Iterative Linear Quadratic Regulator for Quantum Optimal Control

Dirk Heimann[1,2], Felix Wiebe[3], Tahereh Abad[4], Elie Mounzer[3], Tangyou Huang[4],
Frank Kirchner [1,3], and Shivesh Kumar[5]

[1] Robotics Research Group, University of Bremen
[2] Machine Learning Research Group, Bremen University of Applied Science
[3] Robotics Innovation Center, German Research Center for Artificial Intelligence, Bremen
[4] Department of Microtechnology and Nanoscience, Chalmers University of Technology
[5] Division of Dynamics, Department of Mechanics & Maritime Sciences, Chalmers University of Technology

Abstract—Quantum optimal control for gate optimization aims to provide accurate, robust, and fast pulse sequences to achieve gate fidelities on quantum systems below the error correction threshold. Many methods have been developed and successfully applied in simulation and on quantum hardware. In this paper, we establish a connection between the iterative linear quadratic regulator and quantum optimal control by adapting it to gate optimization of quantum systems. We include constraints on the controls and their derivatives to enable smoother pulses. We achieve high-fidelity simulation results for X and cross-resonance gates on one- and two-qubit fixed-frequency transmons simulated with two and three levels.

Index Terms—quantum optimal control, iterative linear quadratic regulator, gate synthesis for superconducting qubits

I. INTRODUCTION

Quantum optimal control (QOC) aims to design and apply precisely shaped pulse sequences of external fields to achieve high-fidelity control of quantum systems as efficiently as possible [1]. While possible applications for QOC range from quantum metrology [2] to quantum sensing [3], we focus on the field of gate optimization for quantum systems [4]. Quantum control algorithms build on the rich field of classical control theory, especially for trajectory optimization: Krotov's method [5] is named after its classical counterpart [6], PRONTO [7] is based on the Projection Operator-based Newton method [8], and PICO [9] is based on direct collocation [10]. Iterative linear quadratic regulator (iLQR) [11] is a value-function-based shooting method for trajectory optimization utilizing an approximation of the Hessian to determine the control updates. In recent years, it has become a state-of-the-art method in robotics, especially in online trajectory optimization for complex robots [12] because it obtains convergence guarantees by Levenberg-Marquardt regularization, is fast because of the Hessian approximation, handles nonlinear system dynamics via line-search, and obeys the system dynamics at any iteration of the optimization. In this paper, we adapt iLQR to the task of quantum gate optimization. Our contributions include:

- Adaptation of iLQR to QOC with smoothed controls
- Providing an in-depth description of the algorithm

- Simulation results for high-fidelity one-qubit and two-qubit gate synthesis for fixed-frequency transmons
- Comparison of our results to solutions obtained by L-BFGS-B GRAPE

This paper is organized as follows: First, in Sec. II, we provide a literature overview of QOC algorithms and recall the one-qubit and two-qubit fixed-frequency transmon Hamiltonian for superconducting systems. Then, in Sec. III, we introduce iLQR with the isomorphic representation and Padé approximation. In Sec. IV, we demonstrate the performance for gate optimization on the transmon Hamiltonians of increasing difficulty and conclude in Sec. V.

II. BACKGROUND

A. Quantum optimal control

State-of-the-art methods for gate optimization for superconducting transmon architectures include analytic and numerical methods. Prominent analytical methods are the Derivative Removal by Adiabatic Gate (DRAG) [13] for single-qubit operation, the cross-resonance (CR) gate [14], and its improved echo CR [15] two-qubit operation. Both have shown remarkable results for fixed-frequency transmon architectures because of the opportunity to fine-tune the pulse shape with system feedback data [16]. However, numerical optimization methods promise to offer more flexibility to adapt to more complex systems, specific control problems, experimental limitations, and uncertainties [4].

Many numerical open-loop control methods have been introduced in the last decades. In particular, two well-established gradient-based methods include gradient ascent pulse engineering (GRAPE) [17] and Krotov's method [5]. GRAPE updates all piecewise constant controls concurrently, i.e., based on the non-updated controls of the previous iteration step, using first-order [17] or second-order approximations by L-BFGS-B [18] of the loss gradients. The L-BFGS-B GRAPE implementation is available in the `qutip-qtrl` package [19] and has been tested on IBM superconducting systems [20]. In addition, Krotov's method [5], with its quasi-Newton extension [21], updates all controls sequentially, i.e., by including updated controls within the same iteration step.

Krotov's method is available for time-continuous [22] and time-discretized control functions [23]. However, the monotonic convergence of Krotov's method is guaranteed only for the time-continuous case [5]. Like Krotov's time-continuous version, PRONTO [7] solves the optimal control problem by sequentially solving ordinary differential equations and is guaranteed to be monotonically convergent. In contrast to previous Krotov's methods, PRONTO employs a Newton descent method to allow for quadratic convergence rates in proximity of local minima by incorporating system dynamics into a modified cost function using a projection operator. In addition to these shooting methods, direct collocation has recently been proposed [9] for quantum gate synthesis that includes both the states and controls at each discrete time step as a decision variable. Hence, in contrast to shooting methods, the number of optimization variables scales with the dimension of the Hilbert space. Additionally, the system dynamics of the trajectory are obeyed up to a certain factor by using them as a constraint during the optimization rather than a strict forward rollout. For a more comprehensive overview of recent progress in QOC, including machine learning and reinforcement learning approaches, we refer the reader to [4].

The iLQR algorithm utilizes first-order derivatives of the system dynamics in a Gauss-Newton approximation step to determine the Hessian for updating piecewise constant controls. Thereby linking it closely to differential dynamic programming (DDP) [24], which requires second-order derivatives for applying a full Newton step. However, the time advantage of the quadratic convergence often diminishes by the longer calculation time needed to determine second-order derivatives [25]. Although iLQR is concurrent in the sense that updates of piecewise constant controls are based on derivatives of the previous iteration step, a fundamental difference in comparison with GRAPE is that iLQR updates the controls based on the value function, which assigns a cost-to-go to each state and hence includes knowledge about future costs of the current control in the update.

