



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

A practical approach to determine minimal quantum gate durations using amplitude-bounded quantum control

S. Guenther, N. A. Petersson

July 21, 2023

AVS Quantum Science

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

A PRACTICAL APPROACH TO DETERMINE MINIMAL QUANTUM GATE DURATIONS USING AMPLITUDE-BOUNDED QUANTUM CONTROLS

Stefanie Günther*, **N. Anders Petersson**
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory, Livermore CA 94550
{guenther5, petersson1}@llnl.gov

November 8, 2023

ABSTRACT

We present an iterative scheme to estimate the minimal duration in which a quantum gate can be realized while satisfying hardware constraints on the control pulse amplitudes. The scheme performs a sequence of unconstrained numerical optimal control cycles that each minimize the gate fidelity for a given gate duration alongside an additional penalty term for the control pulse amplitudes. After each cycle, the gate duration is adjusted based on the inverse of the resulting maximum control pulse amplitudes, by re-scaling the dynamics to a new duration where control pulses satisfy the amplitude constraints. Those scaled controls then serve as an initial guess for the next unconstrained optimal control cycle, using the adjusted gate duration. We provide multiple numerical examples that each demonstrate fast convergence of the scheme towards a gate duration that is close to the quantum speed limit, given the control pulse amplitude bound. The proposed technique is agnostic to the underlying system and control Hamiltonian models, as well as the target unitary gate operation, making the time-scaling iteration an easy to implement and practically useful scheme for reducing the durations of quantum gate operations.

Keywords quantum optimal control · pulse-level control · minimal quantum gate duration, amplitude-bounded quantum control

1 Introduction

Quantum optimal control utilizes numerical optimization tools to design pulses that drive quantum devices. It's potential to improve fundamental operations in Quantum Information Science has by now been demonstrated on various applications, e.g., for quantum state preparation [19], applications to quantum error correction [24, 6], and the realization of logical quantum gate operations [21, 17, 4], see for example [9, 14] for an overview. In the current Noisy Intermediate-Scale Quantum (NISQ) era, it is desirable to design control pulses with minimal gate duration, such that the quantum operation can be performed before decoherent processes collapse the quantum information into a classical state. However, finding the minimal duration in which such an operation is realizable is challenging. When control pulse amplitudes are unlimited, various lower bounds known as intrinsic *quantum speed limits* (QSL) have been established for the time-scale over which a quantum system can evolve, see, e.g., the overview in [5]. Such speed limits can be computed analytically for simple model problems [11, 12, 13], and can be estimated numerically for more complex systems by performing multiple optimal control cycles, sweeping over a range of gate durations. Indeed it has been demonstrated on specific examples that such limits align with the shortest durations for which numerical optimal control techniques converge, when control pulse amplitudes are unlimited [11, 3, 2]. In practice, however, control pulses are generally subjected to hardware constraints, most commonly given in terms of a maximum drive strength as defined by hardware waveform generating devices. When hardware constraints bound the control pulse amplitude, the

*Corresponding author

time-scale required to realize a quantum operation can be significantly slower than the theoretical (intrinsic) QSL. For unitary operations in the closed system setting, an extrinsic speed limit that takes the maximum control pulse amplitude into account has been derived in [1]. While the lower bound matches the minimum gate duration remarkably well for small single-qubit cases, and the extension in [15] to general N-level system demonstrates qualitatively reasonable scaling properties for larger numbers of qubits, the lower bound is not sharp in the sense that the actual gate duration is no longer precisely estimated as system sizes increase. Currently, a general technique to find minimal gate durations when control pulse amplitudes are bounded by hardware constraints is not available. Instead, it is common practice to determine the minimum gate duration by trial and error through multiple optimal control cycles, each using a different pulse duration and a different initial guess for the control pulse. Henceforth, the latter approach will be referred to as the brute-force method.

In [22], a re-seeding technique was proposed that adjusts gate durations based on previous success or failure of the optimal control iterations, combined with a bisection of the resulting gate durations. However, imposing parameter bounds during the optimization updates often leads to non-convex optimization landscapes [7] and slow or stagnant optimization progress, such that optimization cycles need to be performed multiple times for various different initial guesses.

In this paper, we propose an alternative scheme to numerically find minimal time-scales in which a unitary operation can be realized under amplitude-constrained control fields, by re-scaling the gate duration based on the optimized control pulse amplitude. Instead of directly imposing box constraints on the control pulse, we include a penalty term in the objective function that minimizes the control pulse energy during the optimization. After solving the unconstrained and penalized optimization problem for a given gate duration, we re-scale the dynamics to a new duration where control pulses satisfy the amplitude constraints. Those stretched (or squeezed) control pulses then serve as an initial guess for the next optimization cycle for the new gate duration. We present numerical evidence for system of up to three coupled qudits where this iterative scheme rapidly converges to a final gate duration close to the QSL, yielding control pulses that satisfy hardware-specific amplitude bounds while minimizing the gate time duration.

While the approach presented here is agnostic to the underlying Hamiltonian model and optimization strategy, in Section 2 we present the system and control Hamiltonian models that are used in our numerical experiments. Section 3 describes the proposed iterative scheme for finding minimal gate durations and corresponding amplitude-constrained control pulses. Section 4 presents the numerical results, and conclusions are drawn in Section 5.

