

A Projection Operator-based Newton Method for the Trajectory Optimization of Closed Quantum Systems

Jieqiu Shao,¹ Joshua Combes,¹ John Hauser,¹ and Marco M. Nicotra¹

¹*Department of Electrical, Computer and Energy Engineering,
University of Colorado Boulder, Boulder, Colorado 80309, USA*

(Dated: October 11, 2022)

Quantum optimal control is an important technology that enables fast state preparation and gate design. In the absence of an analytic solution, most quantum optimal control methods rely on an iterative scheme to update the solution estimate. At present, the convergence rate of existing solvers is, at most, superlinear. This paper develops a new general purpose solver for quantum optimal control based on the PRojection Operator Newton method for Trajectory Optimization, or PRONTO. Specifically, the proposed approach uses a projection operator to incorporate the Schrödinger equation directly into the cost function, which is then minimized using a Newton descent method. At each iteration, the descent direction is obtained by computing the analytic solution to a linear-quadratic trajectory optimization problem. The resulting method guarantees monotonic convergence at every iteration and quadratic convergence in proximity of the solution. The potential of PRONTO is showcased by solving the optimal state-to-state mapping problem for a qubit and providing comparisons to a state-of-the-art quantum optimal control method.

I. INTRODUCTION

To accomplish the promise of quantum computing and quantum sensing, it is necessary to reliably and accurately control increasingly large quantum systems. Quantum control [1–4] is difficult for many reasons, including: interesting quantum systems are fundamentally nonlinear [5], observations disturb the state of the system being controlled, and, like classical systems, the quantum state space suffers from the curse of dimensionality.

Broadly speaking, the field of quantum optimal control can be divided into model-free, e.g. [6], and model-based methods. In the context of model-based quantum optimal control, the two predominant strategies in modern literature are the GRadient Ascent Pulse Engineering algorithm (GRAPE) [7, 8], which treats the control input as a sequence of piecewise-constant pulses, and Krotov methods [9, 10], which treat the control input as a continuous function. For a detailed comparison between the two strategies, readers are referred to [11]. Because these prior methods use gradient-based descent, their convergence to the optimal solution is predominantly linear (or superlinear in the case of quasi-Newton gradient acceleration schemes [12, 13]). The development of quadratically convergent quantum control methods is an open problem and may lead to the ability to control larger quantum systems.

This paper introduces a new strategy for model-based quantum optimal control by specializing the PRojection Operator-based Newton method for Trajectory Optimization (PRONTO) [14] to quantum systems. An example featuring a prior application of PRONTO to quantum control is discussed briefly in [15].

PRONTO is conceptually similar to Krotov, in the sense that they both solve the optimal control problem directly in function space, and they both do so by sequentially solving backward-in-time and forward-in-time

ordinary differential equations. Their difference lies in how they handle the system dynamics: Krotov enforces them using Lagrange multipliers and employs a primal-dual gradient method to seek the saddle point of the Lagrangian; PRONTO embeds them into a modified cost function using the projection operator, and employs a Newton descent method to seek the minimum of the modified cost function. Thus, the convergence rate of Krotov is linear, whereas the convergence rate of PRONTO can be quadratic.

The paper is structured as follows: Section II states the quantum optimal control problem we wish to solve and reformulates it into a traditional optimal control problem. Section III reviews established results from numerical optimization theory to familiarize the reader with the general concepts used in PRONTO. Section IV provides a detailed description of how to implement PRONTO on quantum systems and includes a pseudo-code summarizing all the steps performed by the method. Finally, Section V provides preliminary comparisons between PRONTO and the state-of-the-art Krotov implementation package [16].

II. PROBLEM STATEMENT

The objective of this paper is to develop a systematic approach for the state-to-state control of quantum systems, meaning that we wish to steer some initial state $|\psi_0\rangle$ to a target state $|\phi\rangle$ over a finite time interval $[0, T]$. In this section, we show how the state-to-state quantum control problem can be reformulated as a classic optimal control problem, which we will then solve using a specialized version of the PRONTO algorithm.

A. Dynamic Model

We consider a system governed by the Hamiltonian

$$\hat{H}[u(t)] = \hat{H}_0 + \sum_{j=1}^m \hat{H}_j f_j[u_j(t)], \quad (1)$$

where $\hat{H}_0 \in \mathbb{C}^{n \times n}$ represents the free evolution of the system, $\hat{H}_j \in \mathbb{C}^{n \times n}$ is the j -th control Hamiltonian and the associated $f_j : \mathbb{R} \rightarrow \mathbb{R}$ is a class \mathcal{C}^2 function of the control input $u_j(t) \in \mathbb{R}$. The dynamics generated by this Hamiltonian are

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H}[u(t)] |\psi(t)\rangle, \quad (2)$$

with $|\psi(t)\rangle \in \mathbb{C}^n$ and $u(t) \in \mathbb{R}^m$. Henceforth, we set $\hbar = 1$ for simplicity.

B. Cost Function

To quantify the effectiveness of a given controller, we will rank the behavior of $(|\psi(t)\rangle, u(t))$ within the time window $t \in [0, T]$, by evaluating the cost function

$$\hat{m}[|\psi(T)\rangle] + \int_0^T \hat{l}[|\psi(t)\rangle, u(t)] dt, \quad (3)$$

where both the *terminal cost* $\hat{m} : \mathbb{C}^n \rightarrow \mathbb{R}_{\geq 0}$ and the *incremental cost* $\hat{l} : (\mathbb{C}^n \times \mathbb{R}^m) \rightarrow \mathbb{R}_{\geq 0}$ are class \mathcal{C}^2 convex functions that can be interpreted as follows:

- The incremental cost penalizes undesirable behaviors within the time window $t \in [0, T]$. A common choice is

$$\hat{l}(|\psi\rangle, u) = \hat{l}_\psi(|\psi\rangle) + l_u(u), \quad (4)$$

with \hat{l}_ψ convex and l_u strongly convex. The latter prevents $u(t)$ from growing unbounded, and is therefore necessary for the well-posedness of the problem. A typical choice for $l_u \neq 0$ is

$$l_u(u) = \frac{1}{2} u^T R(t) u. \quad (5)$$

where $R(t) > 0$, $\forall t \in [0, T]$ is a (potentially time-varying) symmetric matrix. As for the former, $\hat{l}_\psi \neq 0$ can be useful for penalizing the transfer of population onto undesirable states $|\lambda\rangle$, as detailed in [17]. In this case, a suitable choice would be

$$\hat{l}_\psi(|\psi\rangle) = \frac{1}{2} \langle \psi | \hat{P}_\lambda | \psi \rangle, \quad (6)$$

where the operator $\hat{P}_\lambda = |\lambda\rangle\langle\lambda|$ is a projection of $|\psi\rangle$ onto a state $|\lambda\rangle$. If we do not wish to penalize any states during the transient, it is worth noting that $\hat{l}_\psi = 0$ is a suitable, and fairly common [16], alternative.

- The terminal cost drives the final state $|\psi(T)\rangle$ to a desirable target $|\phi\rangle$ by penalizing the deviation between

$|\psi(T)\rangle$ and $|\phi\rangle$. For example, we can perform an arbitrary phase state-to-state transition by assigning

$$\hat{m}(|\psi\rangle) = \langle \psi | \hat{P}_{\neg\phi} | \psi \rangle, \quad (7)$$

where the operator $\hat{P}_{\neg\phi} = 1 - |\phi\rangle\langle\phi|$ is the complement of \hat{P}_ϕ and \hat{m} denotes the squared Hilbert-Schmidt distance between $|\psi\rangle$ and the target state $|\phi\rangle$.

