

Optimal Control in Large Open Quantum Systems: The Case of Transmon Readout and Reset

Ronan Gautier^{1,2,3}, Élie Genois¹, and Alexandre Blais^{1,4}

¹*Institut Quantique and Département de Physique, Université de Sherbrooke, Sherbrooke, Quebec, Canada*

²*Laboratoire de Physique de l'École Normale Supérieure,*

Inria, ENS, Mines ParisTech, Université PSL, Sorbonne Université, Paris, France

³*Alice & Bob, Paris, France*

⁴*Canadian Institute for Advanced Research, Toronto, Ontario, Canada*

 (Received 22 April 2024; revised 9 December 2024; accepted 27 January 2025; published 20 February 2025)

We present a framework that combines the adjoint-state method together with reverse-time back-propagation to solve prohibitively large open-system quantum control problems. Our approach enables the optimization of arbitrary cost functions with fully general controls applied on large open quantum systems described by a Lindblad master equation. It is scalable, computationally efficient, and has a low-memory footprint. We apply this framework to optimize two inherently dissipative operations in superconducting qubits which lag behind in terms of fidelity and duration compared to other unitary operations: the dispersive readout and all-microwave reset of a transmon qubit. Our results show that while standard pulses for dispersive readout are nearly optimal, adding a transmon drive during the protocol can yield $2\times$ improvements in fidelity and duration. We further demonstrate a $2\times$ improvement in reset fidelity and duration through pulse shaping, indicating significant potential for enhancement in reset protocols. Our approach can readily be applied to optimize quantum controls in a vast range of applications such as reservoir engineering, autonomous quantum error correction, and leakage-reduction units.

DOI: [10.1103/PhysRevLett.134.070802](https://doi.org/10.1103/PhysRevLett.134.070802)

Introduction—Quantum optimal control (QOC) provides a framework to design external controls for realizing arbitrary quantum operations with maximal fidelity and minimal time [1–3], crucial requirements of useful quantum error correction [4]. A common assumption of QOC is that minimizing operation time will also reduce the impact of environmental noise, such that only closed quantum systems need to be considered. However, these approaches are limited by the fact that controlling the system's coherent dynamics can drastically alter the impact of some noise sources, as exemplified by dynamical decoupling methods [5–7]. Moreover, closed-system approaches cannot extend to inherently dissipative processes such as qubit readout and reset. Consequently, optimally controlling open quantum systems emerges as an important avenue [8–10]. It addresses both the minimization of decoherence in quantum information processing [11–13] and the design of dissipative protocols [14–16], marking a significant step toward comprehensive quantum control in engineered systems.

Over the last decade, several approaches have emerged for open-system QOC. Closed-loop control methods, such as feedback engineering [17,18] and reinforcement learning [19–22], have seen recent success but are difficult to scale to many parameters. Open-loop methods, including gradient ascent pulse engineering [15,23], Krotov's method [24–26], and automatic differentiation [27–30], circumvent this

limitation by computing numerical gradients, enabling efficient exploration of the parameter space. However, controlling large open systems remains challenging: The sheer problem size prohibits frameworks that require vectorizing the Liouvillian or storing density matrices at each time step. Although methods for low-memory differentiation have been proposed for closed quantum systems [31,32], extending them to open systems while maintaining a favorable memory overhead remains unresolved.

In this Letter, we present a framework enabling the realization of QOC on large open quantum systems with a fully general parametrization over the controls and arbitrary cost functions. Our approach combines the adjoint-state method [33–35] with reverse-time backpropagation [36–39] to reduce the memory cost of differentiation from linear to constant relative to the number of numerical integration steps. This significant reduction enables solving otherwise prohibitively large open-system quantum control problems defined in Lindblad form [40], and makes the scheme ideal for GPU acceleration. Our approach thus ensures precise, fast, and scalable computation of arbitrary gradients. We apply this method to optimize two critical operations for the realization of a fault-tolerant quantum computer based on superconducting circuits: dispersive readout [41–43] and all-microwave reset [44,45] of a transmon qubit [46]. Using the signal-to-noise ratio (SNR) of the readout as a cost function, we find several optimal controls with increasing

experimental complexity and up to $2\times$ improvements in fidelity and duration compared to standard protocols. For reset, we show that pulse shaping alone can halve the operation duration, an improvement of high practical relevance.

