







RESEARCH ARTICLE | JUNE 15 2026

Efficient simulation of non-Markovian path integrals via imaginary time evolution of an effective Hamiltonian

Xiaoyu Yang ; Limin Liu ; Wencheng Zhao ; Jiajun Ren  ; Wei-Hai Fang 



J. Chem. Phys. 164, 234106 (2026)

<https://doi.org/10.1063/5.0328369>



AIP Advances

Why Publish With Us?

-  **21DAYS**
average time
to 1st decision
-  **OVER 4 MILLION**
views in the last year
-  **INCLUSIVE**
scope

[Learn More](#)



Efficient simulation of non-Markovian path integrals via imaginary time evolution of an effective Hamiltonian

Cite as: J. Chem. Phys. 164, 234106 (2026); doi: 10.1063/5.0328369

Submitted: 14 February 2026 • Accepted: 1 June 2026 •

Published Online: 15 June 2026



View Online



Export Citation



CrossMark

Xiaoyu Yang,¹  Limin Liu,¹  Wencheng Zhao,¹  Jiajun Ren,^{1,2,a)}  and Wei-Hai Fang^{1,2} 

AFFILIATIONS

¹Key Laboratory of Theoretical and Computational Photochemistry, Ministry of Education, College of Chemistry, Beijing Normal University, 100875 Beijing, People's Republic of China

²Institute for Advanced Study, Beijing Normal University, 100875 Beijing, People's Republic of China

Note: This paper is part of the Special Topic, Simulation of Excited State Dynamics in Molecular and Condensed-Matter Systems.

^{a)}Author to whom correspondence should be addressed: jjren@bnu.edu.cn

ABSTRACT

Accurately simulating the non-Markovian dynamics of open quantum systems remains a significant challenge. While the recently proposed time-evolving matrix product operator (TEMPO) algorithm based on path integrals successfully circumvents the exponential scaling associated with memory length, its reliance on layer-by-layer tensor contractions and compressions leads to steep scaling with respect to the system Hilbert space dimension. In this work, we introduce the effective Hamiltonian-based TEMPO (EH-TEMPO) algorithm, which reformulates the calculation of the Feynman–Vernon influence functional as an imaginary time evolution governed by an effective Hamiltonian. We demonstrate that this effective Hamiltonian admits a highly compact matrix product operator representation, enabling substantial compression with negligible loss of accuracy. Combining a one-shot global evolution with a backward retrieval approach, EH-TEMPO significantly reduces algorithmic complexity and is naturally suited for GPU acceleration. We benchmark the method by simulating the energy transfer dynamics in the 7-site Fenna–Matthews–Olson complex model and 4-site perylene bisimide model. The results demonstrate that EH-TEMPO achieves numerically exact accuracy with superior efficiency, delivering speedups of up to 17.5× on GPU hardware compared to standard CPU implementations.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0328369>

I. INTRODUCTION

The study of quantum dynamics in condensed-phase systems^{1,2} is central to understanding a wide range of physical and chemical phenomena, including carrier transport in organic semiconductors and excitation energy transfer in photosynthetic complexes. In these systems, the phonon environment affects the electron dynamics through electron–phonon interactions, inducing decoherence and energy dissipation that ultimately govern the overall functional properties.^{3–5}

Various approaches have been developed to simulate such electron–phonon coupled dynamics.⁶ The first category treats the composite system, which comprises both electronic and phononic degrees of freedom, as a closed quantum system by explicitly propagating the full wavefunction according to the Schrödinger

equation. Prominent examples include the time-dependent density matrix renormalization group (TD-DMRG),^{7,8} tree tensor network states (TTNS),^{9,10} and the multilayer multi-configuration time-dependent Hartree (ML-MCTDH) approach.^{11,12} While these methods are highly accurate and robust across broad parameter regimes, their computational cost increases rapidly with the number of bath modes, often rendering long-time simulations of realistic condensed-phase problems prohibitively expensive. The second category focuses on the reduced dynamics of the system of interest by tracing out the phononic degrees of freedom. This includes perturbative approaches based on quantum master equations, such as the Lindblad^{13,14} and Redfield equations,^{15,16} which commonly invoke the Markov approximation. However, capturing non-Markovian memory effects, which are essential in many chemical systems, requires numerically exact methods. The most notable

among these are the hierarchical equations of motion (HEOM),^{17,18} non-Markovian quantum state diffusion (NMQSD),^{19,20} and the quasi-adiabatic propagator path integral (QuAPI).^{21,22}

QuAPI, originally developed by Makri and coworkers, can, in principle, fully retain non-Markovian effects and handle arbitrary bath spectral densities at both zero and finite temperatures.^{21–24} While the QuAPI algorithm scales exponentially with memory time, various techniques have been developed to significantly improve its efficiency in specific parameter regimes, such as path filtering,²⁵ blip decomposition,²⁶ and coarse graining.²⁷ Notably, Makri and coworkers introduced the small matrix path integral (SMatPI) algorithm,²⁸ which offers an efficient strategy to overcome the exponential memory storage scaling of QuAPI. By recursively disentangling the influence functional, SMatPI reduces the necessary memory storage to a small set of fixed-size matrices whose dimensions are equal to the square of system size. Another significant breakthrough was achieved with the introduction of tensor network (TN) algorithms to further enhance path integral methods.^{29–31} Among them, the time-evolving matrix product operator (TEMPO) algorithm proposed by Strathearn *et al.*²⁹ combines the influence functional formalism with TN techniques, specifically matrix product states (MPS) and matrix product operators (MPO).³² It acts as a powerful compression scheme for the exponentially large path space. By expressing the influence functional as the contraction of a two-dimensional TN and utilizing singular value decomposition (SVD) algorithm, TEMPO can systematically compress the influence functional. In broad parameter regimes, this effectively reduces the scaling with respect to memory length from exponential to polynomial.²⁹ Intrinsically, this efficiency is tied to the compressibility of the influence functional; for problems characterized by strong system–bath coupling and long memory, capturing the strong temporal correlations may still require a rapidly growing TN bond dimension. Subsequent developments have significantly extended the generality and capability of TEMPO, including reformulations within the process tensor (PT) framework for multi-time correlation functions,³³ extensions to general off-diagonal or non-commuting system–bath couplings,^{34–37} generalizations to equilibrium initial states,³⁸ and fermionic baths,^{39–41} and combinations with other powerful algorithms, such as SMatPI^{28,42} and high-dimensional tensor networks.⁴³

Despite these successes, the standard TEMPO algorithm still faces significant challenges when applied to multi-state systems. A key step in TEMPO is the iterative contraction and compression of the layered MPOs representing the incremental influence functional.²⁹ In general, the bond dimension of MPOs grows quadratically with the system's Hilbert space dimension d , making standard SVD-based compression prohibitively expensive for multi-state systems, with computational costs scaling as $\mathcal{O}(d^3)$ for tensor contraction and $\mathcal{O}(d^4)$ for tensor decomposition. In addition, while the analytical construction of a compact MPO for a single system–bath coupling operator is well-documented in the literature,²⁹ manually constructing the MPO for general system–bath interactions remains non-trivial and laborious. The existing automated MPO construction algorithms^{44–46} are incompatible with the influence functional form, which lacks a simple sum-of-products (SOP) structure. In this work, we propose an effective Hamiltonian-based TEMPO (EH-TEMPO) algorithm to address the scalability of (PT-)TEMPO for multi-state systems. Instead of the standard layer-by-layer

contraction and compression scheme, we derive the influence functional via the imaginary time evolution of an effective Hamiltonian. This reformulation offers three distinct advantages over the original (PT-)TEMPO algorithm: (i) The effective Hamiltonian takes an SOP form, enabling the use of automated MPO construction algorithms to generate optimal MPO representations for general system–bath couplings with minimal bond dimensions. (ii) Well-established time evolution algorithms, such as those based on the time-dependent variational principle (TDVP),^{47,48} can be employed to obtain the overall influence functional in a one-shot calculation with relatively large imaginary time steps, replacing the costly layer-by-layer heavy calculations. (iii) With TDVP-based time evolution algorithms, the EH-TEMPO algorithm avoids decompositions of large tensors, making it highly amenable to GPU acceleration. We benchmark the EH-TEMPO algorithm by simulating excitation energy transfer dynamics in the 7-site Fenna–Matthews–Olson (FMO) model and 4-site perylene bisimide (PBI) model. By comparing our results with those from the PT-TEMPO algorithm, we demonstrate that EH-TEMPO provides accurate influence functionals with superior computational efficiency for multi-state systems.

The remainder of this paper is organized as follows. Section II briefly reviews the theoretical foundations of the QuAPI and TEMPO algorithms and then details the proposed EH-TEMPO algorithm. In Sec. III, we systematically analyze the accuracy and efficiency of the EH-TEMPO method, comparing its performance against standard benchmarks. Finally, conclusions are presented in Sec. IV.

