


Sublinear Scaling in Non-Markovian Open Quantum Systems Simulations

Moritz Cygorek¹, Jonathan Keeling², Brendon W. Lovett² and Erik M. Gauger¹

¹*SUPA, Institute of Photonics and Quantum Sciences, Heriot-Watt University,
Edinburgh EH14 4AS, United Kingdom*

²*SUPA, School of Physics and Astronomy, University of St Andrews,
St Andrews KY16 9SS, United Kingdom*

 (Received 12 April 2023; revised 11 August 2023; accepted 11 December 2023; published 1 February 2024)

While several numerical techniques are available for predicting the dynamics of non-Markovian open quantum systems, most struggle with simulations for very long memory and propagation times, e.g., due to superlinear scaling with the number of time steps n . Here, we introduce a numerically exact algorithm to calculate process tensors—compact representations of environmental influences—which provides a scaling advantage over previous algorithms by leveraging self-similarity of the tensor networks that represent the environment. It is applicable to environments with Gaussian statistics, such as for spin-boson-type open quantum systems. Based on a divide-and-conquer strategy, our approach requires only $\mathcal{O}(n \log n)$ singular value decompositions for environments with infinite memory. Where the memory can be truncated after n_c time steps, a nominal scaling $\mathcal{O}(n_c \log n_c)$ is found, which is independent of n . This improved scaling is enabled by identifying process tensors with repeatable blocks. To demonstrate the power and utility of our approach, we provide three examples. (1) We calculate the fluorescence spectra of a quantum dot under both strong driving and strong dot-phonon couplings, a task requiring simulations over millions of time steps, which we are able to perform in minutes. (2) We efficiently find process tensors describing superradiance of multiple emitters. (3) We explore the limits of our algorithm by considering coherence decay with a very strongly coupled environment. The observed computation time is not necessarily proportional to the number of singular value decompositions because the matrix dimensions also depend on the number of time steps. Nevertheless, quasilinear and sublinear scaling of computation time is found in practice for a wide range of parameters. While there are instances where existing methods can achieve comparable nominal scaling by precalculating effective propagators for time-independent or periodic system Hamiltonians, process tensors contain all the information needed to extract arbitrary multitime correlation functions of the system when driven by arbitrary time-dependent system Hamiltonians. The algorithm we present here not only significantly extends the scope of numerically exact techniques to open quantum systems with long memory times, but it also has fundamental implications for the simulation complexity of tensor network approaches.

DOI: [10.1103/PhysRevX.14.011010](https://doi.org/10.1103/PhysRevX.14.011010)

Subject Areas: Computational Physics,
Condensed Matter Physics,
Quantum Physics

I. INTRODUCTION

A common challenge in quantum technology is the ubiquity of dephasing and dissipation caused by interactions between quantum systems and their surrounding environment [1]. Thus, understanding environmental influences is crucial for mitigating them [2–6] or even using them to enhance the functionality of quantum devices, as in phonon-assisted state preparation [7,8] or environment-assisted

quantum transport [9–11]. A standard approach for modeling the dynamics of open quantum systems is the Lindblad master equation [1,12]. This equation can be derived using perturbation theory and the Born-Markov approximation, which assumes that the environment is static and memoryless and couples only weakly to the system. However, this is often an oversimplification, and more sophisticated methods are required to accurately model the dynamics in many cases, where the structure of the environment matters or where the coupling is strong. Examples include solid-state quantum dots (QDs) coupled to local phonon environments [4,13–16], charge or excitation transport in biomolecules [10,17,18], ultrafast spin dynamics [19,20], spontaneous emission in photonic structures [21,22], and superconducting qubits in quantum computers [2].

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Beyond the restrictive weak-coupling Born-Markov approximation, the regime of non-Markovian dynamics [23] arises. Here, a quantum system triggers a dynamical evolution of the environment, and the environment's response feeds back to the system at a later point in time, constituting a time-nonlocal memory. Several strategies have been discussed for modeling non-Markovian dynamics [23]. One general strategy is to solve the Schrödinger equation for the total closed system composed of the system of interest and its environment. Because environments typically consist of a large number—often a continuum—of degrees of freedom, this becomes numerically intractable unless the model is somehow reduced to a finite number of effective degrees of freedom. How to choose the relevant degrees of freedom differs from method to method: Mean-field, cumulant, or cluster expansions [24,25] are based on the assumption that higher-order correlations are negligible, and thus one may consider only degrees of freedom within the subspace of product states, or states with low-order correlations. Other methods based on techniques from many-body quantum theory employ efficient representations of the total system such as multilayer multiconfiguration time-dependent Hartree wave functions [26] or matrix product states [18,27,28]. In the reaction coordinate mapping [29], parts of the environment are explicitly taken into account by extending the system of interest. A similar concept is to replace the actual environment by a small set of auxiliary oscillators [30].

A second general strategy for simulating open quantum systems is to keep the description confined to the degrees of freedom of the system but to account for environmental effects via equations of motion that are nonlocal in time. For example, the Nakajima-Zwanzig formalism [1,31] explicitly includes memory effects of the form of an integral over past times. The Feynman-Vernon path-integral formalism [32] encapsulates environmental effects in an influence functional, which acts as a trajectory-dependent weight in a sum over all possible system trajectories. This formalism is the starting point for several practical methods for open quantum systems simulation. For Gaussian environments with exponentially decaying bath correlations, repeated time differentiation of the Feynman-Vernon path sum gives rise to a set of hierarchical equations of motion (HEOM) [33,34]. Alternatively, the path sum can be expressed as a stochastic average with bath correlations encoded in the noise statistics [35], from which one can derive non-Markovian stochastic Schrödinger equations [36–38]. The combination of both ideas leads to equations for a hierarchy of pure states (HOPS) [39]. These concepts can be further generalized, e.g., to stochastic master equations for non-Gaussian environments [40] and to open quantum systems simulations where the environment is continuously measured [41].

Here, we focus on the question of how to predict non-Markovian dynamics for both long propagation times and

long memory times. While some specific features of long time dynamics—such as the nature of the steady state—may sometimes be more accessible straightforwardly, the problem of modeling non-Markovian *dynamics* over long times is important but challenging. For example, one may want to find the general time evolution, or multitime correlations of a system, but in all the methods described above, it is generally hard to track such non-Markovian dynamics over long times. Even though strategies of explicitly representing environment degrees of freedom at first glance seem to scale linearly with the total number of time steps n , i.e., $\mathcal{O}(n)$, they tend to become inaccurate with increasing propagation time. This is due to the increase with time of the number of degrees of freedom of the environment that can be excited and thus entangled with the system. For example, consider the discretization of a continuum of phonon modes by energy intervals of width $\hbar\Delta\omega$. Propagating over a time t_e , energy-time uncertainty implies that a discretization of $\Delta\omega \lesssim 1/t_e$ is required to preclude resolving a difference between the continuum and the discretized model. The need for a finer discretization results in a scaling of $\mathcal{O}(n^2)$ unless a more sophisticated, e.g., adaptive, discretization is used, or the resolution of individual sampling points is blurred by some form of line broadening [27]. In approaches where the environment is mapped to a chain of coupled oscillators [18,27], this statement is related to the observation that the Lieb-Robinson theorem restricts environment excitation to a light cone [42], which increases in size with increasing propagation time t_e . Similarly, stochastic sampling requires more trajectories to reach the same accuracy when the propagation time t_e is increased. Hence, most numerical methods for open quantum systems show superlinear scaling with the number of time steps n .

For direct applications of the Feynman-Vernon influence functional approach, one finds exponential scaling with n . However, a reformulation into the iterative scheme QUAPI [43,44] yields truly linear scaling when the memory time of the environment is finite, so the memory can be truncated after n_c time steps. Such an approach, however, retains exponential scaling with the memory time n_c . It could be expected that $\mathcal{O}(n)$ is the natural lower bound for the scaling, in particular, for general open quantum systems with time-dependent driving, because, at the very least, $\mathcal{O}(n)$ observables have to be obtained in order to extract the full time evolution. However, another way to simulate open quantum system dynamics is to construct effective propagators for the reduced system density matrix before propagating the system using a series of small matrix multiplications, such as in the transfer tensor [45] and the small matrix decomposition of the path-integral (SMatPI) [46] approaches. For the special case of a time-independent or periodically driven system Hamiltonian, time translation invariance allows one to only calculate effective propagators over short times and reuse them for the propagation

over much longer times $n \gg n_c$. While the propagation over n time steps, in principle, scales linearly in n , it does so with a very small prefactor, whereas precalculating the effective propagators is formally constant in n but typically with a larger prefactor, as it has to incorporate the complexity of the environment response within the memory time n_c . Hence, in most practical applications, the latter dominates, which results in sublinear behavior with n . In contrast, for applications involving nonperiodic time-dependent Hamiltonians, the effective propagator must be recalculated for each time step [47].

In this paper, we introduce an algorithm based on the process tensor (PT) framework that demonstrates sublinear scaling of the numerically demanding step. A PT captures environment influences equivalent to the Feynman-Vernon influence functional in a numerically exact way and can be efficiently represented in the form of a matrix product operator (MPO) [48,49]. Once obtained, an open quantum system with an arbitrary time-dependent system Hamiltonian can be propagated in time by simple matrix multiplications on a vector space given by the product of the system Liouville space and the inner dimension of the PT-MPO, which we denote as χ . The bottleneck in MPO techniques is the MPO compression, i.e., the reduction of the inner dimension χ , using rank-reducing operations, often achieved by truncated singular value decomposition (SVD). We show that the self-similarity of time-independent Gaussian environments—which include electromagnetic environments, typical vibrational baths, and the paradigmatic spin-boson model—can be exploited to devise a divide-and-conquer scheme that reduces the number of SVDs from $\mathcal{O}(n^2)$ in the algorithm introduced by Jørgensen and Pollock [49] to $\mathcal{O}(n \log n)$. If the memory time is further limited to n_c time steps, we arrive at a theoretical scaling $\mathcal{O}(n_c \log n_c)$, constant in n . The actual scaling can be larger than this, as the bond dimension χ , and thus the time required for SVD evaluation, can depend on the propagation time. However, we will see from numerical results that sublinear scaling with n is indeed seen in the examples we study.

To test the performance of our algorithm in practice and to demonstrate a sample of new applications it enables, we provide several examples: First, we investigate the fluorescence spectra of semiconductor QDs coupled to a bath of acoustic phonons with strong driving. While being of considerable interest for experiments [50,51], numerically calculating QD fluorescence spectra is a challenging multi-scale problem because the width of the zero-phonon line is determined by radiative lifetimes of the order of nanoseconds; however, typical phonon memory times are of the order of picoseconds, and strong driving forces us to use small time steps on the femtosecond scale. The sublinear scaling of our algorithm enables us to obtain numerically exact spectra from simulations involving a million time steps within minutes on a conventional laptop computer.

Second, we use the capability of our algorithm to deal with very many time steps to study superradiance of multiple emitters without making any rotating wave approximation (RWA). This allows us to describe the breakdown of superradiance in the presence of disorder and dephasing due to interactions with phonon environments. Finally, we discuss coherence decay for a system with strong coupling to an environment with a strongly peaked spectral density. This example illuminates where the limitations of our algorithm arise.

The article is structured as follows: In Sec. II, we introduce and describe our algorithm, where, in Sec. II A, we first summarize the PT formalism, on which our algorithm is based. For comparison with the commonly used sequential algorithm by Jørgensen and Pollock [49], and to introduce quantities also used in our approach, we revise the PT calculation scheme for Gaussian environments in Ref. [49] in Sec. II B. Then, we introduce our divide-and-conquer scheme in Sec. II C and introduce periodic PTs in Sec. II D. The examples of fluorescence spectra, multi-emitter superradiance, and coherence decay are discussed in Secs. III A–III C, respectively. Our results are summarized in Sec. IV.

II. THEORY

A. Process tensors

An open quantum system is defined by dividing the total system into a system of interest S and an environment E . Correspondingly, the total Hamiltonian is decomposed into $H = H_S + H_E$, where H_S is an arbitrary, possibly time-dependent, system Hamiltonian and H_E is the environment Hamiltonian, which also includes the interaction with the system of interest. In this paper, we present an algorithm applicable to environments with Gaussian correlations, focusing on a generalized spin-boson model consisting of a few-level system coupled linearly to a bath of harmonic oscillators, described by

$$H_E = \sum_k \hbar \omega_k b_k^\dagger b_k + \sum_k \hbar (g_k b_k^\dagger + g_k^* b_k) \hat{O}, \quad (1)$$

where \hat{O} is a Hermitian operator acting on the system Hilbert space. The environment Hamiltonian is conveniently characterized by the spectral density $J(\omega) = \sum_k |g_k|^2 \delta(\omega - \omega_k)$. Here, we choose to work in the basis where the coupling is diagonal $\hat{O} = \sum_u \lambda_u |u\rangle\langle u|$ with a set of basis states $|u\rangle$ of the system Hilbert space.

Throughout this article, we use a compact Liouville space notation and assume a uniform time grid $t_j = t_0 + j\Delta t$ with time step Δt . For example, the reduced system density matrix $\bar{\rho}(t) = \text{Tr}_E[\rho(t)]$ at time t_j is denoted by $\bar{\rho}_{\alpha_j} = \bar{\rho}_{s_j, r_j}(t_j) = \langle s_j | \bar{\rho}(t) | r_j \rangle$, where s_j and r_j are system Hilbert space indices, which are combined into the Liouville space index $\alpha_j = (s_j, r_j)$. The time argument t_j

is implied in the index α_j . In this notation, the reduced density matrix $\bar{\rho}_{\alpha_n}$ at a final time t_n can be expressed as [32]

$$\bar{\rho}_{\alpha_n} = \sum_{\alpha_{n-1}, \dots, \alpha_0} \mathcal{F}^{\alpha_n, \dots, \alpha_1} \left(\prod_{l=1}^n \mathcal{M}_{\alpha_l, \alpha_{l-1}} \right) \bar{\rho}_{\alpha_0}, \quad (2)$$

where $\mathcal{M}_{\alpha_l, \alpha_{l-1}} = (e^{\mathcal{L}_S \Delta t})_{\alpha_l, \alpha_{l-1}}$ describes the free evolution of the system during one time step Δt under the free system Liouvillian; e.g., $\mathcal{L}_S[\bar{\rho}] = -(i/\hbar)[H_S, \bar{\rho}]$, and $\mathcal{F}^{\alpha_n, \dots, \alpha_1}$ is the Feynman-Vernon influence functional [32], which exactly captures all environment influences in the limit $\Delta t \rightarrow 0$.

For the generalized spin-boson model, the influence functional is given by [49,52]

$$\mathcal{F}^{\alpha_n, \dots, \alpha_1} = \prod_{i=1}^n \prod_{j=1}^i [b_{i-j}]^{\alpha_i, \alpha_j} = \prod_{j=1}^n \prod_{l=0}^{n-j} [b_l]^{\alpha_{l+j}, \alpha_j}, \quad (3)$$

where the factors $[b_{i-j}]^{\alpha_i, \alpha_j}$ are related to the bath correlation function via

$$[b_{(i-j)}]^{\alpha_i, \alpha_j} = e^{-(\lambda_{s_i} - \lambda_{r_i})(\eta_{i-j} \lambda_{s_j} - \eta_{i-j}^* \lambda_{r_j})} \quad (4)$$

with $\alpha_j = (s_j, r_j)$ and $\lambda_{s,r}$ as defined above, and [53]

$$\eta_{i-j} = \begin{cases} \int_{t_{i-1}}^{t_i} dt' \int_{t_{i-1}}^{t'} dt'' C(t' - t'') & i = j \\ \int_{t_{i-1}}^{t_i} dt' \int_{t_{j-1}}^{t_j} dt'' C(t' - t'') & i \neq j, \end{cases} \quad (5)$$

where

$$C(t) = \int_0^\infty d\omega J(\omega) \left[\coth\left(\frac{1}{2}\beta\hbar\omega\right) \cos(\omega t) - i \sin(\omega t) \right]. \quad (6)$$

The memory of the environment is finite if the bath correlation function $C(t)$ vanishes after n_c time steps, i.e., $C(t \geq n_c \Delta t) \approx 0$, which implies $\eta_{l \geq n_c} \approx 0$ and $[b_{l \geq n_c}]^{\alpha_i, \alpha_j} \approx 1$. In this case, it is sufficient to consider at most n_c terms in the second product in the right-hand side of Eq. (3).

Performing the Feynman-Vernon summation in Eq. (2) is notoriously difficult. For a system Hilbert space of dimension D , the sum involves D^{2n} terms and, thus, scales exponentially with the total number of time steps n . This issue can be addressed [49,52] by representing the influence functional in a more convenient form using matrix product operators (MPOs) [54,55]. It is then referred to as the process tensor matrix product operator (PT-MPO),

$$\mathcal{F}^{\alpha_n, \dots, \alpha_1} = \sum_{d_{n-1}, \dots, d_1} \mathcal{Q}_{1, d_{n-1}}^{\alpha_n} \mathcal{Q}_{d_{n-1}, d_{n-2}}^{\alpha_{n-1}} \dots \mathcal{Q}_{d_2, d_1}^{\alpha_2} \mathcal{Q}_{d_1, 1}^{\alpha_1}, \quad (7)$$

where the monolithic tensor $\mathcal{F}^{\alpha_n, \dots, \alpha_1}$ is decomposed into a set of smaller elements $\mathcal{Q}_{d_l, d_{l-1}}^{\alpha_l}$, which are regarded as

matrices with respect to the inner bond indices d_l . The latter serve the purpose of conveying time-nonlocal information, i.e., memory, over multiple time steps. The maximal dimension χ of the inner bonds d_l strongly depends on the complexity of the environment [48].

