

Low-rank optimal control of quantum devices

Leo Goutte^{*} and Vincenzo Savona[†]

*Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
and Center for Quantum Science and Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*



(Received 23 October 2025; accepted 19 December 2025; published 26 January 2026)

We demonstrate that the control protocols of quantum information devices can be simulated by assuming a low-rank ansatz for the density matrix. The rationale underlying this assumption is that quantum information protocols, by design, operate in a regime of nearly pure quantum states. Within the low-rank assumption, the simulation of these protocols is considerably faster than solving the full Lindblad master equation. This advantage can be used to increase the accuracy of the simulation by avoiding uncontrolled approximations and to streamline protocol optimization. We benchmark our approach on the optimization of the transmon qubit dispersive readout in a realistic transmon-resonator-filter model. With a Hilbert space dimension $N = 2000$, assuming a rank as low as $M = 20$, we achieve a nearly 100-fold speedup compared to full master equation integration while accurately reproducing all relevant observables. By combining the low-rank approximation with a compact pulse parametrization and gradient-free optimization, we obtain state-of-the-art readout assignment errors $\varepsilon_a \approx 1.2 \times 10^{-3}$ for a 40 ns readout pulse schedule, while comfortably running on a laptop and not relying on the rotating-wave approximation. Our approach is broadly applicable to most quantum control protocols, including quantum gates, state preparation, and fast reset operations. This establishes low-rank methods as a general tool for optimal control across diverse quantum platforms.

DOI: [10.1103/v383-2vss](https://doi.org/10.1103/v383-2vss)

I. INTRODUCTION

Quantum computing promises transformative advances in quantum chemistry, materials science, and optimization, offering dramatic speedups for specific problem classes. Realizing this potential demands solving major hardware and software challenges, particularly in the precise control, design, and optimization of quantum devices.

Quantum optimal control (QOC) [1–8] tackles these challenges by designing control schedules for executing elementary operations such as quantum gates [9–27], quantum bit readout [28–38], and state reset [29–31,39,40]. At larger scales, the simultaneous control of many interacting quantum bits requires the orchestration of numerous control pulses and frequencies [37,41–46]. These optimization tasks typically involve numerous evaluations of a loss function that quantifies the performance of the control protocol and depends on a large number of parameters, including pulse shapes, durations, and frequencies. The loss function is then typically evaluated through the accurate simulation of the quantum device dynamics.

State-of-the-art hardware quantum platforms based on superconducting circuits, trapped ions, or Rydberg atoms are

all *open quantum systems*, where the dissipation and decoherence significantly affect the system dynamics. Within the assumption of a Markovian environment, these dynamics are modeled by the Lindblad master equation [47], whose computational complexity scales as $\mathcal{O}(N^2)$. Accurate and predictive simulations often require large values for the Hilbert space dimension N , making the optimization process prohibitively expensive.

To address these challenges, significant effort has been devoted to developing efficient optimization methods, such as the gradient ascent pulse engineering (GRAPE) algorithm [48], the Krotov method [49,50], the chopped random-basis method in Fourier space [51,52], and automatic differentiation [29,31,53]. These methods have been thoroughly extended to QOC of open quantum systems [12,50,54,55]. However, less effort has been directed toward reducing the computational time required to solve the Lindblad master equation itself through controlled approximations that maintain predictive accuracy of the optimization metrics. A typical speedup strategy consists in applying the rotating-wave approximation (RWA), which neglects counterrotating terms in the Hamiltonian and Lindblad operators. This approximation is often justified by the smallness of the driving field amplitude, but it can lead to significant inaccuracies in the prediction of relevant observables [56–59]. Other approaches have simulated the dissipative dynamics using quantum trajectories in place of the full numerical solution of the Lindblad master equation [15,31]. The trade-off is that averaging over sufficiently many trajectories may be computationally expensive and still imply a rough loss-function landscape due to stochastic fluctuations, which would make gradient-based optimization methods less effective. Additionally, especially

^{*}Contact author: leo.goutte@epfl.ch

[†]Contact author: vincenzo.savona@epfl.ch

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.

in low-dissipation settings, it may be statistically difficult to sample rare events, potentially leading to uncontrolled bias in the estimates of relevant quantities.

In this work, we apply a computational framework for QOC that enables efficient, accurate, and scalable optimization for quantum devices. The central feature is a dramatic speedup in the numerical calculation of the loss function with negligible loss of accuracy and predictive capability, which we achieve by adopting a low-rank approximation (LRA) [60–67]. The LRA is successful because, in most quantum operation protocols, the system remains in a nearly pure quantum state throughout the operation schedule. In such regimes, the density matrix of the system contains much less information than implied by its full $N \times N$ dimensionality. By representing the density matrix with rank- M matrices (where $M \ll N$ in high-purity regimes), the number of coupled differential equations to solve is reduced from N^2 to $N \times M$, resulting in dramatic savings in memory and computational time. We benchmark our approach by *deliberately revisiting* the optimization of the dispersive readout of a transmon qubit coupled to a readout resonator and a Purcell filter [33,35,36,68–73] that was very recently studied by Gautier *et al.* [29]. By further combining the LRA with a gradient-free heuristic for the optimization, our approach results in an optimal readout assignment error of 1.2×10^{-3} , on par with the state of the art [29,35], while leading to a nearly 100 times reduction in the computational time required for realistic device simulations.

QOC under the LRA can be applied to any scenario where the quantum state remains weakly mixed along the simulated dynamics, including quantum gates and state preparation. The increased efficiency of this approach can then be leveraged to streamline the optimization process, lift uncontrolled approximations such as the RWA, and extend the scope of QOC to larger and more complex systems and control protocols.

II. LOW-RANK FORMALISM

The density matrix in a Hilbert space of dimension N may always be written as a statistical ensemble of pure states

$$\hat{\rho} = \sum_{j=1}^N q_j |\psi_j\rangle \langle \psi_j|, \quad (1)$$

where q_j are the probabilities (in descending order), with $\sum_{j=1}^N q_j = 1$, and $|\psi_j\rangle$ are the (normalized) states entering the statistical mixture. Note that the decomposition equation (1) is in general not unique.

If the system under consideration is weakly mixed, the density matrix is well approximated by a truncated low-rank ansatz [61–64]:

$$\hat{\rho} \simeq \sum_{j=1}^M p_j |\psi_j\rangle \langle \psi_j|, \quad (2)$$

where $p_j = q_j / \sum_{k=1}^M q_k$ and $M \leq N$ is the rank of $\hat{\rho}$. The validity of this approximation is measured by the fidelity between states (1) and (2), which should be close to unity along the system dynamics. Assuming that the states $|\psi_j\rangle$ form an orthonormal basis, this translates to the minimization

of the truncation error $\epsilon_M = 1 - \sum_{j=1}^M q_j$ (see Sec. II B for more details).

It is worth noting that the truncated expansion (2) is analogous to the structure of locally purified tensor networks [74–78], where the index j acts as the purification bond index arising from an operator-space Schmidt decomposition. Retaining only M terms corresponds to truncating this purification bond dimension while keeping the full entanglement structure of each individual pure state.

A. Dynamics

The time evolution within the low-rank manifold is governed by the so-called nonstochastic matrix Schrödinger equation introduced in Refs. [65,66], which describes both the unitary and dissipative evolution of the states $|\psi_j\rangle$ in a fully deterministic manner. Here, we briefly review only the main results.

We begin by rewriting the density matrix in the diagonal basis as

$$\hat{\rho}(t) = \mathbf{m}(t) \mathbf{m}^\dagger(t), \quad (3)$$

where $\mathbf{m}(t) = (\sqrt{p_1(t)}|\psi_1(t)\rangle \cdots \sqrt{p_M(t)}|\psi_M(t)\rangle)$ is a $N \times M$ matrix, $|\psi_j(t)\rangle$ are the (normalized) states entering the low-rank mixture, and $p_j(t)$ are the corresponding probabilities, with $\sum_{j=1}^M p_j = 1$. Here and in the remainder of this work, we omit writing the explicit time dependence unless ambiguous. The equation governing the dynamics of $\mathbf{m}$ can be derived starting from the master equation ($\hbar = 1$):

$$\dot{\hat{\rho}} = \mathcal{L}[\hat{\rho}] = -i[\hat{H}, \hat{\rho}] + \sum_k \mathcal{D}[\hat{L}_k] \hat{\rho}, \quad (4a)$$

$$\mathcal{D}[\hat{L}_k] \hat{\rho} = \hat{L}_k \hat{\rho} \hat{L}_k^\dagger - \frac{1}{2} \{\hat{L}_k^\dagger \hat{L}_k, \hat{\rho}\}, \quad (4b)$$

where $\hat{L}_k$ are the Lindblad operators describing the dissipation processes. Substituting Eq. (3) into Eq. (4) leads to an equation of motion for the matrix $\mathbf{m}$:

$$\dot{\mathbf{m}} = -i\hat{H}\mathbf{m} + \mathcal{O}[\mathbf{m}]. \quad (5)$$

The first term on the right-hand side of Eq. (5) governs the unitary evolution of $\mathbf{m}$. The second term captures the dissipative nonunitary dynamics and is defined through the relation

$$\sum_k \mathcal{D}[\hat{L}_k] \mathbf{m} \mathbf{m}^\dagger = \mathcal{O}[\mathbf{m}] \mathbf{m}^\dagger + \mathbf{m} \mathcal{O}[\mathbf{m}]^\dagger. \quad (6)$$

Using Eq. (4b), we obtain the specific form of $\mathcal{O}[\mathbf{m}]$,

$$\mathcal{O}[\mathbf{m}] = \frac{1}{2} \sum_k \hat{L}_k \mathbf{m} (\mathbf{m}^{-1} \hat{L}_k \mathbf{m})^\dagger - \hat{L}_k^\dagger \hat{L}_k \mathbf{m}. \quad (7)$$

Note that $\mathbf{m}$ is an $N \times M$ rectangular matrix, whose inverse must be computed as the Moore-Penrose inverse defined as $\mathbf{m}^{-1} = (\mathbf{m}^\dagger \mathbf{m})^{-1} \mathbf{m}^\dagger$. The term low-rank approximation refers to both Eqs. (2) and (5).