II. METHODS

A. QuAPI and TEMPO algorithms

A generic system–bath coupled problem is described by the total Hamiltonian,

$$\hat{H} = \hat{H}_S + \hat{H}_B + \hat{H}_{SB}, \quad (1)$$

$$\hat{H}_B = \sum_i \frac{1}{2} \hat{p}_i^2 + \frac{1}{2} \omega_i^2 \hat{q}_i^2, \quad (2)$$

$$\hat{H}_{SB} = \sum_n \hat{S}_n \otimes \sum_i c_{ni} \hat{q}_i, \quad (3)$$

where \hat{H}_S is the system Hamiltonian, \hat{H}_B is the bath Hamiltonian, and \hat{H}_{SB} describes the system–bath interaction. In this work, we restrict our consideration to diagonal system–bath coupling, for which there exists a basis $\{|s\rangle\}$ that simultaneously diagonalizes all operators \hat{S}_n , with eigenvalues s . For notational conciseness, we consider a single coupling term in Eq. (3) in the following and, thus, omit the summation over n . The time evolution of the total density matrix is given by $\hat{\rho}(t) = e^{-i\hat{H}t} \hat{\rho}(0) e^{i\hat{H}t}$. We assume a factorized initial condition at $t = 0$, $\hat{\rho}(0) = \hat{\rho}_S(0) \otimes \hat{\rho}_B^{\text{eq}}$, where the bath is in thermal equilibrium. The extension to correlated thermal equilibrium initial states has been discussed in Refs. 38, 49, and 50. Applying the standard Trotter splitting of the short-time propagator $e^{-i\hat{H}\Delta t}$ to separate system and bath contributions, the reduced density matrix $\hat{\rho}_S(t) = \text{Tr}_B[\hat{\rho}(t)]$ at time $t = N_t \Delta t$ can be expressed as a discretized path integral,

$$\hat{\rho}_S(N_t \Delta t) = \sum_{s_0^\pm, \dots, s_{N_t-1}^\pm} \langle s_0^+ | \hat{\rho}_S(0) | s_0^- \rangle G(s_0^\pm, \dots, s_{N_t}^\pm) F(s_1^\pm, \dots, s_{N_t}^\pm). \quad (4)$$

Here, $s_0^\pm, s_1^\pm, \dots, s_{N_t}^\pm$ denote the discretized forward (+) and backward (-) system paths. The system propagator G is given by the product of short-time propagators,

$$G(s_0^\pm, \dots, s_{N_t}^\pm) = \prod_{k=1}^{N_t} \langle s_k^+ | e^{-i\hat{H}_0 \Delta t} | s_{k-1}^+ \rangle \langle s_{k-1}^- | e^{i\hat{H}_0 \Delta t} | s_k^- \rangle, \quad (5)$$

where \hat{H}_0 is the QuAPI-type shifted system Hamiltonian along the adiabatic path.²³ The environmental effects are encoded in the Feynman–Vernon influence functional F , which accounts for non-Markovian memory effects.⁵¹ It should be noted that we do not employ any memory truncation in the expressions and calculations in Sec. III; the full non-Markovian history is retained for the entire propagation time. For a Gaussian bath with linear coupling, F takes the analytical form,²¹

$$F(s_1^\pm, \dots, s_{N_t}^\pm) = \prod_{k=1}^{N_t} \prod_{k'=1}^k I_{\Delta k}(s_k^\pm, s_{k'}^\pm), \quad \Delta k = k - k', \quad (6)$$

$$I_{\Delta k}(s_k^\pm, s_{k'}^\pm) = \exp[-(s_k^+ - s_{k'}^+) (\eta_{kk'} s_{k'}^+ - \eta_{kk'}^* s_{k'}^-)].$$

The coefficients $\eta_{kk'}$ are determined by the bath spectral density and the temperature. For simplicity, we present expressions based on first-order Trotter splitting ($e^{-i\hat{H} \Delta t} \approx e^{-i\hat{H}_0 \Delta t} e^{-i(\hat{H} - \hat{H}_0) \Delta t}$ with local error $\mathcal{O}(\Delta t^2)$). The extension to second-order splitting is straightforward ($e^{-i\hat{H} \Delta t} \approx e^{-i\hat{H}_0 \Delta t/2} e^{-i(\hat{H} - \hat{H}_0) \Delta t} e^{-i\hat{H}_0 \Delta t/2}$ with local error $\mathcal{O}(\Delta t^3)$).^{21,22} Computing Eq. (4) exactly incurs an exponential cost with respect to the memory length. To overcome this challenge, the TEMPO algorithm identifies the discrete path variables s_k^\pm as the physical indices of a tensor network. The incremental influence functional in Eq. (7) can be compactly represented as a one-dimensional diagonal MPO (after inserting Kronecker deltas with respect to s_k^\pm),

$$f_k(s_1^\pm, \dots, s_k^\pm) = \prod_{k'=1}^k I_{\Delta k}(s_k^\pm, s_{k'}^\pm), \quad (7)$$




FIG. 1. Tensor network diagrams for the TEMPO algorithms. (a) Triangular tensor network representation of the influence functional used in the (PT)-TEMPO algorithm. (b) Core algorithm of EH-TEMPO, where the influence functional is generated via the imaginary time evolution governed by the effective Hamiltonian MPO (\hat{H}_{eff}). (c) Final tensor contraction to obtain the reduced density matrix, combining the influence functional MPS with the system propagators U_S and the initial state $\rho_S(0)$. The white triangle represents δ_{jk} .

$$\hat{B}_k = \left(\prod_{k'=1}^{k-1} \delta_{i_{k'}^+}^{s_{k'}^+} \delta_{i_{k'}^-}^{s_{k'}^-} \right) f_k = \sum_{\{a\}} B[1]_{a_1}^{s_1^\pm} \dots B[k-1]_{a_{k-2} a_{k-1}}^{s_{k-1}^\pm} B[k]_{a_{k-1}}^{s_k^\pm}. \quad (8)$$

Here, a_i is the index linking adjacent matrices, the size of which is called the bond dimension. The overall influence functional, thus, exhibits a two-dimensional triangular structure, which is contracted as a sequence of layered MPOs to yield a final MPS [see Fig. 1(a)],

$$F(s_1^\pm, \dots, s_{N_t}^\pm) = \hat{B}_{N_t} \hat{B}_{N_t-1} \dots \hat{B}_1, \quad (9)$$

$$\approx \sum_{\{a\}} A_{a_1}^{s_1^\pm} A_{a_1 a_2}^{s_2^\pm} \dots A_{a_{N_t-1}}^{s_{N_t}^\pm}. \quad (10)$$

In practice, after successively contracting adjacent layers, standard SVD-based compression is applied to truncate the bond dimension of MPS by discarding renormalized states with small singular values, thereby circumventing the exponential scaling of computational cost and storage with respect to N_t .

In the original TEMPO algorithm,²⁹ starting from the system initial state $\rho_S(0)$, the one-step system propagator is absorbed into the pairwise influence functional with $\Delta k = 1$ during the contraction. Consequently, instead of obtaining the overall influence functional, the so-called augmented density tensor (ADT), which is the summand in Eq. (4), is obtained directly. The subsequent combination with the PT framework³³ enables the system propagator to be separated from the influence functional construction, resulting in higher accuracy for a given bond dimension in our experience. Moreover, PT-TEMPO allows reuse of the process tensor for different system Hamiltonians and operations and is especially well suited for computing multi-time correlation functions. The trade-off is that the index of each local matrix must in general be doubled, $s_k \rightarrow (s_k, s'_k)$, which inevitably increases the computational cost. Fortunately, this doubling can be avoided for the diagonal system–bath interaction considered in this work, making the computational cost of PT-TEMPO comparable to that of TEMPO. It is important to emphasize that the process tensor encodes the exact same physical information as the foundational Feynman–Vernon influence functional. Mathematically, the process tensor is isomorphic to the time-discretized evaluation of the influence functional, reformulated as a discrete,

multilinear operator to facilitate efficient tensor network operations. Once the overall influence functional is obtained in MPS form, the system reduced density matrix is readily computed by contracting with the system propagators and initial state [see Fig. 1(c)].

We now analyze the computational scaling of TEMPO and PT-TEMPO for diagonal system–bath couplings. Building upon the original algorithm proposed in Ref. 29, our implementation separates the forward and backward index variables into $s_1^+, s_1^-, s_2^+, s_2^-, \dots, s_{N_t}^+, s_{N_t}^-$ rather than $s_1^\pm, s_2^\pm, \dots, s_{N_t}^\pm$. Although this doubles the total number of sites, it significantly reduces the computational cost per site. In the equations and figures presented in this paper, however, we group the forward and backward indices into a single site for clarity. The computational complexity of tensor contraction is $\mathcal{O}(N_t^2 M_S^2 d^3)$, where M_O is the bond dimension of the MPO in Eq. (8), M_S is the bond dimension of the MPS representing the overall influence functional (PT-TEMPO) or ADT (TEMPO), d is the system Hilbert space dimension, and N_t is the number of time steps. After contraction, the scaling of tensor decomposition for canonicalization and compression is $\mathcal{O}(N_t^2 M_S^3 d^4)$.⁴⁸ The detailed analysis of computational complexity is presented in the Appendix. The standard TEMPO and PT-TEMPO algorithms, thus, exhibit steep scaling with respect to the system dimension d , especially in tensor decomposition, which prohibits the efficient treatment of multi-state dynamics. This limitation motivates the development of the EH-TEMPO algorithm presented in Sec. II B.