The PT-MPO representation achieves the reduction of Eq. (2) to

$$\bar{\rho}_{\alpha_n} = \sum_{\substack{\alpha_{n-1}, \dots, \alpha_0 \\ d_{n-1}, \dots, d_1}} \left(\prod_{l=1}^n \mathcal{Q}_{d_l, d_{l-1}}^{\alpha_l} \mathcal{M}_{\alpha_l, \alpha_{l-1}} \right) \bar{\rho}_{\alpha_0}, \quad (8)$$

which can be summed sequentially, one time step at a time, with the complexity of n matrix multiplications of dimension $D^2\chi$. The main challenge for simulating the open quantum system is thus to bring the influence functional in Eq. (3) into the form of the PT-MPO in Eq. (7).

B. Sequential algorithm

Jørgensen and Pollock [49] devised the algorithm that is currently most commonly used to obtain PTs in MPO form for Gaussian environments. To distinguish it from our divide-and-conquer scheme, we refer to the approach in Ref. [49] as the sequential algorithm. The common starting point of both algorithms is the graphical representation of the double product on the right-hand side of Eq. (3) in the form of a triangular tensor network (introduced in Ref. [52]), depicted in Figs. 1(a) and 1(c) for $n = 8$ time steps and memory cutoff $n_c = 6$. The latter limits the maximum length of the rows.

Each node represents a factor $c_{\beta_{l+1}, \beta_l}^{\alpha_l} = \delta_{\beta_{l+1}, \beta_l} [b_l]^{\alpha_{l+j}, \beta_l}$, where system Liouville space indices α_l are represented by upward facing links, while β_{l+1} and β_l are links to the left and right, respectively. Left or right dangling indices are traced out. Note that, in the notation of Ref. [49], the upward facing links are meant to pass by those nodes further up the tensor network and then connect to the same outer index α_l . The more conventional interpretation of tensor networks, where a bottom node is instead connected to its neighbor to the top, is regained by adding a fourth leg with index $\tilde{\alpha}_l$ at the bottom of each node (except those on the bottom row) via another Kronecker delta, i.e., $\tilde{c}_{\beta_{l+1}, \beta_l}^{\alpha_l, \tilde{\alpha}_l} = c_{\beta_{l+1}, \beta_l}^{\alpha_l} \delta_{\alpha_l, \tilde{\alpha}_l}$. In the following, however, we retain the notation of Ref. [49], labeling only three indices on such tensors.

In the sequential algorithm, the PT is written as a product of rows,

$$\mathcal{F}^{\alpha_k, \dots, \alpha_1} = \prod_{j=1}^k \mathcal{C}^{\alpha_k, \dots, \alpha_j} = \prod_{j=1}^k \left[\sum_{\beta_{k-j+1}, \dots, \beta_0} \prod_{l=0}^{k-j} c_{\beta_{l+1}, \beta_l}^{\alpha_{l+j}} \right], \quad (9)$$

where $\mathcal{C}^{\alpha_k, \dots, \alpha_j}$ is the j th row in the tensor network, which, as described above, may be reduced to a product of at most n_c terms if memory truncation is employed. The PT

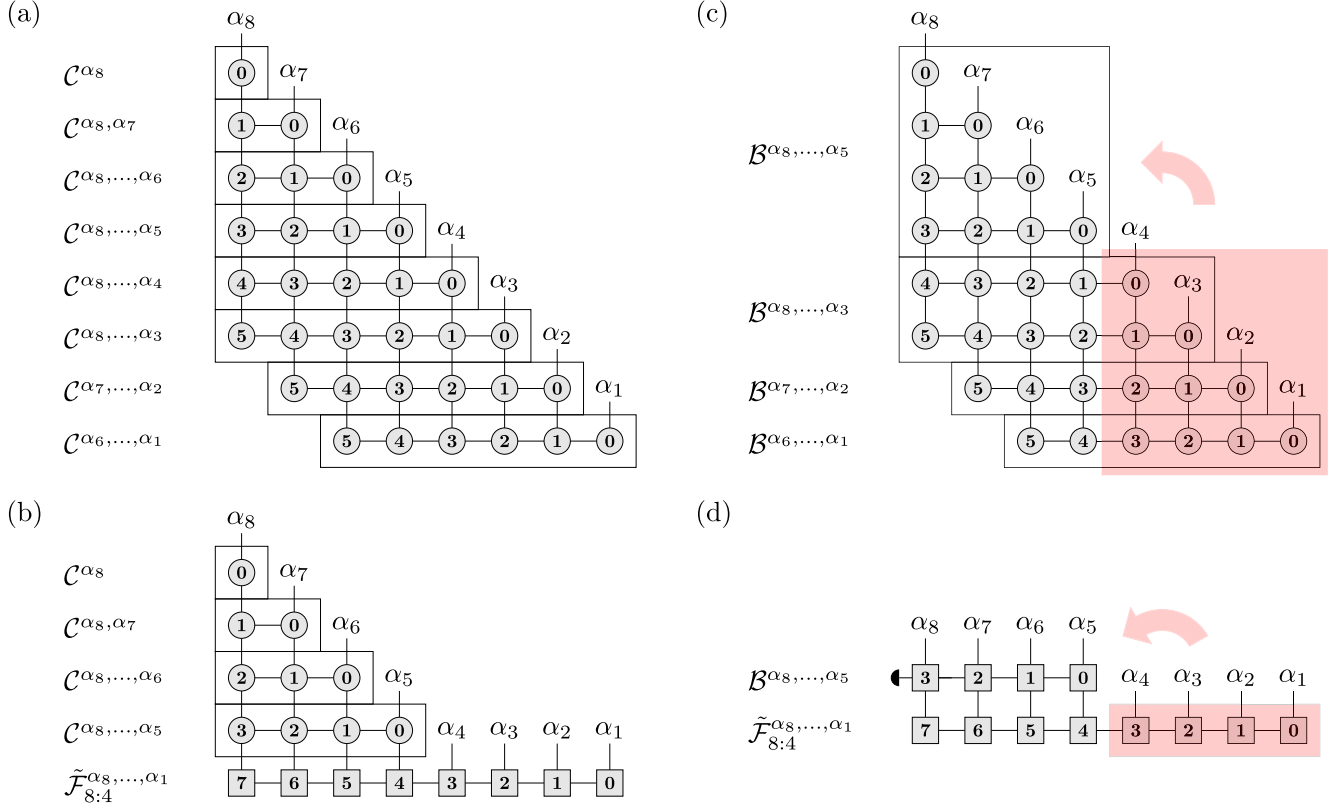


FIG. 1. Tensor network for a PT with $n = 8$ time steps and memory cutoff $n_c = 6$ using the sequential algorithm by Jørgensen and Pollock [49] [(a),(b)] as well as our divide-and-conquer scheme [(c),(d)], respectively. Panels (a) and (c) show the subdivision of the overall tensor network into different blocks (black rectangles) to be contracted from bottom to top. In the sequential algorithm, this is done row by row (b). The divide-and-conquer algorithm is based on the observation that the last block in panel (c), formed by the topmost $n/2$ lines, exactly matches the part of the already-contracted network shaded in red. This enables a contraction of multiple rows at a time as depicted in panel (d). The black semicircle seen here represents a “closure” that describes tracing out dangling bonds; see text for further discussion.

$\mathcal{F}^{\alpha_k, \dots, \alpha_1}$ is constructed by multiplying row after row from bottom to top using the recursion visualized in Fig. 1(b),

$$\tilde{\mathcal{F}}_{(k:j+1)}^{\alpha_k, \dots, \alpha_1} = \mathcal{C}^{\alpha_k, \dots, \alpha_j} \tilde{\mathcal{F}}_{(k:j)}^{\alpha_k, \dots, \alpha_1}, \quad (10)$$

with initial value $\tilde{\mathcal{F}}_{(k:1)}^{\alpha_k, \dots, \alpha_1} = \mathcal{C}^{\alpha_k, \dots, \alpha_1}$. The final PT is identified as $\mathcal{F}^{\alpha_n, \dots, \alpha_1} = \tilde{\mathcal{F}}_{(n:n)}^{\alpha_n, \dots, \alpha_1}$. Crucially, the iteration in Eq. (10) retains the MPO form $\tilde{\mathcal{F}}_{k:j}^{\alpha_k, \dots, \alpha_1} = \sum_{d_k, \dots, d_1} \prod_{l=1}^k [\tilde{f}_{k:j}^{\alpha_l}]_{d_l, d_{l-1}}$, with

$$[\tilde{f}_{k:j+1}^{\alpha_l}]_{d_l, d_{l-1}} = [\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d'_{l-1}} [c^{\alpha_l}]_{\beta_l, \beta_{l-1}}. \quad (11)$$

However, the dimension of inner indices $d_l = (d'_l, \beta_l)$ is expanded to the product of the inner dimension of the previous PT d'_l with that of a single row β_l . If χ' denotes the typical inner dimension of the MPO $\tilde{\mathcal{F}}_{(k:j)}^{\alpha_k, \dots, \alpha_1}$ and the inner dimension of a single line corresponds to D^2 with system Hilbert space dimension D , the inner dimension of the product is now increased to $\chi = \chi' D^2$. To keep the inner

dimensions tractable, the MPO is compressed by sweeping across it and applying truncated SVDs to every element. In the forward direction (increasing time-step indices), one calculates the SVD,

$$[\tilde{f}_{k:j}^{\alpha_l}]_{d_l, d_{l-1}} = \sum_s U_{d_l, s} \sigma_s V_{s, (\alpha_l, d_{l-1})}^\dagger, \quad (12)$$

with non-negative singular values σ_s and matrices U and V with orthogonal columns. Keeping only terms s with significant values $\sigma_s \geq \epsilon \sigma_0$, where ϵ is a given truncation threshold and σ_0 is the largest singular value, one replaces

$$[\tilde{f}_{k:j}^{\alpha_l}]_{s, d_{l-1}} \leftarrow V_{s, (\alpha_l, d_{l-1})}^\dagger, \quad (13a)$$

$$[\tilde{f}_{k:j}^{\alpha_{l+1}}]_{d_{l+1}, s} \leftarrow \sum_{d_l} [\tilde{f}_{k:j}^{\alpha_{l+1}}]_{d_{l+1}, d_l} U_{d_l, s} \sigma_s, \quad (13b)$$

which passes on the singular values—indicators for the importance of the represented degree of freedom—to the next element in the PT while simultaneously reducing

the inner dimension to the number of significant singular values. A forward sweep is followed by an analogous backward sweep. While this scheme to compress the PT-MPO was introduced in Ref. [49], we find empirically that changing the order of sweeps, namely, performing the backward sweep before the forward sweep, leads to smaller inner PT dimensions χ and hence to shorter computation times. For all simulations presented here, we therefore consistently compress PT-MPOs in this reversed order. The combination and backward sweep is diagrammatically represented in Fig. 2(a).

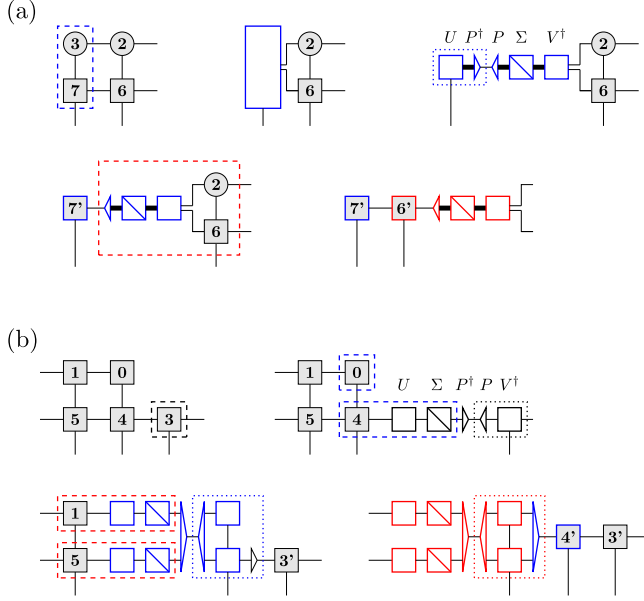


FIG. 2. (a) MPO combination and compression by backward sweep used in the sequential algorithm [see Fig. 1(b)]. The units enclosed by the dashed lines are combined into a single matrix. A truncated SVD $UP^\dagger P\Sigma V^\dagger$ is performed, where P denotes the projection onto the space spanned by singular vectors with singular values $\sigma \geq \epsilon\sigma_0$, which is shown as a triangle, reducing a thick connection (with large dimensions) to a thin connection (with smaller dimensions). The dotted lines indicate units that are combined to form the nodes in the updated PT-MPO. The part $P\Sigma V^\dagger$ is passed onto the nodes corresponding to the previous time step, before another SVD is performed. This process is repeated until the first time step is reached. (b) MPO combination and index preselection for the divide-and-conquer method [see Fig. 1(d)]. A forward sweep with truncated SVDs $U\Sigma P^\dagger P V^\dagger$ is performed on the last matrix of the bottom PT-MPO block before it overlaps with matrices of the top PT-MPO block. The corresponding matrices $U\Sigma$ are passed onto the next MPO matrix. SVDs are performed at the top and bottom blocks independently. Singular values from both SVDs are used to identify the projectors onto the subspaces spanned by singular vectors for which $\sigma^{(1)}\sigma^{(2)} \geq \epsilon\sigma_0^{(1)}\sigma_0^{(2)}$, which is shown as a large triangle with three external legs. Updated blocks of the PT-MPO are formed by compressing the Kronecker product of the corresponding matrices $V^{(1)\dagger}$ and $V^{(2)\dagger}$ using these projectors (dotted boxes). These steps are repeated until the last time step is reached.

Overall, the calculation of a PT-MPO using the sequential algorithm requires $\mathcal{O}(n^2)$ SVDs without memory truncation or $\mathcal{O}(nn_c)$ SVDs with memory truncation. Exact algorithms to perform SVDs of $N \times M$ matrices with $N > M$ generally require $\mathcal{O}(NM^2)$ floating point operations. Denoting by χ' the typical inner bond dimension of the PT-MPO after the previous iteration step, the corresponding matrix dimensions for a single SVD as in Eq. (12) are $M = D^2\chi'$ and $N = D^4\chi'$, respectively, where a factor D^2 in the long matrix dimension N stems from the outer bonds of the PT-MPO while the remaining factors originate from combined inner indices. Consequently, the number of floating point operations required per SVD is $\mathcal{O}(\chi'^3 D^8)$.

C. Divide-and-conquer algorithm

Guided by the visual representation of the PT contraction in Fig. 1(c), we suggest an alternative algorithm based on the observation that the triangular tensor network contains a high degree of self-similarity: The first and second rows are identical up to a shift. Rows three and four together form a shifted and truncated replica of the combination of rows one and two. Similarly, rows five to eight are the same as the part of rows one to four shaded in red in Fig. 1(c).

Concretely, the tensor network can be contracted in blocks of rows, with sizes progressing in powers of two, by iterating:

$$\tilde{\mathcal{F}}_{k:2^{m+1}}^{\alpha_k, \dots, \alpha_1} = \mathcal{B}^{\alpha_k, \dots, \alpha_{(2^m+1)}} \tilde{\mathcal{F}}_{k:2^m}^{\alpha_k, \dots, \alpha_1}, \quad (14)$$

where the block $\mathcal{B}^{\alpha_k, \dots, \alpha_{(2^m+1)}}$ is itself contained in the MPO of $\tilde{\mathcal{F}}_{k:2^m}^{\alpha_k, \dots, \alpha_1}$ and can be written as

$$\mathcal{B}^{\alpha_k, \dots, \alpha_{(2^m+1)}} = \sum_{d_{k-2^m}, \dots, d_0} q_{d_{k-2^m}} \prod_{l=1}^{k-2^m} [\tilde{\mathcal{F}}_{k:2^m}^{\alpha_l}]_{d_l, d_{l-1}}. \quad (15)$$

Here, $q_{d_{k-2^m}}$, represented as a black semicircle in Fig. 1(d), defines objects we call ‘‘closures’’ that describe tracing out dangling bonds. A similar object has been referred to as a ‘‘cap tensor’’ in other works [56]. Such objects are trivial before MPO compression, but after compression, they become nontrivial. However, the closures can be calculated iteratively from PT-MPOs by using trace preservation [57]. For the generalized spin-boson model considered here, one starts with $q_{d_n=1} = 1$ and iterates $q_{d_{i-1}} = \sum_{d_i} q_{d_i} \mathcal{Q}_{d_i, d_{i-1}}^{\alpha_i=(x,x)}$, where x is a fixed but arbitrary system Hilbert space index (e.g., $x = 1$). As can be seen in Eq. (4), setting $\alpha_i = (x, x)$ results in $[b_{(i-j)}]_{\alpha_i=(x,x), \alpha_j} = 1$. This effectively terminates the double product in Eq. (3) at an earlier time step.

The proposed algorithm follows the divide-and-conquer paradigm because the total number of rows of the tensor network is divided into an upper and a lower block, where only the lower block has to be calculated explicitly.

The lower block is then again divided into two halves, and so on. Like other divide-and-conquer schemes, such as the famous fast Fourier transformation algorithm [58,59] or efficient modular path-integral techniques for multisite problems [60], the numerical complexity is reduced from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \log n)$ operations. In our case, these operations are SVDs.

Despite the favorable scaling with n , the divide-and-conquer algorithm faces some challenges. These arise because of the size of the matrices we must combine. In the sequential algorithm, each step combines a MPO with potentially large inner dimension χ' and a single line of relatively small inner dimension D^2 . In contrast, in our divide-and-conquer algorithm, two MPOs of similar inner dimensions χ' are combined to $\chi = \chi'^2$. Because of the scaling of SVD routines with the third power of the combined inner dimension together with a factor D^2 from the outer bond, the direct application of SVDs would lead to a typical complexity $\mathcal{O}(\chi'^6 D^2)$ per SVD, which is prohibitively demanding. Clearly, a different way to combine two MPOs is needed.

