$E_{\text{opt}} = \min_{|\Psi\rangle} \frac{\langle \Psi | H | \Psi \rangle}{\langle \Psi | \Psi \rangle} \quad (15)$$

From the optimized MPS wave function $|\Psi_{\text{opt}}\rangle$, the one- and two-electron reduced density matrices eqs 9 and 10 can be efficiently computed using standard techniques.⁶⁰ One of our main contributions in this work is the utilization of our GPU accelerated method to compute $\tilde{\gamma}_{pq}$ and $\tilde{\Gamma}_{pqrs}$ with substantial speed ups compared to existing implementations, which we put to use for running CAS-SCF calculations at large bond dimensions to obtain stable convergence with high accuracy, and at unprecedented AS sizes of up to 82 electrons in 82 orbitals in about a day or so.

2.3. Conservation of the Total Spin in the DMRG-SCF Procedure. Our hybrid CPU-multiGPU DMRG code can handle both Abelian and non-Abelian quantum numbers,^{44,50} thus a given state with total spin can also be targeted via the DMRG-SCF protocol. In this case, the DMRG bond dimension D refers to the number of multiplets kept to represent the left or the right DMRG block.^{65–70} This in general, leads to a factor of 2 to three reduction in D for a singlet ($S_{\text{tot}} = 0$) state to reach similar accuracy as obtained with a strict U(1) implementation, while for states with higher total spin, this reduction can be even larger. This reduction provides a considerable speedup in the DMRG part of the algorithm. In addition, since no spin contamination can happen, the orbital optimization procedure becomes more

stable and robust. In our implementation, the SU(2) related Clebsch-Gordan layer⁷⁰ is fully decoupled from the MPS layer, thus there is no overhead in the matrix and tensor algebra when massive parallelization is utilized.^{44,71} In the rest of the paper, all bond dimensions D are reported as SU(2) multiplets, with the corresponding U(1) bond dimensions, \tilde{D} , indicated separately where applicable.

As our non-Abelian DMRG implementation treats the Hilbert-space at the multiplet level, the symmetry-coupled form of the one and two electron reduced density matrices are determined, that are built based on the spin- $\frac{1}{2}$ Wigner-Eckart tensor operators formed from creation and annihilation operators:

$$f_p^\dagger = (a_{p\uparrow}^\dagger, a_{p\downarrow}^\dagger), f_p = (a_{p\downarrow}, -a_{p\uparrow}) \quad (16)$$

The spin-coupled form of the one electron reduced density matrix is simply the expectation value of the spin-0 combination of the $f_p^\dagger f_q$ product,

$$\gamma_{pq} = \langle \Psi | [f_p^\dagger f_q]_0 | \Psi \rangle = \langle \Psi | \frac{1}{\sqrt{2}} (f_{p,1}^\dagger f_{q,2} - f_{p,2}^\dagger f_{q,1}) | \Psi \rangle = -\frac{1}{\sqrt{2}} \tilde{\gamma}_{pq} \quad (17)$$

Here and in the following $[AB]_J$ denotes the spin- J tensor operator formed from the product of tensor operators A and B , which is calculated as the linear combination of simple operator products multiplied by the appropriate Clebsch-Gordan coefficients like in eq 17. In case of the two electron reduced density matrix, two channels are formed depending on the total spin created by the $f_p^\dagger f_q$ product,

$$\begin{aligned} \Gamma_{pqrs}^{(0)} &= \langle \Psi | [[f_p^\dagger f_q]_0 [f_r^\dagger f_s]_0] | \Psi \rangle \\ \Gamma_{pqrs}^{(1)} &= \langle \Psi | [[f_p^\dagger f_q]_1 [f_r^\dagger f_s]_1] | \Psi \rangle \end{aligned} \quad (18)$$

and from these the original (eq 10) form of the density matrix is simply

$$\tilde{\Gamma}_{pqrs} = 2\sqrt{3} \Gamma_{pqrs}^{(1)} - \Gamma_{pqrs}^{(0)} \quad (19)$$

Having the $\tilde{\gamma}_{pq}$ and $\tilde{\Gamma}_{pqrs}$ density matrices at hand, CASSCF optimization of ORCA can be called in the same way as in the U(1) symmetric case.

2.4. Error Sources. In the DMRG-SCF method there are various error sources that can accumulate during the procedure and these can even counteract with each other leading to unstable convergence. The CASSCF theory is based on the full-CI, i.e., the exact solution of the CAS, thus convergence is controlled and determined solely by the parameters of the applied gradient descend methods. In the ORCA program package, there are several options to achieve fastest convergence.¹ For the CASSCF procedure the default “TightSCF” setting corresponds to $\epsilon_{\text{ECT}} = \times 10^{-7}$, and $\epsilon_{\text{OGC}} = \times 10^{-3}$.

In contrast to these, the DMRG-SCF procedure brings in another error source since the CAS solution is only approximate and determined mainly by the DMRG bond dimension, D .^{25,35,72,73} For the ab initio DMRG variant, there are also further factors that influence the convergence, thus various algorithmic solutions developed in the past two decades^{20,29,32,34} must be utilized to reach the a priori set error margin with fastest convergence rate. Only if the error of

the DMRG solution is kept below the error settings of the SCF procedure could the DMRG-SCF procedure lead to stable convergence. This, however, usually requires large D values making the DMRG-SCF procedure impractically too long when large CAS spaces are considered without massive parallelization as has been faced in most of the previous attempts. In addition, when only an approximate CAS solution is used instead of the full-CI wave function, then technically active-active rotations should also be utilized. This, however, is usually neglected as they lead to numerical instabilities as the full-CI limit is approached. Therefore, in the current work, we have not utilized active-active rotations.

2.5. Parallelization and Technical Aspects. The wall time of the DMRG-SCF procedure comes from two main components, i.e., the total time spent on the SCF optimization procedure and on the solution of the selected CAS via the DMRG method. Since both scale with system size as N_A^4 where N_A stands for the number of orbitals a massive parallelization is mandatory to keep computational time feasible.

On the one hand, the massively parallel implementation of the SCF module based on message passing interface (MPI) in the ORCA program package can scale up to twenty-four CPU cores easily.¹ On the other hand, a massively parallel hybrid CPU-multiGPU DMRG method,^{44,50} introduced by some of us recently, can reach even a quarter petaflops on a single node by utilizing NVIDIA AI accelerators.⁸⁰ Consequently, the cubic scaling of the wall time with bond dimension in the DMRG algorithm can be reduced to linear scaling by employing parallelization over increasing number of GPUs for a broad range of D values.^{44,50} Similarly, the calculation of the one- and two-particle reduced density matrices has been adapted to multi-GPU systems,⁸¹ showing almost perfect scaling behavior with increasing system size, DMRG bond dimension and number of GPUs. This provides us with the ability to use much larger values of bond dimension D than previously achieved, at a marginal increase in computation time. The opportunity to use such large D values has two key advantages. First, the number of DMRG sweeps to achieve a desired accuracy can be reduced significantly. Since the bulk of the runtime of a DMRG-SCF iteration step is spent in the DMRG-CI solver, this translates into a major speedup for the DMRG-SCF procedure itself. Second, we empirically find that the higher accuracy resulting from large bond dimension calculations significantly improves the convergence behavior of the SCF procedure, resulting in a drastic reduction of required number of SCF iteration steps to achieve convergence. These two key properties allow us to obtain converged DMRG-SCF calculations for CAS sizes far beyond what has so far been reported in the literature, within run times ranging from minutes to ≈ 30 h.

To utilize all benefits of parallelization we perform calculations on a dual AMD EPYC 7702 CPUs with 2×64 cores combined with eight NVIDIA A100-SXM4-40GB devices and on a dual Intel Xeon Platinum 8481C CPUs with 2×52 cores combined with eight NVIDIA H100-HBM3-80GB devices. For the ORCA suite 24–48 MPI processes are allocated, depending on the estimated memory requirement of a single process in the SCF module, while for the DMRG part all available number of threads and GPUs are utilized.

2.6. Numerical Procedure. Our DMRG-SCF method relies on our hybrid CPU-GPU DMRG implementation interfaced to the ORCA program package. For a given CAS

we first perform a few DMRG optimization steps with low bond dimension, $D = \{16, 32, 64\}$ to obtain an optimal orbital ordering.^{82,83} Next, large scale DMRG-SCF is applied by using either fixed $D \in \{128, \dots, 4096\}$, or fixed truncation error ϵ_{TR} using the dynamic block state selection (DBSS) approach.²⁵ Note that for reliable convergence of the SCF procedure, ϵ_{TR} should be at least an order of magnitude smaller than ORCA's ϵ_{OGC} . The DMRG procedure is terminated when the energy difference between three subsequent sweeps falls below an a-priori set error margin of $\epsilon_{\text{sweep}} = 10^{-3}$. The residual error of the Lánczos and Davidson diagonalization methods has been set to 10^{-7} . When only U(1) symmetries are enforced, we use the dynamically extended active space (DEAS) procedure⁸² to speed up convergence. For SU(2) symmetries, the DEAS method is currently being implemented and is expected to lead to a significant reduction in the required number of sweeps, and hence an additional speedup of the DMRG-SCF approach.

In the next section we present results obtained both at fixed bond dimension, as well as at fixed truncation error $\epsilon_{\text{TR}} = 10^{-4}$, using the DBSS method. The number of DMRG sweeps needed to reach convergence typically varies between 10 and 30. SU(2) \times U(1) symmetry is enforced exactly using the spin-adapted DMRG approach, and one- and two-particle reduced density matrices are converted to the U(1) \times U(1) representation before passing them to the ORCA program package as discussed in Section 2.3. For sake of completeness, we also compare results obtained via the nonspin adapted workflow for the Heptacene test system.

