

# Enhanced Krylov Methods for Molecular Hamiltonians: Reduced Memory Cost and Complexity Scaling via Tensor Hypercontraction

Yu Wang,\* Maxine Luo, Matthias Reumann, and Christian B. Mendl\*

Cite This: *J. Chem. Theory Comput.* 2025, 21, 6874–6886


Read Online

ACCESS |

Metrics & More

Article Recommendations

**ABSTRACT:** We introduce an algorithm that is simultaneously memory-efficient and low-scaling for applying ab initio molecular Hamiltonians to matrix-product states (MPS) via the tensor-hypercontraction (THC) format. These gains carry over to Krylov subspace methods, which can find low-lying eigenstates and simulate quantum time evolution while avoiding local minima and maintaining high accuracy. In our approach, the molecular Hamiltonian is represented as a sum of products of four MPOs, each with a bond dimension of only 2. Iteratively applying the MPOs to the current quantum state in MPS form, summing and recompressing the MPS leads to a scheme with the same asymptotic memory cost as the bare MPS and reduces the computational cost scaling compared to the Krylov method using a conventional MPO construction. We provide a detailed theoretical derivation of these statements and conduct supporting numerical experiments to demonstrate the advantage. Our algorithm is highly parallelizable and thus lends itself to large-scale HPC simulations.



## 1. INTRODUCTION

We aim to simulate a molecular Hamiltonian, which is also known as electronic structure Hamiltonian, of the form (with  $L$  the number of electronic spatial orbitals)

$$\begin{aligned}
 H &= T + V \\
 &= \sum_{p,q=1}^L \sum_{\sigma \in \{\uparrow, \downarrow\}} t_{pq} a_{p,\sigma}^\dagger a_{q,\sigma} + \frac{1}{2} \sum_{p,q,r,s=1}^L \sum_{\sigma, \sigma' \in \{\uparrow, \downarrow\}} v_{pqrs} a_{p,\sigma}^\dagger a_{q,\sigma}^\dagger a_{r,\sigma'} a_{s,\sigma'}
 \end{aligned} \quad (1)$$

using tensor network methods, specifically the matrix product state (MPS) formalism.<sup>1–3</sup>  $a_{p,\sigma}^\dagger$  and  $a_{p,\sigma}$  are the Fermionic creation and annihilation operators, respectively, and  $t_{pq}$ ,  $v_{pqrs}$  are coefficients resulting from single- and two-body orbital overlap integrals. To understand molecular properties such as the electronic structure,<sup>4,5</sup> optoelectronic properties,<sup>6</sup> or molecular vibrations,<sup>7</sup> the density matrix renormalization group (DMRG) method<sup>8,9</sup> is widely applied to chemical systems with strong correlations, where traditional density functional theory and coupled cluster approaches face significant challenges.<sup>3,7,10–16</sup> The development of attosecond-level experimental techniques<sup>17–24</sup> motivates the simulation of ultrafast electron dynamics since they determine the formation and breaking of chemical bonds.<sup>21</sup> The time-dependent variational principle (TDVP) is a widely used time evolution method to predict electron dynamics.<sup>5,25–29</sup> However, both of the above are

variational methods in which the MPS evolves locally. It could result in the DMRG method getting trapped in local minima<sup>26,30</sup> and might lead to an inaccurate time evolution simulation by TDVP<sup>31,32</sup> even for simple models.<sup>33</sup>

In contrast, global Krylov subspace methods optimize all the sites globally and simultaneously,<sup>34</sup> offering a reliable alternative method if DMRG or TDVP run into problems. Krylov subspace methods like the Lanczos algorithm<sup>35–37</sup> can compute low-energy eigenstates reliably without local minima. The Lanczos algorithm also has the favorable capability of finding multiple excited states, being less sensitive to the results of lower eigenstates.<sup>38,39</sup> Conversely, using the DMRG algorithm, one has to explicitly project out the lower eigenstates, implying that inaccuracies propagate to the higher ones. To simulate time evolution, the global Krylov method<sup>4</sup> provides high-order error scaling<sup>28,40,41</sup> and works reliably. The value of the Krylov method lies in its ability to ensure accuracy while remaining robust across all models, rather than failing in a few cases like DMRG and TDVP methods.

Received: April 2, 2025

Revised: June 12, 2025

Accepted: June 16, 2025

Published: July 2, 2025




However, the reachable system sizes and MPS bond dimensions in the Krylov methods are relatively small when using the molecular Hamiltonian in conventional matrix product operator (MPO) form. Especially when one chooses high-accuracy MPS compression methods such as singular value decomposition (SVD),<sup>28,42</sup> the restriction results from the core step: applying the Hamiltonian to a quantum state, i.e., computing  $H|\psi\rangle$  in the tensor network formalism and compressing it. Considering a molecular Hamiltonian of the form (1), the maximum bond dimension  $D$  scales as  $O(L^2)$  when using conventional MPO constructions.<sup>11,12,43</sup> One needs intensive memory to store  $H|\psi\rangle$ , whose bond dimension is the product of the MPS and MPO bond dimensions. Moreover, compressing  $H|\psi\rangle$  to MPS form with smaller bond dimensions is essential for further calculations;<sup>34,41</sup> the computational cost is also high. The difficulty arises from the nonlocality of the two-body integral tensor  $v \in \mathbb{R}^{L \times L \times L \times L}$ , which makes the molecular Hamiltonian more complicated than a Hamiltonian containing only local interactions.

In this work, we propose and study an alternative Krylov method based on the tensor hypercontraction (THC) representation of  $v$ .<sup>44–46</sup>

$$v_{pqrs} \approx \sum_{\mu, \nu=1}^N \chi_p^\mu \chi_q^\mu \zeta^{\mu\nu} \chi_r^\nu \chi_s^\nu \quad (2)$$

where  $N$  is the THC rank. This formulation involves only two distinct matrices  $\chi$  and  $\zeta$ , as illustrated in Figure 1. We will show



**Figure 1.** Graphical representation of the THC factorization to approximate the Coulomb (electron repulsion integral) tensor.

that the THC representation allows us to rewrite the electronic Hamiltonian into a sum of sub-Hamiltonians, denoted THC-MPO, where each sub-Hamiltonian can be constructed as the product of four small MPOs with bond dimensions of only 2. Compared to calculations using a conventional MPO, such a small and constant bond dimension enables us to compute and compress  $H|\psi\rangle$  with significantly reduced memory requirements and better complexity scaling; both are reduced by a factor of  $O(L^4)$  asymptotically. We demonstrate the advantages of our THC-MPO by utilizing it for low-lying eigenstates search and time evolution simulations based on Krylov subspace methods, exemplified by the water molecule  $H_2O$ , hydrogen chain with ten atoms  $H_{10}$ , and the Ammonia molecule  $NH_3$ . This allows us to track the accuracy and error sources by comparing them to



**Figure 2.** Graphical tensor network representation of an MPS and MPO. The logical wave function and operator are obtained by contracting the matrices  $A[i]^{n_i}$  and  $W[i]^{m_i}$ , respectively.

results from the full configuration interaction (FCI) or exact diagonalization (ED) method. The numerical experiments show that our method enables us to calculate previously inaccessible system sizes when using Krylov methods with SVD compression; the memory advantages of our method become immediately apparent. Additionally, we will provide a general estimation of the computational complexity of larger molecules to highlight the potential. We will also illustrate why our method is well-suited for parallel computing.

## 2. THEORETICAL BACKGROUND

**2.1. Matrix Product States and Operators.** In the tensor network framework, the wave function  $|\Psi\rangle$  is typically represented as a matrix product state (MPS), also called tensor train<sup>1,2,42,47</sup>

$$|\Psi\rangle = \sum_{n_1, \dots, n_L} A[1]^{n_1} A[2]^{n_2} \dots A[L]^{n_L} |n_1, \dots, n_L\rangle \quad (3)$$

Each  $A[i]$  is a tensor of order three, as shown in Figure 2a. The superscript  $n_i$  is a physical index enumerating the possible states at site  $i$ , and  $A[i]^{n_i}$  is a  $\chi_i \times \chi_{i+1}$  matrix for each  $n_i$ . The variable  $\chi_i$  is the  $i$ -th bond dimension. We denote the maximum MPS bond dimension by  $M$  in the following.


Alongside MPS, there is a corresponding formalism for operators, known as matrix product operators (MPOs),<sup>11,12,42</sup> given by

$$\hat{O} = \sum_{m, n} W[1]^{m_1 n_1} W[2]^{m_2 n_2} \dots W[L]^{m_L n_L} |m_1, \dots, m_L\rangle \langle n_1, \dots, n_L| \quad (4)$$

Each element  $W[i]^{m_i n_i}$  is a matrix of shape  $\beta_i \times \beta_{i+1}$ , where  $\beta_i$  is the  $i$ -th bond dimension, as shown in Figure 2b. We denote the maximum MPO bond dimension by  $D$  in the following. In eq 4, an operator  $\hat{O}$  is represented with respect to computational basis states  $|m_1, \dots, m_L\rangle \langle n_1, \dots, n_L|$ , and the corresponding coefficients are obtained by contracting the  $W$  matrices.

Analogous to how applying a Hamiltonian matrix to a state vector yields a state vector, applying an MPO to an MPS will result in a new MPS by contracting the local physical tensor legs. Such an operation will increase the MPS bond dimension from  $M$  to  $D \cdot M$ . Thus, it is significant to compress this resulting MPS for further calculations, especially for iterative applications appearing in Krylov methods. The combined application and compression procedure is illustrated in Figure 3.

**2.2. The Krylov Methods Based on Matrix Product States.** The complexity of the Hamiltonian in full matrix form scales exponentially with system size; thus, exact diagonalization approaches are restricted to relatively small systems. A possibility to overcome this restriction consists in combining the Lanczos algorithm<sup>35,36</sup> with the MPS representation.<sup>37,48</sup> Namely, the MPS ansatz offers an economical representation of quantum states, and the Lanczos algorithm confines the evolving space to the Krylov subspace spanned by  $\{|\psi\rangle, H|\psi\rangle, H^2|\psi\rangle, \dots$



**Figure 3.** Multiplying an MPO with an MPS and subsequent compression. We first contract the tensors along the physical axis. Then, the MPS is transformed into right-canonical form by QR decompositions. Next, we employ SVDs from left to right to reduce the bond dimension by discarding the smallest singular values and merging  $S$  and  $V^\dagger$  matrices into the next site.

$H^{K-1}|\psi\rangle\}$ . In this work, we construct an orthogonal basis of the Krylov subspace. Starting from some initial state  $|v_0\rangle \equiv |\psi\rangle$ , we compute the next Krylov vector  $|v_{i+1}\rangle$  by applying the Hamiltonian to  $|v_i\rangle$  and orthogonalizing it w.r.t. the previous ones using the Gram-Schmidt algorithm.<sup>34,37,41</sup> It is also possible to use these Krylov vectors without orthogonalization; see ref 34,41 for further details.