Here, we address this issue by preselecting relevant degrees of freedom: Suppose we wish to update an MPO with matrices $[\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d'_{l-1}}$ by multiplying it with an MPO with matrices $[\tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}}$. We first perform SVDs on the individual matrices $[\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d'_{l-1}} = \sum_s U_{d'_l, s}^{(1)} \sigma_s^{(1)} V_{s, (\alpha_l, d'_{l-1})}^{(1)\dagger}$ and $[\tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}} = \sum_t U_{d''_l, t}^{(2)} \sigma_t^{(2)} V_{t, (\alpha_l, d''_{l-1})}^{(2)\dagger}$, respectively. Then, the matrices of the combined MPO can be formally written as

$$\begin{aligned} & [\tilde{f}_{k:j+1}^{\alpha_l}]_{(d'_l, d''_l), (d'_{l-1}, d''_{l-1})} \\ &= [\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d'_{l-1}} [\tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}} \\ &= \sum_{s,t} (U_{d'_l, s}^{(1)} U_{d''_l, t}^{(2)}) \sigma_s^{(1)} \sigma_t^{(2)} (V_{s, (\alpha_l, d'_{l-1})}^{(1)\dagger} V_{t, (\alpha_l, d''_{l-1})}^{(2)\dagger}). \end{aligned} \quad (16)$$

Note again that, to keep the notation similar to the sequential algorithm in Ref. [49], Eq. (16) is considered an elementwise product for each value of α_l . In the established notation for more general tensor networks, Eq. (16) would be considered a contraction over outer indices after including a Kronecker delta $[\tilde{f}_{k:j+1}^{\alpha_l}]_{(d'_l, d''_l), (d'_{l-1}, d''_{l-1})} = \sum_{\tilde{\alpha}_l} [\tilde{f}_{k:j}^{\tilde{\alpha}_l}]_{d'_l, d'_{l-1}} [\delta_{\tilde{\alpha}_l, \alpha_l} \tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}}$.

To reduce the inner dimension, we keep only combinations of indices (s, t) for which the product of singular values $\sigma_s^{(1)} \sigma_t^{(2)} \geq \epsilon_{\text{select}} \sigma_0^{(1)} \sigma_0^{(2)}$ exceeds a value determined by a given threshold ϵ_{select} . In practice, the new matrix is directly set to

$$[\tilde{f}_{k:j+1}^{\alpha_l}]_{(s,t), (d'_{l-1}, d''_{l-1})} \leftarrow V_{s, (\alpha_l, d'_{l-1})}^{(1)\dagger} V_{t, (\alpha_l, d''_{l-1})}^{(2)\dagger} \quad (17)$$

for the subset of (s, t) obeying the condition above, while the remaining terms are passed on to the next matrices of the original MPOs:

$$[\tilde{f}_{k:j}^{\alpha_{l+1}}]_{d'_{l+1}, s} \leftarrow \sum_{d'_l} [\tilde{f}_{k:j}^{\alpha_{l+1}}]_{d'_{l+1}, d'_l} U_{d'_l, s}^{(1)} \sigma_s^{(1)}, \quad (18)$$

$$[\tilde{g}^{\alpha_{l+1}}]_{d''_{l+1}, t} \leftarrow \sum_{d''_l} [\tilde{g}^{\alpha_{l+1}}]_{d''_{l+1}, d''_l} U_{d''_l, t}^{(2)} \sigma_t^{(2)}. \quad (19)$$

The combined forward sweep, selection, and combination process is depicted in Fig. 2(b). It is followed by a backward sweep to further reduce bond dimensions. Over a wide range of different examples, we find that the selection reduces the combined inner dimensions from χ'^2 to $\lambda \chi'$, with an empirical factor λ between 2 and 12. The resulting number of floating point operations per SVD $\mathcal{O}(\lambda^3 \chi'^3 D^2)$ is therefore nominally comparable to that in the sequential algorithm, $\mathcal{O}(\chi'^3 D^2)$. However, the massive reduction of inner bonds by our selection strategy comes at the cost of not guaranteeing an optimal low-rank approximation. This can result in larger bond dimensions χ' and therefore in higher numerical demands per SVD compared to the sequential algorithm. As discussed in more detail in Appendix A, this issue can be ameliorated by choosing different thresholds for the selection and backward sweep compared to the forward sweep, which we parametrize by the ratios $r_s = \epsilon_{\text{select}} / \epsilon_{\text{forward}}$ and $r_b = \epsilon_{\text{backward}} / \epsilon_{\text{forward}}$, while at the same time, we obtain similar accuracies as for the sequential algorithm if we identify $\epsilon_{\text{forward}}$ (the largest of our thresholds) with the nominal threshold ϵ of the sequential algorithm.

Finally, it is noteworthy that our divide-and-conquer algorithm can be implemented in place, i.e., in such a way that only a single copy of the full PT-MPO has to be stored, despite the combination of MPO matrices at different positions in the MPO chain. This reduces the memory footprint of the algorithm. Furthermore, as the line sweeps access the individual MPO matrices in a predictable order, storing the PT-MPO on a hard disk and preloading blocks of matrices can be efficiently implemented, with only a small overhead of about 10% in overall computation time. This process enables the calculation of PTs with many more time steps n than would be possible if required to keep the full MPO in memory.

However, it should be noted that, as discussed above, the time required for each SVD scales as $\mathcal{O}(\chi^3)$. Since χ , in general, depends on the complexity of the environment as well as on the total propagation time in the spirit of the Lieb-Robinson theorem [42], the actual scaling of the computation time with n may be larger than the scaling $\mathcal{O}(n \log n)$ of the number of SVDs in the divide-and-conquer algorithm. Benchmarking the numerical run times, as is done in Sec. III, is therefore essential to assess the overall scaling of the algorithm for typical scenarios in open quantum systems.

D. Finite n_c and periodic process tensors

A useful feature of the sequential algorithm is that it profits significantly from memory truncation. If the memory of the environment becomes negligible after n_c time steps, the number of overlapping columns in subsequent rows in the tensor network is limited to $n_c - 1$, as can be seen, e.g., in the three bottom rows of Fig. 1(a). Only in the overlap region do MPO compression sweeps have to be performed, reducing the number of SVDs to $\mathcal{O}(nn_c)$. This suggests that it is also worthwhile to investigate how memory truncation can be incorporated into the divide-and-conquer algorithm.

To this end, consider an intermediate step in the divide-and-conquer algorithm where a block of $2n_c$ rows is formed by two blocks of n_c rows, as depicted in Fig. 3(b) for $n_c = 4$. In principle, following the same rationale as in the sequential algorithm, sweeping only over the finite overlap of $n_c - 1$ columns would produce an algorithm scaling as $\mathcal{O}(n_c \log n)$. Notably, for $n_c \ll n$, this means that obtaining the total PT involves fewer than n SVDs, which implies that many MPO matrices of the final PT must be exact copies of others.

In fact, one can identify a structure in PTs that can be repeated indefinitely. Such an observation is analogous to that made in introducing repeating tensors that represent the state of spatially infinite systems with translational invariance, in algorithms such as infinite density-matrix renormalization group (iDMRG) [61], infinite time-evolving block decimation (iTEBD) [62], and infinite projected entangled pair states (iPEPS) [63]. We now consider how this can be constructed for the periodic PT. Consider again the extension of the tensor network from n_c to $2n_c$ rows. For this step, the iteration in Eq. (14) becomes

$$\tilde{\mathcal{F}}_{3n_c-1:2n_c}^{\alpha_{3n_c-1}\dots\alpha_1} = \tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_{3n_c-1}\dots\alpha_{n_c+1}} \tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_{2n_c-1}\dots\alpha_1}, \quad (20)$$

as shown in Fig. 3(c). This result has explicit elements:

$$\begin{aligned} & [\tilde{\mathcal{F}}_{3n_c-1:2n_c}^{\alpha_l}]_{(d_l, d'_l), (d_{l-1}, d'_{l-1})} \\ &= \begin{cases} [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l}]_{d_l, d_{l-1}} \delta_{d'_l, 1} \delta_{d'_{l-1}, 1} & l \leq n_c \\ [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l}]_{d_l, d_{l-1}} [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_{l-n_c}}]_{d'_l, d'_{l-1}} & n_c < l \leq 2n_c \\ [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_{l-n_c}}]_{d'_l, d'_{l-1}} \delta_{d_l, 1} \delta_{d_{l-1}, 1} & l > 2n_c, \end{cases} \quad (21) \end{aligned}$$

where we define $[\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l}]_{d_{n_c}, d_{n_c-1}} = \delta_{d_{n_c}, d_{n_c-1}}$ to artificially extend the overlap from $n_c - 1$ to n_c columns. This process achieves a particular partitioning, where the first and last cases in Eq. (21) correspond to the areas shaded in blue and green, respectively, in Figs. 3(b) and 3(c). If shifted and put together, they exactly reproduce the original block $\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_{2n_c-1}\dots\alpha_1}$ consisting of the bottom n_c rows.

The role of the middle section in Eq. (21), visualized as an orange box in Figs. 3(b) and 3(c), becomes clear when the tensor network is further extended by another n_c rows,




FIG. 3. (a) Tensor network with many time steps, $n \gg n_c$, decomposed into blocks of length n_c . The parts of the tensor network highlighted in orange are identical. (b) First $2n_c$ rows of the tensor network subdivided into four blocks, where blocks of the same color are identical. The combination of a blue block stacked on top of a green block forms an orange block, which can be used as a fundamental unit of a periodic PT and can then be repeated indefinitely, as depicted in diagram (a). This result is due to the fact that the left and right interfaces of the orange block, marked by the light blue arrow in diagram (c), seamlessly link together, even after MPO compression of the first n_c rows.

as in Fig. 3(a). The overlap region with the new block of rows again contains n_c columns; however, these are exact copies of the (orange) middle section of the MPO, so they do not have to be calculated anew. Note that the fact that multiple central (orange) blocks fit seamlessly together is nontrivial, as the inner bonds between the subsequent matrices in the MPO have been strongly modified by MPO compression, so their relation to the bonds in the original tensor network is no longer obvious. In particular,

inner bonds at different positions within the MPO generally have different dimensions. Here, however, we have constructed the central (orange) blocks in such a way that the left and right bonds, as shown by the links marked by the light blue arrow in Fig. 3(c), correspond exactly to the bonds between the first (blue) and last (green) sections in Eq. (21). This process allows us to repeat the orange block indefinitely, and we arrive at the periodic PT $\mathcal{F}^{\alpha_n, \dots, \alpha_1} = \sum_{d_n, \dots, d_1} q_{d_n} \prod_{l=1}^n \mathcal{Q}_{d_l, d_{l-1}}^{\alpha_l}$, with

$$\mathcal{Q}_{(d'_l, d''_l), (d'_{l-1}, d''_{l-1})}^{\alpha_l} = \begin{cases} [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l}]_{d'_l, d''_l} \delta_{d''_{l-1}, 1} \delta_{d'_{l-1}, 1} & l \leq n_c \\ [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l \pmod{2n_c}}]_{d'_l, d''_l} [\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l \pmod{2n_c}}]_{d'_{l-1}, d''_{l-1}} & l > n_c, \end{cases} \quad (22)$$

where again the matrices $[\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_l}]_{d'_l, d''_l}$ are those of the MPO $\tilde{\mathcal{F}}_{2n_c-1:n_c}^{\alpha_{2n_c-1}, \dots, \alpha_1}$ describing the bottom n_c lines of the tensor network. In practice, it is useful to compress the periodic part of the PT-MPO once more in a final step using our preselection scheme for combining MPOs with large inner dimensions, described in Eq. (16). In doing so, care must be taken not to modify the left and right interfaces.

To summarize this section, if the memory of a Gaussian environment is finite, a periodic PT can be obtained using only $\mathcal{O}(n_c \log n_c)$ rank-reducing SVDs. Remarkably, this result is constant in the overall propagation time $t_e = n\Delta t$ and again provides a nominal scaling advantage over the sequential algorithm with $\mathcal{O}(nn_c)$. A further advantage of the periodic PT is that the memory requirements for storing it are $\mathcal{O}(n_c)$, much smaller than those for storing a full PT $\mathcal{O}(n)$.

In the following sections, we will showcase the power of our divide-and-conquer scheme with and without memory truncation on a series of applications.

III. APPLICATIONS

A. Fluorescence spectra of quantum dots

1. Model and context

Self-assembled III-V semiconductor QDs are key elements for photonic quantum technologies [64]. Because of their strong interaction with light, they can be used as bright sources of pure single photons [8,65], entangled photon pairs [66,67], and other nonclassical multiphoton states [68,69]. However, electronic excitations in QDs also strongly interact with longitudinal acoustic phonons. The resulting non-Markovian effects are significant, so they typically cannot be fully described by weak-coupling master equations [70]. While polaron master equations accurately account for phonon effects in the limit of weak driving [16], these break down under strong-driving conditions [70]. Therefore, numerically exact path-integral techniques based on QUAPI [43,44] have been applied [8,71] to investigate

the dynamics of driven QDs and other dissipative few-level systems [72].

While path-integral techniques have allowed calculations of some features of quantum dots, a numerically exact calculation of the spectra of strongly driven QDs is difficult. Multitime correlations (as needed to obtain emission spectra) can be calculated with such approaches [73–75], but reaching convergence remains a challenge because of a separation of timescales (see Appendix B for a convergence study): Because of the long radiative lifetimes in QDs, the dynamics has to be propagated to a few nanoseconds in order to avoid artifacts in the Fourier-transformed spectrum. Typical phonon memory times span several picoseconds. When considering driving with high laser intensities, or strongly off-resonant driving [76], the resulting oscillatory dynamics makes it necessary to choose small time steps Δt of the order of a few femtoseconds. Taken together, these separate timescales yield a problem where both the number of time steps n and the memory cutoff n_c must be large to produce accurate results. Moreover, to analyze phonon sidebands, spectra are often presented on a logarithmic scale, which makes numerical errors quite prominent. One should therefore use very small MPO compression thresholds in PT simulations, resulting in sizable inner dimensions χ . Nevertheless, for the reasons explained above, this challenging problem can be addressed using our divide-and-conquer scheme, in particular, when combined with the use of a periodic PT.

In Fig. 4, we present the fluorescence spectra of a QD driven with Rabi frequency Ω and coupled to a bath of phonons, showing how the Mollow triplet [77] is affected by this bath. The phonon-free system evolution with driving and radiative decay with rate κ is given by the Lindblad master equation

$$\frac{\partial}{\partial t} \rho = \mathcal{L}_S \rho = -(i/\hbar)[H_S, \rho] + \kappa \mathcal{D}_{\sigma^-}(\rho), \quad (23)$$

with system Hamiltonian $H_S = \frac{\hbar}{2} \Omega \sigma^x$ and Lindbladian

$$\mathcal{D}_A(\rho) = A\rho A^\dagger - \frac{1}{2}(A^\dagger A\rho + \rho A^\dagger A). \quad (24)$$

Here, the QD is modeled as a two-level system with ground and exciton states $|g\rangle$ and $|e\rangle$, respectively, and we use the conventional notation for operators $\sigma^+ = |e\rangle\langle g|$, $\sigma^- = |g\rangle\langle e|$, and $\sigma^x = \sigma^+ + \sigma^-$. The spectrum is obtained as the Fourier transform of the two-time correlations:

$$S(\omega) = \lim_{t \rightarrow \infty} \lim_{\tau_\infty \rightarrow \infty} \text{Re} \left\{ \int_0^\infty d\tau [\langle \sigma^+(t+\tau) \sigma^-(t) \rangle - \langle \sigma^+(t+\tau_\infty) \sigma^-(t) \rangle] e^{-i\omega\tau} \right\}, \quad (25)$$

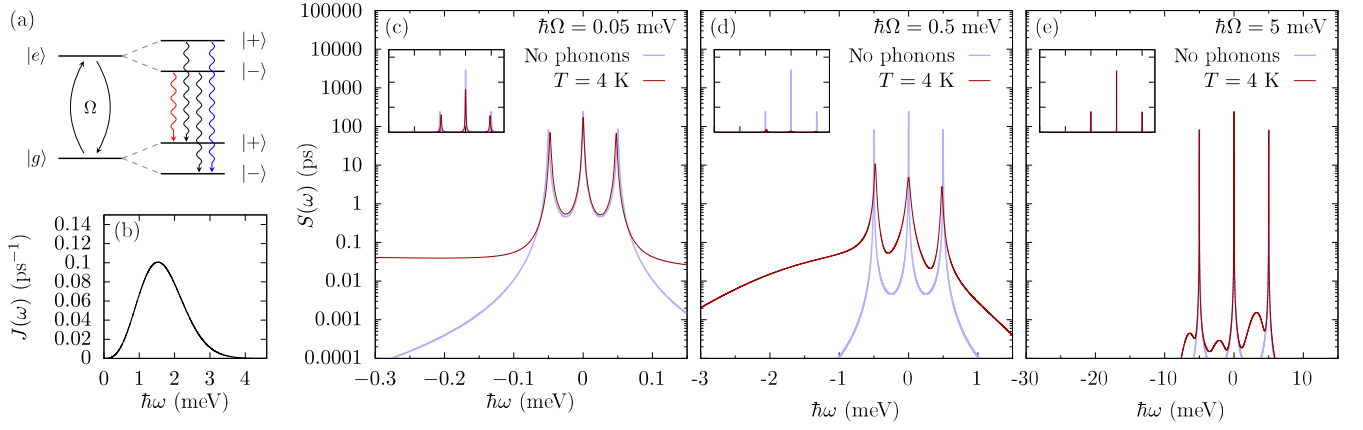


FIG. 4. (a) Cartoon of fluorescence experiment. The external laser drive dresses the ground and excited states of the QD. In the absence of a phonon environment, photon emission can be understood in terms of optical transitions between laser-dressed states. Black arrows indicate transitions contributing to photons at the central frequency of the two-level system, whereas red and blue arrows depict transitions detuned by $-\Omega$ and Ω from the central frequency, respectively. (b) Super-Ohmic phonon spectral density of a QD with electron radius $a_e = 4$ nm. (c)–(e) Fluorescence spectra of driven QDs with different driving strengths Ω , shown on a logarithmic intensity scale. Insets depict the same spectra on a linear scale. The simulations were performed over $n = 2^{21}$ time steps of size $\Delta t = 0.01$ ps using periodic PTs with a memory cutoff at $n_c = 2^{11}$ time steps.

where the coherent scattering contribution (elastic peak) has been subtracted.