3. PREPARATION OF BENCHMARK SYSTEMS

As a first group of test systems we focus on polycyclic aromatic hydrocarbons (PAHs), which have been the subject of extensive previous numerical studies.^{84–90} Specifically, we use linear chains of aromatic rings (see Figure 1; benzene,




Figure 1. Smallest molecules of the PAH series for ring numbers $n = 1, 2$, and 3 .

naphthalene, anthracene, tetracene, etc., $C_{4n+2}H_{2n+4}$) of up to 20 rings. The geometries were optimized with the B3LYP functional^{91,92} and cc-pVDZ basis set.^{93,94} From the corresponding orbitals, a subset containing $4n + 2$ orbitals was chosen manually, with n the number of elementary building blocks (rings). The subset was chosen such as to represent the delocalized π -system over the entire PAH. The selected orbitals formed the initial active space for subsequent

DMRG-SCF calculations. PAH molecules are typically single-reference and hence serve as ideal benchmark systems for method development.

The second group of test systems contains iron–sulfur clusters, which are found in various proteins like ferredoxins, as well as various hydrogenases, dehydrogenases, reductases, and nitrogenases.⁹⁵ The simplest iron–sulfur system is the $2\text{Fe}-2\text{S}$ cluster, which has a planar Fe_2S_2 core and is usually coordinated by four ligands leading to a tetrahedral structure around the iron atoms.

A scalable test system can be made by repeating the Fe_2S_2 motif and closing with $-\text{SH}$ groups. We created the series of $\text{Fe}_n\text{S}_{2n-2}(\text{SH})_4^{n-}$ clusters with $1 \leq n \leq 3$. The first three elements are shown in Figure 2. The first species, $\text{Fe}(\text{SH})_4^-$,




Figure 2. Fe_2S_2 cluster series for $n = 1, 2$, and 3 .

corresponding to $n = 1$ does not contain the Fe_2S_2 structure. For $n = 2$ we have the Fe_2S_2 motif ending with SH ligands. For $n > 2$ the Fe_2S_2 blocks follow each other while the plane is perpendicular to the previous one. In the oxidized state, the Fe ions are in the oxidation state of +3. In a tetrahedral coordination environment, the local spins located in the iron d-based molecular orbitals are all aligned parallel, leading to local $S(\text{Fe}) = 5/2$ fragments. The different local fragments will couple antiferromagnetically to lead to overall low spin states with total spin zero, e.g. $S = 0$ for the dimeric species. However, it should clearly be recognized that this is not a closed shell state but an antiferromagnetic singlet with ten unpaired electrons. Provided that localized orbitals are employed that are properly aligned to belong to the iron ions on the right- and left-hand side, a single configuration state function (CSF) dominates the CAS(10,10) wave function.⁹⁶ This single CSF has a weight of 90%. However, in terms of individual Slater determinants the compactness of this representation is lost and there are 252 Slater determinants belonging to the leading CSF that each show a small CI coefficient on the order of 0.02–0.03. In this work, we demonstrate the power of the DMRG implementation by going to much larger active spaces than the minimal CAS(10,10) that was described above.

Geometries in the Fe_2S_2 series were also optimized with B3LYP functional^{91,92} and cc-pVDZ basis set.^{93,94} Orbitals were selected from quasi-restricted orbitals (QRO). Three different active spaces were defined, labeled as ‘A’, ‘B’, and ‘C’, and their composition is presented in Table 1. ‘A’ contains the 3d, 4s and 4p orbitals of iron, and the 3p orbitals of the sulfur

Table 1. Breakdown of the Active Space of the $\text{Fe}_n\text{S}_{2n-2}(\text{SH})_4^{n-}$ Series^a

species	orbital type	active space	number	electron	orbital
Fe	3s	C	n	2	1
Fe	3p	C	n	6	3
Fe	3d	A,B,C	n	6	5
Fe	4s	A,B,C	n	2	1
Fe	4p	A,B,C	n	0	3
Fe	4d	B,C	n	0	5
S (bridging)	3s	C	$2n - 2$	2	1
S (bridging)	3p	A,B,C	$2n - 2$	4	3
S (terminal)	3s	C	4	2	1
S (terminal)	3p	A,B,C	4	3	2
electron		A,B,C	n	1	0

^aFor the sake of completeness, the electron leading to the correct charge is also included in this list.

atoms, however the p orbitals of sulfurs in the terminal SH groups which make the S–H bonds are excluded. ‘B’ contains also the 4d orbitals of Fe, while ‘C’ includes 3s and 3p orbitals of Fe and 3s orbitals of S-s.

4. NUMERICAL RESULTS

In this section, we present results obtained by various DMRG settings for polycyclic aromatic hydrocarbons and for iron–sulfur clusters. For better readability at certain parts of our analysis we label bond dimension used in the DMRG-SCF procedure with D_{opt} and in DMRG-CI with D while in the figures such different labeling is not applied.

4.1. Polycyclic Aromatic Hydrocarbons (PAHs).

4.1.1. Convergence Analysis. In Figure 3 we show the convergence of the ground state energy E_0 , (offset by 1146 for convenience of plotting) of heptacene ($\text{C}_{30}\text{H}_{18}$) as a function of the SCF iteration steps using the ab initio DMRG as a CAS CI-solver. The left panel shows results obtained using U(1) symmetry only, results in the right panel were obtained using a fully SU(2) symmetric simulation. We show results for increasing U(1) and SU(2) bond dimensions \tilde{D} and D .

As expected, increasing values of \tilde{D} result in decreasing values for the energy in both cases. A key observation from the plots is a significant improvement of the SCF convergence with increasing DMRG accuracy (i.e., increasing bond dimension). For low values of $\tilde{D} = 256$ (left plot) and $D = 128$ (right plot) we observe nonmonotonic behavior of the energy as a function of the SCF iteration step, i.e. the energy between to subsequent SCF steps can increase. This behavior is particularly pronounced for small bond dimensions and during early SCF steps, but can appear as well at later steps and for moderately large bond dimensions (see for example $\tilde{D} = 1024$ (left panel)). For large bond dimensions, we observe fast and smooth convergence of the SCF procedure, with significantly reduced number of SCF iterations required to obtain a desired convergence level. We note that in addition, large values of \tilde{D} typically also lead to fewer DMRG-sweeps required to obtain convergence within the DMRG-CI solver. Consequently, we empirically find a computationally optimal solution regarding accuracy and wall time by setting the truncation error margin $\varepsilon_{\text{TR}} = 10^{-4}$ to be an order of magnitude smaller than ORCA's $\varepsilon_{\text{OGC}} = 10^{-3}$, and employing the DBSS method with a minimal bond dimension $\tilde{D}_{\text{min}} = 2048$ in the DMRG-CI solver.

Enforcing SU(2) spin symmetry leads to further improvements of the convergence behavior of the SCF procedure.

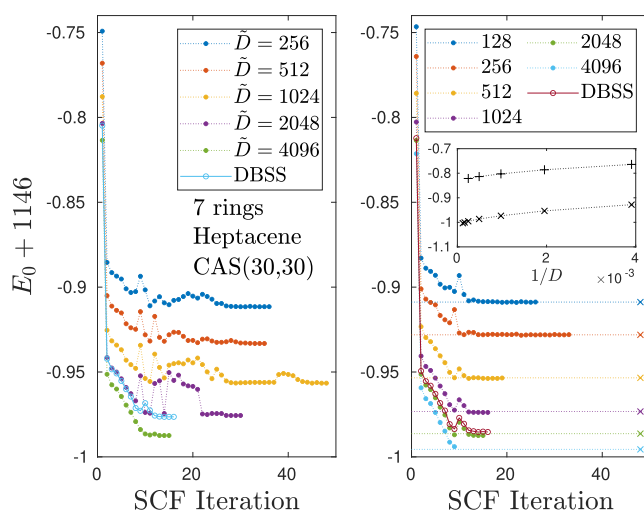


Figure 3. (left) Convergence of ground state energy, E_0 , shifted by 1146 for the heptacene on a CAS(30,30) as a function of the DMRG-SCF macro iterations for various fixed U(1) \tilde{D}_{opt} values and (right) via the spin adapted workflow for various fixed SU(2) bond dimensions, D_{opt} . Data points shown by open circle symbols were obtained via the DBSS procedure by setting $\tilde{D}_{\text{opt,min}} = 2048$ and $D_{\text{opt,min}} = 1024$, $\varepsilon_{\text{TR}} = 10^{-4}$ and maximum number of DMRG sweeps to ten. In the SCF module we set $\varepsilon_{\text{ECT}} = 10^{-7}$ and $\varepsilon_{\text{OGC}} = 10^{-3}$. Symbols X together with horizontal dotted lines stand for post DMRG-SCF energy values obtained for various D values using the basis optimized with $D_{\text{opt}} = 128$. The inset shows the $1/D$ scaling of the shifted ground state energy, using the original nonoptimized and the optimized basis with $D_{\text{opt}} = 128$, i.e. $E_0(D)$ and $E_0(D_{\text{opt}} = 128)$, respectively.

Here we note that even a larger fixed $D = 1024$ calculation can be stuck at a local minimum and several SCF steps were mandatory to reach final convergence. In contrast to this, using the DBSS procedure with $D_{\text{opt,min}} = 1024$ results in a stable and fast convergence (open circles in Figure 3, right panel).