C. Optimal Control Problem

Having identified the system dynamics and a suitable cost function, we can now formulate the well-known [1] quantum optimal control problem

$$\min \hat{m}[|\psi(T)\rangle] + \int_0^T \hat{l}[|\psi(t)\rangle, u(t)] dt \quad (8a)$$

$$\text{s.t. } \frac{\partial}{\partial t} |\psi(t)\rangle = -i\hat{H}[u(t)]|\psi(t)\rangle, \quad |\psi(0)\rangle = |\psi_0\rangle. \quad (8b)$$

To solve this problem, we will start by re-framing it into the form found in conventional control literature [18]. This is done by transforming the complex state vector $|\psi(t)\rangle \in \mathbb{C}^n$ into a larger real vector $x(t) \in \mathbb{R}^{2n}$ by taking advantage of the bijective mapping

$$x(t) = \begin{bmatrix} \text{Re}[|\psi(t)\rangle] \\ \text{Im}[|\psi(t)\rangle] \end{bmatrix}. \quad (9)$$

This mapping, detailed in [19], allows us to transform any operator $\hat{Y} \in \mathbb{C}^{n \times n}$ into a matrix $Y \in \mathbb{R}^{2n \times 2n}$ using

$$Y = \begin{bmatrix} \text{Re}(\hat{Y}) & -\text{Im}(\hat{Y}) \\ \text{Im}(\hat{Y}) & \text{Re}(\hat{Y}) \end{bmatrix}. \quad (10)$$

The Schrödinger equation (2) can therefore be rewritten as the real-valued ordinary differential equation

$$\dot{x}(t) = H[u(t)] x(t) \quad (11)$$

by defining

$$H = \begin{bmatrix} \text{Re}(-i\hat{H}) & -\text{Im}(-i\hat{H}) \\ \text{Im}(-i\hat{H}) & \text{Re}(-i\hat{H}) \end{bmatrix}. \quad (12)$$

Likewise, the bijective mapping (9) can be leveraged to transform the functions $\hat{l}(|\psi\rangle, u)$ and $\hat{m}(|\psi\rangle)$ into their equivalent form $l(x, u)$ and $m(x)$. For example, the cost functions (4)-(7) can be rewritten as

$$l(x, u) = \frac{1}{2} x^T P_\lambda x + \frac{1}{2} u^T R(t) u, \quad (13a)$$

$$m(x) = \frac{1}{2} x^T P_{\neg\phi} x, \quad (13b)$$

where the matrices $P_{\neg\phi}$ and P_λ can be obtained from the operators $\hat{P}_{\neg\phi}$ and \hat{P}_λ using Eq. (10).

Solving the quantum optimal control problem (8) is therefore equivalent to solving the conventional optimal control problem

$$\min m[x(T)] + \int_0^T l[x(t), u(t)] dt \quad (14a)$$

$$\text{s.t. } \dot{x}(t) = H[u(t)] x(t), \quad x(0) = x_0. \quad (14b)$$

Compared to the more general formulation featuring $\dot{x} = f(x, u)$, (14) features two very interesting properties that stem from the nature of closed quantum systems

- The differential equation (14b) is *affine* in the state vector x , meaning that its second derivative satisfies $\nabla_{xx}^2[H(u)x] = 0$;
- The matrix $H(u)$ is *skew-symmetric* for all $u \in \mathbb{R}^m$. This implies $\|x(t)\| = \|x(0)\|$, $\forall t \in [0, T]$.

These properties, plus the general simplicity of computing the partial derivatives of (14) will be leveraged in Sec. IV to design an efficient iterative solver.

III. OPTIMIZATION PRIMER

The goal of this section is to provide the interested reader with an intuition for the main ingredients used in PRONTO. Although most of the information contained in this section can be found in a good textbook on numerical optimization, e.g. [20], the order in which it is presented and the way it is interpreted is both original to this paper and critical for understanding the theory behind PRONTO.

The starting point is Eq. (14), which is a constrained optimization problem defined in *function space*, meaning that its solution $[x(t), u(t)]$ is a pair of functions $x : [0, T] \rightarrow \mathbb{R}^{2n}$ and $u : [0, T] \rightarrow \mathbb{R}^m$. Since readers are more likely familiar with *vector space* optimization, this section describes the general intuition behind the approach used in Sec. IV for the simplified case where $x \in \mathbb{R}^{2n}$ and $u \in \mathbb{R}^m$ are just two vectors.

A. Embedding Constraints into the Cost Function

Given a constrained optimization problem in the form

$$\min h(x, u) \quad (15a)$$

$$\text{s.t. } c(x, u) = 0, \quad (15b)$$

where $h : (\mathbb{R}^{2n} \times \mathbb{R}^m) \rightarrow \mathbb{R}_{\geq 0}$ is a \mathcal{C}^2 convex function and $c : (\mathbb{R}^{2n} \times \mathbb{R}^m) \rightarrow \mathbb{R}$ is a \mathcal{C}^2 function, a systematic method for finding the solution is to search for the saddle point of the Lagrangian $\mathcal{L}(x, u, \chi) = h(x, u) + \chi^T[c(x, u)]$, where $\chi \in \mathbb{R}^{2n}$ is the vector of Lagrange multipliers. This can be done using a primal-dual gradient method (incidentally, if one were to transfer this intuition back into function space, one would obtain something akin to the Krotov method [21]).

Given a \mathcal{C}^2 function $p : \mathbb{R}^m \rightarrow (\mathbb{R}^{2n}, \mathbb{R}^m)$ such that $c(x, u) = 0$ iff $[x, u] = p(u)$, an alternative method would be to solve the unconstrained optimization problem

$$\min g(u), \quad (16)$$

with $g(u) = h[p(u)]$. Clearly, the main challenge associated to this transformation is finding a suitable function

$p(u)$. Doing so, however, enables us to use a standard Newton method for finding the solution.

B. Newton Method

The classic Newton method seeks local minima of (16) by applying an iterative procedure in the form

$$u_{k+1} = u_k + \nu_k, \quad (17)$$

where, given a guess u_k , the **Descent Direction** ν_k is obtained by computing the local quadratic approximation of the cost update, i.e.

$$g(u) \approx g(u_k) + \nabla^T g(u_k)\nu + \frac{1}{2}\nu^T \nabla^2 g(u_k)\nu, \quad (18)$$

and solving a quadratic minimization problem

$$\nu_k = \arg \min \nabla^T g(u_k)\nu + \frac{1}{2}\nu^T \nabla^2 g(u_k)\nu. \quad (19)$$

The interest in this approach is that, unlike the original problem (16), the local quadratic approximation (19) can be solved explicitly iff $\nabla^2 g(u_k) > 0$. Notably, the solution to (19) is

$$\nu_k = -\nabla^2 g(u_k)^{-1} \nabla g(u_k). \quad (20)$$

When a local minimizer of a \mathcal{C}^2 function satisfies $\nabla^2 g(u^*) > 0$, it can be shown that, given a sufficiently small initial error $\|u_0 - u^*\|$, the sequence $\|u_k - u^*\|$ is quadratically convergent to zero. Moreover, since a local minimizer of a \mathcal{C}^2 function satisfy $\nabla g(u^*) = 0$, it is possible to use $\|\nabla g(u_k)\| \leq \text{tol}$, with $\text{tol} > 0$, as an exit condition for the iterative solver.

C. quasi-Newton Method

Even in cases where the requirement $\nabla^2 g(u_k) > 0$ holds for u_k sufficiently close to the optimum, there is no guarantee that $\nabla^2 g(u_k) > 0$, $\forall u_k \in \mathbb{R}^m$. Given $\nabla^2 g(u_k) \not> 0$, the quadratic approximation (19) does not admit a bounded solution, which causes the Newton method to fail. In such instances, the quasi-Newton method can overcome this issue by replacing (19) with

$$\nu_k = \arg \min_{\nu \in \mathbb{R}^m} \nabla g(u_k)\nu + \frac{1}{2}\nu^T G_k \nu, \quad (21)$$

where $G_k > 0$ is a positive-definite matrix. Since $G_k = I$ leads to the gradient descent method $\nu_k = -\nabla g(u_k)$, which is linearly convergent, we note that, if $G_k > 0$ is a suitable approximation of $\nabla^2 g(u_k)$, the quasi-Newton method can achieve superlinear convergence.