For comparison, the spectra for simulations without phonons—as would be relevant for a driven atom—are shown as light blue lines in Figs. 4(c)–4(e) for different Rabi frequencies Ω and fixed radiative decay rate $\kappa = (0.5 \text{ ns})^{-1}$. In this case, one sees the characteristic Mollow triplet [77] with a central peak at the two-level transition frequency and two side peaks at frequencies $\pm\Omega$. The ratio between the heights of the central peak and each side peak is 3:1, while the ratio between the integrated areas is 2:1 [78]. The shape of the atomic Mollow triplet is explained in terms of laser-dressed states, i.e., eigenstates $|\pm\rangle = (|g\rangle \pm |e\rangle)/\sqrt{2}$ of the Hamiltonian H_S with eigenvalues $\pm\hbar\Omega/2$. Optical transitions $|\rightarrow\rangle \rightarrow |\rightarrow\rangle$ and $|\rightarrow\rangle \rightarrow |\rightarrow\rangle$ both contribute to the central peak at the two-level transition frequency, while transitions $|\rightarrow\rangle \rightarrow |\leftarrow\rangle$ and $|\leftarrow\rangle \rightarrow |\rightarrow\rangle$ are responsible for the side peaks at Ω and $-\Omega$, respectively.

2. Simulation results and discussion

In solid-state QDs, emission spectra are strongly affected by interactions with phonons. The QD-phonon interaction is of the form of a spin-boson model, corresponding to Eq. (1) with a coupling operator $\hat{O} = |e\rangle\langle e|$. The phonon spectral density takes the super-Ohmic form

$$J(\omega) = \frac{\omega^3}{4\pi^2 \rho \hbar c_s^5} (D_e e^{-\omega^2 a_e^2 / (4c_s^2)} - D_h e^{-\omega^2 a_h^2 / (4c_s^2)})^2. \quad (26)$$

For a GaAs-based QD, we take parameters given in Ref. [79] corresponding to electron and hole radii $a_e = 4$ nm and $a_h = a_e/1.15$, respectively, and we show

the resulting spectral density in Fig. 4(b). The corresponding polaron shift or reorganization energy $\int_0^\infty d\omega J(\omega)/\omega \approx 0.072$ meV is absorbed into the definition of the excited-state energy. The remaining physical and convergence parameters for our simulations using divide-and-conquer combined with periodic PTs are the initial phonon temperature $T = 4$ K, time discretization $\Delta t = 0.01$ ps, memory time $t_{\text{mem}} = n_c \Delta t$ with $n_c = 2048$, total propagation time $t_e = n \Delta t$ with $n = 2^{21} \approx 2 \times 10^6$, and MPO compression threshold $\epsilon = 10^{-12}$.

Previous works have addressed QD fluorescence spectra in the regime of weak Rabi driving using polaron master equations [15,80], or path-integral calculations [75]. There, two effects have been identified: First, compared to the phonon-free spectra, simulations with phonons reveal a line broadening, which also reduces the heights of the peaks. Second, the side peaks are shifted towards the center due to phonon renormalization of the transition dipole. These effects can be seen in Fig. 4(c), where we show the weak-driving case of $\hbar\Omega = 0.05$ meV.

Importantly, the ability of our divide-and-conquer algorithm to treat problems with many time steps enables us to also investigate the regime of stronger driving, beyond the range of previous work. In Fig. 4(d), we see that on increasing the driving strength to $\hbar\Omega = 0.5$ meV, three changes occur: The spectral lines are broadened, an asymmetric background arises, and there is a notable change in the relative height of the three Mollow peaks, with the low-energy side peak now dominating the emission. We discuss each of these features in turn. The broadening seen is consistent with behavior known in the weak-coupling limit [16], where the linewidth is proportional to the spectral density $J(\omega)$ evaluated at the Rabi frequency $\omega = \Omega$. Since $J(\Omega)$ is significant at

$\hbar\Omega = 0.5$ meV—see Fig. 4(b)—the broadening here becomes large. Regarding the asymmetric background, this feature is similar to that seen in the phonon sideband in an undriven QD [81]. Finally, the relative peak heights can be understood as the result of fast thermalization in the dressed-state basis. Thermalization predominantly occupies the lower dressed state $|-\rangle$ because the energetic splitting between $|+\rangle$ and $|-\rangle$ is larger than the thermal energy $\hbar\Omega > k_B T \approx 0.34$ meV. Hence, emission from $|+\rangle$ is quenched. Moreover, the direct transition $|-\rangle \rightarrow |-\rangle$ emitting photons at the frequency of the central peak also competes with phonon-assisted photoemission processes, where an energy $\hbar\Omega$ is efficiently absorbed into the phonon bath. The photon with the remaining energy contributes instead to the lower energy side peak.

Considering even larger driving with $\hbar\Omega = 5$ meV, as depicted in Fig. 4(e), we find a restoration of the characteristic Mollow triplet with a 3:1 ratio between the heights of central and side peaks just like in the phonon-free simulations. This result is due to dynamical decoupling from phonons [5,82,83], which occurs because the spectral density $J(\omega)$ becomes small for frequencies $\omega \gtrsim 5$ meV/ \hbar ; see Fig. 4(b).

3. Scaling of computation time

We next discuss the computational cost of these calculations for the various different algorithms presented above. In Fig. 5, we show the computation time, defined as the total elapsed time from the start to the end of the program on a conventional laptop computer with an Intel Core i5-8265U processor.

For the sequential algorithm by Jørgensen and Pollock [49], we initially see the expected superlinear increase in computation time up to about $n \sim 1000$ time steps, with a scaling compatible with the $\mathcal{O}(n^2)$ behavior of the number of SVDs in the algorithm [cf. dotted and dashed lines indicating slopes corresponding to $\mathcal{O}(n^2)$ and $\mathcal{O}(n)$ scaling, respectively]. For larger n , where $n > n_c = 2048$, the scaling switches to linear, which again matches the expectation of $\mathcal{O}(nn_c)$ SVDs. The sequential algorithm data are limited to $n \lesssim 10\,000$ time steps due to the computation time required beyond this.

The divide-and-conquer scheme similarly shows an initial superlinear scaling for $n \ll n_c$, but it is approximately 1 order of magnitude faster than the sequential algorithm after $n \sim 100$ time steps. After the kink at $n \sim n_c$, this algorithm scales more slowly, with a slope on the double-logarithmic scale that indicates sublinear behavior.

As discussed in Sec. II, it can be beneficial to use different compression thresholds for selection, backward, and forward sweeps. The computation time for simulations with threshold ratio $r_b = r_s = 0.2$ is also shown in the figure and is seen to be a factor of 3 smaller than that with $r_b = r_s = 1$. We have checked that this change corresponds to a reduction of the maximal inner bond dimension of the

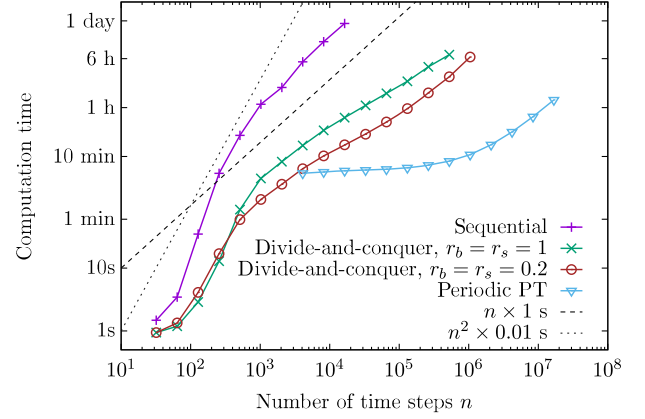


FIG. 5. Computation time as a function of the total number of simulation time steps n , on a double-logarithmic scale. Results are shown for the sequential algorithm, the divide-and-conquer scheme with different ratios $r_b = r_s$ between compression thresholds for backward and forward sweeps, and the periodic PT approach starting from a divide-and-conquer calculation with $r_b = r_s = 0.2$. For reference, linear and quadratic scaling with n are depicted as dashed dark-gray and dotted light-gray lines, respectively.

final PT from $\chi = 158$ to $\chi = 133$. When studying accuracy vs the forward compression threshold ϵ , we find similar accuracies for both values of threshold ratios. These calculations allow us to reach $n \sim 1\,000\,000$ time steps; the ultimate limitation in this case is the storage of the full PT-MPO, which requires over 400 gigabytes, rather than the computation time.

Considering the periodic PT algorithm, we show results from switching to periodic PTs after calculating an intermediate PT of n_c rows using the divide-and-conquer scheme with $r_b = r_s = 0.2$. As anticipated, in this case, the computation time remains constant until about $n \sim 1\,000\,000$ time steps because no more numerical resources have to be spent on PT calculation. At extremely large $n \sim 10^7$, the time becomes dominated by the propagation of the open quantum system in Eq. (8), where PT-MPO is contracted with the set of system propagators $\mathcal{M}_{a_i, a_{i-1}}$ via the outer indices and by the extraction of the observables. Even though this only involves matrix multiplications, which are much less demanding than SVDs, they have to be applied at each time step, so linear scaling arises for extremely large n . The prefactor of this linear scaling is, however, far smaller than that of the other algorithms in Fig. 5, as is clear from the vertical offset on the double-logarithmic graph. Storing the periodic PT requires less than 2 gigabytes, which eliminates the need for writing to and reading from a hard disk.

Finally, we want to point out that there are occasions where PT-MPOs can be employed with genuine sublinear scaling with respect to n . For example, if one is interested in the stationary states of an open quantum system with time-independent or periodic system Hamiltonians, one can

contract the periodic PT with the corresponding system propagators. Thereby, one obtains an effective propagator for the system expanded by the inner dimension of the PT-MPO at the periodic cut, which can be diagonalized. The stationary state(s) is then given by the eigenvector(s) of the effective propagator corresponding to eigenvalue(s) with value 1 (see Ref. [43] for a similar discussion of diagonalizing the QUAPI tensor). In addition, multitime correlations can be calculated directly from considering the other eigenvalues and eigenvectors. Moreover, for time-independent or periodic system Hamiltonians, it would be particularly useful to generate transfer tensor [45] or SMatPI [46] small matrix propagators from PT-MPOs. This would enable a significant reduction in the prefactor of the linear term in the runtime because the propagation then involves only matrix-vector multiplications with dimensions D^2 instead of $D^2\chi$, and thus it would extend the sublinear regime to even longer overall propagation times. On the other hand, the inner bonds of PT-MPOs carry information needed to describe environment responses to more general interventions on the system [2,48], which makes periodic PT-MPOs particularly useful for investigations of systems under pulsed driving [84,85] and multi-environment systems [13,56].

B. Superradiance

1. Model and context

Superradiance is a dramatic consequence of collective quantum behavior [86,87], where N emitters act as one, resulting in spontaneous emission rates that scale super-extensively with N . It occurs when the interemitter distances are much smaller than the wavelength of the emitted

light, making the emitters spatially indistinguishable [see Fig. 6(a)], and when the energies of emitters are the same, making them spectrally indistinguishable. In this case, if all emitters are prepared in their excited state at $t = 0$, the emitted intensity features a characteristic burst at short times and a slowly decaying tail at long times, in distinction to the exponential decay that occurs for independent or distinguishable emitters [87]. Concomitant with superradiance is the existence of subradiant states, which are optically inactive and so can be used to store excitations or quantum information over long periods of time. Because of recent advances in fabrication, it has become possible to explore collective effects of a few spectrally indistinguishable solid-state QDs, including coupling multiple QDs into the same photonic waveguide [88–91]. As in the previous section, consideration of semiconductor QDs raises the issue of phonon effects, and a key theoretical question is thus modeling of imperfections and environmental effects in real-world solid-state devices [13].

The standard model to consider for superradiance is N two-level systems coupled to the electromagnetic field via an electric dipole interaction. We can thus write the full Hamiltonian (i.e., $H_S + H_E$) as [86,87]

$$H = \sum_{j=1}^N \hbar\omega_j \sigma_j^+ \sigma_j^- + \sum_k \hbar\omega_k a_k^\dagger a_k + \sum_{j,k} \hbar(g_{j,k} a_k^\dagger + g_{j,k}^* a_k)(\sigma_j^- + \sigma_j^+), \quad (27)$$

where ω_j is the transition frequency of emitter j , ω_k is the frequency of photon mode k , $g_{j,k}$ are the light-matter

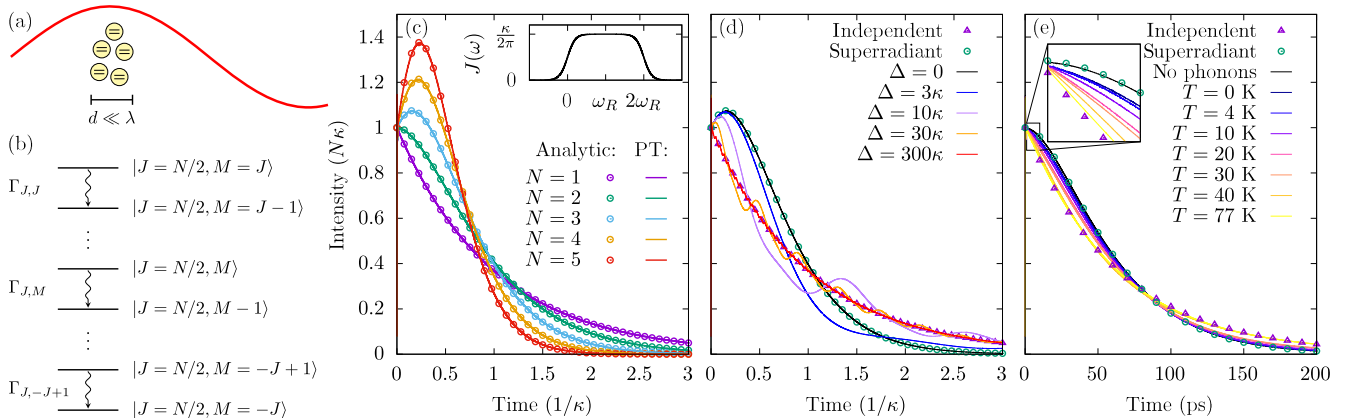


FIG. 6. (a) Cartoon showing $N = 5$ two-level quantum emitters confined to a region much smaller than the wavelength of the emitted light. If the emitters are also spectrally indistinguishable, they effectively couple to the electromagnetic field as a single emitter, resulting in superradiant emission. (b) Superradiant transitions between states in the Dicke ladder. (c) Emitted intensity vs time for $N = 1, \dots, 5$ superradiant identical emitters. Points show analytic results of rate equations (see text); lines represent divide-and-conquer PT simulations without memory truncation. The inset shows the spectral density used to model radiative decay at rate κ . This spectral density is flat around the central frequency ω_R with smooth edges. (d) Breakdown of superradiance for $N = 3$ emitters upon increasing spectral distinguishability. The transition frequencies of the emitters are detuned by $-\Delta/2$, 0 , and $\Delta/2$, respectively, from the central frequency. (e) Effect of phonon baths on superradiance for $N = 2$ degenerate quantum dots, for a variety of temperatures T .

coupling constants, a_k^\dagger and a_k create and destroy a photon in mode k , and σ_j^\dagger and σ_j^- describe the excitation and deexcitation of the j th emitter. Throughout this example, we assume $g_{j,k} = g_k$ is real and independent of j , which implies that all emitters couple to the photon environment with the same phase, as is crucial for superradiance [86,87].

The Hamiltonian written above is a multimode generalization of the Dicke model, without any rotating wave approximation. As such, the time evolution of this model will involve high frequencies, associated with the characteristic frequencies ω_j , ω_k . In many cases, it is preferable to make a RWA, yielding the Tavis-Cummings model, which would then allow one to shift to a frame in which these high frequencies are eliminated. However, doing so changes the system-environment coupling in a way that precludes the direct application of the PT-MPO formalism discussed above: After the RWA, the system-environment coupling no longer takes the simple product form in Eq. (1), and it instead involves two different system operators coupling to the environment. While approaches exist to construct tensor network representations for the evolution of such problems [92], they involve tensors with larger internal dimensions. However, our divide-and-conquer scheme enables us to resolve very many time steps. As such, we are able to calculate a PT-MPO for the photonic environment with the original electric dipole coupling in Eq. (27) without making the RWA. In Appendix C, we discuss an alternative perspective, starting from the Tavis-Cummings model (with the RWA) and reintroducing counterrotating terms as an approximation, controlled by adding a common carrier frequency ω_R to the system frequencies ω_j , ω_k .