Here, we assume the rank M to be constant in time. The evolution according to Eq. (5) therefore requires the multiplication of extremely sparse $N \times N$ matrices (such as $\hat{H}$ and $\hat{L}_k$) and dense $N \times M$ matrices (such as $\mathbf{m}$) at every time step. Conversely, a full evolution according to the master

equation includes, once vectorized [79], dense vectors of size N^2 and sparse matrices of size $N^2 \times N^2$. If $M \ll N$, both the compute time and memory requirements are drastically improved.

When assuming a fixed rank, care must be taken in setting the initial condition for the numerical integration of Eq. (5). If the initial state has a rank $M' < M$, as is the case here when choosing an initial pure state, then the matrix $\mathbf{m}^\dagger \mathbf{m}$ is singular and the Moore-Penrose inverse of $\mathbf{m}$ is not defined. We solve this issue by setting the states $|\psi_j(0)\rangle$, for $j = 2, \dots, M$, to a set of random orthonormalized states and the corresponding probabilities to small values $p_j = \epsilon \ll 1$. The validity of this approach can be easily verified by checking the convergence of the results for decreasing values of ϵ . For the study that follows, we find that $\epsilon = 10^{-5}$ always guarantees convergence.

Note that this addition of small random states is mathematically equivalent to a regularization of the inverse $(\mathbf{m}^\dagger \mathbf{m})^{-1} \rightarrow (\mathbf{m}^\dagger \mathbf{m} + \epsilon \mathbf{1})^{-1}$ [67]. Importantly, this initial regularization bears no influence on the subsequent dynamics, since the equation of motion (5) continuously rotates the columns of $\mathbf{m}$ during the evolution toward the relevant low-rank states. As a result, the low-rank subspace rapidly aligns with the physically relevant part of the state, and the details of the random vectors become irrelevant provided that ϵ remains below the intrinsic truncation error.

B. Validity

Ideally, the validity of the LRA is measured by the truncation error ϵ_M . Computing the low-rank density matrix (3) by solving Eq. (5), however, does not provide access to the probabilities q_j and therefore to ϵ_M . Several control quantities can be used as proxies for the truncation error [67]. Here, we adopt the ratio p_M/p_1 as a control quantity, which is readily computed as a function of time and closely reflects the accuracy of the LRA [67]. We further track the purity $\text{tr}(\hat{\rho}^2)$, which is overestimated by the LRA and is expected to converge to the exact value as the rank M is increased.

To illustrate the validity of the LRA and the reliability of the control quantities, we simulate a typical schedule of the dispersive readout of a transmon qubit, as described in detail in Sec. III. We compare the values of p_M/p_1 and of the purity $\text{tr}(\hat{\rho}^2)$ as obtained from the full integration of the Lindblad master equation and from the LRA. We additionally integrate the Lindblad master equation without applying the LRA but instead applying the RWA, which consists in neglecting all counterrotating terms, in both the coupling and drive terms, in the frame rotating with the driving field [73,80]. The RWA is often adopted in QOC as it leads to a large speedup in computational time due to the elimination of rapidly oscillating terms in the equation of motion [29]. It is, however, an uncontrolled approximation known to produce quantitatively inaccurate results in cases when modes are nonlinear, when driving fields are not weak, or are off-resonant [56–59].

Figure 1 shows the results obtained by integrating the equations of motion for these three levels of approximation: the full Lindblad master equation, the LRA, and the RWA. The dynamics are simulated in the presence of a square driving pulse of duration 40 ns, in the same setting illustrated in Fig. 2. For this simulation, a Hilbert space of total dimension




FIG. 1. Time evolution of (a) the purity $\text{tr}(\hat{\rho}^2)$ and (b) the control quantity p_M/p_1 , as computed from integrating the full master equation [Eq. (9)], the LRA [Eq. (5)] with $M = 20$, and the RWA [Eq. (A26)] for the transmon readout. The solid, dashed, and dotted lines correspond to the full evolution, low-rank, and rotating-wave approximation evolutions, respectively. The red (blue) lines correspond to an initial transmon ground (first excited) state. The relative error between the full and approximate values δ is much smaller for the LRA than for the RWA. (b) The control quantity p_M/p_1 increases as the state becomes mixed, providing a practical indicator of the LRA's accuracy.

$N = 2000$ was considered. For the LRA, $M = 20$ was taken. This choice of parameters coincides with that leading to the results presented in Fig. 3. The purity generally decreases in time as the system evolves toward a mixed state. The comparison presented in Fig. 1 shows how the LRA essentially coincides with the full simulation already for such a small value of M . The RWA, on the other hand, leads to a sizable quantitative inaccuracy. The LRA can thus essentially reproduce the exact results, and hence be predictive for QOC applications, provided the convergence of the results versus M is carefully checked, or alternatively an adaptive low-rank scheme is used [67].




FIG. 2. (a) Schematic of the transmon readout model. The transmon (yellow) is capacitively coupled to a resonator (red), itself coupled to a Purcell filter (blue). The filter is coupled to the environment through a readout line that enables driving and introduces dissipation at a rate κ . (b) Illustrative example of driving pulses $\Omega_d(t)$. Both the square (dark blue) and stepwise (light blue) pulses are smoothed out using a logistic function, as detailed in Appendix A3. The thin lines show the stepwise pulse before smoothing.

III. TRANSMON READOUT

The transmon dispersive readout protocol is an ideal proving ground for the effectiveness of the LRA in QOC. Since a large Hilbert space dimension N is needed to account for the putatively high population of the readout mode, an accurate simulation quickly becomes prohibitively expensive.

A. Model

The setup for the transmon readout is shown in Fig. 2(a), where the transmon qubit is coupled to a readout resonator, which is itself coupled to a Purcell filter. The system is described by the Hamiltonian

$$\begin{aligned} \hat{H} = & 4E_c\hat{n}_t^2 - E_j \cos \hat{\phi}_t + \omega_r \hat{a}^\dagger \hat{a} + \omega_f \hat{f}^\dagger \hat{f} \\ & - J(\hat{a}^\dagger - \hat{a})(\hat{f}^\dagger - \hat{f}) + i g \hat{n}_t (\hat{a}^\dagger - \hat{a}) \\ & + i \Omega_d(t) \sin(\omega_d t) (\hat{f}^\dagger - \hat{f}), \end{aligned} \quad (8)$$

and is governed by a Lindblad master equation

$$\dot{\hat{\rho}} = -i[\hat{H}, \hat{\rho}] + \kappa \mathcal{D}[\hat{f}]\hat{\rho} + \gamma \mathcal{D}[\hat{b}]\hat{\rho}. \quad (9)$$

Here, $\hat{n}_t$ and $\hat{\phi}_t$ are the transmon conjugate charge and phase operators, and E_c and E_j the corresponding charging and Josephson energies. The operators $\hat{b}$, $\hat{a}$, and $\hat{f}$ are annihilation operators of the transmon, resonator, and filter modes, respectively, while ω_r and ω_f are the bare resonator and filter resonant frequencies. The capacitive coupling strengths are g and J , and the dissipation rates are κ for the filter and γ for the transmon. The system is driven through the filter mode by a microwave pulse at frequency ω_d with envelope $\Omega_d(t)$. We initially model the latter as a square pulse of maximum




FIG. 3. (a) The phase-space trajectories $\beta_{g/e} = \text{tr}(\hat{\rho}_{g/e} \hat{f})$ for a readout time $\tau = 40$ ns and a square pulse profile with amplitude $|\Omega_d|/2\pi = 150$ MHz. The full, low-rank with $M = 20$, and RWA evolutions correspond to the full, dashed, and dotted lines, respectively. The red (blue) lines correspond to an initial ground (first excited) transmon state. The thin lines connecting red and blue trajectories highlight equal times on the two curves. (b) The photon occupation $n_{g/e} = \text{tr}(\hat{\rho}_{g/e} \hat{f}^\dagger \hat{f}) = |\beta_{g/e}|^2$ in the filter mode and the relative errors between the full and low-rank (dashed), and full and RWA (dotted) results. The relative error between the full and approximate values, defined for a quantity y as $\delta = |y_{\text{full}} - y_{\text{LRA}}|/y_{\text{full}}$ is much smaller for the LRA than the RWA. The Hilbert space truncations are $N_t = 100$, $N_u = 4$, and $N_f = 5$, resulting in $N = 2000$ states.

amplitude $|\Omega_d|$ with a smooth ramp up and ramp down modulated by a logistic function, as plotted in Fig. 2(b). Importantly, the chosen ramp time of 2.5 ns is slow compared to $\omega_d^{-1} \approx 0.02$ ns so as to avoid spurious photon populations in the driven mode.