The Hamiltonian is projected onto the Krylov subspace, and the elements of the resulting effective Hamiltonian are given by

$$\tilde{H}_{ij} = \langle v_i | H | v_j \rangle \quad (5)$$

where  $\{|v_0\rangle, |v_1\rangle, \dots, |v_{K-1}\rangle\}$  are orthonormalized Krylov vectors forming a basis of the subspace. Typically, we assume this effective Hamiltonian to be tridiagonal so that we only retain the entries with  $|i - j| \leq 1$ .<sup>34,37</sup> Assuming that the subspace dimension  $K$  is sufficiently large, the diagonalization of this effective Hamiltonian  $\tilde{H}$  provides a reliable estimate of the low-lying eigenstates. Such a procedure is the well-known Lanczos algorithm.<sup>35–37</sup> Regarding the Hamiltonian as a linear combination of eigenspace projectors, the Krylov space can only contain components already present in the initial state. Therefore, one can also obtain excited states starting from an initial state orthogonal to the lower eigenstates.

Time evolution can also be simulated based on the Krylov subspace,<sup>34,41,49–51</sup> where the Krylov vectors  $\{|v_0\rangle, |v_1\rangle, \dots, |v_{K-1}\rangle\}$  are built based on the initial state at  $t = 0$ . A general quantum state in such a subspace can be written as

$$|\Psi\rangle = \sum_{i=0}^{K-1} a_i |v_i\rangle \quad (6)$$

We represent the state as  $\vec{a} = (a_0, a_1, \dots, a_{K-1})^T$ , where  $a_i$  refers to the amplitude for each basis vector. With the help of such a formalism, the time-evolved state in the Krylov subspace is formulated as

$$\vec{a}(t = \delta) = e^{-i\delta\tilde{H}} \vec{a}(t = 0) \quad (7)$$

where  $\vec{a}(t = 0)$  is  $(1, 0, 0, \dots, 0)^T$  since the initial state is just  $|v_0\rangle$ . One can explicitly reconstruct the time-evolved quantum state by eq 6. The accuracy of this time evolution method depends on the size of the Krylov subspace, and the error is of order  $O(\delta^K)$  for a single time step  $\delta$  and thus  $O(\delta^{K-1})$  for a fixed duration.<sup>34</sup> The accuracy varies for different models: the upper error bounds are determined by the spectral width  $W$ , the step size  $\delta$ , and the subspace size  $N$ ; see Appendix A for quantitative discussion. As a practical guide, a subspace dimension of 3~10 is typically sufficient to achieve satisfactory accuracy when selecting small time step sizes, as suggested in ref 34.

When using the MPS formalism to implement these algorithms, extra errors are introduced due to the MPS truncation, particularly the loss of orthogonality of the Krylov basis. We employ the strategy proposed in ref 37 to address this issue, and we noted that the canonical orthogonalization method might also be useful.<sup>52</sup> Both these techniques aim at finding a linear combination

$$|\psi_a\rangle = \sum_{i=0}^a C_{ai} |v_i\rangle \quad (8)$$

of current Krylov vectors  $|v_i\rangle$ , so that the resulting vectors  $|\psi_a\rangle$  are well-orthogonalized. Note that we do not change the  $|v_i\rangle$  vectors; instead, we only focus on solving the matrix  $C$ . Consequently, the elements of effective Hamiltonian in the  $|\psi_a\rangle$  basis are given by

$$\tilde{H}_{ab} = \langle \psi_a | H | \psi_b \rangle = \sum_{i=0}^a \sum_{j=0}^b C_{ai}^* C_{bj} \langle v_i | H | v_j \rangle \quad (9)$$

The MPS truncation will still reduce accuracy even though the Krylov vectors are well orthogonalized; we find that restarting the Lanczos algorithm is helpful in improving the convergence. For simulating time evolution, truncation errors become significant only at very small time steps.

The Krylov method's most expensive and memory-intensive part is obtaining the Krylov vectors whose core step is to compute  $H|v_i\rangle$ . Conventionally, one multiplies the Hamiltonian's MPO with an MPS, resulting in an intermediate MPS with a large maximum bond dimension  $O(L^2M)$  for ab initio molecular Hamiltonians. The memory cost to store such an intermediate MPS scales as  $O(L^5M^2)$  for all sites, which could exhaust the available memory for even small-sized systems. Second, the high computational cost of compressing the intermediate MPS of  $H|v_i\rangle$  is another bottleneck of constructing the Krylov subspace. One must compress the bond dimensions of  $H|v_i\rangle$  back to smaller target bond dimensions to avoid the exponential increase in the next multiplications. To achieve fully controllable and highly accurate truncation, one typically employs the SVD method. In this approach, one first brings the intermediate MPS into canonical form using QR decomposition and then truncates the bonds by SVD, as depicted in Figure 3. Given that one can easily read off the Schmidt values from the mixed-canonical form, the truncation can be performed with a desired accuracy.<sup>42,53</sup> The QR decompositions are the main contributors to the cost of compression. The computational complexity of QR decomposition for an MPS tensor with bond dimension  $O(L^2M)$  is  $O(L^6M^3)$ , leading to a total complexity across all sites as high as

$O(L^7M^3)$  which makes it challenging to apply the Krylov method on large molecules.

There are also alternative MPS compression methods. The zip-up method<sup>54</sup> is more efficient, but since the algorithm works on a nonorthogonalized basis, the error is not fully controlled. The variational method requires a proper initial guess. Otherwise, one needs a large number of iterations and sweeps.<sup>42</sup> The recently proposed density matrix method<sup>55</sup> provides another fully controllable compression scheme that merits further study in future research. Our THC-MPO discussed in this paper can improve most of these MPS compression schemes when simulating molecular Hamiltonians, see Appendix B for more details; we focus on the traditional SVD method in this paper. We would like to emphasize that our method is compatible with all the MPS compression methods, and we can always check whether our method can be integrated into them when a newer compression scheme is proposed.

**2.3. The THC Factorization.** Employed widely in the simulation of molecular systems already,<sup>46,56–58</sup> the tensor hypercontraction (THC) proposed by Hohenstein et al.<sup>46,56,57</sup> approximates the two-electron integrals  $v_{pqrs}$  as

$$v_{pqrs} \approx \sum_{\mu,\nu=1}^N \chi_p^\mu \chi_q^\mu \zeta^{\mu\nu} \chi_r^\nu \chi_s^\nu \quad (10)$$

for all  $p, q, r, s \in \{1, \dots, L\}$ , as illustrated in Figure 1.

Currently, several relatively mature algorithms exist to obtain these tensors. The original papers proposed the PF-THC<sup>44</sup> and LS-THC<sup>45</sup> methods as algorithms. Subsequently, the interpolative separable density fitting (ISDF) method<sup>59</sup> enhanced the computational efficiency and improved the approximation accuracy. In density-fitting (DF),<sup>60–62</sup> one approximates the product of two orbitals as

$$\rho_{pq}(r) := \phi_p(r)\phi_q(r) \approx \sum_{\mu=1}^{N_a} C_{pq}^\mu P_\mu(r) \quad (11)$$

where  $P_\mu$  for  $\mu = 1, 2, \dots, N_a$  are auxiliary basis functions. The idea of ISDF is that if we approximate  $\rho_{pq}$  by interpolation, the THC factorization can directly be obtained<sup>59</sup>

$$\rho_{pq}(r) \approx \sum_k \rho_{pq}(r_k) F_k(r) = \sum_k \phi_p(r_k) \phi_q(r_k) F_k(r) \quad (12)$$

where  $r_k$  are selected grid points in the Becke scheme. The selection is implemented by interpolative decomposition, aimed at choosing a limited number of rows to approximate  $\rho_{pq}(r_k)$  interpreted as a  $N_g \times L^2$  matrix, where  $N_g$  is the total number of Becke grid points. Since the row indices represent individual grid points, the procedure can also be interpreted as discarding less important grid points. Their importance is revealed by randomized QR decomposition with column-pivoting.<sup>59,63</sup> We then determine fit functions  $F_k$  after we obtain selected grid points. The fit functions are chosen as auxiliary basis functions  $P_\mu$  ref 59, but in this work, we obtained them following the strategy introduced in LS-THC,<sup>45</sup> as suggested in ref 56.

To improve the accuracy of the THC decomposition, we minimize the relative error

$$\epsilon_V = \frac{\|v_{pqrs} - \sum_{\mu,\nu=1}^N \chi_p^\mu \chi_q^\mu \zeta^{\mu\nu} \chi_r^\nu \chi_s^\nu\|}{\|v_{pqrs}\|} \quad (13)$$

where  $\|v_{pqrs}\| = \sqrt{\sum_{pqrs} |v_{pqrs}|^2}$  denotes the Frobenius norm. It is essentially an optimization problem, and we carried it out using the Adam optimizer,<sup>64</sup> implemented in Optax.<sup>65</sup> Although Adam was introduced for stochastic optimization problems, its adaptive moment mechanism also converges reliably in deterministic settings. In our tests, Adam converged faster than classical optimizers such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm,<sup>66</sup> so we adopted it for the optimization step. It appears that there are two matrices, namely  $\chi$  and  $\zeta$ , to be optimized. However, since  $\zeta$  can be obtained from  $\chi$  by LS-THC, the number of free parameters is reduced to the entries of  $\chi$ . Exemplified by the hydrogen chain, we first carry out the optimization by 1000 rounds with a learning rate of 0.001, followed by another 1000 rounds with a learning rate of 0.0005. In the numerical experiment, we reach the acceptable chemical accuracy (1.6 mHartree) with  $N = 4L$  for the water molecule  $H_2O$ ,  $N = 3L - 3$  for the hydrogen chain  $H_{10}$ , and  $N = 4.5L$  for the Ammonia molecule  $NH_3$ , all in the STO-6G basis set. We ensure the accuracy reaches chemical accuracy by comparing it with the energy obtained via FCI results from PySCF. Regarding our future studies in larger systems, we noticed that previous studies have demonstrated that the THC rank  $N$  typically exhibits near-linear scaling with respect to the system size  $L$ .<sup>44–46,56,59,63,67–69</sup> For instance, there are 76 spatial orbitals needed for the active-space model of the FeMoco system proposed by Li et al.,<sup>70</sup> and the THC rank of 450 is enough to reach the chemical accuracy;<sup>67</sup> for hydrogen chain of  $L$  atoms ( $L$  spatial orbitals in STO-6G basis) with distances of 1.4 Å, a THC rank of  $3L - 3$  is sufficient to achieve the accuracy of  $5 \times 10^{-5}$  Hartree per atom.<sup>71</sup> We remark that this part is a preprocessing step. If a prepared THC decomposition is available (e.g., from other work or a database), we can apply our THC-MPO method directly.

### 3. THE KRYLOV METHOD BASED ON THC

**3.1. Constructing MPOs Using THC.** In this section, we first show how to use the THC factorization to construct a special representation of the molecular Hamiltonian (THC-MPO). Then, we will utilize the THC-MPO in Krylov methods and discuss its advantages. We focus on the challenging Coulomb term here. An MPO of the kinetic term  $T$  can be easily constructed following the strategy introduced in ref 11, and we will also discuss the kinetic term in Section 3.2.

