12th IFAC International Symposium on
Advanced Control of Chemical Processes
July 14-17, 2024. Toronto, Canada

Fast Quantum Gate Control with Trajectory Optimization

Shouliang Hu * **Ming Li** ** **Chunlin Chen** *** **Daoyi Dong** ****

* *School of Engineering, Australian National University, Canberra, Australia (email: Shouliang.Hu@anu.edu.au)*
** *School of Automation, Guangdong University of Technology, Guangdong, China (email: mingli@gdut.edu.cn)*
*** *School of Management and Engineering, Nanjing University, Nanjing, China (email: clchen@nju.edu.cn)*
**** *CIICADA Lab, School of Engineering, Australian National University, Canberra, Australia (email: daoyi.dong@anu.edu.au)*

Abstract: Fast quantum control helps reduce the influence of unavoided disturbances and hence plays a vital role in practical quantum technology and chemical reactions. Instead of optimizing the terminal cost like standard optimal quantum control methods, this paper formulates the problem as a trajectory optimization problem, and implements the sequential quadratic programming algorithm to search for short control fields. The core idea is to minimize the cumulative intermediate error to incentivize early achievement of the designed gate. The numerical result on the Toffoli gate demonstrates the effectiveness of the proposed method.

Keywords: quantum gate, trajectory optimization, sequential quadratic programming.

1. INTRODUCTION

Controlling quantum phenomena is a fundamental task in many areas, such as guiding chemical reactions, initializing quantum computers, and constructing quantum circuits (see Dong and Petersen (2023); Fan et al. (2023); Dong and Petersen (2022); Dong et al. (2020)). In this paper, we focus on searching for short quantum control fields to accomplish fast quantum control. Fast quantum control is crucial for practical quantum technology, as the coherence time of practical qubits is limited and unavoided noises hinder the control performance with a long quantum operation time. The quantum speed limit (QSL) considers the maximum speed at which a quantum system can evolve while satisfying certain constraints (see Deffner and Campbell (2017); Caneva et al. (2009)). As shown in Fig. 1, a quantum system cannot be controlled to evolve to the target state faster than a specific time, which is referred to as the quantum speed limit time. Previous studies have revealed that the QSL time, τ_{QSL}, is a fundamental bound restricted by the intrinsic properties of the system (see Deffner and Campbell (2017)). Mandelstam and Tamm (1945) first found the uncertainty relation between energy and time, given as $\triangle H \cdot \triangle T \geq \hbar$, and derived the expression of QSL as $\tau_{\mathrm{QSL}} = \frac{\pi\hbar}{2\triangle H}$, where $\hbar$ is the reduced Planck constant and $\triangle H$ is the standard deviation of the Hamiltonian H. Later, Margolus and Levitin (1998) proposed another QSL expression, defined as $\tau_{\mathrm{QSL}} = \frac{\pi\hbar}{2<H>}$.

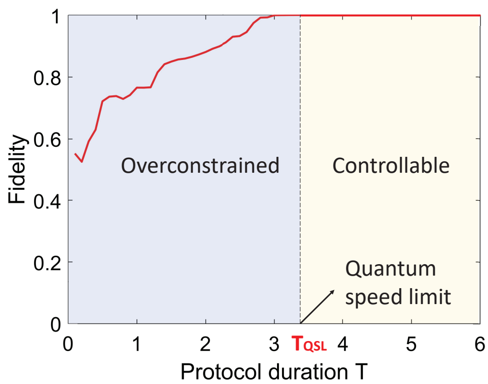

Fig. 1. Fidelity landscape for a three-qubit quantum gate.

There are various methods to estimate the quantum speed limit. For example, Boscain et al. (2014) characterized the time-optimal trajectories for two-level quantum systems based on the Pontryagin Maximum Principle. Jones and Kok (2010) gave a geometrical interpretation of the quantum speed limit. In particular, Caneva et al. (2009) first revealed the link between quantum dynamics and optimal algorithms, that is, the efficiency of the quantum optimal algorithm is governed by the intrinsic property of the system. Caneva et al. (2009) utilized optimal control methods to estimate the QSL. Instead of directly optimizing the evolution time T, they set a fixed value T and implemented the Krotov algorithm to optimize the controls. From an initial set of T, they steadily reduced the value of T and observed the failure of the Krotov algorithm at a certain threshold transfer time τ. The *'collapse time'* τ is proven to be very close to the theoretical estimate of T_{QSL}. Zahedinejad et al. (2014) also found the 'collapse

⋆ This work was supported by the Australian Research Council's Future Fellowship funding scheme under Project FT220100656 and the National Natural Science Foundation of China under Grant 62073160.

property' of the optimal algorithms (also named greedy algorithms), which would be trapped to local optimums with very short transfer time T and very short controls.