Adjoint-state method—Consider a QOC problem for which we seek to find a set of parameters minimizing a cost function $C(\theta, \hat{\rho}(t_0), \dots, \hat{\rho}(t_n))$. This function, in general, depends on both the problem parameters $\theta = (\theta_1, \dots, \theta_m)$ and on the density matrix of the system at a set of times $\hat{\rho}(t_i)$. Gradient-based approaches to optimize the control parameters rely on computing the derivative of the cost function with respect to each parameter $dC/d\theta$. To do so, we apply the adjoint-state method [33] to open quantum systems. In this context, the adjoint state is defined as $\hat{\phi}(t) = dC/d\hat{\rho}(t)$ and represents how a change in the density matrix at time t modifies the cost function. For open quantum systems under the usual Born-Markov approximations [47], the evolution of the density matrix is governed by a Lindblad master equation ($\hbar = 1$),

$$\frac{d\hat{\rho}}{dt} = \mathcal{L}\hat{\rho} \equiv -i[\hat{H}, \hat{\rho}] + \sum_k \mathcal{D}[\hat{L}_k]\hat{\rho}, \quad (1)$$

where \hat{H} is the system Hamiltonian, \hat{L}_k are jump operators, and $\mathcal{D}[\hat{L}]\hat{\rho} = \hat{L}\hat{\rho}\hat{L}^\dagger - \{\hat{L}^\dagger\hat{L}, \hat{\rho}\}/2$. The adjoint state is then subject to a dual ordinary differential equation [48],

$$\frac{d\hat{\phi}}{dt} = -\mathcal{L}^\dagger\hat{\phi} \equiv -i[\hat{H}, \hat{\phi}] - \sum_k \mathcal{D}^\dagger[\hat{L}_k]\hat{\phi}, \quad (2)$$

where $\mathcal{D}^\dagger[\hat{L}]\hat{\phi} = \hat{L}^\dagger\hat{\phi}\hat{L} - \{\hat{L}^\dagger\hat{L}, \hat{\phi}\}/2$. This equation can be integrated numerically over the time interval of interest $[t_0, t_n]$ with initial condition $\hat{\phi}(t_n) = \partial C/\partial\hat{\rho}(t_n)$ computed analytically if a closed form is available, or directly through automatic differentiation. Notably, the overall minus sign in Eq. (2) ensures numerical stability of the integration by generating contracting dynamics in reverse time. The derivative of the cost function with respect to the problem parameters is given by

$$\frac{dC}{d\theta} = \frac{\partial C}{\partial\theta} - \int_{t_n}^{t_0} \partial_\theta \text{Tr} \left[\hat{\phi}^\dagger(t) \mathcal{L}(t, \theta) \hat{\rho}(t) \right] dt. \quad (3)$$

This integral is straightforward to compute using the density matrix and adjoint state at each time $t \in [t_0, t_n]$, as obtained from Eqs. (1) and (2). In particular, the partial derivative with respect to θ can be easily computed from automatic differentiation of the adjoint-state equation by noting that

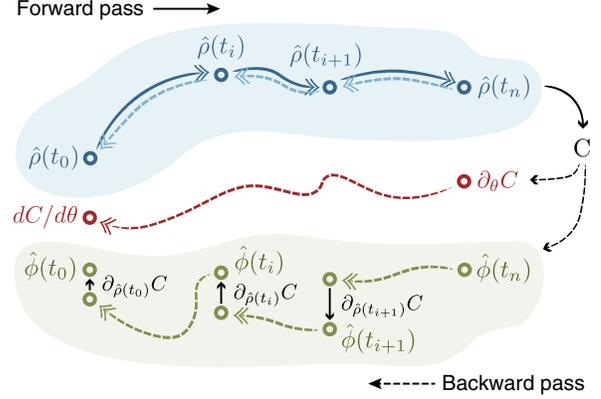


FIG. 1. Adjoint state quantum optimal control. In the forward pass, the master equation is integrated and checkpointed for several time points (dark blue). In the backward pass, the density matrix $\hat{\rho}$ is recomputed in reverse time (light blue) together with the adjoint state $\hat{\phi}$ (green) and with the gradients (red) of the cost function C . When a checkpoint is reached, the density matrix is restored to its forward time trajectory, and the adjoint state updated with the corresponding cost function gradient.

$$\partial_\theta \text{Tr} \left[\hat{\phi}^\dagger \mathcal{L} \hat{\rho} \right] = -\text{Tr} \left[\partial_\theta (d\hat{\phi}/dt)^\dagger \hat{\rho} \right], \quad (4)$$

which has the form of a vector-Jacobian product.

The QOC optimization is illustrated in Fig. 1 and proceeds in two steps. First, the forward pass consists in using the initial set of parameters (e.g., a sequence of discrete pulses) to numerically integrate the master equation from t_0 to t_n while saving the density matrix at each time t_i of interest. The cost function $C(\theta, \hat{\rho}(t_0), \dots, \hat{\rho}(t_n))$ is then evaluated. To lower the memory footprint, the cost function can also be evaluated on the fly during the forward pass such that only a single density matrix needs to be stored. In a second step, the backward pass, both the master and adjoint equations are simultaneously integrated in reverse time, starting from $t = t_n$. During this process, the integral of Eq. (3) is iteratively evaluated, such as to obtain the entire gradients $dC/d\theta$ once the backpropagation is finished. Having access to the gradients of the cost function, we can now iteratively update the control parameters using standard optimization algorithms [58–60].