Due to their coupling, the resonator's frequency undergoes a small dispersive shift that depends on the state of the transmon. This in turn affects the resonator dynamics, allowing one to infer the initial state of the transmon from a heterodyne measurement of the resonator field. The addition of a Purcell filter coupled to the resonator does not fundamentally change this picture, but simply suppresses the spontaneous transmon decay during the readout schedule [70–72]. The coupled resonator-filter subsystem admits two normal modes and either one of these modes can be used for the readout by resonantly driving the filter at the corresponding hybridized normal-mode frequency [35]. Here, we choose the lower normal mode as the readout mode. A detailed description of the readout protocol's numerical implementation and chosen parameter values is provided in Appendix A.

TABLE I. Run time in seconds for the full master equation, RWA, and low-rank solvers with $M = 20$ for the transmon readout with a square pulse of $\tau = 40$ ns duration. The rows correspond to initializing the transmon in the $|g\rangle$ or $|e\rangle$ state, respectively. For the full master equation, we use the `mesolve` solver from QUANTUM-TOOLBOX.JL [82]. The total dimension of the Hilbert space is $N = 2000$ ($N_l = 100$, $N_u = 4$, $N_f = 5$). The times were obtained from CPU calculations on a laptop with an Apple M3 Pro CPU and 36 GB of RAM. We encourage the interested reader to reproduce these and all other results herein with the help of the QUANTUMTOOLBOX.JL package [81,82].

Model	Full	RWA	LRA	LRA + RWA
$ g\rangle$	2441	623	30	21
$ e\rangle$	2837	741	33	20

B. Computational performance

In Fig. 3, we compare the phase-space trajectories, defined on the complex plane by the expectation value of the field in the filter resonator $\beta_{g/e} = \text{tr}(\hat{\rho}_{g/e} \hat{f})$, and the corresponding photon occupation $n_{g/e} = \text{tr}(\hat{\rho}_{g/e} \hat{f}^\dagger \hat{f}) = |\beta_{g/e}|^2$, as obtained with the full master equation, with the LRA, and with the RWA, for a driving field modeled as the smoothed square pulse plotted in Fig. 2(b). This reproduces the square pulse setting studied in Ref. [29].

Already with a modest rank of $M = 20$, the LRA result accurately describes the phase-space trajectories for both choices of initial transmon state. The RWA, on the other hand, results in sizable quantitative differences. For this particular choice of pulse shape and amplitude, this discrepancy still leads to a fairly accurate prediction of the readout assignment error. A full optimization schedule, however, may explore regions in parameter space where the discrepancy brought by the RWA may be larger and possibly lead to suboptimal results. If this issue can be safely excluded by restricting the parameter space in the optimization, then it is possible to simulate the system dynamics by applying both the LRA and RWA, leading to even faster simulation times. Table I shows the wall-clock compute times for the full simulation, the LRA, the RWA, and the two combined. The LRA and combined LRA + RWA results were obtained using a fixed-rank nonstochastic matrix Schrödinger equation solver publicly available through Ref. [81]. The LRA with $M = 20$ achieves an 81 times speedup with respect to the full simulation. This should be compared to the 4 times speedup brought by the RWA. Combining the LRA with $M = 20$ and the RWA, a speedup exceeding 100 times is obtained.

IV. FIDELITY OPTIMIZATION

Having established a method that accurately simulates a single instance of the transmon readout protocol at a minimal computational cost, we proceed by employing it to optimize the readout fidelity. This is a multiobjective optimization problem, where the goal is to find the optimal pulse shape for a given readout time τ , which minimizes the assignment error $\epsilon_a(\tau)$ while also optimizing several other objectives [29].

A. Assignment error

The trajectories of the readout mode in phase space are directly related to the fidelity of the readout measurement. The assignment error of the measurement is determined by the separation between the pointer states of the readout mode [32,37], and is expressed as

$$\epsilon_{\text{sep}}(\tau) = \frac{1}{2} \text{erfc} \left[\frac{\text{SNR}(\tau)}{2} \right]. \quad (10)$$

Here, the signal-to-noise ratio is defined as

$$\text{SNR}(\tau) = \sqrt{2\eta\kappa \int_0^\tau dt |\beta_e(t) - \beta_g(t)|^2}, \quad (11)$$

where η is the measurement efficiency, which we set to be $\eta = 0.6$ to match with experiment [70], and $\beta_{e,g}(t)$ are the expectation values of the field in the readout mode, which are measured in the experiment. Together with the decay of the transmon from its excited to ground state captured by $\epsilon_{\text{decay}}(\tau) = \tau\gamma/2$, we obtain the total assignment error

$$\epsilon_a(\tau) = \epsilon_{\text{sep}}(\tau) + \epsilon_{\text{decay}}(\tau), \quad (12)$$

which is related to the readout fidelity by $\mathcal{F} = 1 - \epsilon_a$ [37,83]. The competition between ϵ_{decay} , which favors short times, and ϵ_{sep} , which favors strong drives and long times, makes for an optimal readout time at which the assignment error is minimized, as seen in Fig. 4(a). For the square pulse, we find a minimum assignment error of $\epsilon_a = 1.6 \times 10^{-3}$ for a readout time of 60 ns. A similar result is obtained with both the LRA and RWA.

B. Optimization parameters

This fidelity is mainly determined by the pulse shape and amplitude of the drive. In Ref. [29], the pulse envelope was represented by 1 ns time bins, resulting in several hundreds of nominally independent parameters to optimize. We argue here that this is an unnecessarily complex parametrization scheme, as the timescale of the system response is considerably longer. In practice, pulse optimization often converges to physically meaningful waveforms that admit simple analytical descriptions with only a few free parameters [38].

In what follows, we therefore parametrize the pulse as a stepwise complex-valued function with a smooth interpolation between its levels, as illustrated in Fig. 2(b). We also assume that the pulse envelope $\Omega_d(t)$ is complex-valued, i.e., that the pulse can feature a slowly varying time-dependent phase. This amounts to a total of $2n_{\text{steps}}$ parameters, where n_{steps} is the number of steps in the pulse. Here, we take $n_{\text{steps}} = 5$, i.e., ten parameters to optimize. The validity of this approach is ultimately determined by the quality of the optimized readout protocol. We do not include the driving frequency ω_d as an optimization parameter, and set it to the frequency of the transmon-lower normal-mode system [35].

Restricting the parameter space to a minimal number of physically relevant parameters brings several advantages. First, it allows for an easier and faster convergence of the optimization algorithm, in particular, by enabling the use of computationally less demanding gradientless optimization schemes, and by increasing the likelihood of finding a global,




FIG. 4. Transmon readout error and ionization for a square pulse compared to the LRA with rank $M = 20$ and RWA. The full, low-rank, and RWA evolutions are plotted in gray, dashed orange, and dotted purple lines, respectively. (a) The assignment error $\epsilon_a(\tau)$ as a function of the readout time. The dashed black line corresponds to the long- τ limit $\epsilon_a \simeq \tau\gamma/2$. The relative error δ , shown in the lower panel, demonstrates that the LRA captures the exact assignment error slightly better than the RWA, though the RWA still performs well due to the relatively weak drive strength. (b) The transmon ionization [Eq. (B1)]. Again, the LRA is a better match to the full dynamics than the RWA, though the latter still manages to capture the transmon ionizations since the drive is relatively weak.

rather than a local minimum. Second, it opens the way for the optimization of more complex protocols where more intricate pulse schedules are required. Here, we found that the simultaneous perturbation stochastic approximation [84,85] leads to excellent optimization results, but we verified that similarly good results were obtained using a truly global optimization scheme such as the particle swarm algorithm [86].