Inspired by Caneva et al. (2009), we are not meant to treat the time as a variable but to reach a local minimum transfer time near the quantum speed limit using the 'collapse' property of the optimal algorithm. In particular, this paper adopts the trajectory optimization technique. From the perspective of motion planning, the quantum gate control problem can be understood as a path planning problem with fixed starting and ending points. Leung et al. (2017) introduced the intermediate cost $J(\boldsymbol{u}) = 1 - \frac{1}{N}\sum_j \mathcal{F}_j$ to penalize deviations from the target gate and approach the time-optimal solution. Propson et al. (2022) utilized trajectory optimization to achieve time-optimal control and suppress errors caused by parameter uncertainties. This paper treats the intermediate error as the main cost instead of a penalty and implements the sequential quadratic programming algorithm to optimize the trajectory. The minimized cumulated deviations contribute to a time-optimal trajectory. Since it is hard to obtain a globally time-optimal trajectory through a single run of optimization, our methods work in an iterative way that adjusts the length of the trajectory and repeatedly implements the optimization to explore the better control trajectory. Numerical simulations of a three-qubit quantum gate demonstrate the effectiveness of the proposed method. The proposed method reduces the control time from an initial guess of $T = 10\mu$s to $T = 3.1\mu$s with few iterations.

The rest of this paper is organized as follows. Section 2 introduces several basic concepts about the quantum gate control task, trajectory optimization and sequential quadratic algorithm. In Section 3, the trajectory-optimal quantum gate control algorithm is presented in detail. The numerical results of a three-qubit quantum gate are presented in Section 4. Concluding remarks are drawn in Section 5.

2. PRELIMINARIES AND PROBLEM FORMULATION

This section provides a brief introduction to quantum gate control problem, trajectory optimization and sequential quadratic programming.

2.1 Control Design of Quantum Gates

Considering a n-level closed quantum system, the state $|\psi\rangle \in \mathbb{C}^n$ (wave function) can be represented with a complex vector (see Nielsen and Chuang (2001)), as

$$|\psi\rangle = [\alpha_1, \alpha_2, \cdots, \alpha_n]^T, \tag{1}$$

where $\{\alpha_j\}$ are amplitudes of orthogonal basis $|j\rangle$, satisfying $\sum_{j=1}^{n} |\alpha_j|^2 = 1$.

The state evolution follows the Schrödinger equation:

$$\frac{d}{dt}|\psi(t)\rangle = -iH(t)|\psi(t)\rangle, \quad |\psi(t_0)\rangle = |\psi_0\rangle. \tag{2}$$

The state transfer can be described by a propagator $U(t, t_0)$:

$$U(t, t_0)|\psi_0\rangle = |\psi(t)\rangle. \tag{3}$$

The propagator $U(t, t_0)$ can also be named as a *quantum gate* in quantum information, and its evolution follows

$$\frac{d}{dt}U(t) = -iH(t)U(t), \quad U(0) = \mathbb{I}^n, \tag{4}$$

where $\mathbb{I}^n = \text{diag}(1,1,\cdots,1)$ is the n-dimensional identity matrix. The Hamintonian $H(t)$ consists of two parts, 1) free Hamitonian H_0 and 2) control Hamitonian $H_c = \sum_{m=1}^{M} u_m(t)H_m$,

$$H(t) = H_0 + H_c[u(t)] \tag{5}$$

Here, $\{u_m(t)\}$ are external controls that interact with the quantum system, e.g., representing the amplitudes and phases of electromagnetic fields.

In quantum engineering, the piece-wise-constant (PWC) technique is often used to discretize the control function $\{u_m(t)\}$, i.e., $u_m(j) = c_j, t \in [(j-1)\triangle t, j\triangle t]$ hence simplifying the computation of propagator U. At the jth time interval, the jth propagator is given as

$$U_j = \exp[-iH(j)\triangle t]U_{j-1}. \tag{6}$$

The quantum gate control task is to design suitable controls $\{u_m(t)\}$ to steer the propagator toward the target gate.

2.2 Trajectory Optimization

Trajectory optimization mainly serves motion planning and control of robotics (see Howell et al. (2019)). Generally, a numerical trajectory optimization problem following discrete dynamics can be formulated as follows:

$$\begin{aligned}
\min_{\boldsymbol{x},\boldsymbol{u}} \quad & l_N(x_k) + \sum_{k=0}^{N-1} l_k(x_k, u_k) \\
s.t. \quad & x_{k+1} = f(x_k, u_k, \triangle t), \\
& g_k(x_k, u_k) \le 0, \\
& h_k(x_k, u_k) = 0.
\end{aligned}$$

Here, k is the time-step index, x_k and u_k are the states and control variables, N is the maximal control-step length, $g(x)$ and $h(x)$ are constraint functions, and $l_N(x_k)$ and $l_k(x_k, u_k)$ are terminal and intermediate costs.