2 The quantum optimal control problem

Given a unitary target gate V^{target} that represents a logical quantum operation, the goal of quantum optimal control is to design optimal pulses that drive any initial quantum state $\phi(0)$ at time $t = 0$ to the unitarily transformed target state $\phi(T) = V^{\text{target}}\phi(0)$ at time $T > 0$. To that end, gradient-based optimization methods are typically applied to minimize a cost function that measures the distance between the target and the driven evolution in terms of the gate infidelity,

$$\min J_{\text{cost}}(U(T)) = 1 - \left| \frac{1}{N} \text{Tr}(U^\dagger(T)V^{\text{target}}) \right|^2, \quad (1)$$

where the columns of $U(T)$ describe the evolved quantum state at final time T for a basis of initial states $U(0) = [e_0, \dots, e_{N-1}]$, where e_i are the canonical basis vectors. We here consider closed quantum systems modeled by Schrodinger's equation, such that the evolved dynamics satisfy

$$\frac{dU(t)}{dt} = -iH(t)U(t), \quad 0 \leq t \leq T, \quad U(0) = I. \quad (2)$$

The Hamiltonian,

$$H(t) = H_{\text{sys}} + H_c(t), \quad (3)$$

decomposes into a time-independent system part, H_{sys} , and a time-varying control part, $H_c(t)$. In the laboratory-frame of reference, the control Hamiltonian is given by $\sum_q f_q(t)(a_q + a_q^\dagger)$, where the real-valued control pulse satisfies $f_q(t) = 2 \text{Re}\{c_q(t)e^{it\omega^d}\}$, where ω^d is the drive frequency. To slow down the time scales of the state dynamics and the control pulse, we apply the rotating wave approximation (RWA) and select the frequency of rotation to equal the drive frequency, $\omega^{\text{rot}} = \omega^d$. In the rotating frame, we consider a general system Hamiltonian for modeling $Q \geq 1$

coupled superconducting qudits,

$$H_{sys} = \sum_{q=1}^Q (\omega_q - \omega^{rot}) a_q^\dagger a_q - \frac{\xi_q}{2} a_q^\dagger a_q^\dagger a_q a_q + \sum_{p>q} J_{pq} (a_p^\dagger a_q + a_p a_q^\dagger), \quad (4)$$

where ω_q denotes the 0-1 transition frequency of qudit q , ξ_q is the self-Kerr coefficient, and J_{pq} denotes the dipole-dipole coupling coefficient between qudits p and q . Further, a_q (a_q^\dagger) denote the lowering (raising) operator for qudit q . The action of external control pulses on qudit q in the rotating frame is given by the control Hamiltonian

$$H_c(t) = \sum_{q=1}^Q c_q(t) a_q + c_q^*(t) a_q^\dagger. \quad (5)$$

where $c_q(t) = p_q(t) + i q_q(t)$ denote the control pulse for qubit q in the rotational frame and $c_q^*(t)$ denotes its complex conjugate. For notational simplicity, we now drop the index q denoting each subsystem (qudit), noting that it is straightforward to apply the proposed method to systems with multiple qudits.

To parameterize the control pulse $c(t)$ in the rotating frame, we here choose B-spline basis functions:

$$c(t) = \sum_{s=1}^{N_s} \alpha_s B_s(t), \quad (6)$$

where $\alpha_s = \alpha_s^{real} + i \alpha_s^{imag} \in \mathbb{C}$ are the control parameters that are being optimized for; $B_s(t)$ are fixed, quadratic, cardinal B-spline basis functions [18] with local support in time. In the following we will refer to the set of control parameters $\vec{\alpha} := \{\alpha_s^{real}, \alpha_s^{imag}\}_{s=1}^{N_s}$ as the control vector. It determines the control pulse $c = c(t; \vec{\alpha})$ in the control Hamiltonian and hence the solution operator $U = U(t; \vec{\alpha})$ of Schroedinger's equation (2).² Parameterizing the control functions using quadratic B-splines provides a compact alternative to discretizing the control functions on the same time step as the underlying dynamical system. The cardinal quadratic B-spline basis functions result in control pulses that are continuously differentiable, see Appendix B. Therefore, they can readily be applied to the control hardware, without the need for smoothing or filtering a stair step control pulse or a bang-bang ansatz. While B-splines have many attractive properties [23], we note that the time-scaling algorithm presented in the following section is not restricted to any particular class of pulse parameterization. The only requirement is that the parameterization must allow the pulses to readily be stretched or compressed to change the pulse duration.

3 Time-scaling iteration to find minimal gate duration under pulse amplitude constraints

When optimal controls are applied in practice, the pulses are typically subject to amplitude bounds as specified by the pulse-generating hardware, such as an Arbitrary Waveform Generator (AWG). Typically those constraints are given in terms of a maximum drive strength b_{max} such that optimized pulses need to satisfy

$$\max_{t \in [0, T]} |c(t)| \leq b_{max}. \quad (7)$$

While the amplitude bound varies depending on the specific AWG, we here consider a maximum amplitude bound of $b_{max}/2\pi = 40$ MHz for the numerical results presented below.³ Standard optimal control approaches incorporate such constraints into the optimization procedure by imposing box-constraints on the control parameter vector, e.g., by projecting the gradient-based update steps onto the given bounds, or by using interior point or barrier methods that augment the objective function and enforce the constraint through penalty terms [16]. However, when control parameters hit the bounds, optimization convergence can deteriorate due to non-convex optimization landscapes [20]. As a result, the constrained optimization problem is often much harder to solve than the unconstrained one. In practice, it is often required to repeat multiple optimization cycles, each for a different randomized initial guesses of the control parameters and different pulse durations, to increase the chances of finding a global optima.

In this section, we propose an iterative scheme to automatically adjust the gate duration based on the energy norm of the optimized control pulse. Instead of explicitly imposing box-constraints on the control parameter vector during the

²In the following, the dependence on $\vec{\alpha}$ will be suppressed for improved readability.

³Because we have scaled Schroedinger's equation to make $\hbar = 1$, the unit of energy becomes angular frequency. As a result, the unit of the control pulse amplitude $|c(t)|$ is also angular frequency.

optimization procedure, we add a penalty term to the objective function that aims to minimize the control pulse energy. The iterative scheme performs a sequence of such unconstrained optimal control iterations, each using an updated gate duration that is selected based on the control pulse energy in the previous iteration.