Inserting the THC factorization eq 10 into the Coulomb term  $V$ , one immediately arrives at

$$V \approx \frac{1}{2} \sum_{\mu,\nu=1}^N \sum_{\sigma,\sigma' \in \{\uparrow, \downarrow\}} G_{\mu\sigma,\nu\sigma'} \quad (14)$$

where  $G_{\mu\sigma,\nu\sigma'}$  is defined as

$$G_{\mu\sigma,\nu\sigma'} = \zeta^{\mu\nu} \left( \sum_{p=1}^L \chi_p^\mu a_{p,\sigma}^\dagger \right) \left( \sum_{q=1}^L \chi_q^\mu a_{q,\sigma'} \right) \left( \sum_{r=1}^L \chi_r^\nu a_{r,\sigma'}^\dagger \right) \left( \sum_{s=1}^L \chi_s^\nu a_{s,\sigma'} \right) \quad (15)$$

Each subterm (exemplified by  $\sum_s \chi_s^\nu a_{s,\sigma'}$ ) in  $G_{\mu\sigma,\nu\sigma'}$  can explicitly be converted to an MPO as follows

$$W[s] = \begin{pmatrix} I & \chi_s^\nu a_{s,\sigma'} \\ 0 & I \end{pmatrix}, \quad s = 2, \dots, L-1 \quad (16a)$$

and the first and last tensors

$$W[1] = (\chi_1^\nu a_{1,\sigma'}), \quad W[L] = \begin{pmatrix} \chi_L^\nu a_{L,\sigma'} \\ I \end{pmatrix} \quad (16b)$$

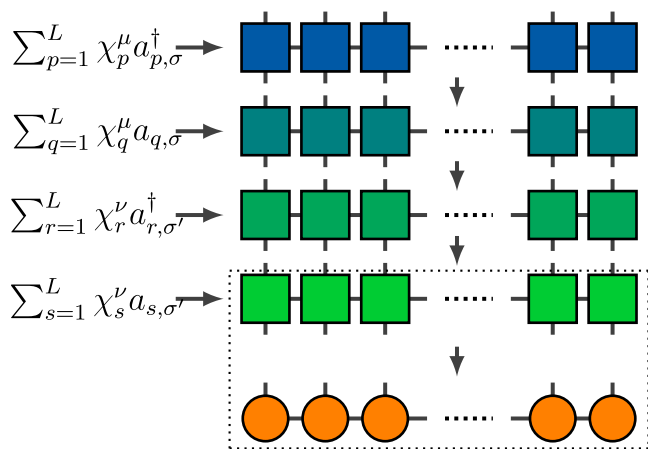
One can contract the  $W$  matrices sequentially to verify the correctness of the construction. It is worth noting that the bond dimension of  $W[s]$  is always only 2, independent of the system size  $L$ .

In this work, we follow the convention of treating each spatial orbital as a single site in the MPS. When implementing a corresponding MPO numerically using the Jordan-Wigner transformation,<sup>72</sup> we replace the Fermionic operators with their bosonic counterparts and substitute the identities in each  $W$  at position (1,1) by Pauli-Z operators (to account for Fermionic sign factors)

$$W[s] = \begin{pmatrix} Z \otimes Z & \chi_s^\nu b_{s,\sigma'} \\ 0 & I_4 \end{pmatrix} \quad (17)$$

where  $I_4$  denotes the identity matrix of size  $4 \times 4$ , and  $b_{s,\sigma'}$  is defined as the local annihilation operator for spin  $\sigma'$  of size  $4 \times 4$ , for which we detail these in Appendix C.

Following the strategy above, one can analogously construct MPOs for the other three subterms:  $\sum_p \chi_p^\mu a_{p,\sigma'}^\dagger$ ,  $\sum_q \chi_q^\mu a_{q,\sigma}$  and  $\sum_r \chi_r^\nu a_{r,\sigma'}^\dagger$ . The entire MPO of  $G_{\mu\sigma,\nu\sigma'}$  is thus the product of the MPOs of these four subterms as shown in Figure 4, and the



**Figure 4.**  $G_{\mu\sigma,\nu\sigma'}$  in eq 15 is composed of four layers of MPOs (square tensors) with bond dimension 2. As specified by the arrows, we contract and compress the layers one at a time with the MPS (orange tensors).

scalar  $\zeta^{\mu\nu}$  can be absorbed into  $W[1]$  at the first site. Therefore, the MPO of  $G_{\mu\sigma,\nu\sigma'}$  likewise has a constant bond dimension. The whole MPO of the Coulomb term is thus the summation of MPOs of sub-Hamiltonians  $G_{\mu\sigma,\nu\sigma'}$ , but to calculate  $V|\psi\rangle$  in compressed MPS form, we will refrain from merging them into a large MPO, see below.

**3.2. Krylov Method Using THC-MPO.** As discussed in Section 2.2, the essential step in Krylov methods is multiplying  $H$  with  $|\psi\rangle$ . Here, we focus on the Coulomb term  $V$  in  $H = T + V$ . We present how to take advantage of our THC-MPO to execute

the multiplication and subsequent compression. With the help of eq 14, we can write  $V|\psi\rangle$  as

$$V|\psi\rangle \approx \frac{1}{2} \sum_{\mu,\nu=1}^N \sum_{\sigma,\sigma' \in \{\uparrow,\downarrow\}} G_{\mu\sigma,\nu\sigma'} |\psi\rangle \quad (18)$$

where we apply each sub-Hamiltonian to  $|\psi\rangle$  and sum the resulting states up. Instead of manipulating large matrices, we compute  $V|\psi\rangle$  via the small MPOs of  $G_{\mu\sigma,\nu\sigma'}$ .

For each term  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$ , we execute multiplication and compression for each elementary MPO (layers shown in Figure 4) sequentially instead of treating  $G_{\mu\sigma,\nu\sigma'}$  as a whole. Each compression returns the bond dimension to  $M$  so that the maximum bond dimension is only  $2M$  during the calculation (since the MPO bond dimension for each layer is 2). Beginning with  $G_{1\uparrow,1\uparrow}|\psi\rangle$ , we add each subsequent  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$ . Such an MPS addition likewise leads to an intermediate MPS of bond dimension  $2M$ , which is still cheap to store and compress. In summary,  $O(LM^2)$  memory is required to store the largest intermediate MPS, which is less by a factor  $O(L^4)$  compared to a conventional MPO algorithm. In addition, the memory for storing  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$  and  $G_{\mu\sigma,\nu\sigma'}$  is immediately released after adding  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$  to others. Implementing eq 18 is flexible regarding the order of additions.

Another optimization can be achieved by reusing intermediate results. We first notice that one can write  $G_{\mu\sigma,\nu\sigma'}$  as

$$G_{\mu\sigma,\nu\sigma'} = \zeta^{\mu\nu} G_{\mu\sigma} G_{\nu\sigma'} \quad (19)$$

where

$$G_{\nu\sigma'} = \left( \sum_{r=1}^L \chi_r^\nu a_{r,\sigma'}^\dagger \right) \left( \sum_{s=1}^L \chi_s^\nu a_{s,\sigma'} \right) \quad (20)$$

and similarly for  $G_{\mu\sigma}$ . It indicates that for two sub-Hamiltonians  $G_{a\tau,\nu\sigma'}$  and  $G_{b\kappa,\nu\sigma'}$  that share the same latter two indices, the term  $G_{\nu\sigma'}$  can be factored out. Therefore, the intermediate state  $G_{\nu\sigma'}|\psi\rangle$ , which is obtained from compressing the two elementary MPOs (layers) in  $G_{\nu\sigma'}|\psi\rangle$ , can be reused. By applying this optimization, we reduce the computational cost by nearly half. Alg. 1 One includes all these steps and illustrates the overall algorithm as pseudocode.

In practice, we must take the kinetic term  $T$  into account as well. The conventional MPO representation of  $T$  has bond dimension  $O(L)$ , which leads to an overall memory requirement of  $O(L^3M^2)$  to store  $T|\psi\rangle$  as MPS (without compression). We can improve on that situation using similar ideas as for the interaction term: We perform a spectral decomposition of  $(t_{pq})$  in eq 1 and construct a sum of products of elementary MPOs with bond dimension 2. Therefore, the memory requirement can be reduced to  $O(LM^2)$  for obtaining a compressed MPS. We explain these steps in detail in Appendix D.

**Algorithm 1** Computing  $V|\psi\rangle$  based on the THC-MPO

```

Input: Initial state  $|\psi\rangle$  as MPS, all sub-Hamiltonians  $G_{\mu\sigma,\nu\sigma'}$  as MPOs
Output:  $V|\psi\rangle$  as compressed MPS
 $|\phi\rangle = 0$ 
for  $\nu \in \{1, \dots, N\}$ ,  $\sigma' \in \{\uparrow, \downarrow\}$ 
  Initialization:  $|\psi_{\nu\sigma'}\rangle = |\psi\rangle$ 
  for each elementary MPO in  $G_{\nu\sigma'}$ 
     $|\psi_{\nu\sigma'}\rangle = \text{elementary MPO} \cdot |\psi_{\nu\sigma'}\rangle$ 
    Compress  $|\psi_{\nu\sigma'}\rangle$ 
  end for
  for  $\mu \in \{1, \dots, N\}$ ,  $\sigma \in \{\uparrow, \downarrow\}$ 
    Initialization:  $|\psi_{\mu\sigma,\nu\sigma'}\rangle = \frac{1}{2}G^{\mu\nu}|\psi_{\nu\sigma'}\rangle$ 
    for each elementary MPO in  $G_{\mu\sigma}$ 
       $|\psi_{\mu\sigma,\nu\sigma'}\rangle = \text{elementary MPO} \cdot |\psi_{\mu\sigma,\nu\sigma'}\rangle$ 
      Compress  $|\psi_{\mu\sigma,\nu\sigma'}\rangle$ 
    end for
     $|\phi\rangle = |\phi\rangle + |\psi_{\mu\sigma,\nu\sigma'}\rangle$ 
    Compress  $|\phi\rangle$ 
  end for
end for
Return  $|\phi\rangle$ 

```

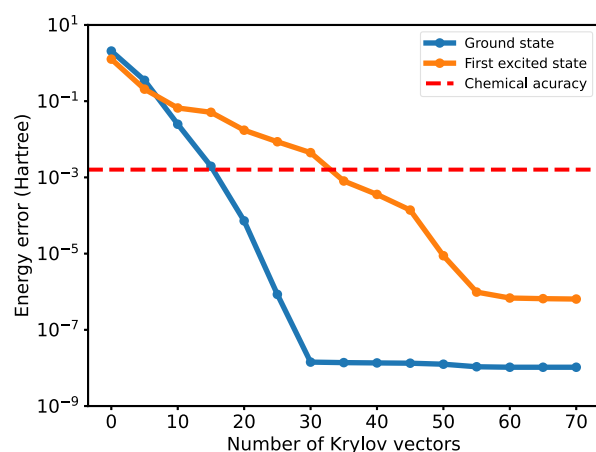
## 4. NUMERICAL RESULTS AND RESOURCE ESTIMATION

**4.1. Ground- and Low-Lying States Finding.** To benchmark the MPS-based Lanczos algorithm using our THC-MPO, we apply our method to the water molecule  $\text{H}_2\text{O}$  and the hydrogen chain  $\text{H}_{10}$  using the STO-6G basis. The electronic integrals and FCI reference are calculated by PySCF,<sup>73,74</sup> the tensor network calculation is implemented with PyTeNet,<sup>75</sup> in which Abelian quantum number conservation laws (electron number and spin) are enforced. We chose these relatively small systems because they allow for easier analysis of error sources and algorithmic behavior. But even so, we will see that the memory advantage has been fully verified. To fully explore our approach's computational complexity capabilities, we plan to switch to high-performance computers and utilize parallel computing to benchmark them for large systems in the future, as discussed in Section 4.5.

