



Heriot-Watt University
Research Gateway

Tree-like process tensor contraction for automated compression of environments

Citation for published version:

Cygorek, M, Lovett, BW, Keeling, J & Gauger, EM 2024, 'Tree-like process tensor contraction for automated compression of environments', *Physical Review Research*.

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Peer reviewed version

Published In:

Physical Review Research

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Tree-like process tensor contraction for automated compression of environments

Moritz Cygorek,^{1,2} Brendon W. Lovett,³ Jonathan Keeling,³ and Erik M. Gauger²

¹*Condensed Matter Theory, Department of Physics, TU Dortmund, 44221 Dortmund, Germany*

²*SUPA, Institute of Photonics and Quantum Sciences,*

Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom

³*SUPA, School of Physics and Astronomy, University of St Andrews, St Andrews KY16 9SS, United Kingdom*

The algorithm “automated compression of environments” (ACE) [Nat. Phys. 18, 662 (2022)] provides a versatile way of simulating an extremely broad class of open quantum systems. This is achieved by encapsulating the influence of the environment, which is determined by the interaction Hamiltonian(s) and initial states, into compact process tensor matrix product operator (PT-MPO) representations. The generality of the ACE method comes at high numerical cost. Here, we demonstrate that orders-of-magnitude improvement of ACE is possible by changing the order of PT-MPO contraction from a sequential to a tree-like scheme. The problem of combining two partial PT-MPOs with large inner bonds is solved by a preselection approach. The drawbacks of the preselection approach are that the MPO compression is suboptimal and that it is more prone to error accumulation than sequential combination and compression. We therefore also identify strategies to mitigate these disadvantages by fine-tuning compression parameters. This results in a scheme that is similar in compression efficiency and accuracy to the original ACE algorithm, yet is significantly faster. Our numerical experiments reach similar conclusions for bosonic and fermionic test cases, suggesting that our findings are characteristic of the combination of PT-MPOs more generally.

I. INTRODUCTION

Understanding how quantum systems are affected by their environments is a challenging but important task, in particular when the system-environment coupling is not weak and the Born-Markov approximation becomes unreliable [1]. For example, strong interactions between system and environment can shake up environment excitations, which dress and screen the quantum system by polaron formation [2] or can stabilize localized spins in fermionic impurity models [3]. The richness of phenomena in open quantum systems is further illustrated in the spin-boson model, which features phase transitions between localized and delocalized states as well as regimes with algebraic versus exponential decay [4, 5], depending on the spectrum of the bosonic bath. These insights play a crucial role in practical applications, e.g., for understanding phonon effects in semiconductor quantum dots with their typical [6] super-Ohmic spectral density with exponential cut-off. This coupling explains the appearance of non-Markovian dynamics [7–11], dynamical decoupling [12–15], phonon-assisted state preparation [16–18], and incomplete decay of coherences resulting in long-lived cooperative emission effects [19, 20]. Even in cases where the decay of excitations or coherences can be described by effective Markovian master equations [21], the corresponding decay rates and Lindblad operators can be challenging to obtain. Because they generally depend on system eigenstates and thus on how the system is driven [2], they are particularly difficult to extract in situations with ultrafast time-dependent driving [17, 22] or when system states are nearly degenerate [21, 23].

To resolve such challenging questions in strongly coupled open quantum systems, it is highly desirable to have powerful numerical tools to simulate complex composite open quantum systems, possibly coupled to several

different environments, numerically exactly; that is, to approximate exact results with high precision controlled by convergence parameters without imposing additional approximations such as perturbation theory or the Born-Markov approximation.

PT-MPOs [24, 25] provide exactly this: They describe the influence of the environment, equivalent to the Feynman-Vernon influence functional [26], in a compact matrix product operator representation. Once obtained, they can be used to efficiently simulate the dynamics of the open quantum system for a given time-dependent system Hamiltonian, e.g., to find optimal driving protocols [27]. Moreover, multiple PT-MPOs can be stacked together to describe quantum systems coupled to multiple environments while still remaining numerically exact [19, 28, 29].

By now, several algorithms to calculate PT-MPOs have been developed. Most are aimed at solving either a generalized spin-boson model [25, 30, 31] or quantum impurities coupled to fermionic baths [32, 33]. In both cases Gaussian path integrals can be employed to derive compact expressions for the nodes of the tensor network to be contracted [34]. PT-MPOs can also be constructed for more general environments using the algorithm *automated compression of environments* (ACE) [28], which starts from a different tensor network: In this case, one need only assume that the environment is composed of independent modes, which could be, e.g., bosons, fermions, spins, or indeed any general non-Gaussian environment. Moreover, these modes could be time-dependently driven and themselves subject to Markovian loss or dephasing. For each mode, a PT-MPO is constructed from the respective Hamiltonian. In ACE, the individual PT-MPOs are then sequentially combined and compressed, yielding a PT-MPO encapsulating the joint influence of all modes. This algorithm is sufficiently abstract that it can be im-

plemented by a computer code that is independent of the specific physical model at hand. Such a code can serve as a general-purpose solver to be utilized as a “black box” by the user. A C++ implementation is available and described in Ref. [35].

The price to pay for the generality of ACE is the relatively high numerical cost in comparison to specialized PT-MPO methods for Gaussian environments. This relative cost is even more significant given recent methodological developments [30, 31] that further reduce the cost for Gaussian environments. Here we show that ideas used in one of these developments [30] can also lead to a significant speed up for ACE. Specifically, we will make use of a scheme introduced there to combine and compress MPOs with moderate to large inner bond dimensions by ‘preselecting’ degrees of freedom based on singular value decomposition (SVD) of the individual MPOs. This ability to combine MPOs with larger bond dimensions enables the contraction of the corresponding tensor network in a different order. In particular, the preselection technique facilitates the partial contraction of blocks of the tensor network and the subsequent combination of those blocks at a later step in the algorithm [30].

In this article, we demonstrate that ACE simulations can be significantly sped up by using a different contraction order of the tensor network. Concretely, we suggest to contract PT-MPOs pairwise, yielding a tree-like contraction scheme. (A simple reordering of the linear sequence in which modes are contracted provides little improvement, see Appendix A.) As in Ref. [30], this scheme is enabled by preselecting degrees of freedom when combining and compressing PT-MPOs. We show several numerical examples for which this tree-like contraction scheme significantly outperforms the previous sequential contraction ACE algorithm [28]. However, we also observe that preselection results in suboptimal PT-MPO contraction, i.e. it produces PT-MPOs with larger bond dimensions than the sequential ACE algorithm. Furthermore, it accumulates larger numerical errors. We show that both of these issues can be addressed by fine-tuning convergence parameters, in particular, adjusting the number of compression sweeps per MPO combination and by choosing a scale-dependent compression threshold.

The article is structured as follows: In Section II we lay out the theory. We start with a brief summary of the ACE tensor network in Section II A before we discuss different schemes to contract this network in Section II B. This includes a brief revision of the sequential contraction scheme of Ref. [28] and of the preselection approach for combining MPOs with medium to large inner bonds from Ref. [30], followed by the introduction of our tree-like contraction scheme as well as the description of strategies for fine-tuning the MPO compression. Numerical results are presented in Section III, where we discuss the model systems in Section III A, the computation times required for simulations using different approaches in Section III B, the corresponding numerical

errors in Section III C, and the PT-MPO compression efficacy in Section III D. Our findings are summarized in the discussion Section IV. Further, there are three appendices: In Appendix A, we briefly investigate the impact of rearranging modes in a different sequence, while in Appendix B we provide an elementary derivation of how to obtain fermionic PT-MPOs using ACE. Finally, convergence is discussed in Appendix C using an alternative measure.

II. THEORY

A. ACE tensor network

We first summarize the tensor network that is the basis of ACE [28]. The dynamics of a quantum system of interest (S) coupled to an environment (E) is determined by the Liouville-von Neumann equation for the total density matrix

$$\frac{\partial}{\partial t}\rho = \mathcal{L}\rho = (\mathcal{L}_S + \mathcal{L}_E)\rho, \quad (1)$$

where we split the Liouvillian into a part acting only on the system degrees of freedom \mathcal{L}_S , and an environment part \mathcal{L}_E , which also includes the interaction with the system. These Liouvillians are given by the respective system and environment Hamiltonians via $\mathcal{L}_{S,E}\rho = -\frac{i}{\hbar}[H_{S,E}, \rho]$. Such a description can be straightforwardly extended to include additional Markovian loss and decoherence channels via Lindbladians as well as time-dependent Hamiltonians. It is convenient to work within a superoperator formalism [36], where the density matrix ρ is viewed as a vector and the Liouvillians become matrices. Then, Eq. (1) is formally solved by the matrix exponential $\rho(t) = e^{\mathcal{L}(t-t_0)}\rho(t_0)$.

The ACE algorithm is obtained by the following steps: (i) We introduce a time grid $t_l = t_0 + l\Delta t$ with time steps Δt and write $e^{\mathcal{L}t_n} = \prod_{l=1}^n e^{\mathcal{L}\Delta t}$. (ii) A Trotter decomposition is employed to split the exponential into system and environment parts $e^{(\mathcal{L}_S + \mathcal{L}_E)\Delta t} = e^{\mathcal{L}_E\Delta t}e^{\mathcal{L}_S\Delta t} + \mathcal{O}(\Delta t^2)$. (iii) We assume that at time t_0 system and environment are uncorrelated $\rho(t_0) = \bar{\rho}(t_0) \otimes \rho^E(t_0)$, and trace out the environment at time t_n . This yields

$$\bar{\rho}(t_n) = \text{Tr}_E \left\{ \left(\prod_{l=1}^n e^{\mathcal{L}_E\Delta t} e^{\mathcal{L}_S\Delta t} \right) (\bar{\rho}(t_0) \otimes \rho^E(t_0)) \right\}, \quad (2)$$

which is cast into the form

$$\bar{\rho}_{\alpha_n} = \sum_{\substack{\alpha_{n-1}, \dots, \alpha_0 \\ \alpha'_n, \dots, \alpha'_1 \\ d_{n-1}, \dots, d_1}} \prod_{l=1}^n \left(\mathcal{Q}_{d_l d_{l-1}}^{(\alpha_l, \alpha'_l)} \mathcal{M}^{\alpha'_l \alpha_{l-1}} \right) \bar{\rho}_{\alpha_0}, \quad (3)$$

where Liouville space indices $\alpha_l = (\nu_l, \mu_l)$ collect left and right (Hilbert space) indices on the reduced system density matrix at time t_l , i.e. $\bar{\rho}_{\alpha_l} = \bar{\rho}_{\nu_l, \mu_l}(t_l)$. Furthermore,

$\mathcal{M}^{\alpha_l, \alpha_{l-1}} = (e^{\mathcal{L}_S \Delta t})_{\alpha_l, \alpha_{l-1}}$ is the free system propagator while the set of matrices $\{\mathcal{Q}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)}\}$ constitute the PT-MPO with inner bonds d_l , which carry information about the state of the environment [36].