In each iteration k of the proposed scheme, we solve the following augmented and unconstrained optimal control problem for a given gate duration T_k ,

$$\min_{\vec{\alpha}} J_{\text{infid}}(U(T_k; \vec{\alpha})) + \gamma J_{\text{energy}}(c(\cdot; \vec{\alpha})). \quad (8)$$

Here, $\gamma > 0$ is a parameter that weights the contribution of the control energy term,

$$J_{\text{energy}}(c(\cdot; \vec{\alpha})) = \frac{1}{T_k} \int_0^{T_k} |c(t; \vec{\alpha})|^2 dt, \quad (9)$$

relative to the infidelity, J_{infid} , in the objective function.

Performing an optimization cycle that minimizes the weighted sum of the infidelity at the final time (1) and the control pulse energy (9) yields control pulses with the smallest energy norm possible, while realizing the given target unitary at time T_k . We can assume that a gradient-based optimal control technique can solve this unconstrained optimization problem, given that the weight γ is chosen appropriately. However, since box-constraints are not imposed during the optimization, the resulting control pulses, while being minimal in terms of their energy, might exceed the given hardware bounds b_{max} . Figure 1 demonstrates the relation between the pulse duration T and the resulting maximum control amplitude

$$c_{max} = \max_{t \in [0, T]} |c_*(t)|, \quad (10)$$

for optimized unconstrained control pulses $c_*(t)$, on a 4-level qudit QFT gate (problem specifications are given in Section 4). Here, each optimization proceeded until the gate infidelity dropped below the threshold 10^{-4} . Note that the maximum control amplitude approximately scales as $1/T$.

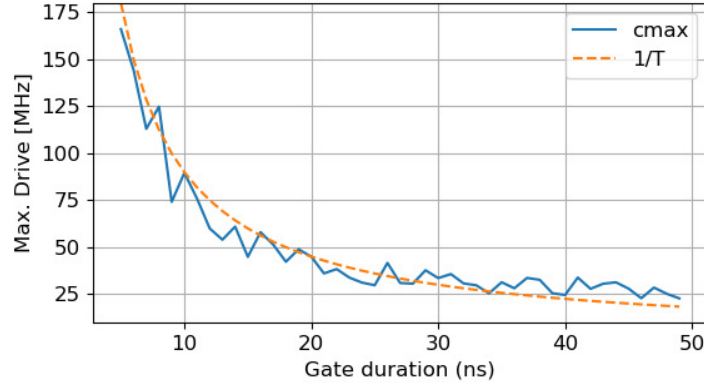


Figure 1: Unconstrained minimization of the infidelity and the control energy (8), illustrating the relation between the maximum control amplitude and the gate duration T , compared with an $1/T$ scaling. Here, the target unitary is a QFT gate on a 4-level qudit.

We make use of the scaling $c_{max} \sim 1/T$ to find a minimal gate duration for given amplitude constraints by scaling the current gate duration T_k by the ratio between the current maximum control amplitude resulting from solving the unconstrained optimization problem (8), and the given hardware constraint b_{max} , i.e., we propose the update scheme

$$T_{k+1} = sT_k, \quad \text{where} \quad s = \frac{c_{max}}{b_{max}}. \quad (11)$$

Consider the scaled time variable

$$\tau(t) = st \quad \Rightarrow \quad \tau \in [0, T_{k+1}] \quad \text{for} \quad t \in [0, T_k], \quad (12)$$

and corresponding scaled evolution $\tilde{U}(\tau) := U(\tau/s) = U(t)$. The scaled unitary $\tilde{U}(\tau)$ satisfies the scaled Schroedinger equation

$$\dot{\tilde{U}}(\tau) = \frac{dU(\tau/s)}{d\tau} = \frac{1}{s}\dot{U}(\tau/s) = -\frac{i}{s}H(\tau/s)\tilde{U}(\tau) \quad (13)$$

for $\tau \in (0, T_{k+1})$,

and, hence, $\tilde{U}(\tau)$ evolves under the scaled Hamiltonian

$$\tilde{H}(\tau) := \frac{1}{s}H(\tau/s) = \frac{1}{s}H_{sys} + \tilde{c}(\tau)a + \tilde{c}(\tau)^*a^\dagger \quad (14)$$

with the control pulse

$$\tilde{c}(\tau) := \frac{1}{s}c(\tau/s) \quad (15)$$

The scaled pulse stretches (if $s > 1$) or compresses (if $s < 1$) the original pulse into the time domain $[0, T_{k+1}]$ while preserving its integral,

$$\int_0^{T_k} |c(t)| dt = \int_0^{T_{k+1}} |\tilde{c}(\tau)| d\tau. \quad (16)$$

Preserving the control pulse integral can be further motivated by solving the control problem analytically for a model problem, see Appendix A. By parameterizing the control functions in terms of B-spline basis functions as in (6), the re-scaled controls can be easily computed by scaling the parameter vector $\tilde{\alpha} = s^{-1}\alpha$, and re-evaluating the B-spline basis functions at the scaled time, hence stretching or compressing their envelopes, see Figure 2. By design (choice of




Figure 2: Re-scaling control pulses: By choosing the scaling parameter s , the maximum drive strength of the stretched ($s > 1$) or compressed ($s < 1$) control pulse $\tilde{c}(t) = \frac{1}{s}c(t/s)$ can be adjusted.

the scaling parameter s in (11)), the scaled control pulses satisfy the hardware bounds exactly with

$$\max_{\tau \in [0, T_{k+1}]} |\tilde{c}(\tau)| = b_{max} \quad (17)$$

Further, if the dynamics had obeyed the scaled system Hamiltonian $\frac{1}{s}H_{sys}$, the scaled controls would have realized the target gate at time T_{k+1} . However, because the actual system Hamiltonian is H_{sys} , we only use the scaled control pulse as an initial guess to solve the optimization problem (8) for the updated gate duration, T_{k+1} . This procedure is repeated until the scaling factor remains constant at $s = 1$, indicating that the minimal gate duration T_* has been found and the corresponding controls pulse $c_*(t)$ satisfies the hardware bounds $|c_*(t)| \leq b_{max}$. Note that as $s \rightarrow 1$, the scaled control pulses provide better and better initial guesses for solving the optimization problem. As a result, fewer and fewer optimization iterations are required.