We first present the results of the water molecule using the STO-6G basis, which leads to 7 spatial orbitals (14 spinor orbitals). In this case, we limit the maximum MPS bond dimension for the Krylov vectors to 30. The THC rank  $N$  for  $\text{H}_2\text{O}$  is set to 28, resulting in the Frobenius norm error  $\|v - v'\| \approx 3 \times 10^{-11}$ , where  $v'$  is the Coulomb term reconstructed by THC tensors according to eq 10.

While the Lanczos algorithm performs well with a random initial state, selecting a proper initial state can significantly speed up convergence. In practice, we start from a state close to the target state, obtained from a heuristic guess or a low-cost algorithm. In this work, we simply use the Hartree–Fock state as the initial state for ground state finding, where paired electrons occupy the five lowest-energy molecular spatial orbitals. Additionally, we excite the highest occupied orbital in the Hartree–Fock state to serve as the initial state for finding the first excited state since the Hartree–Fock state is orthogonal to the ground state. As shown in Figure 5, we obtain the ground state and the first excited state within acceptable chemical accuracy (1.6 mHartree) using only 15 and 35 krylov vectors, respectively. The first excited state energy converges much slower because the gap between the first excited state and the second excited state is smaller than the one between the ground state and the first excited state. The energy error is obtained by comparison with the numerically exact value calculated by the FCI method in PySCF.

However, a high-accurate THC decomposition as above is unnecessary because, first, the total error is also bounded by MPS truncation. Second, a high THC rank will increase the computational cost. A way to reduce computational cost at the



**Figure 5.** Convergence of the water molecule's ground and first excited state calculation using the Lanczos algorithm based on our THC-MPO. We restart the iteration at the 45th step for the first excited state finding.

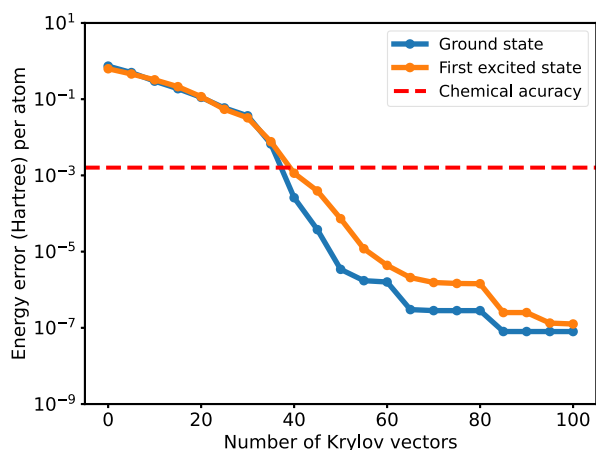
expense of accuracy is using a smaller THC rank  $N$ . To explore this possibility and quantify the resulting error, we study the hydrogen chain of ten atoms  $\text{H}_{10}$  with distances of 1.4 Å in the STO-6G basis, which leads to 10 spatial orbitals. Allowing a ground state energy error of  $3 \times 10^{-6}$  Hartree per atom, its THC rank is as low as 27.

The MPS bond dimensions for representing Krylov vectors are capped at 250. Like the previous example, we again use the Hartree–Fock ground and single-excited states as the initial states for the Krylov method. Interestingly, when using the exact Hamiltonian to calculate the energy expectation value for the approximated ground state

$$E_{\text{avg,Krylov}} = \langle \psi_{\text{approx}} | H_{\text{exact}} | \psi_{\text{approx}} \rangle \quad (21)$$

where  $|\psi_{\text{approx}}\rangle$  is obtained by the THC-MPO-based Krylov method and  $H_{\text{exact}}$  is the exact Hamiltonian, the resulting energy error is smaller than the error introduced by the THC approximation. While the THC approximation leads to an energy error of around  $3 \times 10^{-6}$  Hartree per atom, we can obtain the ground and first excited state with energy error  $\sim 10^{-7}$  Hartree per atom, as illustrated in Figure 6. This indicates that accurate results could still be obtained using the THC-MPO, even when choosing a smaller THC rank that introduces non-negligible errors.

The results also suggest that although a large number of truncations is required to implement eq 18, the MPS truncations introduce only a minor error. Intuitively, assuming that each  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$  term admits a relative error  $\epsilon$ , the summation of them also admits a relative error  $\epsilon$ , especially when allowing larger bond dimensions during reduction (and compress the final bond dimension back to  $M$ ). Despite errors introduced by compressing the terms  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$ , additional errors also arise during the subsequent reduction (summation) process. Again, Assuming that pairwise addition-compression of these terms each results in a relative error  $\epsilon'$ , it follows that each hierarchical reduction level similarly contributes a relative error of  $\epsilon'$ . Given there are  $\log(4N^2) = 2 \log(N) + 2$  reduction levels (as shown in Figure 11), the cumulative error can thus be bounded by  $\epsilon' \sqrt{2 \log(N) + 2}$ , assuming that the errors are all in different “directions”. Furthermore, since the reduction only manipulates the canonical MPS, increasing the bond dimension for this process will significantly enhance the accuracy without costing



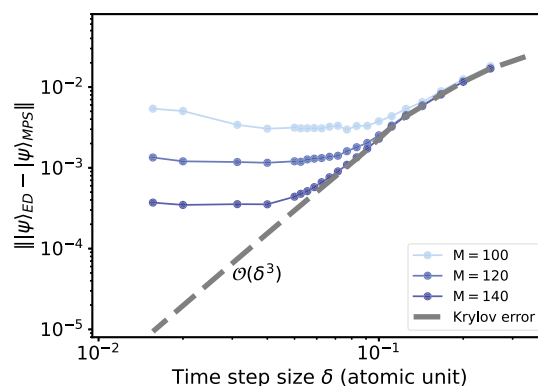
**Figure 6.** Energy convergence for the hydrogen chain  $H_{10}$ . We restart the iteration at the 30th, 60th and 80th steps to improve the convergence. The Krylov space is obtained via THC-MPO, while the resulting energy of the approximated ground- and low-lying states is calculated according to the exact Hamiltonian.

too much computational resources. Additionally, since the final MPS  $|\psi\rangle$  closely approximates the ground state (or other low-lying eigenstates), the bond dimension required to accurately represent the resulting MPO-MPS is expected to remain moderate. Therefore, many subterms should not contribute a much larger error.


**4.2. Time Evolution Using Global Krylov Method.** We also study the Krylov subspace time evolution based on our THC-MPO, where we set the subspace dimension to 4, leading to a single step error  $O(\delta^4)$  and total error  $O(\delta^3)$  for a fixed duration  $T$  for a single time step size  $\delta$ . We apply the global Krylov method to the Ammonia molecule  $NH_3$  in the STO-6G basis, which leads to 8 spatial orbitals (16 spinor orbitals). The THC rank  $N$  for  $NH_3$  is set to 36, resulting in the Frobenius norm error  $\|\nu - \nu'\| \approx 4 \times 10^{-12}$ . The initial state is defined as  $|\psi(t=0)\rangle = a_{3,\uparrow}|\psi_0\rangle$ , where a spin-up electron is annihilated from the third spatial orbital of the ground state. Three factors determine the accuracy: SVD cutoff (bond dimension limitation), time step size  $\delta$ , and the THC error from the THC factorization. The THC error is negligible for the  $NH_3$  molecule since the THC rank  $N = 4.5L$  results in a very accurate approximation.

As depicted in Figure 7, we measure the time evolution error for duration  $T = 1$  atomic unit (a.u.) for different step sizes  $\delta$  and maximum bond dimensions. The behavior of the errors can be explained well: As expected, the Krylov error dominates the overall error for larger time step sizes. Conversely, the  $O(\delta^3)$  scaling leads to small Krylov errors when the step size  $\delta$  is reduced, causing the truncation error to dominate the overall error. To balance efficiency and accuracy, one can reach a sweet spot where the truncation error is comparable to the Krylov error. In Figure 7, it occurs where the total error curve converges with the Krylov error curve. For the case  $N = 4$ , one can observe that when setting  $M = 140$ , a step size of  $\delta \in [0.05, 0.1]$  a.u. appears to be optimal.

It is also meaningful to enlarge the Krylov subspace size and examine whether it would enhance the accuracy as predicted. Specifically, we calculate the time evolution for duration  $T = 41.3$  au (1 fs) with  $M = 140$ , for both subspace size  $N = 4$  and  $N = 5$ ; the step size is set as  $\delta = 0.1$ . As illustrated in Figure 8, the cumulative wave function error is 0.131 for  $N = 4$  but only 0.011



**Figure 7.** Time evolution errors of duration  $T = 1$  atomic unit for  $NH_3$  when using the global Krylov method based on our THC-MPO for various bond dimensions  $M$ , plotted as functions of the time step sizes. The errors are measured by the distance  $\|\psi_{ED} - |\psi\rangle_{MPS}\|$  between the states from our numerical method and the reference time-evolved quantum state obtained by ED. This metric is also used to measure the Krylov errors.

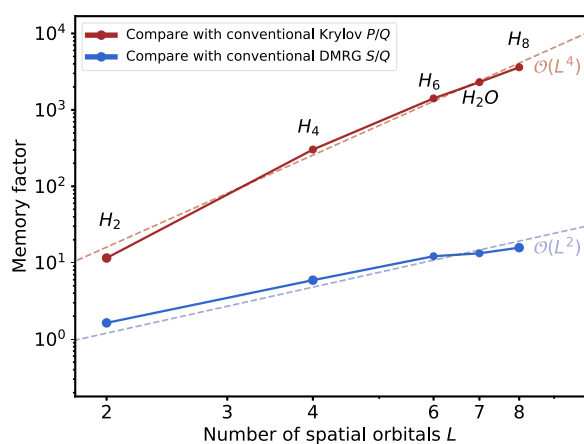


**Figure 8.** Time evolution errors of duration  $T = 41.3$  atomic unit ( $\approx 1$  fs) for  $NH_3$ , plotted as functions of the evolution time  $t$ .

for  $N = 5$ , indicating that the accuracy is enhanced by a factor of 10 when adding one more vector. For this time step size  $\delta = 0.1$ , such a reduction is consistent with the expected error scaling  $O(\delta^N)$ .