B. Transmon Hamiltonian

Closed quantum systems evolve according to the Schrödinger equation. For time-dependent unitary operators $U(t) \in \mathbb{C}^{d \times d}$ and a Hamiltonian H depending on m real-valued controls $u(t) \in \mathbb{R}^m$ the Schrödinger equation reads $i\hbar \dot{U}(t) = H(u(t))U(t)$. We set the reduced Planck's constant $\hbar = 1$ in the following. The solution for N time-discretized controls $u_k \in \mathbb{R}^m$ and a constant time interval Δt is given by

$$U_N(u_{1:N}) = \prod_{k=1}^{N} \exp\left(-iH(u_k)\Delta t\right) U_0. \tag{1}$$

The discrete system dynamics are given by

$$U_{k+1} = f(U_k, u_k) = \exp(-iH(u_k)\Delta t)U_k. \tag{2}$$

Qubits can be physically realized by fixed-frequency transmons dispersively and capacitively coupled to a quantum resonator bus [26] that can be driven by microwave pulses [27].

A single transmon with transition frequency ω between eigenstates $|0\rangle$ and $|1\rangle$ and anharmonicity δ can be described by the Duffing oscillator

$$H_{\text{duffing}} = \omega b^\dagger b + \frac{\delta}{2} b^\dagger b(b^\dagger b - \mathbb{1}) \tag{3}$$

where $b^\dagger$ and b denote the creation and annihilation operators of the quantum harmonic oscillator. The resonator is described by a harmonic oscillator Hamiltonian $\omega_r c^\dagger c$ with transition frequency ω_r for all levels, and frequencies being far detuned, i.e., $\omega_r \gg \omega$.

Transmons can be coupled capacitively to the resonator by the Jaynes-Cummings Hamiltonian with coupling strength g_j in the dispersive regime, i.e. $g_j \ll |\Delta_{jr}|$ where $\Delta_{jr} = \omega_j - \omega_r$ is the detuning. For two transmons coupled to the resonator, the Hamiltonian reads

$$H_{\text{sys}} = \omega_r c^\dagger c + \sum_{j=1}^{2} H_{\text{duffing, j}} + \sum_{j=1}^{2} g_j(b_j^\dagger c + b_j c^\dagger) \tag{4}$$

The first resonator term and the six Duffing terms of the system's Hamiltonian are diagonal. Only the last two interaction terms are non-diagonal. Therefore, in the dispersive regime, the following effective two-transmon Hamiltonian can be derived via a Schrieffer-Wolf transformation and projection on the zero-excitation subspace of the bus [28]

$$H_{\text{sys}}^{\text{eff}} = \sum_{j=1}^{2} \left(\tilde{\omega}_j b_j^\dagger b_j + \frac{\delta_j}{2} b_j^\dagger b_j(b_j^\dagger b_j - \mathbb{1}) \right) + J_{12} \left(b_1^\dagger b_2 + b_1 b_2^\dagger \right) \tag{5}$$

with dressed qubit frequencies $\tilde{\omega}_j = \omega_j + \frac{g_j^2}{\Delta_{jr}}$ and the effective coupling to first-order $J_{12} = \frac{g_1 g_2}{2\Delta_{1r}\Delta_{2r}}(\Delta_{1r} + \Delta_{2r})$.

General microwave drive channels acting on the jth transmon with in-phase quadrature $u_j^X(t)$, off-phase quadrature $u_j^Y(t)$, drive frequency ω_d, and phase ϕ are modeled by the drive Hamiltonian

$$H_{\text{d}}(t) = r_j \left(u_j^X(t) \cos(\omega_d t + \phi) + u_j^Y(t) \sin(\omega_d t + \phi) \right) (b_j^\dagger + b_j) \tag{6}$$

where we include the Rabi strength factor r_j, which acts as a damping factor and is a property of the quantum system.

In the single-transmon case, the transmon is driven with two quadratures in its dressed transition frequency $\omega_d = \tilde{\omega}$ and the total system is described by $H_{\text{duffing}} + H_d(t)$. The transformation $R = \exp\left(-i(\omega_d t + \phi)b^\dagger b\right)$ shifts to the rotating frame,e which leads to the equation of one transmon driven by a microwave pulse with the rotating frame approximation (RWA)

$$H_{\text{1-t}}^{\text{RF}}(t) = \frac{\delta_1}{2} b_1^\dagger b_1(b_1^\dagger b_1 - \mathbb{1})$$
$$+ r_1 \frac{u^X(t)}{2}(b_1^\dagger + b_1) + r_1 \frac{u^Y(t)}{2} i(b_1^\dagger - b_1). \tag{7}$$

In the two-transmon case, a second drive can be added that acts on the second transmon $(b_2^\dagger + b_2)$ at the frequency of the first $\omega_d = \tilde{\omega}_1$. Driving one transmon (the control qubit) with the

frequency of the other (the target qubit) is referred to as cross resonance pulse [14]. Both drives have the same frequency, i.e., $\tilde{\omega}_1$, but act on different transmons. This enables a rotating frame transformation by $R = \exp\left(-i(\tilde{\omega}_1 t + \phi)\sum_{j=1}^2 b_j^\dagger b_j\right)$. With RWA and $\Delta_{21} = \tilde{\omega}_2 - \tilde{\omega}_1$, the two-transmon Hamiltonian is given by

$$
\begin{aligned}
H_{2\text{-}t}^{\text{RF}}(t) = &\Delta_{21} b_2^\dagger b_2 + \sum_{j=1}^2 \frac{\delta_j}{2} b_j^\dagger b_j (b_j^\dagger b_j - \mathbb{1}) + J_{12}\left(b_1^\dagger b_2 + b_1 b_2^\dagger\right) \\
&+ r_1 \frac{u^{X_1}(t)}{2}(b_1^\dagger + b_1) + r_1 \frac{u^{Y_1}(t)}{2}i(b_1^\dagger - b_1) \\
&+ r_1 \frac{u^{X_2}(t)}{2}(b_2^\dagger + b_2) + r_1 \frac{u^{Y_2}(t)}{2}i(b_2^\dagger - b_2).
\end{aligned}
\tag{8}
$$