2.3 Sequential Quadratic Programming

SQP is a widely used optimization tool for nonlinear optimization problems (see Boggs and Tolle (1995); Gill and Wong (2011)). A general form of nonlinear optimization task can be given as

$$\begin{aligned}
& \min_{\boldsymbol{x}} f(\boldsymbol{x}) \\
s.t. \quad & \nabla g_u(\boldsymbol{x}) \le 0 \ (u = 1, 2, \cdots, p), \\
& \nabla h_v(\boldsymbol{x}) = 0 \ (v = 1, 2, \cdots, q).
\end{aligned} \tag{7}$$

The target function $f(\boldsymbol{x})$ and constraint functions $g_u(\boldsymbol{x})$ and $h_v(\boldsymbol{x})$ all could be nolinear. SQP is an iterative search method starting from an initial guess $\boldsymbol{x}^k$ and its core idea is to transfer the nonlinear optimization into a linear quadratic programming problem using Taylor expansion, as

$$\begin{aligned}
\min_{\boldsymbol{x}} & f(\boldsymbol{x}) = \frac{1}{2}[\boldsymbol{x} - \boldsymbol{x}^k]^T \nabla^2 f(\boldsymbol{x}^k)[\boldsymbol{x} - \boldsymbol{x}^k] + \nabla f(\boldsymbol{x}^k)[\boldsymbol{x} - \boldsymbol{x}^k] \\
s.t. & \\
& \nabla g_u(x^k)^T[x - x^k] + g_u(x^k) \le 0 (u = 1, 2, \cdots, p), \\
& \nabla h_v(x^k)^T[x - x^k] + h_v(x^k) = 0 (v = 1, 2, \cdots, q).
\end{aligned}$$

3. TRAJECTORY OPTIMIZATION FOR CONTROL DESIGN OF QUANTUM GATES

We first consider a quantum gate optimal control problem, which mainly focuses on minimizing the terminal distance between the target gate W and the controlled $U(T)$:

$$\begin{aligned} \underset{\boldsymbol{u}}{\text{Minimize}} \quad & J(\boldsymbol{u}) = ||W - U(T)|| \\ s.t. \quad & U_{j+1} = e^{-iH(t_j)dt} U_j, \\ & H(t_j) = H_0 + u(t_j) H_c, \\ & |u(t_j)| \leq u_{\max}. \end{aligned} \tag{8}$$

Only the terminal state $U(T)$ is considered, while the intermediate states are ignored.

The overall control time T is usually treated as a fixed hyper-parameter. As mentioned before, a short protocol duration helps compensate for the effects of unavoided noises. To explore a proper setting of T, a practical way is to implement a numerical QOC algorithm like GRAPE (a first-order gradient descent algorithm) and vary the value of T to find the phase change of the fidelity landscape(see Bukov et al. (2018); Caneva et al. (2009)). But such a method usually be very computationally expensive.

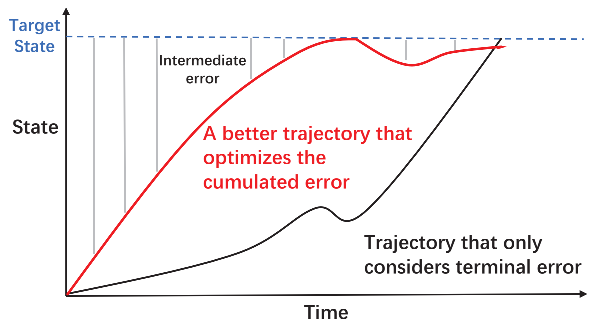

Fig. 2. Trajectory Optimization.

From the perspective of motion planning, the standard quantum gate control problem is a typical trajectory problem where the initial and target states are fixed, and the primary goal is to drift the error to zero with proper control design. Since only the terminal cost is considered, the trajectory obtained by the standard QOC method could be non-optimal. As shown in Fig. 2, the trajectory optimization would take account of all the intermediate errors and give a better trajectory that achieves the target state earlier. Hence, the core idea of this paper is to implement trajectory optimization to achieve the target gate faster.

To implement the trajectory optimization, a proper measure of the distance is of great importance. The squared error can be used to compute the distance between two states or matrices. Propson et al. (2022) mentioned that the squared-difference cost leads to a diagonal Hessian and can speed up the computation.