B. The EH-TEMPO algorithm

The central idea of the EH-TEMPO algorithm is the reformulation of the influence functional as an imaginary time evolution (ITE) governed by an effective Hamiltonian \hat{H}_{eff} , which acts on the space of path indices $|s_1^\pm\rangle \otimes |s_2^\pm\rangle \otimes \dots \otimes |s_{N_t}^\pm\rangle$,

$$\hat{H}_{\text{eff}}^{N_t} = \sum_{k=1}^{N_t} \sum_{k'=1}^k (\hat{s}_k^+ - \hat{s}_k^-) (\eta_{kk'} \hat{s}_{k'}^+ - \eta_{kk'}^* \hat{s}_{k'}^-). \quad (11)$$

Here, the operators \hat{s}_k^\pm are diagonal in the path basis, satisfying $\hat{s}_k^\pm |s_k^\pm\rangle = s_k^\pm |s_k^\pm\rangle$. This effective Hamiltonian describes a one-dimensional lattice of length N_t along the temporal axis with long-range, complex-valued pairwise interactions $\eta_{kk'}$ and $\eta_{kk'}^*$. For example, for a two-level system where s_k^\pm takes values ± 1 , \hat{H}_{eff} is analogous to an Ising chain with long-range couplings. As recognized in early literature, this connection arises because the path integral representation of the spin–boson partition function is exactly equivalent to a classical one-dimensional Ising model with long-range couplings.⁵² Since \hat{H}_{eff} takes the form of a sum of products, its MPO representation can be efficiently constructed using the automated bipartite graph algorithm developed in our previous works.^{10,45} The automated bipartite graph algorithm can convert a general Hamiltonian, expressed as a sum of local operator products, into a highly compact MPO without requiring manual derivation. By mapping the operator strings onto a bipartite graph structure, where nodes represent local operators and edges represent interactions, the algorithm symbolically identifies and merges shared complementary operator prefixes and suffixes across all terms. This graph-based optimization effectively eliminates redundant local operators and maximally reduces the internal MPO bond dimensions (but still exact). Consequently, it provides a highly efficient programmatic

route to construct the MPO representation for complex effective Hamiltonians, directly ensuring optimal memory and computational performance during subsequent tensor network calculations. To compute the state $|\Psi_{\text{IF}}\rangle = \sum_{s_1^\pm, \dots, s_{N_t}^\pm} F(s_1^\pm, \dots, s_{N_t}^\pm) |s_1^\pm, \dots, s_{N_t}^\pm\rangle$, which encodes the overall influence functional from t_1 to t_{N_t} in its amplitudes, we start from an unnormalized initial state $|\Psi_0\rangle$ representing an equal superposition of all path configurations. This state admits an MPS representation with $M_S = 1$ (all elements in each local site are 1, $A_{a_{i-1}a_i}^{s_i^\pm} = 1$),

$$|\Psi_0\rangle = \sum_{s_1^\pm, \dots, s_{N_t}^\pm} |s_1^\pm, \dots, s_{N_t}^\pm\rangle. \quad (12)$$

Then, $|\Psi_{\text{IF}}\rangle$ can be calculated by solving the imaginary time Schrödinger equation from $\tau = 0$ to $\tau = 1$ under the effective Hamiltonian [see Fig. 1(b)],

$$\frac{\partial}{\partial \tau} |\Psi(\tau)\rangle = -\hat{H}_{\text{eff}} |\Psi(\tau)\rangle, \quad |\Psi(0)\rangle = |\Psi_0\rangle, \quad (13)$$

$$|\Psi_{\text{IF}}\rangle = |\Psi(\tau)\rangle = e^{-\tau \hat{H}_{\text{eff}}} |\Psi_0\rangle, \quad \tau = 1. \quad (14)$$

In this work, we employ the standard TDVP-based projector splitting (PS) algorithm to perform the ITE.^{47,48} The TDVP-PS algorithm is an efficient algorithm for integrating the time-dependent Schrödinger equation strictly within the MPS manifold, under the Dirac–Frenkel TDVP $\langle \delta \Psi | i\partial/\partial t - \hat{H} | \Psi \rangle = 0$. Rather than applying the evolution operator globally and subsequently truncating the bond dimension, TDVP directly projects the time evolution operator onto the tangent space of the current MPS. The core mechanism of the splitting approach lies in recognizing that the global tangent space projector can be mathematically decomposed into a sequence of local projectors acting on individual site tensors, alternating with negative projectors acting on the intervening bond matrices. By applying a Trotter decomposition to this operator, the global time evolution is transformed into a robust, DMRG-like sweeping algorithm. During a sweep, each local site tensor (the canonical center) is evolved forward in time, followed by a backward-in-time evolution of the adjacent bond matrix, which correctly shifts the canonical center to the next site while preventing the double-counting of local dynamics. In our calculations, we employ the one-site TDVP-PS algorithm, which ensures that the bond dimension remains fixed throughout the evolution. This practically prevents the exponential growth in computational cost and memory storage associated with the number of time steps. Consequently, systematic convergence checks with respect to the chosen maximum bond dimension must be performed in practical applications.

In (PT-)TEMPO, if one regards each layer of MPO contraction as one time step, obtaining $|\Psi_{\text{IF}}^{N_t}\rangle$ requires N_t steps. By contrast, the number of ITE steps N_τ required in EH-TEMPO can be very small, typically from several to a dozen steps suffice, with $N_\tau \ll N_t$, as shown in Sec. III. This substantial reduction in the number of evolution steps translates directly into computational savings. Moreover, sophisticated adaptive time step algorithms can be employed to further automate the ITE process, which is important for investigating unknown problems. Once $|\Psi_{\text{IF}}^{N_t}\rangle$ is obtained, the reduced density matrix at the final time step t_{N_t} is computed as in PT-TEMPO [see Fig. 1(c)]. However, a direct application of EH-TEMPO provides

only the reduced density matrix at the final time step t_{N_t} . To obtain the full dynamical trajectory from t_1 to t_{N_t} , three schemes can be employed, as follows:

1. **Independent scheme.** Perform ITE independently for $\hat{H}_{\text{eff}}^1, \hat{H}_{\text{eff}}^2, \dots, \hat{H}_{\text{eff}}^{N_t}$ to obtain $|\Psi_{\text{IF}}^1\rangle, |\Psi_{\text{IF}}^2\rangle, \dots, |\Psi_{\text{IF}}^{N_t}\rangle$ separately. Assuming that N_τ steps are required for each evolution to $\tau = 1$, this scheme requires $N_t \cdot N_\tau$ ITE steps in total.
2. **Stepwise scheme.** Decompose $\hat{H}_{\text{eff}}^{N_t}$ into a sum of incremental contributions \hat{h}_{eff}^k defined in Eq. (15) and perform ITE sequentially according to $\hat{h}_{\text{eff}}^1, \hat{h}_{\text{eff}}^2, \dots, \hat{h}_{\text{eff}}^{N_t}$, each from $\tau = 0$ to $\tau = 1$,

$$\hat{h}_{\text{eff}}^k = \sum_{k'=1}^k (\hat{s}_k^+ - \hat{s}_k^-) (\eta_{kk'} \hat{s}_{k'}^+ - \eta_{kk'}^* \hat{s}_{k'}^-). \quad (15)$$

This approach closely parallels (PT-)TEMPO, with the key difference that the incremental influence functional is exactly and explicitly constructed in (PT-)TEMPO, whereas it is implicitly obtained through ITE in EH-TEMPO. Like the independent scheme, this requires $N_t \cdot N_\tau$ ITE steps in total.

3. **Backward retrieval scheme.** All intermediate states $|\Psi_{\text{IF}}^{m < N_t}\rangle$ are successively retrieved from $|\Psi_{\text{IF}}^{N_t}\rangle$ by exploiting a property of the influence functional: when $s_N^+ = s_N^-$, the incremental influence functional f_N equals unity. This allows the influence functional to be retrieved backward,

$$|\Psi_{\text{IF}}^{N_t}\rangle \xrightarrow{s_{N_t}^+ = s_{N_t}^-} |\Psi_{\text{IF}}^{N_t-1}\rangle \xrightarrow{s_{N_t-1}^+ = s_{N_t-1}^-} \dots \xrightarrow{s_2^+ = s_2^-} |\Psi_{\text{IF}}^1\rangle. \quad (16)$$

Combined with this backward retrieval, only a one-shot evolution is needed, requiring just N_τ ITE steps, which is significantly fewer than the preceding two schemes. The retrieval itself incurs negligible additional cost, involving only a single MPS sweep. However, since only $|\Psi_{\text{IF}}^{N_t}\rangle$ is targeted in the one-shot calculation, the retrieval may introduce small errors.

The computational cost of EH-TEMPO is dominated by the ITE. Among the three schemes, the backward retrieval scheme is the most efficient and will be adopted in the numerical calculations presented in Sec. III. Using the TDVP projector splitting method, the tensor contraction scales as $\mathcal{O}(lN_tN_\tau(M_S^2M_O^2d + M_S^3M_Od))$, and the tensor decomposition scales as $\mathcal{O}(N_tN_\tau M_S^3d)$. The detailed analysis of computational complexity is presented in the Appendix. l is the number of tensor contraction demanded by the initial value problem (IVP) solver in updating each local site (typically $l \sim 10$). Compared to (PT-)TEMPO, the computational scaling of tensor decomposition in EH-TEMPO is significantly reduced with respect to both the system size d and the number of time steps N_t . For tensor contraction, the relative efficiency depends on M_O , the bond dimension of the effective Hamiltonian, which is determined by the system–bath coupling operator. For a single bath with one system–bath coupling term, such as in the Caldeira–Leggett model, the exact MPO has $M_O = \mathcal{O}(N_t)$, independent of d . For the Frenkel–Holstein model, where each local two-level system couples to an independent bath, the exact MPO scales as $M_O = \mathcal{O}(N_t d)$. Furthermore, as demonstrated in Sec. III and Appendix, we found that the effective Hamiltonian MPO can be highly compressed with minimal error, regardless of whether the bath possesses a

structureless or structured spectral density. After compression, M_O is substantially smaller than its exact (uncompressed) value. Most importantly, this compressed bond dimension becomes independent of N_t (see Appendix). This intrinsic compressibility, which has also been observed in Ref. 53, yields significant computational savings. In addition, since EH-TEMPO avoids the decomposition of large tensors and the dominant operations are large tensor contractions, it is highly amenable to GPU parallelization, as demonstrated in Ref. 48. This leads to the significant speedups reported in Sec. III.