C. Hardware parameters

On quantum systems, the arbitrary waveform generator (AWG) creates the amplitudes for the pulses as piecewise constant functions $u_k^{X_j}$ and $u_k^{Y_j}$ which are applied for the duration Δt for each timestep $k \in \{1, \ldots, N\}$. We utilize the system parameters of the IBM Eagle processor `ibm_brisbane` for qubit 0 and 1, which can be obtained via the open source library `Qiskit Pulse` [29]. Table I lists the rounded parameters in angular frequencies.

A common native gate set for fixed-frequency transmon architectures consists of single-qubit operations $\sqrt{X}$, X, and $R_Z(\theta)$ and an entangling two-qubit operation like cross resonance, echoed cross resonance, or CNOT gate. The $R_Z(\theta)$ gate can be implemented virtually by changing the phase of the AWG [30]. The single qubit gates $\sqrt{X}$, X and the two-qubit operation are implemented by AWG pulse envelopes. We focus on the X gate for the one-qubit operation because the method works analogously for the $\sqrt{X}$ gate. For the two-qubit operation, we consider the cross-resonance gate $R_{XZ}(\pi/2) = \exp(-i\pi XZ/4)$, which is equivalent to the CNOT gate up to a global phase by applying it together with one-qubit gates $CX_{q_2 q_1} = \mathbb{1} \otimes R_Z(-\pi/2) \cdot R_{XZ}(\pi/2) \cdot R_X(-\pi/2) \otimes \mathbb{1}$. Note that, in our drive channel setup, the second qubit is the control and the first qubit is the target.

III. METHOD

This section outlines the isomorphic representation, the Padé approximation for approximating the matrix exponential, the iLQR method, and its adaptations to obtain smoother control shapes.

TABLE I
FIXED-FREQUENCY TRANSMON PARAMETERS

qubit 1 dressed frequency	$\tilde{\omega}_1/(2\pi)$	4.7219 GHz
qubit 2 dressed frequency	$\tilde{\omega}_2/(2\pi)$	4.8151 GHz
qubit 1 anharmonicity	$\delta_1/(2\pi)$	-0.3120 GHz
qubit 2 anharmonicity	$\delta_2/(2\pi)$	-0.3097 GHz
effective coupling	$J_{12}/(2\pi)$	0.0020 GHz
qubit 1 rabi strength	$r_1/(2\pi)$	0.0921 GHz
qubit 2 rabi strength	$r_2/(2\pi)$	0.0974 GHz
minimal time interval	Δt	0.5 ns

A. Isomorphic representation

We follow [9] and vectorize unitary matrices by an isomorphic representation of complex vector spaces. A complex-valued vector $|\Psi\rangle \in \mathbb{C}^d$ and matrix $U \in \mathbb{C}^{d \times d}$ can be rewritten as the following real-valued vector and matrix

$$
|\tilde{\Psi}\rangle = \begin{pmatrix} \text{Re}\,\Psi \\ \text{Im}\,\Psi \end{pmatrix} \quad \text{and} \quad \tilde{U} = \begin{pmatrix} \text{Re}\,U & -\text{Im}\,U \\ \text{Im}\,U & \text{Re}\,U \end{pmatrix}.
\tag{9}
$$

The state space of complex-valued matrices can, therefore, be expressed by a $2d^2$ real-valued vector $x_k = \left(\text{Re}(U_{00}), \ldots, \text{Re}(U_{dd}), \text{Im}(U_{00}), \ldots, \text{Im}(U_{dd})\right)^T$ where x_k denotes the matrix values at timestep k. The discretized forward step of Eq. (2) based on controls u_k can also be expressed in the isomorphic representation

$$
\tilde{U}_{k+1} = f(\tilde{U}_k, u_k) = \exp\left(-i\tilde{H}(u_k)\Delta t\right)\tilde{U}_k.
\tag{10}
$$

B. Padé approximation

Calculating the matrix exponential in Eq. (2) is essential for the forward propagation. While (quasi-)analytic expressions are available for Lie groups $SU(2)$, $SU(3)$, and $SU(4)$ [31], we decide to use the Padé approximation for matrix exponentials as described in [9] because it can be readily used for higher levels. The Padé approximation is achieved by

$$
\exp(G) \approx B^{-1}(G)F(G)
\tag{11}
$$

where $B(G)$ and $F(G)$ are power series that can be chosen to a certain degree. For example, the eighth-order power series is given by

$$
B(G, \Delta t) = \mathbb{1} - \frac{\Delta_t}{2}G + \frac{3\Delta_t^2}{28}G^2 - \frac{\Delta_t^3}{84}G^3 + \frac{\Delta_t^4}{1680}G^4
\tag{12a}
$$

$$
F(G, \Delta t) = \mathbb{1} + \frac{\Delta_t}{2}G + \frac{3\Delta_t^2}{28}G^2 + \frac{\Delta_t^3}{84}G^3 + \frac{\Delta_t^4}{1680}G^4.
\tag{12b}
$$