A more common way to estimate the distance between quantum states or propagators is to compute the 'Infidelity', which originates in the the squared Hilbert-Schmidt norm. The distance between two matrices can be measured by the squared Hilbert-Schmidt norm, which is defined as (see Palao and Kosloff (2002)):

$$||\boldsymbol{X}||_{\text{HS}}^2 = \text{Tr}[\boldsymbol{X}^\dagger \boldsymbol{X}], \tag{9}$$

Here, $\text{Tr}[\cdot]$ refers to the trace of a matrix and $\dagger$ represents conjugate transpose. Hence, the error between the jth propagator U_j and target gate W is

$$\begin{aligned} ||W - U_j||_{\text{HS}}^2 =& \text{Tr}[(W - U_j)^\dagger (W - U_j)] \\ =& 2n - 2\text{ReTr}[W^\dagger U_j], \end{aligned} \tag{10}$$

where n is the system dimension. Maximizing $\text{ReTr}[W^\dagger U_j]$ can be proved to be the same as maximizing $|\text{Tr}[W^\dagger U_j]|^2$ (see Wu et al. (2016)). We define the 'Fidelity' F:

$$\mathcal{F}(U_j, W) = \frac{1}{n^2} |\text{Tr}[W^\dagger U_j]|^2, \tag{11}$$

which lies at [0,1]. And $(1 - \mathcal{F})$ is referred to as 'infidelity'. When the target gate is perfectly achieved, the error $1 - \mathcal{F}(U_j, W)$ goes to zero.

Since many manipulation tasks of quantum systems demand high precision, i.e., the infidelity less than 10^{-4}, the logarithmic infidelity can be an improved measure and is given as:

$$\mathcal{L}(\boldsymbol{U}_j, W) = \log_{10}(1 - \mathcal{F}) = \log_{10}(1 - \frac{1}{n^2} |\text{Tr}\{W^\dagger \boldsymbol{U}_j\}|^2), \tag{12}$$

which reflects the error magnitude and lies at $(-\infty, 0]$.

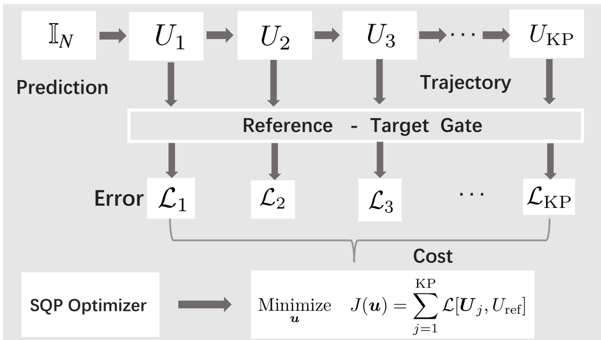

Fig. 3. Optimization framework

The optimization framework is shown in Fig. 3. The gate evolves from the identity matrix $\mathbb{I}^n$ and the cost J is defined as the sum of the intermediate error:

$$J(\boldsymbol{u}) = \sum_{j=1}^{\text{KP}} \mathcal{L}[U_j, W], \tag{13}$$

where KP is the length of the trajectory. We do not tend to optimize the energy cost but bound the drifts as $|u(t_j)| \leq u_{\max}$.

The optimization algorithm is presented in Algorithm 1, and includes two main steps.

1) **Step 1**

In this step, we predict the KP propagators based on a control guess $\boldsymbol{u}_{\text{guess}}$ and then implement the SQP algorithm to optimize the cost.

The optimization problem can be concluded as

Algorithm 1 Trajectory-optimal quantum gate control

Input: Initial guess of trajectory length $\mathrm{KP}^{(0)}$, the PWC duration time dt, termination condition $\mathcal{L}_{\mathrm{ter}}$, precision threshold ϵ for the SQP optimizer.
// **Iteration** //
Step 1: Implement the SQP algorithm to solve the trajectory optimization problem defined in (14) and obtain control $\boldsymbol{u}$.
Step 2: Compute and evaluate the KP-length error trajectory obtained in Step 1.
if $\min(\mathcal{L}_j) \geq \mathcal{L}_{\mathrm{ter}}$ **then**
  End the algorithm since it fails to achieve the target gate.
else
  Update $T = m \cdot dt$ where m is the step index of the propagator U_j that first achieves the target gate. Store the control variables $\boldsymbol{u}(1:m)$.
  Reset the trajectory length KP $= m$ and return to Step 1.
end
Output: Optimal control $\boldsymbol{u}^*$ and control time $T^* = m^* \cdot dt$.