C. Results

We determine the optimal readout schedule by searching for the drive pulse sequence that minimizes the loss function $\Lambda(\tau)$ for a given readout time τ . The loss function is designed so as to balance two main competing effects: minimizing the assignment error while avoiding transmon ionization (see Appendix B 1 for more details). During optimization, the loss was evaluated using the LRA with $M = 20$. The RWA was not used during optimization. Figure 5(a) shows the optimal value of the loss as a function of the assumed pulse duration τ . The optimal value obtained for $\tau = 40$ ns corresponds to

an assignment error $\epsilon_a = 1.2 \times 10^{-3}$, comparable with the current state of the art [29,35]. The inset displays a representative optimization schedule for $\tau = 40$ ns. We observe that the SPSA converges in $\mathcal{O}(100)$ epochs, corresponding to 200 evaluations of the loss, or 400 integrations of the master equation and a total wall-clock time $T_{\text{LRA}} \simeq 1.2 \times 10^4$ s, i.e., less than 3 h on the CPU of a modern laptop.

The optimal pulse envelope obtained for $\tau = 40$ ns is shown in Fig. 5(b). Notably, the real part of the optimized pulse resembles a two-step pulse: a strong initial drive that rapidly populates the readout cavity, followed by a reduced amplitude to limit undesired ionization events [70,87]. This result is obtained using only a single drive on the readout mode, without a secondary drive on the transmon [69,88,89]. Including these features as additional optimization parameters is expected to further reduce the optimal error.

The real and imaginary parts of the field in the filter resonator for the optimal pulse are shown in Fig. 5(c). These were computed with the higher truncations $N_t = 6$, $N_u = 6$, and $N_l = 100$ in order to validate the optimized simulations. The full and LRA results essentially coincide except for a slight discrepancy at the end of the readout schedule for the trajectory initialized with the qubit in the excited state. The convergence of the LRA can be assessed by repeating the calculation for higher M . The inset of Fig. 5(c) shows the same calculation with $M = 80$, which more closely approximates the exact result. Figure 6 reports the dependence of the assignment error on M . The relative error between the LRA and exact values decreases from 5% to 1.5% when increasing M from 20 to 80, confirming the accuracy of the controlled LRA. In light of the analogy between the LRA and locally purified tensor networks, this convergence is equivalent to a convergence with increasing bond dimension. In contrast, the phase-space trajectories computed with the RWA [Fig. 5(c)] differ substantially from the exact ones, despite yielding a comparable assignment error.

Restricting ourselves to controlled approximations—namely, the Hilbert space truncations N_t , N_u , and N_l , and the rank M —allows the optimization to be run with a computationally faster yet more aggressive approximation. The rationale is that even if the loss is not quantitatively exact, it remains correlated with the value obtained under a more conservative approximation, so the resulting optimal parameters are expected to be reliable. To support this, in Appendix B 2 we repeat the full optimization with $M = 80$ and find that they essentially coincide. In contrast, adopting a similar strategy with an uncontrolled approximation such as the RWA is more likely to yield inaccurate loss estimates and less reliable optima.

V. CONCLUSION AND OUTLOOK

In this work, we have demonstrated that the low-rank approximation provides an accurate and fast approach to the optimal control of quantum device protocols. The essential reason is that, by construction, most of these protocols keep the system in high-purity states throughout their schedules. In this regime, a low-rank representation of the density matrix suffices to capture the dynamics with high accuracy. This




FIG. 5. (a) The optimal value of the loss function for various readout times τ between 40 and 100 ns. Inset: a typical variation of the loss function along the optimization schedule for $\tau = 40$ ns. (b) Real (solid) and imaginary (dotted) parts of the optimal pulse envelope (light blue) for $\tau = 40$ ns. The envelope of the square pulse (dark blue) is plotted for comparison. (c) Phase-space trajectories for the optimal readout schedule and $\tau = 40$ ns, for initial states $|g\rangle$ (red) and $|e\rangle$ (blue), computed with $N_r = 5$, $N_u = 4$, and $N_l = 100$. The full and RWA trajectories are taken from a simulation in a Hilbert space of increased dimension $N_r = 6$, $N_u = 6$, and $N_l = 100$. Inset: Difference between $M = 20$ (dashed purple) and $M = 80$ (dashed blue) low-rank evolutions compared to the full result (blue) in the region of phase space where the discrepancy is most pronounced.

enables a drastic reduction in computational time and memory, while preserving predictive power.

We have benchmarked this approach on the optimization of the transmon qubit readout protocol in a realistic transmon-resonator-filter model. Thanks to the LRA's speed and a compact yet physically essential pulse parametrization, we recover a *state-of-the-art assignment error* $\epsilon_a \approx 1.2 \times 10^{-3}$ at 40 ns. This optimization could be carried out seamlessly on a standard laptop in 2 orders of magnitude less optimization time, without relying on uncontrolled approximations such as the RWA.

Beyond this specific application, the computational framework we introduced is broadly applicable. Any protocol whose dynamics remains weakly mixed—including high-fidelity gates and fast reset—can leverage its low rank to

widen the reachable design space, lift *ad hoc* approximations, and scale to larger Hilbert spaces and richer control schedules, offering a pathway to accelerate QOC across diverse quantum platforms. To facilitate its broader adoption, we provide an open-source implementation [81], while a more general adaptive low-rank Lindblad solver is also available as part of QUANTUMTOOLBOX.JL [82].

Looking forward, while we have successfully employed the SPSA optimization method in this work, the LRA framework could naturally be combined with QOC-specific algorithms such as GRAPE or the Krotov method, potentially enhancing performance even further. By making accurate, large-scale device optimization feasible with greatly reduced computational resources, this approach can dramatically broaden access to realistic QOC and accelerate progress toward scalable quantum technologies.

ACKNOWLEDGMENTS

We acknowledge enlightening discussions with Lorenzo Fioroni, Filippo Ferrari, and Alberto Mercurio. We are grateful to Ronan Gautier for making the raw data of Ref. [29] available to us. This work was supported by the Swiss National Science Foundation through Projects No. 200021-227992, No. 20QU-1_215928, and No. 200020_215172.

DATA AVAILABILITY

The data that support the findings of this article are openly available [81].

APPENDIX A: DETAILS OF THE SIMULATION

In what follows, we provide some additional details to guide the reproduction of the results in the main text.




FIG. 6. The assignment error [Eq. (12)] computed for $\tau = 40$ ns and varying rank M . The dashed black line at $\epsilon_a = 1.26 \times 10^{-3}$ corresponds to the assignment error obtained by integrating the full master equation.

1. Hilbert space

To build the Hilbert space basis on which Eq. (8) is numerically expressed, we first diagonalize the transmon Hamiltonian $\hat{H}_t = 4E_c \hat{n}_t^2 - E_j \cos \hat{\phi}_t$ in the flux basis, i.e., by assuming the canonical commutation relation $[\hat{\phi}_t, \hat{n}_t] = i$. The transmon Hamiltonian is then rewritten as $\hat{H}_t = \sum_{j_t=0}^{N_t-1} \omega_{j_t} |j_t\rangle \langle j_t|$, where $|j_t\rangle$ is the j_t th transmon eigenstate and only the lowest N_t eigenstates are retained. We also diagonalize the resonator-filter subsystem in the normal-mode basis $\{|j_l\rangle \otimes |j_u\rangle\}$, where $|j_l\rangle$ and $|j_u\rangle$ are the eigenstates of the lower and upper normal modes, respectively, with frequencies ω_l and ω_u . For readout, we choose the lower normal mode. Correspondingly, the Fock eigenspace of the lower normal mode is truncated to a large number of Fock states, up to $N_l = 100$, in order to account for its potentially large photon number. The other modes are off-resonant and are not expected to reach a large occupation. We therefore opt for smaller cutoffs in the transmon and upper normal modes, up to $N_t = 6$ and $N_u = 6$, respectively. The convergence of all the results has been checked with respect to these cutoffs.

We finally build the full Hilbert space of the coupled system as the tensor product of the transmon and normal-mode eigenstates, i.e., $\{|j_l\rangle \otimes |j_u\rangle \otimes |j_t\rangle\}$, where $j_l = 0, \dots, N_l - 1$, $j_u = 0, \dots, N_u - 1$, and $j_t = 0, \dots, N_t - 1$, and a total Hilbert space dimension $N = N_l \times N_u \times N_t$.

For the readout procedure, the state is initialized in the vacuum of the normal modes and the ground/first excited state of the transmon,

$$\hat{\rho}_{g/e}(0) = |\psi_{g/e}\rangle \langle \psi_{g/e}|, \quad (\text{A1})$$

$$|\psi_{g/e}\rangle = |0_l\rangle \otimes |0_u\rangle \otimes |g/e\rangle. \quad (\text{A2})$$

Note that these states are close to, but in general *different* from the lowest-energy eigenstates of the undriven Hamiltonian. They are exactly equal when the transmon-resonator coupling $g = 0$. Equation (5) is finally numerically integrated for a total readout time τ and expectation values of fields and occupations are evaluated.