C. Iterative Linear Quadratic Regulator

With the isomorphic representation of the forward step in Eq. (10), the Padé approximation in Eq. (11), and the vectorized form x_k of the unitary U_k, the QOC problem can be formulated as an iLQR problem. As a trajectory optimization method, iLQR encodes the search for a control trajectory $u_{1:N-1} = (u_1, \ldots, u_{N-1})$ with a consistent state trajectory $x_{1:N} = (x_1, \ldots, x_N)$ over a horizon N which fulfills constraints in an optimization problem of the shape

$$
\min_{u_{1:N-1}} J(x_{1:N}, u_{1:N-1}) = l_f(x_N) + \sum_{k=1}^{N-1} l(x_k, u_k)
\tag{13a}
$$

$$
\text{subject to:} \quad x_{k+1} = f(x_k, u_k) \quad \forall k \in \{1, \ldots, N-1\}
\tag{13b}
$$

$$
x_1 = \tilde{x}_1
\tag{13c}
$$

with the cost functional J, which usually can be split into a final cost function $l_n(x_N)$ and a running cost function $l(x_k, u_k)$. The constraints enforce the discrete system dynamics between the knot points and fix the initial state.

The value function V is defined as $V_k := V(x_k) = \min_{u_{k:N-1}} J(x_{k:N}, u_{k:N-1})$. Bellman's optimality principle

states that for optimal solutions every individual step has to be optimal, hence yields the recursive relation

$$V_k(x_k) = \min_{u_k} \left[l(x_k, u_k) + V_{k+1}(f(x_k, u_k)) \right]$$
$$:= \min_{u_k} Q_k(x_k, u_k) \tag{14}$$

with $V_N = l_f(x_N)$. The argument of the minimization is the action-value-function $Q(x_k, u_k)$. For small variations, $Q_k(x_k + \delta x_k, u_k + \delta u_k)$ can be approximated with a second-order Taylor expansion, which requires the following terms

$$Q_{x,k} = l_{x,k} + V_{x,k+1} f_{x,k} \tag{15a}$$
$$Q_{u,k} = l_{u,k} + V_{x,k+1} f_{u,k} \tag{15b}$$
$$Q_{xx,k} = l_{xx,k} + f_{x,k}^T V_{xx,k+1} f_{x,k} + V_{x,k+1} f_{xx,k} \tag{15c}$$
$$Q_{uu,k} = l_{uu,k} + f_{u,k}^T V_{xx,k+1} f_{u,k} + V_{x,k+1} f_{uu,k} \tag{15d}$$
$$Q_{ux,k} = l_{ux,k} + f_{u,k}^T V_{xx,k+1} f_{x,k} + V_{x,k+1} f_{ux,k} \tag{15e}$$
$$Q_{xu,k} = Q_{ux,k}^T \tag{15f}$$

where $Q_{x,k} = \partial_x Q|_{x_k, u_k}$, $Q_{xx,k} = \partial_x^2 Q|_{x_k, u_k}$, and accordingly for all other combinations. DDP utilizes the full second-order approximation, while iLQR ignores the second-order derivatives of the dynamics function. With this approximation, the Q-function can be minimized with respect to δu yielding the locally optimal update to the controls

$$\delta u_k^\star = \text{argmin}_{\delta u_k} Q(x_k + \delta x_k, u_k + \delta u_k)$$
$$= -Q_{uu,k}^{-1} Q_{u,k} - Q_{uu,k}^{-1} Q_{ux,k} \delta x_k$$
$$=: \kappa_k + K_k \delta x_k \tag{16}$$

The value function can be updated by $V_{x,k} = Q_{x,k} - Q_{u,k} Q_{uu,k}^{-1} Q_{ux,k}$ and $V_{xx,k} = Q_{xx,k} - Q_{xu,k} Q_{uu,k}^{-1} Q_{ux,k}$. Using a backwards recursion for V_k and V_{k+1} starting from the final condition $V_N = l_f(x_N)$ one can recursively compute updates to the value function and the feedback gains κ_k and K_k. This procedure is called the backward pass. The invertibility of Q_{uu} is ensured by including a Levenberg-Marquardt regularization scheme [32]. During the forward pass, the feedback gains are applied and a new control and state trajectory is calculated starting from the fixed start state $x_1^{\text{new}} = x^1$

$$u_k^{\text{new}} = u_k + \alpha \kappa_k + K_k(x_k^{\text{new}} - x_k) \tag{17}$$
$$x_{k+1}^{\text{new}} = f(x_k^{\text{new}}, u_k^{\text{new}}) \tag{18}$$

A line-search with the parameter α and a Goldstein acceptance criteria [32] prevent steps from going too far away from the reference point around which the system dynamics are linearly approximated. Backward pass and forward pass are alternated until convergence. While the derivatives in the RHS of Eqs. (15) must be calculated backward for each timestep k, these calculations can be performed in parallel.

D. Smooth controls

The original iLQR implementation, as presented in the previous section, does not include explicit constraints and, hence, optimizes over the entire, unconstrained control space. Because of limitations of the AWG abrupt changes in the controls should be avoided and starting and ending at zero is beneficial. To retrieve results with these properties, we choose the derivatives of the pulse envelopes as control variables and extend the state space with the actual pulse values. The derivative $\dot{u}_k$ changes the actual controls via $u_{k+1} = u_k + \dot{u}_k \Delta t$ for all k. With these modifications, we ensure that pulses start at zero by setting the initial condition and enable the penalization of the final pulse value, all intermediate pulse values, and the rate of change for the pulse in the cost function. Besides that, we do not impose further constraints on the controls. The total problem formulation in the structure of Eq. (13a) can be summarized as

$$\min_{\dot{u}_1, \ldots, \dot{u}_{N-1}} \quad \mathcal{L}_N(x_N, u_N) + \sum_{k=1}^{N-1} \mathcal{L}_k(x_k, u_k, \dot{u}_k) \tag{19a}$$
$$\text{subject to:} \quad \tilde{U}_{k+1} = e^{-i\tilde{H}(u_k)\Delta t} \tilde{U}_k \quad \forall k \tag{19b}$$
$$u_{k+1} = u_k + \dot{u}_k \Delta t \quad \forall k \tag{19c}$$
$$\tilde{U}_1 = \mathbb{1} \quad \text{and} \quad u_1 = 0. \tag{19d}$$