$$\begin{aligned} \underset{\boldsymbol{u}}{\text{Minimize}} \quad & J(\boldsymbol{u}) = \sum_{j=1}^{\mathrm{KP}} \mathcal{L}[U_j, W] \\ s.t. \quad & U_{j+1} = e^{-iH(t_j)dt} U_j, \\ & H(t_j) = H_0 + u(t_j) H_c, \\ & |u(t_j)| \leq u_{\max},\ j = 1, \cdots, \mathrm{KP}. \end{aligned} \tag{14}$$

2) **Step 2**

Step 1 will output the KP-step-length optimal control $\boldsymbol{u}^*$. At the second step, based on $\boldsymbol{u}^*$, we compute and evaluate the KP-length error trajectories $\{\mathcal{L}_j, j = 1, 2, 3, \cdots, \mathrm{KP}\}$. The m-th propagator U_m^* that first achieves $\mathcal{L}(U_m^*, W) \leq \mathcal{L}_{\mathrm{ter}}$ denotes a better T setting as $T^* = m \cdot dt$. In the trajectory optimization task, it is hard to obtain a globally optimal trajectory through a single optimization run. A practical way is to optimize the trajectory step by step. Therefore, we store the optimal controls $\boldsymbol{u}(1:m)$, and repeat Step 1 with a new setting of the trajectory length KP$= m$ to further optimize the trajectory.

When all $\mathcal{L}_j$ fail to achieve $\mathcal{L}(U_j^*, W) \leq \mathcal{L}_{\mathrm{ter}}$, it indicates the control time is approaching the QSL limit or the SQP algorithm is trapped to a local optimum. For this case, we stop the iteration and end the algorithm.

Each new iteration gives a better trajectory and better setting of control time T. In the end, the algorithm can find a T that is close to real T_{QSL}. Compared with common ways that implement a QOC algorithm and vary T step by step to explore the 'collapse' time, our method iteratively updates and could drastically save the computation time.

4. NUMERICAL RESULTS

In this section, we implement the proposed method for the control design of the Toffoli gate. The Toffoloi gate is a three-bit gate, also named the Controlled-Controlled-NOT gate. When the first and second qubits are at $|11\rangle$, the Toffoloi gate flips the third state; otherwise, it leaves the third qubit unchanged. The CCNOT gate has the form:

$$\mathrm{W_{CCNOT}} = \begin{pmatrix} 1&0&0&0&0&0&0&0 \\ 0&1&0&0&0&0&0&0 \\ 0&0&1&0&0&0&0&0 \\ 0&0&0&1&0&0&0&0 \\ 0&0&0&0&1&0&0&0 \\ 0&0&0&0&0&1&0&0 \\ 0&0&0&0&0&0&0&1 \\ 0&0&0&0&0&0&1&0 \end{pmatrix}. \tag{15}$$

The system Hamiltonian is considered as (see Ge and Wu (2021)):

$$\begin{aligned} H(t) = & J_{12}\sigma_1^z\sigma_2^z + J_{23}\sigma_2^z\sigma_3^z \\ & + \sum_{k=1}^{3} [u_{kx}(t)\sigma_k^x + u_{ky}(t)\sigma_k^y]. \end{aligned} \tag{16}$$

Here, σ^x, σ^y and σ^z are Pauli matrices:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma_y = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma_z = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \tag{17}$$

and $\sigma_k^{x,y,z}$ denotes the Pauli operator associated with the k-th qubit, e.g., $\sigma_2^x = \mathbb{I}^2 \otimes \sigma^x \otimes \mathbb{I}^2$. $J_{12} = J_{23} = 1$ MHz are the qubit-qubit coupling strengths. The $\{u_{kx}(t), u_{ky}(t), k = 1, 2, 3\}$ are control fields along the x-axis and y-axis and are discretized with duration time $dt = 0.1\mu s$.

In the simulations, we first compare three categories of cost design, following

μ	Cost function design	$J(\boldsymbol{u})$
1	Infidelity	$\sum_{i=1}^{N}[1 - \mathcal{F}(U_j, W)]$
2	Sequared error (L2 norm)	$\sum_{i=1}^{N}[\|\|\mathrm{Re}(W - U_j)\|\|_{\mathrm{L2}}^2 + \|\|\mathrm{Im}(W - U_j)\|\|_{\mathrm{L2}}^2]$
3	Logarithmic infidelity	$\sum_{i=1}^{N}[\mathcal{L}(U_j, W)]$

The three experiments follow the same hyper-parameter settings. We set the pulse length as KP $= 100$. The guessed control parameters are randomly initialized. The 'StepTolerance' for the optimizer (SQP) is 10^{-4}. Fig. 4 displays the cost training curves and control trajectories, namely, the achieved fidelity $\mathcal{F}$ and logarithmic infidelity $\mathcal{L}$ versus the control pulse indexes. As seen from the trajector curves, optimizations based on the $1 - \mathcal{F}$ cost and L2 cost fail to achieve the target gate since the fidelities are always less than 0.99, while the optimization based on $\mathcal{L}$ achieves gate error less than 10^{-4} at time $t = 4.5\mu$s, which indicates a better setting of the protocol duration time.

