

Variational Perturbation Theory in Open Quantum Systems for Efficient Steady State Computation

André Melo, Gaspard Beugnot, and Fabrizio Minganti

Alice & Bob, 53 boulevard du Général Martial Valin, 75015 Paris, France

Determining the steady state of an open quantum system is crucial for characterizing quantum devices and studying various physical phenomena. Often, computing a single steady state is insufficient, and it is necessary to explore its dependence on multiple external parameters. In such cases, calculating the steady state independently for each combination of parameters quickly becomes intractable. Perturbation theory (PT) can mitigate this challenge by expanding steady states around reference parameters, minimizing redundant computations across neighboring parameter values. However, PT has two significant limitations: it relies on the pseudo-inverse—a numerically costly operation—and has a limited radius of convergence. In this work, we remove both of these roadblocks. First, we introduce a variational perturbation theory (VPT) and its multipoint generalization that significantly extends the radius of convergence even in the presence of non-analytic effects such as dissipative phase transitions. Then, we develop two numerical strategies that eliminate the need to compute pseudo-inverses. The first relies on a single LU decomposition to efficiently construct the steady state within the convergence region, while the second reformulates VPT as a Krylov space recycling problem and uses preconditioned iterative methods. We benchmark these approaches across various models, demonstrating their broad applicability and significant improvements over standard PT.

1 Introduction

The interaction between a quantum system and its environment gives rise to unique physical phenomena not found in isolated systems [1, 2]. The Lindblad master equation is a powerful framework for capturing these effects [3, 4] and accurately describes a wide range of experimental platforms, including superconducting circuits [5], polaritons [6], atoms [7], molecules [8] and ions [9].

Fabrizio Minganti: E-mail: fabrizio.minganti@gmail.com

When studying a Lindblad master equation, the short- or long-time regimes are typically of interest. During short-time dynamics, the environment introduces small perturbations to the isolated behavior of the system. Over longer timescales, the system eventually reaches equilibrium with its environment by relaxing towards its steady state. Mathematically, the steady state is described by a stationary density matrix $\hat{\rho}_{ss}$ that satisfies $\mathcal{L}\hat{\rho}_{ss} = 0$, where $\mathcal{L}$ is the Liouvillian governing the dynamics [3]. In experimental settings, steady states are often easier to prepare and measure than transient dynamics and provide insight into diverse physical phenomena, including chaos [10, 11], phase transitions [12–14], and other properties such as sub-Poissonian photon statistics [15]. Therefore, computing steady states is a central task in the study of open quantum systems.

Although a few models are known to have analytical solutions [16–18], most steady states can only be determined numerically. Open quantum systems, however, are particularly challenging to simulate on classical computers due to the exponential scaling of the Hilbert space and the additional quadratic cost of working with density matrices rather than wave functions. To address these challenges, various numerical methods have been developed to solve problems in large Hilbert spaces [19], including the quantum trajectory formalism [20, 21], renormalization techniques [22, 23], tensor [24, 25] and neural network [26, 27] ansätze, phase space methods [28, 29], cluster mean field theory [30], semidefinite relaxation [31] and variational approaches [32].

Less attention has been devoted to intermediate system sizes, where off-the-shelf sparse linear algebra routines (e.g., LU decomposition) are often sufficient to compute the steady state at a single set of parameters but quickly become inefficient for exploring parameter spaces or mapping phase diagrams. Additionally, this approach fails to exploit the continuity of the problem, as small changes in parameters typically result in small changes in the steady state. An alternative strategy relies on perturbation theory for open quantum systems, which extends Rayleigh-Schrödinger perturbation theory to this non-Hermitian case [33–35]. However, as we detail below, perturbation theory requires a *computationally expensive* operation known as the *Moore–Penrose pseudo-inverse* and typically has a *limited radius of conver-*

gence. Beyond perturbative methods, complementary approaches have been proposed to address steady-state computations over extended parameter regions. In Ref. [36], steady states at new parameter values are approximated by projecting onto a reduced basis of solutions computed at selected points in parameter space. Ref. [37] derived an exact expression for the steady state as a function of a single control parameter, yielding a closed-form solution obtained from the unperturbed steady state through an explicit operator inversion. More broadly, perturbative and hybrid expansion strategies have appeared in other contexts, including structural reanalysis of parametric linear systems [38] and perturbational-variational treatments of Schrödinger eigenvalue problems [39, 40]. However, these approaches are formulated for static or Hermitian problems and rely on Taylor expansions or energy-based variational principles, and do not directly extend to Liouvillian steady-state problems in open quantum systems. From a practical perspective, the lack of efficient methods for exploring large parameter regions also impacts experimental efforts. Indeed, when calibrating quantum hardware, steady state measurements obtained while changing multiple controllable parameters are frequently fitted with numerical simulations [41–43], a task that becomes impractical without scalable computational techniques.

In this work, we introduce *variational perturbation theory* (VPT) for open quantum systems. We relax the constraints of standard perturbation theory (PT) and recast the perturbative approximation as a low-dimensional variational problem, allowing more efficient computation of steady states across parameter regions and a *wider convergence radius*. We propose two model-independent methods based on VPT that avoid computing the pseudo-inverse of the Liouvillian, both also enabling the differentiation of the steady state. First, exploiting the properties of the Liouvillian, we compute the steady state and arbitrarily high orders of the perturbative expansion through a single *LU* decomposition. Second, having in mind problems where finding exact matrix factorizations is challenging due to the Hilbert space size, we adapt this approach to construct a variational Krylov space through iterative methods. We demonstrate VPT’s efficiency by applying it to the driven-dissipative Kerr resonator, a two-mode dissipative cat setup, and the dissipative *XYZ* (Heisenberg) model. Compared to direct calculations on these models, our approach reduces the computational cost by up to a factor of a hundred. From a fundamental perspective, we show that efficient generalizations of standard perturbation theory are both possible and practical.

The paper is structured as follows. We begin in Sec. 2 with a brief review of how to determine steady states with matrix factorization techniques, followed by an overview of perturbation theory for open quantum systems. In Sec. 3, we present an efficient method

for computing the perturbative recurrence relation without relying on a pseudo-inverse. We then introduce VPT in Sec. 4 and use it to analyze the driven-dissipative Kerr resonator model in both one- and two-dimensional parameter spaces. In Sec. 5, we demonstrate how to use VPT for efficient parameter estimation and apply it to a dissipative cat qubit model. Finally, in Sec. 6, we extend VPT to an iterative Krylov-based approach and apply it to an *XYZ* model on a 3×3 lattice.

2 Open system steady states and how to compute them

We consider a generic open quantum system described by a time-independent Lindblad master equation of the form ($\hbar = 1$)

$$\partial_t \hat{\rho}(t) = \mathcal{L}\hat{\rho}(t) = -i[\hat{H}, \hat{\rho}(t)] + \sum_{\mu} \kappa_{\mu} \mathcal{D}[\hat{J}_{\mu}] \hat{\rho}(t), \quad (1)$$

where $\hat{H}$ is the Hamiltonian that describes the coherent evolution of the system and $\mathcal{D}[\hat{J}_{\mu}]$ are dissipator superoperators acting at a rate κ_{μ} via the jump operators $\hat{J}_{\mu}$ as

$$\mathcal{D}[\hat{J}_{\mu}] \hat{\rho}(t) = \hat{J}_{\mu} \hat{\rho}(t) \hat{J}_{\mu}^{\dagger} - \frac{1}{2} \left(\hat{J}_{\mu}^{\dagger} \hat{J}_{\mu} \hat{\rho}(t) + \hat{\rho}(t) \hat{J}_{\mu}^{\dagger} \hat{J}_{\mu} \right). \quad (2)$$

Here, $\mathcal{L}$ is the Liouvillian superoperator, generating the Lindblad dynamics. Upon representing $\mathcal{L}$ as a matrix, density matrices are mapped to column vectors that we denote with $|\hat{\rho}(t)\rangle$. We will slightly abuse notation by denoting both the superoperator and its matrix representation as $\mathcal{L}$: $\mathcal{L}\hat{\rho}(t)$ refers to the superoperator form and $\mathcal{L}|\hat{\rho}(t)\rangle$ to the matrix form.

2.1 Computing the steady state

Any open quantum system has at least one steady state $\hat{\rho}_{ss}$, which satisfies $\partial_t \hat{\rho}_{ss} = 0$ and $\hat{\rho}_{ss} = \hat{\rho}(t \rightarrow \infty)$ [3]. Equivalently,

$$\mathcal{L}\hat{\rho}_{ss} = 0, \quad (3)$$

i.e., the steady states span the kernel of $\mathcal{L}$. A common approach to compute $\hat{\rho}_{ss}$ is to directly solve Eq. (3). To avoid converging to the trivial solution $|\hat{\rho}_{ss}\rangle = |0, 0, \dots, 0\rangle$, the equation for the steady state can be modified as

$$\tilde{\mathcal{L}}|\hat{\rho}_{ss}\rangle = (\mathcal{L} + b\mathcal{T})|\hat{\rho}_{ss}\rangle = |b\rangle, \quad (4)$$

where

$$\mathcal{T}|\hat{\rho}_{ss}\rangle = |\text{Tr}(\hat{\rho}_{ss}), 0, \dots, 0\rangle, \quad (5)$$

and $|b\rangle = |b, 0, \dots, 0\rangle$ with b an arbitrary complex number (see Appendix A).

Equation (4) is solvable through standard matrix factorization techniques [44]. Although our results

are agnostic to the choice of factorization, we choose LU decomposition due to its numerical efficiency¹. We get

$$\tilde{\mathcal{L}}|\hat{\rho}_{ss}\rangle = LU|\hat{\rho}_{ss}\rangle = |b\rangle, \quad (6)$$

where L (U) is a lower (upper) triangular matrix. This equation can be solved in two steps (i) solve $L|\eta\rangle = |b\rangle$ to obtain $|\eta\rangle$ (forward substitution); (ii) solve $U|\hat{\rho}_{ss}\rangle = |\eta\rangle$ to obtain $|\hat{\rho}_{ss}\rangle$ (backward substitution). Since both L and U are triangular matrices, solving these systems requires only $\mathcal{O}(N^2)$ operations for a matrix of size $N \times N$. More advanced LU decompositions that preserve the sparsity structure also add steps of initial preconditioning on the $\mathcal{L}$ matrix, such as partial pivoting, further lowering the numerical cost of forward and backward substitutions.