Comparing Eq. (3) with Eq. (2) one can identify

$$\mathcal{Q}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)} = \begin{cases} (e^{\mathcal{L}_E \Delta t} (\mathbb{1} \otimes \rho^E(t_0)))_{(\alpha_l, d_l), (\alpha'_l, d_{l-1})}, & l = 1, \\ (e^{\mathcal{L}_E \Delta t})_{(\alpha_l, d_l), (\alpha'_l, d_{l-1})}, & 1 < l < n \\ (\text{Tr}_E e^{\mathcal{L}_E \Delta t})_{(\alpha_l, d_l), (\alpha'_l, d_{l-1})}, & l = n, \end{cases} \quad (4)$$

with fixed inner bonds at the boundaries $d_n = d_0 = 0$ (where we assume indexing from zero), while the remaining indices d_l run over a complete basis of the environment Liouville space.

Equation (3), which is illustrated in Fig. 1(g), constitutes a series of matrix multiplications, which propagate the total density matrix from one time step to the next. However, direct application of Eq. (3) is typically infeasible because the environment Liouville space is generally extremely large. This issue can be addressed by tensor network techniques: The set of matrices $\{\mathcal{Q}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)}\}$ [gray comb in Fig. 1(g)] has the form of an MPO with inner bonds d_l and outer bonds (α_l, α'_l) . For such a structure, compression techniques are available, where the inner bond dimensions are reduced while not affecting the action of the total MPO on the outer bonds. This can be done by sweeping across the MPO from one end to the other while performing SVDs and eliminating degrees of freedom corresponding to small singular values below a given threshold [28].

Because rank-reducing operations typically scale with the third power $\mathcal{O}(\chi^3)$ of the inner bond dimension χ , it is crucial to start with PT-MPOs with small dimensions and to keep the bond dimensions manageable at all times. The sequential ACE algorithm does this by introducing one environment mode at a time. This implies that this method applies only to environments that are composed of independent modes, where the environment Liouvillian $\mathcal{L}^E = \sum_{k=1}^{N_E} \mathcal{L}_E^{(k)}$ can be decomposed into N_E individual terms $\mathcal{L}_E^{(k)}$. These are assumed to only affect the system of interest as well as the k -th environment mode but they do not modify other degrees of freedom [37]. Similarly, the initial environment density matrix is assumed to factorize $\rho^E(t_0) = \rho^{(1)}(t_0) \otimes \rho^{(2)}(t_0) \otimes \dots \otimes \rho^{(N_E)}(t_0)$. Then, after employing a further Trotter decomposition of the environment propagator

$$e^{\mathcal{L}_E \Delta t} = e^{\mathcal{L}_E^{(1)} \Delta t} e^{\mathcal{L}_E^{(2)} \Delta t} \dots e^{\mathcal{L}_E^{(N_E)} \Delta t} + \mathcal{O}(\Delta t^2), \quad (5)$$

Eq. (3) becomes the two-dimensional tensor network depicted in Fig. 1(a), where each row corresponds to a PT-MPO for an individual environment mode k . If the Hilbert-space dimension of each individual mode is small

enough, the elements of these PT-MPOs can be directly constructed from Eq. (4). It remains to combine all the individual PT-MPOs [the rows in Fig. 1(a)] and to compress the inner bonds, such that one eventually obtains the final PT-MPO in Fig. 1(g), which then accounts for the influences of all environment modes.

B. PT-MPO contraction schemes

Having derived the ACE tensor network, we now discuss schemes to contract the corresponding PT-MPOs. Generally, combining two PT-MPOs (two rows in Fig. 1), whose matrices we denote by $\mathcal{Q}_{e_l, e_{l-1}}^{(\alpha_l, \alpha'_l)}$ and $\mathcal{P}_{f_l, f_{l-1}}^{(\alpha'_l, \alpha'_l)}$, respectively, yields a new PT-MPO with matrices

$$\tilde{\mathcal{Q}}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)} = \sum_{\substack{\alpha'_l \\ e_l, e_{l-1} \\ f_l, f_{l-1}}} \delta_{d_l, (e_l, f_l)} \mathcal{Q}_{e_l, e_{l-1}}^{(\alpha_l, \alpha'_l)} \mathcal{P}_{f_l, f_{l-1}}^{(\alpha'_l, \alpha'_l)} \delta_{d_{l-1}, (e_{l-1}, f_{l-1})}. \quad (6)$$

Here, the Kronecker deltas describe the formation of inner bonds of the combined PT-MPO as the product of inner bonds of the individual PT-MPOs.

To counteract the growth of the inner bonds, one compresses the resulting PT-MPO by sweeping along the MPO chain while performing SVDs. For example, sweeping in the forward direction (starting from $l = 1$), we perform an SVD of the matrix

$$A_{d_l, ((\alpha_l, \alpha'), d_{l-1})} = \tilde{\mathcal{Q}}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)} \longrightarrow \sum_k U_{d_l, k} \sigma_k V_{k, ((\alpha_l, \alpha'), d_{l-1})}^\dagger, \quad (7)$$

where U and V are matrices with orthonormal columns and σ_k are the non-negative singular values in descending order. Compression is achieved by keeping only singular values $\sigma_k \geq \epsilon \sigma_0$ above a certain threshold ϵ with respect to the largest singular value σ_0 . The PT-MPO matrices at steps l and $l + 1$ are updated as

$$\tilde{\mathcal{Q}}_{k, d_{l-1}}^{(\alpha_l, \alpha'_l)} \longleftarrow V_{k, ((\alpha_l, \alpha'), d_{l-1})}^\dagger, \quad (8a)$$

$$\tilde{\mathcal{Q}}_{d_{l+1}, k}^{(\alpha_{l+1}, \alpha'_{l+1})} \longleftarrow \sum_{d_l} \tilde{\mathcal{Q}}_{d_{l+1}, d_l}^{(\alpha_{l+1}, \alpha'_{l+1})} \sum_k U_{d_l, k} \sigma_k. \quad (8b)$$

The same procedure is repeated at the next step (increasing l). After a forward sweep, one performs an analogous sweep in the backward direction.

It should be noted that such a compression is only proven to be optimal if the MPO is of canonical form with orthogonality center at the site to be compressed [38]. This is generally not the case if the MPO is the result of combining two MPOs via Eq. (6). In principle, a sweep along the chain in one direction without truncation brings the MPO in such a form. However, in the spirit of the zip-up algorithm of Ref. [39], it is common practice in

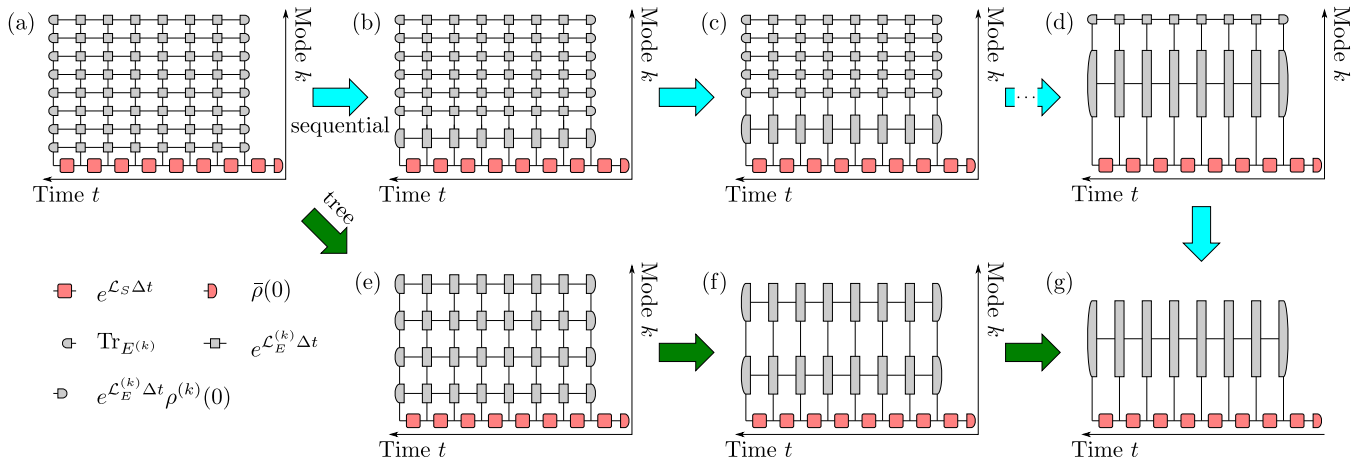


FIG. 1. Contraction of the ACE tensor network (a) using the sequential [(a),(b),(c),(d),(g)] and tree-like [(a),(e),(f),(g)] schemes. After combining two rows, the inner bonds are compressed by sweeps with truncated singular value decompositions. In the sequential scheme, PT-MPOs for individual environment modes are combined one by one, and so contributions with small inner bond dimensions are added to a PT-MPO with large inner bond dimension. In the tree-like scheme, neighboring PT-MPOs are combined, so most combination steps involve the combination of PT-MPOs with small to moderate inner bond dimension. Note that panel (a) represents Eq. (2) with $e^{\mathcal{L}_E \Delta t}$ further decomposed as in Eq. (5). Panel (g) represents Eq. (3) with the contracted gray blocks corresponding to PT-MPO matrices $\mathcal{Q}_{d_l, d_{l-1}}^{(\alpha_l, \alpha_{l-1}^i)}$.