4.1.2. Quality of the Optimized Basis. An important question in the context of DMRG-SCF is in how far the final one-particle basis at the end of the DMRG-SCF procedure depends on the bond dimension used in the DMRG-CI solver during the SCF iterations. The quality and robustness of the single-particle basis obtained at a low bond dimension can be checked by first converging the DMRG-SCF at low bond dimension, obtaining an energy $E_0(D_{\text{opt}} = 128)$, and then performing standard DMRG calculations with the converged basis at increasing bond dimensions, and comparing the resulting energies $E_0(D, D_{\text{opt}} = 128)$ with DMRG-SCF energies obtained using the same bond dimensions $D_{\text{opt}}' = D$. Matching energies indicate that dynamic correlations included via the SCF procedure are captured well already at low-bond-dimensions. On the other hand, a basis error mismatch $\delta_E(D, D_{\text{opt}}) \equiv E_0(D, D_{\text{opt}}) - E_0(D_{\text{opt}}' = D) \neq 0$ indicates missing dynamic correlations at low bond-dimension DMRG-SCF calculations. The results of these tests are shown in the right panel of Figure 3. The horizontal dashed lines show energies obtained by first converging DMRG-SCF at $D_{\text{opt}} = 128$, and then increasing the bond dimension at fixed basis to $D = 256, 512, 1024, 2048$, and 4096. For heptacene we obtain a basis error mismatch of $\delta_E(D, D_{\text{opt}}) \leq 1.2 \times 10^{-3}$ for all values of D , well below chemical accuracy of 1.6 mHartree. It indicates that for single-reference systems, low bond dimension DMRG-SCF calculations already capture the bulk of dynamic

correlations, in alignment with the expectation that single reference systems are typically easier to handle numerically. For strongly correlated systems on the other hand we expect to observe an increasing basis error with increasing bond dimension D . We expect that a systematic analysis of $\delta_E(D, D_{\text{opt}})$ could provide a useful tool for analyzing convergence of DMRG-SCF simulations in these cases.

In the inset of Figure 3(right) we show a $1/D$ scaling analysis of the (shifted) DMRG energies using the optimized basis, $E_0(D, D_{\text{opt}})$, with $D_{\text{opt}} = 128$ (x-symbols) and the (shifted) DMRG energies, $E_0(D)$, obtained without the SCF procedure (+-symbols). A big difference in energy in the order of 200 milliHartree remains for all D values, highlighting the importance of the SCF procedure. Using a second order polynomial fit, the $D \rightarrow \infty$ truncation-free solution can be approximately obtained.

The good quality of the low- D_{opt} optimized basis can be further understood by monitoring the evolution of the orbital occupation number, N_{occ} , via the DMRG-SCF procedure. In Figure 4(left), the occupation number is shown for the initial

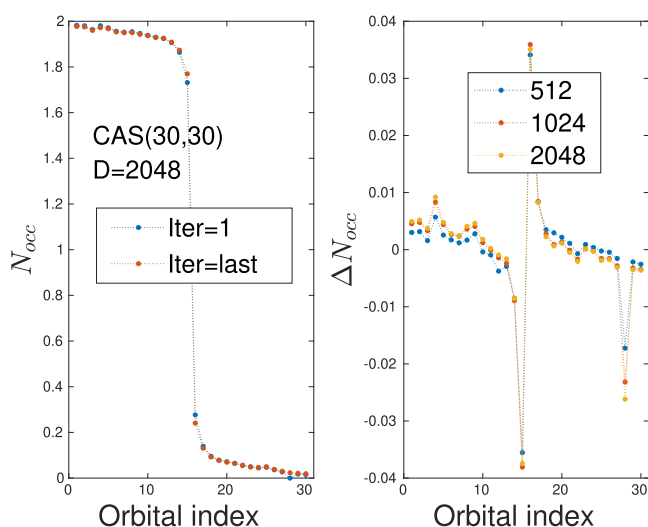


Figure 4. (left) Occupation number, N_{occ} , obtained by the first and last DMRG-SCF macro iteration for $D_{\text{opt}} = 2048$ and (right) the difference between the obtained N_{occ} profile via the first and last DMRG-SCF macro iteration for various D_{opt} values. In the SCF module the convergence on energy was set to 10^{-7} and the gradient to 10^{-3} .

(first) SCF step and for the last one. This can show in general how multireference nature of a given system is changed toward a single reference one by pushing orbital occupation toward two and zero in the vicinity of the Fermi surface. In the current example, however, the quality of the wave function remains almost the same as changes in N_{occ} are less than 4×10^{-2} . For better monitoring, i.e., to see how changes in the profile depend on the bond dimension, in Figure 4(right) we present the difference between the obtained N_{occ} profile via the first and last DMRG-SCF macro iteration for various D_{opt} values. Here a systematic improvement is evident which gets less and less pronounced with increasing D_{opt} values.

4.1.3. Large Active Spaces. In Figure 5 we show similar analyses to the previous section but for larger active spaces, using dodecacene ($\text{C}_{50}\text{H}_{28}$) on a CAS(50,50) and the icosacene ($\text{C}_{82}\text{H}_{44}$) on a CAS(82,82). Again we observe a significant

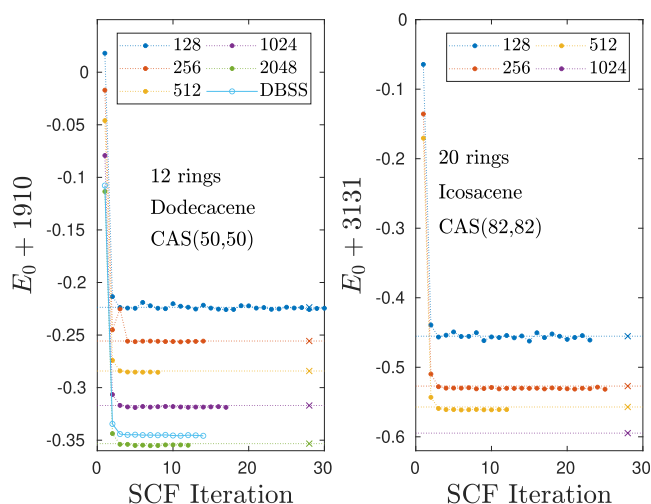


Figure 5. (Left panel) Similar to Figure 3 (right), but for the dodecacene on a CAS(50,50) and (right panel) for the Icosacene on a CAS(82,82).

reduction in SCF iteration number and a more stable convergence with increasing bond dimension D_{opt} .

We repeat the analysis of the previous section of the basis error mismatch by performing DMRG calculations at increasing bond dimensions D using the bases obtained from DMRG-SCF at $D_{\text{opt}} = 128$ for both molecules. We obtain a basis error mismatch in Dodecacene of $\delta_E(D, D_{\text{opt}}) \simeq 1.4 \times 10^{-3}$ for $D = 2048$ and in Icosacene $\simeq 4.1 \times 10^{-3}$ for $D = 512$ already. Therefore, for the larger active space CAS(82,82) the error gets more pronounced. This highlights the importance of basis optimization with high accuracy for increasing CAS sizes.

From a technical point of view, it is also evident from Figures 3 and 5 that the overall accuracy in leading order is determined by the DMRG bond dimension, i.e., by the accuracy of the CAS wave function. Therefore, an optimal protocol to set error margins would rely on the DBSS procedure. First, a DMRG-SCF procedure would be performed starting with a higher truncation error $\epsilon_{\text{tr}} \simeq 10^{-3}$ and by setting ϵ_{OGC} to an order of magnitude larger and setting the accuracy of the residual error of the diagonalization algorithm (e.g., Davidson or Lánczos) to an order of magnitude smaller. After convergence is reached, the DMRG-SCF procedure could be repeated by lowering ϵ_{tr} and adjusting ϵ_{OGC} and the residual error of the diagonalization algorithm accordingly. An extrapolation, as a function of ϵ_{tr} can be performed toward the truncation-free solution just like in conventional DMRG.^{25,72}

To provide further insights, in Table 2 we summarize the total computational time, t_{A100} and t_{H100} obtained on DGX-A100 and DXG-H100 single nodes, respectively, via fixed D and via the DBSS procedure for various parameter sets.

We note that our implementation shows high GPU utilization already at intermediate levels of bond dimension D as demonstrated in refs 44, 50, and 80, hence providing substantial speedups for DMRG-SCF runs already for smaller-scale problems, compared to previous implementations.

4.2. Iron–Sulfur Clusters. Next, we consider a more challenging system, built from iron–sulfur clusters (see Figure 2), where the appearance of metal atoms requires a more elaborate numerical treatment. The ground state is an open shell state with spin $S = 5/2$ or $S = 0$ for even and odd values of

Table 2. Total Number of Electrons, \tilde{N}_e , and Orbitals, \tilde{N} , of the Full Orbital Space, Size of Selected CAS, Minimum DMRG Bond Dimension, D_{\min} , Maximum DMRG Bond Dimension, D_{\max} , Corresponding Maximum DMRG U(1) Bond Dimension, \tilde{D}_{\max} , and Total Time in Hours for the Eight A100 and H100 GPU Accelerated DMRG-SCF, t_{A100} , t_{H100} , Respectively, for the Heptacene, Dodecacene, and Icosacene Obtained with $\epsilon_{\text{TR}} = 10^{-4a}$

\tilde{N}_e	\tilde{N}	CAS	D_{\min}	D_{\max}	\tilde{D}_{\max}	t_{A100} (h)	t_{H100} (h)	energy
198	510	(30,30)	512	512	1606	5.0	3.0	−1146.9536
198	510	(30,30)	1024	1024	3453	7.8	5.1	−1146.9738
198	510	(30,30)	2048	2048	5262	19.7	13.1	−1146.9875
198	510	(30,30)	1024	5363	17,022	15.7	10.3	−1146.9853
328	840	(50,50)	512	512	1622	6.5	4.1	−1910.2853
328	840	(50,50)	1024	1024	2561	24.8	16.3	−1910.3187
328	840	(50,50)	1024	4096	12,826	63.7	25.4	−1910.3458
536	1368	(82,82)	256	256	585	49.4	22.4	−3131.5315
536	1368	(82,82)	512	512	1115	54.6	32.1	−3131.5613

^aThe maximum number of DMRG sweeps was set to ten. For fixed D calculations $D_{\min} = D_{\max}$.

n , respectively. The more challenging nature of these systems requires careful adjustment of the DMRG-SCF hyperparameters, and typically between 20 and 30 sweeps per DMRG run were necessary to achieve stable convergence in the DMRG-SCF runs. Moreover, for sextet ground states ($S = 5/2$, even values of n), the high multiplicity leads to a significant (roughly six times) difference between the SU(2) D and its corresponding U(1) bond dimension \tilde{D} , thus our largest applied $D = 2048$ is equivalent to $\tilde{D} \simeq 12000$. Respecting SU(2) provides a significant computational and memory benefit in these cases.