For example, given $g(u) = h[p(u)]$, we have

$$\nabla^2 g(u) = \nabla p(u)^T \nabla^2 h[p(u)] \nabla p(u) + \sum_{j=1}^m \frac{\partial h}{\partial u_j} \Big|_{p(u)} \nabla^2 p_j(u)$$

which is, in general, not positive definite. If h is a convex function, however, a suitable positive definite approximation of $\nabla^2 g(u_k) \not> 0$ would be

$$G_k = \nabla p(u_k)^T \nabla^2 h[p(u_k)] \nabla p(u_k). \quad (22)$$

D. Damped (quasi-)Newton Method

Since (18), or its quasi-Newton counterpart (21), is only a second-order approximation, there is no guarantee that $g(u_k + \nu_k) < g(u_k)$. This can cause the Newton method to diverge when the initial error $\|u_0 - u^*\|$ is too large. To prevent this, we replace the classic Newton step (17) with the damped Newton step

$$u_{k+1} = u_k + \gamma_k \nu_k, \quad (23)$$

where the **Step Size** $\gamma_k \in (0, 1]$ is chosen to ensure a sufficient decrease in the cost function. Notably, if we define $g_k : (0, 1] \rightarrow \mathbb{R}$ as

$$g_k(\gamma) = g(u_k + \gamma \nu_k), \quad (24)$$

we can compute the Taylor expansion

$$g_k(\gamma) = g_k(0) + g'_k(0)\gamma + \frac{1}{2}g''_k(0)\gamma^2 + o(\gamma^2). \quad (25)$$

By construction of ν_k , we know that the minimum of the quadratic approximation of $g_k(\gamma)$ lies in $\gamma = 1$. As shown in Figure 1, however, the higher order terms $o(\gamma^2)$ may cause $g_k(1)$ to be quite different from its quadratic approximation, thus causing the Newton method to fail. To ensure that the cost decreases by a sufficient amount at each iteration, we use a backtracking linesearch to enforce the Armijo Rule

$$g_k(\gamma_k) \leq g_k(0) - \alpha g'_k(0)\gamma_k, \quad (26)$$

with $\alpha \in (0, 0.5)$. Specifically, given the initial value $\gamma_k = 1$, we check if (26) holds. If it does, the higher-order terms $o(\gamma^2)$ are not dominant. If it doesn't, we assign

$$\gamma_k \leftarrow \beta \gamma_k, \quad (27)$$

with $\beta \in (0, 1)$ and proceed to once again check the Armijo Rule (26). Typical values for the backtracking linesearch are $\alpha = 0.4$ and $\beta = 0.7$. This modification is sufficient to ensure the monotonic convergence property

$$g(u_{k+1}) < g(u_k), \quad (28)$$

thus guaranteeing monotonic convergence to a local minimum for any initial condition such that $\nabla g(u_0) \neq 0$, i.e. for any initial condition that isn't a local maximum or a saddle point.

IV. PRONTO FOR QUANTUM SYSTEMS

Having given a general intuition behind of all the steps used in PRONTO, we now go back to the original optimization problem (14) for which the solution space is $[x(t), u(t)] \in (\mathcal{X} \times \mathcal{U})$, where \mathcal{X} denotes the set of all \mathcal{C}^0 functions $x : [0, T] \rightarrow \mathbb{R}^{2n}$ and \mathcal{U} denotes the set of all \mathcal{C}^0 functions $u : [0, T] \rightarrow \mathbb{R}^m$. Here, the idea is to perform the same steps laid out in Sec. III using functional analysis instead of vector calculus. The PRONTO pseudocode obtained by following all these steps is provided in Algorithm 1 at the end of this section.

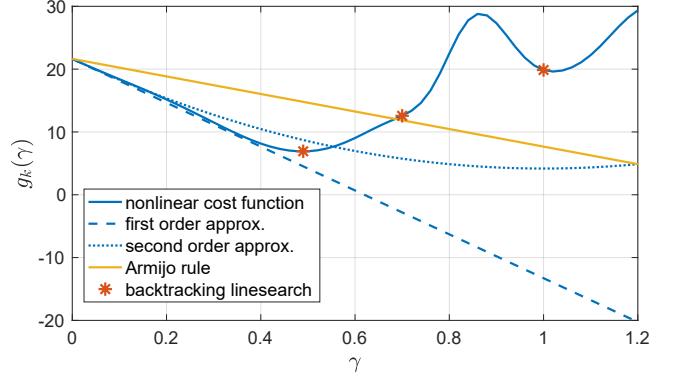


FIG. 1. Visual representation of the backtracking linesearch applied to a nonlinear cost function $g_k(\gamma)$. For $\gamma = 1$ and $\gamma = 0.7$, the nonlinearities are too strong to be captured by a second order approximation. Given $\gamma = 0.49$, the Armijo rule is satisfied, meaning that the decrease in cost is acceptable. Note that, for $\gamma \rightarrow 0$, the nonlinear function $g_k(\gamma)$ tends to coincide with its first and second order approximations.

A. The Projection Operator

The first step to solve the constrained optimization problem (14) is to reformulate it as an *unconstrained* optimization problem by embedding the equality constraints into the cost function. To this end, we define $\xi(t) = [x(t), u(t)]$ which means we can rewrite (14a) as

$$h(\xi) = m[x(T)] + \int_0^T l[x(t), u(t)]dt, \quad (29)$$

where $h : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is a \mathcal{C}^2 convex function. Since the Schrödinger equation (14b) imposes the restriction $\xi \in \mathcal{T}$, where

$$\mathcal{T} = \{\xi \in (\mathcal{X} \times \mathcal{U}) \mid \dot{x} = \mathcal{H}(u)x, \quad x(0) = x_0\}, \quad (30)$$

we define the simplified¹ projection operator $\mathcal{P} : \mathcal{U} \rightarrow \mathcal{T}$ as the solution to

$$\mathcal{P}(\mu) = \xi : \begin{cases} \dot{x} = H(u)x, & x_0 = \bar{x} \\ u = \mu. \end{cases} \quad (31)$$

This operator simply maps a specific control field $\mu(t)$ to an abstract state $\xi(t)$ consisting of both the control field and the dynamic response of the state under that control field. Then, given $g(u) = h(\mathcal{P}(u))$, the optimal control

¹ As detailed in [14], the projection operator should formally map $(\mathcal{X} \times \mathcal{U}) \rightarrow \mathcal{T}$. Given $\alpha \in \mathcal{X}$, we can turn (31) into a “proper” projection operator $\mathcal{P}(\alpha, \mu)$ by defining $u = \mu + K_r(\alpha - x)$, where K_r is a time-varying feedback gain that stabilizes the system trajectories. Since closed quantum systems evolve on the unit sphere (and are therefore inherently stable), we assign $K_r = 0$ for the sake of simplicity. The design of suitable $K_r \neq 0$, which is likely to improve convergence, is left to future work.

problem (14) can be rewritten as an *unconstrained* optimization problem in the form

$$\min g(u), \quad (32)$$

which we now solve using a damped (quasi-)Newton method in the function space \mathcal{U} .