PT-MPO algorithms [25] to perform truncation in every sweep because this reduces computation times while often providing close-to-optimal compression.

The central question we address in this article is the performance of different schemes to contract the PT-MPO for the network in Fig. 1(a). We consider the following variants:

1. Original (sequential) ACE

In the original ACE algorithm in Ref. [28], the tensor network is contracted by including the influence of one mode at a time sequentially into a growing PT-MPO. This sequence is depicted in Fig. 1 along panels (a), (b), (c), (d), and (g). After including each additional environment mode, a single forward and backward sweep is performed to compress the inner bonds with a fixed compression threshold ϵ .

The sequential order was chosen in Ref. [28] because of the scaling of SVDs with the third power of the matrix dimension. Together with the fact that combining two PT-MPO matrices as in Eq. (6) results in multiplication of the respective inner bond dimensions, brute-force calculations of SVDs become intractable if both matrices have large inner dimension (in practice the product of inner bonds each $\gtrsim 20$ prove challenging).

2. Preselected singular values

The problem of combining two PT-MPOs, both with moderate to large inner bond dimensions also arose in the

divide-and-conquer scheme for PT-MPOs with Gaussian (bosonic) reservoirs Ref. [30]. In that work we addressed this issue by ‘preselecting’ relevant parts of matrices before performing the most demanding SVDs.

This preselection is done by first performing independent forward sweeps of the two PT-MPOs to be combined. The corresponding singular values $\sigma_{e_l}^{(1)}$ and $\sigma_{f_l}^{(2)}$ are stored and used as indicators for the numerical relevance of the row e_l and f_l , respectively. The PT-MPOs are then combined during the subsequent backward sweep, where we keep only those indices d_l for which the products of singular values $\sigma_{e_l}^{(1)} \sigma_{f_l}^{(2)} \geq \epsilon \sigma_0^{(1)} \sigma_0^{(2)}$ exceed a threshold. This threshold is defined relative to the product of the maximal individual singular values (and likewise for d_{l-1}). This preselection step dramatically reduces the workload for the backward sweep SVD that follows the combination.

The combination of two moderately-sized PT-MPOs via the preselection approach enables new algorithms. However, this approach removes information from the MPO before compression, and so it can lead to a less effective compression than found for the zip-up approach alone. Thus, numerical tests are needed to investigate the practicality of preselecting degrees of freedom, which may depend on the specific example considered [30].

In this article we explore one variant of tensor network contraction in which we follow the same sequential contraction order as in the original ACE algorithm, yet combine pairs of PT-MPOs using preselection based on prior individual SVDs. This comparison lets us separate the effects of the MPO-MPO contraction from the effects of contraction order.

3. Tree-like contraction

The original sequential ACE contraction always only adds a small influence—resulting from one individual mode—at a time. However, this implies that it requires operations involving one PT-MPO with large inner dimensions for most of the combination and compression steps. For example, in the penultimate step in Fig. 1(d), the inclusion of the last mode, which by itself should have a small effect, requires operations on a PT-MPO with inner bonds of the size comparable to the full final PT-MPO. This can in principle be avoided by other contraction orders.

The preselection method now allows us to explore a different contraction order, where we first combine neighboring modes, assuming a sequence where these have similar influences (e.g., when modes are listed in order of increasing mode energies; see Appendix A for discussion of alternate orderings). The resulting PT-MPOs are then further combined, leading to a tree-like overall contraction order depicted in Fig. 1 as the sequence along panels (a), (e), (f), and (g). While both, sequential ACE and the tree-like contraction, require $N_E - 1$ combinations and compressions, the distribution of inner bond dimensions is drastically different. In particular, in the tree-like scheme, the first $N_E/2 - 1$ contractions involve combinations with the smallest possible inner bond dimensions. If neighboring modes have similar physical influences, the combined PT-MPOs should be compressible to small dimensions, which also makes the next layer of $N_E/4 - 1$ combinations and compressions efficient. Only the last few combination and compression steps involve PT-MPOs with large inner bond dimensions. This is the key argument why the tree-like contraction scheme can be expected to outperform contraction in the sequential order.

Finally, it is noteworthy that the tree-like scheme is in principle parallelizable except for the last step. As our goal is to compare the algorithms themselves, we do not make use of this in our present study.

4. Fine-tuning of MPO compression

The preselection method can in general result in sub-optimal MPO compression [30]. We therefore explore two additional strategies to mitigate this. The first strategy is based on the following observation: Even when MPOs are locally optimally compressed by sweeps with truncated SVDs starting from MPOs in canonical form, truncations further into the chain depend on earlier truncations closer to the starting edge. Hence, better compression may be achieved by sweeping multiple times $n_{\text{SW}} > 1$ after combining two PT-MPOs. In many-body quantum physics, this is often done using iterative, e.g., variational, methods [38], which are faster than repeated SVD sweeps [40]. Motivated by this, we check for potential improvements by sweeping multiple times $n_{\text{SW}} > 1$ with truncating

SVDs after combining two PT-MPOs.

The second strategy is to reduce the accumulation of rounding errors by choosing scale-dependent compression, where different compression thresholds are used for different layers of the tree-like scheme. Concretely, we start with a compression threshold ϵ_{min} for the first layer and exponentially increase the threshold for each successive layer, such that the last compression is done with the maximal threshold ϵ_{max} . This corresponds to equidistant sampling of $\ln \epsilon$ in the range $[\ln \epsilon_{\text{min}}, \ln \epsilon_{\text{max}}]$. This way, excess degrees of freedom resulting from suboptimal compression at earlier stages of the algorithm have a chance to be eliminated in later steps, while computation time is increased mainly at the first few steps, which are the fastest steps in the tree-like scheme. We parameterize the scale-dependent compression by the threshold range factor $r = \epsilon_{\text{max}}/\epsilon_{\text{min}}$ and identify $\epsilon_{\text{max}} = \epsilon$ with the nominal compression threshold ϵ , e.g., for comparison with sequential ACE simulations.

C. Reducing Trotter errors

For clarity, we have so far described the ACE tensor network assuming a first-order (asymmetric) Trotter split between system and environment as well as between different environment modes. However, the Trotter error can be reduced [41] from $\mathcal{O}(\Delta t^2)$ to $\mathcal{O}(\Delta t^3)$ using a second-order (symmetric) Trotter decomposition of the form

$$e^{(A+B)\Delta t} = e^{A\Delta t/2} e^{B\Delta t} e^{A\Delta t/2} + \mathcal{O}(\Delta t^3). \quad (9)$$

The system-environment split can be straightforwardly brought into symmetric form without changing the PT-MPO by calculating system propagators over half time steps $e^{\mathcal{L}_S \Delta t/2}$.

The Trotter decomposition between different environment modes in Eq. (5) can also be easily implemented using the symmetric split in the sequential contraction scheme, where the PT-MPO is expanded by including the environment propagator for a single environment mode. There, one can replace Eq. (6) by a symmetric version where half-step propagators $e^{\mathcal{L}_E^{(k)} \Delta t/2}$ are multiplied from the left as well as from the right to the current PT-MPO (see Ref. [28] for details).

In contrast, in the tree-like scheme, implementing the symmetric Trotter splitting is more challenging because the time steps of PT-MPOs in intermediate layers of the algorithm are fixed. We solve this issue by changing the order of combination [36]: At odd time steps (starting from $l = 1$), we combine PT-MPO matrices in the same order as in Eq. (6), whereas at even time steps, we change

the order (with respect to the outer indices α) to

$$\tilde{\mathcal{Q}}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)} = \sum_{\substack{\alpha'_l \\ e_l, e_{l-1} \\ f_l, f_{l-1}}} \delta_{d_l, (e_l, f_l)} \mathcal{P}_{f_l, f_{l-1}}^{(\alpha_l, \alpha'_l)} \mathcal{Q}_{e_l, e_{l-1}}^{(\alpha'_l, \alpha'_l)} \delta_{d_{l-1}, (e_{l-1}, f_{l-1})}. \quad (10)$$

Then, the propagation over two steps effectively amounts to a symmetric Trotter splitting, where the lower-order Trotter error introduced by an odd time step is canceled by the corresponding term in the subsequent even step.

III. RESULTS

A. Model systems

We test the performance of the different variants for contracting the ACE tensor network on two examples.

First, we consider a two-level system coupled to a bath of harmonic oscillators, i.e. the paradigmatic spin-boson model. For this model, very efficient alternative PT-MPO construction algorithms are available [25, 30, 31], which can be used as a reference. These make use of the Gaussian statistics of bath correlations, which facilitates the PT-MPO construction using a different tensor network derived by path integration [34], which in turn is based on previous path integration approaches [42, 43]. Concretely, the environment Hamiltonian is given by $H_E = \sum_k H_E^{(k)}$ with

$$H_E^{(k)} = \hbar\omega_k b_k^\dagger b_k + \hbar g_k (b_k^\dagger + b_k) |e\rangle\langle e| + \hbar \frac{g_k^2}{\omega_k} |e\rangle\langle e|, \quad (11)$$

where b_k^\dagger and b_k are bosonic creation and annihilation operators, $\hbar\omega_k$ is the energy of the k -th environment mode, g_k are the coupling constants and $|e\rangle\langle e|$ is the projection onto the excited state of the two-level spin system. The last term is added for convenience to cancel to reorganization energy or polaron shift.