We use the quadratic costs

$$\mathcal{L}_k(x_k, u_k, \dot{u}_k) = \dot{u}_k^T R_d \dot{u}_k + u_k^T R_c u_k \tag{20}$$
$$\mathcal{L}_N(x_N, u_N) = (x_N - x_g)^T Q_f (x_N - x_g) + u_N^T R_f u_N \tag{21}$$

with the cost matrices $Q_f \in \mathbb{R}^{2d^2 \times 2d^2}$ and $R_d, R_c, R_f \in \mathbb{R}^{m \times m}$. R_d regularizes the rate of change for the pulses, R_c the pulse amplitudes and R_f the final control values u_N. Q_f penalizes differences between the final unitary $\tilde{U}_N = \prod_{k=1}^{N-1} \exp(-i\tilde{H}(u_k)\Delta t) \tilde{U}_1$ and the goal matrix $\tilde{U}_g$ computed in their vectorized form $x_N - x_g$.

IV. RESULTS

This section presents the simulation results of four gate optimization examples based on superconducting transmons restricted to two and three levels. We utilize the system parameters of the IBM_brisbane Eagle processor as summarized in table I and keep the minimal time interval $\Delta t = 0.5$ ns constant. In addition, we start our numerical experiments by selecting the number of timesteps slightly shorter than the gate times implemented on the IBM_brisbane system for their one- and two-qubit gates. We choose $N = 80$ time steps for our 1-qubit experiments and $N = 480$ for our 2-qubit experiments. We initialize all pulses with random values between $[-0.01, \ldots 0.01]$ and performed hyperparameter tuning of the cost matrices. As a first step, we manually find four diagonal values for the cost matrices Q_f, R_d, R_c, and R_f, which lead to good fidelities and relatively smooth pulses. For the one-qubit and three-level model, as well as for both two-qubit models, we perform a grid search by multiplying the four diagonal values by $\{1/10, 1/2, 1, 5, 10\}$ to search through 625 different parameter combinations. For the two-qubit model with three level, we select the best run of the previous gridsearch, fix R_c, and define a finer grid $\{1/10, 1/4, 1/2, 3/4, 1, 2.5, 5, 7.5, 10\}$ for the remaining diagonal values of the three cost matrices. The remaining section presents the best results regarding gate fidelity and the smoothness of the envelopes.
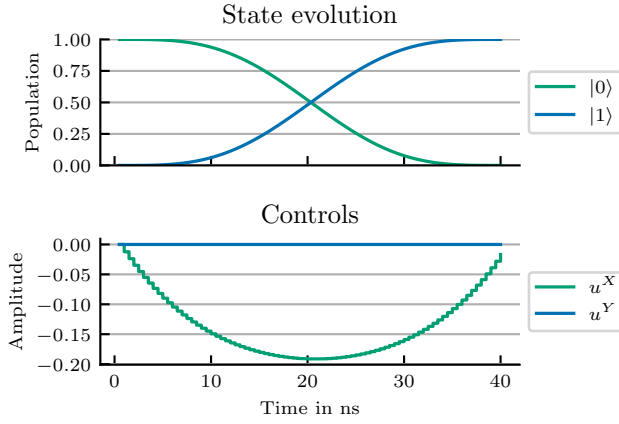
Fig. 1. Simulation results for one transmon qubit with two levels with gate infidelity of $4.0 \cdot 10^{-9}$. Population probability obtained by state evolution of the basis state $|0\rangle$ (**top**) under the optimal control sequence for the X-gate (**bottom**).
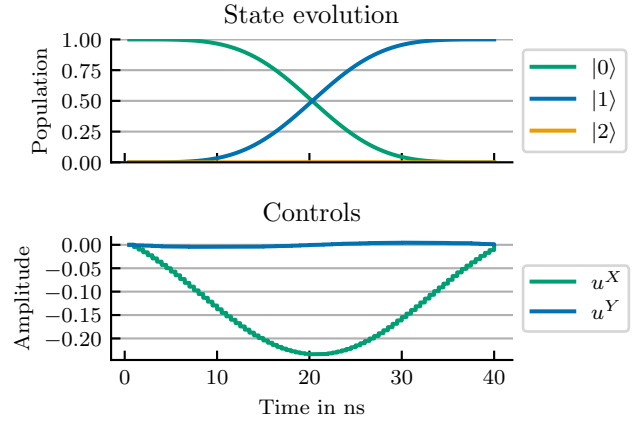

Fig. 2. Simulation results for one transmon qubit with three levels with gate infidelity of $2.1 \cdot 10^{-7}$. Population probability obtained by state evolution of basis state $|0\rangle$ (**top**) under the optimal control sequence for the X-gate (**bottom**).

A. One qubit with two levels

As a first example, we consider the two-level time-discretized Hamiltonian of Eq. (7) with piecewise constant controls. We choose $U_g = i\sigma_X \in SU(2)$ which equals the X-gate up to a global phase. This task is analytically solvable because only the $u^X \sigma_X$ operation is required [13]: if $U_g = i\sigma_X \overset{!}{=} U_N = \prod_{k=0}^{N-1} \exp\left(-i\frac{r_1}{2}u_k^X \Delta t \sigma_X\right)$ than $\sum_k u_k^X \Delta t = -\frac{\pi}{r_1}$ because of Euler's formula for $SU(2)$. Without including the derivatives of the controls, iLQR reproduces the analytic *bang-bang* solution with constant controls $u_k^X = -0.135722$ for all $N = 80$ time steps k leading to a simulation rollout infidelity of $1.3 \cdot 10^{-13}$. The sum over all controls matches the analytic factor up to $3 \cdot 10^{-7}$. However, rapid changes of the AWG's amplitude, as they can, for example, appear if the controls start and end with values different from zero, should be avoided in practice. That is why smoother controls are preferred in practice.