B. Newton Descent Direction

In analogy to Sec. III B, given the current solution estimate $u_k(t) \in \mathcal{U}$, we now wish to compute a descent

direction $\nu_k(t) \in \mathcal{U}$ for our next estimate by minimizing the local quadratic approximation of (32). This can be done by solving the optimization problem

$$\nu_k = \arg \min Dg(u_k) \circ \nu + \frac{1}{2} D^2 g(u_k) \circ (\nu, \nu), \quad (33)$$

where D and D^2 are the first and second Fréchet derivatives of $g(u)$. As detailed in Appendix A, the evaluation of (33) leads in the following Linear-Quadratic Optimal Control Problem (LQ-OCP)

$$\begin{aligned} \nu_k(t) = \arg \min & \pi_k^T z(T) + \frac{1}{2} z(T)^T \Pi_k z(T) + \int_0^T \begin{bmatrix} q_k(\tau) \\ r_k(\tau) \end{bmatrix}^T \begin{bmatrix} z(\tau) \\ \nu(\tau) \end{bmatrix} + \frac{1}{2} \begin{bmatrix} z(\tau) \\ \nu(\tau) \end{bmatrix}^T \begin{bmatrix} Q_k(\tau) & S_k(\tau) \\ S_k^T(\tau) & R_k(\tau) \end{bmatrix} \begin{bmatrix} z(\tau) \\ \nu(\tau) \end{bmatrix} d\tau \\ \text{s.t. } & \dot{z}(t) = A_k(t)z(t) + B_k(t)\nu(t), \quad z(0) = 0, \end{aligned} \quad (34)$$

where

$$\begin{aligned} q_k(t) &= \nabla_x l(x_k(t), u_k(t)), \\ r_k(t) &= \nabla_{ul} l(x_k(t), u_k(t)), \\ \pi_k &= \nabla_x m(x_k(T)), \end{aligned} \quad (35)$$

capture the first-order contributions to the cost and

$$\begin{aligned} Q_k(t) &= \nabla_{xx}^2 l(x_k(t), u_k(t)), \\ R_k(t) &= \nabla_{uu}^2 l(x_k(t), u_k(t)) + \tilde{R}_k(t), \\ S_k(t) &= \nabla_{xu}^2 l(x_k(t), u_k(t)) + \tilde{S}_k(t), \\ \Pi_k &= \nabla_{xx}^2 m(x_k(T)), \end{aligned} \quad (36)$$

capture the second-order contributions, with

$$\begin{aligned} \tilde{S}_k(t) &= [H_1(t)\chi_k(t) \dots H_m(t)\chi_k(t)] \\ \tilde{R}_k(t) &= \begin{bmatrix} \chi_k^T(t)H_{11}(t)x_k(t) & \dots & \chi_k^T(t)H_{m1}(t)x_k(t) \\ \vdots & \ddots & \vdots \\ \chi_k^T(t)H_{1m}(t)x_k(t) & \dots & \chi_k^T(t)H_{mm}(t)x_k(t) \end{bmatrix}. \end{aligned}$$

Here, the brackets indicate a matrix concatenation of column vectors for \tilde{S}_k and scalars for \tilde{R}_k , the co-state² $\chi_k \in \mathcal{X}$ is obtained by solving the differential equation

$$-\dot{\chi}_k = H(u_k(t))^T \chi_k + q_k(t), \quad \chi(T) = \pi_k, \quad (37)$$

and, given the vector $w \in \mathbb{R}^m$, we define

$$H_i(t) = \frac{\partial H(w)}{\partial w_i} \Big|_{w=u_k(t)} \quad H_{ij}(t) = \frac{\partial^2 H(w)}{\partial w_i \partial w_j} \Big|_{w=u_k(t)}.$$

Finally, the matrices

$$\begin{aligned} A_k(t) &= H(u_k(t)), \\ B_k(t) &= [H_1(t)x_k(t) \dots H_m(t)x_k(t)], \end{aligned} \quad (38)$$

capture the linearization of system dynamics $\dot{x} = H(u)x$ around the current trajectory $\xi_k = (x_k(t), u_k(t))$.

Although somewhat daunting, (34) is a well-known optimal control problem [18], which admits an explicit solution (see Appendix B). Unfortunately, depending on the matrices in (36), there is no guarantee that (34) the solution exists and is unique. This issue is solved in the next subsection by using a positive definite approximation of the cost function whenever necessary.

C. quasi-Newton Descent Direction

If, at a given iteration k , the LQ-OCP (34) does not admit a solution, it is possible to replace (36) with

$$\begin{aligned} Q_k(t) &= \nabla_{xx}^2 l(x_k(t), u_k(t)), \\ R_k(t) &= \nabla_{uu}^2 l(x_k(t), u_k(t)), \\ S_k(t) &= \nabla_{xu}^2 l(x_k(t), u_k(t)), \\ \Pi_k &= \nabla_{xx}^2 m(x_k(T)), \end{aligned} \quad (39)$$

By construction, these matrices satisfy the positive semi-definite properties $\Pi_k \geq 0$, $R_k(t) > 0$, $\forall t \in [0, T]$, and $[Q_k(t) \ S_k(t); S_k(t)^T \ R_k(t)] \geq 0$, $\forall t \in [0, T]$. This is sufficient to ensure that the new LQ-OCP has a unique minimizer, making it always possible to compute a descent direction ν_k .

It is very interesting, and somewhat counter-intuitive, to note that the co-state $\chi_k(t)$ is not required to compute the quasi-Newton descent direction since $\chi_k(t)$ only enters the problem through the matrices \tilde{S}_k , \tilde{R}_k in (36).

² Coincidentally, it is worth noting that the co-state dynamics (37), which stem from the computation of $D^2 g(u_k)(\nu, \nu)$, detailed in Appendix A, are identical to what is found by the Krotov method [16, Eq. (30)] using Lagrange multipliers.

D. Maximum Step Size

In analogy to Sec. III D, since the descent direction $\nu_k(t)$ was computed using a *local* second order approximation, there is no guarantee that $u_{k+1}(t) = u_k(t) + \nu_k(t)$ satisfies $g(u_{k+1}) < g(u_k)$. To ensure a monotonically convergent sequence, we therefore define the next estimate

$$u_{k+1}(t) = u_k(t) + \gamma_k \nu_k(t), \quad (40)$$

where the step size $\gamma_k \in (0, 1]$ is chosen using a backtracking linesearch to enforce the Armijo rule

$$g(u_k + \gamma_k \nu_k) \leq g_k(u_k) - \alpha \gamma_k Dg(u_k) \circ \nu_k. \quad (41)$$

Details on how to compute $Dg(u_k) \circ \nu_k$ are also provided in Appendix B. A possible issue with the backtracking linesearch is that, if the Armijo rule is only satisfied for very small γ , the algorithm may perform a large number of checks before selecting a suitable step size.

To prevent unnecessary calculations, we introduce a heuristic that upper bounds the initial value of the backtracking linesearch whenever the norm of the descent direction is too large. To this end, we note that, although $x_{k+1} \in \mathcal{X}$ is obtained from the projection operator (31), the LQR problem (B2) approximates the state update as

$$x_{k+1} \approx x_k + \gamma_k z_k. \quad (42)$$

Since (14b) is a closed quantum system such that $\|x_{k+1}(t)\| = \|x_k(t)\| = \|x_0\|, \forall t \in [0, T]$, the approximation (42) is valid only if $\gamma_k \|z_k(t)\|$ is “sufficiently small” $\forall t \in [0, T]$. This motivates the step size upper bound

$$\gamma_k \leq \frac{\delta \|x_0\|}{\max_{t \in [0, T]} \|z_k(t)\|}, \quad (43)$$

where $\delta \in (0, 1)$ ensures that the norm of the update is, at most, comparable to the norm of the state. A reasonable heuristic for this upper bound is $\delta = 0.6$. Note that, once $\|z_k(t)\| \leq \delta \|x_0\|, \forall t \in [0, T]$, i.e. once the updates are sufficiently small, we return to the full range $\gamma_k \in (0, 1]$.

Although PRONTO only guarantees monotonic convergence to a local minimum, it is worth noting that there exists a large class of quantum optimal control problems for which all local minima share the same cost [22].