2. Comparison of numerical and analytic results

Starting from Eq. (27), we directly apply our algorithm to the environment Hamiltonian $H_E = H - H_S$ with system Hamiltonian $H_S = \sum_j \omega_j \sigma_j^\dagger \sigma_j^-$. The form of the light-matter coupling is advantageous, as degeneracies allow us to treat relatively large N exactly. Specifically, $\sum_j \sigma_j^\dagger$ has a spectrum with only $N + 1$ different eigenvalues $-N, -N - 2, \dots, N$, allowing us to construct PTs with only $(N + 1)^2$ instead of 2^{2N} different outer indices α [52,71]. As a result, the PT-MPO calculation is so efficient that it is the final contraction—describing the propagation of the N -emitter system in a 2^{2N} -dimensional space—rather than the PT calculation that currently limits the number of emitters N in our numerically exact simulations.

The spectral density used for the PT-MPO calculation, shown in the inset to Fig. 6(c), is chosen as

$$J(\omega) = \frac{\kappa}{2\pi} \frac{1}{1 + e^{-\omega/w}} \frac{1}{1 + e^{(2\omega_R - \omega)/w}}. \quad (28)$$

This expression has a wide, flat region around $\omega \approx \omega_R$ with a constant value $\kappa/(2\pi)$, so in the Markov limit, a single

emitter would spontaneously decay with rate κ . The edges of this box-shaped spectral density are rounded using logistic functions centered at $\omega = 0$ and $\omega = 2\omega_R$, respectively, with transition widths $w = \omega_R/10$. We find such rounding produces PT-MPOs with smaller inner dimensions than for a sharp cutoff. Throughout this discussion, we choose $\omega_R = 1000\kappa$, time steps $\Delta t = 1/(8192\kappa)$, total number of time steps $n = t_e/\Delta t = 32768$, and MPO compression threshold $\epsilon = 10^{-9}$. The divide-and-conquer algorithm is applied without memory truncation. The emitted intensity $I(t) = -(\partial/\partial t) \sum_j n_j(t)$ is extracted using central differences $I(t) \approx -\sum_j [n_j(t + \Delta t) - n_j(t - \Delta t)]/(2\Delta t)$, where $n_j(t)$ denotes the occupation of the j th emitter. The results of such calculations for identical emitters, with $\omega_j = \omega_R$, are shown by solid lines in Fig. 6(c).

For comparison to these simulations, Fig. 6 also shows analytic results. These results are found using the standard approach of considering transitions between the ladder of Dicke states, as depicted in Fig. 6(b), and considering the result in the RWA (i.e., Tavis-Cummings model). The validity of the RWA is controlled by the size of ω_R compared to other parameters. The Dicke states $|J, M\rangle$ are eigenstates of the collective angular momentum operator $\hat{J}_{x/y/z} = \sum_{j=1}^N \sigma_j^{x/y/z}$, specifically, the total angular momentum $\hat{J}^2 |J, M\rangle = J(J + 1) |J, M\rangle$ and its z component $\hat{J}_z |J, M\rangle = M |J, M\rangle$. Starting from the fully excited state, $|J = N/2, M = N/2\rangle$, the dynamics is restricted to the states with $J = N/2$ and $M = -J, \dots, J$. The light-matter interaction involves the collective lowering operator \hat{J}^- , which gives rise to optical transitions between states $|J, M\rangle$ and $|J, M - 1\rangle$. The rates $\Gamma_{J,M}$ for these transitions can be found using Fermi's golden rule, which yields $\Gamma_{J,M} = \kappa(J + M)(J - M + 1)$ [87]. The sequential photon emission processes can be solved analytically; the solutions are listed explicitly in Appendix D and are shown by points in Fig. 6(c).

We see in Fig. 6(c) that the emitted intensities from $N = 1, \dots, 5$ indistinguishable ($\omega_j = \omega_R$ for all j) quantum emitters obtained from the PT-MPO simulations very closely match the analytical results. As noted above, the PT-MPO calculations are for the Dicke model (without RWA), while the analytic results are for the Tavis-Cummings model, with the RWA in the light-matter coupling; this confirms the RWA would be valid for these parameters. While not clearly visible in the figure, there are discrepancies between the two calculations at early times, $t \lesssim 1/\omega_R$. These differences are due to the difference of models with and without the RWA: There are transient effects of turning on the matter-light coupling, including the counterrotating terms at $t = 0$. We have checked that, as expected, such effects can be reduced by increasing ω_R (not shown).

3. Effects of distinguishability and decoherence

While the results above demonstrate the ability of the PT-MPO approach to recover superradiant dynamics for the

Dicke model, the fact that all results were recoverable from analytic calculations suggests such an approach was not needed. However, the PT-MPO approach includes the full description of environment influences and so can be used in situations where no analytic solutions are available. Here, we use it to present results on how superradiance is destroyed by spectral distinguishability and decoherence.

Figure 6(d) shows the breakdown of superradiance for $N = 3$ quantum emitters when they become spectrally distinguishable. This breakdown is implemented by changing the system Hamiltonian by setting different values for the two-level transition frequencies $\omega_1 = \omega_R - \Delta/2$, $\omega_2 = \omega_R$, and $\omega_3 = \omega_R + \Delta/2$. We note that this can be done with a single calculation of the PT-MPO and by just changing the system propagator with which it is then contracted. Indeed, with increasing frequency detuning between emitters, Δ , the superradiant intensity burst is suppressed and replaced by oscillations around an exponentially decaying intensity. At sufficiently large Δ , one recovers monoexponential decay with rate κ as is expected for independent emitters.

The PT-MPO approach also allows us to investigate how superradiance is affected when the emitters are additionally coupled to other environments, such as phonons. These can affect the emission dynamics by dephasing interemitter coherences and by accumulating which-path information, i.e., which of the emitters is excited and which is not. In Fig. 6(e), we show the corresponding intensity for $N = 2$ spectrally indistinguishable semiconductor QDs ($\omega_j = \omega_R$), with each QD coupled to a bath of longitudinal acoustic phonons. For this bath, we use the same spectral density and parameters as in Fig. 4; we thus present results in physics units and require a specific choice of $\kappa = 1/64 \text{ ps}^{-1}$. We find that superradiant emission is already slightly suppressed due to QD-phonon interactions for baths at initial temperature $T = 0 \text{ K}$; similar behavior persists up to temperatures of $T = 4 \text{ K}$. Signatures of superradiance start to significantly decrease from about $T = 10 \text{ K}$ and almost vanish at liquid-nitrogen temperature $T = 77 \text{ K}$.

C. Coherence decay with a strongly peaked spectral density

1. Motivation and model

To determine the practical limits of what our divide-and-conquer approach can achieve, we consider strong coupling to a bath with strongly peaked spectral density. Such environments can be addressed by approximate methods such as the reaction coordinate approach [93,94] or by specialized numerically exact techniques like hierarchical equations of motion (HEOM) [33]. However, such environments have been challenging for PT-MPO-based approaches due to long memory times and correspondingly large inner PT-MPO dimensions. As such, they can serve as a critical test for new PT-MPO algorithms.

To explore the behavior of such models, we consider the free decay dynamics of the spin-boson model. In other words, we consider a two-level system, with vanishing system Hamiltonian $H_S = 0$ and system-environment coupling operator $\hat{O} = |e\rangle\langle e|$. We then consider the decay of initially prepared coherences, i.e., the time evolution of $\langle \sigma_x(t) \rangle$ for initial state $\bar{\rho}(0) = \frac{1}{2}(|g\rangle + |e\rangle)(\langle g| + \langle e|)$. This setup is particularly well suited for studying the convergence, as the coherences react much more sensitively to the environment than occupations. Furthermore, for $H_S = 0$, the model reduces to the independent boson model, for which analytical solutions are available for comparison. Moreover, as $[H_S, H_E] = 0$, the Trotter error, i.e., the error due to the finiteness of the time step Δt , vanishes.

We take the spectral density to have the form

$$J_{\text{QDT}}(\omega) = \frac{\eta\omega\Omega^4}{(\Omega^2 - \omega^2)^2 + 4\omega^2\gamma^2} \quad (29)$$

defined by interaction strength η , central frequency Ω , and damping γ . This form of spectral density is commonly used to describe quantum dissipative tunneling [18,93], such as in charge or excitation transfer in biological or chemical molecular systems. It can be derived from a hierarchical model, where the two-level system couples to a single nuclear coordinate with frequency Ω , which in turn is damped by an Ohmic bath [93]. In this interpretation, η is set by the Ohmic damping strength, and $\gamma = \eta/(2M)$, where M is the nuclear mass. Throughout this example, we set $\hbar = 1$ and use dimensionless parameters $\Omega = 10$ and $\gamma = 1$, while we vary the coupling strength η . Note that $J_{\text{QDT}}(\omega)$ is depicted in the inset of Fig. 7(a). The initial bath temperature is set to $T = 0$.

2. Numerical results and computation time

Figures 7(a) and 7(b) show the time evolution of the coherence, calculated using the sequential and the divide-and-conquer algorithms, for $\eta = 0.01$, $\Delta t = 1/32$ and $\eta = 0.1$, $\Delta t = 1/128$, respectively. For the case $H_S = 0$ that we consider, exact results can be obtained by polaron transformation [16],

$$\langle \sigma_x \rangle = \exp \left\{ \int_0^\infty d\omega \frac{J(\omega)}{\omega^2} (\cos \omega t - 1) \coth(\beta\omega/2) \right\}, \quad (30)$$

where $\coth(\beta\omega/2) \rightarrow 1$ for zero temperature $T \rightarrow 0$. These results are shown for comparison in Fig. 7. For the weaker coupling, we see that both methods are capable of recovering the exact result when the compression threshold ϵ is small enough. For the stronger coupling strength, convergence is not reached, as discussed further below.

Of particular interest here is how the computation time for the construction of the PT-MPOs scales as a function of

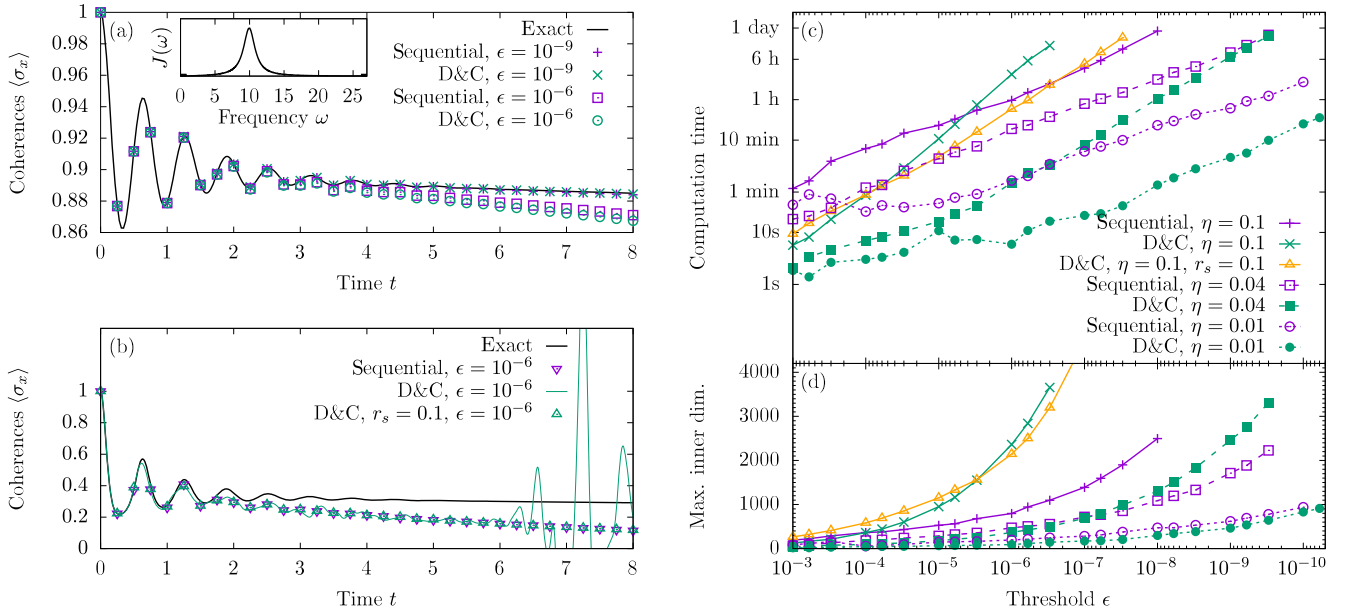


FIG. 7. (a) Free decay of coherences in a two-level system with spectral density $J_{\text{QDT}}(\omega)$ [cf. Eq. (29)] for parameters $\eta = 0.01$, $\Omega = 10$, $\gamma = 1$. Note that $J_{\text{QDT}}(\omega)$ is shown in the inset. Exact results for the time evolution of coherences are shown by the solid black line, and results for the different PT-MPO algorithms are indicated by points. Although an internal time step $\Delta t = 1/32$ is used, data points are only shown at times $t_j = (1/4)j$ for clarity. (b) Time evolution for $\eta = 0.1$ and $\Delta t = 1/128$, where the PT-MPO results are not converged. The divide-and-conquer algorithm with equal thresholds ϵ for forward and backward sweeps (shown, in this case, as a thin green solid line) converges less regularly than the sequential algorithm; this can be ameliorated by choosing a nonunity threshold ratio $r_s = 0.1$ in the preselection stage of the divide-and-conquer block combination. (c) Computation times of the sequential and the divide-and-conquer algorithms for different coupling strengths η as a function of the truncation threshold ϵ . (d) Maximal inner PT dimension encountered at intermediate steps in the respective algorithm.

the compression threshold ϵ . This case is presented in Fig. 7(c) for baths with coupling strengths $\eta = 0.01$, $\eta = 0.04$, and $\eta = 0.1$ and time steps $\Delta t = 1/32$. For weak coupling $\eta = 0.01$, we find the divide-and-conquer scheme to be about 1 order of magnitude faster than the sequential algorithm for all relevant compression thresholds. As shown in Fig. 7(a) for this coupling, both methods reproduce exact results for small thresholds $\epsilon = 10^{-9}$, while they incur comparable numerical errors for large thresholds $\epsilon = 10^{-6}$.

Increasing the coupling to $\eta = 0.04$, the computation time required for the divide-and-conquer scheme is still shorter than for the sequential algorithm when the threshold ϵ remains large enough. However, the divide-and-conquer computation time increases faster with decreasing threshold than the time for the sequential algorithm. This increase eliminates the computational advantage over the sequential approach at about $\epsilon \lesssim 10^{-9}$. This increase in computation time corresponds to an increase in the maximal inner dimension of the PT-MPO at intermediate steps of the algorithm, as is shown in Fig. 7(d). From this figure, we may note that the bond dimension resulting from the divide-and-conquer calculation is notably larger than for the sequential algorithm at the same cutoff. This suggests some inefficiency in how the truncation is applied in the

divide-and-conquer algorithm, so it may be possible to improve performance further with more sophisticated approaches to truncation [95,96].

Increasing the coupling further to $\eta = 0.1$, we again see a crossover in computation times on reducing ϵ , but in this case, it occurs at about $\epsilon = 3 \times 10^{-6}$. Once again, this increased computation time is associated with a larger maximal inner dimension. Moreover, for this coupling strength, Fig. 7(b) shows that the convergence behavior of the divide-and-conquer algorithm can be less regular than that of the sequential algorithm: We see large errors at late times due to the fact that the preselection of products of singular values does not provide the locally optimal low-rank approximation, as discussed in detail in Appendix A. There, we also demonstrate that this can be combatted by using a finer threshold for the selection process, characterized by the ratio $r_s = \epsilon_{\text{select}}/\epsilon_{\text{forward}}$. Setting $r_s = 0.1$ in simulations in Fig. 7(b), we see the results of the divide-and-conquer approach match those of the sequential algorithm at $\epsilon = 10^{-6}$. However, we note that neither algorithm has converged to the exact result for such a value of ϵ .

Moreover, we find that the choice of $r_s = 0.1$ can lead to an order-of-magnitude speedup of the divide-and-conquer algorithm for simulations with strong system-environment coupling $\eta = 0.1$ depicted in Fig. 7(c), rendering the

computation time of the divide-and-conquer scheme close to that of the sequential algorithm, even for the most challenging parameter regime with small truncation thresholds.

The identification of the main drawback of our divide-and-conquer algorithm suggests that even more challenging open quantum systems with strong system-environment coupling and very long memory times can be tackled by developing and implementing higher-quality low-rank approximations for the combination of PT-MPO blocks with large inner dimensions.

IV. SUMMARY

We have introduced an approach for numerically exact simulations of non-Markovian open quantum systems with a scaling advantage over established techniques. Most techniques that account for non-Markovian effects scale quadratically $\mathcal{O}(n^2)$ with the number of time steps n , or linearly $\mathcal{O}(nn_c)$ if the memory can be truncated after n_c time steps. In contrast, we arrive at a method that calculates the PT-MPO—a compact representation of the Feynman-Vernon influence functional—in $\mathcal{O}(n \log n)$ or $\mathcal{O}(n_c \log n_c)$ rank-reducing operations for situations without and with memory truncation, respectively. Once the PT-MPO is obtained, the system dynamics are simulated with linear complexity $\mathcal{O}(n)$ with a small prefactor. We achieve this enhancement by exploiting the high degree of self-similarity in the tensor network representing the influence functional, which enables two novel developments: a divide-and-conquer strategy to contract PT-MPOs with a large number of time steps, and the construction of periodic PTs with blocks that can be repeated indefinitely. To demonstrate that the theoretical scaling advantage can be realized in practice, we have applied our approach to several examples of challenging multiscale problems of technological interest and compared it with the established sequential PT-MPO calculation scheme introduced by Jørgensen and Pollock [49].