4.2.1. Fe_1S_4 , $n = 1$ Case. In Figure 6(left) we show the convergence of the DMRG-SCF procedure for three different

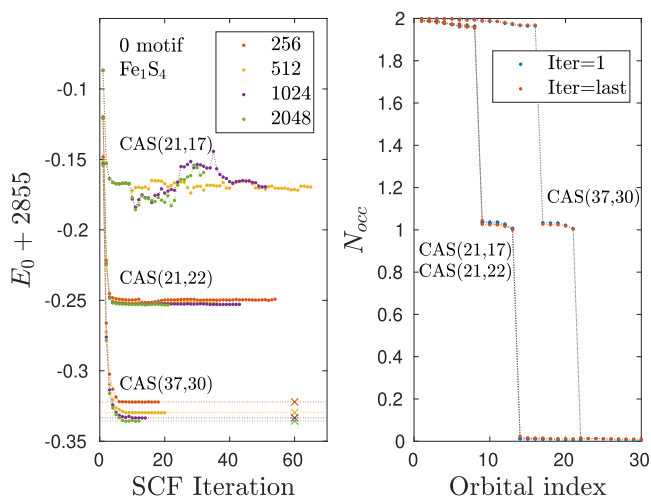


Figure 6. (left) Convergence profiles of the DMRG-SCF method for the Fe_1S_4 cluster for various bond dimension values and for three different active spaces. Symbols X together with horizontal dotted lines stand for post DMRG-SCF energy values obtained for the D values shown in the legend, but using the basis optimized with $D_{\text{opt}} = 256$. (right) occupation number distribution of active space orbitals calculated by the first and last DMRG-SCF iterations.

active spaces of the Fe_1S_4 model constructed from orbitals given in Table 1, labeled by ‘A’, ‘B’, and ‘C’, i.e. for CAS(21,17), CAS(21,22), and CAS(37,30), respectively, to study the effect of proper CAS selection. It is obvious that for a given CAS, a lower energy is reached with increasing D values as has also been shown for the PAH series. For too small $D = 64$, however, we found a highly oscillating profile, indicating

that the DMRG wave function was not accurate enough (not shown for the larger CAS configurations). With regard to the CAS dependence, we confirmed that the energy decreases with increasing CAS size, as expected.

However, it is important to note that for the smallest CAS, the DMRG-SCF procedure did not converge, even employing a very large bond dimension $D = 2048$, characterized by a very small truncation error of the order of 10^{-7} . Consequently, the DMRG result can be considered as a full-CI solution, and the numerical instability is due to the insufficiency of the truncated active space, namely due to the exclusion of the 4d orbitals of the Fe atom. We have further confirmed this by additional calculations using different orbital optimization techniques in the ORCA program, i.e., applying SuperCI+DIIS with level shift, instead of the default SuperCI(PT). Here, we got ~ 6 mE_h lower energy, but still experienced convergence issues. More importantly, orbitals corresponding to the 4p orbitals of Fe have been replaced by the 4d orbitals of Fe, signaling the poor choice of the selected CAS, and the importance of 4d orbitals. In contrast to this, for the larger active space ‘B’, including the 4d orbitals of the Fe atom as well, we found a stable convergence, but still, a very large D value had to be enforced to reduce the number of SCF iteration steps (for $D = 2048$ the optimization procedure terminated after 23 SCF iterations). In contrast to this, for the largest active space ‘C’, which includes all chemically relevant orbitals of the Fe and S atoms except the sigma bonds of the SH ligands, we found a very fast and stable convergence. The DMRG-SCF procedure terminated after 13–24 iterations. According to tight error settings in the CASSCF module of ORCA, the change in energy in the last three SCF iterations for $D \geq 512$ was in the range of 10^{-4} . The different numbers of data points shown for the various D values refer to converged DMRG-SCF results.

In Figure 6(right) we show the occupation number distribution for the largest active space CAS(37,30), which includes the additional 3s orbitals of the S and 3s and 3p orbitals of Fe atom. The five singly occupied d-orbitals of the Fe atom are well visible in the figure, remaining orbitals are either almost fully occupied or empty. The obtained profiles for the smaller active spaces are also shown for the sake of completeness, but they almost overlap in the given scale. It is clearly visible, however, that the additional occupied orbitals between the largest two active spaces are characterized with an occupation number larger than 1.99. Regarding the effect of orbital optimization on the occupation number distribution, we note that the changes in the profiles measured between the

Table 3. Total Number of Electrons, \tilde{N}_e , and Orbitals, \tilde{N} , of the Full Orbital Space, Size of Selected CAS, Minimum DMRG Bond Dimension, D_{\min} , Maximum DMRG Bond Dimension, D_{\max} , Corresponding Maximum DMRG U(1) Bond Dimension, \tilde{D}_{\max} and Total Time for the Eight A100 and H100 GPU Accelerated DMRG-SCF for the Iron–Sulfur Clusters with Increasing Number of Fe_2S_2 Motifs, n

\tilde{N}_e	\tilde{N}	CAS	D_{\min}	D_{\max}	\tilde{D}_{\max}	t_{A100} (h)	t_{H100} (h)	energy
95	135	(37,30)	256	256	1551	12.7	7.3	−2855.3221
95	135	(37,30)	512	512	3164	26.2	13.5	−2855.3297
95	135	(37,30)	1024	1024	6218	29.5	17.4	−2855.3334
95	135	(37,30)	2048	2048	12,478	54.1	33.6	−2855.3356
154	214	(66,56)	256	256	769	23.9	14.3	−4913.0073
154	214	(66,56)	512	512	1591	50.7	23.1	−4913.0699
154	214	(66,56)	1024	1024	3197	80.6	47.4	−4913.1147
213	293	(55,62)	256	256	1671	NaN	NaN	NaN
213	293	(55,62)	512	512	3467	123	55.7	−6970.4836

first and last DMRG-SCF iterations have been found to be of order 10^{-2} (see the small differences between the red and blue data points for each CAS space). Moreover, the plateau of the d orbitals has been preserved; thus, the initial configuration has not been altered by the SCF procedure. This, together with the small changes in the occupation number profile confirms that this system is single reference in nature.

We conclude our analysis, by comparing the best DMRG-SCF energy (see Table 3) to the CCSD(T) energy, $E = -2856.16129$. This indicates that there is still a discrepancy of 825 mHartree. Therefore, although the SCF procedure lead to an improvement of 182 mHartree, additional post-DMRG methods like DMRG-TCC,^{52,97,98} DMRG-RAS-X,^{79,99} or perturbation theory^{34,47,100–103} are required to capture the still missing dynamic correlations. The application of these methods to this and similar systems for recovering these final missing dynamic correlations will be part of a future publication. Finally, we highlight the fact that DMRG-SCF provides the best variational reference energy, and can be systematically improved with increasing bond dimension and CAS sizes.

4.2.2. One Fe_2S_2 Motif, $n = 2$ Case. A similar analysis for $n = 2$, i.e., for $\text{Fe}_2\text{S}_6\text{H}_4^{-2}$ corresponding to one Fe_2S_2 motif has also been performed, leading to a very similar conclusion. Our results are summarized in Figure 7.

Again, we obtained faster and more stable convergence with increasing CAS sizes and D values. For the smallest CAS the oscillating profile is due to the insufficiency of the active space, while for configuration “B” even $D = 256$ was not enough to reach stable convergence. Therefore, for more complex systems, larger bond dimensions in the SCF procedure are mandatory, which ultimately calls for the hybrid CPU-GPU parallel implementation. The importance of the inclusion of the 4d orbitals of the Fe atoms in the active spaces is reflected by the more stable convergence profiles for the two larger CASs. Regarding the quality of the optimized basis obtained via truncated bond dimension, we have performed again post-DMRG-SCF calculations with $D = 256, 512, 1024, 2048$ using the basis optimized with $D_{\text{opt}} = 256$. As can be seen, results shown by filled circles together with dotted lines are very close to those where the basis has been optimized via larger, corresponding, bond dimension values. However, for this more complex molecular system, the basis error gets more pronounced with increasing D , that for $D = 1024$ we found to be $\delta_E(D, D_{\text{opt}}) \approx 16$ mHartree. This is beyond chemical accuracy and indicates that larger bond dimensions are required in the SCF procedure for more complex problems

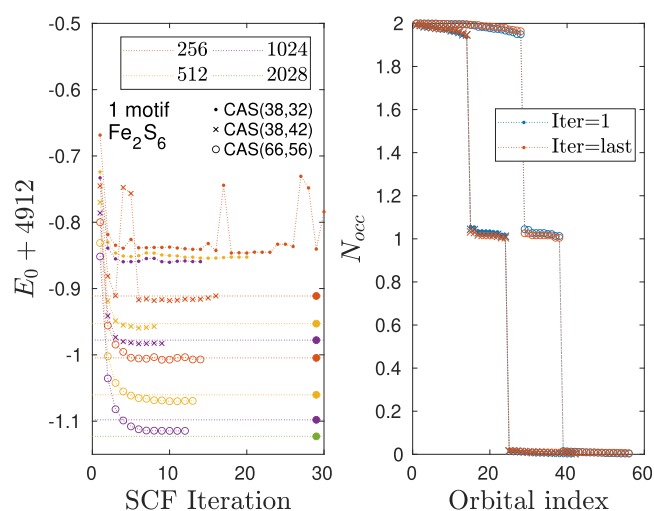


Figure 7. Similar to Figure 6 but for the $\text{Fe}_2\text{S}_6\text{H}_4^{-2}$ cluster corresponding to one Fe_2S_2 motif, parametrized with $n = 1$. Symbols filled circle together with horizontal dotted lines stand for post-DMRG-SCF energy values obtained for various $D = \{256, 512, 1024, 2048\}$ values using the basis optimized with $D_{\text{opt}} = 256$.

to reduce basis error. On the other hand, using post-DMRG-SCF data sets and a proper extrapolation via $1/D$ or the truncation error allows us to reach chemical accuracy.