As mentioned above, solving the trajectory optimization problem defined in (14) could achieve the target gate and drastically reduce the control time. To obtain an optimal setting of T that is close to QSL, we can run the proposed approach to optimize the trajectory iteratively. We run the algorithm ten times and present the learning curves for optimal control time, i.e., the optimal control time denoted by the trajectory versus the iteration index within a single algorithm run. As shown in Fig. 5, starting from an initial guess of $T = 10\mu$s, the best-estimated control time T is 3.1 μs, while the worst estimation is $T = 3.4\mu$s. To verify the real QSL time, we implement the widely used QOC algorithm GRAPE (see Khaneja et al. (2005)), varying the control duration time T from 10 μs to 0.1 μs. For the GRAPE algorithm, the learning rate is 0.02 and the maximum iteration number is 50000.

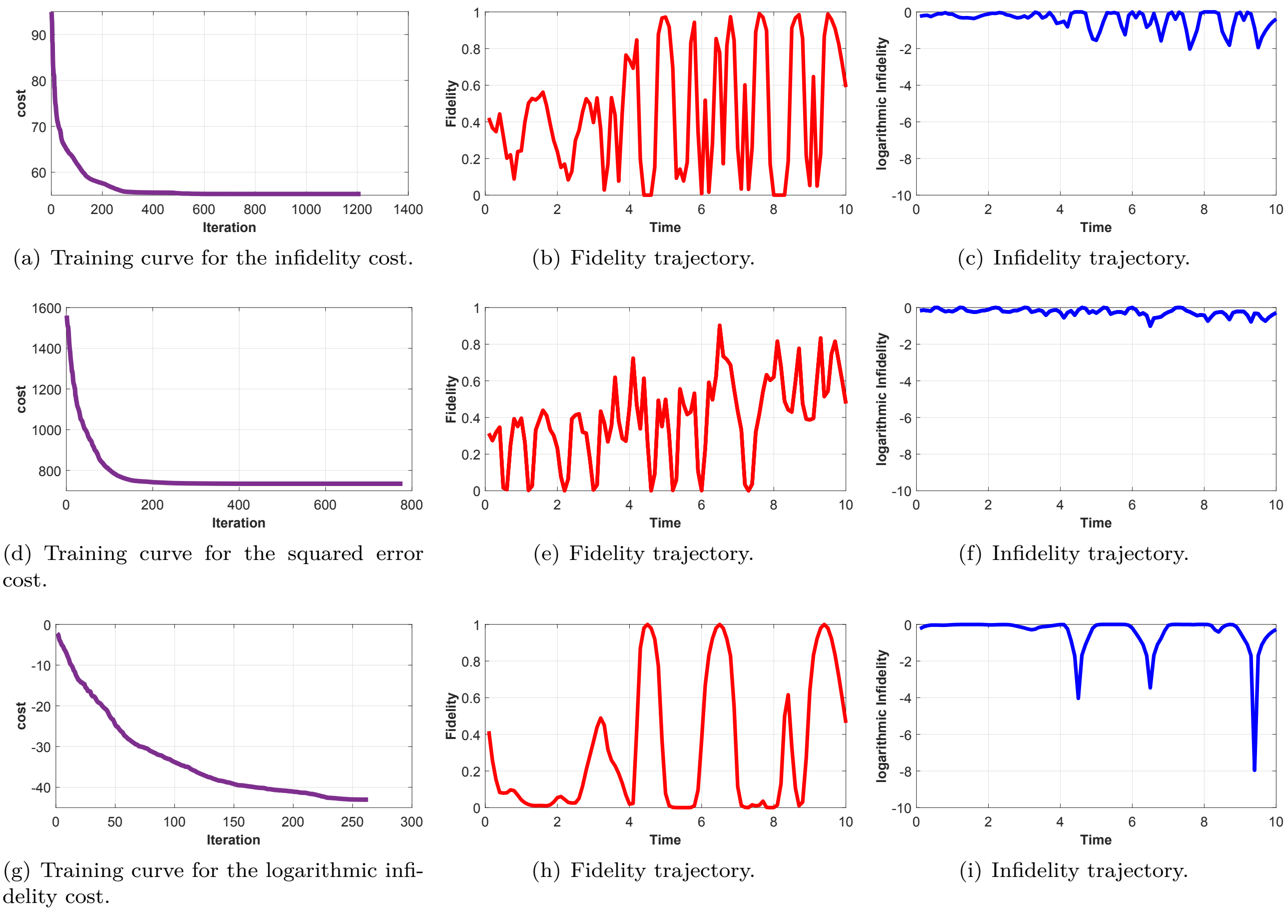

(a) Training curve for the infidelity cost.
(b) Fidelity trajectory.
(c) Infidelity trajectory.
(d) Training curve for the squared error cost.
(e) Fidelity trajectory.
(f) Infidelity trajectory.
(g) Training curve for the logarithmic infidelity cost.
(h) Fidelity trajectory.
(i) Infidelity trajectory.

Fig. 4. The comparison of trajectory optimization using different cost design.

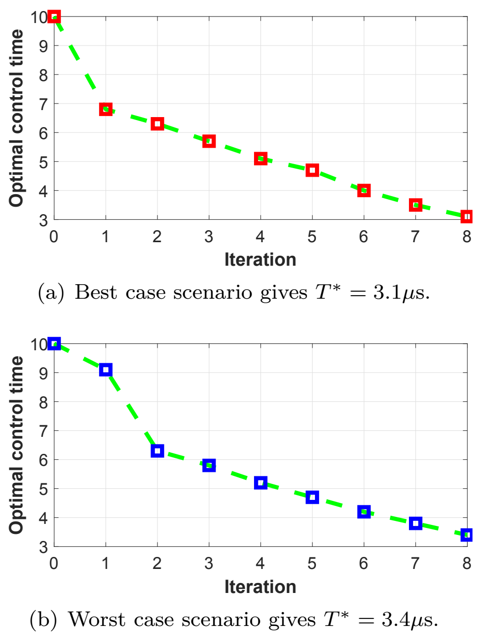

(a) Best case scenario gives $T^* = 3.1\mu s$.

(b) Worst case scenario gives $T^* = 3.4\mu s$.

Fig. 5. Iteration curves for exploring the optimal control time, where an iteration refers to solving a trajectory optimization problem defined in (14).