In the proposed time-scaling scheme, the gate duration increases whenever the control pulse exceeds the amplitude bound, and decreases whenever the pulse is below the same bound. However, many iterations of the time-scaling scheme would be needed to precisely determine the minimal gate duration such that the maximum control amplitude equals the prescribed bound. To obtain a practically useful stopping criteria, we therefore introduce a *range* of acceptable maximum control amplitudes $[b_{max} - \delta_b, b_{max}]$, for $\delta_b > 0$. The time-scaling iteration then terminates if $c_{max} \in [b_{max} - \delta_b, b_{max}]$; otherwise the gate duration and control coefficients are re-scaled as described above. The resulting time-scaling iteration is summarized in Algorithm 1.

Algorithm 1 Minimizing the gate duration

Require: Specify the range of acceptable maximum control amplitudes: $[b_{max} - \delta_b, b_{max}]$.

Require: Pick an initial gate duration $T_0 > 0$ and select an initial guess for the control vector $\vec{\alpha}_0$.

1: **for** $k = 0, 1, 2, \dots$ **do**

2: Given $\vec{\alpha}_k$, solve the optimization problem

$$\vec{\alpha}_* = \arg \min_{\vec{\alpha}} J_{\text{infid}}(U(T_k; \vec{\alpha})) + \gamma J_{\text{energy}}(c(\cdot; \vec{\alpha}))$$

such that $U(T_k; \vec{\alpha})$ solves $\dot{U}(t; \vec{\alpha}) = -iH(t; \vec{\alpha})U(t; \vec{\alpha})$ with $U(0; \vec{\alpha}) = I$.

3: **if** $\max_{t \in [0, T_k]} |c(t, \vec{\alpha}_*)| \in [b_{max} - \delta_b, b_{max}]$ **then**

4: **Success!**

5: **return** $T_k, \vec{\alpha}_*$

6: **else**

7: Update the gate duration and the initial guess for next control vector:

$$T_{k+1} = sT_k, \quad \text{and} \quad \vec{\alpha}_{k+1} = \frac{1}{s}\vec{\alpha}_*,$$

$$\text{where } s = \frac{\max_{t \in [0, T_k]} |c(t; \vec{\alpha}_*)|}{b_{max}}$$

8: **end if**

9: **end for**

4 Numerical results

In this section, we verify on multiple test cases that the iterative time-scaling scheme yields gate durations that are close to the smallest possible time in which the target gate can be realized, when hardware constraints are in place on the control pulse amplitude. We here choose a maximum allowable amplitude bound of $b_{max}/2\pi = 40$ MHz in the rotating frame, and accept the resulting control pulses (and final times) whenever the maximum control pulse amplitudes are within the range of $[35, 40]$ MHz. The test cases are described in Table 1, consisting of two single-qudit cases with 3 and 4 energy levels, one two-qubit case with 2 energy levels per qubit, and two three-qubit cases with 2 energy levels each. The qubits are coupled to their nearest neighbor(s) in a one-dimensional chain with dipole-dipole coupling strength $J = 5$ MHz. System frequencies for each test case are shown in Table 2.

We solve the inner optimization problem in Algorithm 1 using L-BFGS as implemented in the quantum control software Quandary [10]. In order to ensure that the inner optimization cycle generates control pulses with minimal energy norm, we choose the penalty coefficient to be $\gamma = 1$. While γ can be considered a tunable parameter in the algorithm, we motivate this value of γ through a numerical experiment. Here we study the effect of varying γ on the maximum control drive amplitude and infidelity that results from minimizing the objective (8). Because the optimization minimizes the sum of two terms, it is to be expected that increasing γ would put more emphasis on minimizing the control energy and less emphasis on the infidelity. Indeed, the results shown in Fig. 3 verify the expected behavior. In the numerical tests presented below we aim the final infidelity to be less than 10^{-4} . In the case tested here, $\gamma = 1$ meets the criteria for the infidelity, while also significantly reducing the maximum control drive amplitude.

In order to ensure that an optimal point with minimal energy norm has been reached, we note that it is important to base the optimization stopping criterion on the norm of the gradient, rather than on a small gate infidelity. In the following, we set the gradient threshold to be 10^{-5} . To stabilize the inner optimization convergence, we further add a Tikhonov regularization term $\gamma_1 \|\vec{\alpha}\|^2$, with parameter $\gamma_1 = 10^{-2}$. The purpose of the Tikhonov term is to regularize the optimization problem by adding a convex term that shifts the eigenvalues of the Hessian of the objective function by a small constant factor γ_1 . This factor should be chosen small enough to not alter the optimization outcome [16]. We remark that, when the controls are parameterized by B-spline basis functions, the control energy term (9) can be interpreted as a weighted Tikhonov regularization term, see Appendix C. In addition, to encourage a state evolution where the populations vary slowly in time, we also penalize the second time-derivative of the state populations. To this end, we add a penalty term of the form

$$\frac{\gamma_2}{2T} \int_0^T \sum_{i,j} \left(\frac{\partial^2}{\partial t^2} |U_{i,j}(t)|^2 \right)^2 dt, \quad (18)$$

Name	System	Target gate
QFT ₄	one qudit, 4 levels	$\frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{pmatrix} \in \mathbb{C}^{4 \times 4}$
SWAP02	one qudit, 3 levels	$\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \in \mathbb{R}^{3 \times 3}$
CNOT	two qubits coupling 1 ↔ 2 at 5MHz	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \in \mathbb{R}^{4 \times 4}$
CCNOT	three-qubit chain coupling 1 ↔ 2 at 5MHz coupling 2 ↔ 3 at 5MHz	$\begin{pmatrix} 1 & & & & & & & \\ & \ddots & & & & & & \\ & & 1 & & & & & \\ & & & 0 & 1 & & & \\ & & & 1 & 0 & & & \end{pmatrix} \in \mathbb{R}^{8 \times 8}$
SWAP chain	three-qubit chain coupling 1 ↔ 2 at 5MHz coupling 2 ↔ 3 at 5MHz	$\begin{pmatrix} 1 & & & & & & & \\ & 0 & & & 1 & & & \\ & & 1 & & & & & \\ & & & 0 & & & 1 & \\ & 1 & & & 0 & & & \\ & & & & & 1 & & \\ & & & & 1 & & 0 & \\ & & & & & & & 1 \end{pmatrix} \in \mathbb{R}^{8 \times 8}$

Table 1: Test case specifications and target gates.