For calculations of the fluorescence spectrum of a semiconductor QD strongly coupled to a super-Ohmic phonon bath, we indeed find good agreement between the theoretically expected scaling behaviors. We observe sublinear scaling over a wide range of practically relevant propagation times with $n = 10^3$ to $n = 10^6$ time steps. The combination of divide-and-conquer with periodic PTs enables well-converged calculations of the resonance fluorescence spectra of driven QDs, which requires simulations of quantum dynamics over millions of time steps. Our method produces these spectra within minutes, whereas the current standard PT algorithm of Ref. [49] would require weeks to months to achieve the same results.

In the example of superradiant emitters, we demonstrate several additional aspects enabled by our algorithm: First, the possibility of resolving many time steps facilitates the treatment of problems in quantum optics beyond the RWA. The PT-MPO approach allows combination of multiple environments, which facilitates investigations of the

breakdown of superradiance with additional couplings to vibrational environments. Such numerically exact studies including effects of local environments are relevant for current efforts by several research groups to realize cooperative emission [88,97] in spectrally tunable semiconductor QDs [89–91], as well as in organic systems [98]. The ability of the divide-and-conquer scheme to explicitly model both optical decay and phonon dephasing promises tremendous acceleration of investigations of a whole class of challenging scenarios, for example, where photonic environments as well as phonon baths are strongly structured and nonadditive cross interactions between different environments are important [99,100].

While our first two examples clearly show the significant practical utility of our novel approach, a detailed analysis of performance hints at a potential disadvantage of the divide-and-conquer algorithm. It requires the combination of PTs with large inner dimensions. We have shown how to significantly reduce the cost associated with this by preselecting degrees of freedom based on SVDs of the individual PTs. This combination strategy, however, does not lead to a locally optimal rank reduction like a SVD of the combined PT, which can result in the appearance of a larger number of singular values above the compression threshold than are found for the PT-MPO in the sequential algorithm of Ref. [49]. Consequently, the inner dimension χ of the PT may be increased, and thereby the computation time needed to perform each SVD may also increase. This limitation of the divide-and-conquer approach is illustrated in our final example: free coherence decay with a model spectral density consisting of a narrow peak. This choice is motivated by the expectation of large inner dimensions χ , which are controlled by the bath coupling strength in such a case. Indeed, we find that on increasing the coupling strength, there is a cross-over from a regime where the divide-and-conquer algorithm outperforms the sequential algorithm to a regime where the sequential algorithm performs better. Simple strategies to mitigate the extra singular values appearing in the divide-and-conquer scheme, e.g., by using different compression thresholds at different steps in the algorithm, are promising.

Such numerical experiments support a picture that the optimal strategy for PT-MPO calculations is a balance between reducing the number of SVDs (or alternative rank-reducing operations) and reducing the numerical effort of each individual operation. The divide-and-conquer algorithm provides a scaling advantage for the number of operations but at the cost of increasing the effort of each operation. However, it should be noted that the increased inner PT dimension of the divide-and-conquer algorithm is not due to the divide-and-conquer strategy *per se*: It is a consequence of the suboptimal compression after joining two PT-MPOs with large inner dimensions. It would thus be worthwhile for future work to explore other strategies to combine large PT-MPOs, e.g., alternate criteria for selecting singular values, or other methods for low-rank matrix

approximations [95,96]. Alternatively, one may also construct approaches that interpolate between the sequential and the divide-and-conquer approaches: for example, merging fixed-size blocks of the tensor network instead of single lines in the sequential algorithm, or by sweeping and compressing PT-MPOs multiple times after combining into blocks and allowing a different compression threshold for each sweep. We also note that periodic PTs can be readily calculated starting from blocks obtained with the sequential algorithm. This could lead to a sizable reduction of computation time when propagating open quantum systems coupled to environments with finite memory times but very strong system-environment couplings.

Finally, it is noteworthy that, even though we have implemented the divide-and-conquer approach only for bosonic environments, as defined in Eq. (1), in principle, it can be extended to the more general class of Gaussian environments [92,99], which includes, e.g., fermionic environments of impurity problems [101,102]. Moreover, PTs describing environments coupled to small quantum systems can be reused for simulations for which the system of interest is coupled to another subsystem. Extending the system Hilbert space is accommodated simply by modifying the outer bonds of the PT-MPO [57,71]. Thus, the PT-MPO calculated in our first example can be readily employed to obtain spectra of QDs embedded in microcavities, which is another current topic of interest [103]. A variant of this approach was used in Ref. [13] to investigate cooperative emission from two QDs, each additionally strongly coupled to a local phonon environment, providing numerically exact predictions including possible cross interactions between different environments. Furthermore, Ref. [56] delivers a proof of principle for scaling up PT-MPO-based numerically exact approaches to small networks of open quantum systems. Given this wider context, progress in methods for constructing PT-MPOs has immediate implications for a large class of topical problems in open quantum systems.

Note added.— Recently, a related preprint [104] by Link *et al.* showed how periodic PTs can also be calculated using iTEBD methods.

The data presented in this article is available via [105].

ACKNOWLEDGMENTS

M. C. and E. M. G. acknowledge funding from EPSRC Grant No. EP/T01377X/1. B. W. L. and J. K. were supported by EPSRC Grant No. EP/T014032/1. We also thank Gerald Fux for fruitful discussions.

APPENDIX A: DIFFERENT THRESHOLDS FOR SELECTION, BACKWARD, AND FORWARD SWEEPS

We now discuss how choosing different thresholds for forward and backward sweeps as well as for the selection of

products of singular values can affect the performance of the divide-and-conquer algorithm. To this end, we first review the role of canonical forms for the optimality of matrix product state (MPS) compression [55,106]. We then discuss the combination based on selecting products of singular values before providing numerical examples.

1. Locally optimal MPS truncation

A general pure state $|\psi\rangle$ of a one-dimensional quantum system with L sites, each described by a Hilbert space of dimension d , can be expanded as

$$|\psi\rangle = \sum_{\sigma_1, \dots, \sigma_L} c_{\sigma_1, \dots, \sigma_L} |\sigma_1, \dots, \sigma_L\rangle, \quad (\text{A1})$$

where $|\sigma_1, \dots, \sigma_L\rangle$ are products of local basis states and $c_{\sigma_1, \dots, \sigma_L}$ are the d^L expansion coefficients. They can be exactly expressed as matrix products

$$c_{\sigma_1, \dots, \sigma_L} = \sum_{a_1=1}^{\chi_1} \dots \sum_{a_{L-1}=1}^{\chi_{L-1}} M_{a_0, a_1}^{\sigma_1} M_{a_1, a_2}^{\sigma_2} \dots M_{a_{L-1}, a_L}^{\sigma_L}, \quad (\text{A2})$$

where $a_0 = a_L = 1$ are dummy indices. The dimensions χ_L of indices a_l , in general, grow exponentially towards the center, e.g., $\chi_{L/2} \leq d^{L/2}$ for even L . To approximate the state $|\psi\rangle$ by an MPS $|\psi_{\text{trunc}}\rangle$ with truncated bond dimension D between sites l and $l+1$, it is advisable to first use the gauge freedom of the inner bonds to bring the MPS into the mixed-canonical form [55]

$$c_{\sigma_1, \dots, \sigma_L} = \sum_{a_1, \dots, a_{L-1}} A_{a_0, a_1}^{\sigma_1} A_{a_1, a_2}^{\sigma_2} \dots A_{a_{l-1}, a_l}^{\sigma_l} S_{a_l} B_{a_l, a_{l+1}}^{\sigma_{l+1}} B_{a_{l+1}, a_L}^{\sigma_L}, \quad (\text{A3})$$

with $\sum_{\sigma} A^{\sigma} A^{\sigma\dagger} = \mathbb{1}$ and $\sum_{\sigma} B^{\sigma} B^{\sigma\dagger} = \mathbb{1}$. With these conditions, the Schmidt decomposition at the bond between sites l and $l+1$ is given by

$$|\psi\rangle = \sum_{a_l=1}^{\chi_l} s_{a_l} |a_l\rangle_A |a_l\rangle_B, \quad (\text{A4})$$

with bases for the left and right blocks,

$$|a_l\rangle_A = \sum_{\sigma_1, \dots, \sigma_l} (A^{\sigma_1} \dots A^{\sigma_l})_{1, a_l} |\sigma_1, \dots, \sigma_l\rangle, \quad (\text{A5})$$

$$|a_l\rangle_B = \sum_{\sigma_{l+1}, \dots, \sigma_L} (B^{\sigma_{l+1}} \dots B^{\sigma_L})_{a_l, 1} |\sigma_{l+1}, \dots, \sigma_L\rangle, \quad (\text{A6})$$

respectively, where the orthonormality ${}_A\langle a_l | a'_l \rangle_A = {}_B\langle a_l | a'_l \rangle_B = \delta_{a_l, a'_l}$ of the Schmidt basis follows directly from the orthonormality of A^{σ_l} and B^{σ_l} . Given the Schmidt decomposition (A4), the locally optimal approximation to

$|\psi\rangle$ by an MPS with reduced dimension D , guaranteed by the Eckart-Young-Mirsky theorem [107], is obtained by neglecting the smallest singular values

$$|\psi_{\text{trunc}}\rangle = \sum_{a_l=1}^D s_{a_l} |a_l\rangle_A |a_l\rangle_B, \quad (\text{A7})$$

which is equivalent to restricting the sum over a_l in the mixed-canonical MPS representation in Eq. (A3) to the first D terms. The truncation error at a single site is

$$E_l(D) = \|\psi - |\psi_{\text{trunc}}\rangle\|_2^2 = \sum_{a_l=D+1}^{\chi_l} s_{a_l}^2. \quad (\text{A8})$$

A general MPS as in Eq. (A2) can be brought into the mixed-canonical form of Eq. (A3) by sweeping, e.g., in the forward direction (from $l = 1$ to $l = L - 1$) while performing SVDs. There, the MPS matrices are replaced by

$$M_{a_{l-1}, a_l}^{\sigma_l} \rightarrow U_{(\sigma_l, a_{l-1}), \tilde{a}_l} s_{\tilde{a}_l} V_{\tilde{a}_l, a_l}^\dagger, \quad (\text{A9a})$$

$$A_{a_{l-1}, \tilde{a}_l}^{\sigma_l} \equiv M_{a_{l-1}, \tilde{a}_l}^{\sigma_l} \leftarrow U_{(\sigma_l, a_{l-1}), \tilde{a}_l}, \quad (\text{A9b})$$

$$M_{\tilde{a}_l, a_{l+1}}^{\sigma_{l+1}} \leftarrow \sum_{a_l} s_{\tilde{a}_l} V_{\tilde{a}_l, a_l}^\dagger M_{a_l, a_{l+1}}^{\sigma_{l+1}}. \quad (\text{A9c})$$

The normalization property $\sum_{\sigma} A^{\sigma\dagger} A^{\sigma} = \mathbb{1}$ for MPS matrices to the left of the link l during the forward sweep follows from the orthonormality of the columns of $U_{(\sigma_l, a_{l-1}), \tilde{a}_l}$. Analogously, a backward sweep (from $l = L$ to $l = 2$) creates the orthonormality condition $\sum_{\sigma} B^{\sigma} B^{\sigma\dagger} = \mathbb{1}$ for all MPS matrices to the right of the link l during the sweep.

To summarize, the locally optimal compression of a noncanonical MPS, in principle, requires a first sweep along one direction without truncation to restore the canonical form, before the inner bond dimensions are reduced in a subsequent sweep along the opposite direction with truncated SVDs. The overall error is then bounded by [106] $\|\psi - |\psi_{\text{trunc}}\rangle\|_2^2 \leq 2 \sum_{l=1}^L E_l(D)$. These results for MPS similarly apply to MPOs with only subtle differences, such as that MPOs are generally not normalized to unity and may require rescaling [108].

2. Backward versus forward sweep thresholds

Despite the fact that the PT-MPO is not of canonical form after incorporating another row of the tensor network into the PT-MPO, it is standard practice for the sequential algorithm [49] to truncate small singular values in every sweep. Empirically, this is justified by the observation that the potentially suboptimal low-rank approximation—and thus larger bond dimensions of PT-MPOs after compression for the same accuracy—is typically overcompensated by the reduced computation time due to much smaller bond dimensions during the second sweep.

Here, we investigate whether a similar practice is also advisable for the divide-and-conquer algorithm. There, the combination of blocks takes place before the backward sweep; i.e., the backward sweep is responsible for restoring the canonical form. By choosing different compression thresholds for backward and forward sweeps, which we characterize by the ratio $r_b = \epsilon_{\text{backward}}/\epsilon_{\text{forward}}$, we interpolate between the choices of equal thresholds $r_b = 1$, as is the common practice for the sequential algorithm [49], and $r_b \rightarrow 0$, where the MPO is only brought to canonical form so that the truncation during the subsequent forward sweep is locally optimal. For a fixed forward sweep threshold $\epsilon_{\text{forward}} = \epsilon$, the latter can lead to smaller inner bond dimensions compared to calculations with $r_b = 1$ and, hence, to an overall faster algorithm. The optimal choice for r_b in different scenarios is explored below in numerical experiments.

3. Preselection thresholds

The preselection process in Eq. (16) consists of combining MPO matrices $[\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d''_{l-1}}$ and $[\tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}}$ by first constructing SVDs of the individual matrices $[\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d''_{l-1}} = \sum_s U_{d'_l, s}^{(1)} \sigma_s^{(1)} V_{s, (\alpha_l, d''_{l-1})}^{(1)\dagger}$ and $[\tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}} = \sum_t U_{d''_l, t}^{(2)} \sigma_t^{(2)} \times V_{t, (\alpha_l, d''_{l-1})}^{(2)\dagger}$. Then, combined indices (s, t) corresponding to small products of singular values $\sigma_s^{(1)} \sigma_t^{(2)} < \epsilon_{\text{select}} \sigma_0^{(1)} \sigma_0^{(2)}$ are disregarded. This process does not result in the locally optimal low-rank approximation for a given bond dimension either. However, it is noteworthy that if the outer bonds α_l on the combined MPO matrices acted on different spaces described by different indices α_l and α'_l , respectively, the decomposition of the matrix

$$\begin{aligned} M_{(d'_l, d''_l), ((\alpha_l, d''_{l-1}), (\alpha'_l, d''_{l-1}))} &= [\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d''_{l-1}} [\tilde{g}^{\alpha'_l}]_{d''_l, d''_{l-1}} \\ &= \sum_{s,t} (U_{d'_l, s}^{(1)} U_{d''_l, t}^{(2)}) \sigma_s^{(1)} \sigma_t^{(2)} (V_{s, (\alpha_l, d''_{l-1})}^{(1)\dagger} V_{t, (\alpha'_l, d''_{l-1})}^{(2)\dagger}) \end{aligned} \quad (\text{A10})$$

is—up to reordering of singular values—the SVD $M = U \Sigma V^\dagger$, with $U = U^{(1)} \otimes U^{(2)}$, $V^\dagger = V^{(1)\dagger} \otimes V^{(2)\dagger}$, and $\Sigma = \Sigma^{(1)} \otimes \Sigma^{(2)}$, with $\Sigma^{(i)} = \text{diag}(\sigma^{(i)})$ the diagonal matrices containing the individual singular values. Hence, by the Eckart-Young-Mirsky theorem [107], our truncation procedure based on products of singular values is optimal in this expanded space. The combined matrix $[\tilde{f}_{k:j+1}^{\alpha_l}]_{(d'_l, d''_l), (d''_{l-1}, d''_{l-1})} = [\tilde{f}_{k:j}^{\alpha_l}]_{d'_l, d''_{l-1}} [\tilde{g}^{\alpha_l}]_{d''_l, d''_{l-1}}$ in Eq. (16) can be obtained by eliminating one outer bond

$$\begin{aligned} [\tilde{f}_{k:j+1}^{\alpha_l}]_{(d'_l, d''_l), (d''_{l-1}, d''_{l-1})} &= \sum_{\alpha'_l} \delta_{\alpha_l, \alpha'_l} M_{(d'_l, d''_l), ((\alpha_l, d''_{l-1}), (\alpha'_l, d''_{l-1}))}. \end{aligned} \quad (\text{A11})$$

Hence, what makes our selection procedure suboptimal is that it neglects the fact that only the elements of M with $\alpha'_l = \alpha_l$ have to be reproduced by the low-rank approximation. Again, one strategy to compensate for this is to choose a smaller threshold ϵ_{select} for the selection of products of singular values, which we characterize by the ratio $r_s = \epsilon_{\text{select}}/\epsilon_{\text{forward}}$.

4. Numerical study of optimal threshold ratios

We now investigate how the overall computation time of PT-MPO simulations is affected by the threshold ratios r_b and r_s . Figure 8(a) shows results for PT-MPOs for a quantum dot coupled to longitudinal acoustic phonons as in Sec. III A. The PT-MPO is calculated by the divide-and-conquer algorithm using the same spectral density as in Fig. 4, threshold $\epsilon_{\text{forward}} = \epsilon = 10^{-12}$, time step $\Delta t = 0.01$ ps, and total propagation time $t_e = 20.48$ ps.

First, we find that reducing only the backward sweep threshold ratio r_b already results in faster PT-MPO construction with an optimal ratio around $r_b \sim 0.7$. The computation time can be further reduced by simultaneously decreasing the threshold used for preselection via r_s , where less than half the computation time is required for the choice $r_s = r_b = 0.2$ compared to equal thresholds $r_s = r_b = 1$. In contrast, here, decreasing only r_s does not lead to a better overall performance. This suggests that, in the example studied here, the preselection is already nearly optimal in selecting relevant degrees of freedom, while the MPO after selection and combination deviates considerably from its canonical form. The former can be explained by the fact that what dominates the computation time are the last few iterations of the divide-and-conquer algorithm, where the inner dimensions are largest. In this example, this involves the combination of elements that are shifted by a number of time steps of the order of the memory time n_c , where the new block overlaps with the long tail of the MPO with elements that are nearly independent of the outer index α_l , as $[b_l^{\alpha_l, \alpha_j}] \rightarrow 1$ for $l \rightarrow n_c$, irrespective of the values of α_l and α_j . In this case, it can be expected that the difference between Eqs. (A10) and (A11) becomes less relevant. Moreover, the fact that decreasing r_s along with r_b is more efficient than the decrease of r_b alone can be attributed to consistency: When $\epsilon_{\text{select}} < \epsilon_{\text{backward}}$, terms are eliminated in a way that is not locally optimal during the preselection that would otherwise allow the canonicalization during the backward sweep to result in a more compact and accurate form.