2. Parameters

Throughout this work and unless otherwise stated, we set the transmon parameters to $E_c/2\pi = 315$ MHz and $E_j/E_c = 51$. This yields a qubit frequency $\omega_t/2\pi = 6$ GHz and anharmonicity $\alpha/2\pi = -349$ MHz. We take the couplings $g/2\pi = 150$ MHz and $J/2\pi = 30$ MHz. The frequency of the resonator is $\omega_r/2\pi = 7.2$ GHz, while that of the filter is $\omega_f/2\pi = 7.21$ GHz. The dissipation rates are $\kappa/2\pi = 30$ MHz and $\gamma/2\pi = 8$ kHz, for the transmon and filter, respectively. The frequency of the drive is taken to be $\omega_d/2\pi = 7.18$ GHz.

3. Drive pulse shape

As mentioned in the main text, a microwave drive pulse is applied to the Purcell filter modes at a frequency ω_d . The drive pulse may generally have a time-dependent, complex envelope given by $\Omega_d(t)$. In this work, we begin by setting $\Omega_d(t)$ to a square pulse shape and later generalize it to contain many steps.

For the initial simulations with a square pulse shape, $\Omega_d(t)$ is given by

$$\Omega_d(t) = |\Omega_d| \times f(t, t_0, \sigma) \times [1 - f(t, t_0 + T, \sigma)], \quad (\text{A3})$$

where

$$f(t, t_0, \sigma) = \frac{1}{1 + e^{-\frac{t-t_0}{\sigma}}} \quad (\text{A4})$$

is the logistic function [90]. The parameters t_0 and σ set the curve's center and steepness, respectively. The time it takes to go from an amplitude of zero to $|\Omega_d|$, referred to in the main text as ramp-up time, is given by $\sim 5\sigma$. In the present work, we choose $t_0 = 3$ ns and $\sigma = 0.5$ ns, corresponding to a ramp-up/down time of 2.5 ns.

The square pulse is then straightforwardly generalized to a smoothed stepwise function. For a given list of n_{steps} heights h and at a time t , the smoothed stepwise function outputs $\sum_{j=1}^{n_{\text{steps}}} (h_{j+1} - h_j) \times f(t, j\tau/n_{\text{steps}}, \sigma)$. An illustrative example of both the square and smoothed stepwise functions is shown in Fig. 2(b).

4. Diagonalization of $\hat{H}_{\text{rf}}$

We seek two normal modes $\hat{c}_l$ and $\hat{c}_u$ such that

$$\hat{H}_{\text{rf}} = \omega_r \hat{a}^\dagger \hat{a} + \omega_f \hat{f}^\dagger \hat{f} - J(\hat{a}^\dagger - \hat{a})(\hat{f}^\dagger - \hat{f}) \quad (\text{A5})$$

can be written as

$$\hat{H}_{\text{rf}} = \omega_l \hat{c}_l^\dagger \hat{c}_l + \omega_u \hat{c}_u^\dagger \hat{c}_u. \quad (\text{A6})$$

Therefore, we manifestly have

$$[\hat{H}_{\text{rf}}, \hat{c}_i] = -\omega_i \hat{c}_i, \quad (\text{A7})$$

where $i = \{l, u\}$. We define the normal modes in terms of the bare operators $\hat{a}$ and $\hat{f}$ as follows:

$$\hat{c}_i = \alpha_i \hat{a} + \beta_i \hat{f} + \gamma_i \hat{a}^\dagger + \delta_i \hat{f}^\dagger. \quad (\text{A8})$$

Note that we must account for the particle-hole nature of the c_i modes due to the presence the terms $-J\hat{a}^\dagger \hat{f}^\dagger - J\hat{a} \hat{f}$. This in turn yields the symmetry $\hat{c}_i \rightarrow \hat{c}_i^\dagger$ bringing $\omega \rightarrow -\omega$ following directly from Eq. (A7).

Substituting Eq. (A8) into Eq. (A7), we find four equations for the four coefficients of $\hat{c}_i$ by imposing that the terms attached to each operator $\hat{a}$, $\hat{f}$, $\hat{a}^\dagger$, and $\hat{f}^\dagger$ must match. These equations can be succinctly written in matrix form

$$\mathbf{\Omega} = \begin{pmatrix} \omega_r & J & 0 & J \\ J & \omega_f & J & 0 \\ 0 & -J & -\omega_r & -J \\ -J & 0 & -J & -\omega_f \end{pmatrix}. \quad (\text{A9})$$

Diagonalizing $\mathbf{\Omega}$ yields the two normal modes ω_l and ω_u along with their respective negative counterparts due to the aforementioned symmetry.

We now seek the transformation from $\hat{a}$, $\hat{f}$ to $\hat{c}_l$, $\hat{c}_u$ in order to express our Hamiltonian in terms of the normal modes. To do so, we write the system of equations

$$\mathbf{c} = U^\top \mathbf{a}, \quad (\text{A10})$$

which express the vector of operators $\mathbf{c} = (\hat{c}_u^\dagger, \hat{c}_l^\dagger, \hat{c}_l, \hat{c}_u)^\top$ as linear combinations of $\mathbf{a} = (\hat{a}, \hat{f}, \hat{a}^\dagger, \hat{f}^\dagger)^\top$. Note the ordering

of these operators is done to match the ordered list of frequencies, from $-\omega_u$ to ω_u , with $\omega_u > \omega_l$. The matrix U contains in each of its columns the eigenvectors of $\mathbf{\Omega}$ normalized subject to the condition

$$|\alpha|^2 + |\beta|^2 - |\gamma|^2 - |\delta|^2 = 1, \quad (\text{A11})$$

given by the commutation relations of the normal operators $[c_i, c_j^\dagger] = \delta_{ij}$. Note the difference with respect to the naive normalization $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$. One can verify that the commutation relations are preserved by this transformation via the equality

$$S \mathbf{\Sigma}_c S^\dagger = \mathbf{\Sigma}_a, \quad (\text{A12})$$

where we have defined $S = \text{inv } U^\top$ such that $\mathbf{a} = S\mathbf{c}$, and

$$\begin{aligned} (\mathbf{\Sigma}_a)_{ij} &= [a_i, a_j] \Rightarrow \begin{pmatrix} 0 & \mathbf{1} \\ -\mathbf{1} & 0 \end{pmatrix}, \\ (\mathbf{\Sigma}_c)_{ij} &= [c_i, c_j] \Rightarrow \begin{pmatrix} 0 & -\sigma_x \\ \sigma_x & 0 \end{pmatrix}. \end{aligned} \quad (\text{A13})$$

The coefficients of the transformation from the bare to the normal operators are the elements of the matrix S . Letting $(S)_{ij} = s_{ij}$, we can in particular denote

$$\hat{a} = s_{11}\hat{c}_u^\dagger + s_{12}\hat{c}_l^\dagger + s_{13}\hat{c}_l + s_{14}\hat{c}_u, \quad (\text{A14})$$

$$\hat{f} = s_{21}\hat{c}_u^\dagger + s_{22}\hat{c}_l^\dagger + s_{23}\hat{c}_l + s_{24}\hat{c}_u, \quad (\text{A15})$$

and so on for $a^\dagger$ and $f^\dagger$. Without loss of generality, we take s_{ij} to be fully real. The quantities of interest in the coupling and drive, $a^\dagger - a$ and $f^\dagger - f$, respectively, take the following forms:

$$a^\dagger - a = (s_{14} - s_{11})(\hat{c}_u^\dagger - \hat{c}_u) + (s_{13} - s_{12})(\hat{c}_l^\dagger - \hat{c}_l), \quad (\text{A16})$$

$$f^\dagger - f = (s_{24} - s_{21})(\hat{c}_u^\dagger - \hat{c}_u) + (s_{23} - s_{22})(\hat{c}_l^\dagger - \hat{c}_l). \quad (\text{A17})$$

We can finally define the constants $\mu_u = s_{14} - s_{11}$, $\mu_l = s_{13} - s_{12}$, $\nu_u = s_{24} - s_{21}$, and $\nu_l = s_{23} - s_{22}$.

5. Rotating-wave approximation on $\hat{H}$

We seek to isolate and eliminate the fast-rotating terms from our Hamiltonian equation (8),

$$\hat{H} = \hat{H}_t + \hat{H}_{\text{rf}} + ig\hat{n}_t(\hat{a}^\dagger - \hat{a}) + i\Omega_d(t)\sin(\omega_d t)(\hat{f}^\dagger - \hat{f}). \quad (\text{A18})$$

To identify these terms, we begin by expressing $\hat{H}$ in terms of the normal modes $\hat{c}_l$ and $\hat{c}_u$:

$$\begin{aligned} \hat{H} &= \hat{H}_t + \hat{H}_{\text{rf}} + ig\hat{n}_t[\mu_u(\hat{c}_u^\dagger - \hat{c}_u) + \mu_l(\hat{c}_l^\dagger - \hat{c}_l)] \\ &\quad + i\Omega_d(t)\sin(\omega_d t)[\nu_u(\hat{c}_u^\dagger - \hat{c}_u) + \nu_l(\hat{c}_l^\dagger - \hat{c}_l)]. \end{aligned} \quad (\text{A19})$$

We then rotate into the frame of the drive by use of the unitary

$$\hat{U}_d = \hat{U}_l \otimes \hat{U}_u \otimes \hat{U}_t, \quad (\text{A20})$$

where $\hat{U}_{l/u} = \exp(i\omega_d t \hat{c}_{l/u}^\dagger \hat{c}_{l/u})$ and

$$\hat{U}_t = \exp \left[i\omega_d t \sum_{j=1}^{N_t} (j-1) |\psi_j\rangle \langle \psi_j| + i\omega_1 t \right], \quad (\text{A21})$$

and ω_1 is the ground-state energy of the transmon Hamiltonian H_t . The Hamiltonian transforms according to $\hat{H} \rightarrow \hat{U}_d^\dagger \hat{H} \hat{U}_d - i\hat{U}_d^\dagger \dot{\hat{U}}_d$.