Before closing this section, we briefly contrast EH-TEMPO with recent, closely related methods based on the TEMPO framework.^{37,53} In Ref. 37, Shi and coworkers derived a partial differential equation to compute the incremental influence functional. While technically similar to the stepwise scheme introduced above, our approach differs by adopting an effective Hamiltonian interpretation to calculate the overall influence functional in a one-shot calculation. This effectively reduces the computational scaling from $\mathcal{O}(N_t^2)$ to $\mathcal{O}(N_\tau N_t)$, where $N_\tau \ll N_t$. In Ref. 53, Guo and coworkers employ Prony fitting to approximate $\eta_{\Delta k}$ and construct the MPO for the exponent of fermionic influence functional by exploiting the time-translational invariance of $\eta_{\Delta k}$. The subsequent time evolution is then obtained using the $W^{1/|I|}$ and exponential thermal tensor renormalization group (XTRG) algorithms, which is also a one-shot calculation. Besides the differences in the time-evolution algorithms, EH-TEMPO utilizes a different MPO construction strategy: it maps the effective Hamiltonian to an exact MPO and subsequently compresses it using SVD without assuming a specific function form, which is inherently more numerically robust and stable than Prony fitting.

III. RESULTS AND DISCUSSION

To evaluate the performance of the EH-TEMPO algorithm, we apply it to the widely used testbed, the excitation energy transfer dynamics in molecular aggregates described by the Frenkel–Holstein model. The Hamiltonian is as follows:

$$\hat{H}_S = \sum_i \varepsilon_i \hat{a}_i^\dagger \hat{a}_i + \sum_{i \neq j} J_{ij} \hat{a}_i^\dagger \hat{a}_j, \quad (17)$$

$$\hat{H}_B + \hat{H}_{SB} = \sum_{i,\xi} \omega_\xi \hat{b}_{i\xi}^\dagger \hat{b}_{i\xi} + \sum_{i,\xi} g_\xi \omega_\xi \hat{a}_i^\dagger \hat{a}_i (\hat{b}_{i\xi}^\dagger + \hat{b}_{i\xi}), \quad (18)$$

where ε_i is the site energy of molecule i and J_{ij} is the excitonic coupling between molecules i and j . Each molecule has an independent phonon bath with frequency ω_ξ and dimensionless electron–phonon coupling constant g_ξ .

For the Frenkel–Holstein model, the influence functional generalizes Eq. (6) by taking the product over all site indices:

$$F(s_1^\pm, \dots, s_{N_t}^\pm) = \prod_{i=1}^d \exp \left\{ - \sum_{k=1}^{N_t} \sum_{k'=1}^k (n_k^{i+} - n_k^{i-}) (\eta_{kk'} n_{k'}^{i+} - \eta_{kk'}^* n_{k'}^{i-}) \right\}, \quad (19)$$

where $|s_k^\pm\rangle$ takes the local electronic state $|i\rangle$ with $i = 1, \dots, d$ (d is the number of molecules) and $n_k^{i\pm}$ is the occupation number of site i at step k . The corresponding effective Hamiltonian is

$$\hat{H}_{\text{eff}} = \sum_{i=1}^d \sum_{k=1}^{N_i} \sum_{k'=1}^k (\hat{n}_k^{i+} - \hat{n}_k^{i-}) (\eta_{kk'} \hat{n}_{k'}^{i+} - \eta_{kk'}^* \hat{n}_{k'}^{i-}), \quad (20)$$

where $\hat{n}_k^{i\pm}$ is the occupation number operator satisfying $\hat{n}_k^{i\pm} |s_k^{\pm}\rangle = \delta_{is} |s_k^{\pm}\rangle$.

We benchmarked our approach using two model systems: the 7-site Fenna–Matthews–Olson complex with a Debye spectral density, and a 4-site perylene bisimide aggregate coupled to discrete bath modes. All calculations were performed using the Renormalizer package.

A. FMO model

For the FMO model, the parameters for \hat{H}_S are from Refs. 54 and 55. The bath is characterized by a Debye spectral density, $J(\omega) = 2\lambda \frac{\omega\omega_c}{\omega^2 + \omega_c^2}$, where λ is the reorganization energy and ω_c is the cutoff frequency. We benchmarked EH-TEMPO against the numerically exact HEOM method and the PT-TEMPO algorithm. For the Debye spectral density, the HEOM method serves as an ideal reference. The accuracy of the population dynamics is quantified by the time-averaged cumulative deviation,




FIG. 2. Exciton population dynamics of the 7-site FMO complex at $T = 77$ K. The solid colored lines represent the results obtained via the EH-TEMPO algorithm (bond dimension $M_S = 128$), while the dashed black lines denote the numerically exact HEOM reference. The two datasets are in excellent agreement, with the EH-TEMPO results accurately reproducing the detailed coherent oscillations.




FIG. 4. Compressibility of the effective Hamiltonian MPO. (a) The maximum bond dimension of \hat{H}_{eff} as a function of the SVD compression threshold ϵ . The bond dimension decreases drastically from the exact value (1752) to compact sizes (<100) with non-zero thresholds. (b) The time-dependent calculation error introduced by compression, relative to the uncompressed case. The error remains negligible (order of 10^{-3}) even for aggressive compression ($\epsilon = 10^{-4}$).




FIG. 3. Time evolution of the cumulative error $\mathcal{E}(t)$ for the 7-site FMO complex. The performance of EH-TEMPO is compared against PT-TEMPO across three bond dimensions ($M_S = 30, 60, 100$), represented by different colors. The line styles distinguish the algorithmic schemes: solid lines for EH-TEMPO with the backward retrieval scheme, dash-dotted lines for the independent scheme, and dashed lines for PT-TEMPO.

$$\mathcal{E}(\Delta t N_t) = \frac{1}{7N_t} \sum_{i=1}^7 \sum_{k=1}^{N_i} \left| P^{(i)}(k) - P_{\text{ref}}^{(i)}(k) \right|, \quad (21)$$

where $P^{(i)}$ denotes the population of site i . We first verify the correctness of the EH-TEMPO algorithm. Figure 2 displays the population dynamics at $T = 77$ K with $\lambda = 35$ cm^{-1} and $\omega_c = 1/50$ fs^{-1} . The EH-TEMPO results employ a real-time step size $\Delta t = 4$ fs, an adaptive imaginary time step (with relative tolerance 10^{-5}), and a bond dimension $M_S = 128$. The results show excellent agreement with the HEOM reference, capturing all oscillatory features of the coherent energy transfer dynamics.

In Fig. 3, we compare the error accumulation $\mathcal{E}(t)$ of EH-TEMPO against PT-TEMPO. Two schemes are tested for EH-TEMPO: the backward retrieval scheme and the independent scheme. Both EH-TEMPO schemes maintain errors on the order of 10^{-3} , comparable to PT-TEMPO at the same bond dimension M_S . Crucially, the backward retrieval scheme achieves accuracy similar to the much more expensive independent scheme, confirming that with a proper bond dimension, the efficient backward retrieval introduces negligible additional error.

A key feature of EH-TEMPO is the compressibility of the effective Hamiltonian. This allows the MPO representation of \hat{H}_{eff} to be

compressed to save computational cost without significant loss of accuracy. Figure 4 demonstrates this behavior. In panel (a), without compression ($\varepsilon = 0$), the bond dimension of MPO grows to 1752. By imposing an SVD truncation threshold ε , the bond dimension is drastically reduced (e.g., to 20 for $\varepsilon = 10^{-4}$). Panel (b) shows that this aggressive compression incurs a minimal error penalty: even at $\varepsilon = 10^{-4}$, the deviation from the uncompressed result remains negligible, with errors maintained on the order of 10^{-3} .

The compressibility is further analyzed in Fig. 5 across different parameter regimes. As expected, slower bath dynamics (left panel, ω_c from 1/10 to 1/100 fs^{-1}) and lower temperatures (right panel, T from 300 to 10 K) lead to larger MPO bond dimensions, reflecting stronger memory effect. Nevertheless, in all cases, the compressed bond dimensions are orders of magnitude smaller than the uncompressed values (dashed lines), highlighting the effectiveness of compressing the effective Hamiltonian.