**4.3. Memory Consumption Comparison.** One of the advantages of our THC-MPO method is the significantly reduced memory cost by a factor  $O(L^4)$ . We separately monitored memory consumption to store intermediate MPS in the Krylov algorithm based on the conventional MPO and the THC-MPO to test this prediction in our numerical experiments. We denote memory consumption when using conventional MPOs as  $P$ , and when using THC-MPOs as  $Q$ . Figure 9 shows the quotient  $P/Q$  for the water molecule and hydrogen chains. By studying systems of different sizes, one clearly observes the predicted scaling difference. For example, considering the hydrogen chain of eight atoms, the memory required for storing an intermediate state  $H|\psi\rangle$  calculated with the conventional MPO amounts to 12,586 MB. In contrast, only 3.49 MB is needed when using the THC-MPO method, leading to a factor  $P/Q$  as large as 3606. We measure the memory cost by saving these intermediate MPS in HDF5 files and directly accessing their sizes. This case's maximum bond dimension is 80, and we utilized double-precision complex numbers. Such a large memory usage is even too large to apply the Krylov methods on the  $H_8$  molecule. Therefore, in this respect, the  $H_{10}$  example has already proved the advantage over the original Krylov methods.

The Krylov method based on THC-MPO also outperforms the DMRG algorithm in terms of memory usage. Theoretically,



**Figure 9.** Comparison of memory consumption for the Krylov method based on the conventional MPO versus the THC-MPO (red), as well as DMRG algorithm versus the Krylov method based on the THC-MPO (blue). The maximum bond dimensions are 4, 16, 60, 70, and 80, respectively. The dotted lines representing  $O(L^4)$  and  $O(L^2)$  demonstrate that the scaling of  $P/Q$  and  $S/Q$  aligns well with the theoretical prediction.

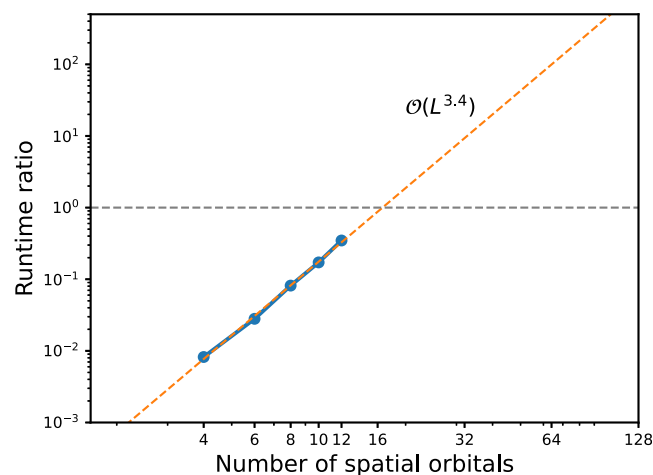
the DMRG algorithm requires  $O(L^3M^2)$  memory (to store the left and right environment blocks), which is  $O(L^2)$  times larger than the THC-MPO-based Krylov method. As shown in Figure 9, we numerically compare  $Q$  with the memory consumption  $S$  for the DMRG algorithm, using the same MPS bond dimensions. The results suggest that our method requires significantly less memory than the DMRG algorithm, and the observed values agree with the theoretically predicted  $O(L^2)$  scaling. Since the TDVP method could be implemented within a framework similar to the DMRG algorithm, our method also outperforms TDVP in terms of memory consumption when simulating time evolution. We do not continue to increase the system size to measure more cases since the memory usage for conventional MPO methods rapidly exceeds our available memory (32 GB), and the results shown in Figure 9 are sufficient to demonstrate the memory advantage of our method. Due to memory constraints, the runtime comparison for large systems is also infeasible.

**4.4. Computational Complexity Estimation.** Besides memory consumption, the global Krylov methods based on our THC-MPO also perform better in terms of computational cost scaling than global Krylov methods using conventional MPOs. Here, we only present the summary; see Appendix E for a detailed derivation.

The primary contributor to the overall cost is obtaining compressed Krylov vectors. When using conventional MPO construction, renormalization is the most expensive step in compression. Since we have to handle the intermediate MPS with bond dimension  $O(L^2M)$ , the renormalization has an overall complexity of  $O(L^7M^3)$  for all sites. In contrast, for global Krylov methods utilizing THC-MPO, we only need to deal with MPS with bond dimension  $O(M)$  since the bond dimension of each layer in the subterms  $G_{\mu\sigma,\nu\sigma}$  is only 2. Therefore, it costs  $O(LM^3)$  to obtain  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$  as compressed MPS, leading to an overall cost of  $O(L^3M^3)$  for all  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$  (when assuming that the THC rank  $N$  scales linearly with  $L$ ). This computational cost has a large prefactor; it could be around  $10^5$  when taking hydrogen chains as an example. The large prefactor leads to

longer run times for small molecules; for example, on a 13th-generation Intel Core i7-1355U CPU (12 cores, 1.70 GHz base frequency), the average wall-time runtime is 56 s for computing a Krylov vector for our  $H_2O$  case, and 26 min is needed for a Krylov vector for our  $H_{10}$  case. The relatively slow performance can be attributed to our current implementation, which is not yet performance-oriented, and much of the multicore capacity is left idle. Switching to a more efficient programming language and applying further optimizations should substantially improve the runtime. We are actively developing a more efficient implementation.<sup>76</sup>

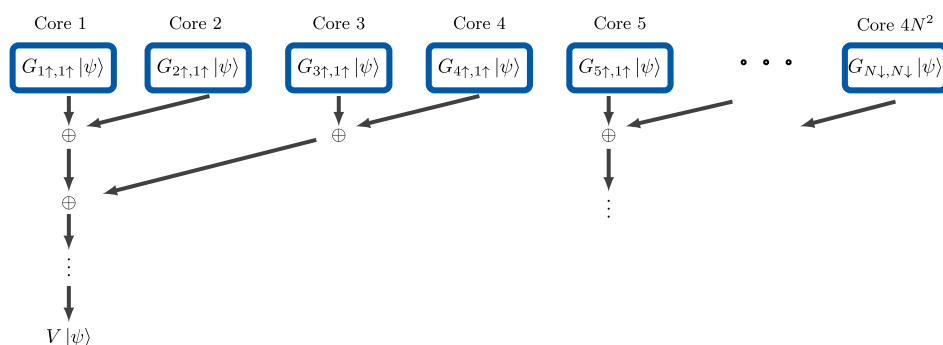
Nevertheless, we expect an advantage for medium- and large-sized molecules due to their promising scaling gap  $O(L^4)$ . To quantify the computational benefit of the THC-MPO method, we benchmark the wall-clock time required to compute the compressed MPS  $H|\psi\rangle$ . Denote the runtime with the explicit (original) Hamiltonian by  $t_o$  and that with THC-MPO by  $t_{\text{THC}}$ , we report the runtime ratio  $t_{\text{THC}}/t_o$ , i.e., how many times faster the THC-MPO method is on hydrogen chains comprising 4–12 orbitals, as shown in Figure 10. Even though we only use a small



**Figure 10.** Log–log plot of the runtime ratio  $t_{\text{THC}}/t_o$  as a function of the number of spatial orbitals. The data points (blue markers) are fitted with a linear least-squares regression (orange dotted line). While the overall fitted slope is 3.4, the slope extracted from the last two data points rises to 3.8, indicating that the scaling trend converges toward our analytical prediction of 4 as the system size increases.

bond dimension  $M = 30$ , the benchmark using conventional MPOs is still infeasible when  $L \geq 14$  due to the large memory requirements. The curve admits a scaling  $O(L^{3.4})$ , which is not as perfect accurate as our predicts, but the obtained runtime ratio can clearly reveal the ratio trend, that is, our THC-MPO method should be more efficient when the system size goes beyond 20.


**4.5. A Natural Scalable Parallelization Scheme.** Parallel computing has been effectively integrated into DMRG algorithms for quantum chemistry to take advantage of high-performance computing platforms. This integration has significantly enhanced the ability to study large molecular systems; various parallel schemes were proposed,<sup>11,77–82</sup> and notable open-source packages like Block2 were developed.<sup>83</sup> The Krylov method based on our THC-MPO can straightforwardly use the potential of parallel computing: to obtain  $H|\psi\rangle$  following eq 18, each of the  $4N^2$  subterms can be calculated and compressed independently, and the summation of these subterms can also be performed in parallel by a reduction.



**Figure 11.** Parallelization scheme for applying the Coulomb operator  $V$  to a state  $|\psi\rangle$  in MPS form according to eq 18. Each core is first assigned the subtask to compute and compress an intermediate state  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$  as MPS. These are then aggregated through a reduction process. For simplicity, we assume that the high-performance computer is able to perform at least  $4N^2$  cores; otherwise, a single core would handle several of the  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$  states.

More specifically, we propose a parallelism scheme as illustrated in Figure 11. For each core, we first assign the task of computing and compressing one (or several) subterms  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$ . The power of multiple cores can be perfectly utilized for this part. After this step, we add and compress these terms pairwise in parallel. It appears that some computational resources are idling during such a process, but the compression can utilize multiple cores for parallel computation when using packages like multithreaded LAPACK implementations.<sup>84</sup> Because the SVD and QR decomposition can be significantly sped up by parallel computing,<sup>85,86</sup> the reduction part can utilize the power of parallel computing as well. Also, since each compressed term  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$  has already been canonical form (up to a factor), the MPS addition-compression process does not contain the QR decomposition, which makes the reduction inexpensive. Another bottleneck of parallel computing is communication;<sup>77</sup> an extra advantage of our parallel scheme is that communication only occurs during the reduction.

As a preliminary demonstration (with more advanced systems and larger molecules planned for future work), we benchmark our parallel scheme on a 112-core node using OpenMP.<sup>87</sup> Each thread (core) is tasked with calculating and compressing  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$ . As displayed in Figure 12, we compare the runtime for computing  $H|\psi\rangle$  with a single thread versus  $K$  threads. The



**Figure 12.** Speedup is calculated by comparing the runtime using  $K$  threads with that of a single thread. Although slightly imperfect due to factors such as cache contention, the results indicate that our parallel scheme efficiently leverages available computational resources.

results indicate near-ideal speedups when multiple threads are utilized. Extending this approach to multiple nodes should also yield near-linear scaling because each node can achieve this speedup independently, and communication among nodes is only required once all nodes have completed their tasks.

Due to such an efficient and scalable parallel scheme, the parallel runtime scales as  $O(LM^3)$  under ideal parallelization conditions with  $4N^2$  available cores efficiently. Currently, advanced tensor network methods in quantum chemistry (e.g., DMRG) can utilize thousands of cores efficiently,<sup>77,82</sup> the Krylov methods (Lanczos algorithm and Krylov time evolution method) based on THC-MPO have the potential to leverage cores scaling as  $4N^2$  with high efficiency, making it possible to surpass the current state of the art in CPU utilization. Note that each subtask shown in Figure 11 can also be implemented by multiple cores (e.g., a node), thereby further increasing the number of cores we can efficiently utilize and decreasing the reduction depth.