We also implement the sequential quadratic programming algorithm, varying the control duration time T from 6 μs to 0.1 μs. We run GRAPE and SQP three times for each T value. The threshold for the logarithmic infidelity to determine the QSL time is -4. As shown in Fig. 6, both the GRAPE algorithm and SQP denote the optimal control time $T^* = 3.1\mu s$.

Fig. 7 plots the optimal control trajectory obtained by our method. Our proposed method can achieve a QSL time estimation that is pretty close to the real T_{QSL}, while guaranteeing high control precision $\mathcal{L} \leq -8$. The simulation experiments are implemented with Matlab and run on a 6-core 3.0 GHz CPU with 16 GB memory. The computation time for the ten trial runs varies from 10 to 27 minutes, and the average consuming time is 20 minutes. One run of the GRAPE algorithm takes 147 minutes. Hence, our method is more computational-friendly and can stand out for large-scale quantum systems since the dimension grows exponentially with the number of qubits.

5. CONCLUSION

In this paper, we utilized the trajectory optimization technique to search for fast quantum control fields. The core idea is to minimize the cumulative intermediate gate error to incentivize early achievement of the desired gate. The proposed method could approach the QSL by optimizing the trajectory iteratively. Numerical results demonstrate that the proposed scheme can efficiently explore a proper control time setting and guarantee high gate precision. Compared with the widely used GRAPE algorithm, our scheme can achieve a very close QSL time

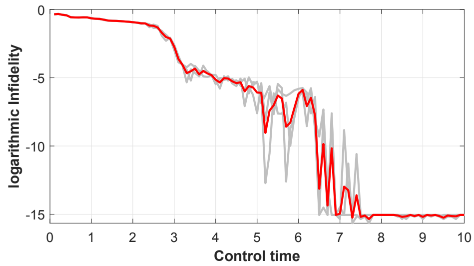

(a) The landscape curve obtained by GRAPE.

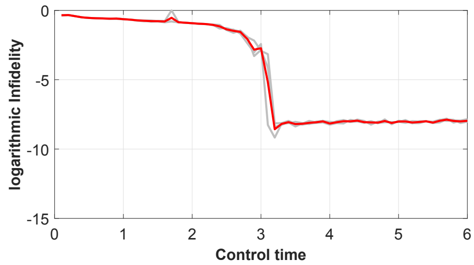

(b) The landscape curve obtained by SQP.

Fig. 6. The landscape curve, i.e., logarithm gate infidelity versus control time. The red lines plot the average data among different runs (plotted by gray lines).