In this work, we are interested in finding the steady state over a wide range of parameters, parameterized by the Liouvillian $\mathcal{L}(\theta)$ [cf. Fig. 1(a)]. We will benchmark the methods developed below against LU decomposition, whose number of operations to obtain the steady state of a Liouvillian of size $N \times N$ scales as $\mathcal{O}(PN^3)$, where P is the number of parameter combinations.

2.2 Perturbation theory

Rather than solving the problem for each parameter combination, one can use perturbation theory (PT) to approximate the steady state over contiguous regions in parameter space. Closely following Refs. [33–35], let us call $\mathcal{L}_0 = \mathcal{L}(\bar{\theta})$ for a specific choice $\theta = \bar{\theta}$. We introduce the Liouvillian in a neighborhood $\varepsilon = \theta - \bar{\theta}$ as

$$\mathcal{L}(\varepsilon) = \mathcal{L}_0 + \varepsilon \mathcal{L}_1, \quad (7)$$

where both $\mathcal{L}_0$ and $\mathcal{L}_1$ are Liouvillians. PT assumes that the density matrix $\hat{\rho}_{ss}$ can be expanded as

$$\hat{\rho}_{ss}^{PT}(\varepsilon) = \frac{1}{\mathcal{N}} \sum_{n=0}^{\infty} \varepsilon^n \hat{\rho}_{ss}^{(n)}, \quad (8)$$

where $\mathcal{N}$ is a normalization coefficient ensuring $\text{Tr}[\hat{\rho}_{ss}^{PT}(\varepsilon)] = 1$. As discussed in the Appendix A and sketched in Fig. 1(b), each term can be obtained through the recurrence relation,

$$\mathcal{L}_0 \hat{\rho}_{ss}^{(0)} = 0, \quad \mathcal{L}_0 \hat{\rho}_{ss}^{(n)} + \mathcal{L}_1 \hat{\rho}_{ss}^{(n-1)} = 0. \quad (9)$$

Since $\mathcal{L}_0$ has a nullspace spanned by $\hat{\rho}_{ss}$, Eq. (9) has an infinite number of solutions. The expansion in Eq. (8) corresponds to the choice

$$\hat{\rho}_{ss}^{(n)} = -\mathcal{L}_0^{-1} \mathcal{L}_1 \hat{\rho}_{ss}^{(n-1)}, \quad (10)$$

where here $\mathcal{L}_0^{-1}$ is the (Moore-Penrose) pseudo-inverse. Note that this choice of $\hat{\rho}_{ss}^{(n)}$ gives the solution

¹Another common choice is the QR decomposition. We use LU since its complexity scales as $\mathcal{O}(2N^3/3)$ for a system of size N whereas QR scales as $\mathcal{O}(4N^3/3)$.

with minimum norm and

$$\langle \hat{\rho}_{ss}^{(0)} | \hat{\rho}_{ss}^{(n)} \rangle = 0, \quad (11)$$

i.e., the pseudo-inverse removes the nullspace component from $\hat{\rho}_{ss}^{(n)}$. Within the radius of convergence of the series in Eq. (8), we can approximate the steady state by truncating the expansion to a desired order M .

It is straightforward to generalize Eq. (8) to multi-parameter Liouvillians $\mathcal{L}(\theta_1, \theta_2, \dots)$. For instance,

$$\mathcal{L}(\varepsilon_1, \varepsilon_2) = \mathcal{L}_0 + \varepsilon_1 \mathcal{L}_1 + \varepsilon_2 \mathcal{L}_2, \quad (12)$$

with $\mathcal{L}_0 = \mathcal{L}(\bar{\theta}_1, \bar{\theta}_2)$ and $\varepsilon_{1,2} = \theta_{1,2} - \bar{\theta}_{1,2}$, leads to a two-dimensional perturbative corrections grid [see Fig. 1(c)] defined by

$$\begin{aligned} \mathcal{L}_0 \hat{\rho}_{ss}^{(0,0)} &= 0, \\ \mathcal{L}_0 \hat{\rho}_{ss}^{(n,m)} + \mathcal{L}_1 \hat{\rho}_{ss}^{(n-1,m)} + \mathcal{L}_2 \hat{\rho}_{ss}^{(n,m-1)} &= 0. \end{aligned} \quad (13)$$

A crucial limitation of this approach is that constructing the pseudo-inverse is a numerically expensive operation, as it requires diagonalizing $\mathcal{L}_0$ or computing its singular value decomposition. We solve this issue in Sec. 3 and in Sec. 6. Furthermore, as we show in Secs. 4.1.1 and 4.1.2 and depict in Fig. 1(d), the convergence radius of the perturbation series is fairly limited especially in the proximity of critical phenomena such as phase transitions. In Sec. 4, we introduce a Variational Perturbation Theory (VPT) that circumvents these limitations.

3 Efficiently computing the recursive relation

Suppose the LU decomposition at a point $\bar{\theta}$ has been computed to obtain $\hat{\rho}_{ss}(\bar{\theta})$. We show how to reuse this factorization to solve the perturbative recurrence relation in Eq. (9). First, we introduce the modified recurrence relation

$$\tilde{\mathcal{L}}_0 |\hat{\rho}^{(0)}\rangle = |b\rangle, \quad \tilde{\mathcal{L}}_0 |\hat{\sigma}^{(n)}\rangle + \mathcal{L}_1 |\hat{\rho}_{ss}^{(n-1)}\rangle = 0. \quad (14)$$

Crucially, we can reuse L and U in Eq. (6) to solve this equation at a marginal numerical cost of $\mathcal{O}(N^2)$. Since time-evolution is trace-preserving, any Liouvillian admits the identity as a left eigenvector $\langle \mathbb{1} |$ with zero eigenvalue: $\langle \mathbb{1} | \mathcal{L} = 0$. Applying $\langle \mathbb{1} |$ to Eq. (14) and using $\langle \mathbb{1} | \mathcal{L}_{0,1} = 0$ [see also Eq. (48)], we obtain

$$b \langle \mathbb{1} | \mathcal{T} | \hat{\sigma}^{(n)} \rangle = 0. \quad (15)$$

Since $b \neq 0$, this implies $\text{Tr}[\hat{\sigma}^{(n)}] = 0$ and therefore $\mathcal{T} |\hat{\sigma}^{(n)}\rangle = 0$. Substituting back into Eq. (14) eliminates the $b\mathcal{T}$ contribution to $\tilde{\mathcal{L}}_0$, leaving

$$\mathcal{L}_0 |\hat{\sigma}^{(n)}\rangle + \mathcal{L}_1 |\hat{\rho}_{ss}^{(n-1)}\rangle = 0.$$




Figure 1: Depiction of the working principle of standard perturbation theory (PT), variational PT (VPT), and of the perturbative recursion relation. We are interested in finding steady states of a parameterized Liouvillian $\mathcal{L}(\varepsilon) = \mathcal{L}_0 + \varepsilon \mathcal{L}_1$ for many values of ε [or, more generally, $\mathcal{L}(\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n) = \mathcal{L}_0 + \sum \varepsilon_j \mathcal{L}_j$]. (a) Sketch of the parameter space. At the point $\varepsilon = 0$, the steady state $\hat{\rho}_{ss}(\varepsilon = 0)$ of $\mathcal{L}(\varepsilon = 0) = \mathcal{L}_0$ is computed. To compute the steady state in neighboring points in the parameter space, we resort to PT and find the set of matrices $\hat{\rho}_{ss}^{(n)}$ defined in Eq. (9). We develop methods that do that efficiently for (b) 1D parameter spaces, (c) 2D parameter spaces and, more generally, for arbitrarily high dimensional parameter space. These methods are based on either re-applying the *LU* decomposition used to compute $\hat{\rho}_{ss}(\varepsilon = 0)$, or using approximate methods to build the basis of matrices defining the steady state. (d) Having obtained the series of perturbation matrices $\hat{\rho}_{ss}^{(n)}$, one can construct the state $\hat{\rho}_{ss}(\varepsilon)$ by weighting each matrix according to the power series in Eq. (8) within the radius of convergence. (e) Allowing for a more expressive ansatz in the form of Eq. (18), one can use $\hat{\rho}_{ss}^{(n)}$ to describe the solution on a much wider range of parameters. This makes VPT more efficient than PT, as the “heavy” operation of computing the steady state at one point via exact numerical methods becomes less frequent.

We conclude that, up to a nullspace component, $|\hat{\rho}_{ss}^{(n)}\rangle$ and $|\hat{\sigma}^{(n)}\rangle$ are identical. Using Eq. (11) we finally obtain

$$|\hat{\rho}^{(n)}\rangle = |\hat{\sigma}^{(n)}\rangle - \langle \hat{\rho}^{(0)} | \hat{\sigma}^{(n)} \rangle |\hat{\rho}^{(0)}\rangle. \quad (16)$$

In summary, although the initial *LU* factorization of $\hat{\mathcal{L}}_0$ is an expensive operation, it enables us to compute the steady state $\hat{\rho}_{ss}^{(0)}$ and all perturbative corrections $\hat{\rho}_{ss}^{(n)}$ at minimal additional cost without computing the pseudo-inverse of $\mathcal{L}_0$. Using this approach, the cost to compute the steady state across a phase space of a Liouvillian of size $N \times N$ scales as

$$\mathcal{O} \left[\frac{P}{R_{PT}} (N^3 + M^d N^2) \right] \simeq \mathcal{O} \left(\frac{P}{R_{PT}} N^3 \right), \quad (17)$$

where R_{PT} is the typical convergence volume of the PT series, P is the number of points to compute, and M is the perturbation order. The right-hand side follows in the limit of $M^d \ll N$ and highlights the importance of maximizing the radius of convergence.