We emphasize how each density matrix (blue) is computed twice: once during the forward pass, and once during the backward pass. This enables a low-memory footprint for the overall scheme, with at most a single density matrix and adjoint state needed to be stored at any given moment. The memory footprint of the method thus scales as $\mathcal{O}(N^2)$ with N the Hilbert space dimension. This is in stark contrast with methods based on automatic differentiation [28], for which the density matrix needs to be stored at each time point of the numerical integration, thus scaling as $\mathcal{O}(nN^2)$, with n the number of numerical integration steps [32]. Such

memory requirements can quickly become prohibitive, even for open quantum systems of intermediate sizes $N \gtrsim 100$.

Note that this large gain in memory comes at the cost of trading off some numerical run-time. Overall, the scheme requires the integration of four differential equations in total [48], against only two for automatic differentiation. In addition, the reverse-time integration of Eq. (1) can be numerically unstable due to the expansive dynamics of the system. This can however be fully resolved with checkpointing the quantum states during the forward pass [39], thus effectively trading back some memory for numerical stability. In practice, checkpointing at the timescale of the largest dissipation operator is sufficient to ensure numerical stability without adding significant complexity.

We have implemented this optimization scheme using PyTorch [61], taking advantage of its automatic differentiation capabilities and GPU support. This framework allows us to run optimization problems for an open quantum system with hundreds of parameters, arbitrary cost functions, and for Hilbert space dimensions of up to $N \sim 5000$ while running on a single GPU with 24 GB of memory. Our code is available through the DYNAMIQS open-source library [49], simplifying replication of this Letter and its application to various QOC problems. We now demonstrate the usefulness of this method by optimizing readout and reset of a transmon, two operations that inherently rely on dissipation.

Transmon model—Let us consider the experimentally realistic model depicted in Fig. 2(a) of a transmon coupled to a readout resonator and Purcell filter [50]

$$\frac{d\hat{\rho}}{dt} = -i[\hat{H}, \hat{\rho}] + \gamma\mathcal{D}[\hat{b}]\hat{\rho} + \kappa\mathcal{D}[\hat{f}]\hat{\rho}, \quad (5)$$

with transmon relaxation rate γ and filter relaxation rate κ , and where

$$\begin{aligned} \hat{H} = & 4E_C\hat{n}_t - E_J\cos(\hat{\phi}_t) + \omega_r\hat{a}^\dagger\hat{a} + \omega_f\hat{f}^\dagger\hat{f} \\ & - ig\hat{n}_t(\hat{a} - \hat{a}^\dagger) - J(\hat{a} - \hat{a}^\dagger)(\hat{f} - \hat{f}^\dagger) \\ & + \Omega_r\hat{n}_t\sin(\omega_{d,r}t) - i\Omega_f(\hat{f} - \hat{f}^\dagger)\sin(\omega_{d,f}t). \end{aligned} \quad (6)$$

The first two terms denote the free transmon Hamiltonian with charging energy E_C and Josephson energy E_J , with \hat{n}_t and $\hat{\phi}_t$ the charge and phase operators, and with \hat{b} the corresponding annihilation operator in the diagonal basis. The resonator and filter modes are denoted by \hat{a} and \hat{f} , with respective frequencies ω_r and ω_f . These three modes are capacitively coupled in series with coupling strengths $g \gg J$. The system can be driven using a capacitive coupling either through the transmon with a microwave pulse at frequency $\omega_{d,r}$ and envelope $\Omega_r(t)$, or through the Purcell filter at frequency $\omega_{d,f}$ and envelope $\Omega_f(t)$.

For numerical simulation of this model, we first diagonalize the free transmon Hamiltonian and identify

the lowest energy eigenstates. We also diagonalize the resonator-filter subsystem yielding two normal modes, each coupled to the transmon. Finally, we apply the rotating-wave approximation on couplings and drives. This allows for larger numerical time steps by eliminating fast oscillating dynamics thereby simplifying master equation integration. However, this also implies that not all of the chaotic or transmon ionization dynamics are captured [62–64]. To avoid probing these regimes, we limit the maximum amplitudes of control drives, e.g., to 200 MHz for transmon readout.

We use typical device parameters corresponding to a critical photon number of $\bar{n}_{\text{crit}} = (\Delta/2g)^2 = 16$ [41] and dispersive rates of $\chi/2\pi = 3.8$ and 8.1 MHz with the lower and higher normal modes, respectively. The filter loss rate is $\kappa/2\pi = 30$ MHz, and the transmon relaxation time is $T_1 = 20 \mu\text{s}$. The remaining system parameters can be found in [48].