We choose environment parameters g_k, ω_k to reproduce a particular form of the spectral density $J(\omega) = \sum_k |g_k|^2 \delta(\omega - \omega_k)$. Specifically, we choose parameters such that $J(\omega)$ takes the form of a well established model [6] for a quantum dot coupled to longitudinal acoustic phonons

$$J(\omega) = \omega^3 \left(c_e e^{-\omega^2/\omega_e^2} - c_h e^{-\omega^2/\omega_h^2} \right)^2. \quad (12)$$

For a GaAs quantum dot with radius 4 nm relevant parameters are $c_e = 0.1271 \text{ ps}^{-1}$, $c_h = -0.0635 \text{ ps}^{-1}$, $\omega_e = 2.555 \text{ ps}^{-1}$, and $\omega_h = 2.938 \text{ ps}^{-1}$.

For the ACE tensor network, we discretize the frequency range $\omega = [0, \omega_{\max}]$ uniformly, with $\omega_{\max} = 7 \text{ ps}^{-1}$ and using $N_E = 64$ modes. Thus, the mode frequencies are then $\omega_k = (k - \frac{1}{2})\omega_{\max}/N_E$ and the coupling strengths are $g_k = \sqrt{J(\omega_k)\omega_{\max}/N_E}$ for $k = 1, \dots, N_E$.

The initial state of each mode is given by a thermal distribution with temperature $T = 4 \text{ K}$. We check convergence with respect to the Hilbert-space dimension per mode, and find $M = 4$ (up to three bosonic excitations) is sufficient. A total propagation time of $t_e = 20 \text{ ps}$ is discretized with steps of size $\Delta t = 0.1 \text{ ps}$ or $\Delta t = 0.01 \text{ ps}$. The system Hamiltonian H_S does not enter the PT-MPO calculation algorithm.

The second example we consider is a fermionic resonant level model. This is relevant, e.g., for impurity problems [33]. The environment Hamiltonians are given by

$$H_E^{(k)} = \hbar\omega_k c_k^\dagger c_k + \hbar g_k (c_k^\dagger c_0 + c_0^\dagger c_k), \quad (13)$$

where again $\hbar\omega_k$ and g_k are mode energies and coupling constants, respectively. The operators c_k^\dagger and c_k are fermionic creation and annihilation operators, where the mode $k = 0$ is the system mode while modes $k > 0$ constitute the environment.

Here, we model the hopping of electrons from metallic leads to a central site. If the band width of the environment density of states is finite, the dynamics is in general non-Markovian. For concreteness, we assume that the environment density of states is centered around the system mode energy, which we shift to zero—such a shift has no effects for particle-number-conserving models. We thus discretize the (dimensionless) energy range $\omega \in [-32, 32]$ uniformly with $N_E = 128$ modes. The coupling constants $g_k = \sqrt{64/(2\pi N_E)}$ are chosen such that the hopping rate in the Markov limit is $\Gamma = 1$. All environment modes are initially in the occupied state, which describes a situation where the central site is energetically far below the bath's Fermi level. A time grid spanning the time $t_e = 5$ in steps of size $\Delta t = 0.1$ is used for the calculations.

It should be noted that the derivation of ACE assumes that operators corresponding to different environment modes commute. By contrast, fermionic operators anticommute. In Appendix B, we derive how one can nevertheless obtain PT-MPOs properly accounting for the fermionic anticommutation relations using ACE. This only requires a slight modification of the input to the algorithm. This strategy can be applied irrespective of whether the contraction order is sequential or tree-like.

B. Computation times

In Fig. 2 we show the computation times needed to calculate PT-MPOs for the spin-boson model with larger ($\Delta t = 0.1 \text{ ps}$, panel a) and smaller ($\Delta t = 0.01 \text{ ps}$, panel b) time steps, as well as for the fermionic model (panel c). The different variants of contracting the ACE tensor network described in the theory section are used. The computation times are extracted as the total elapsed time on a conventional laptop computer with AMD Ryzen 7 5825U processor.

In both models, we obtain similar trends, which remain robust over the range of thresholds shown in Fig. 2: se-

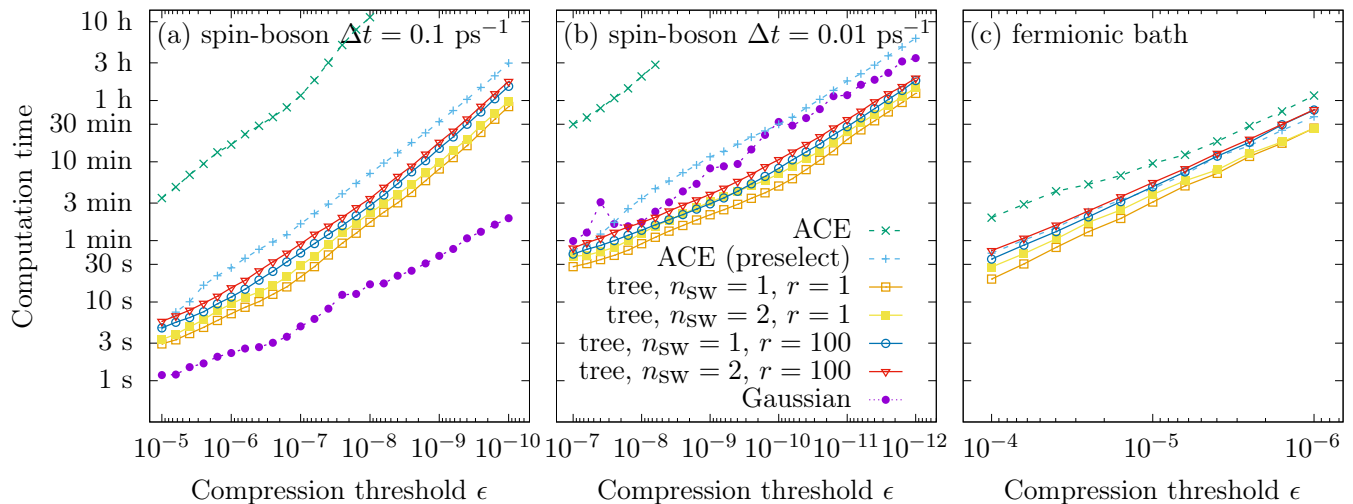


FIG. 2. Computation times of PT-MPO contraction for a spin-boson model with larger (a) and smaller (b) time steps, and for a fermionic environment (c) using different methods: The original sequential ACE algorithm [28] (green), the sequential algorithm using the preselection approach for PT-MPO combination and compression (light blue), and the tree-like scheme using $n_{SW} = 1$ or 2 SVD sweeps per combination and a scale-dependent compression threshold with range factor $r = 1$ and $r = 100$, respectively. For the spin-boson model, we compare with the algorithm of Jørgensen and Pollock in Ref. [25] which assumes Gaussian environment correlations (purple). Note that the scale for both computation times and compression thresholds ϵ is logarithmic. Note also that thresholds ϵ decrease from left to right on the axis, so that accuracy improves from left to right.

quential ACE simulations using the original algorithm of Ref. [28] require by far the longest computation times and become impractical when high accuracies (small thresholds ϵ) are needed. Using the sequential algorithm but with the preselection approach for combining rows of the tensor network already leads to massive savings in computation times. In the spin-boson example in Fig. 2(a) the computation time at $\epsilon = 10^{-7}$ is reduced by a factor of about 40, whereas we find a speed-up by a factor 2 in the fermionic model in Fig. 2(c) at $\epsilon = 10^{-5}$. This is due to the fact that the larger Liouville space per mode in the spin-boson model ($M^2 = 16$ in the spin-boson model versus $M^2 = 4$ for fermions) has more potential for reduction by a prior forward sweep before combination.

Switching to the tree-like scheme, we find a further reduction of computation time with the fastest simulations being those without any fine-tuning of convergence parameters, i.e., sweeping only once ($n_{SW} = 1$) per combination and using a fixed compression threshold $r = 1$. For the spin-boson example, the advantage over the original ACE algorithm is roughly two orders of magnitude. Compared to the method by Jørgensen and Pollock [25]—which is restricted to Gaussian environments, hence labeled as ‘Gaussian’ in Fig. 2—the more general tree-like ACE scheme requires longer computation times for those simulations with larger time steps [Fig. 2(a)]. However, when time steps are smaller [Fig. 2(b)], sequential ACE is faster than the specialized Jørgensen and Pollock algorithm. This is because the latter scales quadratically with the number of time steps, while the tree-like scheme is linear in n [44].

C. Numerical error

The numerical error due to MPO compression is analyzed in Fig. 3(a). To do this, we simulate the dynamics of a two-level system driven with a system Hamiltonian $H_S = (1 \text{ ps}^{-1}) \frac{\hbar}{2} \sigma_x$ and coupled to the bosonic environment described by the PT-MPOs calculated in Fig. 2(a). We define the compression error as the maximal (over all time steps) difference between occupations $n_e(t_j) = \langle (|e\rangle\langle e|) \rangle_{t_j}$ from simulations with compression threshold ϵ and those from simulations with the smallest threshold in each data series ($\epsilon = 10^{-8}$ for the sequential ACE without preselection and $\epsilon = 10^{-10}$ for all other methods).

In Fig. 3(a), it can be seen that, for the same compression threshold ϵ , the tree-like contraction scheme results in larger errors than the Gaussian algorithm unless a scale-dependent compression threshold ($r = 100$) is used. The fact that a similarly large error is found in the sequential ACE scheme with preselection indicates that this error is due to the preselection rather than the contraction order.