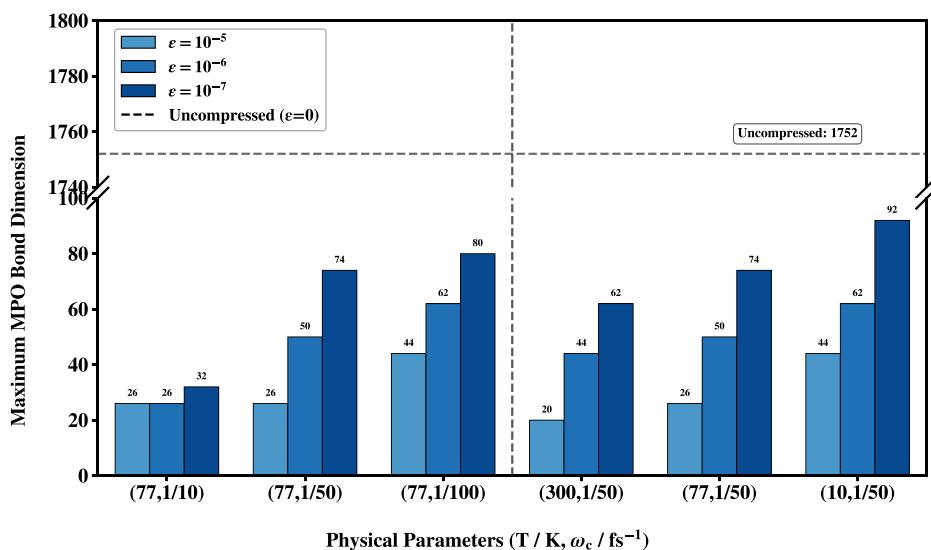


FIG. 5. Comparison of the maximum bond dimension of the effective Hamiltonian MPO with various physical parameters. The x axis labels represent the parameter pairs ($T/\text{K}, \omega_c/\text{fs}^{-1}$). The dashed line indicates the uncompressed bond dimension (1752), which is constant across these parameters for the given discretization. The colored bars show the significantly reduced bond dimensions after compression with thresholds $\varepsilon = 10^{-5}$, 10^{-6} and 10^{-7} , highlighting the efficiency of the compression across different temperatures and cutoff frequencies.

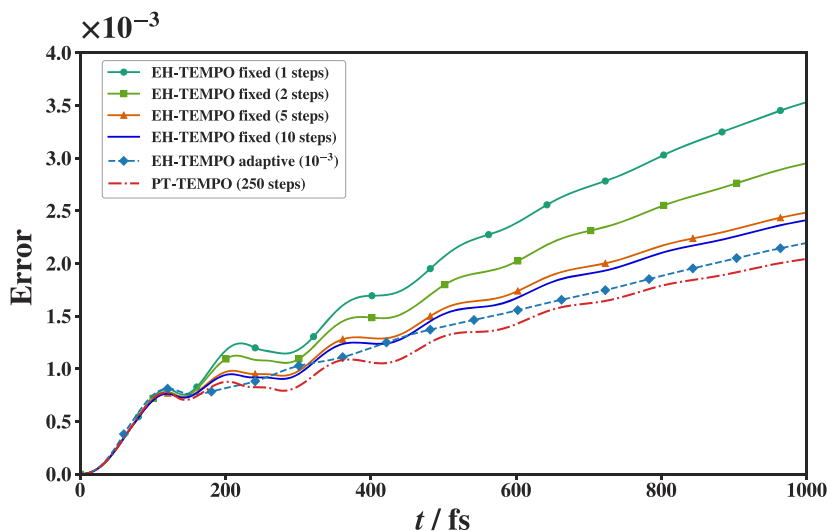


FIG. 6. Convergence analysis of EH-TEMPO with respect to imaginary time evolution step size. The cumulative error $\mathcal{E}(t)$ is shown for EH-TEMPO using fixed step sizes ($N_t = 1, 2, 5, 10$) and an adaptive step-size strategy (relative tolerance 10^{-3}), compared against the PT-TEMPO reference ($N_t = 250$). Notably, even a single evolution step ($N_t = 1$) yields reasonable accuracy, demonstrating the efficiency of the global evolution approach. All calculations used a bond dimension of $M_S = 60$.

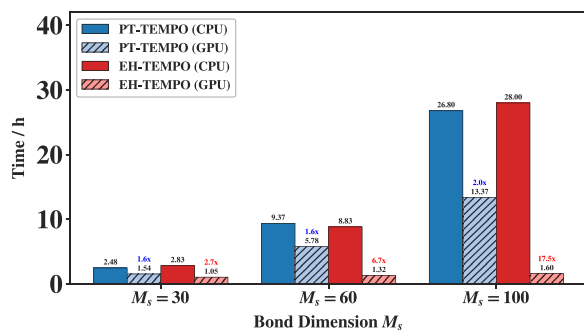


FIG. 7. Computational wall-clock time for the PT-TEMPO and EH-TEMPO algorithms on CPU and GPU platforms as a function of the MPS bond dimension M_S . The numbers above the bars indicate the specific speedup factor achieved by the GPU implementation relative to the CPU. While CPU performance is comparable between the two methods, EH-TEMPO exhibits superior scalability on GPU hardware, achieving a dramatic 17.5 \times speedup at $M_S = 100$ due to free of large tensor decomposition, compared to the modest gains observed for PT-TEMPO.

PT-TEMPO. It should be noted that in the PT-TEMPO implementation, the MPO of the incremental influence functional, whose bond dimension M_O would be $7^2 = 49$, is decomposed into seven separate layers of MPOs, each representing one independent bath with $M_O = 3$ (see Appendix). This decomposition significantly accelerates the PT-TEMPO calculation; without it, simulations with $M_S = 100$ would be computationally prohibitive. However, since EH-TEMPO avoids large tensor decompositions, it is exceptionally well suited for GPU acceleration. On an NVIDIA A100 GPU, EH-TEMPO achieves speedups of 2.7 \times , 6.7 \times , and 17.5 \times for bond dimensions $M_S = 30$, 60, and 100, respectively. In contrast, PT-TEMPO, which relies heavily on SVD and QR decompositions—operations that are inherently difficult to parallelize on GPUs—shows limited benefit from GPU acceleration. With GPU, EH-TEMPO can be one order of magnitude faster than PT-TEMPO for $M_S = 100$. This makes EH-TEMPO the superior choice for high-accuracy simulations requiring large bond dimensions.

B. PBI model

To further demonstrate the robustness of the EH-TEMPO algorithm, particularly for benchmarking against systems with larger reorganization energies and more structured environments, we investigate the excitation energy transfer dynamics in a 4-site perylene bisimide aggregate model.

In contrast to the weakly coupled FMO complex, the PBI model places the system in a notoriously challenging intermediate-to-strong coupling regime. The environment is explicitly modeled using 28 discrete undamped intramolecular modes rather than a continuous, heavily damped spectral density, with parameters adopted from Ref. 56. This structured, undamped bath is further complicated by a substantial reorganization energy of $\lambda = 879.1 \text{ cm}^{-1}$, relative to a nearest-neighbor excitonic coupling of $J = -500 \text{ cm}^{-1}$. This choice of environment represents a highly demanding test case; the absence of bath relaxation prevents the decay of the correlation function, leading to infinitely long-lived memory effects that typically cause rapid exponential scaling or convergence failures in standard path-integral-based methods.

Here, we used EH-TEMPO to simulate the dynamics. The time step size is $\Delta t = 1 \text{ fs}$. The temperature is 300 K. To evaluate the accuracy of EH-TEMPO, the results are compared against numerically exact reference data generated using the TD-DMRG approach (propagating the full system–bath wavefunction with a bond dimension of $M_S = 64$). Figure 8(a) illustrates the convergence of the population at the initial excitation site, $n_0(t)$, with different bond dimension M_S . Because the memory effect is extremely strong, small bond dimensions (e.g., $M_S = 64$ –256) lack the capacity to capture the persistent temporal correlations, yielding trajectories that deviate significantly from the TD-DMRG reference. As M_S is systematically increased, the EH-TEMPO trajectories smoothly converge. At highly expanded bond dimensions ($M_S = 1024$ and 1536), the EH-TEMPO dynamics overlap with the numerically exact TD-DMRG reference.

This application underscores a critical advantage of the EH-TEMPO algorithm. While realistic systems with large reorganization energies and highly structured environmental modes inevitably require much larger bond dimensions to converge than the weakly coupled FMO model, reaching these necessary dimensions ($M_S > 1000$) is computationally prohibitive for standard TEMPO due to its bottleneck in tensor decomposition. EH-TEMPO, by transforming the problem into a “one-shot” imaginary time evolution and utilizing highly parallelizable tensor contractions on GPUs, makes accessing these extreme bond dimensions feasible.

Notably, this model has been successfully solved using the modular path integral approach developed by Makri and coworkers, which mitigates exponential scaling of QuAPI by partitioning the multiple baths into independent modules that are evaluated sequentially.^{57,58} EH-TEMPO can leverage similar ideas to accelerate computations for larger aggregates: local baths can be evaluated via EH-TEMPO, and the multiple baths and system propagators can then be contracted as a 2D tensor network, akin to the strategies employed in Refs. 41 and 43.

However, it is important to note a fundamental difference in computational efficiency between EH-TEMPO and TD-DMRG in this specific undamped regime. While EH-TEMPO successfully converges, the computational cost required to reach $M_S > 1000$ is substantial. In contrast, TD-DMRG converges much faster and with a significantly smaller bond dimension ($M_S = 64$). This discrepancy arises from the underlying mathematical structures of the two approaches. EH-TEMPO relies on compressing memory effects along the time axis; an undamped bath produces an infinitely long-lived correlation function, which severely penalizes the compressibility of the influence functional and demands massive bond dimensions to retain un-decaying temporal correlations. Conversely, TD-DMRG propagates the joint system–bath wavefunction explicitly. For a limited number of discrete modes over a short time, the entanglement entropy between the system and the bath modes remains relatively constrained, allowing wavefunction-based methods to capture the dynamics highly efficiently. Therefore, while EH-TEMPO successfully extends the reach of influence functional methods to the strong-coupling regime, explicit wavefunction methods, such as TD-DMRG, are better suited for environments dominated by discrete, undamped modes.