Transmon readout—The readout of transmon qubits is realized through the dispersive coupling to a resonator [41]. In this case, the resonator frequency is shifted by the average occupancy in the transmon, and can be measured by driving the resonator at its bare frequency and monitoring the output field; see Fig. 2(b). In the presence of a Purcell filter, either normal mode of the hybridized resonator-filter subsystem can be used for readout [51].

The metric we use to maximize the measurement fidelity is the SNR. Accounting for optimal weighting functions [65], it reads [66]

$$\text{SNR}(\tau_m) = \sqrt{2\eta\kappa \int_0^{\tau_m} dt |\beta_e(t) - \beta_g(t)|^2}, \quad (7)$$

where $\eta \in [0, 1]$ is the measurement efficiency, τ_m is the readout integration time, and $\beta_{e/g} = \text{Tr}[\hat{f}\hat{\rho}_{g/e}]$ is the average field value in the filter mode, with $\hat{\rho}_{g/e}$ the density matrix obtained after initializing the transmon in the $|g/e\rangle$ state. To obtain results that can be compared to experiments, we use $\eta = 0.6$ [42]. The optimization objective is to maximize the SNR, and thus maximize the distance between the pointer states $|\beta_e - \beta_g|$ in the shortest possible time. Further assuming that the pointer states $\beta_{g,e}$ are Gaussian, one can link the SNR and the transmon lifetime to the readout assignment error [48,52].

To optimize the transmon readout, we discretize the control pulse envelopes $\Omega(t)$ with 1 ns time bins and use a 250 MHz Gaussian filter to interpolate between these pixels during numerical integration and to model realistic experimental distortions [53]. In addition to the discretized drive amplitudes, the optimization parameters θ include the carrier frequency ω_d of each drive. Contrary to the drive amplitudes, the latter are kept constant throughout the pulse duration, in accordance with typical experiments. The cost function used to optimize the transmon readout is

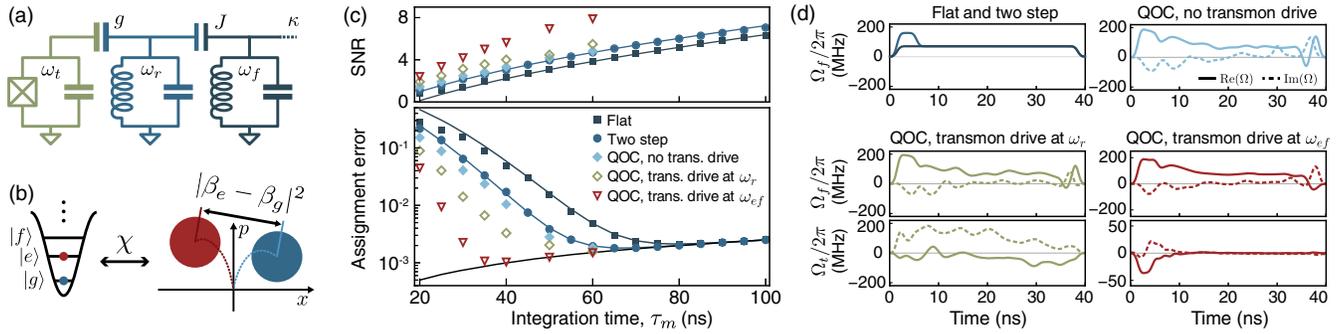


FIG. 2. (a) Lumped-element model of a transmon coupled to a readout resonator and Purcell filter. The transmon and filter are driven, and the filter output field is measured through its transmission line. (b) Dispersive readout of a transmon. The mean field in the filter depends on the transmon state. The SNR of the readout increases with the integrated difference of mean fields. (c) SNR and assignment error of the transmon readout for several drive envelopes: a flat envelope, a 4 ns two-step envelope [42], and optimized envelopes (QOC) with optional additional drives on the transmon at frequencies $\omega_{d,t} \simeq \omega_r$ (green) and $\omega_{d,t} \simeq \omega_{ef}$ (red). The flat and two-step data points are fitted according to Eq. (8). The black line shows the T_1 limit given by $\tau_m/2T_1$. (d) Reference and optimized pulse envelopes at 40 ns of integration time.

principally composed of the SNR of Eq. (7), with additional cost terms constraining the control pulses in order to regularize the optimization and avoid out-of-model dynamics. For example, we limit the number of photons in the hybridized resonator-filter modes, penalize unwanted transitions to higher excited transmon states, and limit the maximal available pulse amplitudes. The full cost function is detailed in [48]. We perform gradient descent using ADAM [59] and use the adjoint-state method previously described to compute gradients.