Algorithm 1 Q-PRONTO

Initialization

1: $Dg = -10\text{tol}$
Initial Control Guess
2: $u_k(t) = u_0(t)$
Solve Schrödinger Equation and Compute Cost
3: $\xi_k \leftarrow \mathcal{P}(u_k(t))$ ▷ F.i.t. (31)
4: $g_k \leftarrow h(\xi_k)$

Main Loop

5: **while** $-Dg \geq \text{tol}$ **do**
Compute Linear-Quadratic Approximation
6: $[A_k(t), B_k(t), q_k(t), r_k(t), \pi_k] \leftarrow [(38) - (35)]$
Perform Newton Step
7: Compute $\chi_k(t)$ ▷ B.i.t (37)
8: $[Q_k(t), S_k(t), R_k(t), \Pi_k] \leftarrow (36)$ ▷ B.i.t (B1)
9: Compute $[K_o(t), v_o(t)]$ ▷ B.i.t (B1)

Perform quasi-Newton Step

10: **if** (B1) failed to converge **then**
11: $[Q_k(t), S_k(t), R_k(t), \Pi_k] \leftarrow (39)$ ▷ B.i.t (B1)
12: Compute $[K_o(t), v_o(t)]$ ▷ B.i.t (B1)
13: **end if**

Compute Descent Direction and Cost Gradient

14: Compute $[z_k(t), \nu_k(t), \eta_k(T)]$ ▷ F.i.t. (B2)
15: $Dg \leftarrow \pi_k^T z_k(T) + \eta_k(T)$

Apply Armijo Rule for Step Size Selection

16: $\gamma_k \leftarrow \min(1, \delta \|x_0\| / \max(\|z_k(t)\|))$
17: $\xi_{k+1} \leftarrow \mathcal{P}(u_k(t) + \nu_k(t))$ ▷ F.i.t. (31)
18: $g_{k+1} \leftarrow h(\xi_{k+1})$
19: **while** $g_{k+1} > g_k - \alpha Dg \circ \gamma_k$ **do**
20: $\gamma_k \leftarrow \beta \gamma_k$
21: $\xi_{k+1} \leftarrow \mathcal{P}(u_k(t) + \gamma_k \nu_k(t))$ ▷ F.i.t. (31)
22: $g_{k+1} \leftarrow h(\xi_{k+1})$
23: **end while**

Proceed to Next Iteration

24: $u_k(t) \leftarrow u_k(t) + \gamma_k \nu_k(t)$
25: $\xi_k \leftarrow \xi_{k+1}$
26: $g_k \leftarrow g_{k+1}$
27: **end while**

Output

28: **return** $u_k(t)$

F.i.t. = Integrate forward in time

B.i.t. = Integrate backward in time

V. EXAMPLE: QUBIT STATE-TO-STATE CONTROL

We now apply PRONTO to a canonical quantum control problem and compare it to one of the leading quantum control techniques, the Krotov method [9, 10].

Specifically, we wish to perform a state-to-state transition $|0\rangle \rightarrow |1\rangle$ on a qubit evolving under the Schrödinger equation

$$|\dot{\psi}\rangle = -i(\hat{H}_0 + u\hat{H}_1)|\psi\rangle, \quad |\psi(0)\rangle = |0\rangle, \quad (44)$$

where $\hat{H}_0 = -\frac{\omega}{2}\hat{\sigma}_z$, $\hat{H}_1 = \hat{\sigma}_x$, $\hat{\sigma}_i$ are the usual Pauli

matrices and $\omega = 1$. To achieve this goal while limiting the fluence of the control input $u(t)$, we minimize the cost function

$$\frac{1}{2} \langle \psi(T) | \hat{P}_{\neg 1} | \psi(T) \rangle + \int_0^T \frac{\vartheta(t)}{2} \|u(t)\|^2 dt, \quad (45)$$

where $\hat{P}_{\neg 1} = I - |1\rangle\langle 1| = |0\rangle\langle 0|$ is a projection operator, $\vartheta(t) > 0$, $\forall t \in [0, T]$ is a time-varying penalty on the energy of the control input, and $T = 5$ s is the control horizon. Using the bijective mapping in (9) and (10), this optimization problem can then be rewritten as

$$\min \frac{1}{2} x(T)^T P_{\neg 1} x(T) + \int_0^T \frac{\vartheta(t)}{2} \|u(t)\|^2 dt \quad (46a)$$

$$\text{s.t. } \dot{x} = (H_0 + uH_1)x, \quad x = x_0. \quad (46b)$$

To completely define the optimal control problem, we now have to specify $\vartheta(t)$. For ease of comparison³ with the benchmark in the Krotov Package, we assign the same input weight function featured in [16], i.e.

$$\vartheta(t) = \begin{cases} \frac{1+\epsilon}{\mathcal{B}_{.6}(t)+\epsilon} & \forall t \in [0, 0.3], \\ 1 & \forall t \in (0.3, 4.7], \\ \frac{1+\epsilon}{\mathcal{B}_{.6}(5-t)+\epsilon} & \forall t \in [4.7, 5], \end{cases} \quad (47)$$

where $\epsilon = 10^{-6}$ and

$$\mathcal{B}_{.6}(t) = \frac{1}{2} \left(.84 - \cos \left(\frac{2\pi t}{.6} \right) + .16 \cos \left(\frac{4\pi t}{.6} \right) \right) \quad (48)$$

is a Blackman window of length 0.6. Note that $\epsilon > 0$ ensures that $\vartheta(t)$ remains bounded for $t = 0$ and $t = 5$. The control cost $\vartheta(t)$ in (47) is high at the beginning and end, and constant throughout the majority of the control time interval. This is used to ensure that the input starts at zero and ends at zero, thereby making the control input better suited for hardware implementation.

Finally, to initialize our iterative solver, we use the same initial guess featured in [16], i.e.

$$u_0(t) = \begin{cases} 0.2 \mathcal{B}_{.6}(t) & \forall t \in [0, 0.3], \\ 0.2 & \forall t \in (0.3, 4.7), \\ 0.2 \mathcal{B}_{.6}(5-t) & \forall t \in [4.7, 5]. \end{cases} \quad (49)$$

A. Benchmark Comparison, Part I

To perform an initial comparison between PRONTO and the Krotov method implemented in [16], we begin by

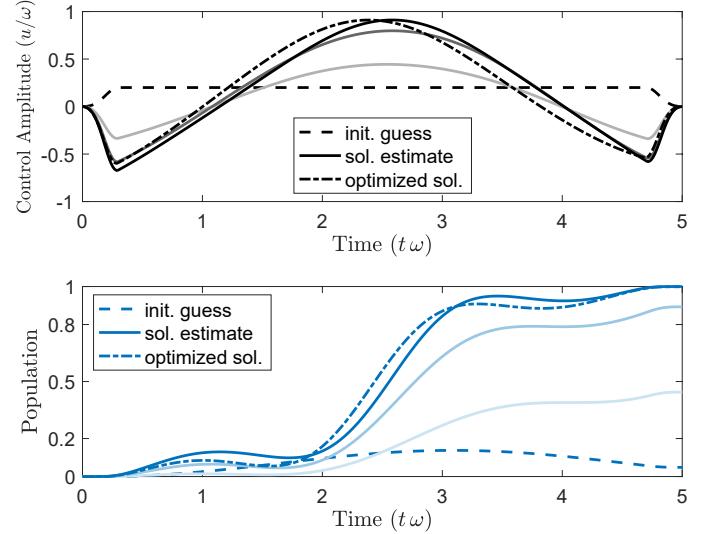


FIG. 2. **Top:** Control input $u(t)$ and **Bottom:** Population $P_1(t) = |\langle \psi(t) | 1 \rangle|^2$ for the qubit benchmark comparison. The dashed lines are the initial guess, the solid lines are the solution estimate satisfying $\text{tol} \leq 10^{-2}$, and the dotted lines are the solution obtained using Krotov. Intermediate iterations of PRONTO are represented using semi-transparent lines.

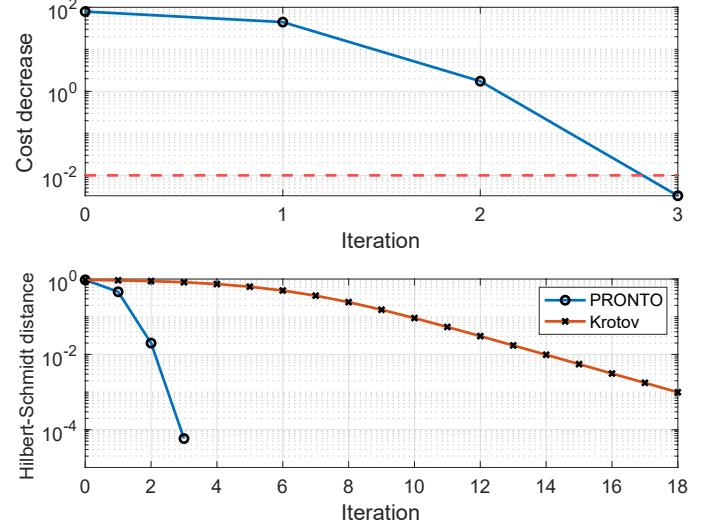


FIG. 3. **Top:** Value of the cost decrease $-Dg(u_k) \circ \nu_k$ at each iteration of PRONTO. The red dashed line is the exit condition $\text{tol} = 10^{-2}$. **Bottom:** Value of the squared Hilbert-Schmidt distance $1 - |\langle \psi(T) | 1 \rangle|^2$ at each iteration of PRONTO and of the Krotov method.

solving the quantum optimal control problem (46) using Algorithm 1 subject to $\text{tol} = 10^{-2}$.