In addition to the purely undamped discrete modes, we also investigated the system by replacing the environment with a Brownian oscillator spectral density, which describes the discrete modes coupled to a secondary Ohmic bath,

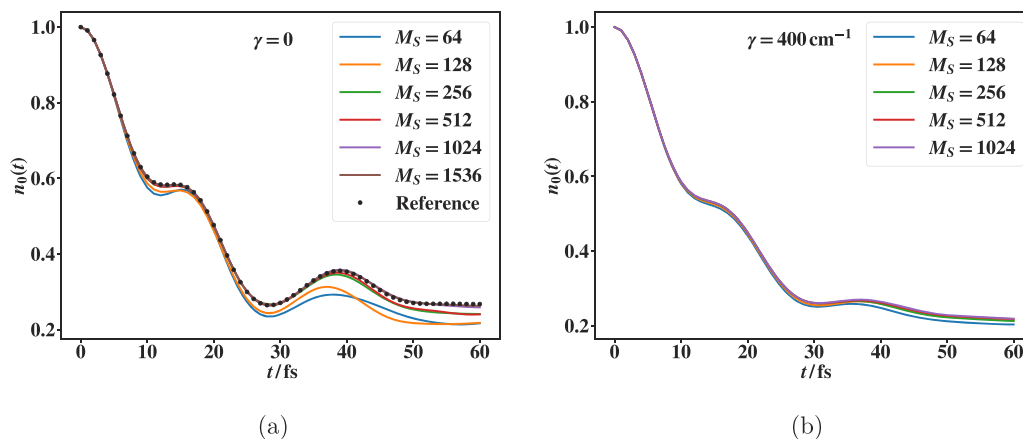


FIG. 8. Population dynamics $n_0(t)$ of the 4-site PBI aggregate model with (a) an undamped discrete mode bath ($\gamma = 0$) and (b) a damped Brownian bath ($\gamma = 400 \text{ cm}^{-1}$). The solid lines show the EH-TEMPO results converging as the bond dimension M_S increases. The black dots represent the numerically exact TD-DMRG reference at $\gamma = 0$.

$$J(\omega) = \sum_{i=1}^{28} \frac{2\lambda_i \gamma_i \Omega_i^2 \omega}{(\Omega_i^2 - \omega^2)^2 + \gamma_i^2 \omega^2}, \quad (22)$$

where Ω_i is the frequency of the discrete mode with reorganization energy λ_i and γ_i is the damping factor. We applied a uniform damping parameter $\gamma_i = \gamma = 400 \text{ cm}^{-1}$ across all modes. Consequently, modes with frequencies $\Omega_i < 400 \text{ cm}^{-1}$ are overdamped, whereas those with $\Omega_i > 400 \text{ cm}^{-1}$ are underdamped. The discrete mode case discussed above corresponds to $\gamma = 0$. The dynamics are presented in Fig. 8(b). Compared to the undamped case, convergence is achieved with significantly smaller bond dimensions M_S . This behavior is expected, as the introduction of the damping factor accelerates the decay of the bath correlation function, thereby weakening the long-lived memory effects.

IV. CONCLUSION

In this work, we have introduced the effective Hamiltonian-based time-evolving matrix product operator (EH-TEMPO) algorithm as a highly scalable algorithm for simulating the non-Markovian dynamics of open quantum systems. The central innovation of our approach lies in reformulating the Feynman–Vernon influence functional into the imaginary time evolution of an effective Hamiltonian. This exact mathematical mapping naturally casts the effective Hamiltonian into a sum-of-products form, which seamlessly enables the use of automated, optimal MPO construction algorithms. Consequently, EH-TEMPO successfully replaces the computationally prohibitive iterative contraction-and-compression schemes of standard TEMPO and PT-TEMPO with a highly efficient projector-splitting time evolution algorithm based on time-dependent variational principle.

We benchmarked the accuracy and efficiency of EH-TEMPO on the excitation energy transfer dynamics of the 7-site Fenna–Matthews–Olson complex and a 4-site perylene bisimide aggregate. Across these diverse environments, ranging from weakly coupled continuous Debye spectral densities to strongly coupled discrete undamped modes, EH-TEMPO accurately reproduced the

numerically exact population dynamics obtained from HEOM and TD-DMRG. A pivotal finding of our analysis is the intrinsic compressibility of the effective Hamiltonian’s MPO representation; we demonstrated that its bond dimensions can be aggressively truncated without compromising numerical accuracy.

A defining practical advantage of EH-TEMPO is its ability to circumvent the severe bottleneck of large tensor decompositions inherent to the original TEMPO algorithms. By relying primarily on tensor contractions, our method is exceptionally well-suited for massive GPU parallelization. In our numerical tests, EH-TEMPO achieved dramatic performance gains, delivering speedups of up to 17.5× on a single GPU for typical bond dimensions ($M_S = 100$) compared to standard CPU implementations. These results indicate that EH-TEMPO is uniquely positioned to tackle complex, multi-state quantum systems that were previously computationally intractable.

In future work, exploring the synergy between our EH-TEMPO algorithm and the SMatPI algorithm presents a highly promising avenue. In particular, EH-TEMPO is exceptionally well-equipped to efficiently generate the influence functional data required within the memory time to initialize the SMatPI recursion. Integrating these methodologies could further alleviate memory bottlenecks, minimize tensor dimensions, and drastically reduce the overall computational scaling for simulating long-time non-Markovian dynamics in large-scale systems.

ACKNOWLEDGMENTS

The authors thank Chu Guo for the insightful discussions. This work is supported by the National Natural Science Foundation of China (Grant Nos. 22422301 and 22273005), the Quantum Science and Technology-National Science and Technology Major Project (Grant No. 2023ZD0300200), and the NSAF (Grant No. U2330201).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Xiaoyu Yang: Investigation (lead); Methodology (equal); Writing – original draft (lead). **Limin Liu:** Investigation (supporting); Methodology (equal). **Wencheng Zhao:** Investigation (supporting). **Jiajun Ren:** Conceptualization (lead); Methodology (equal); Supervision (lead); Writing – review & editing (lead). **Wei-Hai Fang:** Supervision (supporting).

DATA AVAILABILITY

The data that support the findings of this study are available within the article.

APPENDIX: COMPUTATIONAL COMPLEXITY ANALYSIS

A detailed analysis of the computational complexity, specifically tensor contraction and tensor decomposition of the (PT-)TEMPO algorithms compared to the proposed EH-TEMPO algorithm for diagonal system–bath couplings, is provided in the [Appendix](#).

For clarity, we define the following variables:

- N_t : total number of real-time steps.
- d : system Hilbert space dimension.
- M_S : bond dimension of the MPS representing the overall influence functional.
- M_O : bond dimension of the MPO representing the incremental influence functional (in TEMPO) or the effective Hamiltonian (in EH-TEMPO).
- p : the physical dimension of the local MPS and MPO sites.

Note that the following scaling analysis considers retaining full memory without truncation.

1. Standard and optimized TEMPO algorithms

The standard MPO \times MPS tensor contraction at a single site scales as $\mathcal{O}(M_S^2 M_O p^2)$. However, if the MPO is diagonal in the physical indices, this contraction can be optimized via the Hadamard product, reducing the scaling to $\mathcal{O}(M_S^2 M_O p)$. The subsequent tensor decomposition (for canonicalization and compression) of the resulting enlarged tensor scales as $\mathcal{O}(M_S^3 M_O p)$.⁴⁸

There are several strategies to construct the incremental influence functional in the standard TEMPO framework, as follows:

- Grouped path variables (original): Following the implementation in Ref. 29, the forward and backward path variables for a single time step are grouped into a single site. This

results in a physical dimension $p = d^2$. The single-layer MPO representing the incremental influence functional, $\exp[-\sum_{k'} (s_k^+ - s_k^-)(\eta_{kk'} s_{k'}^+ - \eta_{kk'}^* s_{k'}^-)]$, has a bond dimension $M_O = \mathcal{O}(d^2)$. Utilizing the diagonal property of the MPO, the contraction complexity per site is $\mathcal{O}(M_S^2 d^6)$, and the decomposition complexity is $\mathcal{O}(M_S^3 d^8)$. To simulate the dynamics over N_t time steps, there are $\mathcal{O}(N_t)$ layers, and each layer contains $\mathcal{O}(N_t)$ sites. Thus, the cumulative scaling becomes $\mathcal{O}(N_t^2 M_S^2 d^6)$ for contraction and $\mathcal{O}(N_t^2 M_S^3 d^8)$ for decomposition.

- Separated path variables (optimized): A significant optimization can be achieved by separating the forward and backward variables into distinct sites. This reduces the physical dimension per site to $p = d$. The incremental MPO can also be split into forward $\exp[-\sum_{k'} s_k^+ (\eta_{kk'} s_{k'}^+ - \eta_{kk'}^* s_{k'}^-)]$ and backward $\exp[+\sum_{k'} s_k^- (\eta_{kk'} s_{k'}^+ - \eta_{kk'}^* s_{k'}^-)]$ components, yielding two MPOs with $M_O = \mathcal{O}(d)$. Consequently, the scaling per site is reduced to $\mathcal{O}(M_S^2 d^3)$ for contraction and $\mathcal{O}(M_S^3 d^4)$ for decomposition. Over N_t time steps, the overall complexity becomes $\mathcal{O}(N_t^2 M_S^2 d^3)$ and $\mathcal{O}(N_t^2 M_S^3 d^4)$, respectively. This scaling matches the results reported in Ref. 42 (Table I). Although this optimized implementation doubles both the number of sites and the number of layers, the computational scaling with respect to the system size d is dramatically reduced.