Figure 2(c) shows the SNR and the assignment error as a function of the integration time τ_m obtained by our approach, and Fig. 2(d) shows the corresponding pulse envelopes for $\tau_m = 40$ ns optimizations. As a point of comparison, we first consider the two nonoptimized reference pulses labeled “flat” and “two step.” The former consists of a constant pulse with 2 ns ramp-up and ramp-down times (dark blue squares), and the latter of a two-step pulse meant to rapidly populate the readout mode (blue circles) [42]. In both cases, the amplitude is calibrated to reach $\bar{n} = \bar{n}_{\text{crit}}$ photons in the steady state. The SNR versus τ_m for these two pulses is fitted with the function (full blue and dark blue lines) [50]

$$\text{SNR}(\tau_m) = \alpha \sqrt{2\eta\kappa} (\sqrt{\tau_m} - \sqrt{\tau_{m,0}}), \quad (8)$$

where $\alpha = 2|\Omega_f \sin(2\phi)|/\kappa$ is the effective resonator displacement in the steady state, with $\phi = \arctan(2\chi/\kappa)$ and χ the dispersive shift obtained from exact diagonalization of Eq. (6). In this expression, $\sqrt{\tau_{m,0}}$ accounts for an initial delay for the resonator to populate, and is numerically fitted to $\tau_{m,0} = 19$ ns and $\tau_{m,0} = 13$ ns for the flat and two-step pulses, respectively. As the integration time increases, the SNR (assignment error) of both reference pulses increases (decreases) up until the transmon T_1 limit is reached (solid black line). Minimum assignment errors of 2.1×10^{-3} and

1.8×10^{-3} are obtained at 80 and 65 ns, respectively. This is similar performance to state-of-the-art readout experiments [42,43,51,67], as expected from our choice of realistic experimental parameters. Our objective is now to obtain smaller assignment errors in shorter measurement times.

The light blue symbols in Fig. 2(c) are obtained by optimizing the pulse envelope and drive frequency using our QOC approach. The gain is modest and mainly limited by the dispersive coupling with the transmon. Interestingly, the optimized pulses follow a two-step-like shape with a strong initial drive and a weaker subsequent drive; see Fig. 2(d). We attribute the small oscillations in the envelope to the rotational gauge freedom of the resonators, which the optimizer is arbitrarily choosing.

Significant improvements are, however, obtained by adding a drive on the transmon concurrently to the readout drive on the resonator. Interestingly, the optimizer converges on two distinct frequencies for the transmon drive. The first strategy found by the optimizer is to drive the transmon at a frequency close to the resonator frequency (green symbols). In that case, the assignment errors decrease faster with integration time than with the above approaches, leading to a minimal assignment error of 1.6×10^{-3} at 60 ns. The effectiveness of this optimized readout strategy stems from the fact that driving the qubit at the resonator frequency creates a longitudinal-like interaction that can be combined with the usual dispersive interaction to improve readout, as demonstrated in Refs. [68–70].

The second strategy found by the optimizer employs a transmon drive at the (ac-Stark shifted) $|e\rangle$ - $|f\rangle$ transition frequency (red symbols). Given that the cavity response differs more significantly between the transmon states $|g\rangle$ and $|f\rangle$ than between $|g\rangle$ and $|e\rangle$ [50], transferring population into the $|f\rangle$ state leads to a significant improvement of the assignment error, which reaches 1.0×10^{-3} in 40 ns. Interestingly, this shelving approach has already been used

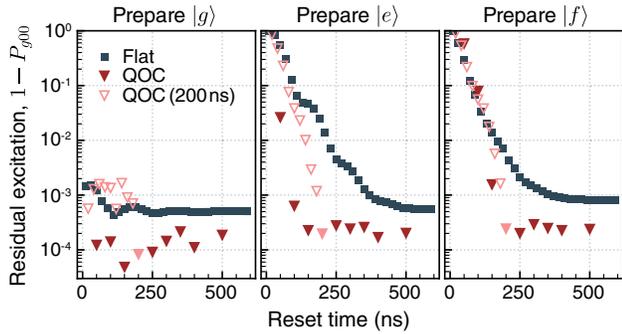


FIG. 3. Residual excitation out of $|g00\rangle$ after $f0$ - $g1$ reset for a calibrated pulse (blue) and an optimized pulse (red) for a system prepared in $|g00\rangle$, $|e00\rangle$, and $|f00\rangle$. Each filled red marker corresponds to a different optimization. Lighter hollow markers illustrate the population dynamics at shorter times for the optimized pulse with 200 ns duration.

to improve readout in circuit QED [71–74]. There, a π pulse between $|e\rangle$ and $|f\rangle$ is applied to the transmon followed by the measurement drive. In contrast, the optimized strategy found here applies the π pulse while the cavity is loaded with measurement photons leading to a considerable reduction in the measurement time; see Fig. 2(d). This is possible because the optimizer accounts for the time-dependent ac-Stark shift. The optimized π pulse features an envelope akin to Derivative Removal by Adiabatic Gate (DRAG) [75] and achieves a gate fidelity over 99% in less than 10 ns, even while the readout mode is being strongly driven. Importantly, we note that this approach could achieve significantly higher fidelities by increasing the modest transmon lifetime of 20 μ s used here, as shown by the high SNR in Fig. 2(c).