Regarding the occupation number profiles, now the ten d-orbitals of the two Fe atoms are singly occupied as shown in Figure 7 (right). The difference between the two larger active spaces is marginal except that the Fermi surface is shifted. The occupation numbers of all added orbitals have again been found to be larger than 1.99. Since these almost fully occupied orbitals have only marginal contributions to correlation effects, for larger n values we performed calculations employing only CAS configuration “B”.

We close our analysis by remarking on the difficulty of the underlying calculations. We have found that, due to the close-lying eigenstates, the DMRG-SCF is very sensitive to accuracy settings. Therefore, if a not fully converged DMRG reference wave function is used, the SCF procedure can lead to unreliable results. For example, for a higher residual error in the Lanczos diagonalization (10^{-6}) or for not properly optimized orbital ordering, or not enough sweeps we could not reproduce the plateau of the d orbitals, and usually 3–4 scattered data points of occupation numbers appeared in the range of 1.4–1.8 and 0.2–0.6 even employing a large $D = 2048$. Moreover, performing post-DMRG-SCF calculations,

starting with a badly converged SCF result, despite using very large bond dimension values up to $D \sim 6000$, the correct occupation number profile could not be recovered. Therefore, in practice, it is not enough to monitor the convergence of the energy alone, but further properties of the wave function must also be analyzed. This follows from the fact that in a fixed-rank MPS manifold, several local minima with close energy values can correspond to wave functions with different properties.¹⁰⁴

For completeness, we also remark that if the d-orbitals were not localized initially, then the proper shape of the plateau could have been obtained with very large $D \geq 2048$ bond dimension values only even for the smallest active space. This also highlights the importance of the preparation of initial orbitals and their significant effect on the convergence speed and quality of the final result of DMRG-SCF.

4.2.3. Several Fe_2S_2 Motifs with $n \geq 3$. For larger n values, i.e., for more Fe_2S_2 motifs, we have obtained similar results employing CAS configuration “B”. Convergence profiles for $n = 3$, i.e., for two motifs, are shown in Figure 8. Here, for $n = 3$

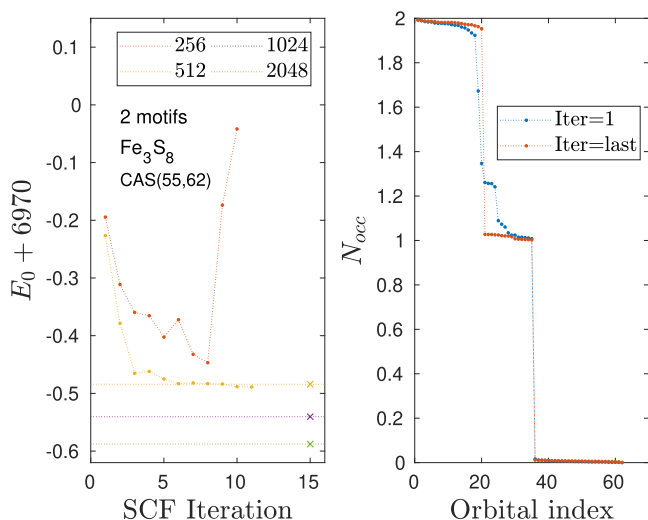


Figure 8. Similar to Figure 6 but for the $\text{Fe}_3\text{S}_8\text{H}_4^{3-}$ cluster corresponding to two Fe_2S_2 motifs parametrized with $n = 3$. Symbols X together with horizontal dotted lines stand for post DMRG-SCF energy values obtained for various D values using the basis optimized with $D_{\text{opt}} = 512$.

corresponding to $\text{Fe}_3\text{S}_8\text{H}_4^{3-}$ using CAS(55,62) we recovered the 15 singly occupied d-orbitals, after careful orbital ordering optimization. For tutorial purposes, we show the result for $n = 3$ starting DMRG-SCF from a not fully optimized ordering configuration with non fully converged wave function. As can be seen, even if the initial wave function was inadequate (see occupation number distribution at the plateau), DMRG-SCF can correct it, but as discussed before, there is no guarantee for that. Using a fully optimized initial DMRG configuration, we have a higher chance of getting the expected distribution of occupation numbers, which are usually preserved via the DMRG-SCF procedure. The main difficulty is due to the fact that only the total spin is conserved, but not the spin states of the individual iron atoms. Moreover, due to the larger computational complexity of the problem, larger bond dimensions are mandatory. As can be seen, for $D = 256$ DMRG-SCF can even lose the target after a few iteration steps when the large CAS error starts to influence the stability

of the SCF procedure. Such non converging profile can be seen in Figure 8.

For $n = 4$, for $\text{Fe}_4\text{S}_{10}\text{H}_4^{4-}$, we faced further serious difficulties even to recover the plateau of the 20 singly occupied d-orbitals on the applied CAS(72,82) model space. In spite of performing orbital optimizations via a large number of DMRG sweeps up to 50–60 together with bond dimensions up to 1024, we have always found a few scattered data points in the range of 1.4–1.8 and 0.2–0.6 of occupation. Therefore, finding the correct target state for DMRG is a hard task as the number of singly occupied orbitals increases, for which open shell problems, treated with large active spaces, further investigation is necessary. A promising direction would be a combination with restricted open-shell configuration interaction singles (ROCIS) family of methods.¹⁰⁵

To provide further insights in Table 3, we summarize the computational time via fixed D values for various parameter sets.

5. CONCLUSIONS

In this work, we have presented an efficient orbital optimization procedure utilizing the massively parallel spin adapted ab initio density matrix renormalization group (DMRG) method via the complete active space self-consistent field (CASSCF) framework on high performance computing infrastructures building on state-of-the-art hardware and software technologies. Our DMRG-SCF method employs the message passing interface (MPI) based parallel implementation of the SCF procedure accessible via the ORCA program package¹ which works in a perfect synergy with our hybrid CPU-multiGPU DMRG implementation.^{44,50,51,71} Substantial speed ups compared to existing CPU-based implementations is also due to the utilization of a GPU accelerated method to compute the one- and two-particle reduced density matrices, which we have put to use for running CASSCF calculations at large bond dimensions to obtain stable convergence with high accuracy and at unprecedented AS sizes of up to 82 electrons in 82 orbitals.

Our results for the polycyclic aromatic hydrocarbons (PAHs) obtained with various DMRG bond dimension values demonstrate the crucial need for accurate CAS solutions to reduce the number of SCF iterations and to reach converged results with an accuracy in the order of milliHartree. Moreover, we provided a detailed error analysis as a function of the DMRG bond dimension and DMRG truncation error. The quality of the optimized basis obtained with lower accuracy, i.e., larger truncation error, has been analyzed, indicating that basis error with larger active space sizes increases systematically. Nevertheless, by performing extrapolations with inverse bond dimension as part of a post-SCF DMRG procedure, one can approximate the truncation-free solution. On the other hand, much larger bond dimensions are needed if a nonoptimized basis is used.

Our analysis on iron–sulfur clusters reveals the importance of proper CAS selection and constraints on minimal CAS sizes to achieve stable convergence via the DMRG-SCF procedure. Numerical challenges related to such open-shell systems have been discussed in detail, along with the reported higher computational complexity and total wall time. Nevertheless, the singly occupied d-orbitals of the Fe atoms have been recovered, and the corresponding plateau in the occupation number profile has been argued to be an important quantity to monitor besides ground state energy. A natural extension of

our work would be toward Fe(II) porphyrin,⁵⁷ which is a particularly interesting system where treating all π orbitals as well as selected sigma electrons, double d-shell, and semicore orbitals may finally lead to a converged result.

Regarding computational time, we have also presented the measured wall time of full DMRG-SCF protocol on NVIDIA DGX-A100 and DGX-H100 hardware. These indicate that by considering large active spaces for systems composed of several hundreds of electrons over thousands of orbitals, the CASSCF-based orbital optimization can be revolutionized by taking advantage of the underlying computational power in AI accelerators available via graphics process units. In addition, for intermediate bond dimension values, i.e., for $D \leq 2000$, we measured a factor of 1.6–2.3 speedup by switching from an A100 to an H100 node. For larger D values, this speedup can even be larger based on our previous DMRG CAS benchmark calculations.⁸⁰

As future possibilities, we note that performing large-scale DMRG calculations on the optimized CAS space with even higher accuracy thresholds, together with the tailored coupled cluster (TCC) implementation, a significant part of the dynamic correlation can also be recovered.⁵² The latter also supports LPNO⁵³ and DLPNO treatments^{54,55} in the ORCA software package suited for several thousands of orbitals.^{56,57} Therefore, our approach paves the way for simulating strongly correlated multireference molecular systems with several thousands of electrons and orbitals as a routine daily task. In addition, extensions toward multi-Node/multi-GPU DMRG implementations reaching petaFLOPS performance is straightforward, however the details of this implementation will be discussed in a follow-up work.⁵⁸

As mentioned in Section 2, we have not utilized active–active rotations in the current work since they usually lead to numerical instabilities as the full-CI limit is approached. A way to circumvent this problem is to utilize more advanced DMRG protocols based on Fermionic mode optimization,^{74–78} which also optimizes the underlying one-particle basis in the CAS. This method also reduces the level of entanglement encoded in the CAS wave function, thus same accuracy can be achieved with significantly lower DMRG bond dimension. In addition, it is very stable and robust and even stationary conditions can be fulfilled.⁷⁹ After convergence is reached the resulting one- and two-particle reduced density matrices can be back rotated to the original CAS basis^{71,75} and passed to the SCF solver.