Figure 2 illustrates the behavior of the control input (top panel) and the population dynamics (bottom panel) at each iteration. In the top panel it is evident that the control amplitude remains bounded and does not contain any discontinuities.

Figure 3 (Top) illustrates the value of the exit condi-

³ Unfortunately, the optimal control problem featured in [16] is different from Eq. (46). The Krotov method penalizes the control update, i.e. $\int (u_k(t) - u_{k-1}(t))^2 dt$, to achieve convergence, whereas the Bolza-type cost functional (14) penalizes the control effort, i.e. $\int u(t)^2 dt$, as in [23, 24]. See [25, Chap. 1] for further discussion on Bolza-type optimal control formulations.

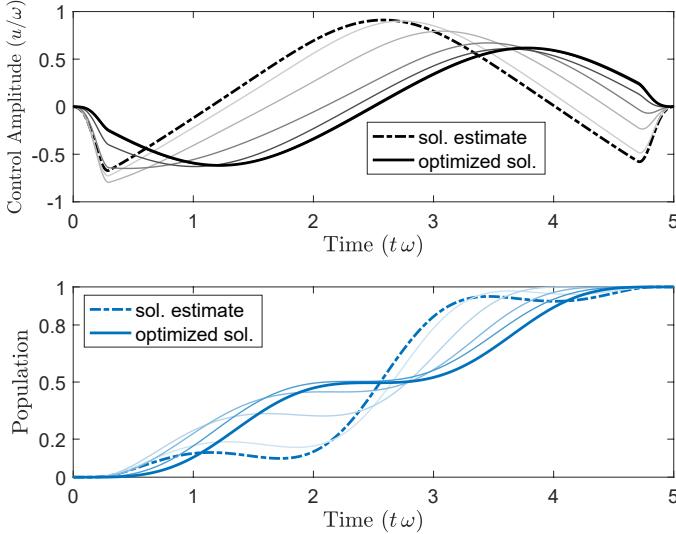


FIG. 4. **Top:** Control input $u(t)$ and **Bottom:** Population $P_1(t) = |\langle \psi(t) | 1 \rangle|^2$ for the qubit benchmark comparison. The dash-dotted lines denote the solution estimate satisfying $\text{tol} \leq 10^{-2}$ and the solid lines denote the optimized solution satisfying $\text{tol} \leq 10^{-8}$. Intermediate iterations of PRONTO are represented using semi-transparent lines.

tion $-Dg(u_k) \circ \nu_k \leq \text{tol}$ at each iteration. The method meets the desired tolerance after just 3 iterations. In Figure 3 (Bottom) we compare PRONTO to the Krotov method [16] by plotting the value of the Hilbert-Schmidt distance $1 - |\langle \psi(T) | 1 \rangle|^2$ obtained by each method at each iteration. For this particular example, all PRONTO iterations ended up using the Newton step as opposed to the quasi-Newton step. Although the two methods converge to similar results, the convergence rate is quadratic in the case of PRONTO and linear in the case of Krotov. This is not surprising since the former is a Newton method, whereas the latter is a gradient descent method.

B. Benchmark Comparison, Part II

We now continue to solve the quantum optimal control problem (46) up to $\text{tol} = 10^{-8}$. Intuitively, if the estimates obtained in the previous section were sufficiently close to the solution of (46), the method would only require a few additional iterations to meet the desired tolerance and there would be no perceivable changes to the estimate. Instead, something curious happens: the method departs from the trajectory obtained in [16] and converges to an entirely different solution.

Figure 4 illustrates the behavior of the control input (top panel) and the population dynamics (bottom panel) at every third iteration of Algorithm 1. Looking at the population dynamics, we note that every iteration achieves the target objective $1 - |\langle \psi_k(T) | 1 \rangle|^2 \approx 0$. Looking at the control input, we note that the overall amplitude of $u_k(t)$ tends to decrease at every iteration.

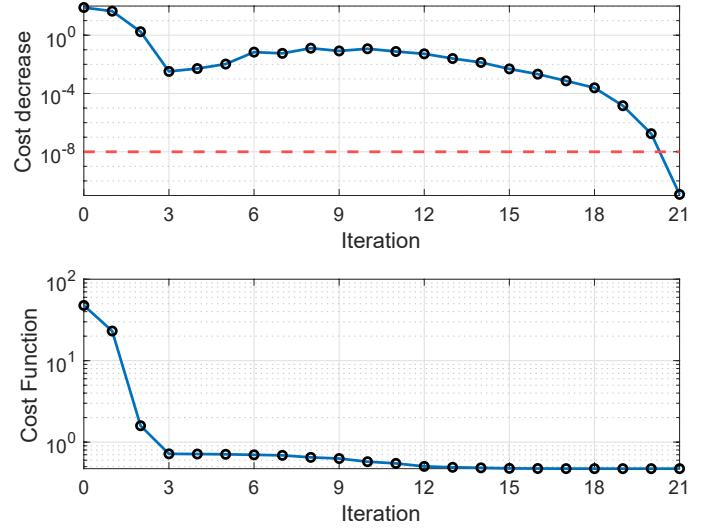


FIG. 5. **Top:** Value of the cost decrease $-Dg(u_k) \circ \nu_k$ at each iteration of PRONTO. The red dashed line is the exit condition $\text{tol} = 10^{-8}$. **Bottom:** Value of the cost function (45) at each iteration of PRONTO.

This can be interpreted as PRONTO initially moving in the direction that minimizes the cost function (thereby obtaining a solution estimate similar to Krotov), but then refining its solution to *also* minimize the control effort. This is not surprising since the quantum optimal control problem (46) penalizes both the Hilbert-Schmidt distance *and* the fluence of the control input, as opposed to the Krotov method, which only penalizes the Hilbert-Schmidt distance.

Figure 5 illustrates the value of the exit condition $-Dg(u_k) \circ \nu_k \leq \text{tol}$ (top panel) and of the cost function (3) (bottom panel) at each iteration. The cost function is monotonically decreasing and is lower-bounded by the value at the local minimizer. Quadratic convergence to the solution is achieved in proximity of the minimizer.

In this case, the advantage of using PRONTO over Krotov is not in terms of requiring a lower computational effort, but in terms of obtaining an arguably “better” control input (i.e., the same output result is obtained with a smaller fluence $\int_0^T \vartheta(t) \|u(t)\|^2 dt$).