Fig. 1 shows the resulting controls and state propagation for basis states $|0\rangle$ and $|1\rangle$ obtained in simulation by iLQR including the derivatives of the controls for the smoothening effect. The simulation rollout infidelity is $4 \cdot 10^{-9}$ and the sum of the controls matches the factor up to $5.7 \cdot 10^{-5}$.

Similar results can be obtained with L-BFGS-B GRAPE. We modified the code slightly to allow different initial pulses for the controls. Starting with a Gaussian initial pulse type for u^X and zeros for u^Y leads to an almost ideal Gauss-like solution with the sum being equal to the factor up to $3.9 \cdot 10^{-7}$ and a simulation rollout infidelity of $1.9 \cdot 10^{-13}$. The Gaussian form, compared to the solution obtained by iLQR, has the advantage of smoother start and end regions.

B. One qubit with three levels

In the next example, we enable the possibility of leakage into a higher level by utilizing the time-discretized Hamilto-nian of Eq. (7) with piecewise constant controls and three-dimensional creation and annihilation operators. We choose

$$U_g = i\left(|1\rangle\langle 0| + |0\rangle\langle 1|\right) + |2\rangle\langle 2| \in SU(3) \qquad (22)$$

which acts on the basis states $|0\rangle$ and $|1\rangle$ as an ordinary X-gate up to a global phase. We use the derivative control to ensure a smoother rise and decline in the amplitudes. Fig. 2 shows the resulting controls and state propagation for basis states $|0\rangle$ obtained in simulation by iLQR. The infidelity is $2.1 \cdot 10^{-7}$ and the sum over all controls $\sum_k u_k^X \Delta t$ equals the factor up to $3.8 \cdot 10^{-4}$. The state propagation shows almost no leakage into the higher excitation state $|2\rangle$. The off-phase quadrature amplitude u^Y is roughly proportional to the derivative of the u^X amplitude, i.e. $u^Y = -\delta_1 \dot{u}^X$, which is, up to the factor, in agreement with the DRAG procedure [13]. In our iLQR setup, we force both in-phase and off-phase quadratures to go towards zero at the start and end. For many hyperparameters and system values, iLQR provides u^X solutions that are more parabola- than Gaussian-shaped. For parabola solutions for u^X, the derivatives at the start and end are highest, which means that u^Y, especially at the beginning and end, is not proportional to the derivative.

As a comparison, when initializing L-BFGS-B GRAPE with a Gaussian pulse shape for u^X and $u^Y = -\delta_1 \dot{u}^X$, we obtain solutions that mainly stay in their initial shapes with a gate infidelity of $1.6 \cdot 10^{-10}$.

C. Two qubits with two levels

For the two-qubit gate experiments, we start by restricting the Hamiltonian in Eq. (8) to two levels and choose to optimize for the cross-resonance gate

$$U_g = \exp\left(-i\frac{\pi}{4}\sigma_X \otimes \sigma_Z\right) \in SU(4) \qquad (23)$$

where the first qubit is the target qubit and the second is the control qubit. Fig. 3 shows the resulting controls and state propagation for basis state $|00\rangle$ obtained in simulation by
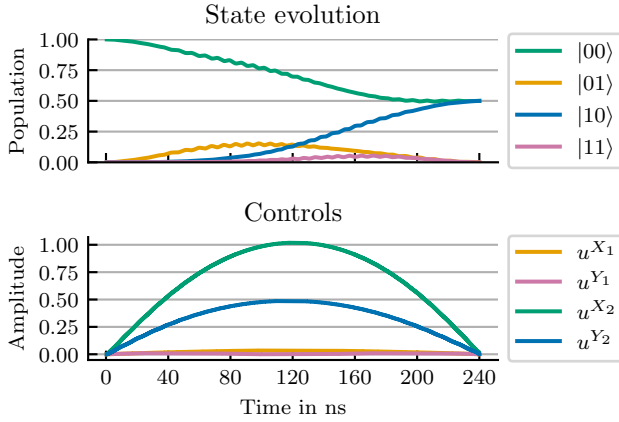
Fig. 3. Simulation results for two transmon qubits with two levels with gate infidelity of $1.1 \cdot 10^{-8}$. Population probability obtained by state evolution of basis state $|00\rangle$ (**top**) under the optimal control sequence for the 240 ns XZ-gate (**bottom**).
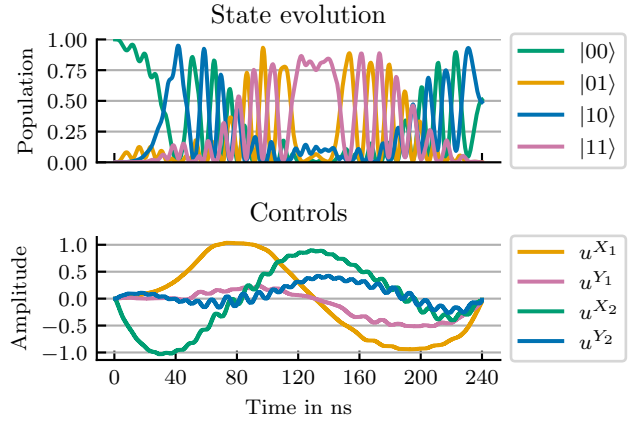

Fig. 4. Simulation results for two transmon qubits with three levels with gate infidelity of $5.9 \cdot 10^{-5}$. Population probability obtained by state evolution of basis state $|00\rangle$ (**top**) under the optimal control sequence for the 240 ns XZ-gate (**bottom**). In the top plot, only the 2-level states are depicted because the contribution of the third level was much smaller with an average probability of $5.8 \cdot 10^{-3}$ and a max value of $5.8 \cdot 10^{-2}$ of all five higher-level basis states.

iLQR. The infidelity is $1.1 \cdot 10^{-8}$ with a gate duration of 240 ns. Four configurations in the grid search lead to lower infidelities than 10^{-7}. They have slightly less smooth controls with varying amplitudes for all four controls such that the state's population probability evolves more periodically until it reaches its goal.