The analysis above is universal for any diagonal system–bath coupling. However, for models such as the Frenkel–Holstein model, the d states can be represented as d two-level sites, each coupled to an independent bath. Here, the incremental influence functional can be factorized into d independent components. Therefore, each MPO layer can also be factorized into d sub-layers applied successively to the MPS. In this case, M_O effectively becomes $\mathcal{O}(1)$ per sub-layer. Following the optimized TEMPO algorithm, the single-site scaling reduces to $\mathcal{O}(M_S^2 d)$ for contraction and $\mathcal{O}(M_S^3 d)$ for decomposition. Because the simulation involves $\mathcal{O}(N_t d)$ sub-layers with $\mathcal{O}(N_t)$ sites each, the total scaling over the entire simulation is $\mathcal{O}(N_t^2 M_S^2 d^2)$ for contraction and $\mathcal{O}(N_t^2 M_S^3 d^2)$ for decomposition.

2. EH-TEMPO algorithm

EH-TEMPO fundamentally alters the computational bottleneck by framing the problem as an imaginary time evolution governed by the time-dependent variational principle, thereby avoiding the direct decomposition of heavily enlarged $M_S \times M_O$ tensors.

In the one-site TDVP-PS evolution scheme, applying the effective Hamiltonian MPO to the MPS involves projecting the action of the MPO onto the tangent space of the MPS. Generally, the contraction complexity per site is $\mathcal{O}(M_S^2 M_O^2 p^2 + M_S^3 M_O p)$.⁴⁸ By

TABLE I. Summary of overall computational complexity for maintaining full memory across N_t time steps. N_t is the number of imaginary time steps, and l is the number of IVP solver iterations per site (typically ~ 10).

Algorithm	Contraction scaling	Decomposition scaling
Standard TEMPO	$\mathcal{O}(N_t^2 M_S^2 d^6)$	$\mathcal{O}(N_t^2 M_S^3 d^8)$
Optimized TEMPO	$\mathcal{O}(N_t^2 M_S^2 d^3)$	$\mathcal{O}(N_t^2 M_S^3 d^4)$
Optimized TEMPO (Holstein-specific)	$\mathcal{O}(N_t^2 M_S^2 d^2)$	$\mathcal{O}(N_t^2 M_S^3 d^2)$
EH-TEMPO	$\mathcal{O}(l N_t N_t (M_S^2 M_O^2 d + M_S^3 M_O d))$	$\mathcal{O}(l N_t N_t M_S^3 d)$

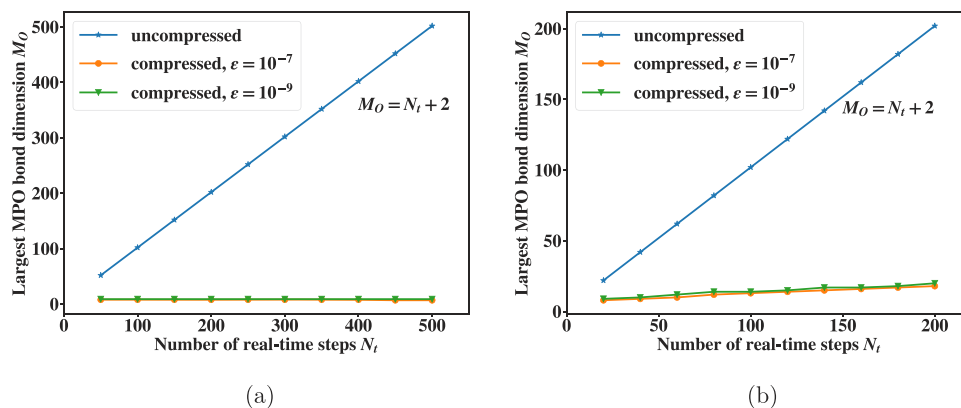


FIG. 9. Largest uncompressed and compressed MPO bond dimensions, M_O , of the effective Hamiltonian for a single local bath as a function of the number of real-time steps N_t . The results are presented for (a) the 7-site FMO complex ($\Delta t = 4$ fs) and (b) the 4-site PBI aggregate model ($\Delta t = 1$ fs). Two SVD compression thresholds, 10^{-7} and 10^{-9} , are compared. For the FMO model, the largest compressed bond dimensions remain constant at $M_O = 8$ (for 10^{-7}) and $M_O = 9$ (for 10^{-9}) across all values of N_t . For the PBI model, the largest compressed bond dimensions at $N_t = 20$ are $M_O = 8$ (10^{-7}) and $M_O = 9$ (10^{-9}); at $N_t = 200$, these values increase only moderately to $M_O = 18$ and $M_O = 20$, respectively.

exploiting the diagonality of the effective Hamiltonian MPO, this reduces to $\mathcal{O}(M_S^2 M_O^2 p + M_S^3 M_O p)$. Unlike standard TEMPO, updating each local site in EH-TEMPO requires l tensor contractions (matrix-vector multiplications, $H \cdot c$) as demanded by the initial value problem (IVP) solver (typically $l \sim 10$). Crucially, the decomposition step within TDVP only scales as $\mathcal{O}(M_S^3 p)$ per site, because it operates directly on the MPS tensors rather than the combined MPO \times MPS tensors.

Using separated path variables ($p = d$), a full TDVP sweep over the entire chain of length $\mathcal{O}(N_t)$ across a total of N_τ imaginary time evolution steps results in an overall scaling of $\mathcal{O}(l N_t N_\tau (M_S^2 M_O^2 d + M_S^3 M_O d))$ for contraction and $\mathcal{O}(N_t N_\tau M_S^3 d)$ for decomposition.

For a single bath, the exact effective Hamiltonian MPO exhibits a bond dimension of $M_O = \mathcal{O}(N_t)$. Fortunately, as demonstrated in the main text, this effective Hamiltonian MPO is highly compressible. Regardless of whether the system is coupled to a broad continuous Debye spectral density or to discrete undamped modes, where the memory length spans the entire simulation time ($N_t \Delta t$), the actual compressed bond dimension M_O becomes effectively independent of N_t , as illustrated in Fig. 9. This intrinsic compressibility was similarly observed in Ref. 53, which employed the Prony fitting algorithm to construct the effective Hamiltonian MPO. That work revealed that M_O is fundamentally governed by the number of terms required to decompose the memory coefficient $\eta_{\Delta k}$ into a sum of complex exponentials, $\eta_{\Delta k} \approx \sum_l \alpha_l \lambda_l^{\Delta k}$ (where α_l and λ_l are the expansion parameters). The required number of fitting terms is significantly smaller than the total number of time steps.

For systems coupled to d independent local baths, such as the Frenkel-Holstein model, the total effective Hamiltonian is a direct summation of the individual bath contributions. Consequently, while the MPO for each individual bath component can be aggressively compressed, the overall MPO blocks cannot be interleaved or further compressed across different bath sites. As a result, the total bond dimension scales additively with the number of independent baths, yielding a compressed scaling of $M_O = \mathcal{O}(d)$.

3. Summary of computational complexity

Table I summarizes the computational complexities for maintaining full memory dynamics across N_t time steps. As shown, EH-TEMPO significantly mitigates the harsh tensor decomposition bottleneck associated with standard TEMPO algorithms.

REFERENCES

- A. Nitzan, *Chemical Dynamics in Condensed Phases: Relaxation, Transfer and Reactions in Condensed Molecular Systems* (Oxford University Press, 2006).
- U. Weiss, *Quantum Dissipative Systems* (World Scientific, 2012).
- V. May and O. Kühn, *Charge and Energy Transfer Dynamics in Molecular Systems* (John Wiley & Sons, 2023).
- Y. Yan, Y. Liu, T. Xing, and Q. Shi, "Theoretical study of excitation energy transfer and nonlinear spectroscopy of photosynthetic light-harvesting complexes using the nonperturbative reduced dynamics method," *WIREs Comput. Mol. Sci.* **11**, e1498 (2021).
- S. J. Jang and B. Mennucci, "Delocalized excitons in natural light-harvesting complexes," *Rev. Mod. Phys.* **90**, 035003 (2018).
- M. Xu, V. Vadimov, J. T. Stockburger, and J. Ankerhold, "Colloquium: Simulating non-Markovian dynamics in open quantum systems," *Rev. Mod. Phys.* **98**, 021002 (2026).
- H. Ma, Z. Luo, and Y. Yao, "The time-dependent density matrix renormalisation group method," *Mol. Phys.* **116**, 854–868 (2018).
- J. Ren, W. Li, T. Jiang, Y. Wang, and Z. Shuai, "Time-dependent density matrix renormalization group method for quantum dynamics in complex systems," *WIREs Comput. Mol. Sci.* **12**, e1614 (2022).
- F. A. Y. N. Schröder, D. H. P. Turban, A. J. Musser, N. D. M. Hine, and A. W. Chin, "Tensor network simulation of multi-environmental open quantum dynamics via machine learning and entanglement renormalisation," *Nat. Commun.* **10**, 1062 (2019).
- W. Li, J. Ren, H. Yang, H. Wang, and Z. Shuai, "Optimal tree tensor network operators for tensor network simulations: Applications to open quantum systems," *J. Chem. Phys.* **161**, 054116 (2024).
- H. Wang and M. Thoss, "Multilayer formulation of the multiconfiguration time-dependent Hartree theory," *J. Chem. Phys.* **119**, 1289–1299 (2003).
- H. Wang, "Multilayer multiconfiguration time-dependent Hartree theory," *J. Phys. Chem. A* **119**, 7951–7965 (2015).
- G. Lindblad, "On the generators of quantum dynamical semigroups," *Commun. Math. Phys.* **48**, 119–130 (1976).