Transmon reset—As a second demonstration of the adjoint-state method, we consider the optimization of the $f0$ - $g1$ reset of a transmon [44,45,54]. This is an all-microwave reset protocol based on a Raman transition between states $|f00\rangle$ and $|g01\rangle$. For the ket $|ijk\rangle$, $|i\rangle$ stands for the qubit state and $|jk\rangle$ the resonator-filter normal modes. Given the large photon loss rate of the filter, the state $|g01\rangle$ quickly decays to $|g00\rangle$, thus ensuring a fast reset of the transmon $|f\rangle$ state. An additional drive at the $|e\rangle$ - $|f\rangle$ transition frequency allows us to reset both $|e\rangle$ and $|f\rangle$ states of the transmon. We use the adjoint-state method to find optimal controls for both the $f0$ - $g1$ and e - f drives simultaneously, in a similar fashion to optimizing the readout. The cost function is now principally maximizing the transmon population in the $|g00\rangle$ state at the end of the protocol, along with smaller contributions for regularizing the pulses; see Ref. [48] for details.

The results of the reset optimization are summarized in Fig. 3. The three panels show the residual excitation out of $|g00\rangle$ against the reset time for a reference flat pulse (blue) and an optimized pulse (red) for different initial transmon states. The reference pulse is composed of two constant drives at the $f0$ - $g1$ and e - f transitions, where amplitudes

and frequencies are calibrated numerically in a similar fashion to what is done in experiments; see Ref. [48]. The QOC pulse is obtained by optimizing the carrier frequency and envelopes of both drives, for several total reset times. The optimized pulses show significant improvement over the reference, with a residual excitation of less than 0.05% at 100 ns (200 ns) for the $|e\rangle$ ($|f\rangle$) state preparation. Note that this delay in the $|f\rangle$ reset time is due to a larger relative weight for the reset of $|e\rangle$ chosen in the cost function and could be adjusted to achieve the most experimentally relevant reset scheme. This represents a notable improvement over the reference pulses, which reach a steady state after more than 300 ns with larger residual excitations of about 0.07%. Our results also favorably compare to state-of-the-art experimental realizations of this protocol that reach 1.7% residual excitations in 100 ns [43], or 0.3% in 300 ns [45].

Conclusion—We obtained a fully general framework to optimize open quantum system dynamics in large Hilbert spaces by combining the adjoint-state method and reverse-time backpropagation. We have demonstrated the applicability of this method to complex open-system optimization problems using the example of superconducting transmon readout and reset. We stress that our method can readily be applied to optimizing a wide range of quantum control problems where the dissipative dynamics play a significant role such as reservoir (dissipation) engineering [76,77], autonomous QEC [78,79], leakage-reduction units [80], quantum cooling, and more. We encourage readers to apply this framework on their own optimal control problems using the open-source library DYNAMIQS [49].

Acknowledgments—We sincerely thank Alain Sarlette, Ross Shillito, Cristobal Lledó, Pierre Guilmin, and Adrien Bocquet for useful discussions. R. G. extends his gratitude to Alain Sarlette for helping organize and support this long-term academic exchange. This work was supported by Grant No. ANR-18-CE47-0005, NSERC, the Canada First Research Excellence Fund, and the Ministère de l’Économie et de l’Innovation du Québec. The numerical simulations were performed using HPC resources at Institut quantique and Inria Paris.

-
- [1] A. P. Peirce, M. A. Dahleh, and H. Rabitz, *Phys. Rev. A* **37**, 4950 (1988).
 - [2] J. Werschnik and E. Gross, *J. Phys. B* **40**, R175 (2007).
 - [3] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, *Eur. Phys. J. Quantum Technol.* **9**, 19 (2022).
 - [4] B. M. Terhal, *Rev. Mod. Phys.* **87**, 307 (2015).
 - [5] L. Viola, E. Knill, and S. Lloyd, *Phys. Rev. Lett.* **82**, 2417 (1999).
 - [6] K. Khodjasteh and D. A. Lidar, *Phys. Rev. Lett.* **95**, 180501 (2005).