A somewhat different picture unfolds for the example of the spin-boson model with strongly peaked spectral density with coupling strength $\eta = 0.01$ and total propagation time $t_e = 8$ discussed in Sec. III C, for which we plot the computation time for varying threshold ratios r_b and r_s in Fig. 8(b). There, decreasing only r_b increases the computation time, while decreasing r_s is found to be more efficient. The main difference from the situation in Fig. 8(a)

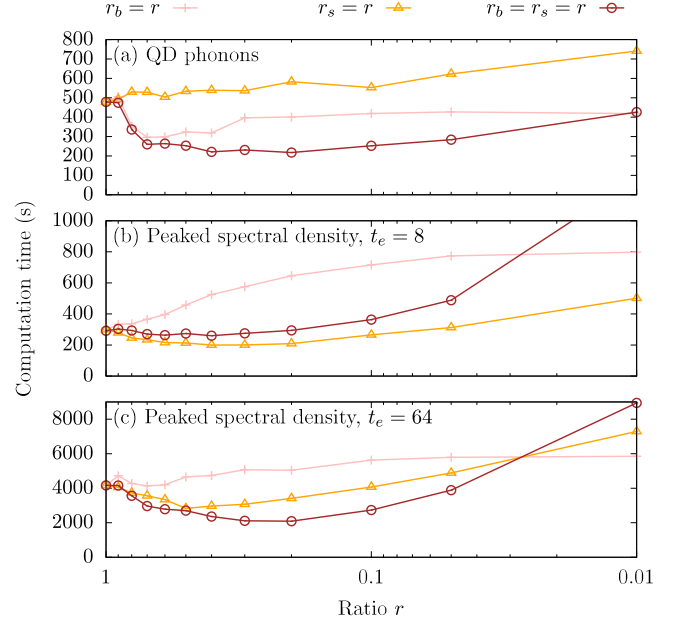


FIG. 8. Computation time for PT calculation as a function of ratios between backward versus forward sweep thresholds $r = r_b = \epsilon_{\text{backward}}/\epsilon_{\text{forward}}$, selection threshold versus forward sweep threshold $r = r_s = \epsilon_{\text{select}}/\epsilon_{\text{forward}}$, and both, $r = r_b = r_s$. Results are shown for (a) a quantum dot coupled to phonons as in Sec. III A and strongly peaked spectral densities as in Sec. III C with coupling strength $\eta = 0.01$ for total propagation times $t_e = 8$ (b) and $t_e = 64$ (c), respectively.

is that, in the parameter regime considered in Fig. 8(b), the final time t_e is still well within the memory time of the environment, and the preselection combines MPO matrices whose elements depend significantly on the outer indices. As such, optimal compression of Eq. (A10) does not necessarily compress Eq. (A11) optimally, and the preselection based on products of singular values is not a perfect proxy for the SVD of the product of the MPO matrices. As small r_s increases the inner bond dimension after the selection process, it appears that truncating sooner rather than later, i.e., already during the backward sweep by keeping $r_b = 1$, is the best strategy to minimize the overall computation time in this situation.

To test the hypothesis that the truncation strategy should be chosen differently depending on whether the total propagation time exceeds the physical memory time of the environment, we show in Fig. 8(c) the computation times for simulations as in Fig. 8(b) but with larger total propagation time $t_e = 64$. In this case, we indeed find that the optimal strategy is to reduce both threshold ratios r_s and r_b , as the situation is more comparable to the one in Fig. 8(a).

APPENDIX B: CONVERGENCE OF FLUORESCENCE SPECTRA

To demonstrate that calculating fluorescence spectra of quantum dots in the strong driving limit indeed puts strong

requirements on the convergence parameters, we compare in Fig. 9 the spectra obtained for $\hbar\Omega = 5$ meV in Fig. 4(e) with simulations where a single convergence parameter is varied.

The impact of a too-coarse time discretization can be observed in Figs. 9(a) and 9(b), where panel (a) shows the spectra on a logarithmic scale while panel (b) shows the left Mollow peak enlarged and on a linear scale. The overall memory time $n_c\Delta t = 20.48$ ps is kept fixed; i.e., n_c is varied along with the time steps Δt . For larger time steps $\Delta t = 320$ fs, one already finds many features of the spectra well reproduced on the logarithmic scale. Note, however, that sizable deviations are still found on the linear scale in panel (b) for $\Delta t = 160$, suggesting that our choice for $\Delta t = 10$ fs in the main text is indeed of the required order of magnitude for convergence. In additional calculations (not shown), we kept this time step $\Delta t = 10$ fs while varying the memory cutoff n_c . We found that the algorithm leads to the fastest results for $n_c = 2048$ memory time steps; i.e., no gain is made by truncating the memory earlier.

In Figs. 9(c) and 9(d), we show that reducing the overall number of time steps from $n = 2^{21}$ to $n = 2^{19}$ results in

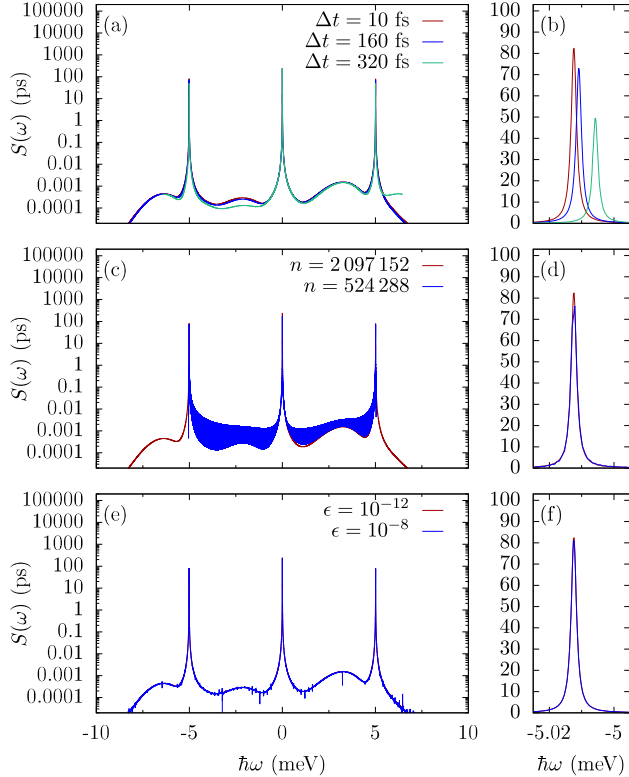


FIG. 9. Convergence of fluorescence spectra with respect to time step Δt (a),(b), total number of time steps n (c),(d), and truncation threshold ϵ (e),(f). Left panels (a),(c),(e) use the logarithmic scale; right panels (b),(d),(f) show a zoom into the left peak on a linear scale. The remaining parameters are the same as for Fig. 4(e).

oscillatory artifacts in the spectra, as the total propagation time $t_\epsilon = n\Delta t$ is too short for all itinerant excitations in the two-time correlation function to decay. One can also see, on the linear scale in panel (d), that the low-energy peak is not sampled sufficiently.

Finally, the impact of the compression threshold is analyzed in Figs. 9(e) and 9(f). This seems less severe, as it mainly leads to small blips on top of the otherwise well-reproduced spectra. Note, however, that we find deviations of the two-time correlation functions in the time domain by less than or around 10^{-4} when comparing simulations with $\epsilon = 10^{-12}$ to those with $\epsilon = 10^{-11}$ (not shown). This result indicates that such small thresholds are required for well-converged dynamics, yet small errors in temporal data may be averaged out by the Fourier transform.

In any case, the chosen convergence parameters $\Delta t = 10$ fs, $n_c = 2^{11}$, and $n = 2^{21}$ used for the simulations in Fig. 4 are indeed not too far from the minimal requirement for well-converged spectra.

APPENDIX C: REVERSE ROTATING WAVE APPROXIMATION

In this appendix, we discuss how one may consider Eq. (27) as a “reverse rotating wave approximation” of the standard model of superradiance, which uses the Tavis-Cummings model. In this picture, we consider using the reverse RWA to convert a system-environment coupling that is not of the product form required to match Eq. (1) into such a form. In such a picture, we consider ω_R as a convergence parameter that determines the quality of the approximation.

We start by summarizing the standard RWA, starting from Eq. (27). Such an approximation is valid if the typical value of the emitter frequencies, ω_j , is much larger than the couplings g_k . The RWA is achieved by first applying a time-dependent unitary to change to a rotating frame,

$$U(t) = e^{i\omega_R(\sum_j \sigma_j^+ \sigma_j^- + \sum_k a_k^\dagger a_k)t}. \quad (\text{C1})$$

Here, ω_R is a reference frequency, typically chosen to be similar to the emitter frequencies ω_j . If a wave function $|\Psi(t)\rangle$ obeys the Schrödinger equation $i\hbar(\partial/\partial t)|\Psi(t)\rangle = H|\Psi(t)\rangle$, then $|\Psi'(t)\rangle = U(t)|\Psi(t)\rangle$ obeys

$$\begin{aligned} i\hbar \frac{\partial}{\partial t} |\Psi'(t)\rangle &= \left(U H U^\dagger + i\hbar \frac{\partial U}{\partial t} U^\dagger \right) |\Psi'(t)\rangle \\ &= H' |\Psi'(t)\rangle, \end{aligned} \quad (\text{C2})$$

with

$$\begin{aligned}
 H' = & \sum_{j=1}^N \hbar(\omega_j - \omega_R) \sigma_j^+ \sigma_j^- + \sum_k \hbar(\omega_k - \omega_R) a_k^\dagger a_k \\
 & + \sum_{j,k} \hbar g_k (a_k^\dagger \sigma_j^- + a_k \sigma_j^+) \\
 & + \sum_{j,k} \hbar g_k (a_k^\dagger \sigma_j^+ e^{i2\omega_R t} + a_k \sigma_j^- e^{-i2\omega_R t}). \quad (\text{C3})
 \end{aligned}$$

The RWA is completed by neglecting the counterrotating terms, i.e., the terms oscillating with frequency $\pm 2\omega_R$, which is justified if the reference frequency ω_R is much larger than any other frequency in the system. The impact of these terms on the dynamics on a coarse-grained time-scale T is of the order

$$\frac{1}{T} \int_0^T dt e^{\pm i2\omega_R t} = \frac{e^{\pm i2\omega_R T} - 1}{\pm i2\omega_R T} = \mathcal{O}\left(\frac{1}{\omega_R T}\right). \quad (\text{C4})$$

With the definition of relative frequencies $\tilde{\omega}_j = \omega_j - \omega_R$ and $\tilde{\omega}_k = \omega_k - \omega_R$, the RWA turns the multimode Dicke model of Eq. (27) into the multimode Tavis-Cummings model

$$\begin{aligned}
 \tilde{H} = & \sum_{j=1}^N \hbar \tilde{\omega}_j \sigma_j^+ \sigma_j^- + \sum_k \hbar \tilde{\omega}_k a_k^\dagger a_k \\
 & + \sum_{j,k} \hbar g_k (a_k^\dagger \sigma_j^- + a_k \sigma_j^+). \quad (\text{C5})
 \end{aligned}$$

Equation (C5) has the advantage that only comparatively slow oscillations with frequencies of the order $\tilde{\omega}_j$ and g_k have to be resolved; fast oscillations with frequencies ω_R are eliminated. One may also note that after this process, the value ω_R does not appear: As long as the RWA is valid, models with different values of ω_R will all map to the same approximate model.

We may now consider this process in reverse. If one requires a solution of the multimode Tavis-Cummings model in Eq. (C5), a reverse rotating wave approximation can be applied by introducing a reference frequency ω_R , which now takes the role of an additional convergence parameter. As seen in Eq. (C4), the error introduced by the reverse RWA can be systematically reduced by increasing the frequency ω_R . Our divide-and-conquer algorithm is key to dealing with the small time steps Δt needed to resolve the fast oscillations $\omega_R \Delta t \ll 1$. As noted above, the value of ω_R does not enter the calculation as long as it is large enough. Thus, one may also use the above arguments to model systems where the physical ω_R is too large even for the divide-and-conquer approach. A RWA can be made to eliminate the original common frequency, and then the reverse RWA can be made to recover a model of the form of Eq. (1), allowing the use of the PT-MPO approaches described in this paper.

APPENDIX D: ANALYTIC EXPRESSION FOR PHOTON INTENSITIES FROM IDEAL SUPERRADIANT EMITTERS

Here, we derive analytic expressions for the intensity of photons emitted from N ideal superradiant emitters coupled to a flat (Markovian) bath of photon modes. As described in the main text, the dynamics is governed by a sequence of photon emission processes between the Dicke states $|J, M\rangle$ and $|J, M-1\rangle$ with rates $\Gamma_{J,M} = \kappa(J+M)(J-M+1)$, where $J = N/2$ and $M = J, J-1, \dots, -J$. We denote by $p_M^{[N]}$ the populations of the Dicke state $|N/2, M\rangle$. The corresponding rate equations are

$$\frac{\partial}{\partial t} p_{N/2}^{[N]} = -\Gamma_{N/2, N/2} p_{N/2}^{[N]} \quad (\text{D1a})$$

for the Dicke state with maximal excitation $M = N/2$ and

$$\frac{\partial}{\partial t} p_M^{[N]} = \Gamma_{N/2, M+1} p_{M+1}^{[N]} - \Gamma_{N/2, M} p_M^{[N]} \quad (\text{D1b})$$

for the remaining Dicke states with $M < N/2$. These equations are solved by

$$p_{N/2}^{[N]}(t) = e^{-\Gamma_{N/2, N/2} t} \quad (\text{D2a})$$

and

$$p_M^{[N]}(t) = \Gamma_{N/2, M+1} \int_0^t dt e^{-\Gamma_{N/2, M}(t-\tau)} p_{M+1}^{[N]}(\tau), \quad (\text{D2b})$$

respectively, where we have already incorporated the initial condition of initially maximally excited emitters $p_{N/2}^{[N]}(0) = 1, p_{M < N/2}^{[N]}(0) = 0$. It can be seen that the solution $p_M^{[N]}(t)$ can be calculated by integrating over the solutions $p_{M+1}^{[N]}(t)$ after multiplying with a simple exponential kernel. Hence, all Dicke state populations can be obtained one at a time by Eq. (D2b) starting from $p_{N/2}^{[N]}(t)$ given in Eq. (D2a).

Because the Dicke state $|J, M\rangle$ carries $M + J$ excitations, the emitted photon intensity, i.e., the negative change of the total emitter excitation, is

$$I^{[N]} = -\frac{\partial}{\partial t} \sum_{M=-N/2}^{N/2} (M + N/2) p_M^{[N]}. \quad (\text{D3})$$

For up to $N = 5$, we thus find

$$\frac{I^{[1]}}{\kappa} = e^{-\kappa t}, \quad (\text{D4a})$$

$$\frac{I^{[2]}}{2\kappa} = (2\kappa t + 1) e^{-2\kappa t}, \quad (\text{D4b})$$

$$\frac{I^{[3]}}{3\kappa} = (12\kappa t - 7)e^{-3\kappa t} + 8e^{-4\kappa t}, \quad (\text{D4c})$$

$$\frac{I^{[4]}}{4\kappa} = (36\kappa t - 23)e^{-4\kappa t} + (18\kappa t + 24)e^{-6\kappa t}, \quad (\text{D4d})$$

$$\frac{I^{[5]}}{5\kappa} = \left(80\kappa t - \frac{143}{3}\right)e^{-5\kappa t} + \left(128\kappa t - \frac{16}{3}\right)e^{-8\kappa t} + 54e^{-9\kappa t}. \quad (\text{D4e})$$