4 Multipoint variational perturbation theory for open quantum systems

PT has a fairly limited radius of convergence due to the overly stringent condition on coefficients that leads to Eq. (8) [see also the discussion in Appendix B] and not because of a lack of expressivity of the vectors $\hat{\rho}_{ss}^{(n)}$. Indeed, they are the Krylov basis of $\mathcal{L}_0^{-1} \mathcal{L}_1$, and thus form a basis for any finite-dimensional density matrix except in pathological cases. We thus generalize the series in Eq. (8) to

$$\hat{\rho}_{ss}^{PT}(\varepsilon) = \frac{1}{\mathcal{N}} \sum_{n=0}^{\infty} c_n(\varepsilon) \hat{\rho}_{ss}^{(n)}, \quad (18)$$

with $\mathcal{N}$ the normalization term. Within the region where perturbation theory is valid, Eqs. (8) and (18) must give the same result. Therefore, a natural choice is to fix $\hat{\rho}_{ss}^{(n)}$ to be identical in both expressions. In turn, this implies that $\hat{\rho}_{ss}^{(n)}$ can be efficiently determined using Eq. (9) and exploiting the strategy de-

veloped in Sec. 3. This leaves the task of determining the coefficients $c_n(\varepsilon)$. While within the radius of convergence of PT we expect $c_n(\varepsilon) \approx \varepsilon^n$ to recover the results of PT, the general form of the coefficients could be determined by assuming c_n to be smooth functions of ε , and then solving

$$\left\| \tilde{\mathcal{L}}(\varepsilon) \sum_{n=0}^{\infty} c_n(\varepsilon) \left| \hat{\rho}_{\text{ss}}^{(n)} \right\rangle - |b\rangle \right\|^2 = 0. \quad (19)$$

order-by-order. This is, however, a cumbersome procedure. Instead, we adopt a variational approach and approximate $c_n(\varepsilon)$ solving up to order M the least-squares problem

$$\min_{c_n} \left\| \tilde{\mathcal{L}}(\varepsilon) \sum_{n=0}^M c_n(\varepsilon) \left| \hat{\rho}_{\text{ss}}^{(n)} \right\rangle - |b\rangle \right\|^2. \quad (20)$$

That is, we determine the optimal low-dimensional approximation of the steady state within the span of $\{\hat{\rho}_{\text{ss}}^{(0)}, \hat{\rho}_{\text{ss}}^{(1)}, \dots, \hat{\rho}_{\text{ss}}^{(M)}\}$. The variational nature of Eq. (20) ensures that this *variational* PT (VPT) provides a better approximation than standard PT for any finite order M and hence has a larger convergence radius, as we illustrate in Figs. 1(a) and (e).

As the number of basis elements M increases, $\left| \hat{\rho}_{\text{ss}}^{(M)} \right\rangle$ defined in eq. (10) converges to the top eigenvector of the update operator; this is the power method. Numerically, it is thus convenient to work in a basis obtained by orthonormalizing $\{\left| \hat{\rho}_{\text{ss}}^{(0)} \right\rangle, \left| \hat{\rho}_{\text{ss}}^{(1)} \right\rangle, \dots, \left| \hat{\rho}_{\text{ss}}^{(M)} \right\rangle\}$. Denoting the matrix of orthonormal operators as $\mathcal{Q}$, we rewrite the problem as

$$\min_{q_n} \left\| \tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q} - |b\rangle \right\|^2, \quad (21)$$

where $\vec{q} = (q_0, q_1, \dots, q_M)^T$ and q_n are the coefficients in the orthonormal basis.

Minimizing Eq. (21) is faster than finding the LU decomposition to solve $\mathcal{L}(\theta) \hat{\rho}_{\text{ss}}(\theta) = 0$, but can still be numerically costly since the matrix $\tilde{\mathcal{L}}(\varepsilon) \mathcal{Q}$ has size $N \times (M+1)$ for a Liouvillian of size $N \times N$. To further reduce computational costs we observe that

$$\begin{aligned} \left\| \tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q} - |b\rangle \right\|^2 &= \left\| (\mathcal{I} - \mathcal{Q} \mathcal{Q}^\dagger) (\tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q} - |b\rangle) \right\|^2 \\ &\quad + \left\| \mathcal{Q}^\dagger (\tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q} - |b\rangle) \right\|^2, \end{aligned} \quad (22)$$

where $\mathcal{I}$ is the identity superoperator. Equation Eq. (22) is an equality because $\mathcal{I} - \mathcal{Q} \mathcal{Q}^\dagger$ and $\mathcal{Q} \mathcal{Q}^\dagger$ are orthogonal projectors onto complementary subspaces, and we used $\|\mathcal{Q} \mathcal{Q}^\dagger v\| = \|\mathcal{Q}^\dagger v\|$, which follows from the orthonormality of the columns of $\mathcal{Q}$. Rather than minimizing the full residual, we choose to null only the second term $\|\mathcal{Q}^\dagger (\tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q} - |b\rangle)\|$, which vanishes exactly when

$$\tilde{\mathcal{L}}_{\text{eff}}(\varepsilon) \vec{q}_{\text{eff}} = \vec{b}_{\text{eff}}, \quad (23)$$

where $\tilde{\mathcal{L}}_{\text{eff}}(\varepsilon) = \mathcal{Q}^\dagger \tilde{\mathcal{L}}(\varepsilon) \mathcal{Q}$ has dimension $(M+1) \times (M+1)$ and $\vec{b}_{\text{eff}} = \mathcal{Q}^\dagger |b\rangle$. Substituting $\vec{q}_{\text{eff}}$ into Eq. (22), the residual reduces to its component orthogonal to $\text{range}(\mathcal{Q})$:

$$\begin{aligned} &\left\| \tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q}_{\text{eff}} - |b\rangle \right\|^2 \\ &= \left\| (\mathcal{I} - \mathcal{Q} \mathcal{Q}^\dagger) (\tilde{\mathcal{L}}(\varepsilon) \mathcal{Q} \vec{q}_{\text{eff}} - |b\rangle) \right\|^2. \end{aligned} \quad (24)$$

In the limit in which $\mathcal{Q}$ is a good basis to represent the steady state, this residual is small, and $\vec{q}_{\text{eff}}$ provides a good approximation to the solution of Eq. (21). Notably, since M is typically much smaller than the Hilbert space dimension, solving this reduced equation introduces a significant speedup at the price of finding a slightly suboptimal solution. In practice, having solved the reduced system for $\vec{q}_{\text{eff}}$, we lift the solution back to the full space as $|\hat{\rho}_{\text{ss}}^{\text{PT}}(\varepsilon)\rangle = \mathcal{Q} \vec{q}_{\text{eff}} / \mathcal{N}$, where $\mathcal{N}$ is the normalization coefficient ensuring $\text{Tr}[\hat{\rho}_{\text{ss}}^{\text{PT}}(\varepsilon)] = 1$, and directly evaluate the residual of the original problem with a single matrix-vector product against the full Liouvillian:

$$\|\mathcal{L}(\varepsilon) |\hat{\rho}_{\text{ss}}^{\text{PT}}(\varepsilon)\rangle\| \leq \text{tol}. \quad (25)$$

This is a common type of criterion, also used in other variational approaches [19, 32]. The tolerance should be carefully chosen, especially in the presence of critical phenomena (see Appendix C), as lower tolerance implies smaller convergence regions, but greater precision.

Using VPT the cost of computing a phase diagram in a d -dimensional parameter space with P points for a $N \times N$ Liouvillian scales as

$$\begin{aligned} &\mathcal{O} \left[\frac{P}{R_{\text{VPT}}} (N^3 + M^d N^2) + P M^{3d} \right] \\ &\simeq \mathcal{O} \left(\frac{P}{R_{\text{VPT}}} N^3 \right), \end{aligned} \quad (26)$$

where R_{VPT} is the typical size of the region where VPT up to order M is valid, and the latter follows under the assumption $M^d \ll N$. Since $R_{\text{VPT}} \geq R_{\text{PT}}$, using VPT leads to an increase in performance.

4.1 Variational vs standard perturbation theory in the driven Kerr resonator

We consider a driven-dissipative Kerr resonator, sketched in Fig. 2(a), whose Hamiltonian, in the frame rotating at the pump frequency, reads

$$\hat{H} = -\Delta \hat{a}^\dagger \hat{a} - K/2 \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} + F(\hat{a} + \hat{a}^\dagger). \quad (27)$$

Here, $\hat{a}$ ($\hat{a}^\dagger$) is the bosonic annihilation (creation) operator. Δ is the pump-to-cavity detuning, K is the Kerr nonlinearity, and F is the pump amplitude. The system is subject to single-photon loss events occurring at a rate κ , with a Liouvillian reading

$$\mathcal{L} \hat{\rho}(t) = -i[\hat{H}, \hat{\rho}(t)] + \kappa \mathcal{D}[\hat{a}] \hat{\rho}(t). \quad (28)$$




Figure 2: Comparison of standard PT and VPT in the study of the driven-dissipative Kerr resonator described by Eq. (28) and sketched in (a). (b) Average photon number $\langle \hat{a}^\dagger \hat{a} \rangle$ as a function of the detuning Δ determined using an exact solution (black line), standard PT (yellow line), and VPT (blue line). Both PTs have been performed starting from $\Delta = 0$, as indicated by the dashed vertical line. The solid background denotes the region where the perturbative expansion has converged up to $\|\mathcal{L}(\varepsilon)\hat{\rho}_{ss}^{PT}(\varepsilon)\| < 10^{-2}$. In the yellow region, both standard PT and VPT are convergent. In the blue region, only VPT reached convergence. (c) Error $\|\mathcal{L}(\varepsilon)\hat{\rho}_{ss}^{PT}(\varepsilon)\|$ as a function of the maximal order of perturbation M for standard PT. (d) Same as (c), but for VPT. (e-h) Coefficients obtained by standard [c.f. Eq. (8)] and variational [c.f. Eq. (18)] PTs. Parameters: $K/\kappa = 10$ and $F/\kappa = 10$. Fock space truncation: 30 photons.