Especially relevant for practical applications is the trade-off between compression error and computation time for each of the methods considered here, which is depicted in Fig. 3(b), (c), and (d) for the same simulations as in Fig. 2(a), (b), and (c), respectively. It is clearly seen that, first of all, the tree-like contraction outperforms the sequential versions of ACE with and without preselection, and second, using a scale-dependent compression threshold leads to improved accuracy for the

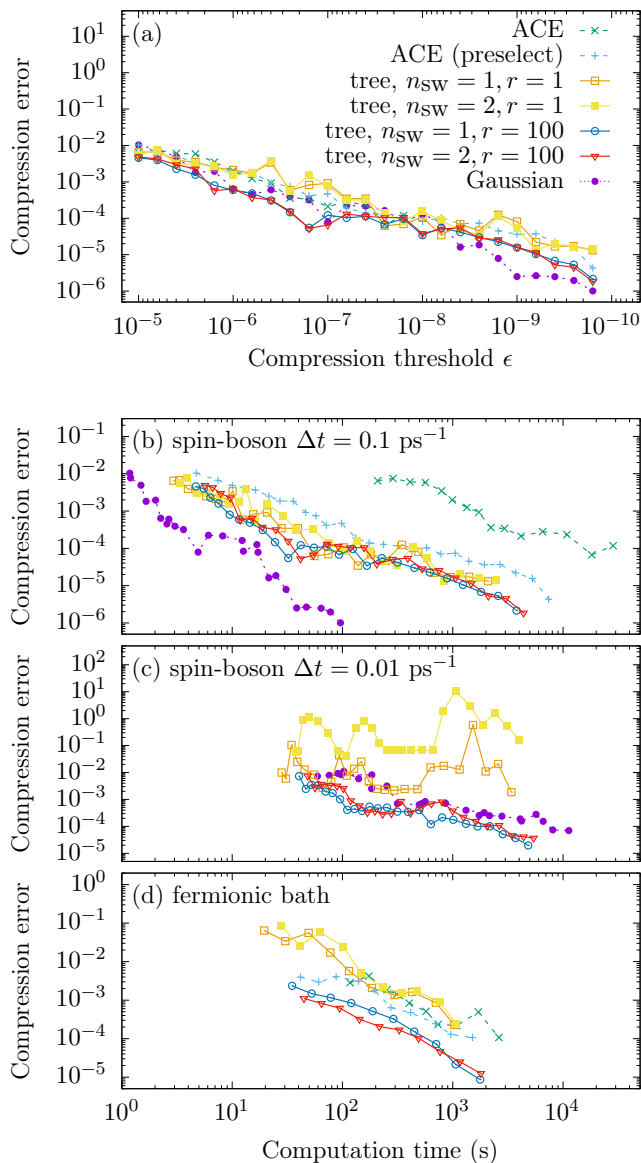


FIG. 3. (a): Compression error (see main text) as a function of the compression threshold for the same simulations as in Fig. 2(a). (b), (c), and (d): Compression error versus computation time for the same simulations as in Fig. 2(a), (b), and (c), respectively.

same computation time in the tree-like scheme. Moreover, the Gaussian algorithm outperforms all ACE variants when the time steps are large [Fig. 3(b)], but the tree-like ACE scheme with scale-dependent compression threshold becomes more efficient than the Gaussian algorithm for small time steps [Fig. 3(c)].

It should be noted that very small time steps Δt can introduce numerical artifacts. This is because in environment propagators $e^{\mathcal{L}_E \Delta t} \approx \mathbb{1} + \mathcal{L}_E \Delta t + \dots$ the information from \mathcal{L}_E is scaled by Δt , which results in numerical values much smaller than the unity elements of the identity matrix and thus to numerical cancellation ef-

fects. Moreover, for the same overall propagation length, calculations with small time steps require many more individual combination and compression steps, and so numerical errors can accumulate. This can lead to unstable, nonconverging behavior of the tree-like scheme when error accumulation is not mitigated by a scale-dependent compression threshold [see Fig. 3(c)]. Additionally, due to the long computation times, we cannot extract reliable estimates for the compression error using sequential ACE without preselection. With preselection, the sequential ACE scheme suffers even worse numerical instabilities than the tree-like scheme, leading to unphysical results with occupations greater than unity for almost all thresholds ϵ considered [not shown in Fig. 3(c)]. We have checked that also the sequential ACE with preselection produces reliable results when an scale-dependent compression threshold $r = 100$ is used. The accuracy per computation time for the fermionic bath shown in Fig. 3(d) behaves qualitative similar to that of the spin-boson model and also benefits from the use of a scale-dependent compression threshold.

In Appendix C, we analyze the same situation as in Fig. 3(a) and (b) but determining the convergence error via a distance measure of MPOs. This measure is independent of the choice of system Hamiltonian and yields smoother convergence curves, but numerical cancellation prevents its use very close to convergence.

D. Compression efficacy

The final performance measure we discuss is how effective the reduction of environment degrees of freedom is, which can be measured by the inner bond dimension χ of the PT-MPO for a given compression threshold ϵ . To study this we analyze the distribution of singular values of the matrix $\mathcal{Q}_{d_l, d_{l-1}}^{(\alpha_l, \alpha'_l)}$ in the middle ($l = n/2$) of the final PT-MPOs that result from different contraction algorithms. This is shown in Fig. 4 for the spin-boson model with large time steps $\Delta t = 0.1 \text{ ps}$ at compression threshold $\epsilon = 10^{-7}$, for the spin-boson model with small time steps $\Delta t = 0.01 \text{ ps}$ at threshold $\epsilon = 10^{-9}$, and for the fermionic environment at threshold $\epsilon = 10^{-5}$.

As a reference, in Fig. 4(a) we also show singular values obtained by the Gaussian algorithm of Ref. [25], which typically results in close-to-optimal compression [31]. Here, we obtain 26 singular values above the threshold. In contrast, the original sequential ACE algorithm results in 41 singular values, where the largest ones agree with those from the Gaussian algorithm. Using ACE with preselection yields 60 singular values and the tree-like contraction algorithm without fine-tuning produces 124. Singular values obtained from tree-like contraction deviate visibly from the close-to-optimal ones already from the fourth value (index $i = 3$), and the tree-like algorithm produces many singular values lying in a long tail within about one order of magnitude of the truncation threshold.

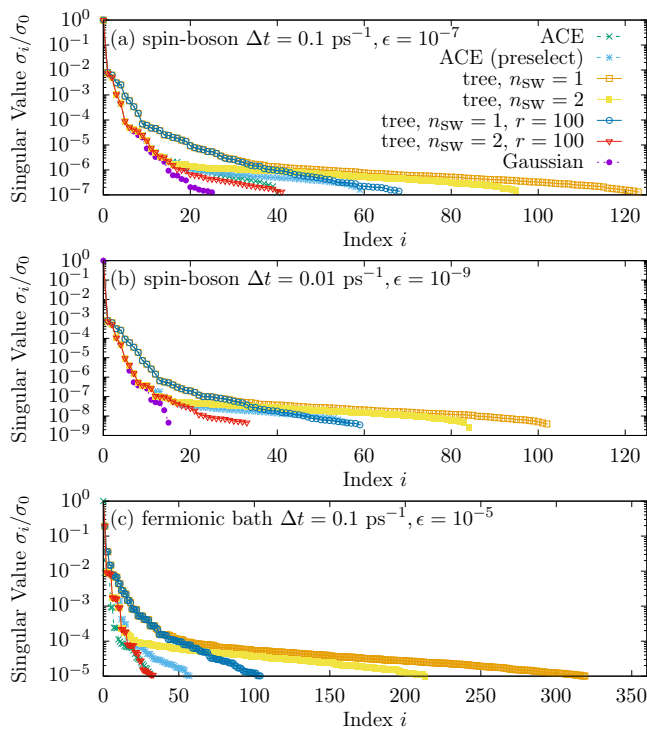


FIG. 4. Normalized singular values at the center of the final PT-MPO for the spin-boson model with large (a) and small (b) time steps, and (c) for the fermionic environment, corresponding to calculations in Fig. 2(a), (b), and (c), respectively.

As can be seen in Fig. 4(a), the larger singular values can be brought close to the values of the Gaussian method by sweeping multiple times $n_{\text{SW}} > 1$ per PT-MPO combination. Furthermore, by comparing singular values from the tree-like scheme without and with scale-dependent thresholds ($r = 100$), we find that the latter strongly suppresses the long tail of small singular values. Along with the results on the numerical error, this suggests that the small singular values are spurious and a result of error accumulation. However, the scale-dependent threshold does not affect the large singular values, which still do not match the close-to-optimal values. This shows that both strategies for fine-tuning compression parameters address different aspects of suboptimal PT-MPO compression in the tree-like contraction scheme. Indeed, bringing both strategies together—i.e. sweeping twice and dynamically increasing the threshold—yields a similar number of singular values as the original ACE algorithm. As shown in Fig. 2(a), the computation times are however still reduced by one to two orders of magnitude compared to sequential ACE. Similar results are found for both the spin-boson model with small time steps in Fig. 4(b) and the fermionic environment in Fig. 4(c).

Finally, Fig. 5 shows in more detail how inner bond dimensions and computation times depend on the fine-tuning parameters, i.e. the number of sweeps n_{SW} and the threshold range factor r , on the example of the

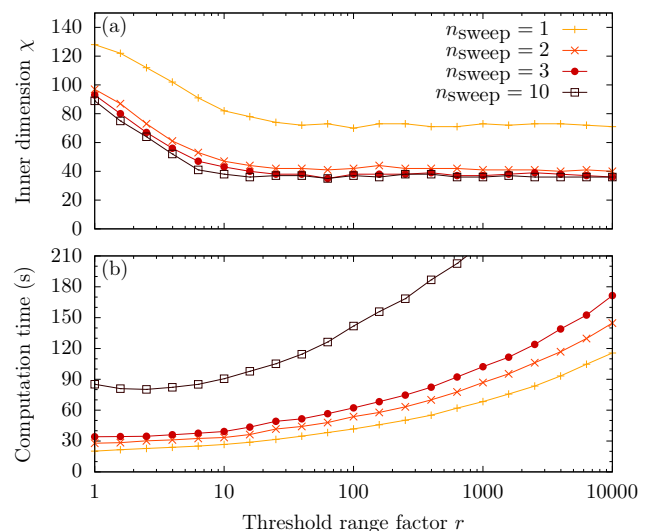


FIG. 5. Inner dimension χ (a) and computation time (b) for different numbers of MPO compression sweeps n_{SW} per combination as a function of the threshold range factor $r = \epsilon_{\text{max}}/\epsilon_{\text{min}}$ for the scale-dependent compression threshold. Simulations were done for the same parameters as in Fig. 2(a) with $\epsilon_{\text{min}} = \epsilon = 10^{-7}$.

spin-boson model from Fig. 2(a). We find that sweeping twice per PT-MPO combination significantly reduces bond dimensions, but increasing the number of sweeps n_{SW} further has only marginal effect on the compression. Similarly, increasing the factor r controlling the scale-dependent compression threshold initially reduces the bond dimensions, but these then saturate when $r \approx 10$. The computation time increases with n_{SW} and with r , yet the increase with r is relatively slow. Thus, for a given environment, experimenting with the large values of r to find a good trade-off between accuracy, compression, and computation time is manageable.