Finally, we remark that results presented in the current work have been obtained without employing the dynamically extended active space (DEAS) procedure^{83,106} since its SU(2) spin-adapted version is not available yet. The lack of such an initialization protocol is the reason for the high number of DMRG sweeps required in the DMRG-SCF approach. This is, however, part of our current developments, thus a factor of 2 to three reduction in sweeping is expected just like in our U(1) implementation. Further reduction in wall time is also expected by synchronizing the various error terms in DMRG and SCF even dynamically. Therefore, in the near future, our hybrid CPU-multiGPU DMRG code interfaced with the ORCA program package has the potential to target complex strongly correlated molecular clusters, including several transition metal centers, as a routinely applied daily method.

AUTHOR INFORMATION

Corresponding Authors

Örs Legeza – Strongly Correlated Systems Lendület Research Group, Wigner Research Centre for Physics, H-1525 Budapest, Hungary; Institute for Advanced Study, Technical University of Munich, Germany, 85748 Garching, Germany; Parmenides Stiftung, 82343 Pöcking, Germany; orcid.org/0000-0002-9746-5802; Email: legeza.ors@wigner.hu

Jeff Hammond – NVIDIA Helsinki Oy, 00180 Helsinki, Finland; orcid.org/0000-0003-3181-8190; Email: jeffpapers@nvidia.com

Sotiris S. Xantheas – Advanced Computing, Mathematics, and Data Division, Pacific Northwest National Laboratory, Richland, Washington 99354, United States; Department of Chemistry, University of Washington, Seattle, Washington 98195, United States; orcid.org/0000-0002-6303-1037; Email: Sotiris.Xantheas@pnnl.gov

Martin Ganahl – SandboxAQ, Palo Alto, California 94304, United States; Email: martin.ganahl@sandboxaq.com

Frank Neese – Max-Planck Institut für Kohlenforschung, D-45470 Mülheim an der Ruhr, Germany; Email: neese@kofo.mpg.de

Authors

Andor Menczer – Strongly Correlated Systems Lendület Research Group, Wigner Research Centre for Physics, H-1525 Budapest, Hungary; Eötvös Loránd University, 1117 Budapest, Hungary

Ádám Ganyecz – Strongly Correlated Systems Lendület Research Group, Wigner Research Centre for Physics, H-1525 Budapest, Hungary; orcid.org/0000-0002-8828-0283

Miklós Antal Werner – Strongly Correlated Systems Lendület Research Group, Wigner Research Centre for Physics, H-1525 Budapest, Hungary

Kornél Kapás – Strongly Correlated Systems Lendület Research Group, Wigner Research Centre for Physics, H-1525 Budapest, Hungary

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.jctc.5c00571>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work has been supported by the Hungarian National Research, Development and Innovation Office (NKFIH) through Grant Nos. K134983, PD146265 and TKP2021-NVA-04, by the Quantum Information National Laboratory of Hungary. Ö.L. acknowledges financial support by the Hans Fischer Senior Fellowship programme funded by the Technical University of Munich – Institute for Advanced Study. Ö.L. and S.S.X. acknowledge support from the Center for Scalable and Predictive methods for Excitation and Correlated phenomena (SPEC), funded as part of the Computational Chemical Sciences Program FWP 70942 by the U.S. Department of Energy (DOE), Office of Science, Office of Basic Energy Sciences, Division of Chemical Sciences, Geosciences, and Biosciences at Pacific Northwest National Laboratory. M.A.W. has also been supported by the Janos Bolyai Research Scholarship of the Hungarian Academy of Sciences. The simulations were performed on the national supercomputer HPE Apollo Hawk at the High Performance Computing

Center Stuttgart (HLRS) under the grant number MPTNS/44246, at the Wigner Scientific Computing Laboratory (WSCLAB), and on Google Cloud.

REFERENCES

- (1) Neese, F.; Wennmohs, F.; Becker, U.; Riplinger, C. The ORCA quantum chemistry program package. *J. Chem. Phys.* **2020**, *152*, 224108.
- (2) Aprà, E.; et al. NWChem: Past, present, and future. *J. Chem. Phys.* **2020**, *152*, 184102.
- (3) Werner, H.-J.; Knowles, P. J.; Knizia, G.; Manby, F. R.; Schütz, M. Molpro: a general-purpose quantum chemistry program package. *WIREs Computational Molecular Science* **2012**, *2*, 242.
- (4) Frisch, M. J. et al. *Gaussian 16 Revision C.01*; Gaussian Inc.: Wallingford CT, 2016.
- (5) Shao, Y.; et al. Advances in molecular quantum chemistry contained in the Q-Chem 4 program package. *Mol. Phys.* **2015**, *113*, 184.
- (6) te Velde, G.; Bickelhaupt, F. M.; Baerends, E. J.; Fonseca Guerra, C.; van Gisbergen, S. J. A.; Snijders, J. G.; Ziegler, T. Chemistry with ADF. *J. Comput. Chem.* **2001**, *22*, 931.
- (7) Kresse, G.; Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **1996**, *54*, 11169.
- (8) Sun, Q.; et al. Recent developments in the PySCF program package. *J. Chem. Phys.* **2020**, *153*, No. 024109.
- (9) Seritan, S.; Bannwarth, C.; Fales, B. S.; Hohenstein, E. G.; Isborn, C. M.; Kokkila-Schumacher, S. I. L.; Li, X.; Liu, F.; Luehr, N.; Snyder, J. W., Jr; Song, C.; Titov, A. V.; Ufimtsev, I. S.; Wang, L.-P.; Martínez, T. J. TeraChem: A graphical processing unit-accelerated electronic structure package for large-scale ab initio molecular dynamics. *WIREs Comput. Mol. Sci.* **2021**, *11*, No. e1494.
- (10) Giannozzi, P.; Baroni, S.; Bonini, N.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Chiarotti, G. L.; Cococcioni, M.; Dabo, I.; et al. Quantum ESPRESSO: a modular and open-source software project for quantum simulations of materials. *J. Phys.: Condens. Matter* **2009**, *21*, 395502.
- (11) Giannozzi, P.; Andreussi, O.; Brumme, T.; Bunau, O.; Buongiorno Nardelli, M.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Cococcioni, M.; et al. Advanced capabilities for materials modelling with Quantum ESPRESSO. *J. Phys.: Condens. Matter* **2017**, *29*, 465901.
- (12) Giannozzi, P.; Baseggio, O.; Bonfá, P.; Brunato, D.; Car, R.; Carnimeo, I.; Cavazzoni, C.; de Gironcoli, S.; Delugas, P.; Ferrari Ruffino, F.; et al. Quantum ESPRESSO toward the exascale. *J. Chem. Phys.* **2020**, *152*, 154105.
- (13) Pederson, R.; Kozłowski, J.; Song, R.; Beall, J.; Ganahl, M.; Hauru, M.; Lewis, A. G. M.; Yao, Y.; Mallick, S. B.; Blum, V.; Vidal, G. Large scale quantum chemistry with tensor processing units. *J. Chem. Theory Comput.* **2023**, *19*, 25.
- (14) Li Manni, G.; Alavi, A. Understanding the mechanism stabilizing intermediate spin states in Fe(II)-porphyrin. *J. Phys. Chem. A* **2018**, *122*, 4935.
- (15) Reiher, M. A theoretical challenge: Transition-metal compounds. *CHIMIA* **2009**, *63*, 140.
- (16) Li, Z.; Li, J.; Dattani, N. S.; Umrigar, C. J.; Chan, G. K.-L. The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations. *J. Chem. Phys.* **2019**, *150*, No. 024302.
- (17) Evangelista, F. A. Perspective: Multireference coupled cluster theories of dynamical electron correlation. *J. Chem. Phys.* **2018**, *149*, No. 030901.
- (18) Cao, Y.; Romero, J.; Olson, J. P.; Degroote, M.; Johnson, P. D.; Kieferová, M.; Kivlichan, I. D.; Menke, T.; Peropadre, B.; Sawaya, N. P. D.; Sim, S.; Veis, L.; Aspuru-Guzik, A. Quantum chemistry in the age of quantum computing. *Chem. Rev.* **2019**, *119*, 10856.
- (19) Bauer, B.; Bravyi, S.; Motta, M.; Chan, G. K.-L. Quantum algorithms for quantum chemistry and quantum materials science. *Chem. Rev.* **2020**, *120*, 12685.
- (20) Baiardi, A.; Reiher, M. The density matrix renormalization group in chemistry and molecular physics: Recent developments and new challenges. *J. Chem. Phys.* **2020**, *152*, No. 040903.
- (21) Affleck, I.; Kennedy, T.; Lieb, E. H.; Tasaki, H. Rigorous results on valence-bond ground states in antiferromagnets. *Phys. Rev. Lett.* **1987**, *59*, 799.
- (22) White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **1992**, *69*, 2863.
- (23) White, S. R.; Martin, R. L. Ab initio quantum chemistry using the density matrix renormalization group. *J. Chem. Phys.* **1999**, *110*, 4127.
- (24) Chan, G. K.-L.; Head-Gordon, M. Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group. *J. Chem. Phys.* **2002**, *116*, 4462.
- (25) Legeza, Ö.; Röder, J.; Hess, B. A. Controlling the accuracy of the density-matrix renormalization-group method: The dynamical block state selection approach. *Phys. Rev. B* **2003**, *67*, No. 125114.
- (26) Legeza, Ö.; Noack, R. M.; Sólyom, J.; Tincani, L. Applications of quantum information in the density-matrix renormalization group. In *Computational Many-Particle Physics*, Fehske, H.; Schneider, R.; Weisse, A., Eds.; Lecture Notes in Physics; Springer: Berlin, Heidelberg, 2008; Vol. 739, p. 653.
- (27) Chan, G. K.-L.; Dorando, J. J.; Ghosh, D.; Hachmann, J.; Neuscamman, E.; Wang, H.; Yanai, T. An introduction to the density matrix renormalization group ansatz in quantum chemistry. In *Frontiers in Quantum Systems in Chemistry and Physics*, Wilson, S.; Grout, P. J.; Maruani, J.; Delgado-Barrio, G.; Piecuch, P., Eds.; Progress in Theoretical Chemistry and Physics; Springer: Netherlands, 2008, p. 18.
- (28) Yanai, T.; Kurashige, Y.; Ghosh, D.; Chan, G. K.-L. Accelerating convergence in iterative solution for large-scale complete active space self-consistent-field calculations. *Int. J. Quantum Chem.* **2009**, *109*, 2178.
- (29) Kurashige, Y.; Yanai, T. High-performance ab initio density matrix renormalization group method: Applicability to large-scale multireference problems for metal compounds. *J. Chem. Phys.* **2009**, *130*, 234114.
- (30) Marti, K. H.; Reiher, M. The density matrix renormalization group algorithm in quantum chemistry. *Zeitschrift für Physikalische Chemie* **2010**, *224*, 583.
- (31) Wouters, S.; Poelmans, W.; Ayers, P. W.; Van Neck, D. CheMP2S2: A free open-source spin-adapted implementation of the density matrix renormalization group for ab initio quantum chemistry. *Comput. Phys. Commun.* **2014**, *185*, 1501.
- (32) Szalay, Sz; Pfeffer, M.; Murg, V.; Barcza, G.; Verstraete, F.; Schneider, R.; Legeza, Ö. Tensor product methods and entanglement optimization for ab initio quantum chemistry. *Int. J. Quantum Chem.* **2015**, *115*, 1342.
- (33) Chan, G. K.-L.; Keselman, A.; Nakatani, N.; Li, Z.; White, S. R. Matrix product operators, matrix product states, and ab initio density matrix renormalization group algorithms. *J. Chem. Phys.* **2016**, *145*, No. 014102.
- (34) Cheng, Y.; Xie, Z.; Ma, H. Post-density matrix renormalization group methods for describing dynamic electron correlation with large active spaces. *J. Phys. Chem. Lett.* **2022**, *13*, 904.
- (35) Schollwöck, U. The density-matrix renormalization group. *Rev. Mod. Phys.* **2005**, *77*, 259.
- (36) Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics* **2011**, *326*, 96.
- (37) Verstraete, F.; Cirac, J. I. Renormalization algorithms for quantum-many body systems in two and higher dimensions. *arXiv:cond-mat/0407066v1* 2004.
- (38) Corboz, P.; Vidal, G. Fermionic multiscale entanglement renormalization ansatz. *Phys. Rev. B* **2009**, *80*, No. 165129.
- (39) Vidal, G. Entanglement renormalization. *Phys. Rev. Lett.* **2007**, *99*, No. 220405.