- ¹⁴V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, "Completely positive dynamical semigroups of N-level systems," *J. Math. Phys.* **17**, 821–825 (1976).
- ¹⁵A. G. Redfield, "On the theory of relaxation processes," *IBM J. Res. Dev.* **1**, 19–31 (1957).
- ¹⁶A. G. Redfield, "The theory of relaxation processes," in *Advances in Magnetic and Optical Resonance* (Elsevier, 1965), Vol. 1, pp. 1–32.
- ¹⁷Y. Tanimura and R. Kubo, "Time evolution of a quantum system in contact with a nearly Gaussian–Markoffian noise bath," *J. Phys. Soc. Jpn.* **58**, 101–114 (1989).
- ¹⁸Y. Tanimura, "Numerically 'exact' approach to open quantum dynamics: The hierarchical equations of motion (HEOM)," *J. Chem. Phys.* **153**, 020901 (2020).
- ¹⁹L. Diósi, N. Gisin, and W. T. Strunz, "Non-Markovian quantum state diffusion," *Phys. Rev. A* **58**, 1699 (1998).
- ²⁰D. Suess, A. Eisfeld, and W. T. Strunz, "Hierarchy of stochastic pure states for open quantum system dynamics," *Phys. Rev. Lett.* **113**, 150403 (2014).
- ²¹N. Makri and D. E. Makarov, "Tensor propagator for iterative quantum time evolution of reduced density matrices. I. Theory," *J. Chem. Phys.* **102**, 4600–4610 (1995).
- ²²N. Makri and D. E. Makarov, "Tensor propagator for iterative quantum time evolution of reduced density matrices. II. Numerical methodology," *J. Chem. Phys.* **102**, 4611–4618 (1995).
- ²³N. Makri, "Improved Feynman propagators on a grid and non-adiabatic corrections within the path integral framework," *Chem. Phys. Lett.* **193**, 435–445 (1992).
- ²⁴N. Makri, "Quantum dissipative dynamics: A numerically exact methodology," *J. Phys. Chem. A* **102**, 4414–4427 (1998).
- ²⁵E. Sim, "Quantum dynamics for a system coupled to slow baths: On-the-fly filtered propagator method," *J. Chem. Phys.* **115**, 4450–4456 (2001).
- ²⁶N. Makri, "Iterative blip-summed path integral for quantum dynamics in strongly dissipative environments," *J. Chem. Phys.* **146**, 134101 (2017).
- ²⁷R. Ovcharenko, X. Xu, and B. P. Fingerhut, "Scalable distributed memory implementation of the quasi-adiabatic propagator path integral," *J. Chem. Theory Comput.* **22**, 3810–3826 (2026).
- ²⁸N. Makri, "Small matrix disentanglement of the path integral: Overcoming the exponential tensor scaling with memory length," *J. Chem. Phys.* **152**, 041104 (2020).
- ²⁹A. Strathearn, P. Kirton, D. Kilda, J. Keeling, and B. W. Lovett, "Efficient non-Markovian quantum dynamics using time-evolving matrix product operators," *Nat. Commun.* **9**, 3322 (2018).
- ³⁰E. Ye and G. K.-L. Chan, "Constructing tensor network influence functionals for general quantum dynamics," *J. Chem. Phys.* **155**, 044104 (2021).
- ³¹M. Cygorek, M. Cosacchi, A. Vagov, V. M. Axt, B. W. Lovett, J. Keeling, and E. M. Gauger, "Simulation of open quantum systems by automated compression of arbitrary environments," *Nat. Phys.* **18**, 662–668 (2022).
- ³²U. Schollwöck, "The density-matrix renormalization group in the age of matrix product states," *Ann. Phys.* **326**, 96–192 (2011).
- ³³M. R. Jørgensen and F. A. Pollock, "Exploiting the causal tensor network structure of quantum processes to efficiently simulate non-Markovian path integrals," *Phys. Rev. Lett.* **123**, 240602 (2019).
- ³⁴M. Richter and S. Hughes, "Enhanced TEMPO algorithm for quantum path integrals with off-diagonal system-bath coupling: Applications to photonic quantum networks," *Phys. Rev. Lett.* **128**, 167403 (2022).
- ³⁵D. Gribben, D. M. Rouse, J. Iles-Smith, A. Strathearn, H. Maguire, P. Kirton, A. Nazir, E. M. Gauger, and B. W. Lovett, "Exact dynamics of nonadditive environments in non-Markovian open quantum systems," *PRX Quantum* **3**, 010321 (2022).
- ³⁶R. Chen, "Path integral formalism of open quantum systems with non-diagonal system-bath coupling," *Commun. Theor. Phys.* **76**, 115701 (2024).
- ³⁷S. Zhang and Q. Shi, "Time-evolving matrix product operator method in a non-diagonal basis set based on the derivative of the path-integral expression," *Phys. Rev. A* **111**, 062210 (2025).
- ³⁸A. Bose, "Quantum correlation functions through tensor network path integral," *J. Chem. Phys.* **159**, 214110 (2023).
- ³⁹N. Ng, G. Park, A. J. Millis, G. K.-L. Chan, and D. R. Reichman, "Real-time evolution of Anderson impurity models via tensor network influence functionals," *Phys. Rev. B* **107**, 125103 (2023).
- ⁴⁰R. Chen, X. Xu, and C. Guo, "Grassmann time-evolving matrix product operators for quantum impurity models," *Phys. Rev. B* **109**, 045140 (2024).
- ⁴¹R. Chen, X. Xu, and C. Guo, "Grassmann time-evolving matrix product operators for equilibrium quantum impurity problems," *New J. Phys.* **26**, 013019 (2024).
- ⁴²S. Kundu and N. Makri, "PathSum: A C++ and Fortran suite of fully quantum mechanical real-time path integral methods for (multi-) system + bath dynamics," *J. Chem. Phys.* **158**, 224801 (2023).
- ⁴³A. Bose and P. L. Walters, "A multisite decomposition of the tensor network path integrals," *J. Chem. Phys.* **156**, 024101 (2022).
- ⁴⁴C. Hubig, I. P. McCulloch, and U. Schollwöck, "Generic construction of efficient matrix product operators," *Phys. Rev. B* **95**, 035129 (2017).
- ⁴⁵J. Ren, W. Li, T. Jiang, and Z. Shuai, "A general automatic method for optimal construction of matrix product operators using bipartite graph theory," *J. Chem. Phys.* **153**, 084118 (2020).
- ⁴⁶H. Çakır, R. M. Milbradt, and C. B. Mendl, "Optimal symbolic construction of matrix product operators and tree tensor network operators," *Phys. Rev. B* **112**, 035101 (2025).
- ⁴⁷J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, "Unifying time evolution and optimization with matrix product states," *Phys. Rev. B* **94**, 165116 (2016).
- ⁴⁸W. Li, J. Ren, and Z. Shuai, "Numerical assessment for accuracy and GPU acceleration of TD-DMRG time evolution schemes," *J. Chem. Phys.* **152**, 024127 (2020).
- ⁴⁹J. Shao and N. Makri, "Iterative path integral calculation of quantum correlation functions for dissipative systems," *Chem. Phys.* **268**, 1–10 (2001).
- ⁵⁰J. Shao and N. Makri, "Iterative path integral formulation of equilibrium correlation functions for quantum dissipative systems," *J. Chem. Phys.* **116**, 507–514 (2002).
- ⁵¹R. P. Feynman and F. L. Vernon, Jr., "The theory of a general quantum system interacting with a linear dissipative system," *Ann. Phys.* **281**, 547–607 (2000).
- ⁵²A. J. Leggett, S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger, "Dynamics of the dissipative two-state system," *Rev. Mod. Phys.* **59**, 1 (1987).
- ⁵³C. Guo and R. Chen, "Efficient construction of the Feynman–Vernon influence functional as matrix product states," *SciPost Phys. Core* **7**, 063 (2024).
- ⁵⁴A. Ishizaki and G. R. Fleming, "Theoretical examination of quantum coherence in a photosynthetic system at physiological temperature," *Proc. Natl. Acad. Sci. U. S. A.* **106**, 17255–17260 (2009).
- ⁵⁵J. Adolphs and T. Renger, "How proteins trigger excitation energy transfer in the FMO complex of green sulfur bacteria," *Biophys. J.* **91**, 2778–2797 (2006).
- ⁵⁶D. Ambrosek, A. Köhn, J. Schulze, and O. Kühn, "Quantum chemical parametrization and spectroscopic characterization of the Frenkel exciton Hamiltonian for a j-aggregate forming perylene bisimide dye," *J. Phys. Chem. A* **116**, 11451–11458 (2012).
- ⁵⁷S. Kundu and N. Makri, "Modular path integral for finite-temperature dynamics of extended systems with intramolecular vibrations," *J. Chem. Phys.* **153**, 044124 (2020).
- ⁵⁸S. Kundu and N. Makri, "Efficient matrix factorisation of the modular path integral for extended systems," *Mol. Phys.* **119**, e1797200 (2021).