The normal-mode operators will simply acquire a time-dependent phase: $\hat{c}_{l/u} \rightarrow \hat{c}_{l/u} e^{-i\omega_d t}$. To elucidate the transformation of $\hat{n}_t$, it is useful to decompose it in terms of its counterclockwise- and clockwise-rotating terms:

$$\hat{n}_t = \sum_{j>k} \hat{n}_t^{jk} |j\rangle \langle k| + \sum_{j<k} \hat{n}_t^{jk} |j\rangle \langle k| = \hat{n}_t^- + \hat{n}_t^+. \quad (\text{A22})$$

The case $j=k$ is omitted since $\hat{n}_t$ contains only zeros on its diagonal. We then make the approximation $\hat{n}_t^+ \simeq \sum_j \hat{n}_t^{j,j+1} |j\rangle \langle j+1|$ and similarly for $\hat{n}_t^-$. In doing so, $\hat{n}_t^+$ now transforms in the familiar way $\hat{n}_t^+ \rightarrow \hat{n}_t^+ e^{-i\omega_d t}$.

The total Hamiltonian now transforms in a straightforward way. First, the bare Hamiltonians become

$$\hat{H}'_t = \sum_{j=1}^{N_t} [\omega_j - \omega_1 - \omega_d(j-1)] |j\rangle \langle j|, \quad (\text{A23})$$

$$\hat{H}'_{\text{rf}} = (\omega_l - \omega_d) \hat{c}_l^\dagger \hat{c}_l + (\omega_u - \omega_d) \hat{c}_u^\dagger \hat{c}_u. \quad (\text{A24})$$

The Hamiltonian in this rotated frame is then

$$\begin{aligned} \hat{H}_{\text{rot}} &= \hat{H}'_t + \hat{H}'_{\text{rf}} \\ &\quad + ig\mu_u(\hat{n}_t^+ \hat{c}_u^\dagger - \hat{n}_t^- \hat{c}_u - \hat{n}_t^+ \hat{c}_u e^{-2i\omega_d t} + \hat{n}_t^- \hat{c}_u^\dagger e^{2i\omega_d t}) \\ &\quad - \frac{\nu_u}{2} \Omega_d(t)(\hat{c}_u^\dagger + \hat{c}_u - \hat{c}_u^\dagger e^{2i\omega_d t} - \hat{c}_u e^{-2i\omega_d t}) \\ &\quad + (u \leftrightarrow l), \end{aligned} \quad (\text{A25})$$

where the last line should be understood as repeating the second and third lines but with the index u replaced by l . Equation (A25) is the Hamiltonian we use for the full and LRA simulations in the main text.

To obtain the RWA Hamiltonian, we simply ignore all explicitly time-dependent terms in the coupling and drive and arrive at the slow Hamiltonian

$$\begin{aligned} \hat{H}_{\text{rwa}} &= \hat{H}'_t + \hat{H}'_{\text{rf}} + ig\mu_l(\hat{n}_t^+ \hat{c}_l^\dagger - \hat{n}_t^- \hat{c}_l) \\ &\quad - \frac{\nu_l}{2} \Omega_d(t)(\hat{c}_l + \hat{c}_l^\dagger) + ig\mu_u(\hat{n}_t^+ \hat{c}_u^\dagger - \hat{n}_t^- \hat{c}_u) \\ &\quad - \frac{\nu_u}{2} \Omega_d(t)(\hat{c}_u + \hat{c}_u^\dagger). \end{aligned} \quad (\text{A26})$$

We draw particular attention to the factor of $1/2$ appearing in the driving terms proportional to $\Omega_d(t)$, originating from the sine function in the Hamiltonian (A18). This departs from the usual convention of writing an RWA driving Hamiltonian as $\hat{H}_d^{\text{rwa}}(t) = \Omega_d^{\text{rwa}}(t)(\hat{f} + \hat{f}^\dagger)$. In our convention, therefore, values of $\Omega_d(t)$ are twice as large as those quoted in Ref. [29].




FIG. 7. Optimization of the measurement fidelity performed with a rank $M = 80$. (a) The optimal value of the loss function for various readout times τ between 40 and 100 ns. Inset: A typical variation of the loss function along the optimization schedule for $\tau = 40$ ns. A minimum is reached in roughly 100 epochs. (b) Real (solid blue) and imaginary (dotted blue) parts of the optimal pulse envelope (light blue) for $\tau = 40$ ns. The envelope of the square pulse (dark blue) is plotted for comparison. Thin lines show the stepwise pulse before smoothing. (c) Phase-space trajectories for the optimal readout schedule and $\tau = 40$ ns, computed with $N_l = 5$, $N_u = 4$, and $N_l = 100$. The full and RWA trajectories are taken from a simulation in a Hilbert space of increased dimension $N_l = 6$, $N_u = 6$, and $N_l = 100$. Inset: Difference between $M = 20$ (dashed purple) and $M = 80$ (dashed blue) low-rank evolutions compared to the full result (blue) in the region of phase space where the discrepancy is most pronounced.

APPENDIX B: DETAILS OF THE FIDELITY OPTIMIZATION

1. Loss function

We define a loss function that balances two competing effects: minimizing the assignment error while avoiding transmon ionization. The assignment error is presented in Sec. IV A. The ionization penalty, for its part, is quantified by the population in higher transmon levels [29]:

$$T(\tau) = \frac{1}{\tau} \sum_{j=g,e} \sum_{k \geq 2}^{N_l} \int_0^\tau dt \operatorname{tr}[|k_t\rangle\langle k_t| \hat{\rho}_j(t)]. \quad (\text{B1})$$

Here, $\hat{\rho}_j(t)$ is the time-evolved density matrix of the full system, obtained for an initial transmon state $|j = g, e\rangle$, respectively. The transmon ionization during readout with a square pulse is shown in Fig. 4(b). The loss function to minimize is finally

$$\Lambda(\tau) = \log_{10} [\epsilon_a(\tau) + T(\tau)], \quad (\text{B2})$$

where the logarithm is taken to ensure that the loss function is of order 1.

We restrict the loss function to only these two contributions, as they are the most relevant, especially for increasing drive strength. A more complete multiobjective optimization may include additional contributions to penalize populating the undriven normal mode or approaching the critical photon number [29]. We find that the contribution of these terms to the loss function is minimal for the case studied here, and we therefore expect the minima of Eq. (B2) to remain close to those of a more comprehensive multiobjective loss function. Nevertheless, in order to safeguard against any further unaccounted ionization in the transmon and undriven normal mode, we impose an upper bound on the driving field strength

as $|\Omega_d^{\max}|/2\pi < 400$ MHz. This value coincides with that used in Ref. [29]. We also verify that the optimal result is characterized by very small values of these additional contributions.

2. Optimization process

We optimize the measurement fidelity by minimizing the loss function [Eq. (B2)] for a given readout time τ . The optimization parameters are the drive pulse heights, which generally consist of both real and imaginary parts. For n_{steps} steps, there are therefore $2n_{\text{steps}}$ parameters to optimize.

To avoid probing regimes that are not physical, we impose that the drive must be off at $t = 0$ and $t = \tau$. We also add a simple smoothing between neighboring steps, as discussed in Appendix A 3. Finally, we impose a strict upper bound on the amplitude of each step as $|\Omega_d^{\max}|/2\pi < 400$ MHz. Due to a different definition of $\Omega_d(t)$, this upper bound coincides with the one used in Ref. [29], as detailed above.

The drive pulse is initialized in the square pulse shape, i.e., with all real pulse steps equal and of amplitude $\Omega_d/2\pi = 150$ MHz, while all imaginary steps are zero. Each optimization step is then performed using a simultaneous perturbation stochastic approximation algorithm that approximates the full gradient as a finite difference directional derivative in a random direction [84,85]. Each step then requires only two function calls, which in turn each require two evolutions: one each for β_g and β_e . These evolutions are performed with the LRA for a rank of $M = 20$. The optimized results are validated with a single additional simulation in an enlarged Hilbert space, with the full Lindblad equation [Eq. (9)]. For these simulations, we also increase the rank to $M = 80$ in order to further verify the convergence of our final results with respect to the rank. We find no significant difference in the final pulse shapes when performing the optimization with $M = 80$ as opposed to the faster $M = 20$. The results of this optimization with $M = 80$ are displayed in Fig. 7.