IV. DISCUSSION

We have introduced and analyzed a novel process tensor contraction scheme, with which the ACE algorithm [28] for the simulation of open quantum systems can be considerably improved. The new, tree-like contraction order requires a way to combine and compress PT-MPOs with large inner dimensions, for which we use the preselection approach described in Ref. [30]. Our scheme achieves about two orders of magnitude reduction in computation time compared to the original sequential ACE algorithm. The reason for this speed-up is the typical distribution of inner bond dimensions in different steps of the algorithm: most steps in the sequential scheme involve contracting an object with a small bond dimension with a PT-MPO with a large inner dimension; in contrast the tree-like scheme combines influences of similar dimensions, which are on average

smaller than those of the final PT-MPO. However, the compression—as defined by the the reduction of bond dimension and numerical error—is suboptimal for the tree-like algorithm, unless the compression parameters are fine-tuned. We have explored multiple line sweeps with SVDs after combining two PT-MPOs as well as the use of a scale-dependent compression threshold; we indeed find that these strategies give a significant reduction of inner bond dimensions and of numerical errors.

Even though the exact value of optimal fine-tuning parameters may depend on the concrete physical problem, the general trends found in our examples suggest the following rules for choosing them: As calculations without fine-tuning ($n_{\text{SW}} = 1, r = 1$) are fastest, one should first perform such calculations, from which one can also extract and analyze the distribution of singular values in the final PT-MPO. If these show a long tail of small singular values near the truncation threshold, it is advisable to increase r in steps corresponding to orders of magnitude (i.e., $r = 10, 100, 1000, \dots$) until no significant reduction in the number of singular values in the final PT-MPO is found. Sweeping multiple times $n_{\text{SW}} > 1$ per PT-MPO combination is often not necessary if one is only interested in numerical accuracy and computation time. However, sweeping twice $n_{\text{SW}} = 2$ can help if close-to-optimal MPO compression is sought. Close-to-optimal compression is important in several cases; these include: when the simulation complexity encoded in the inner bond dimension χ is of interest, when the PT-MPO is reused for many calculations, or when quantum systems with multiple environments described by multiple PT-MPOs will be simulated. In any case, sweeping more than twice is typically not profitable.

Summarizing, we find that the tree-like contraction scheme clearly outperforms the ACE algorithm with sequential contraction over all test cases considered, yet it remains generally applicable. Thus, we expect that our new scheme also significantly accelerates investigations of complex open quantum systems including those coupled to non-Gaussian environments, or those which are manipulated with time-dependent driving and are lossy. Moreover, all insights gained here relate to the combination of multiple PT-MPOs in general. This is also important in different setting, e.g., for combining two PT-MPOs describing different environments [19] and their non-additive cross effects [29, 45]. Furthermore, the algorithms described here can be useful to parallelize PT-MPO schemes, e.g., where different parts of the ACE tensor network are contracted on different computing nodes before they are combined to form the final PT-MPO.

ACKNOWLEDGMENTS

M.C. acknowledges funding by the Return Program of the State of North Rhine-Westphalia. M.C. and E.M.G. acknowledge funding from EPSRC grant no. EP/T01377X/1. B.W.L. and J.K. acknowledge funding

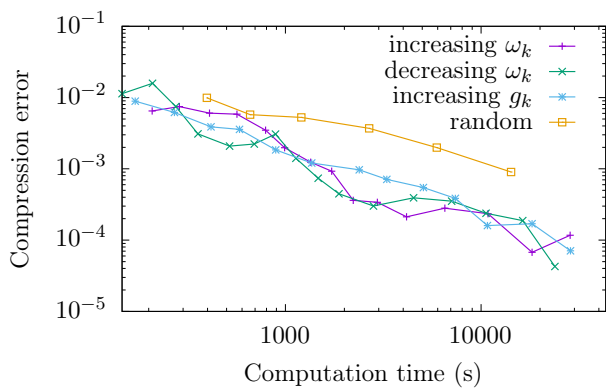


FIG. 6. Compression error vs. computation time for sequential ACE simulations of the same spin-boson model as in Fig. 3(b), where the environment modes are sorted in order of increasing frequency, decreasing frequency, and increasing coupling strength, respectively. For comparison we also show results found from a single random ordering of modes.

from EPSRC grant no. EP/T014032/1.

Appendix A: Impact of sequential ordering of modes

In the main text, we discuss the improvement of the ACE algorithm by switching from a sequential to a tree-like order of contractions. The question naturally arises whether another aspect of contraction order, namely reordering the individual modes within the sequential scheme, has the potential for improving the efficiency of the ACE algorithm.

To investigate this question, we perform convergence studies of the sequential ACE algorithm for the spin-boson model discussed in Fig. 3(a) and (b), where we sort the environment modes in order of increasing mode frequency ω_k , decreasing mode frequency ω_k , and increasing coupling strength g_k . Additionally, we compare to the situation with randomly enumerated modes. The results are shown in Fig. 6 and indicate no consistent preference for choosing any particular of the ordered sequences. However, the results do indicate that randomly ordering the modes results in much longer computation times for a given compression error. This observation can be understood by noting that random ordering inhibits efficient compression of modes with similar influences at intermediate steps of the algorithm. For definiteness, in the simulations presented in the main text of the article, we choose to label the modes in increasing order in frequency, i.e. $\omega_{k_1} < \omega_{k_2}$ for $k_1 < k_2$.

Appendix B: Fermionic PT-MPOs

In the main text, we consider the fermionic open quantum system with total Hamiltonian $H_E = \sum_{k=1}^{N_E} H_E^{(k)}$,

where

$$H_E^{(k)} = \hbar\omega_k c_k^\dagger c_k + \hbar g_k (c_k^\dagger c_0 + c_0^\dagger c_k). \quad (\text{B1})$$

Here, c_j^\dagger and c_j are fermionic creation and annihilation operators, respectively, which obey the canonical anti-commutator relations $\{c_i, c_j^\dagger\} = \delta_{ij}$ and $\{c_i, c_j\} = 0$. We can compare this to the spin-like system with two-level environment modes described by

$$\tilde{H}_E^{(k)} = \hbar\omega_k \sigma_k^+ \sigma_k^- + \hbar g_k (\sigma_k^+ \sigma_0^- + \sigma_0^+ \sigma_k^-), \quad (\text{B2})$$

where σ_j^+ and σ_j^- are two-level (spin-1/2) ladder operators. These also satisfy $\{\sigma_j^-, \sigma_j^+\} = 1$ for a single mode j , but spin ladder operators for different modes commute $[\sigma_i^-, \sigma_{j \neq i}^\pm] = 0$. That is, no sign change occurs when changing the order of spin ladder operators, while it does for fermions.

While the spin-like system is directly in a form that allows a numerically exact treatment using ACE, this is not straightforward for a proper fermionic model, where anticommutation relations have to be enforced. Here, we show how the fermionic model can be mapped to a spin-like model, which turns out to be a slight variation of Eq. (B2). This provides an alternative approach to recent derivations of fermionic tensor network techniques using Grassmann variables [32, 33, 46].

We first employ the Jordan-Wigner transformation to replace fermionic operators by spin operators

$$c_j^\dagger = P^{[0, j-1]} \sigma_j^+, \quad (\text{B3a})$$

$$c_j = P^{[0, j-1]} \sigma_j^-, \quad (\text{B3b})$$

where

$$P^{[j_1, j_2]} = \prod_{j=j_1}^{j_2} (-\sigma_j^z) \quad (\text{B4})$$

will be referred to as the parity operator for modes j_1 to j_2 . It is straightforward to show that the respective sets of fermionic and spin operators obey the corresponding (anti)commutation relations. Applying this transformation to the fermionic environment Hamiltonian, we obtain

$$H_E^{(k)} = \hbar\omega_k \sigma_k^+ \sigma_k^- + \hbar g_k P^{[1, k-1]} (\sigma_k^+ \sigma_0^- + \sigma_0^+ \sigma_k^-). \quad (\text{B5})$$

In contrast to the spin system defined by $\tilde{H}_E^{(k)}$, here, the parity operator $P^{[1, k-1]}$ introduces a non-local coupling between different environment modes. However, progress can be made by realizing that both Hamiltonians are connected by

$$H_E^{(k)} = (P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-} \tilde{H}_E^{(k)} (P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-}, \quad (\text{B6})$$

which can be seen by noting that $(P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-}$ commutes with σ_k^\pm and $(P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-} \sigma_0^\pm (P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-} = P^{[1, k-1]} \sigma_0^\pm$. Because $[(P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-}]^2 = 1$, we may see that $(P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-}$ is a unitary transform. Hence, we can relate the time evolution operators for both models by

$$e^{\pm \frac{i}{\hbar} H_E^{(k)} \Delta t} = (P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-} e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(k)} \Delta t} (P^{[1, k-1]})^{\sigma_0^+ \sigma_0^-}. \quad (\text{B7})$$