- [7] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, *Nature (London)* **458**, 996 (2009).
- [8] R. Schmidt, A. Negretti, J. Ankerhold, T. Calarco, and J. T. Stockburger, *Phys. Rev. Lett.* **107**, 130404 (2011).
- [9] T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. Glaser, *J. Phys. B* **44**, 154013 (2011).
- [10] C. P. Koch, *J. Phys. Condens. Matter* **28**, 213001 (2016).
- [11] M. H. Goerz, D. M. Reich, and C. P. Koch, *New J. Phys.* **16**, 055012 (2014).
- [12] Z. An, H.-J. Song, Q.-K. He, and D. L. Zhou, *Phys. Rev. A* **103**, 012404 (2021).
- [13] T. Propson, B. E. Jackson, J. Koch, Z. Manchester, and D. I. Schuster, *Phys. Rev. Appl.* **17**, 014036 (2022).
- [14] D. J. Egger and F. K. Wilhelm, *Phys. Rev. A* **90**, 052331 (2014).
- [15] S. Boutin, C. K. Andersen, J. Venkatraman, A. J. Ferris, and A. Blais, *Phys. Rev. A* **96**, 042315 (2017).
- [16] D. Basilewitsch, F. Cosco, N. L. Gullo, M. Möttönen, T. Ala-Nissilä, C. P. Koch, and S. Maniscalco, *New J. Phys.* **21**, 093054 (2019).
- [17] H. Chen, H. Li, F. Motzoi, L. Martin, K. B. Whaley, and M. Sarovar, *New J. Phys.* **22**, 113014 (2020).
- [18] R. Porotti, V. Peano, and F. Marquardt, *PRX Quantum* **4**, 030305 (2023).
- [19] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret, *Phys. Rev. X* **12**, 011059 (2022).
- [20] V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. Brock, A. Ding, L. Frunzio *et al.*, *Nature (London)* **616**, 50 (2023).
- [21] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. R. Carvalho, and M. J. Biercuk, *PRX Quantum* **2**, 040324 (2021).
- [22] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, B. M. Niedzielski, A. Melville, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, J. A. Grover, K. Serniak, and W. D. Oliver, *Phys. Rev. X* **13**, 031035 (2023).
- [23] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *J. Magn. Reson.* **172**, 296 (2005).
- [24] V. Krotov, *Global Methods in Optimal Control Theory* (CRC Press, Boston, 1995), Vol. 195.
- [25] M. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich, and C. Koch, *SciPost Phys.* **7**, 080 (2019).
- [26] D. Basilewitsch, C. P. Koch, and D. M. Reich, *Adv. Quantum Technol.* **2**, 1800110 (2019).
- [27] H. Jirari, *Europhys. Lett.* **87**, 40003 (2009).
- [28] N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, *Phys. Rev. A* **95**, 042318 (2017).
- [29] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, *J. Mach. Learn. Res.* **18**, 1 (2018), <http://jmlr.org/papers/v18/17-468.html>.
- [30] M. Abdelhafez, D. I. Schuster, and J. Koch, *Phys. Rev. A* **99**, 052327 (2019).
- [31] M. H. Goerz, S. C. Carrasco, and V. S. Malinovsky, *Quantum* **6**, 871 (2022).
- [32] Y. Lu, S. Joshi, V. San Dinh, and J. Koch, *J. Phys. Comm.* **8**, 025002 (2024).
- [33] L. Pontryagin, V. Boltyanskii, R. Gamkrelidze, and E. Mishechenko, *The Mathematical Theory of Optimal Processes* (CRC Press, Boca Raton, 1962).
- [34] N. B. Dehaghani and A. P. Aguiar, [arXiv:2302.09143](https://arxiv.org/abs/2302.09143).
- [35] U. Boscain, M. Sigalotti, and D. Sugny, *PRX Quantum* **2**, 030203 (2021).
- [36] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, [arXiv:1806.07366](https://arxiv.org/abs/1806.07366).
- [37] P. Kidger, On neural differential equations, Ph.D. thesis, Mathematical Institute, University of Oxford, 2022, [arXiv:2202.02435](https://arxiv.org/abs/2202.02435).
- [38] J. Somló, V. A. Kazakov, and D. J. Tannor, *Chem. Phys.* **172**, 85 (1993).
- [39] S. H. K. Narayanan, T. Propson, M. Bongarti, J. Hüchelheim, and P. Hovland, in *International Conference on Computational Science* (Springer, New York, 2022), pp. 129–142.
- [40] H.-P. Breuer, F. Petruccione *et al.*, *The Theory of Open Quantum Systems* (Oxford University Press on Demand, 2002).
- [41] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, *Phys. Rev. A* **69**, 062320 (2004).
- [42] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, *Phys. Rev. Appl.* **7**, 054020 (2017).
- [43] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura, *Phys. Rev. Appl.* **17**, 044016 (2022).
- [44] M. Pechal, L. Huthmacher, C. Eichler, S. Zeytinoğlu, A. A. Abdumalikov, S. Berger, A. Wallraff, and S. Filipp, *Phys. Rev. X* **4**, 041010 (2014).
- [45] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff, *Phys. Rev. Lett.* **121**, 060502 (2018).
- [46] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Phys. Rev. A* **76**, 042319 (2007).
- [47] C. Gardiner and P. Zoller, *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics* (Springer Science & Business Media, New York, 2004).
- [48] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.134.070802> for more information, which includes Refs. [36,40,41,44,45,49–57].
- [49] P. Guilmin, R. Gautier, A. Bocquet, and E. Genois (to be published), library available at github.com/dynamiqs/dynamiqs.
- [50] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [51] F. Swiadek, R. Shillito, P. Magnard, A. Remm, C. Hellings, N. Lacroix, Q. Ficheux, D. C. Zanuz, G. J. Norris, A. Blais *et al.*, *PRX Quantum* **5**, 040326 (2024).
- [52] J. Gambetta, W. A. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, *Phys. Rev. A* **76**, 012325 (2007).
- [53] F. Motzoi, J. M. Gambetta, S. T. Merkel, and F. K. Wilhelm, *Phys. Rev. A* **84**, 022307 (2011).
- [54] S. Zeytinoğlu, M. Pechal, S. Berger, A. A. Abdumalikov, A. Wallraff, and S. Filipp, *Phys. Rev. A* **91**, 043846 (2015).
- [55] V. V. Albert, [arXiv:1802.00010](https://arxiv.org/abs/1802.00010).
- [56] J. R. Dormand and P. J. Prince, *J. Comput. Appl. Math.* **6**, 19 (1980).
- [57] C. Le Bris and P. Rouchon, *Phys. Rev. A* **87**, 022125 (2013).
- [58] C. G. Broyden, *IMA J. Appl. Math.* **6**, 76 (1970).