- [1] A. P. Peirce, M. A. Dahleh, and H. Rabitz, Optimal control of quantum-mechanical systems: Existence, numerical approximation, and applications, *Phys. Rev. A* **37**, 4950 (1988).
- [2] J. Werschnik and E. K. U. Gross, Quantum optimal control theory, *J. Phys. B: At. Mol. Opt. Phys.* **40**, R175 (2007).
- [3] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Phys. Rev. A* **84**, 022305 (2011).
- [4] C. P. Koch, U. Boschain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, *et al.*, Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe, *EPJ Quantum Technol.* **9**, 19 (2022).
- [5] C. P. Koch, Controlling open quantum systems: tools, achievements, and limitations, *J. Phys.: Condens. Matter* **28**, 213001 (2016).
- [6] R. Schmidt, A. Negretti, J. Ankerhold, T. Calarco, and J. T. Stockburger, Optimal control of open quantum systems: Cooperative effects of driving and dissipation, *Phys. Rev. Lett.* **107**, 130404 (2011).
- [7] P. Doria, T. Calarco, and S. Montangero, Optimal control technique for many-body quantum dynamics, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [8] S. Lloyd and S. Montangero, Information theoretical analysis of quantum optimal control, *Phys. Rev. Lett.* **113**, 010502 (2014).
- [9] F. Sauvage and F. Mintert, Optimal control of families of quantum gates, *Phys. Rev. Lett.* **129**, 050507 (2022).
- [10] B. Sarma and M. J. Hartmann, Designing fast quantum gates using optimal control with a reinforcement-learning ansatz, *Phys. Rev. Appl.* **23**, 014015 (2025).
- [11] S. Jandura and G. Pupillo, Time-optimal two- and three-qubit gates for Rydberg atoms, *Quantum* **6**, 712 (2022).
- [12] T. Schulte-Herbrüggen, A. Spoerl, N. Khaneja, and S. J. Glaser, Optimal control for generating quantum gates in open dissipative systems, *J. Phys. B: At. Mol. Opt. Phys.* **44**, 154013 (2011).
- [13] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, *npj Quantum Inf.* **7**, 14 (2021).
- [14] R. Gautier, M. Mirrahimi, and A. Sarlette, Designing high-fidelity zeno gates for dissipative cat qubits, *PRX Quantum* **4**, 040316 (2023).
- [15] T. Propson, B. E. Jackson, J. Koch, Z. Manchester, and D. I. Schuster, Robust quantum optimal control with trajectory optimization, *Phys. Rev. Appl.* **17**, 014036 (2022).
- [16] R. Blümel, N. Grzesiak, N. H. Nguyen, A. M. Green, M. Li, A. Maksymov, N. M. Linke, and Y. Nam, Efficient stabilized two-qubit gates on a trapped-ion quantum computer, *Phys. Rev. Lett.* **126**, 220503 (2021).
- [17] A. Rahman, D. J. Egger, and C. Arenz, Learning how to dynamically decouple by optimizing rotational gates, *Phys. Rev. Appl.* **22**, 054074 (2024).
- [18] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, *et al.*, High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler, *Phys. Rev. X* **13**, 031035 (2023).
- [19] E. Hyppä, A. Vepsäläinen, M. Papič, C. F. Chan, S. Inel, A. Landra, W. Liu, J. Luus, F. Marxer, C. Ockeloen-Korppi, *et al.*, Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes, *PRX Quantum* **5**, 030353 (2024).
- [20] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard, A. Greene, G. O. Samach, C. McNally, D. Kim, *et al.*, Realization of high-fidelity CZ and ZZ-free iSWAP gates with a tunable coupler, *Phys. Rev. X* **11**, 021058 (2021).
- [21] V. Tripathi, H. Chen, M. Khezri, K.-W. Yip, E. M. Levenson-Falk, and D. A. Lidar, Suppression of crosstalk in superconducting qubits using dynamical decoupling, *Phys. Rev. Appl.* **18**, 024068 (2022).
- [22] D. J. Egger and F. K. Wilhelm, Optimized controlled-Z gates for two superconducting qubits coupled through a resonator, *Supercond. Sci. Technol.* **27**, 014001 (2014).
- [23] R. Blümel, N. Grzesiak, N. Pienti, K. Wright, and Y. Nam, Power-optimal, stabilized entangling gate between trapped-ion qubits, *npj Quantum Inf.* **7**, 147 (2021).
- [24] M. Mohan, R. de Keijzer, and S. Kokkelmans, Robust control and optimal Rydberg states for neutral atom two-qubit gates, *Phys. Rev. Res.* **5**, 033052 (2023).
- [25] M. Kuzmanović, I. Moskalenko, Y.-H. Chang, O. Stanisavljević, C. Warren, E. Hogedal, A. Aggarwal, I. Ahmad, J. Biznárová, M. Dahiya, *et al.*, Neural-network-based design and implementation of fast and robust quantum gates, [arXiv:2505.02054](https://arxiv.org/abs/2505.02054).
- [26] E. Genois, N. J. Stevenson, N. Goss, I. Siddiqi, and A. Blais, Quantum optimal control of superconducting qubits based on machine-learning characterization, *Phys. Rev. Appl.* **24**, 034073 (2025).
- [27] H. N. Nguyen, F. Motzoi, M. Metcalf, K. B. Whaley, M. Bukov, and M. Schmitt, Reinforcement learning pulses for transmon qubit entangling gates, *Mach. Learn.: Sci. Technol.* **5**, 025066 (2024).
- [28] J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, *et al.*, Rapid high-fidelity multiplexed readout of superconducting qubits, *Phys. Rev. Appl.* **10**, 034040 (2018).
- [29] R. Gautier, E. Genois, and A. Blais, Optimal control in large open quantum systems: The case of transmon readout and reset, *Phys. Rev. Lett.* **134**, 070802 (2025).
- [30] L. Chen, S. P. Fors, Z. Yan, A. Ali, T. Abad, A. Osman, E. Moschandreou, B. Lienhard, S. Kosen, H.-X. Li, *et al.*, Fast unconditional reset and leakage reduction in fixed-frequency transmon qubits, [arXiv:2409.16748](https://arxiv.org/abs/2409.16748).
- [31] M. Abdelhafez, D. I. Schuster, and J. Koch, Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation, *Phys. Rev. A* **99**, 052327 (2019).
- [32] J. Gambetta, W. A. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement, *Phys. Rev. A* **76**, 012325 (2007).
- [33] P. D. Kurilovich, T. Connolly, C. G. L. Böttcher, D. K. Weiss, S. Hazra, V. R. Joshi, A. Z. Ding, H. Nho, S. Diamond, V. D. Kurilovich, *et al.*, High-frequency readout free from transmon multi-excitation resonances, [arXiv:2501.09161](https://arxiv.org/abs/2501.09161).