While the fermionic and spin Hamiltonians for each mode k are related by a unitary transform, this transform is different for each k , so the total fermionic and spin Hamiltonians are not unitarily equivalent. However, we can nonetheless combine the time evolution operators for the different environment modes as follows. The standard Trotter decomposition gives

$$\begin{aligned} e^{\pm \frac{i}{\hbar} \sum_k H_E^{(k)} \Delta t} &\approx e^{\pm \frac{i}{\hbar} H_E^{(NE)} \Delta t} \dots e^{\pm \frac{i}{\hbar} H_E^{(3)} \Delta t} e^{\pm \frac{i}{\hbar} H_E^{(2)} \Delta t} e^{\pm \frac{i}{\hbar} H_E^{(1)} \Delta t} \\ &= \dots (P^{[1, 2]})^{\sigma_0^+ \sigma_0^-} e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(3)} \Delta t} (P^{[1, 2]})^{\sigma_0^+ \sigma_0^-} (P^{[1, 1]})^{\sigma_0^+ \sigma_0^-} e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(2)} \Delta t} (P^{[1, 1]})^{\sigma_0^+ \sigma_0^-} e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(1)} \Delta t} \\ &= (P^{[1, NE]})^{\sigma_0^+ \sigma_0^-} \left[(P^{[NE, NE]})^{\sigma_0^+ \sigma_0^-} e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(NE)} \Delta t} \right] \dots \left[(P^{[1, 1]})^{\sigma_0^+ \sigma_0^-} e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(1)} \Delta t} \right]. \end{aligned} \quad (\text{B8})$$

For the step from the second to the last line, we use the fact that $P^{[1, j+1]} P^{[1, j]} = P^{[j+1, j+1]}$. This results in cancellation of most of the non-local parity terms, so that the full time evolution almost completely decou-

ples into a product of ‘‘local’’ terms, each acting only on a single environment mode as well as the system mode. Each of these local terms—which are collected into square brackets—have the form of the time evolution

for a single mode spin-like system, $e^{\pm \frac{i}{\hbar} \tilde{H}_E^{(k)} \Delta t}$ multiplied by $(P^{[k,k]})^{\sigma_0^+ \sigma_0^-}$, the local parity operator. This parity operator is -1 if both environment and system mode are simultaneously occupied and 1 otherwise.

The only remaining non-local term is the parity operator of the full environment $(P^{[1, N_E]})^{\sigma_0^+ \sigma_0^-}$. This term can however be eliminated using a Trotter decomposition with alternating order over two time steps. That is, for even time steps n (starting from $n = 0$), we use the Trotter decomposition

$$e^{\pm \frac{i}{\hbar} (H_E + H_S) \Delta t} \approx e^{\pm \frac{i}{\hbar} H_E^{(N_E)} \Delta t} \dots e^{\pm \frac{i}{\hbar} H_E^{(1)} \Delta t} e^{\pm \frac{i}{\hbar} H_S \Delta t}, \quad (\text{B9})$$

as above, while for odd time steps n we use

$$e^{\pm \frac{i}{\hbar} (H_E + H_S) \Delta t} \approx e^{\pm \frac{i}{\hbar} H_S \Delta t} e^{\pm \frac{i}{\hbar} H_E^{(1)} \Delta t} \dots e^{\pm \frac{i}{\hbar} H_E^{(N_E)} \Delta t}. \quad (\text{B10})$$

Then, the factor $(P^{[1, N_E]})^{\sigma_0^+ \sigma_0^-}$ appearing on the left of the decomposition for even n (see Eq. (B8)) exactly cancels with the corresponding factor appearing on the right of the subsequent (odd n) time step. A side effect of the alternating order of Trotter decompositions is that it also cancels the leading-order Trotter error [36]. The key point of this approach is that the time evolution over two time steps now completely decouples into terms involving only single environment modes. This makes the fermionic problem amenable to a treatment with the ACE algorithm.

Having established the mapping above, we conclude this appendix by discussing the practical implementation of the resulting fermionic PT-MPOs. For the construction of the fermionic PT-MPOs, it should be noted that the PT-MPO matrices can be understood as the environment propagator acting on the full environment Liouville space compressed to the most relevant subspace of environment excitation, where the lossy compression matrices and their pseudoinverses are denoted by \mathcal{T} and \mathcal{T}^{-1} , respectively. These matrices are determined implicitly by the combination and rank-reducing compression [36] in the sequential or the tree-like ACE algorithm. For the spin-like system, using the superoperator formalism of Ref. [36], the PT-MPO take the form

$$\begin{aligned} \tilde{\mathcal{Q}} &= \mathcal{T} e^{\tilde{\mathcal{L}}_E \Delta t} \mathcal{T}^{-1} = \mathcal{T} \left[e^{-\frac{i}{\hbar} \tilde{H}_E \Delta t} \otimes (e^{\frac{i}{\hbar} \tilde{H}_E \Delta t})^T \right] \mathcal{T}^{-1} \\ &= \mathcal{T} \left[\prod_{k=1}^{N_E} e^{-\frac{i}{\hbar} \tilde{H}_E^{(k)} \Delta t} \otimes (e^{\frac{i}{\hbar} \tilde{H}_E^{(k)} \Delta t})^T \right] \mathcal{T}^{-1}, \end{aligned} \quad (\text{B11})$$

where $(\dots)^T$ indicates transpose. For the fermionic system, using the Jordan-Wigner transform, the PT-MPO matrices instead take the form:

$$\begin{aligned} \mathcal{Q} &= \mathcal{T} \left[\prod_{k=1}^{N_E} \left((P^{[k,k]})^{\sigma_0^+ \sigma_0^-} e^{-\frac{i}{\hbar} \tilde{H}_E^{(k)} \Delta t} \right) \right. \\ &\quad \left. \otimes \left(e^{\frac{i}{\hbar} \tilde{H}_E^{(k)} \Delta t} (P^{[k,k]})^{\sigma_0^+ \sigma_0^-} \right)^T \right] \mathcal{T}^{-1}. \end{aligned} \quad (\text{B12})$$

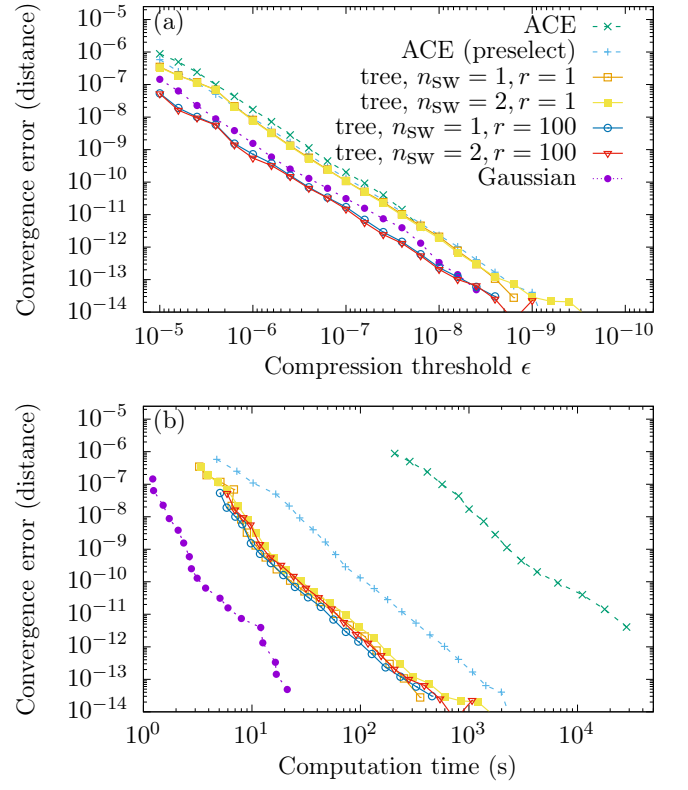


FIG. 7. Convergence error in terms of the tensor distance as a function of (a) compression threshold ϵ and (b) computation time. Parameters are chosen as in Fig. 3(a) and (b).

By comparison, we see that the fermionic PT-MPO can be obtained in the same way as the PT-MPO of the corresponding spin system. One simply has to multiply the main input of the algorithm—the environment propagators for the individual environment modes $e^{\tilde{\mathcal{L}}_E^{(k)} \Delta t}$ —by the local parity operator $(P^{[k,k]})^{\sigma_0^+ \sigma_0^-}$ in both, forward and backward propagation direction.

Appendix C: Convergence using tensor distance

In the main text, we analyzed the convergence error for a specific system observable and system Hamiltonian. Alternatively, one may be interested in a global measure for convergence that is independent of the choice of observable and system propagator. Given two PT-MPOs A and B , we consider the normalized distance

$$d(A, B) = \frac{|A - B|^2}{|A||B|} = \frac{\langle A|A \rangle + \langle B|B \rangle - 2 \operatorname{Re}\{\langle A|B \rangle\}}{\sqrt{\langle A|A \rangle \langle B|B \rangle}}, \quad (\text{C1})$$

where, labeling the PT-MPO matrices of A and B by $\mathcal{Q}^{[A]}$ and $\mathcal{Q}^{[B]}$, respectively, the overlap is calculated by

contracting over the outer bonds

$$\langle A|B \rangle = \sum_{\substack{\alpha_{n-1}, \dots, \alpha_0 \\ \alpha'_n, \dots, \alpha'_1}} \sum_{\substack{d_{n-1}, \dots, d_1 \\ \tilde{d}_{n-1}, \dots, \tilde{d}_1}} \prod_{l=1}^n \left(\mathcal{Q}_{d_l \tilde{d}_{l-1}}^{[A](\alpha_l, \alpha'_l)} \mathcal{Q}_{\tilde{d}_l d_{l-1}}^{[B](\alpha_l, \alpha'_l)} \right). \quad (\text{C2})$$

Due to the MPO structure of the PTs, the distance can be calculated in $\mathcal{O}(nD^4(\chi_A^2\chi_B^2 + \chi_A\chi_B^2))$ floating point operations, where χ_A and χ_B are the inner bonds of PT-MPOs A and B , respectively, and D is the system Hilbert space dimension.