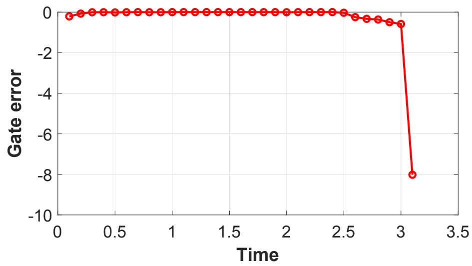

Fig. 7. Optimal control trajectory via the trajectory optimization, where the gate error is referred to the logarithmic infidelity. The control time $T = 3.1\mu s$.

estimate. The typical way to estimate QSL utilizing the optimal algorithm must vary the protocol duration time for many trial runs. Our scheme is less time-consuming and stands out for large dimensional quantum control optimization tasks.

REFERENCES

Boggs, P.T. and Tolle, J.W. (1995). Sequential quadratic programming. *Acta numerica*, 4, 1–51.

Boscain, U., Grönberg, F., Long, R., and Rabitz, H. (2014). Minimal time trajectories for two-level quantum systems with two bounded controls. *Journal of Mathematical Physics*, 55(6), 062106.

Bukov, M., Day, A.G., Weinberg, P., Polkovnikov, A., Mehta, P., and Sels, D. (2018). Broken symmetry in a two-qubit quantum control landscape. *Physical Review A*, 97(5), 052114.

Caneva, T., Murphy, M., Calarco, T., Fazio, R., Montangero, S., Giovannetti, V., and Santoro, G.E. (2009). Optimal control at the quantum speed limit. *Physical Review Letters*, 103(24), 240501.

Deffner, S. and Campbell, S. (2017). Quantum speed limits: from heisenberg's uncertainty principle to optimal quantum control. *Journal of Physics A: Mathematical and Theoretical*, 50(45), 453001.

Dong, D. and Petersen, I.R. (2022). Quantum estimation, control and learning: opportunities and challenges. *Annual Reviews in Control*, 54, 243–251.

Dong, D. and Petersen, I.R. (2023). Learning and robust control in quantum technology. Springer Nature, Switzerland.

Dong, D., Shu, C.C., Chen, J., Xing, X., Ma, H., Guo, Y., and Rabitz, H. (2020). Learning control of quantum systems using frequency-domain optimization algorithms. *IEEE Transactions on Control Systems Technology*, 29(4), 1791–1798.

Fan, L.B., Shu, C.C., Dong, D., He, J., Henriksen, N.E., and Nori, F. (2023). Quantum coherent control of a single molecular-polariton rotation. *Physical Review Letters*, 130(4), 043604.

Ge, X. and Wu, R.B. (2021). Risk-sensitive optimization for robust quantum controls. *Physical Review A*, 104(1), 012422.

Gill, P.E. and Wong, E. (2011). Sequential quadratic programming methods. In *Mixed integer nonlinear programming*, 147–224. Springer.

Howell, T.A., Jackson, B.E., and Manchester, Z. (2019). Altro: A fast solver for constrained trajectory optimization. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 7674–7679. IEEE.

Jones, P.J. and Kok, P. (2010). Geometric derivation of the quantum speed limit. *Physical Review A*, 82(2), 022107.

Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbrüggen, T., and Glaser, S.J. (2005). Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2), 296–305.

Leung, N., Abdelhafez, M., Koch, J., and Schuster, D. (2017). Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Physical Review A*, 95(4), 042318.

Mandelstam, L. and Tamm, I. (1945). The uncertainty relation between energy and time in non-relativistic quantum mechanics. *J. Phys.*, (9), 115–123.

Margolus, N. and Levitin, L.B. (1998). The maximum speed of dynamical evolution. *Physica D: Nonlinear Phenomena*, 120(1-2), 188–195.

Nielsen, M.A. and Chuang, I.L. (2001). Quantum computation and quantum information. *Phys. Today*, 54(2), 60.

Palao, J.P. and Kosloff, R. (2002). Quantum computing by an optimal control algorithm for unitary transformations. *Physical Review Letters*, 89(18), 188301.

Propson, T., Jackson, B.E., Koch, J., Manchester, Z., and Schuster, D.I. (2022). Robust quantum optimal control with trajectory optimization. *Physical Review Applied*, 17(1), 014036.

Wu, C., Qi, B., Chen, C., and Dong, D. (2016). Robust learning control design for quantum unitary transformations. *IEEE Transactions on Cybernetics*, 47(12), 4405–4417.

Zahedinejad, E., Schirmer, S., and Sanders, B.C. (2014). Evolutionary algorithms for hard quantum control. *Physical Review A*, 90(3), 032310.