D. Two qubits with three levels

As our last example, we consider the two-transmon time-discretized Hamiltonian of Eq. (8) restricted to three levels. We choose to optimize the following extended CR gate

$$U_g = \sum_{j \in \{0,1,3,4\}} \frac{1}{\sqrt{2}} |j\rangle \langle j| + \sum_{j \in \{2,5,6,7,8\}} |j\rangle \langle j|$$
$$+ \frac{-i}{\sqrt{2}} (|3\rangle \langle 0| + |0\rangle \langle 3|) + \frac{i}{\sqrt{2}} (|4\rangle \langle 1| + |1\rangle \langle 4|) \quad (24)$$

which is an element of $SU(9)$. Fig. 4 shows the resulting controls and state propagation for basis state $|00\rangle$ obtained in simulation by iLQR. The infidelity is $5.9 \cdot 10^{-5}$ with a gate duration of 240 ns. The finer grid search includes three runs with infidelities better than 10^{-4}. None of them have smoother controls or more direct state population probabilities.

V. Conclusion

In this work, we demonstrated that iLQR can be used in quantum optimal control, especially for superconducting gate synthesis. We have focused on the single- and two-qubit operations, X and CR gates, among the basis gates for fixed-frequency transmon qubits. The algorithm optimizes the derivatives indirectly influencing the controls to ensure smoother piecewise constant controls. Additionally, this allows us to constrain the start and end controls closer to zero. In particular, the optimizer finds smooth solutions with high gate fidelities for the X gate 40 ns pulses, and the CR gate 240 ns pulse controls. The design of the quadratic cost has had more impact on the results than the initial pulse shapes. Hence, our

iLQR setup leads to high-fidelity solutions without requiring knowledge about an approximate solution. In contrast, the shape of the L-BFGS-B GRAPE solution depends heavily on the input pulse. While L-BFGS-B GRAPE finds solutions with even higher fidelities for a wide range of initial pulses, for reasonable initial pulses, the solutions tend to stay close to the initial pulse shape. Hence, we see the potential to use iLQR for solving QOC problems where no analytical solution is known a priori.

The iLQR gates are still relatively long, so optimizing for shorter pulses under realistic AWG constraints is a crucial next step. This might be made possible by using more sophisticated cost functions, for example, time-dependent running control costs for the controls R_c and their derivatives R_d to enforce different pulse envelopes. Additional work is needed to incorporate the fidelity directly in the cost function, which could further improve the solution's fidelity. Moreover, our proposed method can be applied to different hardware system architectures. One example of another superconducting system is fixed-frequency transmons coupled to a tunable transmon, which naturally allows for iSWAP or CZ gates [33] as their fundamental two-qubit operation. Open-loop numerical methods like the one presented in this work rely on accurate model and control signal descriptions, which are especially difficult to satisfy for superconducting transmon quantum systems. One approach to deal with the model's inaccuracy is to incorporate hardware feedback [34]. As a next step, we will investigate how iLQR can integrate hardware feedback to obtain pulses specialized to specific quantum systems.

Acknowledgment

References

[1] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, *et al.*, "Training schrödinger's cat: Quantum optimal control: Strategic report on current status, visions and goals for research in europe," *The European Physical Journal D*, vol. 69, pp. 1–24, 2015. [Online]. Available: https://doi.org/10.1140/epjd/e2015-60464-1

[2] J. Liu, M. Zhang, H. Chen, L. Wang, and H. Yuan, "Optimal scheme for quantum metrology," *Advanced Quantum Technologies*, vol. 5, no. 1, p. 2100080, 2022. [Online]. Available: https://doi.org/10.1002/qute.202100080

[3] P. Rembold, N. Oshnik, M. M. Müller, S. Montangero, T. Calarco, *et al.*, "Introduction to quantum optimal control for quantum sensing with nitrogen-vacancy centers in diamond," *AVS Quantum Science*, vol. 2, no. 2, p. 024701, 06 2020. [Online]. Available: https://doi.org/10.1116/5.0006785

[4] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, *et al.*, "Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in europe," *EPJ Quantum Technology*, vol. 9, no. 1, p. 19, 2022. [Online]. Available: https://doi.org/10.1140/epjqt/s40507-022-00138-x

[5] D. M. Reich, M. Ndong, and C. P. Koch, "Monotonically convergent optimization in quantum control using Krotov's method," *The Journal of Chemical Physics*, vol. 136, no. 10, p. 104103, 03 2012. [Online]. Available: https://doi.org/10.1063/1.3691827

[6] V. F. Krotov, *Global Methods in Optimal Control Theory*. Boston, MA: Birkhäuser Boston, 1993, pp. 74–121. [Online]. Available: https://doi.org/10.1007/978-1-4612-0349-0_3

[7] J. Shao, J. Combes, J. Hauser, and M. M. Nicotra, "Projection-operator-based newton method for the trajectory optimization of closed quantum systems," *Phys. Rev. A*, vol. 105, p. 032605, 03 2022. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.105.032605

[8] J. Hauser, "A projection operator approach to the optimization of trajectory functionals," *IFAC Proceedings Volumes*, vol. 35, no. 1, pp. 377–382, 2002, 15th IFAC World Congress. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667015387334