Name	$\dim(H)$	$\omega_q/2\pi$ [GHz]	$\xi_q/2\pi$ [GHz]	$J_{pq}/2\pi$ [GHz]	$\omega^{rot}/2\pi$ [GHz]	Δ_B [ns]
QFT ₄	$d = 4$	4.914	0.33	-	4.584	0.3
SWAP02	$d = 3$	5.12	0.34	-	4.78	0.3
CNOT	$d = 2 \times 2$	5.12, 5.06	0.34, 0.34	0.005	5.09	1.65
CCNOT	$d = 2 \times 2 \times 2$	5.18, 5.12, 5.06	0.34, 0.34, 0.34	0.005, 0.0, 0.005	5.12	1.65
SWAP chain	$d = 2 \times 2 \times 2$	5.18, 5.12, 5.06	0.34, 0.34, 0.34	0.005, 0.0, 0.005	5.12	1.65

Table 2: System parameters.

where $U_{i,j}$ denotes element (i,j) in the matrix U . Again, $\gamma_2 \geq 0$ should be chosen small enough to not dominate the infidelity and the control energy terms during the optimization. In the following, we use $\gamma_2 = 10^{-2}$. While the qualitative results of the minimal-gate-duration scheme are not expected to be sensitive to the precise value of γ_2 , a thorough analysis of this term remains to be done.

During the initialization phase of Algorithm 1 ($k = 0$), we assign the elements of the initial control vector from a uniform probability distribution $\vec{\alpha}_0 \sim \mathcal{U}(-0.9b_{max}, 0.9b_{max})$.

For each test case, we deploy Algorithm 1 for a wide range of initial gate durations, T_0 . Figure 4 shows the convergence history of the intermediate gate durations T_k (left column), as well as the corresponding maximum drive amplitude of the optimized pulses (right column) at each iteration. Across all test cases, we observe that an optimal gate duration can be found within very few iterations (often only two, but never more than 8 iterations), meaning that only a small number of optimal control cycles are required to determine the minimal gate duration, for a given bound on the control amplitude. This provides a clear advantage over the standard (brute-force) approach of sweeping over a wide range of gate durations, which can easily require tens to hundreds of optimal control cycles to be completed. The resulting final gate times at the last cycle of the time-scaling scheme are indicated by gray bands in Figure 4, and are also provided in Table 3. The width of these bands is determined from the user-prescribed range of allowable maximum amplitudes




Figure 3: Maximum control drive amplitude and infidelity when optimizing for the QFT_4 gate for various values of the tunable parameter γ in the objective function (8). Here, the gate duration is fixed to $T = 20$ ns.

$[b_{max} - \delta_b, b_{max}]$, which serves as the stopping criterion for the minimal time iteration. Note that we here consider $b_{max} = 40\text{MHz}$ and $\delta_b = 5\text{MHz}$, allowing for a 12.5% deviation from the maximum bound. The variations of final minimal times as presented in Table 3 are of the same order. Importantly, we observe that the iterations converge to this band for a wide range of initial gate durations, T_0 .

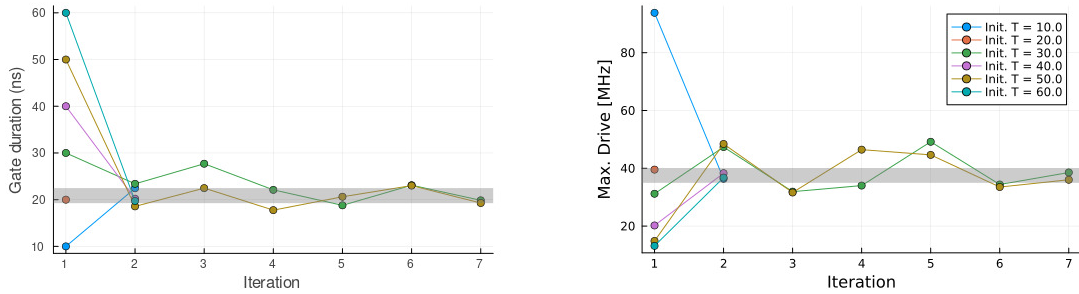
To evaluate how well the iterative time-scaling scheme recovers the actual minimal gate duration for a given amplitude bound, we compare the resulting gate durations with a brute-force approach that performs a sweep of optimizations for a wide range of (fixed) gate durations, while enforcing box constraints on the control vector using a projected L-BFGS method. Figure 5 shows statistics of optimized gate fidelities, each gathered over 10 different optimizations starting from random initial control vectors. Here, variations in the resulting gate fidelities can be attributed to deterioration of the optimization progress, related to the imposed box-constraints. Nevertheless, we compare the resulting minimal gate durations as observed in Figure 5 with those produced by Algorithm 1 as in Figure 4. Table 3 summarizes the resulting gate durations, demonstrating that the time-scaling scheme converges to final gate durations that are remarkably close to those obtained with the brute force method. The range of resulting gate durations as produced by Algorithm 1 is a direct consequence of the stopping criteria based on a prescribed range of allowed amplitudes $[b_{max} - \delta_b, b_{max}]$. The range in gate duration hence depends on the user-controlled parameter δ_b . In the last column, we compare the shortest gate durations with the ones obtained from the brute-force optimization approach. We observe a relative difference below 7% for all test cases.