4.1.1 1D example: varying detuning

First, we vary Δ and fix all other parameters. We choose $\Delta = 0$ as the initial point for PT and VPT. In Fig. 2(b) we compare the photon number in the steady state $\langle \hat{a}^\dagger \hat{a} \rangle = \text{Tr}[\hat{a}^\dagger \hat{a} \hat{\rho}_{ss}]$ obtained from the exact solution to the problem [18], standard PT, and VPT. The shaded backgrounds indicates the region where $\|\mathcal{L}\hat{\rho}_{ss}^{PT}\| < \text{tol}$ for each PT approach. We see that VPT recovers the exact solution in a broader region than standard PT, from the vacuum at large detuning, to the coherent-like state at $\Delta \simeq 0$, passing through the multiphoton resonances at intermediate detuning. Notice also that VPT is capable of finding the solutions in disconnected regions of the parameter space. In Figs. 2(c-d) we investigate the effect of increasing the order M of PT. While for standard PT larger M result in a modest gain in precision and in the region of validity [c.f. Fig. 2(c)], for VPT the region of validity significantly broadens with M , as shown in Fig. 2(d). Finally, in Figs. 2(e-h) we show the first coefficients of the PT series. As expected,

VPT coincides with standard PT close to $\Delta = 0$.

This example highlights the advantages of variational perturbation theory, which will become even more pronounced for larger Hilbert spaces. Indeed, the effective Liouvillian $\tilde{\mathcal{L}}_{\text{eff}}(\Delta)$ in Fig. 2(b) has size 21×21 . Achieving the same accuracy with a standard truncation of the Hilbert space in the Fock basis requires keeping states up to $|n = 8\rangle$, i.e., a Liouvillian of size 81×81 .

4.1.2 2D example: simultaneously varying detuning and drive amplitude

We now apply VPT to compute steady states in a region where Δ and F change. Compared to the previous case, we consider a larger value of κ/K , showcasing the efficiency of VPT in a more dissipative configuration. We plot the average photon number in the steady state in Fig. 3(a). At large positive detuning, the steady state is approximately the vacuum. Upon decreasing Δ , the system smoothly transitions out of the vacuum phase and gets populated. At




Figure 3: Comparison of the performance of perturbation theory methods to compute the steady state of the driven-dissipative Kerr resonator defined in Eq. (28) in a two-dimensional parameter space. In (a) we plot the average photon number $\langle a^\dagger a \rangle$ in the steady state as a function of detuning Δ and drive strength F along with the boundaries of the convergence regions of multipoint VPT. In (b) we plot the points where we exactly computed the steady state through LU decomposition. At each point we computed a grid of perturbative corrections up to $M = (10, 10)$, totaling 121 perturbation vectors per point, and set the error tolerance to $\|\mathcal{L}(\varepsilon)\hat{\rho}_{ss}^{PT}(\varepsilon)\| < 10^{-7}$. Parameters: $K/\kappa = \frac{1}{2}$. Fock space truncation: 70 photons.

large negative detuning, the system abruptly transitions from a high-photon phase to a vacuum-like one. As the drive amplitude increases, this transition becomes more pronounced, signaling the emergence of a dissipative phase transition [14, 45, 46]. Compared to the previous example, accurately capturing the phase transition will require a significantly smaller tolerance (see the discussion in Appendix C).

To cover the parameter space, we first select a random point $(\bar{\Delta}, \bar{F})$ and compute the steady state and its perturbative corrections up to order M , introducing the perturbation parameters $\varepsilon_1 = \Delta - \bar{\Delta}$ and $\varepsilon_2 = F - \bar{F}$. Next, we use (V)PT to construct approximate steady states in the neighborhood of the point, ensuring the error remains below the specified tolerance for $\|\mathcal{L}\hat{\rho}_{ss}^{PT}(\varepsilon_1, \varepsilon_2)\|$. Once we have found the convergence region of (V)PT, we randomly select a new point and repeat the process until the parameter space is completely covered.

In Fig. 3(b) we compare the performance of PT and VPT. Specifically, we plot all the points where we recompute the exact steady state and its perturbative corrections. We observe that standard PT requires a




Figure 4: Performance of perturbation theory methods with increasing perturbation order M . We study the same system as in Fig. 3 and apply (V)PT at 40 randomly selected points in parameter space. For each resulting convergence region, we compute the optimal low-rank basis (see Appendix D). In both panels, blue violins correspond to VPT and yellow violins to PT, with scatter points indicating individual samples. (a) Violin plot of the convergence region area as a function of the VPT basis size $(M+1)^2$. (b) Violin plot of $(M+1)^2/M_{\text{SVD}}$ as a function of $(M+1)^2$, where $(M+1)^2$ is the number of perturbation vectors used in (V)PT, and M_{SVD} is the size of the optimal low-rank basis. A higher ratio indicates a more efficient representation relative to the optimal basis. Violin plots illustrate data distributions, with the width representing the density of data points at each value, estimated using a Gaussian kernel.

high density of points near the phase transition region. This is a consequence of the non-analytic behavior of the system near the critical points where the Liouvillian gap closes. Strikingly, VPT significantly outperforms PT and requires seven times fewer points to map the whole phase space [cf. also the inset in Fig. 3(a)]. To provide a more quantitative analysis of the performance of VPT and PT, we randomly select 40 points in parameter space where we apply (V)PT and compute the size of the convergence region with increasing perturbation order M . The resulting size distributions, histogrammed in Fig. 4(a), indicate that VPT consistently yields significantly larger regions of convergence than PT. In Fig. 4(b) we estimate how efficiently (V)PT compresses the steady states by comparing $(M+1)^2$ (the number of perturbation vectors used in (V)PT), to M_{SVD} —the minimal number of vectors required to span the same convergence region with the same precision as VPT (see details in Appendix D). Our results indicate that VPT achieves a compression efficiency close to the theoretical limit. As the perturbation order M increases,

both PT and VPT become less efficient in terms of compression. However, a higher M also extends the coverage of the parameter space, allowing the method to reach a broader range of solutions. This reveals an inherent trade-off: increasing the perturbation order enhances parameter space coverage but reduces the optimality of the low-rank representation.

4.2 Multipoint variational perturbation theory

The performance of VPT deteriorates in the vicinity of critical points. To gain intuition into why it breaks down near critical points, we consider a system with a Liouvillian parameterized by a single parameter θ that transitions at $\theta = 0$. Suppose we computed the steady state and the perturbative series at two points: $\theta_- < 0$ and $\theta_+ > 0$ on the opposite sides of the critical point. Using standard PT, we have

$$\hat{\rho}_{ss}^{PT, \pm}(\varepsilon_{\pm}) = \sum_n \varepsilon_{\pm}^n \hat{\rho}_{ss}^{(n)}(\theta_{\pm}), \quad (29)$$

where $\varepsilon_{\pm} = (\theta - \theta_{\pm})$. None of the states $\hat{\rho}_{ss}^{PT, \pm}$ will individually describe points near the phase transition, because, roughly speaking, the steady state switches between $\hat{\rho}_{ss}(\theta < 0)$ and $\hat{\rho}_{ss}(\theta > 0)$. The ansatz

$$\hat{\rho}_{ss}(\theta) \approx A(\varepsilon_-) \hat{\rho}_{ss}^{PT, -}(\varepsilon_-) + B(\varepsilon_+) \hat{\rho}_{ss}^{PT, +}(\varepsilon_+), \quad (30)$$

instead, can do it, with $A(\varepsilon_-)$ and $B(\varepsilon_+)$ the coefficients of the left and right steady states, respectively.

Leveraging this intuition, we propose a multipoint variational perturbation theory (m-VPT) based on the perturbative vectors computed at several points:

$$\hat{\rho}_{ss}^{PT}(\theta) = \sum_{n=0}^{\infty} \sum_i c_n^{(i)} (\theta - \theta_i) \hat{\rho}_{ss}^{(n)}(\theta_i). \quad (31)$$

We repeat the steady state calculations in Fig. 3 where for each point we consider a basis $\mathcal{Q}$ that contains the local PT correction vectors but also those of the two nearest points. Compared to VPT, we find that this approach achieves a further three-fold reduction in the number of points where $\hat{\rho}_{ss}$ is calculated through LU , as shown in the inset of Fig. 3(a). In the same panel, we also plot the edges of the convergence regions for this m-VPT, showing they cross the boundary between the high- and low-photon-number phases.

5 Parameter estimation through efficient gradient computation

A common task in operating quantum devices is fitting numerical simulations to experimental results to estimate system parameters. Here we show how to efficiently use VPT for this task.

Mathematically, we can frame the problem using a Liouvillian $\mathcal{L}(\phi, \theta)$, where ϕ represents an unknown parameter and θ is the controllable parameter of the

experiment (the following equations can be easily generalized to an arbitrary number of fixed and controllable parameters). We assume that we have access to a set of noisy measurements $\{o_{\text{exp}}(\theta)\}$ associated with the operator $\hat{o}$, whose expectation value $\langle \hat{o}(\phi, \theta) \rangle = \text{Tr}[\hat{\rho}_{ss}(\phi, \theta) \hat{o}]$ we can compute numerically. The objective of a fit routine is then to find ϕ^{fit} that minimizes the difference between the experimental and numerical data:

$$\phi^{\text{fit}} = \min_{\phi} \sum_{\theta} |\langle \Delta \hat{o}(\phi, \theta) \rangle|^2 = \min_{\phi} \mathcal{C}(\phi), \quad (32)$$

where $\langle \Delta \hat{o}(\phi, \theta) \rangle = \langle \hat{o}(\phi, \theta) \rangle - o_{\text{exp}}(\theta)$ and $\mathcal{C}(\phi)$ is the cost function. One can use gradient-based methods to iteratively find ϕ^{fit} : at each step n of the method, a guess ϕ_n is updated based on $\partial_{\phi} \mathcal{C}(\phi)$. Introducing $\delta = \phi - \phi_n$ such that

$$\mathcal{L}(\theta, \phi) = \mathcal{L}(\bar{\theta}, \phi_n) + \varepsilon \mathcal{L}_1 + \delta \mathcal{L}_2, \quad (33)$$

we get

$$\begin{aligned} \left. \frac{\partial \mathcal{C}(\phi)}{\partial \phi} \right|_{\phi=\phi_n} &= \left. \frac{\partial \mathcal{C}(\phi)}{\partial \delta} \right|_{\delta=0} \\ &= \sum_{\theta} \langle \Delta \hat{o}(\phi_n, \theta) \rangle \text{Tr} \left[\left. \frac{\partial \hat{\rho}_{ss}(\phi_n + \delta, \theta)}{\partial \delta} \right|_{\delta=0} \hat{o} \right]. \end{aligned} \quad (34)$$

We thus need to determine $\partial_{\delta} \hat{\rho}_{ss}(\varepsilon, \delta)|_{\delta=0}$.