## 5. CONCLUSIONS

The THC-MPO approach allows us to implement Krylov subspace methods, e.g., the Lanczos algorithm and the global Krylov method for time evolution, with reduced memory usage and lower computational cost scaling. When compared to the Krylov method based on the conventional MPO representation, the memory advantage of THC-MPO is apparent, even for the smallest molecules. Moreover, it outperforms popular methods like DMRG and TDVP in terms of memory consumption, suggesting that THC-MPO can potentially enable simulations of even larger systems than currently reachable by DMRG or TDVP. While the benefit of computational cost is not immediate for small molecules due to large prefactors, we expect that the improvement will become significant for moderate and large molecular systems. We emphasize that the THC-MPO is essentially a proper decomposition. A group of MPOs with a small bond dimension can also be achieved by building an MPO for each term in the molecular Hamiltonian, but the final computational complexity will be as large as  $O(L^7)$  if we do so.

A cornerstone of our work is the compressed THC representation of the two-body integral tensor  $v$ . A promising research direction (complementary to the present study) could be the exploitation of sparsity structures of  $v$ , for example, due to localized orbitals or wavelet-type orbitals supported on a fine grid. Also, the tensor  $v_{pqrs}$  (which originates from overlap integrals) is symmetric with respect to interchanges of  $p \leftrightarrow q$ ,  $r \leftrightarrow s$ , and  $(p, q) \leftrightarrow (r, s)$ . The symmetries are passed on to the

THC representation. It is worth exploring how to exploit such symmetries in our approach. We also noticed that our THC-MPO could help enable the computation of spectral functions<sup>52,88,89</sup> for large-size molecular Hamiltonians when combining with the Chebyshev expansions, where multiplication-compression  $H|\psi\rangle$  also remains a major bottleneck. With our THC-MPO method and optimized parallel computing implementations, our primary plan is to explore these ideas and the reachable system sizes in future works.

## APPENDIX

### A. Error Estimation for Krylov Subspace Time Evolution

As discussed by Hochbruck and Lubich,<sup>90</sup> the error  $\varepsilon_N$  incurred by the  $N$ th-order Krylov approximation of the time-evolution operator  $e^{-iHt}$  is bounded by

$$\varepsilon_N \leq \begin{cases} 10e^{-4N^2/(SW\delta)}, & \text{if } \sqrt{W\delta} \leq N \leq \frac{W\delta}{2}, \\ 10\left(\frac{W\delta}{4}\right)^{-1} e^{-W\delta/4} \left(\frac{eW\delta}{4N}\right)^N, & \text{if } N \geq \frac{W\delta}{2}, \end{cases} \quad (22)$$

where  $W$  is the spectral width, and  $\delta$  is the time step size. Even though the original formula is derived for a Hermitian negative semidefinite matrix  $A$  with eigenvalues in the range  $[-W, 0]$ , one may equivalently simulate the shifted operator  $A + \frac{W}{2}I$ , which only results in a global phase  $e^{-iW\delta/2}$ .

We can estimate the theoretical error upper bounds for different Krylov subspace sizes according to eq 22. A satisfactory accuracy (error  $<10^{-4}$ ) for a single-step evolution can be obtained for  $N = 4$  with step size  $W\delta < 0.2$ , or for  $N = 5$  when  $W\delta < 0.5$ . Higher precision can be achieved either by enlarging the subspace or by reducing the time step size. For instance, the accuracy level  $10^{-7}$  can be achieved by time step size  $W\delta = 0.1$  when  $N = 5$ , or  $W\delta = 1$  when  $N = 8$ . Note that even though the Krylov vector could reach very high accuracy, the algorithm could suffer from the MPS truncation error. The most efficient way is to simulate the time evolution using a time step size or subspace size where the error incurred by the Krylov vector is comparable to the MPS truncation error.

### B. Alternative MPS Compression Methods

The density matrix method<sup>55</sup> is a newly proposed error-controlled method. Its computational complexity is  $O(L^5M^3)$  for a conventional MPO, and it can be reduced to  $O(L^3M^3)$  when employing our THC-MPO method. Its memory cost is to store the  $L_i$  matrices, and it can be reduced to  $O(LM^2)$  by our THC-MPO method.

The zip-up method<sup>54</sup> trades computational efficiency for controllability of the approximation error. The reason is that the zip-up method uses a nonorthogonal basis.<sup>55</sup> Therefore, its cost is reduced, but the magnitude of the error is not explicitly known. Applied to the molecular Hamiltonian, its computational complexity scales as  $O(L^5M^2 + L^3M^3)$ , which can be reduced to  $O(L^3M^3)$  when employing our THC-MPO method. Its memory cost is  $O(L^2M^2)$ , reducible to  $O(M^2)$  by our THC-MPO method. Furthermore, the memory required to store the conventional MPO can grow prohibitively large due to the memory cost scaling  $O(L^4)$ , especially for a large number of spatial orbitals (e.g., larger than 100). By contrast, storing a single THC-MPO term  $G_{\mu\sigma,\nu\sigma}|\psi\rangle$  demands only  $O(L)$  memory.

The variational method<sup>42</sup> is prone to getting trapped in local minima, and hence, a proper initial guess is needed. One may use an inaccurate compression, such as the zip-up method, for this initial guess, and then perform a further variational optimization, as suggested in.<sup>34</sup> In this scenario, the dominant cost will not stem from the variational method itself.

We would like to emphasize that our THC-MPO method does not engage in a competition between the SVD method and others, but provides potential improvements, including memory cost, computational complexity, and parallelizability, for all of the MPS compression methods. As more innovative MPS compression methods will be proposed in the future, we can always check whether our THC-MPO method can be combined with them.

### C. Local Annihilation and Creation Operators

Two spinor orbitals are contained in each spatial orbital; therefore, the Pauli-Z operator is also needed inside each site when employing the Jordan-Wigner transformation. Ordering the spin-up before the spin-down, the local annihilation operators for site  $s$  are written as

$$\begin{aligned} b_{s,\uparrow} &= c \otimes I_2, \\ b_{s,\downarrow} &= Z \otimes c, \end{aligned} \quad (23)$$

where  $c$  is defined as the  $2 \times 2$  matrix  $c \equiv \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ , and  $I_2$  denotes the identity matrix of size  $2 \times 2$ . The creation operators can be readily obtained by taking the conjugate transpose of the annihilation operators.

### D. Decomposition of the Kinetic Term

The matrix  $(t_{pq})$  of one-body integrals is real symmetric and can thus be diagonalized via an orthogonal matrix  $(u_{pi})$  of eigenvectors and corresponding eigenvalues  $\lambda_i$

$$t_{pq} = \sum_{i=1}^L u_{pi}\lambda_i u_{qi} \quad \text{for all } p, q = 1, \dots, L \quad (24)$$

Inserted into the kinetic term in eq 1, we directly obtain

$$T = \sum_{i=1}^L \sum_{\sigma \in \{\uparrow, \downarrow\}} \lambda_i \underbrace{\left( \sum_{p=1}^L u_{pi} a_{p,\sigma}^\dagger \right) \left( \sum_{q=1}^L u_{qi} a_{q,\sigma} \right)}_{=T_{i,\sigma}} \quad (25)$$

For each subterm  $T_{i,\sigma}$ , one can construct elementary MPOs for  $\sum_{p=1}^L u_{pi} a_{p,\sigma}^\dagger$  and  $\sum_{q=1}^L u_{qi} a_{q,\sigma}$  in the same way as in eqs 16a and 16b. Therefore,  $T_{i,\sigma}$  is a product of two MPOs with individual bond dimensions 2. Since the kinetic term is the sum of the subterms  $T_{i,\sigma}$ , the operation  $T|\psi\rangle$  can be performed analogously to the Coulomb interaction by sequential summation and compression of the states  $T_{i,\sigma}|\psi\rangle$  in MPS form. In total, there are  $2L$  subterms for the kinetic part, which is relatively small compared to the  $O(L^2)$  subterms arising from the Coulomb interaction. Moreover, note that the spectral decomposition (24) is numerically exact, while the THC representation of the Coulomb overlap integrals in eq 2 is an approximation in general.

### E. Computational Cost Estimate for Computing Krylov Vectors

There are three primary steps to compute  $H|\psi\rangle$  in compressed MPS form: multiplying  $H$  with  $|\psi\rangle$ , renormalization, and truncation by SVD. Regarding the conventional MPO-based method, the most expensive step is renormalization, which

contains QR decompositions and the subsequent absorption of the  $R$  matrices from the QR decomposition into the next site. The QR decomposition on tensors of shape  $(2, L^2M, L^2M)$  leads to a cost of  $\sim \frac{10}{3}L^6M^3$  floating point operations if the Householder reflection method is utilized.<sup>91</sup> Such a decomposition results in an  $R$  matrix of shape  $(L^2M, L^2M)$ , absorbing it into the next site costs  $\sim 2L^6M^3$ . The leading term in computational cost is  $\sim \frac{32}{3}L^7M^3$  floating point operations for all  $2L$  sites (spin-orbitals). The SVD cost is minor: starting from the very left or right side, one of the two MPS virtual bonds has already been reduced to  $M$ , leading to a tensor of shape  $(2, M, L^2M)$ . Applying an SVD of such a tensor only costs  $O(L^2M^3)$ , which is much smaller than the cost from the QR decomposition. The asymptotic scaling  $O(L^7M^3)$  is a significant hurdle when applying global methods to large systems. Typically, the maximum MPO bond dimensions exceed  $L^2$ , so we provide only a rough estimation to offer some intuition.

For THC-MPO, we first discuss the computational complexity of evaluating a subterm  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$ , for which we execute multiplication and compression layer by layer as discussed in 3.2. After multiplying a layer with the current MPS, the shape of the resulting temporary MPS tensors is  $(2, 2M, 2M)$  since the MPO bond dimension for each layer is 2. Thus, employing a single-site QR decomposition costs  $\sim \frac{80}{3}M^3$  floating point operations. Subsequently, absorbing the  $R$  matrix of shape  $(2M, 2M)$  into the next site costs  $\sim 16M^3$ . Finally, we apply an SVD to truncate the intermediate MPS. Similarly to the conventional case, one of the two MPS virtual bonds has already been reduced to  $M$ , leading to a tensor of shape  $(2, M, 2M)$ . Applying an SVD of such a tensor costs  $\sim 56M^3$  using the divide-and-conquer method implemented in LAPACK,<sup>92,93</sup> and absorbing the obtained matrix into the next site costs  $\sim 8M^3$ . In summary, for each of the four layers in  $G_{\mu\sigma,\nu\sigma'}$  it costs  $\sim 214LM^3$  to compress the intermediate MPS for all  $2L$  sites, leading to  $\sim 10^3LM^3$  floating point operations to obtain  $G_{\mu\sigma,\nu\sigma'}|\psi\rangle$  in compressed MPS form. To implement eq 18, one needs to execute  $4N^2$  times multiplication-compression where  $N$  is the THC rank which scales linearly with system size. Taking the hydrogen chain with  $N = 3L - 3$  as an example, this leads to a total cost of  $\sim 10^4L^3M^3$  floating point operations, considering the optimization in which we reuse the later half of  $G_{\mu\sigma,\nu\sigma'}$ , as mentioned in Section 3.2. We also need to implement  $4N^2 - 1$  times addition-compression, but its cost is negligible in comparison since QR decomposition is not necessary.