In Fig. 7, we perform the analogous convergence analysis as for the spin-boson model in Fig. 3(a) and (b) but using the distance between PT-MPOs calculated with nominal thresholds ϵ and the PT-MPO with the smallest threshold in the series. The convergence is similar to

that found in Fig. 3, where the tree-like scheme provides orders-of-magnitude speed-up over sequential ACE, and using a scale-dependent threshold results in improved accuracy for the same nominal compression threshold ϵ . Compared to Fig. 3, the analysis using the tensor distance produces much smoother convergence curves. However, we find that at small thresholds there is significant numerical cancellation, which results in an unreliable measure. The values can even become negative, here for $\epsilon \gtrsim 10^{-9}$, which is why no data points are shown for these values in Fig. 7.

Moreover, it should be noted that the distance between two PT-MPOs makes no quantitative statement about the accuracy of any physical quantity. We thus suggest that, despite the dependence on a concrete system propagator, analyzing the convergence behavior of a given physical observable as done in the main text is more useful for determining the accuracy of PT-MPOs in practical applications.

-
- [1] I. de Vega and D. Alonso, Dynamics of non-markovian open quantum systems, *Rev. Mod. Phys.* **89**, 015001 (2017).
- [2] A. Nazir and D. P. S. McCutcheon, Modelling exciton-phonon interactions in optically driven quantum dots, *J. Phys.: Condens. Matter* **28**, 103002 (2016).
- [3] P. W. Anderson, Localized magnetic states in metals, *Phys. Rev.* **124**, 41 (1961).
- [4] A. 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).
- [5] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002).
- [6] B. Krummheuer, V. M. Axt, T. Kuhn, I. D’Amico, and F. Rossi, Pure dephasing and phonon dynamics in gas- and gan-based quantum dot structures: Interplay between material parameters and geometry, *Phys. Rev. B* **71**, 235329 (2005).
- [7] D. E. Reiter, T. Kuhn, and V. M. Axt, Distinctive characteristics of carrier-phonon interactions in optically driven semiconductor quantum dots, *Adv. Phys.: X* **4**, 1655478 (2019).
- [8] A. Carmele and S. Reitzenstein, Non-markovian features in semiconductor quantum optics: quantifying the role of phonons in experiment and theory, *Nanophotonics* **8**, 655 (2019).
- [9] D. Mogilevtsev, A. P. Nisovtsev, S. Kilin, S. B. Cavalcanti, H. S. Brandi, and L. E. Oliveira, Non-markovian damping of rabi oscillations in semiconductor quantum dots, *J. Phys.: Condens. Matter* **21**, 055801 (2009).
- [10] A. Trushechkin, Long-term behaviour in an exactly solvable model of pure decoherence and the problem of markovian embedding, *Mathematics* **12**, 10.3390/math12010001 (2024).
- [11] M. Cosacchi, T. Seidelmann, M. Cygorek, A. Vagov, D. E. Reiter, and V. M. Axt, Accuracy of the quantum regression theorem for photon emission from a quantum dot, *Phys. Rev. Lett.* **127**, 100402 (2021).
- [12] A. Vagov, M. D. Croitoru, V. M. Axt, T. Kuhn, and F. M. Peeters, Nonmonotonic field dependence of damping and reappearance of rabi oscillations in quantum dots, *Phys. Rev. Lett.* **98**, 227403 (2007).
- [13] E. V. Denning, M. Bundgaard-Nielsen, and J. Mørk, Optical signatures of electron-phonon decoupling due to strong light-matter interactions, *Phys. Rev. B* **102**, 235303 (2020).
- [14] T. Kaldewey, S. Lüker, A. V. Kuhlmann, S. R. Valentin, J.-M. Chauveau, A. Ludwig, A. D. Wieck, D. E. Reiter, T. Kuhn, and R. J. Warburton, Demonstrating the decoupling regime of the electron-phonon interaction in a quantum dot using chirped optical excitation, *Phys. Rev. B* **95**, 241306 (2017).
- [15] A. Ramachandran, G. R. Wilbur, S. O’Neal, D. G. Deppe, and K. C. Hall, Suppression of decoherence tied to electron-phonon coupling in telecom-compatible quantum dots: low-threshold reappearance regime for quantum state inversion, *Opt. Lett.* **45**, 6498 (2020).
- [16] A. M. Barth, S. Lüker, A. Vagov, D. E. Reiter, T. Kuhn, and V. M. Axt, Fast and selective phonon-assisted state preparation of a quantum dot by adiabatic undressing, *Phys. Rev. B* **94**, 045306 (2016).
- [17] M. Cosacchi, F. Ungar, M. Cygorek, A. Vagov, and V. M. Axt, Emission-frequency separated high quality single-photon sources enabled by phonons, *Phys. Rev. Lett.* **123**, 017403 (2019).
- [18] C. Gustin and S. Hughes, Efficient pulse-excitation techniques for single photon sources from quantum dots in optical cavities, *Adv Quantum Technol.* **3**, 1900073 (2020).
- [19] J. Wiercinski, E. M. Gauger, and M. Cygorek, Phonon coupling versus pure dephasing in the photon statistics of cooperative emitters, *Phys. Rev. Res.* **5**, 013176 (2023).
- [20] Z. X. Koong, M. Cygorek, E. Scerri, T. S. Santana, S. I. Park, J. D. Song, E. M. Gauger, and B. D. Gerardot, Coherence in cooperative photon emission from indistinguishable quantum emitters, *Sci. Adv.* **8**, eabm8171 (2022).
- [21] G. McCauley, B. Cruikshank, D. I. Bondar, and K. Ja-

- cobs, Accurate lindblad-form master equation for weakly damped quantum systems across all regimes, *npj Quantum Inf.* **6**, 74 (2020).
- [22] K. Boos, S. K. Kim, T. Bracht, F. Sbresny, J. M. Kaspari, M. Cygorek, H. Riedl, F. W. Bopp, W. Rauhaus, C. Calcagno, J. J. Finley, D. E. Reiter, and K. Müller, Signatures of dynamically dressed states, *Phys. Rev. Lett.* **132**, 053602 (2024).
- [23] M. Cygorek, E. D. Scerri, T. S. Santana, Z. X. Koong, B. D. Gerardot, and E. M. Gauger, Signatures of cooperative emission in photon coincidence: Superradiance versus measurement-induced cooperativity, *Phys. Rev. A* **107**, 023718 (2023).
- [24] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, Non-markovian quantum processes: Complete framework and efficient characterization, *Phys. Rev. A* **97**, 012127 (2018).
- [25] M. 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).
- [26] R. Feynman and F. Vernon, The theory of a general quantum system interacting with a linear dissipative system, *Ann. Phys. (N.Y.)* **24**, 118 (1963).
- [27] G. E. Fux, E. P. Butler, P. R. Eastham, B. W. Lovett, and J. Keeling, Efficient exploration of hamiltonian parameter space for optimal control of non-markovian open quantum systems, *Phys. Rev. Lett.* **126**, 200401 (2021).
- [28] M. 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 (2022).
- [29] D. 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).
- [30] M. Cygorek, J. Keeling, B. W. Lovett, and E. M. Gauger, Sublinear scaling in non-markovian open quantum systems simulations, *Phys. Rev. X* **14**, 011010 (2024).
- [31] V. Link, H.-H. Tu, and W. T. Strunz, Open quantum system dynamics from infinite tensor network contraction, *Phys. Rev. Lett.* **132**, 200403 (2024).
- [32] J. Thoenniss, M. Sonner, A. Lerose, and D. A. Abanin, Efficient method for quantum impurity problems out of equilibrium, *Phys. Rev. B* **107**, L201115 (2023).
- [33] N. 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).
- [34] A. 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).
- [35] M. Cygorek and E. M. Gauger, ACE: A general-purpose non-Markovian open quantum systems simulation toolkit based on process tensors, *The Journal of Chemical Physics* **161**, 074111 (2024).
- [36] M. Cygorek and E. M. Gauger, Understanding and utilizing the inner bonds of process tensors (2024), [arXiv:2404.01287 \[quant-ph\]](https://arxiv.org/abs/2404.01287).
- [37] A modification of the ACE method can also in principle be applied to environments where each mode couples to “nearest neighbor” modes, such as arise from chain mapping techniques [47, 48].
- [38] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (N.Y.)* **326**, 96 (2011).
- [39] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New J. Phys.* **12**, 055026 (2010).
- [40] Nevertheless, most applications in quantum many-body physics still require an initial SVD line sweep to provide a starting point for an iterative scheme [38].
- [41] If a fixed final time t is discretized on a grid with n steps of size $\Delta t = t/n$, the total Trotter error scales as $\mathcal{O}(1/n)$ for the first-order and $\mathcal{O}(1/n^2)$ for the second-order Trotter formula, respectively, because Trotter errors occur at each of the n time steps.
- [42] N. Makri and D. E. Makarov, Tensor propagator for iterative quantum time evolution of reduced density matrices. I. Theory, *J. Chem. Phys.* **102**, 4600 (1995).
- [43] N. Makri and D. E. Makarov, Tensor propagator for iterative quantum time evolution of reduced density matrices. II. Numerical methodology, *J. Chem. Phys.* **102**, 4611 (1995).
- [44] Note however that there exist quasilinear [30] and linear-scaling [31] PT-MPO methods for Gaussian environments.
- [45] H. Maguire, J. Iles-Smith, and A. Nazir, Environmental nonadditivity and franck-condon physics in nonequilibrium quantum systems, *Phys. Rev. Lett.* **123**, 093601 (2019).
- [46] R. Chen, X. Xu, and C. Guo, Grassmann time-evolving matrix product operators for quantum impurity models, *Phys. Rev. B* **109**, 045140 (2024).
- [47] A. W. Chin, A. Datta, F. Caruso, S. F. Huelga, and M. B. Plenio, Noise-assisted energy transfer in quantum networks and light-harvesting complexes, *New J. Phys.* **12**, 065002 (2010).
- [48] D. Tamascelli, A. Smirne, J. Lim, S. F. Huelga, and M. B. Plenio, Efficient simulation of finite-temperature open quantum systems, *Phys. Rev. Lett.* **123**, 090402 (2019).