Case	Algorithm 1	Brute force	Smallest rel. diff.
QFT ₄	19ns to 23ns	18ns to 21ns	5%
SWAP02	18ns to 23ns	18ns	0%
CNOT	68ns to 78ns	70ns	-2%
CCNOT	200ns to 225ns	190ns	5%
SWAP chain	205ns to 239ns	190ns	7%

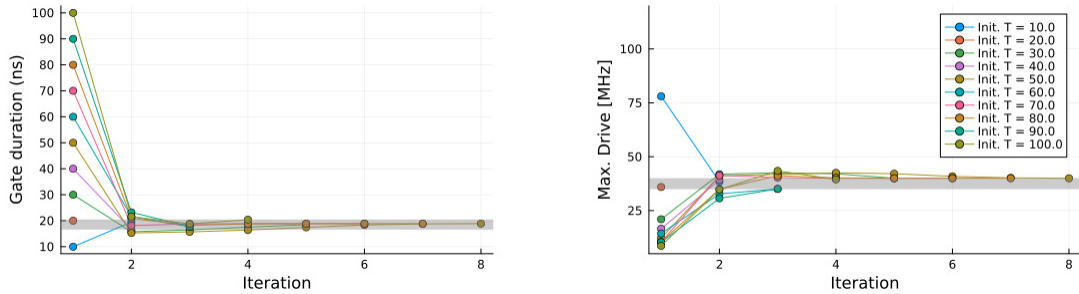
Table 3: Minimal gate durations with gate fidelity $\geq 99.9\%$, as achieved from the time-scaling scheme in Algorithm 1 (Figure 4), and brute-force optimization with box constraints on the control vector (Figure 5).

5 Conclusion


This paper presents a practical scheme for approximating the minimal duration for realizing a unitary gate while satisfying given amplitude bounds on the control pulses. The scheme performs a sequence of unconstrained optimal control cycles, each minimizing the gate fidelity alongside an additional penalty term for the control pulse amplitudes. After each cycle, the gate duration is scaled to a new value for which the scaled controls satisfy the amplitude bounds. We provide numerical evidence of convergence to a final gate duration that is close to the shortest achievable duration obtained from sweeping over a large number of gate durations. Updating the gate duration based on the previous control amplitudes converges to the minimal gate duration in a few optimization cycles, drastically reducing the computational costs for practical usage of time-optimal quantum control, which often requires tens to hundreds



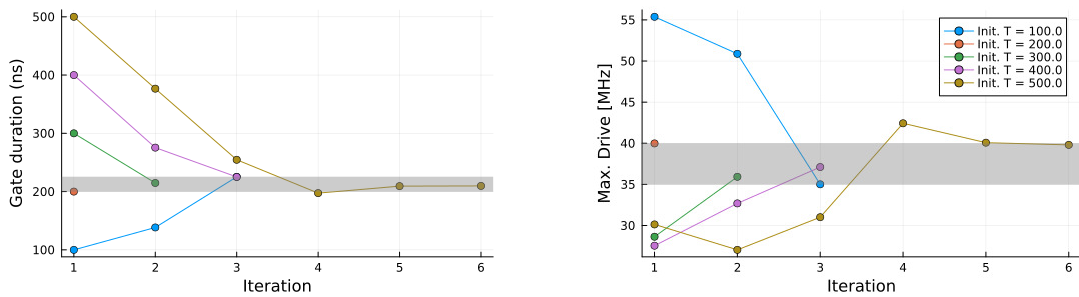
(a) QFT₄ gate on a single qubit.



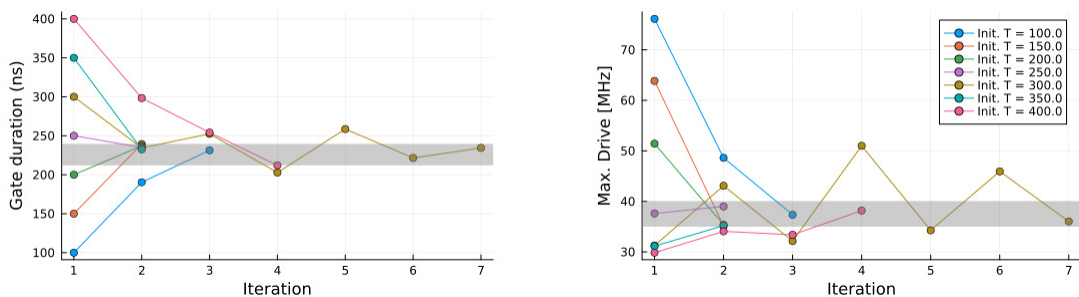
(b) SWAP02 gate on a single qubit



(c) CNOT gate, coupling $J_{12} = 5\text{MHz}$



(d) CCNOT (Toffoli) gate on a three-qubit chain



(e) SWAP qubit 1 and 3 on a three-qubit chain

Figure 4: Optimization history for the time-scaling algorithm (Alg. 1): Gate durations (left column) and maximum amplitudes (right column). The gray areas denote the resulting band of final gates times, and the prescribed range of acceptable amplitudes, respectively.