- (40) Verstraete, F.; Nishino, T.; Schollwöck, U.; Bañuls, M. C.; Chan, G. K.; Stoudenmire, M. E. Density matrix renormalization group, 30 years on. *Nat. Rev. Phys.* **2023**, *5*, 273–276.
- (41) Siegbahn, P.; Heiberg, A.; Roos, B.; Levy, B. A comparison of the super-CI and the Newton-Raphson scheme in the complete active space SCF method. *Phys. Scr.* **1980**, *21*, 323.
- (42) Roos, B. O.; Taylor, P. R.; Siegbahn, P. E. M. A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach. *Chem. Phys.* **1980**, *48*, 157.
- (43) Siegbahn, P. E. M.; Almlöf, J.; Heiberg, A.; Roos, B. O. The complete active space SCF (CASSCF) method in a Newton–Raphson formulation with application to the HNO molecule. *J. Chem. Phys.* **1981**, *74*, 2384.
- (44) Menczer, A.; Legeza, Ö. Tensor network state algorithms on AI accelerators. *J. Chem. Theory Comput.* **2024**, *20*, 8897.
- (45) Xiang, C.; Jia, W.; Fang, W.-H.; Li, Z. A distributed multi-GPU ab initio density matrix renormalization group algorithm with applications to the P-cluster of nitrogenase. *J. Chem. Theory Comput.* **2024**, *20*, 775.
- (46) Zgid, D.; Nooijen, M. The density matrix renormalization group self-consistent field method: Orbital optimization with the density matrix renormalization group method in the active space. *J. Chem. Phys.* **2008**, *128*, 144116.
- (47) Liu, F.; Kurashige, Y.; Yanai, T.; Morokuma, K. Multireference ab initio density matrix renormalization group (DMRG)-CASSCF and DMRG-CASPT2 study on the photochromic ring opening of spiropyran. *J. Chem. Theory Comput.* **2013**, *9*, 4462.
- (48) Wouters, S.; Bogaerts, T.; Van Der Voort, P.; Van Speybroeck, V.; Van Neck, D. Communication: DMRG-SCF study of the singlet, triplet, and quintet states of oxo-Mn(Salen). *J. Chem. Phys.* **2014**, *140*, 241103.
- (49) Ding, L.; Knecht, S.; Schilling, C. Quantum information-assisted complete active space optimization (QICAS). *J. Phys. Chem. Lett.* **2023**, *14*, 11022.
- (50) Menczer, A.; Legeza, Ö. Massively parallel tensor network state algorithms on hybrid CPU-GPU based architectures. *J. Chem. Theory Comput.* **2025**, *21*, 1572.
- (51) Menczer, A.; Legeza, Ö. Cost optimized ab initio tensor network state methods: industrial perspectives. *arXiv:2412.04676*, 2024.
- (52) Veis, L.; Antalík, A.; Brabec, J.; Neese, F.; Legeza, Ö.; Pittner, J. Coupled cluster method with single and double excitations tailored by matrix product state wave functions. *J. Phys. Chem. Lett.* **2016**, *7*, 4072.
- (53) Demel, O.; Pittner, J.; Neese, F. A local pair natural orbital-based multireference Mukherjee's coupled cluster method. *J. Chem. Theory Comput.* **2015**, *11*, 3104.
- (54) Brabec, J.; Lang, J.; Saitow, M.; Pittner, J.; Neese, F.; Demel, O. Domain-based local pair natural orbital version of Mukherjee's state-specific coupled cluster method. *J. Chem. Theory Comput.* **2018**, *14*, 1370.
- (55) Lang, J.; Brabec, J.; Saitow, M.; Pittner, J.; Neese, F.; Demel, O. Perturbative triples correction to domain-based local pair natural orbital version of Mukherjee's state-specific coupled cluster method. *Phys. Chem. Chem. Phys.* **2019**, *21*, 5022.
- (56) Lang, J.; Antalík, A.; Veis, L.; Brabec, J.; Legeza, Ö.; Pittner, J. Near-linear scaling in DMRG-based tailored coupled clusters: An implementation of DLPNO-TCCSD and DLPNO-TCCSD(T). *J. Chem. Theory Comput.* **2020**, *16*, 3028.
- (57) Antalík, A.; Nachtigallová, D.; Lo, R.; Matoušek, M.; Lang, J.; Legeza, Ö.; Pittner, J.; Hobza, P.; Veis, L. Ground state of the Fe(II)-porphyrin model system corresponds to the quintet state: DFT, DMRG-TCCSD and DMRG-TCCSD(T) computations. *Phys. Chem. Chem. Phys.* **2020**, *22*, 17033.
- (58) Menczer, A.; Legeza, Ö. Petaflops density matrix renormalization group method, unpublished, 2023.
- (59) Kollmar, C.; Sivalingam, K.; Helmich-Paris, B.; Angeli, C.; Neese, F. A perturbation-based super-CI approach for the orbital optimization of a CASSCF wave function. *J. Comput. Chem.* **2019**, *40*, 1463.
- (60) Zgid, D.; Nooijen, M. Obtaining the two-body density matrix in the density matrix renormalization group method. *J. Chem. Phys.* **2008**, *128*, 144115.
- (61) Khedkar, A.; Roemelt, M. Active space selection based on natural orbital occupation numbers from n-electron valence perturbation theory. *J. Chem. Theory Comput.* **2019**, *15*, 3522.
- (62) Stein, C. J.; Reiher, M. Automated selection of active orbital spaces. *J. Chem. Theory Comput.* **2016**, *12*, 1760.
- (63) Sayfutyarova, E. R.; Sun, Q.; Chan, G. K.-L.; Knizia, G. Automated construction of molecular active spaces from atomic valence orbitals. *J. Chem. Theory Comput.* **2017**, *13*, 4063.
- (64) Östlund, S.; Rommer, S. Thermodynamic limit of density matrix renormalization. *Phys. Rev. Lett.* **1995**, *75*, 3537.
- (65) McCulloch, I. P.; Gulácsi, M. The non-Abelian density matrix renormalization group algorithm. *Europhys. Lett.* **2002**, *57*, 852.
- (66) Tóth, A. I.; Moca, C. P.; Legeza, Ö.; Zaránd, G. Density matrix numerical renormalization group for non-Abelian symmetries. *Phys. Rev. B* **2008**, *78*, No. 245109.
- (67) Sharma, S.; Chan, G. K.-L. Spin-adapted density matrix renormalization group algorithms for quantum chemistry. *J. Chem. Phys.* **2012**, *136*, 124121.
- (68) Keller, S.; Reiher, M. Spin-adapted matrix product states and operators. *J. Chem. Phys.* **2016**, *144*, 134101.
- (69) Gunst, K.; Verstraete, F.; Van Neck, D. Three-legged tree tensor networks with SU(2) and molecular point group symmetry. *J. Chem. Theory Comput.* **2019**, *15*, 2996.
- (70) Werner, M. A.; Moca, C. P.; Legeza, Ö.; Zaránd, G. Quantum quench and charge oscillations in the SU(3) Hubbard model: A test of time evolving block decimation with general non-abelian symmetries. *Phys. Rev. B* **2020**, *102*, No. 155108.
- (71) Menczer, A.; Kapás, K.; Werner, M. A.; Legeza, Ö. Two-dimensional quantum lattice models via mode optimized hybrid CPU-GPU density matrix renormalization group method. *Phys. Rev. B* **2024**, *109*, No. 195148.
- (72) Legeza, Ö.; Fátih, G. Accuracy of the density-matrix renormalization-group method. *Phys. Rev. B* **1996**, *53*, 14349.
- (73) Noack, R. M.; Manmana, S. R. Diagonalization- and numerical renormalization-group-based methods for interacting quantum systems. *AIP Conf. Proc.* **2005**, *789*, 93.
- (74) Krumnow, C.; Veis, L.; Legeza, Ö.; Eisert, J. Fermionic orbital optimization in tensor network states. *Phys. Rev. Lett.* **2016**, *117*, No. 210402.
- (75) Krumnow, C.; Veis, L.; Eisert, J.; Legeza, Ö. Effective dimension reduction with mode transformations: Simulating two-dimensional fermionic condensed matter systems with matrix-product states. *Phys. Rev. B* **2021**, *104*, No. 075137.
- (76) Máté, M.; Petrov, K.; Szalay, Sz.; Legeza, Ö. Compressing multireference character of wave functions via fermionic mode optimization. *J. Math. Chem.* **2023**, *61*, 362.
- (77) Petrov, K.; Ganyecz, Á.; Benedek, Z.; Olsasz, A.; Barcza, G.; Legeza, Ö. Low-cost generation of optimal molecular orbitals for multireference CI expansion: natural orbitals versus Rényi entropy minimized orbitals provided by the density matrix renormalization group. In *Advances in Methods and Applications of Quantum Systems in Chemistry, Physics, and Biology*; Grabowski, I.; Słowik, K.; Maruani, J.; Brändas, E. J., Eds.; Springer: Cham, 2024.
- (78) Werner, M. A.; Menczer, A.; Legeza, Ö. Tensor network state methods and quantum information theory for strongly correlated molecular systems. *arXiv:2501.18263*, 2025.
- (79) Friesecke, G.; Barcza, G.; Legeza, Ö. Predicting the FCI energy of large systems to chemical accuracy from restricted active space density matrix renormalization group calculations. *J. Chem. Theory Comput.* **2024**, *20*, 87.
- (80) Menczer, A.; van Damme, M.; Rask, A.; Huntington, L.; Hammond, J.; Xantheas, S. S.; Ganahl, M.; Legeza, Ö. Parallel implementation of the density matrix renormalization group method achieving a quarter petaFLOPS performance on a single DGX-H100 GPU node. *J. Chem. Theory Comput.* **2024**, *20*, 8397.