Even though the complexity estimation here is approximate since we treat all bond dimensions as  $M$  for simplicity and might need larger  $M$  for desired accuracy, the asymptotic scaling gap  $O(L^4)$  is faithfully captured. Comparing the cost  $\sim 10L^7M^3$  for a conventional MPO with  $\sim 10^4L^3M^3$  for the THC-MPO method, the crossover point is estimated to occur when  $L$  is in the range of a few tens.

## AUTHOR INFORMATION

### Corresponding Authors

**Yu Wang** – Department of Computer Science, Technical University of Munich, CIT, 85748 Garching, Germany; [orcid.org/0009-0004-3972-4388](https://orcid.org/0009-0004-3972-4388); Email: [18yu.wang@tum.de](mailto:18yu.wang@tum.de)

**Christian B. Mendl** – Department of Computer Science, Technical University of Munich, CIT, 85748 Garching,

Germany; Institute for Advanced Study, Technical University of Munich, 85748 Garching, Germany; Email: [christian.mendl@tum.de](mailto:christian.mendl@tum.de)

### Authors

**Maxine Luo** – Max Planck Institute of Quantum Optics, 85748 Garching, Germany; Munich Center for Quantum Science and Technology, 80799 München, Germany

**Matthias Reumann** – Department of Computer Science, Technical University of Munich, CIT, 85748 Garching, Germany; [orcid.org/0009-0009-5589-4391](https://orcid.org/0009-0009-5589-4391)

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.5c00525>

### Notes

The authors declare no competing financial interest.

### ACKNOWLEDGMENTS

We thank Haibo Ma, Zhaoxuan Xie, Ke Liao, and Philipp Seitz for their helpful discussion. C. M. acknowledges funding by the Munich Quantum Valley initiative, supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

### ADDITIONAL NOTE

<sup>†</sup>There is also a “local version” for the Krylov method used in DMRG and time evolution simulation;<sup>34</sup> in this paper, we solely focus on the global Krylov method introduced in Section 2.2.

### REFERENCES

- (1) Östlund, S.; Rommer, S. Thermodynamic Limit of Density Matrix Renormalization. *Phys. Rev. Lett.* **1995**, *75*, 3537–3540.
- (2) Rommer, S.; Östlund, S. Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group. *Phys. Rev. B* **1997**, *55*, 2164–2181.
- (3) Ma, H.; Schollwöck, U.; Shuai, Z. *Density Matrix Renormalization Group (DMRG)-Based Approaches in Computational Chemistry*; Elsevier, 2022.
- (4) Chan, G. K.-L.; Sharma, S. The density matrix renormalization group in quantum chemistry. *Annu. Rev. Phys. Chem.* **2011**, *62*, 465–481.
- (5) Xu, Y.; Xie, Z.; Xie, X.; Schollwöck, U.; Ma, H. Stochastic adaptive single-site time-dependent variational principle. *JACS Au* **2022**, *2*, 335–340.
- (6) Ren, J.; Shuai, Z.; Kin-Lic Chan, G. Time-dependent density matrix renormalization group algorithms for nearly exact absorption and fluorescence spectra of molecular aggregates at both zero and finite temperature. *J. Chem. Theory Comput.* **2018**, *14*, 5027–5039.
- (7) Baiardi, A.; Stein, C. J.; Barone, V.; Reiher, M. Vibrational density matrix renormalization group. *J. Chem. Theory Comput.* **2017**, *13*, 3764–3777.
- (8) White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **1992**, *69*, 2863–2866.
- (9) White, S. R. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* **1993**, *48*, 10345–10356.
- (10) Xu, Y.; Cheng, Y.; Song, Y.; Ma, H. New density matrix renormalization group approaches for strongly correlated systems coupled with large environments. *J. Chem. Theory Comput.* **2023**, *19*, 4781–4795.
- (11) Chan, G. K.; 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*, 014102.
- (12) Keller, S.; Dolfi, M.; Troyer, M.; Reiher, M. An efficient matrix product operator representation of the quantum chemical Hamiltonian. *J. Chem. Phys.* **2015**, *143*, 244118.

- (13) Szalay, S.; Pfeffer, M.; Murg, V.; Barcza, G.; Verstraete, F.; Schneider, R.; Legeza, O. Tensor product methods and entanglement optimization for ab initio quantum chemistry. *Int. J. Quantum Chem.* **2015**, *115*, 1342–1391.
- (14) Friesecke, G.; Barcza, G.; Legeza, O. 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*, 87–102.
- (15) Cheng, Y.; Ma, H. Renormalized-Residue-Based Multireference Configuration Interaction Method for Strongly Correlated Systems. *J. Chem. Theory Comput.* **2024**, *20*, 1988–2009.
- (16) Xie, Z.; Song, Y.; Peng, F.; Li, J.; Cheng, Y.; Zhang, L.; Ma, Y.; Tian, Y.; Luo, Z.; Ma, H. Kynin 1.0: An ab-initio density matrix renormalization group quantum chemistry program. *J. Comput. Chem.* **2023**, *44*, 1316–1328.
- (17) Klünder, K.; Dahlström, J. M.; Gisselbrecht, M.; Fordell, T.; Swoboda, M.; Guénot, D.; Johnsson, P.; Caillat, J.; Mauritsson, J.; Maquet, A.; Taïeb, R.; L'Huillier, A. Probing single-photon ionization on the attosecond time scale. *Phys. Rev. Lett.* **2011**, *106*, 143002.
- (18) Krausz, F.; Ivanov, M. Attosecond physics. *Rev. Mod. Phys.* **2009**, *81*, 163–234.
- (19) Lépine, F.; Ivanov, M. Y.; Vrakking, M. J. Attosecond molecular dynamics: fact or fiction? *Nat. Photonics* **2014**, *8*, 195–204.
- (20) Goulielmakis, E.; Loh, Z.-H.; Wirth, A.; Santra, R.; Rohringer, N.; Yakovlev, V. S.; Zherebtsov, S.; Pfeifer, T.; Azzeer, A. M.; Kling, M. F.; Leone, S. R.; Krausz, F. Real-time observation of valence electron motion. *Nature* **2010**, *466*, 739–743.
- (21) Nisoli, M.; Decleva, P.; Calegari, F.; Palacios, A.; Martin, F. Attosecond electron dynamics in molecules. *Chem. Rev.* **2017**, *117*, 10760–10825.
- (22) Baltuška, A.; Udem, T.; Uiberacker, M.; Hentschel, M.; Goulielmakis, E.; Gohle, C.; Holzwarth, R.; Yakovlev, V. S.; Scrinzi, A.; Hänsch, T. W.; Krausz, F. Attosecond control of electronic processes by intense light fields. *Nature* **2003**, *421*, 611–615.
- (23) Borrego-Varillas, R.; Lucchini, M.; Nisoli, M. Attosecond spectroscopy for the investigation of ultrafast dynamics in atomic, molecular and solid-state physics. *Rep. Prog. Phys.* **2022**, *85*, 066401.
- (24) Midorikawa, K. Progress on table-top isolated attosecond light sources. *Nat. Photonics* **2022**, *16*, 267–278.
- (25) Haegeman, J.; Lubich, C.; Oseledets, I.; Vandereycken, B.; Verstraete, F. Unifying time evolution and optimization with matrix product states. *Phys. Rev. B* **2016**, *94*, 165116.
- (26) Baiardi, A.; Reiher, M. The density matrix renormalization group in chemistry and molecular physics: Recent developments and new challenges. *J. Chem. Phys.* **2020**, *152* (4), 040903.
- (27) Baiardi, A. Electron dynamics with the time-dependent density matrix renormalization group. *J. Chem. Theory Comput.* **2021**, *17*, 3320–3334.
- (28) Ren, J.; Li, W.; Jiang, T.; Wang, Y.; Shuai, Z. Time-dependent density matrix renormalization group method for quantum dynamics in complex systems. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2022**, *12*, No. e1614.
- (29) Xu, Y.; Liu, C.; Ma, H. Kynin-V: An open-source package calculating the dynamic and spectroscopic properties of large systems. *J. Chem. Phys.* **2024**, *161*, 052501.
- (30) Hu, W.; Chan, G. K.-L. Excited-State Geometry Optimization with the Density Matrix Renormalization Group, as Applied to Polyenes. *J. Chem. Theory Comput.* **2015**, *11* (7), 3000–3009.
- (31) Kloss, B.; Lev, Y. B.; Reichman, D. Time-dependent variational principle in matrix-product state manifolds: Pitfalls and potential. *Phys. Rev. B* **2018**, *97*, 024307.
- (32) Goto, S.; Danshita, I. Performance of the time-dependent variational principle for matrix product states in the long-time evolution of a pure state. *Phys. Rev. B* **2019**, *99*, 054307.
- (33) Yang, M.; White, S. R. Time-dependent variational principle with ancillary Krylov subspace. *Phys. Rev. B* **2020**, *102*, 094315.
- (34) Paeckel, S.; Köhler, T.; Swoboda, A.; Manmana, S. R.; Schollwöck, U.; Hubig, C. Time-evolution methods for matrix-product states. *Ann. Phys.* **2019**, *411*, 167998.
- (35) Haydock, R. The recursive solution of the Schrödinger equation. *Solid State Phys.* **1980**, *35*, 215–294.
- (36) Lanczos, C. An Iteration Method for the Solution of the Eigenvalue Problem of Linear Differential and Integral Operators. *J. Res. Natl. Bur. Stand.* **1950**, *45*, 255–282.
- (37) Dargel, P. E.; Woellert, A.; Honecker, A.; McCulloch, I.; Schollwöck, U.; Pruschke, T. Lanczos algorithm with matrix product states for dynamical correlation functions. *Phys. Rev. B* **2012**, *85*, 205119.
- (38) Wang, Y.; Yang, Z.; Mendl, C. B. Improved Lanczos Algorithm using Matrix Product States. **2025**. arXiv:2504.21786.
- (39) Dektor, A.; DelMastro, P.; Ye, E.; Van Beeumen, R.; Yang, C. Inexact subspace projection methods for low-rank tensor eigenvalue problems. **2025**. arXiv:2502.19578.
- (40) Ronca, E.; Li, Z.; Jimenez-Hoyos, C. A.; Chan, G. K.-L. Time-step targeting time-dependent and dynamical density matrix renormalization group algorithms with ab initio Hamiltonians. *J. Chem. Theory Comput.* **2017**, *13*, 5560–5571.
- (41) Frahm, L.-H.; Pfannkuche, D. Ultrafast ab initio quantum chemistry using matrix product states. *J. Chem. Theory Comput.* **2019**, *15*, 2154–2165.
- (42) Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **2011**, *326*, 96–192.
- (43) Ren, J.; Li, W.; Jiang, T.; Shuai, Z. A general automatic method for optimal construction of matrix product operators using bipartite graph theory. *J. Chem. Phys.* **2020**, *153*, 084118.
- (44) Hohenstein, E. G.; Parrish, R. M.; Martínez, T. J. Tensor hypercontraction density fitting. I. Quartic scaling second- and third-order Møller-Plesset perturbation theory. *J. Chem. Phys.* **2012**, *137*, 044103.
- (45) Parrish, R. M.; Hohenstein, E. G.; Martínez, T. J.; Sherrill, C. D. Tensor hypercontraction. II. Least-squares renormalization. *J. Chem. Phys.* **2012**, *137*, 224106.
- (46) Hohenstein, E. G.; Parrish, R. M.; Sherrill, C. D.; Martínez, T. J. Communication: Tensor hypercontraction. III. Least-squares tensor hypercontraction for the determination of correlated wavefunctions. *J. Chem. Phys.* **2012**, *137*, 221101.
- (47) Bridgeman, J. C.; Chubb, C. T. Hand-waving and interpretive dance: An introductory course on tensor networks. *J. Phys. A: Math. Theor.* **2017**, *50*, 223001.
- (48) Dargel, P. E.; Honecker, A.; Peters, R.; Noack, R. M.; Pruschke, T. Adaptive Lanczos-vector method for dynamic properties within the density matrix renormalization group. *Phys. Rev. B* **2011**, *83*, 161104.
- (49) Park, T. J.; Light, J. Unitary quantum time evolution by iterative Lanczos reduction. *J. Chem. Phys.* **1986**, *85*, 5870–5876.
- (50) Kosloff, R. Time-dependent quantum-mechanical methods for molecular dynamics. *J. Phys. Chem.* **1988**, *92*, 2087–2100.
- (51) Hochbruck, M.; Lubich, C. On Krylov subspace approximations to the matrix exponential operator. *SIAM J. Numer. Anal.* **1997**, *34*, 1911–1925.
- (52) Jiang, T.; Ren, J.; Shuai, Z. Chebyshev matrix product states with canonical orthogonalization for spectral functions of many-body systems. *J. Phys. Chem. Lett.* **2021**, *12*, 9344–9352.
- (53) Hauschild, J.; Pollmann, F. Efficient numerical simulations with tensor networks: Tensor Network Python (TeNPy). *SciPost Phys. Lect. Notes* **2018**, *5*, 5.
- (54) Stoudenmire, E. M.; White, S. R. Minimally entangled typical thermal state algorithms. *New J. Phys.* **2010**, *12*, 055026.
- (55) Ma, L.; Fishman, M.; Stoudenmire, E. M.; Solomonik, E. Approximate contraction of arbitrary tensor networks with a flexible and efficient density matrix algorithm. *Quantum* **2024**, *8*, 1580.
- (56) Lee, J.; Lin, L.; Head-Gordon, M. Systematically improvable tensor hypercontraction: Interpolative separable density-fitting for molecules applied to exact exchange, second- and third-order Møller-Plesset perturbation theory. *J. Chem. Theory Comput.* **2020**, *16*, 243–263.
- (57) Parrish, R. M.; Sherrill, C. D.; Hohenstein, E. G.; Kokkila, S. I.; Martínez, T. J. Communication: Acceleration of coupled cluster singles