C. Multiple Control Inputs

We now wish to perform a state-to-state transition $|0\rangle \rightarrow |1\rangle$ in the multi-input case of a qubit described by the Schrödinger equation

$$\hbar |\dot{\psi}\rangle = -i(\hat{H}_0 + u_1 \hat{H}_1 + u_2 \hat{H}_2) |\psi\rangle, \quad |\psi(0)\rangle = |0\rangle, \quad (50)$$

where $\hat{H}_2 = \hat{\sigma}_y$ and all other parameters are the same as (44). Given $u = [u_1, u_2]^T$, the optimization problem (46) becomes

$$\min \frac{1}{2} x(T)^T P_{-1} x(T) + \int_0^T \frac{\vartheta(t)}{2} \|u(t)\|^2 dt \quad (51a)$$

$$\text{s.t. } \dot{x} = (H_0 + u_1 H_1 + u_2 H_2)x, \quad x = x_0. \quad (51b)$$

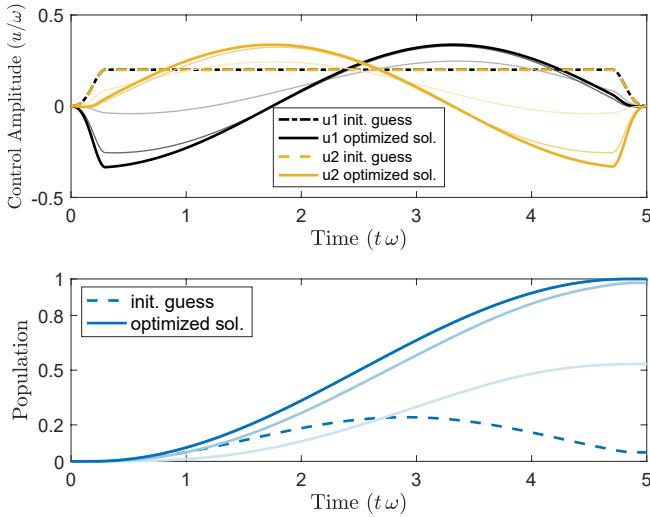


FIG. 6. **Top:** Control inputs $u_1(t)$, $u_2(t)$ and **Bottom:** Population $P_1(t) = |\langle \psi(t) | 1 \rangle|^2$ for the multi-input qubit. The dashed lines are the initial guess and the solid lines are the optimized solution satisfying $\text{tol} \leq 10^{-8}$. Intermediate iterations are represented using semi-transparent lines.

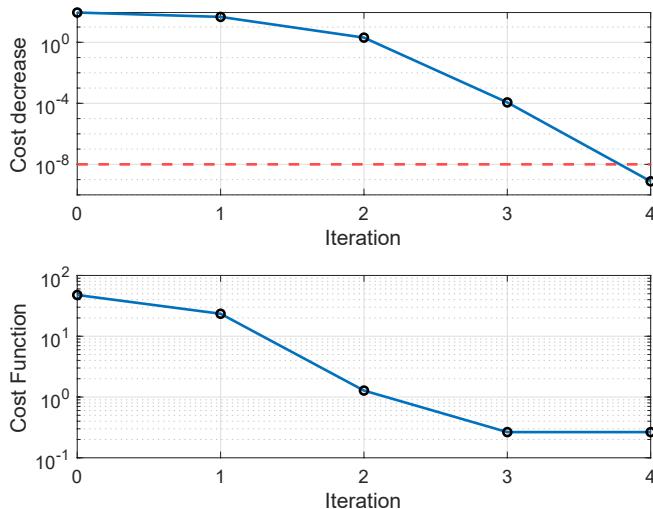


FIG. 7. **Top:** Value of the cost decrease $-Dg(u_k) \circ \nu_k$ at each iteration of PRONTO. The red dashed line is the exit condition $\text{tol} = 10^{-8}$. **Bottom:** Value of the cost function (45) at each iteration of PRONTO.

The initial guess (49) is used for both u_1 and u_2 . The quantum optimal control problem (51) is solved using Algorithm 1 subject to $\text{tol} = 10^{-8}$.

Figure 6 illustrates the behavior of the control input (top panel) and the population dynamics (bottom panel) at each iteration. In this example, it is interesting to note that the shape of $u_1(t)$ is similar to the one found in the previous section, but the amplitude is lower due to the contribution of $u_2(t)$.

Figure 7 illustrates the value of the exit condition $-Dg(u_k) \circ \nu_k \leq \text{tol}$ at each iteration. The method quadratically converges to the optimizer in 4 iterations. Once again, the cost function decreases monotonically and quadratically converges to a lower bound.

VI. CONCLUSIONS

This paper specialized the Projection Operator-based Newton method for Trajectory Optimization (PRONTO) to quantum control problems. The method is guaranteed to be monotonically convergent at all times and features a quadratic convergence rate in proximity of local minima.

There are many directions for future work. The original PRONTO method [14, 15] included a “regulator” that is not featured in the quantum extension presented here. Thus, an important technical extension of the quantum PRONTO is to design a quantum-specific regulator to be used in the projection operator. This will likely enable faster convergence for higher-dimensional quantum systems. Additional extensions include the design of unitary gates, control of open quantum systems, and in-depth comparisons with existing optimization methods. Eventually, we plan to release an open source toolkit/package for quantum control using PRONTO.

Acknowledgements The authors would like to thank Liang-Ying Chih and Murray Holland for the helpful discussions. The work by Jieqiu Shao and Marco M. Nicotra is supported by the NSF QII-TAQS award number 1936303.

[1] C. Brif, R. Chakrabarti, and H. Rabitz, Control of quantum phenomena: past, present and future, *New J. of Phys.* **12**, 075008 (2010).

[2] D. Dong and I. Petersen, Quantum control theory and applications: a survey, *EIT Control Theory Appl.* **4**, 2651 (2010).

[3] C. Altafini and F. Ticozzi, Modeling and control of quantum systems: An introduction, *IEEE Trans. Autom. Control* **57**, 1898 (2012).

[4] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Training Schrödinger’s cat: quantum optimal control, *European Physics Journal D* **69**, 279 (2015).

[5] S. Lloyd and S. L. Braunstein, Quantum computation over continuous variables, *Phys. Rev. Lett.* **82**, 1784 (1999).

[6] P. Doria, T. Calarco, and S. Montangero, Optimal control technique for many-body quantum dynamics, *Phys. Rev. Lett.* **106**, 190501 (2011).

- [7] 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. Magnet. Res.* **172**, 296 (2005).
- [8] N. Khaneja, R. Brockett, and S. J. Glaser, Time optimal control in spin systems, *Phys. Rev. A* **63**, 032308 (2001).
- [9] S. E. Sklarz and D. J. Tannor, Loading a Bose-Einstein condensate onto an optical lattice: an application of optimal control theory to the nonlinear Schrödinger equation, *Phys. Rev. A* **66**, 053619 (2002).
- [10] D. M. Reich, M. Ndong, and C. P. Koch, Monotonically convergent optimization in quantum control using Krotov's method, *J. Chem. Phys.* **136**, 104103 (2012).
- [11] F. Wilhelm, S. Kirchhoff, S. Machnes, N. Wittler, and D. Sugny, An introduction into optimal control for quantum technologies, *arXiv: Quantum Physics* (2020).
- [12] R. Eitan, M. Mundt, and D. J. Tannor, Optimal control with accelerated convergence: combining the Krotov and quasi-Newton methods, *Phys. Rev. A* **83**, 053426 (2011).
- [13] P. de Fouquieres, S. G. Schirmer, S. J. Glaser, and I. Kuprov, Second order gradient ascent pulse engineering, *J. of Magnetic Resonance* **212**, 412 (2011).
- [14] J. Hauser, A projection operator approach to the optimization of trajectory functionals, *IFAC Proc. Vol.* **35**, 377 (2002).
- [15] J. Hauser, On the computation of optimal state transfers with application to the control of quantum spin systems, *Proc. of ACC* **3**, 2169 (2003).
- [16] 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**, 80 (2019).
- [17] J. P. Palao, R. Kosloff, and C. P. Koch, Protecting coherence in optimal control theory: state-dependent constraint approach, *Phys. Rev. A* **77**, 063412 (2008).
- [18] B. D. O. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*, Dover Books on Engineering (2007).
- [19] J. P. Palao and R. Kosloff, Optimal control theory for unitary transformations, *Phys. Rev. A* **68**, 062308 (2003).
- [20] J. Nocedal and S. J. Wright, *Numerical optimization*, Springer (2006).
- [21] V. Krotov, *Global Methods in Optimal Control Theory*, CRC Press (1995).
- [22] H. A. Rabitz, M. M. Hsiem, and C. M. Rosenthal, Quantum optimally controlled transition landscapes, *Science* **303**, 1998 (2008).
- [23] D. D'Alessandro and M. Dahleh, Optimal control of two-level quantum systems, *IEEE Trans. Autom. Control* **46**, 866 (2001).
- [24] S. Grivopoulos and B. Bamieh, Optimal population transfers in a quantum system for large transfer time, *IEEE Trans. Autom. Control* **53**, 980 (2008).
- [25] L. Cesari, Problems of optimization—a general view, in *Optimization—Theory and Applications: Problems with Ordinary Differential Equations* (Springer New York, New York, NY, 1983) pp. 1–23.