A first way to do this is to assume

$$|\hat{\rho}_{ss}^{PT}(\varepsilon, \delta)\rangle = \frac{1}{\mathcal{N}} \sum_n c_n(\varepsilon) \left(\hat{\rho}_{ss}^{(n,0)} + \delta \hat{\rho}_{ss}^{(n,1)} \right). \quad (35)$$

That is, we use the VPT ansatz to describe the dependence in θ and assume that the coefficients $c_n(\varepsilon)$ are independent of δ . We get

$$\begin{aligned} \frac{\partial \hat{\rho}_{ss}(\phi_n + \delta, \theta)}{\partial \delta} &= \frac{\mathcal{N} \sum_n c_n(\varepsilon) \hat{\rho}_{ss}^{(n,1)} - (\partial_{\delta} \mathcal{N}) |\hat{\rho}_{ss}^{PT}(\theta, \delta)\rangle}{\mathcal{N}^2}, \end{aligned} \quad (36)$$

where

$$\partial_{\delta} \mathcal{N} = \text{Tr} \left[\sum_n c_n(\varepsilon) \hat{\rho}_{ss}^{(n,1)} \right]. \quad (37)$$

A better estimate of the gradient is obtained by allowing all coefficients to simultaneously depend on δ and ε , namely

$$\begin{aligned} |\hat{\rho}_{ss}^{PT}(\varepsilon, \delta)\rangle &= \frac{1}{\mathcal{N}} \left(\sum_n c_{n,0}(\varepsilon, \delta) \hat{\rho}_{ss}^{(n,0)} + c_{n,1}(\varepsilon, \delta) \hat{\rho}_{ss}^{(n,1)} \right). \end{aligned} \quad (38)$$

Finding the gradient then amounts to solving the following reduced equation, obtained by implicitly differentiating the linear constraint $\tilde{\mathcal{L}} |\hat{\rho}_{ss}\rangle = |b\rangle$ through

the VPT ansatz (see Appendix E for a derivation, and [47] for background on implicit differentiation in numerical analysis):

$$\begin{aligned} (\mathcal{L}_0^{\text{eff}} + \varepsilon \mathcal{L}_1^{\text{eff}}) \frac{\partial \tilde{q}(\delta, \varepsilon)}{\partial \delta} \bigg|_{\delta=0} &= -\mathcal{L}_2^{\text{eff}} \tilde{q}(0, \varepsilon) \\ \Rightarrow \frac{\partial |\hat{\rho}_{\text{ss}}^{PT}(\delta, \varepsilon)\rangle}{\partial \delta} \bigg|_{\delta=0} &= \mathcal{Q}_\delta \frac{\partial \tilde{q}(\delta, \varepsilon)}{\partial \delta} \bigg|_{\delta=0}, \end{aligned} \quad (39)$$

where $\mathcal{Q}_\delta$ is the basis obtained by orthonormalizing $\{\hat{\rho}_{\text{ss}}^{(n,0)}, \hat{\rho}_{\text{ss}}^{(n,1)}\}$ and $\mathcal{L}_j^{\text{eff}} = \mathcal{Q}_j^\dagger \mathcal{L}_j \mathcal{Q}_j$.

Estimating the parameters of a Schrödinger cat by fitting buffer spectroscopy measurements

We apply VPT and the gradient estimate in Eq. (39) on numerically generated data to estimate parameters in a superconducting device that hosts Schrödinger cat states [43]. The system, which we sketch in Fig. 5(a), consists of a memory mode $\hat{a}$ and a buffer mode $\hat{b}$ that are parametrically coupled through a term g_2 that converts a photon in the buffer into two photons in the memory. The Hamiltonian reads

$$\begin{aligned} \hat{H} = & -\Delta_a \hat{a}^\dagger \hat{a} - \Delta_b \hat{b}^\dagger \hat{b} + g_2 (\hat{a}^2 \hat{b}^\dagger + \text{h.c.}) + (F \hat{b} + \text{h.c.}) \\ & - K_a \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} - K_b \hat{b}^\dagger \hat{b}^\dagger \hat{b} \hat{b} + \chi \hat{a}^\dagger \hat{a} \hat{b}^\dagger \hat{b}, \end{aligned} \quad (40)$$

where Δ_a (Δ_b) is the detuning between the drive and memory (buffer), K_a (K_b) is the Kerr nonlinearity, and F is the drive amplitude acting on $\hat{b}$. Both modes dissipate photons at rates κ_a and κ_b , and the Lindblad master equation reads

$$\mathcal{L}\hat{\rho}(t) = -i[\hat{H}, \hat{\rho}(t)] + \kappa_a \mathcal{D}[\hat{a}] + \kappa_b \mathcal{D}[\hat{b}]. \quad (41)$$

In an ideal configuration to generate cats, $\Delta_a = \Delta_b = 0$, and $1/\kappa_b$ is the shortest timescale. Schrödinger cat states persist over times determined by the loss of parity at a rate $\kappa_a \langle \hat{a}^\dagger \hat{a} \rangle$ [48, 49].

Several parameters in Eq. (40) can be directly measured in experiments [43]. The detuning Δ_a and Δ_b are controlled by the relative frequencies of the driving mechanisms. K_b and κ_b can be inferred by probing the $\hat{b}$ mode that has a dedicated measurement and feedline. Since it is possible to prepare the states $|0\rangle$ or $|1\rangle$ in the memory through the combined action of drives and dissipation, χ can be inferred by comparing spectroscopy measurements on the $\hat{b}$ mode with the memory in states $|0\rangle$ or $|1\rangle$. Similarly, κ_a can be extracted from the decay rate of a Fock state $|1\rangle$. The remaining parameters are not straightforward to measure and require more advanced fitting procedures. Although all these parameters affect the spectroscopic response of the buffer mode, there are no closed or analytical expressions that one can use to fit the system's response. This motivates us to pursue

a simulation-based approach to fit the response of the system and indirectly extract these parameters.

In experiments, a common measurement is the spectroscopic response of the buffer given by $S_{11} = 1 - i\kappa_b \langle \hat{b} \rangle / F$ [43]. First, we pick a set of ground-truth values for g_2 , F , and K_a and numerically generate a dataset of S_{11} as a function of Δ_a and Δ_b . To mimic experimental imperfections, we add Gaussian noise to the data, as shown in Fig. 5(b). Then, starting from an initial random guess of g_2 , F , and K_a and assuming the remaining parameters are known, we try to recover the parameters used in the initial simulation using the scipy implementation [50] of the L-BFGS optimizer [51]. At each step, we minimize the cost function in Eq. (32) using the VPT gradient estimate provided in Eq. (39). We show the map of S_{11} computed at various steps of the optimization in Figs. 5(c-e). As shown in Figs. 5(f-i), we recover the parameters with high accuracy within 15 iterations, although noise introduces small deviations and prevents the loss function from reaching zero.

6 Variational perturbation theory as a preconditioned Krylov method

Computing exact matrix factorizations of the Liouvillian becomes infeasible for large Hilbert spaces. An alternative approach is to use iterative procedures that rely solely on matrix-vector products, such as Krylov-subspace-based methods. Starting from an initial guess $|q_0\rangle$, these methods iteratively build the Krylov basis of the Liouvillian $\{|q_k\rangle\} = \{\tilde{\mathcal{L}}^k |q_0\rangle\}$ for $k = 0, 1, \dots, M$ and search for an approximate solution of Eq. (4) in the form

$$|\hat{\rho}_{\text{ss}}\rangle = \sum_{k=0}^M c_k |q_k\rangle. \quad (42)$$

As a rule of thumb, the closer $|q_0\rangle$ to the actual steady state, the smaller M needs to be. Common iterative solvers include GMRES [52] and BiCGSTAB [53], and Conjugate Gradient descent when the linear operator is Hermitian [54].

The convergence rate of these iterative methods depends on the ratio between the largest and smallest singular value of $\tilde{\mathcal{L}}$, a quantity known as the condition number. For ill-conditioned problems, one can apply a *preconditioner* $\mathcal{C}$ and solve

$$\mathcal{C} \tilde{\mathcal{L}} |\hat{\rho}_{\text{ss}}\rangle = \mathcal{C} |b\rangle. \quad (43)$$

A good preconditioner improves the condition number while being inexpensive to compute and evaluate. Commonly used preconditioners include band-limited matrices and sparse approximate inverses (see [55] for a general review and [56] for an in-depth discussion on open quantum systems). Since the Liouvillian $\tilde{\mathcal{L}}_0$ is often sparse, a natural choice is to use incomplete




Figure 5: Parameter estimation using VPT and gradient-based optimization for a memory-buffer system sketched in (a) and described by the Liouvillian in Eq. (41). (b) Synthetic data of the steady state reflection coefficient S_{11} as a function of Δ_a and Δ_b . To mimic experimental imperfections, we add Gaussian noise with average $\mu = 0$ and standard deviation $\sigma = 0.02$. (c-e) Reflection coefficient at various stages of the optimization. The algorithm starts from a random guess of g_2 , F , and K_a and uses the L-BFGS algorithm to optimize them, assuming the remaining parameters are known. In (f) we plot the loss as a function of the iteration number, while (g-i) show how the parameters evolve along the optimization process. The dashed lines indicate the true parameters used to generate the data. Parameters (in MHz): $g_2 = 2$, $F = 2$, $K_a = 0.1$, $K_b = 0.3$, $\kappa_a = 0.1$, $\kappa_b = 10$. We fix the Fock space truncation at 10 photons for mode $\hat{a}$ and 6 for $\hat{b}$.

LU (iLU), that performs an approximate LU decomposition restricting the outcome to a sparsity pattern closely matching that of $\tilde{\mathcal{L}}_0$ (i.e., some matrix elements are set to zero).

6.1 Method and connection with Variational Perturbation Theory

Our second approach for solving Eq. (4) is to use $\tilde{\mathcal{L}}_0^{-1}(\varepsilon = 0)$ as a preconditioner and the corresponding steady state $\hat{\rho}_{ss}(\varepsilon = 0)$ as $|q_0\rangle$. For a given M , we construct the Krylov basis and use it to compute the steady states for the neighboring parameters. We repeat this procedure on all points where convergence within a fixed tolerance is reached. Then, we recompute both the steady state and the preconditioner $\tilde{\mathcal{L}}_0^{-1}(\varepsilon' = 0)$.