- [59] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [60] H. Robbins and S. Monro, *Ann. Math. Stat.* **22**, 400 (1951).
- [61] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimsheine, L. Antiga *et al.*, *Adv. Neural Inf. Process. Syst.* **32** (2019).
- [62] J. Cohen, A. Petrescu, R. Shillito, and A. Blais, *PRX Quantum* **4**, 020312 (2023).
- [63] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. M. Lewis, G. Vidal, and A. Blais, *Phys. Rev. Appl.* **18**, 034031 (2022).
- [64] M. F. Dumas, B. Groleau-Paré, A. McDonald, M. H. Muñoz-Arias, C. Lledó, B. D’Anjou, and A. Blais, *Phys. Rev. X* **14**, 041023 (2024).
- [65] C. A. Ryan, B. R. Johnson, J. M. Gambetta, J. M. Chow, M. P. da Silva, O. E. Dial, and T. A. Ohki, *Phys. Rev. A* **91**, 022118 (2015).
- [66] C. C. Bultink, B. Tarasinski, N. Haandbæk, S. Poletto, N. Haider, D. Michalak, A. Bruno, and L. DiCarlo, *Appl. Phys. Lett.* **112**, 092601 (2018).
- [67] Y. Sunada, K. Yuki, Z. Wang, T. Miyamura, J. Ilves, K. Matsuura, P. A. Spring, S. Tamate, S. Kono, and Y. Nakamura, *PRX Quantum* **5**, 010307 (2024).
- [68] J. Ikonen, J. Goetz, J. Ilves, A. Keränen, A. M. Gunyho, M. Partanen, K. Y. Tan, D. Hazra, L. Grönberg, V. Vesterinen, S. Simbierowicz, J. Hassel, and M. Möttönen, *Phys. Rev. Lett.* **122**, 080503 (2019).
- [69] S. Touzard, A. Kou, N. E. Frattini, V. V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. H. Devoret, *Phys. Rev. Lett.* **122**, 080502 (2019).
- [70] M. H. Muñoz Arias, C. Lledó, and A. Blais, *Phys. Rev. Appl.* **20**, 054013 (2023).
- [71] F. Mallet, F. R. Ong, A. Palacios-Laloy, F. Nguyen, P. Bertet, D. Vion, and D. Esteve, *Nat. Phys.* **5**, 791 (2009).
- [72] B. D’Anjou and W. A. Coish, *Phys. Rev. A* **96**, 052321 (2017).
- [73] S. S. Elder, C. S. Wang, P. Reinhold, C. T. Hann, K. S. Chou, B. J. Lester, S. Rosenblum, L. Frunzio, L. Jiang, and R. J. Schoelkopf, *Phys. Rev. X* **10**, 011001 (2020).
- [74] 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.*, *npj Quantum Inf.* **9**, 26 (2023).
- [75] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [76] J. F. Poyatos, J. I. Cirac, and P. Zoller, *Phys. Rev. Lett.* **77**, 4728 (1996).
- [77] P. M. Harrington, E. Mueller, and K. Murch, *Nat. Rev. Phys.* **4**, 660 (2022).
- [78] B. Royer, S. Singh, and S. M. Girvin, *Phys. Rev. Lett.* **125**, 260509 (2020).
- [79] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, *Nature (London)* **590**, 243 (2021).
- [80] P. Aliferis and B. M. Terhal, *Quantum Inf. Comput.* **7**, 139 (2007).