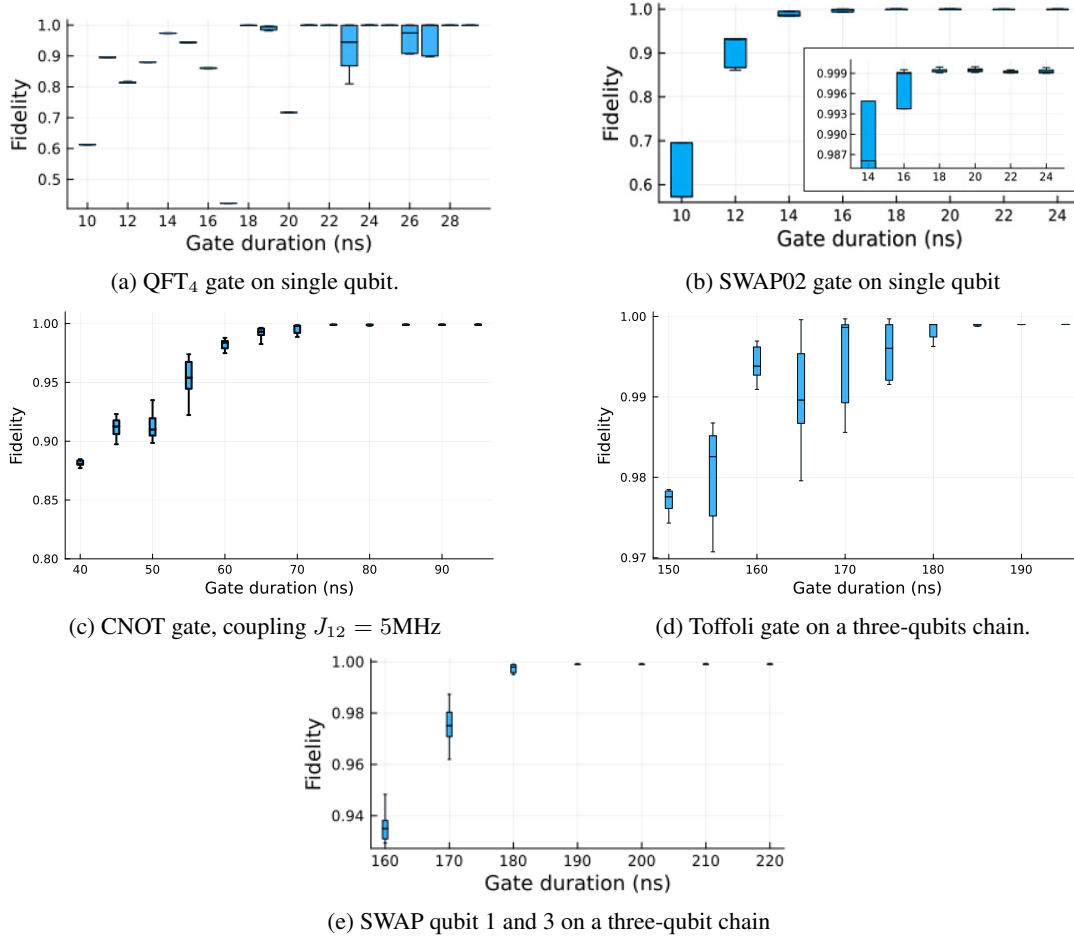


Figure 5: Estimating the quantum speed limit with the brute-force method: Optimized fidelities sweeping over various gate durations, while enforcing hardware constraints directly during the optimization process. For each gate duration, statistics are gathered by performing 10 optimization cycles starting from randomly sampled initial control vectors.

of optimization cycles to sweep over a range of gate durations in a trial and error fashion. We demonstrate that the proposed scheme converges for a wide range of initial gate durations, indicating that almost no prior knowledge of the minimal gate duration is required. The proposed technique is straightforward to be implemented on top of an existing quantum optimal control code, where we note that it is beneficial to parameterize the control pulses by basis functions such that the pulse amplitudes can easily be adjusted. Further, the proposed method is agnostic to the underlying system and control Hamiltonian models, as well as the target unitary gate operation, making the time-scaling iteration an easy to implement and practically useful scheme for reducing the durations of quantum gate operations.

As illustrated in Figure 1, the control pulse amplitude can be a non-monotonic function of the gate duration. We conjecture that this phenomena is the underlying cause of the non-monotonic convergence (over- and under-shooting the gate duration) of our algorithm, and we conjecture that convergence may be further improved by damping the updates of the duration. This will be subject of further investigations.

Acknowledgments

The authors would like to acknowledge Dr. Jonathan Dubois at Lawrence Livermore National Laboratory, CA, who provided expertise and partial funding resources during the preparation of this paper.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-JRNL-851967.

A The infidelity is invariant to the control pulse integral in a qubit model problem

To gain further insight into the essential properties of a control pulse, we consider the model problem of a driven two-level system in the rotating frame,

$$H(t) = c(t)a^\dagger + c^*(t)a, \quad c(t) : \mathbb{R} \mapsto \mathbb{C}. \quad (19)$$

We consider the displacement transformation [8],

$$U(t) = D(\beta(t))\tilde{U}(t), \quad \beta(t) : \mathbb{R} \mapsto \mathbb{C}, \quad (20)$$

where the unitary transformation operator is

$$D(\beta) = \exp(\beta a^\dagger - \beta^* a), \quad D(0) = I. \quad (21)$$

Thus, by taking $\beta(0) = 0$, we have $U(0) = \tilde{U}(0)$ and the transformed unitary solution operator satisfies

$$\dot{\tilde{U}}(t) = -i\tilde{H}(t)\tilde{U}(t), \quad \tilde{U}(0) = I, \quad (22)$$

$$\tilde{H}(t) = \left(c(t) - i\frac{d\beta(t)}{dt} \right) a^\dagger + \left(c^*(t) + i\frac{d\beta^*(t)}{dt} \right) a, \quad (23)$$

where we have omitted terms of the form $\delta(t)I$ in $\tilde{H}(t)$ because they only change the global phase in \tilde{U} . We conclude that $\tilde{U}(t) = I$, for all $t \geq 0$, if $\beta(t)$ satisfies the ordinary differential equation

$$\frac{d\beta(t)}{dt} = -ic(t), \quad \beta(0) = 0, \quad \Leftrightarrow \quad \beta(t) = -i \int_0^t c(\tau) d\tau. \quad (24)$$

In this case, we have $U(T) = D(\beta(T))\tilde{U}(T) = D(\beta(T))$, and the infidelity at time $t = T$ satisfies

$$J_{\text{infid}} = 1 - \left| \frac{1}{N} \text{Tr} \left(V_{\text{target}}^\dagger D(\beta(T)) \right) \right|^2. \quad (25)$$

For a given duration T , we note that the infidelity only depends on the integral $\beta(T) = -i \int_0^T c(\tau) d\tau$. As a result, any control function $\tilde{c}(t)$ with the same integral over time yields the same infidelity.