- and doubles via orbital-weighted least-squares tensor hypercontraction. *J. Chem. Phys.* **2014**, *140*, 181102.
- (58) Schutski, R.; Zhao, J.; Henderson, T. M.; Scuseria, G. E. Tensor-structured coupled cluster theory. *J. Chem. Phys.* **2017**, *147*, 184113.
- (59) Lu, J.; Ying, L. Compression of the electron repulsion integral tensor in tensor hypercontraction format with cubic scaling cost. *J. Comput. Phys.* **2015**, *302*, 329–335.
- (60) Whitten, J. L. Coulombic potential energy integrals and approximations. *J. Chem. Phys.* **1973**, *58*, 4496–4501.
- (61) Dunlap, B. L.; Connolly, J.; Sabin, J. On some approximations in applications of  $X\alpha$  theory. *J. Chem. Phys.* **1979**, *71*, 3396–3402.
- (62) Werner, H.-J.; Manby, F. R.; Knowles, P. J. Fast linear scaling second-order Møller–Plesset perturbation theory (MP2) using local and density fitting approximations. *J. Chem. Phys.* **2003**, *118*, 8149–8160.
- (63) Hu, W.; Lin, L.; Yang, C. Interpolative separable density fitting decomposition for accelerating hybrid density functional calculations with applications to defects in silicon. *J. Chem. Theory Comput.* **2017**, *13*, 5420–5431.
- (64) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. 2015; Published online at <https://iclr.cc/archive/www/doku.php%3Fid=iclr2015:main.html> (accessed June 28, 2025).
- (65) <https://github.com/google-deeppmind/optax> (accessed May 2025).
- (66) Nocedal, J.; Wright, S. J. *Numerical Optimization*, 2nd ed.; Springer, 2006.
- (67) Lee, J.; Berry, D. W.; Gidney, C.; Huggins, W. J.; McClean, J. R.; Wiebe, N.; Babbush, R. Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction. *PRX Quantum* **2021**, *2*, 030305.
- (68) Matthews, D. A. Improved grid optimization and fitting in least squares tensor hypercontraction. *J. Chem. Theory Comput.* **2020**, *16*, 1382–1385.
- (69) Dong, K.; Hu, W.; Lin, L. Interpolative separable density fitting through centroidal voronoi tessellation with applications to hybrid functional electronic structure calculations. *J. Chem. Theory Comput.* **2018**, *14*, 1311–1320.
- (70) Li, Z.; Li, J.; Dattani, N. S.; Umrigar, C.; Chan, G. K. The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations. *J. Chem. Phys.* **2019**, *150*, 024302.
- (71) Luo, M.; Cirac, J. I. Efficient simulation of quantum chemistry problems in an enlarged basis set. *PRX Quantum* **2025**, *6*, 010355.
- (72) Jordan, P.; Wigner, E. Über das Paulische Äquivalenzverbot. *Z. Phys.* **1928**, *47*, 631.
- (73) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; Wouters, S.; Chan, G. K.-L. PySCF: the Python-based simulations of chemistry framework. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2018**, *8*, No. e1340.
- (74) Sun, Q.; Zhang, X.; Banerjee, S.; Bao, P.; Barbry, M.; Blunt, N. S.; Bogdanov, N. A.; Booth, G. H.; Chen, J.; Cui, Z.-H.; et al. Recent developments in the PySCF program package. *J. Chem. Phys.* **2020**, *153*, 024109.
- (75) Mendl, C. B. PyTeNet: A concise Python implementation of quantum tensor network algorithms. *J. Open Source Softw.* **2018**, *3*, 948.
- (76) <https://github.com/qc-tum/chemtensor> (accessed: May 2025).
- (77) Zhai, H.; Chan, G. K. Low communication high performance ab initio density matrix renormalization group algorithms. *J. Chem. Phys.* **2021**, *154*, 224116.
- (78) Xiang, C.; Jia, W.; Fang, W.-H.; Li, Z. Distributed multi-GPU ab initio density matrix renormalization group algorithm with applications to the P-cluster of nitrogenase. *J. Chem. Theory Comput.* **2024**, *20*, 775–786.
- (79) Olivares-Amaya, R.; Hu, W.; Nakatani, N.; Sharma, S.; Yang, J.; Chan, G. K.-L. The ab-initio density matrix renormalization group in practice. *J. Chem. Phys.* **2015**, *142*, 034102.
- (80) Chan, G. K.-L. An algorithm for large scale density matrix renormalization group calculations. *J. Chem. Phys.* **2004**, *120*, 3172–3178.
- (81) Wouters, S.; Van Neck, D. The density matrix renormalization group for ab initio quantum chemistry. *Eur. Phys. J. D* **2014**, *68*, 272.
- (82) Brabec, J.; Brandeje, J.; Kowalski, K.; Xantheas, S.; Legeza, O.; Veis, L. Massively parallel quantum chemical density matrix renormalization group method. *J. Comput. Chem.* **2021**, *42*, 534–544.
- (83) Zhai, H.; Larsson, H. R.; Lee, S.; Cui, Z.-H.; Zhu, T.; Sun, C.; Peng, L.; Peng, R.; Liao, K.; Tölle, J.; Yang, J.; Li, S.; Chan, G. K.-L. Block2: A comprehensive open source framework to develop and apply state-of-the-art DMRG algorithms in electronic structure and beyond. *J. Chem. Phys.* **2023**, *159*, 234801.
- (84) Blackford, L. S.; Demmel, J.; Dongarra, J.; Duff, I.; Hammarling, S.; Henry, G.; Heroux, M.; Kaufman, L.; Lumsdaine, A.; Petit, A.; Pozo, R.; Remington, K.; Whaley, R. C. An updated set of basic linear algebra subprograms (BLAS). *ACM Trans. Math. Software* **2002**, *28*, 135–151.
- (85) Demmel, J.; Grigori, L.; Hoemmen, M.; Langou, J. Communication-optimal parallel and sequential QR and LU factorizations. *SIAM J. Sci. Comput.* **2012**, *34*, A206–A239.
- (86) Jessup, E. R.; Sorensen, D. C. A parallel algorithm for computing the singular value decomposition of a matrix. *SIAM J. Matrix Anal. Appl.* **1994**, *15*, 530–548.
- (87) Dagum, L.; Menon, R. OpenMP: an industry standard API for shared-memory programming. *IEEE Comput. Sci. Eng.* **1998**, *5*, 46–55.
- (88) Yang, Y.; Iblisdir, S.; Cirac, J. I.; Banuls, M. C. Probing thermalization through spectral analysis with matrix product operators. *Phys. Rev. Lett.* **2020**, *124*, 100602.
- (89) Holzner, A.; Weichselbaum, A.; McCulloch, I. P.; Schollwöck, U.; von Delft, J. Chebyshev matrix product state approach for spectral functions. *Phys. Rev. B* **2011**, *83*, 195115.
- (90) Hochbruck, M.; Lubich, C. On Krylov Subspace Approximations to the Matrix Exponential Operator. *SIAM J. Numer. Anal.* **1997**, *34*, 1911–1925.
- (91) Golub, G. H.; Van Loan, C. F. *Matrix Computations*; Johns Hopkins University Press, 2013.
- (92) Gu, M.; Eisenstat, S. C. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.* **1996**, *17*, 848–869.
- (93) Anderson, E.; Bai, Z.; Bischof, C.; Blackford, L. S.; Demmel, J.; Dongarra, J.; Du Croz, J.; Greenbaum, A.; Hammarling, S.; McKenney, A.; Sorensen, D. *LAPACK Users' Guide*; Society for Industrial and Applied Mathematics, 1999.