To understand why $\tilde{\mathcal{L}}_0^{-1}$ and $|\hat{\rho}_{ss}(\varepsilon = 0)\rangle$ are good candidates for preconditioning and Krylov subspaces construction, we note that

$$\begin{aligned} & \text{Span} \left\{ \left[\tilde{\mathcal{L}}_0^{-1}(\tilde{\mathcal{L}}_0 + \varepsilon \mathcal{L}_1) \right]^k |\hat{\rho}_{ss}(\varepsilon = 0)\rangle \right\} \\ &= \text{Span} \left\{ (\mathcal{I} + \varepsilon \tilde{\mathcal{L}}_0^{-1} \mathcal{L}_1)^k |\hat{\rho}_{ss}(\varepsilon = 0)\rangle \right\} \quad (44) \\ &= \text{Span} \left\{ (\tilde{\mathcal{L}}_0^{-1} \mathcal{L}_1)^k |\hat{\rho}_{ss}(\varepsilon = 0)\rangle \right\}. \end{aligned}$$

Therefore, this preconditioned Krylov method and VPT in Eq. (10) give the same result. This connects VPT with the larger body of work on recycled Krylov methods, widely used for finite element analysis [57–59]. They differ in that VPT recycles the whole subspace, while those methods recycle it portion by portion.

In practice, one rarely has access to $\tilde{\mathcal{L}}_0^{-1}(\varepsilon = 0)$, and computing it through LU decomposition would defy the purpose of using iterative methods. We instead compute the iLU decomposition of the Liouvillian that, despite being a worse preconditioner, is significantly cheaper to compute.

Finally, to further speed up the convergence of the Krylov space, one can *warm start* the algorithm by updating $|q_0\rangle$ using the nearest computed steady state solution. However, this comes at the cost of re-computing the Krylov subspace. Despite not being as costly as computing $\mathcal{L}_0^{-1}$, this still poses a trade-off between the convergence rate and the advantage of reusing a previous Krylov subspace.




Figure 6: Study of the dissipative XYZ model using preconditioned Krylov methods. (a) Pictographic representation of the XYZ model described by Eq. (46). Two-level systems interact with their nearest neighbours, with an anisotropic coupling along the three main axes of the Bloch sphere. Dissipative events can flip any spin and push it to point in the negative z direction. (b) steady state phase diagram of the magnetization $m_z = \langle \hat{\sigma}_z \rangle$ for a 3×3 lattice.

6.2 Phase diagram of the dissipative XYZ model

The XYZ model describes a set of two-level systems arranged in a square lattice and interacting according to the anisotropic Heisenberg Hamiltonian. In dimension $D = 2$, it reads

$$\hat{H} = \sum_{\langle \vec{n}, \vec{m} \rangle} (J_x \hat{\sigma}_n^x \hat{\sigma}_m^x + J_y \hat{\sigma}_n^y \hat{\sigma}_m^y + J_z \hat{\sigma}_n^z \hat{\sigma}_m^z), \quad (45)$$

where each spin is indexed with a vector $\vec{n} = (n_x, n_y)$, $\hat{\sigma}_i^\alpha$ ($\alpha = x, y, z$) are the Pauli matrices acting on the $\vec{n}$ -th site, and the summation includes only nearest neighbor spin pairs $\langle \vec{n}, \vec{m} \rangle$. The coefficients J_α denote the spin-spin interaction strengths. Dissipation manifests as incoherent spin flips at a rate γ that force spins towards the negative z -axis direction, with

$$\frac{\partial \hat{\rho}(t)}{\partial t} = -i [\hat{H}, \hat{\rho}(t)] + \gamma \sum_{\vec{n}} \mathcal{D}[\hat{\sigma}_{\vec{n}}^-] \hat{\rho}(t), \quad (46)$$

where $\hat{\sigma}_{\vec{n}}^\pm = (\hat{\sigma}_{\vec{n}}^x \pm i \hat{\sigma}_{\vec{n}}^y)/2$ are the operators raising and lowering the n -th spin in the z direction.

The Lindblad master equation (46) remains unchanged under a π rotation of all spins about the z -axis (transforming $\hat{\sigma}_{\vec{n}}^x \rightarrow -\hat{\sigma}_{\vec{n}}^x$ and $\hat{\sigma}_{\vec{n}}^y \rightarrow -\hat{\sigma}_{\vec{n}}^y$). This is a Z_2 symmetry, meaning that $[\mathcal{P}, \mathcal{L}] = 0$, with $\mathcal{P}\hat{\rho} = \hat{P}\hat{\rho}\hat{P}$ and $\hat{P} = \prod_{\vec{n}} \hat{\sigma}_{\vec{n}}^z$ the parity superoperator and operator, respectively. In the thermodynamic limit of an infinite lattice and dimension $D \geq 2$, the model spontaneously breaks this symmetry and a ferromagnetic order emerges [23, 30, 60–62]. In this study, we concentrate on a specific parameter range where mean-field theory predicts this transition to occur.

We consider a 3×3 system with periodic boundary conditions along x and y lattice directions so that the system is translationally invariant. The Lindblad master equation is then invariant with respect to the transformations $\hat{\sigma}_{(n_x, n_y)}^\alpha \rightarrow \hat{\sigma}_{(n_x+1, n_y)}^\alpha$ and $\hat{\sigma}_{(n_x, n_y)}^\alpha \rightarrow \hat{\sigma}_{(n_x, n_y+1)}^\alpha$. Defining $\mathcal{T}_x$ and $\mathcal{T}_y$ as superoperators associated with the translational symmetry along the two directions we have $[\mathcal{T}_x, \mathcal{L}] = [\mathcal{T}_y, \mathcal{L}] = [\mathcal{P}, \mathcal{L}] = 0$ and $[\mathcal{T}_x, \mathcal{T}_y] = [\mathcal{T}_x, \mathcal{P}] = [\mathcal{T}_y, \mathcal{P}] = 0$. Let $\hat{\eta}_j$ be a simultaneous eigenbasis of these operators, such that

$$\mathcal{T}_x \hat{\eta}_j = e^{i\kappa_x^{(j)}} \hat{\eta}_j, \quad \mathcal{T}_y \hat{\eta}_j = e^{i\kappa_y^{(j)}} \hat{\eta}_j, \quad \mathcal{P} \hat{\eta}_j = e^{i\pi z^{(j)}} \hat{\eta}_j, \quad (47)$$

with $\kappa_x^{(j)}$, $\kappa_y^{(j)}$, and $z^{(j)}$ representing quantum numbers conserved along the dynamics. For the 3×3 system considered below, $\kappa_{x,y}^{(j)} \in \{0, 2\pi/3, 4\pi/3\}$ while $z^{(j)} \in \{0, 1\}$. The Liouvillian written in this basis is block diagonal. Furthermore, one has that $\mathcal{T}_x \hat{\rho}_{ss} = \mathcal{T}_y \hat{\rho}_{ss} = \mathcal{P} \hat{\rho}_{ss} = \hat{\rho}_{ss}$. Therefore, the steady state belongs to the symmetry sector with eigenvalues $\kappa_x = \kappa_y = 0$ and $z = 0$.

We can thus apply the preconditioned Krylov method to the Liouvillian block containing the steady state. The full Liouvillian has dimension $2^{18} \approx 2.6 \times 10^5$, a size at which repeated LU factorizations across a phase diagram become prohibitively expensive. Restricting to the symmetry sector containing the steady state reduces this to $\approx 1.5 \times 10^4$, where LU remains tractable but costly, making it a natural target for the iterative approach. We plot the phase diagram of the magnetization m_z in Fig. 6(b). We correctly capture all the expected features, including the transition from the paramagnetic to the ferromagnetic regime. Similarly to VPT, we observe that the largest concentration of points where the preconditioner is re-computed is in the proximity of the phase transition.

7 Conclusions

In this work, we introduced variational perturbation theory (VPT), a generalization of standard perturbation theory for open quantum systems. We demonstrated its advantages and resilience to detrimental

non-analytic behavior across several examples. To make it numerically viable, we introduced two strategies, one based on exact LU decomposition of the Liouvillian superoperator, and one exploiting incomplete LU to investigate larger system sizes. Our methods allow rapid exploration of phase diagrams, as well as rapid estimation of steady state gradients, making them ideal tools for parameter fitting routines.

Our method is agnostic to the underlying details of the model under consideration, as we demonstrated. As such, it can be combined with other techniques such as cluster expansions [30, 63] and renormalization method [22, 23] to efficiently investigate more complex problems. We also plan to further extend VPT to time-domain simulations in systems with time-scale separation where a quasi-steady state regime is present. Beyond cat states already discussed in this work, typical examples of models with this feature include generic error-corrected quantum systems, metastable configurations, or systems where degrees of freedom can be adiabatically eliminated.

Author contributions

A.M. and F.M. conceived the project and developed the variational perturbation theory and its multipoint generalization. G.B. developed the iterative formulation of the method, establishing its equivalence with preconditioned Krylov-subspace solvers and its relation to subspace recycling.

All authors discussed the results and contributed to writing the manuscript.

Large language models were used to proofread the manuscript and to assist in writing parts of the code implementing the methods presented here. All authors have reviewed and verified the resulting text and code, and take full responsibility for the content of this work.

Acknowledgments

We acknowledge support and useful discussion with colleagues at Alice & Bob. We are grateful to the theoreticians for useful comments, and the experimentalists for providing measurements to test the efficiency of parameter fitting routines. In particular, we acknowledge Nicolas Didier, Nathanael Cottet, and their teams.