- [34] A. A. Chapple, O. Benhayoune-Khadraoui, S. Richer, and A. Blais, Balanced cross-Kerr coupling for superconducting qubit readout, *Phys. Rev. Lett.* **135**, 256002 (2025).
- [35] F. Swiadek, R. Shillito, P. Magnard, A. Remm, C. Hellings, N. Lacroix, Q. Ficheux, D. C. Zanuz, G. J. Norris, A. Blais, *et al.*, Enhancing dispersive readout of superconducting qubits through dynamic control of the dispersive shift: Experiment and theory, *PRX Quantum* **5**, 040326 (2024).
- [36] L. Chen, H.-X. Li, Y. Lu, C. W. Warren, C. J. Križan, S. Kosen, M. Rommel, S. Ahmed, A. Osman, J. Biznárová, *et al.*, Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier, *npj Quantum Inf.* **9**, 26 (2023).
- [37] A. Bengtsson, A. Opremcak, M. Khezri, D. Sank, A. Bourassa, K. J. Satzinger, S. Hong, C. Erickson, B. J. Lester, K. C. Miao, *et al.*, Model-based optimization of superconducting qubit readout, *Phys. Rev. Lett.* **132**, 100603 (2024).
- [38] A. Chatterjee, J. Schwinger, and Y. Y. Gao, Enhanced qubit readout via reinforcement learning, *Phys. Rev. Appl.* **23**, 054057 (2025).
- [39] D. Egger, M. Werninghaus, M. Ganzhorn, G. Salis, A. Fuhrer, P. Müller, and S. Filipp, Pulsed reset protocol for fixed-frequency superconducting qubits, *Phys. Rev. Appl.* **10**, 044030 (2018).
- [40] M. McEwen, D. Kafri, Z. Chen, J. Atalaya, K. J. Satzinger, C. Quintana, P. V. Klimov, D. Sank, C. Gidney, A. G. Fowler, *et al.*, Removing leakage-induced correlated errors in superconducting quantum error correction, *Nat. Commun.* **12**, 1761 (2021).
- [41] P. V. Klimov, A. Bengtsson, C. Quintana, A. Bourassa, S. Hong, A. Dunsworth, K. J. Satzinger, W. P. Livingston, V. Sivak, M. Y. Niu, *et al.*, Optimizing quantum gates towards the scale of logical qubits, *Nat. Commun.* **15**, 2442 (2024).
- [42] P. V. Klimov, J. Kelly, J. M. Martinis, and H. Neven, The snake optimizer for learning quantum processor control parameters, *arXiv:2006.04594*.
- [43] V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, *et al.*, Real-time quantum error correction beyond break-even, *Nature (London)* **616**, 50 (2023).
- [44] R. Acharya, D. A. Abanin, L. Aghababaie-Beni, I. Aleiner, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, N. Astrakhantsev, *et al.*, Quantum error correction below the surface code threshold, *Nature (London)* **638**, 920 (2025).
- [45] V. Sivak, M. Newman, and P. Klimov, Optimization of decoder priors for accurate quantum error correction, *Phys. Rev. Lett.* **133**, 150603 (2024).
- [46] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, *et al.*, Suppressing quantum errors by scaling a surface code logical qubit, *Nature (London)* **614**, 676 (2023).
- [47] H. P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2007).
- [48] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [49] V. F. Krotov, *Global Methods in Optimal Control Theory* (Birkhäuser, Boston, MA, 1993), pp. 74–121.
- [50] M. H. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich, and C. P. Koch, Krotov: A Python implementation of Krotov’s method for quantum optimal control, *SciPost Phys.* **7**, 080 (2019).
- [51] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [52] M. M. Müller, S. Gherardini, T. Calarco, S. Montangero, and F. Caruso, Information theoretical limits for quantum optimal control solutions: Error scaling of noisy control channels, *Sci. Rep.* **12**, 21405 (2022).
- [53] N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, Speedup for quantum optimal control from automatic differentiation based on graphics processing units, *Phys. Rev. A* **95**, 042318 (2017).
- [54] S. Boutin, C. K. Andersen, J. Venkatraman, A. J. Ferris, and A. Blais, Resonator reset in circuit QED by optimal control for large open quantum systems, *Phys. Rev. A* **96**, 042315 (2017).
- [55] D. J. Egger and F. K. Wilhelm, Optimal control of a quantum measurement, *Phys. Rev. A* **90**, 052331 (2014).
- [56] A. G. Kofman and G. Kurizki, Unified theory of dynamically suppressed qubit decoherence in thermal baths, *Phys. Rev. Lett.* **93**, 130406 (2004).
- [57] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. M. Lewis, G. Vidal, and A. Blais, Dynamics of transmon ionization, *Phys. Rev. Appl.* **18**, 034031 (2022).
- [58] M. F. Dumas, B. Groleau-Paré, A. McDonald, M. H. Muñoz-Arias, C. Lledó, B. D’Anjou, and A. Blais, Measurement-induced transmon ionization, *Phys. Rev. X* **14**, 041023 (2024).
- [59] D. Sank, Z. Chen, M. Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, *et al.*, Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation, *Phys. Rev. Lett.* **117**, 190503 (2016).
- [60] C. Le Bris and P. Rouchon, Low-rank numerical approximations for high-dimensional Lindblad equations, *Phys. Rev. A* **87**, 022125 (2013).
- [61] G. McCaul, K. Jacobs, and D. I. Bondar, Fast computation of dissipative quantum systems with ensemble rank truncation, *Phys. Rev. Res.* **3**, 013017 (2021).
- [62] Y.-T. Chen, C. Farquhar, and R. M. Parrish, Low-rank density-matrix evolution for noisy quantum circuits, *npj Quantum Inf.* **7**, 61 (2021).
- [63] K. Donatella, Z. Denis, A. Le Boité, and C. Ciuti, Continuous-time dynamics and error scaling of noisy highly entangling quantum circuits, *Phys. Rev. A* **104**, 062407 (2021).
- [64] S. Santos, X. Song, and V. Savona, Low-rank variational quantum algorithm for the dynamics of open quantum systems, *Quantum* **9**, 1620 (2025).
- [65] L. Joubert-Doriol, I. G. Ryabinkin, and A. F. Izmaylov, Non-stochastic matrix Schrödinger equation for open systems, *J. Chem. Phys.* **141**, 234112 (2014).
- [66] L. Joubert-Doriol and A. F. Izmaylov, Problem-free time-dependent variational principle for open quantum systems, *J. Chem. Phys.* **142**, 134107 (2015).
- [67] L. Gravina and V. Savona, Adaptive variational low-rank dynamics for open quantum systems, *Phys. Rev. Res.* **6**, 023072 (2024).
- [68] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation, *Phys. Rev. A* **69**, 062320 (2004).

- [69] S. Touzard, A. Kou, N. E. Frattini, V. V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. H. Devoret, Gated conditional displacement readout of superconducting qubits, *Phys. Rev. Lett.* **122**, 080502 (2019).
- [70] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, *et al.*, Rapid high-fidelity single-shot dispersive readout of superconducting qubits, *Phys. Rev. Appl.* **7**, 054020 (2017).
- [71] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura, Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic purcell filter, *Phys. Rev. Appl.* **17**, 044016 (2022).
- [72] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, Fast reset and suppressing spontaneous emission of a superconducting qubit, *Appl. Phys. Lett.* **96**, 203110 (2010).
- [73] A. Blais, A. L. Grimsmo, S. Girvin, and A. Wallraff, Circuit quantum electrodynamics, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [74] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Matrix product density operators: Simulation of finite-temperature and dissipative systems, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [75] J. Cui, J. I. Cirac, and M. C. Bañuls, Variational matrix product operators for the steady state of dissipative quantum systems, *Phys. Rev. Lett.* **114**, 220601 (2015).
- [76] A. 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).
- [77] Y. Guo and S. Yang, Locally purified density operators for noisy quantum circuits, *Chin. Phys. Lett.* **41**, 120302 (2024).
- [78] Y. Guo, J.-H. Zhang, H.-R. Zhang, S. Yang, and Z. Bi, Locally purified density operators for symmetry-protected topological phases in mixed states, *Phys. Rev. X* **15**, 021060 (2025).
- [79] D. Manzano, A short introduction to the Lindblad master equation, *AIP Adv.* **10**, 025106 (2020).
- [80] M. Khezri, A. Opremcak, Z. Chen, K. C. Miao, M. McEwen, A. Bengtsson, T. White, O. Naaman, D. Sank, A. N. Korotkov, *et al.*, Measurement-induced state transitions in a superconducting qubit: Within the rotating-wave approximation, *Phys. Rev. Appl.* **20**, 054008 (2023).
- [81] L. Goutte, GitHub (2025), https://github.com/leogoutte/low_rank.
- [82] A. Mercurio, Y.-T. Huang, L.-X. Cai, Y.-N. Chen, V. Savona, and F. Nori, QuantumToolbox.jl: An efficient Julia framework for simulating open quantum systems, *Quantum* **9**, 1866 (2025).
- [83] C. C. Bultink, B. Tarasinski, N. Haandbæk, S. Poletto, N. Haider, D. J. Michalak, A. Bruno, and L. DiCarlo, General method for extracting the quantum efficiency of dispersive qubit readout in circuit QED, *Appl. Phys. Lett.* **112**, 092601 (2018).
- [84] J. Spall, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, *IEEE Trans. Autom. Control* **37**, 332 (1992).
- [85] S. Bhatnagar, H. L. Prasad, and L. A. Prashanth, *Stochastic Recursive Algorithms for Optimization: Simultaneous Perturbation Methods* (Springer, London, 2013).
- [86] J. Kennedy and R. Eberhart, in *Proceedings of the ICNN'95: International Conference on Neural Networks* (ICNN, Perth, WA, Australia, 1995), Vol. 4, pp. 1942–1948.
- [87] D. T. McClure, H. Paik, L. S. Bishop, M. Steffen, J. M. Chow, and J. M. Gambetta, Rapid driven reset of a qubit readout resonator, *Phys. Rev. Appl.* **5**, 011001(R) (2016).
- [88] M. H. Muñoz Arias, C. Lledó, and A. Blais, Qubit readout enabled by qubit cloaking, *Phys. Rev. Appl.* **20**, 054013 (2023).
- [89] J. Ikonen, J. Goetz, J. Ilves, A. Keränen, A. M. Gunyho, M. Partanen, K. Y. Tan, D. Hazra, L. Grönberg, V. Vesterinen, *et al.*, Qubit measurement by multichannel driving, *Phys. Rev. Lett.* **122**, 080503 (2019).
- [90] G. B. Arfken, H. J. Weber, and F. E. Harris, *Mathematical Methods for Physicists: A Comprehensive Guide*, 7th ed. (Academic Press, Waltham, MA, 2013).