If we scale time, $t' = st$, and define $\beta'(t') = \beta(t)$, the infidelity with respect to V_{target} becomes invariant to the scaling as long as $\beta'(T') = \beta(T)$. The scaling implies

$$\beta'(T') = -i \int_0^{T'} c'(\tau') d\tau' = -i \int_0^T c'(s\tau) s d\tau. \quad (26)$$

Hence, $\beta'(sT) = \beta(T)$ if the scaled control pulse satisfies

$$c'(st) = \frac{c(t)}{s}. \quad (27)$$

For example, if $s > 1$ the duration of the control pulse is increased and its amplitude is decreased, such that the time integral of the control pulse remains unchanged.

B Control parameterization via B-spline basis functions

The control function $c(t)$ is parameterized in terms of a weighted sum of B-spline basis functions, according to (6). It is convenient to split the complex-valued function into its real and imaginary parts, $c(t) = p(t) + iq(t)$, where

$$p(t) = \sum_{s=1}^{N_s} \alpha_s^{\text{real}} B_s(t), \quad q(t) = \sum_{s=1}^{N_s} \alpha_s^{\text{imag}} B_s(t). \quad (28)$$

Here, $N_s \geq 1$ equals the number of terms in $p(t)$ and $q(t)$, specified by the real-valued design variables $\{\alpha_s^{\text{real}}, \alpha_s^{\text{imag}}\}_{s=1}^{N_s}$. In this work we let $B_s(t)$ be quadratic B-spline basis functions centered on a uniform grid in

- [3] Tommaso Caneva, Michael Murphy, Tommaso Calarco, Rosario Fazio, Simone Montangero, Vittorio Giovannetti, and Giuseppe E Santoro. Optimal control at the quantum speed limit. *Physical review letters*, 103(24):240501, 2009.
- [4] Yujin Cho, Kristin M Beck, Alessandro R Castelli, Kyle A Wendt, Bram Evert, Matthew J Reagor, and Jonathan L DuBois. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits. *arXiv preprint arXiv:2303.04261*, 2023.
- [5] Sebastian Deffner and Steve Campbell. Quantum speed limits: from heisenberg’s uncertainty principle to optimal quantum control. *Journal of Physics A: Mathematical and Theoretical*, 50(45):453001, 2017.
- [6] Frank Gaitan. *Quantum error correction and fault tolerant quantum computing*. CRC Press, 2008.
- [7] Xiaozhen Ge, Re-Bing Wu, and Herschel Rabitz. The optimization landscape of hybrid quantum–classical algorithms: From quantum control to nisq applications. *Annual Reviews in Control*, 2022.
- [8] Christopher Gerry and Peter Knight. *Introductory Quantum Optics*. Cambridge University Press, 2004.
- [9] Steffen J Glaser, Ugo Boscain, Tommaso Calarco, Christiane P Koch, Walter Köckenberger, Ronnie Kosloff, Ilya Kuprov, Burkhard Luy, Sophie Schirmer, Thomas Schulte-Herbrüggen, 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*, 69:1–24, 2015.
- [10] Stefanie Günther, N Anders Petersson, and Jonathan L DuBois. Quandary: an open-source c++ package for high-performance optimal control of open quantum systems. In *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)*, pages 88–98. IEEE, 2021.
- [11] Gerhard C Hegerfeldt. Driving at the quantum speed limit: optimal control of a two-level system. *Physical review letters*, 111(26):260501, 2013.
- [12] Navin Khaneja, Roger Brockett, and Steffen J. Glaser. Time optimal control in spin systems. *Phys. Rev. A*, 63:032308, Feb 2001.
- [13] Navin Khaneja, Steffen J. Glaser, and Roger Brockett. Sub-riemannian geometry and time optimal control of three spin systems: Quantum gates and coherence transfer. *Phys. Rev. A*, 65:032301, Jan 2002.
- [14] Christiane P Koch. Controlling open quantum systems: tools, achievements, and limitations. *Journal of Physics: Condensed Matter*, 28(21):213001, 2016.
- [15] Juneseo Lee, Christian Arenz, Herschel Rabitz, and Benjamin Russell. Dependence of the quantum speed limit on system size and control complexity. *New Journal of Physics*, 20(6):063002, 2018.
- [16] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [17] José P. Palao and Ronnie Kosloff. Quantum computing by an optimal control algorithm for unitary transformations. *Phys. Rev. Lett.*, 89:188301, Oct 2002.
- [18] Les Piegl and Wayne Tiller. *The NURBS Book*. Springer-Verlag, New York, NY, USA, second edition, 1996.
- [19] Katharina Rojan, Daniel M. Reich, Igor Dotsenko, Jean-Michel Raimond, Christiane P. Koch, and Giovanna Morigi. Arbitrary-quantum-state preparation of a harmonic oscillator via optimal control. *Phys. Rev. A*, 90, 2014.
- [20] Benjamin Russell, Herschel Rabitz, and Rebing Wu. Quantum control landscapes are almost always trap free. *arXiv preprint arXiv:1608.06198*, 2016.
- [21] Sonia Schirmer. Implementation of quantum gates via optimal control. *Journal of Modern Optics*, 56(6):831–839, 2009.
- [22] Lennart Maximilian Seifert, Jason Chadwick, Andrew Litteken, Frederic T Chong, and Jonathan M Baker. Time-efficient qudit gates through incremental pulse re-seeding. In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 304–313. IEEE, 2022.
- [23] M. Unser. Ten good reasons for using spline wavelets. *Wavelet Applications in Signal and Image Processing V*, 3169, 09 2002.
- [24] Gerald Waldherr, Y Wang, S Zaiser, M Jamali, T Schulte-Herbrüggen, H Abe, T Ohshima, J Isoya, JF Du, P Neumann, et al. Quantum error correction in a solid-state hybrid spin register. *Nature*, 506(7487):204–207, 2014.