A Modified Liouvillian

Numerically solving the system of linear equations $\mathcal{L}|\hat{\rho}_{ss}\rangle = 0$ outputs the trivial solution $|\hat{\rho}_{ss}\rangle = |0, 0, \dots, 0\rangle$. To avoid this issue, one can construct a modified Liouvillian by adding a rank-1 matrix to $\mathcal{L}$, i.e. $\tilde{\mathcal{L}} = \mathcal{L} + |b\rangle\langle c|$, with b and c two vectors. This operator is full-rank provided that c has non-zero overlap

with the kernel of $\mathcal{L}$. A natural choice for c is then the identity eigenvector $\langle 1|$, since $\langle 1|\mathcal{L} = 0$. Indeed, any Liouvillian admits identity as a zero left eigenvector because the Lindblad master equation is trace-preserving. The vector $|b\rangle$ can be arbitrary; for convenience, in the main text we set it to $|b\rangle = |b, 0, \dots, 0\rangle$ with b an arbitrary complex number. The equation (4) follows from this choice, and

$$\mathcal{T}|\hat{\rho}_{ss}\rangle = |1\rangle\langle 1|\hat{\rho}_{ss}\rangle = |\text{Tr}(\hat{\rho}_{ss}), 0, \dots, 0\rangle. \quad (48)$$

B Derivation of PT

To derive PT in open quantum systems, we expand both the eigenvalues λ_μ and eigenoperators of $\mathcal{L} = \mathcal{L}_0 + \varepsilon\mathcal{L}_1$ as

$$\lambda_\mu = \sum_{j=0}^{\infty} \varepsilon^j \lambda_\mu^{(j)}, \quad \hat{\rho}_\mu = \sum_{j=0}^{\infty} \varepsilon^j \hat{\rho}_\mu^{(j)}, \quad (49)$$

where $\lambda_\mu^{(j)}$ and $\hat{\rho}_\mu^{(j)}$ are the j -th-order terms for eigenvalues and eigenvectors, respectively. The recursive relations for $\lambda_\mu^{(j)}$ and $\hat{\rho}_\mu^{(j)}$ can be obtained by expanding

$$\begin{aligned} \mathcal{L}(\varepsilon)\hat{\rho}_\mu(\varepsilon) &= (\mathcal{L}_0 + \varepsilon\mathcal{L}_1) \sum_{n=0}^{\infty} \varepsilon^n \hat{\rho}_\mu^{(n)} \\ &= \sum_{n=0}^{\infty} \varepsilon^n \left(\mathcal{L}_0 \hat{\rho}_\mu^{(n)} + \mathcal{L}_1 \hat{\rho}_\mu^{(n-1)} \right) \\ &= \sum_{n=0}^{\infty} \varepsilon^n \lambda_\mu^{(n)} \sum_{m=0}^{\infty} \varepsilon^m \hat{\rho}_\mu^{(m)}. \end{aligned} \quad (50)$$

Assuming an order-by-order resolution of this equation we get

$$\left(\mathcal{L}_0 - \lambda_\mu^{(0)} \right) \hat{\rho}_\mu^{(j)} = -\mathcal{L}_1 \hat{\rho}_\mu^{(j-1)} + \sum_{k=1}^j \lambda_\mu^{(k)} \hat{\rho}_\mu^{(j-k)}. \quad (51)$$

Since we are interested in the steady state $\hat{\rho}_{ss}$, $\lambda_\mu^{(0)} = 0$. Furthermore, the identity $\mathbb{1}$ is the left eigenoperator with zero eigenvalue for every Liouvillian and $\langle \mathbb{1} | \mathcal{L}_{0,1} = 0$. Thus,

$$\sum_{k=1}^j \lambda_\mu^{(k)} \langle \mathbb{1} | \hat{\rho}_\mu^{(j-k)} \rangle = 0, \quad (52)$$

for all orders j . We conclude that $\lambda_\mu^{(j)} = 0$, from which we obtain the recursion relation in Eq. (9).

C Tolerance criterion

Tolerance should be fixed with respect to the smallest nonzero Liouvillian eigenvalue. Indeed, introducing

the spectrum of the Liouvillian through the eigenoperators $\hat{\rho}_j$ and eigenvalues λ_j defined by $\mathcal{L}\hat{\rho}_j = \lambda_j\hat{\rho}_j$, we get

$$\begin{aligned}\hat{\rho}_{\text{ss}}^{\text{PT}}(\varepsilon) &= \sum_j c_j \hat{\rho}_j \\ \Rightarrow \|\mathcal{L}(\varepsilon) |\hat{\rho}_{\text{ss}}^{\text{PT}}(\varepsilon)\rangle\|^2 &= \left\| \sum_j c_j \lambda_j |\hat{\rho}_j\rangle \right\|^2 \\ &\sim \sum_j |c_j \lambda_j|^2,\end{aligned}\quad (53)$$

We conclude that a safe criterion to find the appropriate steady state is $\lambda_1 \gg \text{tol}$ where $\lambda_1 = \min_{j>0} |\lambda_j|$ is the Liouvillian gap, representing the slowest rate in the dynamics of an open quantum system, whose closure indicates the onset of critical phenomena.

D Building an optimal low-rank basis using singular value decomposition

To evaluate the optimality of (V)PT's representation, in the main text, we compare it to the best possible low-dimensional approximation of the steady states within the region of validity of the perturbative approximation. First, we compute the steady states for all parameter combinations. We then stack these state vectors as columns to form a matrix and perform a singular value decomposition (SVD). The left singular vectors corresponding to the M_{SVD} largest singular values form the optimal rank- M_{SVD} basis for representing the steady states. The optimal rank M_{SVD} is then chosen to match the accuracy of (V)PT.

E Efficient computation of the gradient

To determine $\partial_\delta \hat{\rho}_{\text{ss}}(\varepsilon, \delta)|_{\delta=0}$, we first notice that, differentiating Eq. (4), we get

$$\partial_\delta [\mathcal{L}(\phi, \theta) |\hat{\rho}_{\text{ss}}(\phi, \theta)\rangle]_{\delta=0} = \partial_\delta |b\rangle_{\delta=0} = 0. \quad (54)$$

Expanding the derivative, we conclude that the gradient can be computed by solving

$$(\mathcal{L}_0 + \varepsilon \mathcal{L}_1) \frac{\partial |\hat{\rho}_{\text{ss}}(\phi_n + \delta, \theta)\rangle}{\partial \delta} \Big|_{\delta=0} = -\mathcal{L}_2 |\hat{\rho}_{\text{ss}}(\phi_n, \theta)\rangle. \quad (55)$$

But since $\hat{\rho}_{\text{ss}}^{\text{PT}}(\delta, \varepsilon)$ and $\partial_\delta |\hat{\rho}_{\text{ss}}^{\text{PT}}(\delta, \varepsilon)\rangle$ are approximately spanned by $\{\hat{\rho}_{\text{ss}}^{(n,0)}, \hat{\rho}_{\text{ss}}^{(n,1)} | n \in [0, M]\}$, Eq. (55) can be solved with the same techniques as those used to compute $\vec{q}$. Namely, we build the orthonormal basis $\mathcal{Q}_\delta$ by orthonormalizing the basis of $\{\hat{\rho}_{\text{ss}}^{(n,0)}, \hat{\rho}_{\text{ss}}^{(n,1)}\}$.

Defining $\mathcal{L}_j^{\text{eff}} = \mathcal{Q}_\delta^\dagger \mathcal{L}_j \mathcal{Q}_\delta$, we get the reduced equation

$$\begin{aligned}(\mathcal{L}_0^{\text{eff}} + \varepsilon \mathcal{L}_1^{\text{eff}}) \frac{\partial \vec{q}(\delta, \varepsilon)}{\partial \delta} \Big|_{\delta=0} &= -\mathcal{L}_2^{\text{eff}} \vec{q}(0, \varepsilon) \\ \Rightarrow \frac{\partial |\hat{\rho}_{\text{ss}}^{\text{PT}}(\delta, \varepsilon)\rangle}{\partial \delta} \Big|_{\delta=0} &= \mathcal{Q}_\delta \frac{\partial \vec{q}(\delta, \varepsilon)}{\partial \delta} \Big|_{\delta=0}\end{aligned}\quad (56)$$

Let us also briefly comment on the complexity of computing the gradient for d controllable parameters and f to differentiate. Since we can compute each partial derivative independently along the f parameters, the additional cost of computing the gradient having the phase diagram is

$$\mathcal{O} \left[\frac{Pf}{R_{\text{VPT}}} (M^d N^2) + Pf(2M)^{3d} \right]. \quad (57)$$

References

- [1] S. Haroche and J. M. Raimond. “Exploring the quantum: Atoms, cavities, and photons”. [Oxford University Press](#). Oxford (2006).
- [2] H. M. Wiseman and G. J. Milburn. “Quantum measurement and control”. [Cambridge University Press](#). Cambridge (2010).
- [3] H. P. Breuer and F. Petruccione. “The theory of open quantum systems”. [Oxford University Press](#). Oxford (2007).
- [4] R. Fazio, J. Keeling, L. Mazza, and M. Schiró. “Many-body open quantum systems”. [SciPost Phys. Lect. Notes](#) **99** (2025).
- [5] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff. “Circuit quantum electrodynamics”. [Rev. Mod. Phys.](#) **93**, 025005 (2021).
- [6] I. Carusotto and C. Ciuti. “Quantum fluids of light”. [Rev. Mod. Phys.](#) **85**, 299–366 (2013).
- [7] H. Ritsch, P. Domokos, F. Brennecke, and T. Esslinger. “Cold atoms in cavity-generated dynamical optical potentials”. [Rev. Mod. Phys.](#) **85**, 553–601 (2013).
- [8] L. D. Carr, D. DeMille, R. V. Krems, and J. Ye. “Cold and ultracold molecules: science, technology and applications”. [New J. Phys.](#) **11**, 055049 (2009).
- [9] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. “Quantum dynamics of single trapped ions”. [Rev. Mod. Phys.](#) **75**, 281–324 (2003).
- [10] F. Ferrari, L. Gravina, D. Eeltink, P. Scarlino, V. Savona, and F. Minganti. “Dissipative quantum chaos unveiled by stochastic quantum trajectories”. [Phys. Rev. Res.](#) **7**, 013276 (2025).
- [11] D. Dahan, G. Arwas, and E. Grosfeld. “Classical and quantum chaos in chirally-driven, dissipative Bose-Hubbard systems”. [npj Quantum Inf.](#) **8** (2022).