-
- [1] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002).
- [2] G. A. L. White, C. D. Hill, F. A. Pollock, L. C. L. Hollenberg, and K. Modi, *Demonstration of non-Markovian process characterisation and control on a quantum processor*, *Nat. Commun.* **11**, 6301 (2020).
- [3] 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).
- [4] S. Lüker and D. E. Reiter, *A review on optical excitation of semiconductor quantum dots under the influence of phonons*, *Semicond. Sci. Technol.* **34**, 063002 (2019).
- [5] 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(R) (2017).
- [6] G. R. Wilbur, A. Binai-Motlagh, A. Clarke, A. Ramachandran, N. Milson, J. P. Healey, S. O'Neal, D. G. Deppe, and K. C. Hall, *Notch-filtered adiabatic rapid passage for optically driven quantum light sources*, *APL Photonics* **7**, 111302 (2022).
- [7] J. H. Quilter, A. J. Brash, F. Liu, M. Glässl, A. M. Barth, V. M. Axt, A. J. Ramsay, M. S. Skolnick, and A. M. Fox, *Phonon-assisted population inversion of a single InGaAs/GaAs quantum dot by pulsed laser excitation*, *Phys. Rev. Lett.* **114**, 137401 (2015).
- [8] 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).
- [9] C. Maier, T. Brydges, P. Jurcevic, N. Trautmann, C. Hempel, B. P. Lanyon, P. Hauke, R. Blatt, and C. F. Roos, *Environment-assisted quantum transport in a 10-qubit network*, *Phys. Rev. Lett.* **122**, 050501 (2019).
- [10] 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).
- [11] M. B. Plenio and S. F. Huelga, *Dephasing-assisted transport: quantum networks and biomolecules*, *New J. Phys.* **10**, 113019 (2008).
- [12] G. Lindblad, *On the generators of quantum dynamical semigroups*, *Commun. Math. Phys.* **48**, 119 (1976).
- [13] 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).
- [14] M. Bundgaard-Nielsen, J. Mørk, and E. V. Denning, *Non-Markovian perturbation theories for phonon effects in strong-coupling cavity quantum electrodynamics*, *Phys. Rev. B* **103**, 235309 (2021).
- [15] C. Roy and S. Hughes, *Polaron master equation theory of the quantum-dot Mollow triplet in a semiconductor cavity-QED system*, *Phys. Rev. B* **85**, 115309 (2012).
- [16] A. Nazir and D. P. S. McCutcheon, *Modelling exciton-phonon interactions in optically driven quantum dots*, *J. Phys. Condens. Matter* **28**, 103002 (2016).
- [17] S. Kundu, R. Dani, and N. Makri, *Tight inner ring architecture and quantum motion of nuclei enable efficient energy transfer in bacterial light harvesting*, *Sci. Adv.* **8**, eadd0023 (2022).
- [18] 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).
- [19] F. Ungar, M. Cygorek, and V. M. Axt, *Trend reversal in the magnetic-field dependence of exciton spin-transfer rates in diluted magnetic semiconductors due to non-Markovian dynamics*, *Phys. Rev. B* **97**, 045210 (2018).
- [20] M. Cygorek, F. Ungar, P. I. Tamborenea, and V. M. Axt, *Influence of nonmagnetic impurity scattering on spin dynamics in diluted magnetic semiconductors*, *Phys. Rev. B* **95**, 045204 (2017).
- [21] U. Hoeppe, C. Wolff, J. Küchenmeister, J. Niegemann, M. Drescher, H. Benner, and K. Busch, *Direct observation of non-Markovian radiation dynamics in 3D bulk photonic crystals*, *Phys. Rev. Lett.* **108**, 043603 (2012).
- [22] S. John and T. Quang, *Spontaneous emission near the edge of a photonic band gap*, *Phys. Rev. A* **50**, 1764 (1994).
- [23] I. de Vega and D. Alonso, *Dynamics of non-Markovian open quantum systems*, *Rev. Mod. Phys.* **89**, 015001 (2017).
- [24] F. Rossi and T. Kuhn, *Theory of ultrafast phenomena in photoexcited semiconductors*, *Rev. Mod. Phys.* **74**, 895 (2002).
- [25] V. M. Axt and A. Stahl, *A dynamics-controlled truncation scheme for the hierarchy of density matrices in semiconductor optics*, *Z. Phys. B* **93**, 195 (1994).
- [26] H. Wang and M. Thoss, *Quantum dynamical simulation of electron-transfer reactions in an anharmonic environment*, *J. Phys. Chem. A* **111**, 10369 (2007).
- [27] A. D. Somoza, O. Marty, J. Lim, S. F. Huelga, and M. B. Plenio, *Dissipation-assisted matrix product factorization*, *Phys. Rev. Lett.* **123**, 100502 (2019).
- [28] M. Brenes, J. J. Mendoza-Arenas, A. Purkayastha, M. T. Mitchison, S. R. Clark, and J. Goold, *Tensor-network method to simulate strongly interacting quantum thermal machines*, *Phys. Rev. X* **10**, 031040 (2020).
- [29] C. McConnell and A. Nazir, *Strong coupling in thermo-electric nanojunctions: A reaction coordinate framework*, *New J. Phys.* **24**, 025002 (2022).
- [30] F. Mascherpa, A. Smirne, A. D. Somoza, P. Fernández-Acebal, S. Donadi, D. Tamascelli, S. F. Huelga, and

- M. B. Plenio, *Optimized auxiliary oscillators for the simulation of general open quantum systems*, *Phys. Rev. A* **101**, 052108 (2020).
- [31] S. Nakajima, *On quantum theory of transport phenomena: Steady diffusion*, *Prog. Theor. Phys.* **20**, 948 (1958).
- [32] 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).
- [33] Y. 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 (1989).
- [34] Y. Tanimura, *Nonperturbative expansion method for a quantum system coupled to a harmonic-oscillator bath*, *Phys. Rev. A* **41**, 6676 (1990).
- [35] W. T. Strunz, *Linear quantum state diffusion for non-Markovian open quantum systems*, *Phys. Lett. A* **224**, 25 (1996).
- [36] L. Diósi and W. T. Strunz, *The non-Markovian stochastic Schrödinger equation for open systems*, *Phys. Lett. A* **235**, 569 (1997).
- [37] L. Diósi, N. Gisin, and W. T. Strunz, *Non-Markovian quantum state diffusion*, *Phys. Rev. A* **58**, 1699 (1998).
- [38] J. T. Stockburger and H. Grabert, *Exact c -number representation of non-Markovian quantum dissipation*, *Phys. Rev. Lett.* **88**, 170407 (2002).
- [39] D. Suess, A. Eisfeld, and W. T. Strunz, *Hierarchy of stochastic pure states for open quantum system dynamics*, *Phys. Rev. Lett.* **113**, 150403 (2014).
- [40] C.-Y. Hsieh and J. Cao, *A unified stochastic formulation of dissipative quantum dynamics. I. Generalized hierarchical equations*, *J. Chem. Phys.* **148**, 014103 (2018).
- [41] V. Link, K. Müller, R. G. Lena, K. Luoma, F. m. c. Damanet, W. T. Strunz, and A. J. Daley, *Non-Markovian quantum dynamics in strongly coupled multimode cavities conditioned on continuous measurement*, *PRX Quantum* **3**, 020348 (2022).
- [42] M. P. Woods and M. B. Plenio, *Dynamical error bounds for continuum discretisation via Gauss quadrature rules—A Lieb-Robinson bound approach*, *J. Math. Phys. (N.Y.)* **57**, 022105 (2016).
- [43] 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).
- [44] 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).
- [45] J. Cerrillo and J. Cao, *Non-markovian dynamical maps: Numerical processing of open quantum trajectories*, *Phys. Rev. Lett.* **112**, 110401 (2014).
- [46] N. Makri, *Small matrix disentanglement of the path integral: Overcoming the exponential tensor scaling with memory length*, *J. Chem. Phys.* **152**, 041104 (2020).
- [47] N. Makri, *Small matrix path integral for driven dissipative dynamics*, *J. Phys. Chem. A* **125**, 10500 (2021).
- [48] 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).
- [49] 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).
- [50] S. M. Ulrich, S. Ates, S. Reitzenstein, A. Löffler, A. Forchel, and P. Michler, *Dephasing of triplet-sideband optical emission of a resonantly driven InAs/GaAs quantum dot inside a microcavity*, *Phys. Rev. Lett.* **106**, 247402 (2011).
- [51] A. Ulhaq, S. Weiler, S. M. Ulrich, R. Roßbach, M. Jetter, and P. Michler, *Cascaded single-photon emission from the Mollow triplet sidebands of a quantum dot*, *Nat. Photonics* **6**, 238 (2012).
- [52] 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).
- [53] Note that the time integrals in Eq. (5) can be solved analytically, and the remaining integral over ω in Eq. (6) can be expressed as a Fourier transform. Thus, n values of η_{i-j} are efficiently obtained in $\mathcal{O}(n \log n)$ time.
- [54] R. Orús, *A practical introduction to tensor networks: Matrix product states and projected entangled pair states*, *Ann. Phys. (N.Y.)* **349**, 117 (2014).
- [55] U. Schollwöck, *The density-matrix renormalization group in the age of matrix product states*, *Ann. Phys. (N.Y.)* **326**, 96 (2011).
- [56] G. E. Fux, D. Kilda, B. W. Lovett, and J. Keeling, *Tensor network simulation of chains of non-Markovian open quantum systems*, *Phys. Rev. Res.* **5**, 033078 (2023).
- [57] 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).
- [58] J. W. Cooley and J. W. Tukey, *An algorithm for the machine calculation of complex Fourier series*, *Math. Comput.* **19**, 297 (1965).
- [59] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. (Cambridge University Press, Cambridge, USA, 1992).
- [60] S. Kundu and N. Makri, *Efficient matrix factorisation of the modular path integral for extended systems*, *Mol. Phys.* **119**, e1797200 (2021).
- [61] S. Östlund and S. Rommer, *Thermodynamic limit of density matrix renormalization*, *Phys. Rev. Lett.* **75**, 3537 (1995).
- [62] G. Vidal, *Classical simulation of infinite-size quantum lattice systems in one spatial dimension*, *Phys. Rev. Lett.* **98**, 070201 (2007).
- [63] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, *Classical simulation of infinite-size quantum lattice systems in two spatial dimensions*, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [64] P. Lodahl, *Quantum-dot based photonic quantum networks*, *Quantum Sci. Technol.* **3**, 013001 (2018).
- [65] S. E. Thomas, M. Billard, N. Coste, S. C. Wein, Priya, H. Ollivier, O. Krebs, L. Tazaïrt, A. Harouri, A. Lemaitre, I. Sagnes, C. Anton, L. Lanco, N. Somaschi, J. C. Loredo, and P. Senellart, *Bright polarized single-photon source*

- based on a linear dipole, *Phys. Rev. Lett.* **126**, 233601 (2021).
- [66] T. Seidelmann, F. Ungar, A. M. Barth, A. Vagov, V. M. Axt, M. Cygorek, and T. Kuhn, *Phonon-induced enhancement of photon entanglement in quantum dot-cavity systems*, *Phys. Rev. Lett.* **123**, 137401 (2019).
- [67] R. M. Stevenson, C. L. Salter, J. Nilsson, A. J. Bennett, M. B. Ward, I. Farrer, D. A. Ritchie, and A. J. Shields, *Indistinguishable entangled photons generated by a light-emitting diode*, *Phys. Rev. Lett.* **108**, 040503 (2012).
- [68] C. S. Muñoz, E. del Valle, A. G. Tudela, K. Müller, S. Lichtmannecker, M. Kaniber, C. Tejedor, J. J. Finley, and F. P. Laussy, *Emitters of N -photon bundles*, *Nat. Photonics* **8**, 550 (2014).
- [69] M. Cosacchi, J. Wiercinski, T. Seidelmann, M. Cygorek, A. Vagov, D. E. Reiter, and V. M. Axt, *On-demand generation of higher-order Fock states in quantum-dot-cavity systems*, *Phys. Rev. Res.* **2**, 033489 (2020).
- [70] D. P. S. McCutcheon, N. S. Dattani, E. M. Gauger, B. W. Lovett, and A. Nazir, *A general approach to quantum dynamics using a variational master equation: Application to phonon-damped Rabi rotations in quantum dots*, *Phys. Rev. B* **84**, 081305(R) (2011).
- [71] M. Cygorek, A. M. Barth, F. Ungar, A. Vagov, and V. M. Axt, *Nonlinear cavity feeding and unconventional photon statistics in solid-state cavity QED revealed by many-level real-time path-integral calculations*, *Phys. Rev. B* **96**, 201201(R) (2017).
- [72] M. Grifoni and P. Hänggi, *Driven quantum tunneling*, *Phys. Rep.* **304**, 229 (1998).
- [73] J. Shao and N. Makri, *Iterative path integral calculation of quantum correlation functions for dissipative systems*, *Chem. Phys.* **268**, 1 (2001).
- [74] J. Shao and N. Makri, *Iterative path integral formulation of equilibrium correlation functions for quantum dissipative systems*, *J. Chem. Phys.* **116**, 507 (2002).
- [75] M. Cosacchi, M. Cygorek, F. Ungar, A. M. Barth, A. Vagov, and V. M. Axt, *Path-integral approach for non-equilibrium multi-time correlation functions of open quantum systems coupled to Markovian and non-Markovian environments*, *Phys. Rev. B* **98**, 125302 (2018).
- [76] T. K. Bracht, M. Cosacchi, T. Seidelmann, M. Cygorek, A. Vagov, V. M. Axt, T. Heindel, and D. E. Reiter, *Swing-up of quantum emitter population using detuned pulses*, *PRX Quantum* **2**, 040354 (2021).
- [77] B. R. Mollow, *Power spectrum of light scattered by two-level systems*, *Phys. Rev.* **188**, 1969 (1969).
- [78] H. J. Kimble and L. Mandel, *Theory of resonance fluorescence*, *Phys. Rev. A* **13**, 2123 (1976).
- [79] B. Krummheuer, V. M. Axt, T. Kuhn, I. D'Amico, and F. Rossi, *Pure dephasing and phonon dynamics in GaAs- and GaN-based quantum dot structures: Interplay between material parameters and geometry*, *Phys. Rev. B* **71**, 235329 (2005).
- [80] Y.-J. Wei, Y. He, Y.-M. He, C.-Y. Lu, J.-W. Pan, C. Schneider, M. Kamp, S. Höfling, D. P. S. McCutcheon, and A. Nazir, *Temperature-dependent Mollow triplet spectra from a single quantum dot: Rabi frequency renormalization and sideband linewidth insensitivity*, *Phys. Rev. Lett.* **113**, 097401 (2014).
- [81] 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).
- [82] 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).
- [83] 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).
- [84] 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).
- [85] K. Boos, S. K. Kim, T. Bracht, F. Sbresny, J. Kaspari, M. Cygorek, H. Riedl, F. W. Bopp, W. Rauhaus, C. Calcagno, J. J. Finley, D. E. Reiter, and K. Mueller, *Signatures of dynamically dressed states*, arXiv:2305.15827.
- [86] R. H. Dicke, *Coherence in spontaneous radiation processes*, *Phys. Rev.* **93**, 99 (1954).
- [87] M. Gross and S. Haroche, *Superradiance: An essay on the theory of collective spontaneous emission*, *Phys. Rep.* **93**, 301 (1982).
- [88] 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).
- [89] J. Q. Grim, A. S. Bracker, M. Zalalutdinov, S. G. Carter, A. C. Kozen, M. Kim, C. S. Kim, J. T. Mlack, M. Yakes, B. Lee, and D. Gammon, *Scalable in operando strain tuning in nanophotonic waveguides enabling three-quantum-dot superradiance*, *Nat. Mater.* **18**, 963 (2019).
- [90] J.-H. Kim, S. Aghaeimeibodi, C. J. K. Richardson, R. P. Leavitt, and E. Waks, *Super-radiant emission from quantum dots in a nanophotonic waveguide*, *Nano Lett.* **18**, 4734 (2018).
- [91] A. Tiranov, V. Angelopoulou, C. J. van Diepen, B. Schriniski, O. A. D. Sandberg, Y. Wang, L. Midolo, S. Scholz, A. D. Wieck, A. Ludwig, A. S. Sørensen, and P. Lodahl, *Collective super- and subradiant dynamics between distant optical quantum emitters*, *Science* **379**, 389 (2023).
- [92] A. Strathearn, *Modelling Non-Markovian Quantum Systems Using Tensor Networks*, Springer Theses (Springer Cham, Cham, 2020), 10.1007/978-3-030-54975-6.
- [93] A. Garg, J. N. Onuchic, and V. Ambegaokar, *Effect of friction on electron transfer in biomolecules*, *J. Chem. Phys.* **83**, 4491 (1985).
- [94] J. Iles-Smith, N. Lambert, and A. Nazir, *Environmental dynamics, correlations, and the emergence of noncanonical equilibrium states in open quantum systems*, *Phys. Rev. A* **90**, 032114 (2014).
- [95] J. Gray and S. Kourtis, *Hyper-optimized tensor network contraction*, *Quantum* **5**, 410 (2021).
- [96] M. C. Bañuls, *Tensor network algorithms: A route map*, *Annu. Rev. Condens. Matter Phys.* **14**, 173 (2023).
- [97] 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).
- [98] Y. Luo, G. Chen, Y. Zhang, L. Zhang, Y. Yu, F. Kong, X. Tian, Y. Zhang, C. Shan, Y. Luo, J. Yang, V. Sandoghdar, Z. Dong, and J. G. Hou, *Electrically driven single-photon superradiance from molecular chains in a plasmonic nanocavity*, *Phys. Rev. Lett.* **122**, 233901 (2019).
- [99] 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).
- [100] H. Maguire, J. Iles-Smith, and A. Nazir, *Environmental nonadditivity and Franck-Condon physics in nonequilibrium quantum systems*, *Phys. Rev. Lett.* **123**, 093601 (2019).
- [101] 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).
- [102] 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).
- [103] C. Roy and S. Hughes, *Influence of electron–acoustic-phonon scattering on intensity power broadening in a coherently driven quantum-dot–cavity system*, *Phys. Rev. X* **1**, 021009 (2011).
- [104] V. Link, H.-H. Tu, and W. T. Strunz, *Open quantum system dynamics from infinite tensor network contraction*, *arXiv:2307.01802*.
- [105] M. Cygorek, J. Keeling, B. W. Lovett, and E. Gauger, *Datafiles for “Sublinear scaling in non-Markovian open quantum systems simulations”*, 2024, Heriot-Watt University, [10.17861/6666a00a-887c-4d91-887d-c19bf3450180](https://doi.org/10.17861/6666a00a-887c-4d91-887d-c19bf3450180)
- [106] F. Verstraete and J. I. Cirac, *Matrix product states represent ground states faithfully*, *Phys. Rev. B* **73**, 094423 (2006).
- [107] C. Eckart and G. Young, *The approximation of one matrix by another of lower rank*, *Psychometrika* **1**, 211 (1936).
- [108] C. Hubig, I. P. McCulloch, and U. Schollwöck, *Generic construction of efficient matrix product operators*, *Phys. Rev. B* **95**, 035129 (2017).