[9] A. Trowbridge, A. Bhardwaj, K. He, D. I. Schuster, and Z. Manchester, "Direct collocation for quantum optimal control," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Los Alamitos, CA, USA: IEEE Computer Society, 09 2023, pp. 1278–1285. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/QCE57702.2023.00144

[10] C. R. Hargraves and S. W. Paris, "Direct trajectory optimization using nonlinear programming and collocation," *AIAA J. Guidance*, vol. 10, pp. 338–342, 1987. [Online]. Available: https://doi.org/10.2514/3.20223

[11] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *First International Conference on Informatics in Control, Automation and Robotics*, vol. 2. SciTePress, 2004, pp. 222–229. [Online]. Available: https://doi.org/10.5220/0001143902220229

[12] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, *et al.*, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2024. [Online]. Available: https://doi.org/10.1109/TRO.2023.3324580

[13] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, "Simple pulses for elimination of leakage in weakly nonlinear qubits," *Phys. Rev. Lett.*, vol. 103, p. 110501, 9 2009. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.103.110501

[14] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, *et al.*, "Simple all-microwave entangling gate for fixed-frequency superconducting qubits," *Phys. Rev. Lett.*, vol. 107, p. 080502, 8 2011. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.107.080502

[15] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, "Procedure for systematically tuning up cross-talk in the cross-resonance gate," *Phys. Rev. A*, vol. 93, p. 060302, 6 2016. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.93.060302

[16] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, *et al.*, "Demonstration of quantum volume 64 on a superconducting quantum computing system," *Quantum Science and Technology*, vol. 6, no. 2, p. 025020, 3 2021. [Online]. Available: https://dx.doi.org/10.1088/2058-9565/abe519

[17] 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," *Journal of magnetic resonance*, vol. 172, no. 2, pp. 296–305, 2005. [Online]. Available: https://doi.org/10.1016/j.jmr.2004.11.004

[18] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, *et al.*, "Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework," *Phys. Rev. A*, vol. 84, p. 022305, 8 2011. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.84.022305

[19] J. Johansson, P. Nation, and F. Nori, "Qutip 2: A python framework for the dynamics of open quantum systems," *Computer Physics Communications*, vol. 184, no. 4, pp. 1234–1240, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0010465512003955

[20] E. S. Matekole, Y.-L. L. Fang, and M. Lin, "Methods and results for quantum optimal pulse control on superconducting qubit systems," in *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2022, pp. 600–606. [Online]. Available: https://doi.org/10.1109/IPDPSW55747.2022.00102

[21] R. Eitan, M. Mundt, and D. J. Tannor, "Optimal control with accelerated convergence: Combining the krotov and quasi-newton methods," *Phys. Rev. A*, vol. 83, p. 053426, 05 2011. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.83.053426

[22] D. J. Tannor, V. Kazakov, and V. Orlov, *Control of Photochemical Branching: Novel Procedures for Finding Optimal Pulses and Global Upper Bounds*. Boston, MA: Springer US, 1992, pp. 347–360. [Online]. Available: https://doi.org/10.1007/978-1-4899-2326-4_24

[23] M. H. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, *et al.*, "Krotov: A python implementation of krotov's method for quantum optimal control," *SciPost Phys.*, vol. 7, p. 080, 2019. [Online]. Available: https://scipost.org/10.21468/SciPostPhys.7.6.080

[24] D. Mayne, "A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems," *International Journal of Control*, vol. 3, no. 1, pp. 85–95, 1966. [Online]. Available: https://doi.org/10.1080/00207176608921369

[25] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1168–1175.

[26] J. Majer, J. Chow, J. Gambetta, J. Koch, B. Johnson, *et al.*, "Coupling superconducting qubits via a cavity bus," *Nature*, vol. 449, no. 7161, pp. 443–447, 2007. [Online]. Available: https://doi.org/10.1038/nature06184

[27] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, *et al.*, "A quantum engineer's guide to superconducting qubits," *Applied physics reviews*, vol. 6, no. 2, 2019. [Online]. Available: https://doi.org/10.1063/1.5089550

[28] E. Magesan and J. M. Gambetta, "Effective hamiltonian models of the cross-resonance gate," *Phys. Rev. A*, vol. 101, p. 052308, 5 2020. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.101.052308

[29] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, *et al.*, "Qiskit pulse: programming quantum computers through the cloud with pulses," *Quantum Science and Technology*, vol. 5, no. 4, p. 044006, 2020. [Online]. Available: https://doi.org/10.1088/2058-9565/aba404

[30] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient z gates for quantum computing," *Phys. Rev. A*, vol. 96, p. 022330, 08 2017. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.96.022330

[31] N. Kaiser, "Solving the matrix exponential function for the lie groups su(3), su(4) and sp(2)," *The European Physical Journal A*, vol. 58, no. 170, 09 2022. [Online]. Available: https://doi.org/10.1140/epja/s10050-022-00816-5

[32] C. Mastalli, W. Merkt, J. Marti-Saumell, H. Ferrolho, J. Solà, *et al.*, "A feasibility-driven approach to control-limited ddp," *Autonomous*

Robots, vol. 46, no. 8, pp. 985–1005, 2022. [Online]. Available: https://doi.org/10.1007/s10514-022-10061-w

[33] J. Stehlik, D. M. Zajac, D. L. Underwood, T. Phung, J. Blair, *et al.*, "Tunable coupling architecture for fixed-frequency transmon superconducting qubits," *Phys. Rev. Lett.*, vol. 127, p. 080505, 08 2021. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.127.080505

[34] R.-B. Wu, B. Chu, D. H. Owens, and H. Rabitz, "Data-driven gradient algorithm for high-precision quantum control," *Phys. Rev. A*, vol. 97, p. 042122, 04 2018. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.97.042122