Appendix A: Fréchet Derivatives

This appendix shows how to obtain the LQR problem (34) by formally computing the first and second Fréchet derivatives featured in (33).

First Derivatives

Using the chain rule, we obtain

$$Dg(u_k) \circ \nu = Dh(\mathcal{P}(u_k)) \circ D\mathcal{P}(u_k) \circ \nu. \quad (\text{A1})$$

Following from the definition of the projection operator (31), we have $\mathcal{P}(u_k) = \xi_k$. Moreover, its local derivative $\zeta(t) = (z(t), v(t))$ satisfies

$$D\mathcal{P}(u_k) \circ \nu = \zeta : \begin{cases} \dot{z} = A_k(t)z + B_k(t)v, & z(0) = 0 \\ v = \nu, \end{cases}$$

with $A_k(t)$ and $B_k(t)$ defined in (38). Thus, the first derivative can be rewritten as

$$Dg(u_k) \circ \nu = Dh(\xi_k) \circ \zeta. \quad (\text{A2})$$

It then follows from (29) that

$$Dh(\xi_k) \circ \zeta = \pi_k^T z(T) + \int_0^T q_k^T(\tau)z(\tau) + r_k^T(\tau)\nu(\tau)d\tau,$$

with

$$\begin{aligned} q_k(t) &= \nabla_x l(x_k(t), u_k(t)), \\ r_k(t) &= \nabla_{ul}(x_k(t), u_k(t)), \\ \pi_k &= \nabla_x m(x_k(T)), \end{aligned} \quad (\text{A3})$$

which translate into (35) once all the partial derivatives are evaluated.

Second Derivatives

By applying the chain rule to (A1), we obtain

$$\begin{aligned} D^2g(u_k) \circ (\nu, \nu) &= D^2h(\mathcal{P}(u_k)) \circ (D\mathcal{P}(u_k) \circ \nu, D\mathcal{P}(u_k) \circ \nu) \\ &\quad + Dh(\mathcal{P}(u_k)) \circ D^2\mathcal{P}(u_k) \circ (\nu, \nu). \end{aligned}$$

Substituting $\mathcal{P}(u_k) = \xi_k$ and $D\mathcal{P}(u_k) \circ \nu = \zeta$ then implies

$$\begin{aligned} D^2g(u_k) \circ (\nu, \nu) &= D^2h(\xi_k) \circ (\zeta, \zeta) + \dots \\ &\quad Dh(\xi_k) \circ D^2\mathcal{P}(u_k) \circ (\nu, \nu). \end{aligned} \quad (\text{A4})$$

It then follows from (29) that the first term satisfies

$$\begin{aligned} D^2h(\xi_k) \circ (\zeta, \zeta) &= z(T)^T \Pi_k z(T) + \dots \\ &\quad \int_0^T [z(\tau)]^T [Q_k(\tau) \quad S_k(\tau)] [z(\tau)] d\tau, \end{aligned}$$

with

$$\begin{aligned} Q_k(t) &= \nabla_{xx}^2 l(x_k(t), u_k(t)) \\ S_k(t) &= \nabla_{xu}^2 l(x_k(t), u_k(t)) \\ R_k(t) &= \nabla_{uu}^2 l(x_k(t), u_k(t)) \\ \Pi_k &= \nabla_{xx}^2 m(x_k(T)), \end{aligned} \quad (\text{A5})$$

which yield (39) when evaluated.

As for the second term, we note that the second derivative of the projection operator (31) yields

$$D^2\mathcal{P}(u_k) \circ (\nu, \nu) = \begin{cases} \dot{y} = A_k(t)y + B_k(t)w + \phi_k(t) \\ w = 0, \end{cases}$$

with $y(0) = 0$ and

$$\phi_k(t) = \sum_{i=1}^m \nu_i(t) \mathcal{H}_i(t) z(t) + \sum_{i,j=1}^m \nu_i(t) \nu_j(t) \mathcal{H}_{ij}(t) x_k(t).$$

Given the state transition matrix $\Phi(t, \tau)$ satisfying

$$\frac{\partial}{\partial t} \Phi(t, \tau) = A(t) \Phi(t, \tau), \quad (\text{A6})$$

we can write

$$y(t) = \int_0^t \Phi(t, s) \phi_k(s) ds. \quad (\text{A7})$$

Thus, we have

$$\begin{aligned} & Dh(\xi_k) \circ D^2\mathcal{P}(u_k) \circ (\nu, \nu) \\ &= \pi_k^T y(T) + \int_0^T q_k^T(\tau) y(\tau) d\tau \\ &= \pi_k^T y(T) + \int_0^T q_k^T(\tau) \int_0^\tau \Phi(\tau, s) \phi_k(s) ds d\tau \\ &= \pi_k^T y(T) + \int_0^T \int_s^T q_k^T(\tau) \Phi(\tau, s) d\tau \phi_k(s) ds \\ &= \int_0^T \left(\pi_k^T \Phi(T, s) + \int_s^T q_k^T(\tau) \Phi(\tau, s) d\tau \right) \phi_k(s) ds \end{aligned}$$

Given

$$\chi_k(s) = \Phi(T, s)^T \pi_k + \int_s^T \Phi(\tau, s)^T q_k(\tau) d\tau, \quad (\text{A8})$$

it follows from the properties of state transition matrices that $\chi : [0, T] \rightarrow \mathbb{R}^{2n}$ can be obtained by solving the

differential equation (37). Thus, we obtain

$$Dh(\xi_k) \circ D^2\mathcal{P}(u_k) \circ (\nu, \nu) = \int_0^T \chi_k^T(s) \phi_k(s) ds,$$

where it is possible to show that

$$\chi_k^T(s) \phi_k(s) = \begin{bmatrix} z(s) \\ \nu(s) \end{bmatrix}^T \begin{bmatrix} 0 & \tilde{S}_k(s) \\ \tilde{S}_k^T(s) & \tilde{R}_k(s) \end{bmatrix} \begin{bmatrix} z(s) \\ \nu(s) \end{bmatrix},$$

with $\tilde{R}_k(\cdot)$ and $\tilde{S}_k(\cdot)$ the same as in (36).

Appendix B: Linear-Quadratic Optimal Control

The optimal control problem (34) is a special class of trajectory optimization problems for which it is possible to compute an explicit solution [18]. To do so, we first solve the backwards in time Differential Riccati Equation

$$\begin{cases} -\dot{P} = A_k^T P + P A_k - K_o^T R_k K_o + Q_k, & P(T) = \Pi_k, \\ -\dot{p} = (A_k - B_k K_o)^T p - K_o^T r_k + q_k, & p(T) = \pi_k, \\ K_o = R_k^{-1} (B_k^T P + S_k^T), \\ v_o = R_k^{-1} (B_k^T p + r_k). \end{cases} \quad (\text{B1})$$

Having solved for $K_o(t)$ and $v_o(t)$, it is then possible to obtain $\nu_k(t)$ by solving

$$\begin{cases} \dot{\eta}_k = q_k^T z_k + r_k^T \nu_k, & \eta(0) = 0, \\ \dot{z}_k = A_k z_k + B_k \nu_k, & z(0) = 0, \\ \nu_k = -v_o - K_o z_k, \end{cases} \quad (\text{B2})$$

where $z_k \in \mathcal{X}$ is a local approximation of the state update and the running cost $\eta_k : [0, T] \rightarrow \mathbb{R}_{\geq 0}$ is used to compute

$$Dg(u_k) \circ \nu_k = \pi_k^T z_k(T) + \eta_k(T), \quad (\text{B3})$$

which is needed to a) determine the step size γ_k via the Armijo rule (41), and b) verify the exit condition for the iterative solver $-Dg(u_k) \cdot \nu_k \leq \text{tol}$.