- (81) Menczer, A.; Legeza, Ö. Efficient calculation of one- and two-particle reduced density matrices on GPUs via the density matrix renormalization group method. unpublished, 2024.
- (82) Legeza, Ö.; Röder, J.; Hess, B. A. QC-DMRG study of the ionic-neutral curve crossing of LiF. *Mol. Phys.* **2019**, *101*, 2019.
- (83) Barcza, G.; Legeza, Ö.; Marti, K. H.; Reiher, M. Quantum-information analysis of electronic states of different molecular structures. *Phys. Rev. A* **2011**, *83*, No. 012508.
- (84) Hammond, J. R.; Kowalski, K.; deJong, W. A. Dynamic polarizabilities of polyaromatic hydrocarbons using coupled-cluster linear response theory. *J. Chem. Phys.* **2007**, *127*, 144105.
- (85) Hajgató, B.; Szieberth, D.; Geerlings, P.; De Proft, F.; Deleuze, M. S. A benchmark theoretical study of the electronic ground state and of the singlet-triplet split of benzene and linear acenes. *J. Chem. Phys.* **2009**, *131*, 224321.
- (86) Zade, S. S.; Bendikov, M. Heptacene and beyond: The longest characterized acenes. *Angew. Chem., Int. Ed.* **2010**, *49*, 4012.
- (87) Pelzer, K.; Greenman, L.; Gidofalvi, G.; Mazziotti, D. A. Strong correlation in acene sheets from the active-space variational two-electron reduced density matrix method: Effects of symmetry and size. *J. Phys. Chem. A* **2011**, *115*, 5632.
- (88) Aiga, F. Theoretical study on oligoacenes and polycyclic aromatic hydrocarbons using the restricted active space self-consistent field method. *J. Phys. Chem. A* **2012**, *116*, 663.
- (89) Rivero, P.; Jiménez-Hoyos, C. A.; Scuseria, G. E. Entanglement and polyradical character of polycyclic aromatic hydrocarbons predicted by projected Hartree-Fock theory. *J. Phys. Chem. B* **2013**, *117*, 12750.
- (90) Plasser, F.; Pašalić, H.; Gerzabek, M. H.; Libisch, F.; Reiter, R.; Burgdörfer, J.; Müller, T.; Shepard, R.; Lischka, H. The multiradical character of one- and two-dimensional graphene nanoribbons. *Angew. Chem., Int. Ed.* **2013**, *52*, 2581.
- (91) Becke, A. D. Density-functional thermochemistry. I. the effect of the exchange-only gradient correction. *J. Chem. Phys.* **1992**, *96*, 2155.
- (92) Becke, A. D. A new mixing of Hartree-Fock and local density-functional theories. *J. Chem. Phys.* **1993**, *98*, 1372.
- (93) Dunning, T. H., Jr. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.* **1989**, *90*, 1007.
- (94) Woon, D. E.; Dunning, T. H., Jr. Gaussian basis sets for use in correlated molecular calculations. III. The atoms aluminum through argon. *J. Chem. Phys.* **1993**, *98*, 1358.
- (95) Lippard, S. J.; Berg, J. M. *Principles of Bioinorganic Chemistry*; University Science Books, 1994.
- (96) Izák, R.; Ivanov, A. V.; Blunt, N. S.; Holzmann, N.; Neese, F. Measuring electron correlation: The impact of symmetry and orbital transformations. *J. Chem. Theory Comput.* **2023**, *19*, 2703.
- (97) Faulstich, F. M.; Laestadius, A.; Legeza, Ö.; Schneider, R.; Kvaal, S. Analysis of the tailored coupled-cluster method in quantum chemistry. *SIAM Journal on Numerical Analysis* **2019**, *57*, 2579.
- (98) Liao, K.; Ding, L.; Schilling, C. Quantum information orbitals (QIO): Unveiling intrinsic many-body complexity by compressing single-body triviality. *J. Phys. Chem. Lett.* **2024**, *15*, 6782.
- (99) Larsson, H. R.; Zhai, H.; Gunst, K.; Chan, G. K.-L. Matrix product states with large sites. *J. Chem. Theory Comput.* **2022**, *18*, 749.
- (100) Stolarczyk, L. Z. Complete active space coupled-cluster method. Extension of single-reference coupled-cluster method using the CASSCF wavefunction. *Chem. Phys. Lett.* **1994**, *217*, 1.
- (101) Kurashige, Y.; Yanai, T. Second-order perturbation theory with a density matrix renormalization group self-consistent field reference function: Theory and application to the study of chromium dimer. *J. Chem. Phys.* **2011**, *135*, No. 094104.
- (102) Pulay, P. A perspective on the CASPT2 method. *Int. J. Quantum Chem.* **2011**, *111*, 3273.
- (103) Sharma, S.; Chan, G. K.-L. Communication: A flexible multi-reference perturbation theory by minimizing the Hylleraas functional with matrix product states. *J. Chem. Phys.* **2014**, *141*, 111101.
- (104) Boguslawski, K.; Marti, K. H.; Reiher, M. Construction of CASCI-type wave functions for very large active spaces. *J. Chem. Phys.* **2011**, *134*, 224101.
- (105) Leyser da Costa Gouveia, T.; Maganas, D.; Neese, F. General spin-restricted open-shell configuration interaction approach: Application to metal K-edge X-ray absorption spectra of ferro- and antiferromagnetically coupled dimers. *J. Phys. Chem. A* **2024**, *129*, 330.
- (106) Legeza, Ö.; Sólyom, J. Optimizing the density-matrix renormalization group method using quantum information entropy. *Phys. Rev. B* **2003**, *68*, No. 195116.



CAS BIOFINDER DISCOVERY PLATFORM™

**PRECISION DATA
FOR FASTER
DRUG
DISCOVERY**

CAS BioFinder helps you identify
targets, biomarkers, and pathways

Unlock insights

CAS
A division of the
American Chemical Society