- [12] E. M. Kessler, G. Giedke, A. Imamoglu, S. F. Yelin, M. D. Lukin, and J. I. Cirac. “Dissipative phase transition in a central spin system”. *Phys. Rev. A* **86**, 012116 (2012).
- [13] H. J. Carmichael. “Breakdown of photon blockade: A dissipative quantum phase transition in zero dimensions”. *Phys. Rev. X* **5**, 031028 (2015).
- [14] F. Minganti, A. Biella, N. Bartolo, and C. Ciuti. “Spectral theory of Liouvillians for dissipative phase transitions”. *Phys. Rev. A* **98**, 042118 (2018).
- [15] M. J. Hartmann, F. G. S. L. Brandão, and M. B. Plenio. “Strongly interacting polaritons in coupled arrays of cavities”. *Nat. Phys.* **2**, 849–855 (2006).
- [16] D. Roberts and A. A. Clerk. “Driven-dissipative quantum Kerr resonators: New exact solutions, photon blockade and quantum bistability”. *Phys. Rev. X* **10**, 021022 (2020).
- [17] F. Minganti, N. Bartolo, J. Lolli, W. Casteels, and C. Ciuti. “Exact results for Schrödinger cats in driven-dissipative systems and their feedback control”. *Sci. Rep.* **6**, 26987 (2016).
- [18] P. D. Drummond and D. F. Walls. “Quantum theory of optical bistability. i. nonlinear polarisability model”. *J. Phys. A: Math. Gen.* **13**, 725–741 (1980).
- [19] H. Weimer, A. Kshetrimayum, and R. Orús. “Simulation methods for open quantum many-body systems”. *Rev. Mod. Phys.* **93**, 015008 (2021).
- [20] M. B. Plenio and P. L. Knight. “The quantum-jump approach to dissipative dynamics in quantum optics”. *Rev. Mod. Phys.* **70**, 101 (1998).
- [21] A. J. Daley. “Quantum trajectories and open many-body quantum systems”. *Adv. Phys.* **63**, 77–149 (2014).
- [22] S. Finazzi, A. Le Boité, F. Storme, A. Bak-sic, and C. Ciuti. “Corner-space renormalization method for driven-dissipative two-dimensional correlated systems”. *Phys. Rev. Lett.* **115**, 080604 (2015).
- [23] R. Rota, F. Storme, N. Bartolo, R. Fazio, and C. Ciuti. “Critical behavior of dissipative two-dimensional spin lattices”. *Phys. Rev. B* **95**, 134431 (2017).
- [24] A. Kshetrimayum, H. Weimer, and R. Orús. “A simple tensor network algorithm for two-dimensional steady states”. *Nat. Commun.* **8**, 1291 (2017).
- [25] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero. “Positive tensor network approach for simulating open quantum many-body systems”. *Phys. Rev. Lett.* **116**, 237201 (2016).
- [26] F. Vicentini, A. Biella, N. Regnault, and C. Ciuti. “Variational neural-network ansatz for steady states in open quantum systems”. *Phys. Rev. Lett.* **122**, 250503 (2019).
- [27] N. Yoshioka and R. Hamazaki. “Constructing neural stationary states for open quantum many-body systems”. *Phys. Rev. B* **99**, 214306 (2019).
- [28] P. Deuar and P. D. Drummond. “Correlations in a BEC collision: first-principles quantum dynamics with 150 000 atoms”. *Phys. Rev. Lett.* **98**, 120402 (2007).
- [29] A. Polkovnikov. “Phase space representation of quantum dynamics”. *Ann. Phys.* **325**, 1790–1852 (2010).
- [30] J. Jin, A. Biella, O. Viyuela, L. Mazza, J. Keeling, R. Fazio, and D. Rossini. “Cluster mean-field approach to the steady-state phase diagram of dissipative spin systems”. *Phys. Rev. X* **6**, 031011 (2016).
- [31] G. Robichon and A. Tilloy. “Bootstrapping the stationary state of bosonic open quantum systems”. *Phys. Rev. A* **113**, 062440 (2026).
- [32] H. Weimer. “Variational principle for steady states of dissipative quantum many-body systems”. *Phys. Rev. Lett.* **114**, 040402 (2015).
- [33] A. C. Y. Li, F. Petruccione, and J. Koch. “Perturbative approach to Markovian open quantum systems”. *Sci. Rep.* **4**, 4887 (2014).
- [34] F. Benatti, A. Nagy, and H. Narnhofer. “Asymptotic entanglement and Lindblad dynamics: a perturbative approach”. *J. Phys. A: Math. Theor.* **44**, 155303 (2011).
- [35] A. C. Y. Li, F. Petruccione, and J. Koch. “Resummation for nonequilibrium perturbation theory and application to open quantum lattices”. *Phys. Rev. X* **6**, 021037 (2016).
- [36] H. Christiansen, V. V. Baran, and J. Paaske. “Reduced basis method for driven-dissipative quantum systems”. *Phys. Rev. Lett.* **135**, 236503 (2025).
- [37] O. Nagib and T. G. Walker. “Exact steady state of perturbed open quantum systems”. *Phys. Rev. Res.* **7**, 033076 (2025).
- [38] A. K. Noor and H. E. Lowder. “Approximate techniques of structural reanalysis”. *Comput. Struct.* **4**, 801–812 (1974).
- [39] J. N. Silverman and J. C. van Leuven. “Perturbational-variational approach to the calculation of variational wave functions. i. theory”. *Phys. Rev.* **162**, 1175 (1967).
- [40] L. Garrigue and B. Stamm. “On reduced basis methods for eigenvalue problems, and on its coupling with perturbation theory” (2025). [arXiv:2408.11924](https://arxiv.org/abs/2408.11924).
- [41] V. Gebhart, R. Santagati, A. A. Gentile, E. M. Gauger, D. Craig, N. Ares, L. Banchi, F. Marquardt, L. Pezze, and C. Bonato. “Learning quantum systems”. *Nat. Rev. Phys.* **5**, 141–156 (2023).

- [42] G. Beaulieu, F. Minganti, S. Frasca, V. Savona, S. Felicetti, R. Di Candia, and P. Scarlino. “Observation of first- and second-order dissipative phase transitions in a two-photon driven Kerr resonator”. *Nat. Commun.* **16**, 1951 (2025).
- [43] C. Berdou, A. Murani, U. Reglade, W. C. Smith, M. Villiers, J. Palomo, M. Rosticher, A. Denis, P. Morfin, M. Delbecq, et al. “One hundred second bit-flip time in a two-photon dissipative oscillator”. *PRX Quantum* **4**, 020350 (2023).
- [44] L. N. Trefethen and D. Bau. “Numerical linear algebra”. *SIAM*. (1997).
- [45] N. Bartolo, F. Minganti, W. Casteels, and C. Ciuti. “Exact steady state of a Kerr resonator with one- and two-photon driving and dissipation: Controllable Wigner-function multimodality and dissipative phase transitions”. *Phys. Rev. A* **94**, 033841 (2016).
- [46] W. Casteels, F. Storme, A. Le Boité, and C. Ciuti. “Power laws in the dynamic hysteresis of quantum nonlinear photonic resonators”. *Phys. Rev. A* **93**, 033824 (2016).
- [47] S. G. Krantz and H. R. Parks. “The implicit function theorem: history, theory, and applications”. *Springer*. (2002).
- [48] Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, M. J. Hatridge, M. Reagor, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret. “Confining the state of light to a quantum manifold by engineered two-photon loss”. *Science* **347**, 853–857 (2015).
- [49] R. Lescanne, M. Villiers, T. Peronnin, A. Sarlette, M. Delbecq, B. Huard, T. Kontos, M. Mirrahimi, and Z. Leghtas. “Exponential suppression of bit-flips in a qubit encoded in an oscillator”. *Nat. Phys.* **16**, 509–513 (2020).
- [50] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. *Nat. Methods* **17**, 261–272 (2020).
- [51] D. C. Liu and J. Nocedal. “On the limited memory BFGS method for large scale optimization”. *Math. Program.* **45**, 503–528 (1989).
- [52] Y. Saad and M. H. Schultz. “GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems”. *SIAM J. Sci. Stat. Comput.* **7**, 856–869 (1986).
- [53] H. A. Van der Vorst. “Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems”. *SIAM J. Sci. Stat. Comput.* **13**, 631–644 (1992).
- [54] M. R. Hestenes, E. Stiefel, et al. “Methods of conjugate gradients for solving linear systems”. *J. Res. Natl. Bur. Stand.* **49**, 409 (1952).
- [55] Y. Saad. “Iterative methods for sparse linear systems”. *SIAM*. (2003).
- [56] P. D. Nation. “Steady-state solution methods for open quantum optical systems” (2015). [arXiv:1504.06768](https://arxiv.org/abs/1504.06768).
- [57] M. L. Parks, E. de Sturler, G. Mackey, D. D. Johnson, and S. Maiti. “Recycling Krylov subspaces for sequences of linear systems”. *SIAM J. Sci. Comput.* **28**, 1651–1674 (2006).
- [58] K. M. Soodhalter, D. B. Szyld, and F. Xue. “Krylov subspace recycling for sequences of shifted linear systems”. *Appl. Numer. Math.* **81**, 105–118 (2014).
- [59] K. M. Soodhalter, E. de Sturler, and M. E. Kilmer. “A survey of subspace recycling iterative methods”. *GAMM-Mitt.* **43**, e202000016 (2020).
- [60] T. E. Lee, S. Gopalakrishnan, and M. D. Lukin. “Unconventional magnetism via optical pumping of interacting spin systems”. *Phys. Rev. Lett.* **110**, 257204 (2013).
- [61] R. Rota, F. Minganti, A. Biella, and C. Ciuti. “Dynamical properties of dissipative XYZ Heisenberg lattices”. *New J. Phys.* **20**, 045003 (2018).
- [62] D. Huybrechts, F. Minganti, F. Nori, M. Wouters, and N. Shammah. “Validity of mean-field theory in a dissipative critical system: Liouvillian gap, $\mathbb{PT}$ -symmetric antigap, and permutational symmetry in the XYZ model”. *Phys. Rev. B* **101**, 214302 (2020).
- [63] A. Biella, J. Jin, O. Viyuela, C. Ciuti, R. Fazio, and D. Rossini. “Linked cluster expansions for open quantum systems on a lattice”. *Phys. Rev. B* **97**, 035103 (2018).